

# User Documentation for IDAS v6.1.0

SUNDIALS v7.1.0

Radu Serban<sup>1</sup>, Cosmin Petra<sup>1</sup>, Alan C. Hindmarsh<sup>1</sup>, Cody J. Balos<sup>1</sup>,  
David J. Gardner<sup>1</sup>, Daniel R. Reynolds<sup>2</sup>, and Carol S. Woodward<sup>1</sup>

<sup>1</sup>*Center for Applied Scientific Computing, Lawrence Livermore National Laboratory*

<sup>2</sup>*Department of Mathematics, Southern Methodist University*

June 20, 2024



UCRL-SM-208112

### **DISCLAIMER**

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

## **CONTRIBUTORS**

The SUNDIALS library has been developed over many years by a number of contributors. The current SUNDIALS team consists of Cody J. Balos, David J. Gardner, Alan C. Hindmarsh, Daniel R. Reynolds, and Carol S. Woodward. We thank Radu Serban for significant and critical past contributions.

Other contributors to SUNDIALS include: James Almgren-Bell, Lawrence E. Banks, Peter N. Brown, George Byrne, Rujeko Chinomona, Scott D. Cohen, Aaron Collier, Keith E. Grant, Steven L. Lee, Shelby L. Lockhart, John Loffeld, Daniel McGreer, Yu Pan, Slaven Peles, Cosmin Petra, Steven B. Roberts, H. Hunter Schwartz, Jean M. Sexton, Dan Shumaker, Steve G. Smith, Shahbaj Sohal, Allan G. Taylor, Hilari C. Tiedeman, Chris White, Ting Yan, and Ulrike M. Yang.



# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Changes to SUNDIALS in release 6.1.0 . . . . .	2
1.2	Reading this User Guide . . . . .	4
1.3	SUNDIALS License and Notices . . . . .	4
<b>2</b>	<b>Mathematical Considerations</b>	<b>7</b>
2.1	Initial Condition . . . . .	7
2.2	IVP solution . . . . .	8
2.3	Preconditioning . . . . .	12
2.4	Rootfinding . . . . .	12
2.5	Pure quadrature integration . . . . .	13
2.6	Forward sensitivity analysis . . . . .	14
2.7	Adjoint sensitivity analysis . . . . .	17
2.8	Second-order sensitivity analysis . . . . .	20
<b>3</b>	<b>Code Organization</b>	<b>23</b>
<b>4</b>	<b>Getting Started</b>	<b>25</b>
4.1	Data Types . . . . .	26
4.2	The SUNContext Type . . . . .	28
4.3	Error Checking . . . . .	33
4.4	Status and Error Logging . . . . .	36
4.5	Performance Profiling . . . . .	40
4.6	Getting Version Information . . . . .	43
4.7	Fortran Interface . . . . .	44
4.8	Features for GPU Accelerated Computing . . . . .	53
<b>5</b>	<b>Using IDAS</b>	<b>57</b>
5.1	Using IDAS for IVP Solution . . . . .	57
5.2	Integration of pure quadrature equations . . . . .	108
5.3	Preconditioner modules . . . . .	115
5.4	Using IDAS for Forward Sensitivity Analysis . . . . .	120
5.5	Using IDAS for Adjoint Sensitivity Analysis . . . . .	147
<b>6</b>	<b>Vector Data Structures</b>	<b>187</b>
6.1	Description of the NVECTOR Modules . . . . .	187
6.2	Description of the NVECTOR operations . . . . .	195
6.3	NVECTOR functions used by IDAS . . . . .	208
6.4	The NVECTOR_SERIAL Module . . . . .	209
6.5	The NVECTOR_PARALLEL Module . . . . .	212
6.6	The NVECTOR_OPENMP Module . . . . .	216
6.7	The NVECTOR_PTHREADS Module . . . . .	219
6.8	The NVECTOR_PARHYP Module . . . . .	223

6.9	The NVECTOR_PETSC Module . . . . .	225
6.10	The NVECTOR_CUDA Module . . . . .	227
6.11	The NVECTOR_HIP Module . . . . .	232
6.12	The NVECTOR_SYCL Module . . . . .	237
6.13	The NVECTOR_RAJA Module . . . . .	242
6.14	The NVECTOR_KOKKOS Module . . . . .	245
6.15	The NVECTOR_OPENMPDEV Module . . . . .	248
6.16	The NVECTOR_TRILINOS Module . . . . .	251
6.17	The NVECTOR_MANYVECTOR Module . . . . .	252
6.18	The NVECTOR_MPIMANYVECTOR Module . . . . .	255
6.19	The NVECTOR_MPIPLUSX Module . . . . .	259
6.20	NVECTOR Examples . . . . .	260
<b>7</b>	<b>Matrix Data Structures</b>	<b>265</b>
7.1	Description of the SUNMATRIX Modules . . . . .	265
7.2	Description of the SUNMATRIX operations . . . . .	268
7.3	The SUNMATRIX_DENSE Module . . . . .	269
7.4	The SUNMATRIX_MAGMADENSE Module . . . . .	272
7.5	The SUNMATRIX_ONEMKLDENSE Module . . . . .	277
7.6	The SUNMATRIX_BAND Module . . . . .	281
7.7	The SUNMATRIX_CUSPARSE Module . . . . .	287
7.8	The SUNMATRIX_SPARSE Module . . . . .	290
7.9	The SUNMATRIX_SLUNRLOC Module . . . . .	296
7.10	The SUNMATRIX_GINKGO Module . . . . .	298
7.11	The SUNMATRIX_KOKKOSDENSE Module . . . . .	300
7.12	SUNMATRIX Examples . . . . .	303
7.13	SUNMatrix functions used by IDAS . . . . .	304
<b>8</b>	<b>Linear Algebraic Solvers</b>	<b>305</b>
8.1	The SUNLinearSolver API . . . . .	306
8.2	IDAS SUNLinearSolver interface . . . . .	319
8.3	The SUNLinSol_Band Module . . . . .	320
8.4	The SUNLinSol_Dense Module . . . . .	322
8.5	The SUNLinSol_KLU Module . . . . .	323
8.6	The SUNLinSol_LapackBand Module . . . . .	327
8.7	The SUNLinSol_LapackDense Module . . . . .	329
8.8	The SUNLinSol_MagmaDense Module . . . . .	330
8.9	The SUNLinSol_OneMklDense Module . . . . .	332
8.10	The SUNLinSol_PCG Module . . . . .	333
8.11	The SUNLinSol_SPBCGS Module . . . . .	337
8.12	The SUNLinSol_SPGMR Module . . . . .	340
8.13	The SUNLinSol_SPGMR Module . . . . .	344
8.14	The SUNLinSol_SPTFQMR Module . . . . .	348
8.15	The SUNLinSol_SuperLUDIST Module . . . . .	351
8.16	The SUNLinSol_SuperLUMT Module . . . . .	354
8.17	The SUNLinSol_cuSolverSp_batchQR Module . . . . .	357
8.18	The SUNLINEARSOLVER_GINKGO Module . . . . .	358
8.19	The SUNLINEARSOLVER_KOKKOSDENSE Module . . . . .	361
8.20	SUNLinearSolver Examples . . . . .	363
<b>9</b>	<b>Nonlinear Algebraic Solvers</b>	<b>365</b>
9.1	The SUNNonlinearSolver API . . . . .	365
9.2	IDAS SUNNonlinearSolver interface . . . . .	374
9.3	The SUNNonlinSol_Newton implementation . . . . .	378

9.4	The SUNNonlinSol_FixedPoint implementation . . . . .	381
9.5	The SUNNonlinSol_PetscSNES implementation . . . . .	384
<b>10</b>	<b>Tools for Memory Management</b>	<b>387</b>
10.1	The SUNMemoryHelper API . . . . .	387
10.2	The SUNMemoryHelper_Cuda Implementation . . . . .	392
10.3	The SUNMemoryHelper_Hip Implementation . . . . .	394
10.4	The SUNMemoryHelper_Sycl Implementation . . . . .	396
<b>11</b>	<b>Acquiring SUNDIALS</b>	<b>399</b>
<b>12</b>	<b>Building and Installing with CMake</b>	<b>401</b>
12.1	Configuring, building, and installing on Unix-like systems . . . . .	401
12.2	Configuration options . . . . .	405
12.3	Configuration examples . . . . .	417
12.4	Working with external Libraries . . . . .	418
12.5	Testing the build and installation . . . . .	423
12.6	Building and Running Examples . . . . .	423
12.7	Configuring, building, and installing on Windows . . . . .	423
12.8	Installed libraries and exported header files . . . . .	424
12.9	Using SUNDIALS in your project . . . . .	424
12.10	Using SUNDIALS as a Third Party Library in other CMake Projects . . . . .	425
12.11	Table of SUNDIALS libraries and header files . . . . .	426
12.12	Installing SUNDIALS on HPC Clusters . . . . .	429
12.13	Building with SUNDIALS Addons . . . . .	430
<b>13</b>	<b>IDAS Constants</b>	<b>431</b>
13.1	IDAS input constants . . . . .	431
13.2	IDAS output constants . . . . .	432
<b>14</b>	<b>Release History</b>	<b>435</b>
<b>15</b>	<b>Changelog</b>	<b>437</b>
15.1	Changes to SUNDIALS in release 7.1.0 . . . . .	437
15.2	Changes to SUNDIALS in release 7.0.0 . . . . .	439
15.3	Changes to SUNDIALS in release 6.7.0 . . . . .	442
15.4	Changes to SUNDIALS in release 6.6.2 . . . . .	443
15.5	Changes to SUNDIALS in release 6.6.1 . . . . .	443
15.6	Changes to SUNDIALS in release 6.6.0 . . . . .	443
15.7	Changes to SUNDIALS in release 6.5.1 . . . . .	444
15.8	Changes to SUNDIALS in release 6.5.0 . . . . .	444
15.9	Changes to SUNDIALS in release 6.4.1 . . . . .	445
15.10	Changes to SUNDIALS in release 6.4.0 . . . . .	445
15.11	Changes to SUNDIALS in release 6.3.0 . . . . .	446
15.12	Changes to SUNDIALS in release 6.2.0 . . . . .	447
15.13	Changes to SUNDIALS in release 6.1.1 . . . . .	450
15.14	Changes to SUNDIALS in release 6.1.0 . . . . .	450
15.15	Changes to SUNDIALS in release 6.0.0 . . . . .	450
15.16	Changes to SUNDIALS in release 5.8.0 . . . . .	456
15.17	Changes to SUNDIALS in release 5.7.0 . . . . .	457
15.18	Changes to SUNDIALS in release 5.6.1 . . . . .	457
15.19	Changes to SUNDIALS in release 5.6.0 . . . . .	457
15.20	Changes to SUNDIALS in release 5.5.0 . . . . .	458
15.21	Changes to SUNDIALS in release 5.4.0 . . . . .	458
15.22	Changes to SUNDIALS in release 5.3.0 . . . . .	460

15.23 Changes to SUNDIALS in release 5.2.0 . . . . .	461
15.24 Changes to SUNDIALS in release 5.1.0 . . . . .	462
15.25 Changes to SUNDIALS in release 5.0.0 . . . . .	463
15.26 Changes to SUNDIALS in release 4.1.0 . . . . .	466
15.27 Changes to SUNDIALS in release 4.0.2 . . . . .	467
15.28 Changes to SUNDIALS in release 4.0.1 . . . . .	467
15.29 Changes to SUNDIALS in release 4.0.0 . . . . .	467
15.30 Changes to SUNDIALS in release 3.2.1 . . . . .	470
15.31 Changes to SUNDIALS in release 3.2.0 . . . . .	470
15.32 Changes to SUNDIALS in release 3.1.2 . . . . .	471
15.33 Changes to SUNDIALS in release 3.1.1 . . . . .	471
15.34 Changes to SUNDIALS in release 3.1.0 . . . . .	472
15.35 Changes to SUNDIALS in release 3.0.0 . . . . .	472
15.36 Changes to SUNDIALS in release 2.7.0 . . . . .	474
15.37 Changes to SUNDIALS in release 2.6.2 . . . . .	476
15.38 Changes to SUNDIALS in release 2.6.1 . . . . .	476
15.39 Changes to SUNDIALS in release 2.6.0 . . . . .	477
15.40 Changes to SUNDIALS in release 2.5.0 . . . . .	478
15.41 Changes to SUNDIALS in release 2.4.0 . . . . .	479
15.42 Changes to SUNDIALS in release 2.3.0 . . . . .	480
15.43 Changes to SUNDIALS in release 2.2.0 . . . . .	480
15.44 Changes to SUNDIALS in release 2.1.1 . . . . .	481
15.45 Changes to SUNDIALS in release 2.1.0 . . . . .	481
15.46 Changes to SUNDIALS in release 2.0.2 . . . . .	481
15.47 Changes to SUNDIALS in release 2.0.1 . . . . .	481
15.48 Changes to SUNDIALS in release 2.0.0 . . . . .	482
<b>Bibliography</b>	<b>483</b>
<b>Index</b>	<b>487</b>



# Chapter 1

## Introduction

IDAS is part of a software family called SUNDIALS: SUite of Nonlinear and Differential/ALgebraic equation Solvers [41]. This suite consists of CVODE, ARKODE, KINSOL, and IDAS, and variants of these with sensitivity analysis capabilities, CVODES and IDAS.

IDAS is a general purpose solver for the initial value problem (IVP) for systems of differential-algebraic equations (DAEs). The name IDAS stands for Implicit Differential-Algebraic solver with Sensitivity capabilities. IDAS is an extension of the IDA solver within SUNDIALS, itself based on on DASPK [18, 19], but is written in ANSI-standard C rather than Fortran77. Its most notable features are that, (1) in the solution of the underlying nonlinear system at each time step, it offers a choice of Newton/direct methods and a choice of Inexact Newton/Krylov (iterative) methods; and (2) it is written in a *data-independent* manner in that it acts on generic vectors and matrices without any assumptions on the underlying organization of the data. Thus IDAS shares significant modules previously written within CASC at LLNL to support the ordinary differential equation (ODE) solvers CVODE [26, 44] and PVODE [22, 23], and also the nonlinear system solver KINSOL [45].

At present, IDAS may utilize a variety of Krylov methods provided in SUNDIALS that can be used in conjunction with Newton iteration: these include the GMRES (Generalized Minimal RESidual) [60], FGMRES (Flexible Generalized Minimum RESidual) [59], Bi-CGStab (Bi-Conjugate Gradient Stabilized) [65], TFQMR (Transpose-Free Quasi-Minimal Residual) [35], and PCG (Preconditioned Conjugate Gradient) [38] linear iterative methods. As Krylov methods, these require little matrix storage for solving the Newton equations as compared to direct methods. However, the algorithms allow for a user-supplied preconditioner, and, for most problems, preconditioning is essential for an efficient solution.

For very large DAE systems, the Krylov methods are preferable over direct linear solver methods, and are often the only feasible choice. Among the Krylov methods in SUNDIALS, we recommend GMRES as the best overall choice. However, users are encouraged to compare all options, especially if encountering convergence failures with GMRES. Bi-CGStab and TFQMR have an advantage in storage requirements, in that the number of workspace vectors they require is fixed, while that number for GMRES depends on the desired Krylov subspace size. FGMRES has an advantage in that it is designed to support preconditioners that vary between iterations (e.g. iterative methods). PCG exhibits rapid convergence and minimal workspace vectors, but only works for symmetric linear systems.

IDAS is written with a functionality that is a superset of that of IDA. Sensitivity analysis capabilities, both forward and adjoint, have been added to the main integrator. Enabling forward sensitivity computations in IDAS will result in the code integrating the so-called *sensitivity equations* simultaneously with the original IVP, yielding both the solution and its sensitivity with respect to parameters in the model. Adjoint sensitivity analysis, most useful when the gradients of relatively few functionals of the solution with respect to many parameters are sought, involves integration of the original IVP forward in time followed by the integration of the so-called *adjoint equations* backward in time. IDAS provides the infrastructure needed to integrate any final-condition ODE dependent on the solution of the original IVP (in particular the adjoint system).

## 1.1 Changes to SUNDIALS in release 6.1.0

### Major Features

Created shared user interface functions for ARKODE to allow more uniform control over time-stepping algorithms, improved extensibility, and simplified code maintenance. The corresponding stepper-specific user-callable functions are now deprecated and will be removed in a future major release.

Added CMake infrastructure that enables externally maintained addons/plugins to be *optionally* built with SUNDIALS. See [Contributing](#) for details.

### New Features and Enhancements

Added support for Kokkos Kernels v4.

Added the following Runge-Kutta Butcher tables

- `ARKODE_FORWARD_EULER_1_1`
- `ARKODE_RALSTON_EULER_2_1_2`
- `ARKODE_EXPLICIT_MIDPOINT_EULER_2_1_2`
- `ARKODE_BACKWARD_EULER_1_1`
- `ARKODE_IMPLICIT_MIDPOINT_1_2`
- `ARKODE_IMPLICIT_TRAPEZOIDAL_2_2`

Added the following MRI coupling tables

- `ARKODE_MRI_GARK_FORWARD_EULER`
- `ARKODE_MRI_GARK_RALSTON2`
- `ARKODE_MRI_GARK_RALSTON3`
- `ARKODE_MRI_GARK_BACKWARD_EULER`
- `ARKODE_MRI_GARK_IMPLICIT_MIDPOINT`
- `ARKODE_IMEX_MRI_GARK_EULER`
- `ARKODE_IMEX_MRI_GARK_TRAPEZOIDAL`
- `ARKODE_IMEX_MRI_GARK_MIDPOINT`

Added `ARKodeButcherTable_ERKIDToName()` and `ARKodeButcherTable_DIRKIDToName()` to convert a Butcher table ID to a string representation.

Added the function `ARKodeSetAutonomous()` in ARKODE to indicate that the implicit right-hand side function does not explicitly depend on time. When using the trivial predictor, an autonomous problem may reuse implicit function evaluations across stage solves to reduce the total number of function evaluations.

Users may now disable interpolated output in ARKODE by passing `ARK_INTERP_NONE` to `ARKodeSetInterpolantType()`. When interpolation is disabled, rootfinding is not supported, implicit methods must use the trivial predictor (the default option), and interpolation at stop times cannot be used (interpolating at stop times is disabled by default). With interpolation disabled, calling `ARKodeEvolve()` in `ARK_NORMAL` mode will return at or past the requested output time (setting a stop time may still be used to halt the integrator at a specific time). Disabling interpolation will reduce the memory footprint of an integrator by two or more state vectors (depending on the interpolant type and degree) which can be beneficial when interpolation is not needed e.g., when integrating to a final time without output in between or using an explicit fast time scale integrator with an MRI method.

Added “Resize” capability to ARKODE’s `SPRKStep` time-stepping module.

Enabled the Fortran interfaces to build with 32-bit `sunindextype`.

## Bug Fixes

Updated the CMake variable HIP\_PLATFORM default to amd as the previous default, hcc, is no longer recognized in ROCm 5.7.0 or newer. The new default is also valid in older version of ROCm (at least back to version 4.3.1).

Renamed the DPCPP value for the `SUNDIALS_GINKGO_BACKENDS` CMake option to SYCL to match Ginkgo's updated naming convention.

Changed the CMake version compatibility mode for SUNDIALS to `AnyNewerVersion` instead of `SameMajorVersion`. This fixes the issue seen [here](#).

Fixed a CMake bug that caused an MPI linking error for our C++ examples in some instances. Fixes [GitHub Issue #464](#).

Fixed the runtime library installation path for windows systems. This fix changes the default library installation path from `CMAKE_INSTALL_PREFIX/CMAKE_INSTALL_LIBDIR` to `CMAKE_INSTALL_PREFIX/CMAKE_INSTALL_BINDIR`.

Fixed conflicting `.lib` files between shared and static libs when using MSVC on Windows

Fixed invalid `SUNDIALS_EXPORT` generated macro when building both shared and static libs.

Fixed a bug in some Fortran examples where `c_null_ptr` was passed as an argument to a function pointer instead of `c_null_funptr`. This caused compilation issues with the Cray Fortran compiler.

Fixed a bug in the HIP execution policies where `WARP_SIZE` would not be set with ROCm 6.0.0 or newer.

Fixed a bug that caused error messages to be cut off in some cases. Fixes [GitHub Issue #461](#).

Fixed a memory leak when an error handler was added to a `SUNContext`. Fixes [GitHub Issue #466](#).

Fixed a bug where `MRISStepEvolve()` would not handle a recoverable error produced from evolving the inner stepper.

Added missing `SetRootDirection` and `SetNoInactiveRootWarn` functions to ARKODE's `SPRKStep` time-stepping module.

Fixed a bug in `ARKodeSPRKTable_Create()` where the coefficient arrays were not allocated.

Fix bug on LLP64 platforms (like Windows 64-bit) where `KLU_INDEXTYPE` could be 32 bits wide even if `SUNDIALS_INT64_T` is defined.

Check if size of `SuiteSparse_long` is 8 if the size of `sunindextype` is 8 when using KLU.

Fixed several build errors with the Fortran interfaces on Windows systems.

## Deprecation Notices

Numerous ARKODE stepper-specific functions are now deprecated in favor of ARKODE-wide functions.

Deprecated the `ARKStepSetOptimalParams` function. Since this function does not have an ARKODE-wide equivalent, instructions have been added to the user guide for how to retain the current functionality using other user-callable functions.

The unsupported implementations of `N_VGetArrayPointer` and `N_VSetArrayPointer` for the *hypr*e and PETSc vectors are now deprecated. Users should access the underlying wrapped external library vector objects instead with `N_VGetVector_ParHyp` and `N_VGetVector_Petsc`, respectively.

For changes in prior versions of SUNDIALS see [§15](#).

## 1.2 Reading this User Guide

The structure of this document is as follows:

- In Chapter §2, we give short descriptions of the numerical methods implemented by IDAS for the solution of initial value problems for systems of DAEs, along with short descriptions of preconditioning (§2.3) and rootfinding (§2.4).
- The following chapter describes the software organization of the IDAS solver (§3).
- Chapter §5.1 is the main usage document for IDAS for simulation applications. It includes a complete description of the user interface for the integration of DAE initial value problems. Readers that are not interested in using IDAS for sensitivity analysis can then skip the next two chapters.
- Chapter §5.4 describes the usage of IDAS for forward sensitivity analysis as an extension of its IVP integration capabilities. We begin with a skeleton of the user main program, with emphasis on the steps that are required in addition to those already described in Chapter §5.1. Following that we provide detailed descriptions of the user-callable interface routines specific to forward sensitivity analysis and of the additional optional user-defined routines.
- Chapter §5.5 describes the usage of IDAS for adjoint sensitivity analysis. We begin by describing the IDAS checkpointing implementation for interpolation of the original IVP solution during integration of the adjoint system backward in time, and with an overview of a user's main program. Following that we provide complete descriptions of the user-callable interface routines for adjoint sensitivity analysis as well as descriptions of the required additional user-defined routines.
- Chapter §6 gives a brief overview of the generic `N_Vector` module shared among the various components of SUNDIALS, as well as details on the `N_Vector` implementations provided with SUNDIALS.
- Chapter §7 gives a brief overview of the generic `SUNMatrix` module shared among the various components of SUNDIALS, and details on the `SUNMatrix` implementations provided with SUNDIALS.
- Chapter §8 gives a brief overview of the generic `SUNLinearSolver` module shared among the various components of SUNDIALS. This chapter contains details on the `SUNLinearSolver` implementations provided with SUNDIALS. The chapter also contains details on the `SUNLinearSolver` implementations provided with SUNDIALS that interface with external linear solver libraries.
- Chapter §9 describes the `SUNNonlinearSolver` API and nonlinear solver implementations shared among the various components of SUNDIALS.
- Finally, in the appendices, we provide detailed instructions for the installation of IDAS, within the structure of SUNDIALS (Appendix §11), as well as a list of all the constants used for input to and output from IDAS functions (Appendix §13).

## 1.3 SUNDIALS License and Notices

All SUNDIALS packages are released open source, under the BSD 3-Clause license. The only requirements of the license are preservation of copyright and a standard disclaimer of liability. The full text of the license and an additional notice are provided below and may also be found in the LICENSE and NOTICE files provided with all SUNDIALS packages.

---

**Note:** If you are using SUNDIALS with any third party libraries linked in (e.g., LAPACK, KLU, SuperLU\_MT, PETSc, or *hypre*), be sure to review the respective license of the package as that license may have more restrictive terms than the SUNDIALS license. For example, if someone builds SUNDIALS with a statically linked KLU, the build is

subject to terms of the more-restrictive LGPL license (which is what KLU is released with) and *not* the SUNDIALS BSD license anymore.

---

### 1.3.1 BSD 3-Clause License

Copyright (c) 2002-2024, Lawrence Livermore National Security and Southern Methodist University.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the copyright holder nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

### 1.3.2 Additional Notice

This work was produced under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

This work was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights.

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or Lawrence Livermore National Security, LLC.

The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

### **1.3.3 SUNDIALS Release Numbers**

LLNL-CODE-667205 (ARKODE)

UCRL-CODE-155951 (CVODE)

UCRL-CODE-155950 (CVODES)

UCRL-CODE-155952 (IDA)

UCRL-CODE-237203 (IDAS)

LLNL-CODE-665877 (KINSOL)

## Chapter 2

# Mathematical Considerations

IDAS solves the initial-value problem (IVP) for a DAE system of the general form

$$F(t, y, \dot{y}) = 0, \quad y(t_0) = y_0, \quad \dot{y}(t_0) = \dot{y}_0 \quad (2.1)$$

where  $y$ ,  $\dot{y}$ , and  $F$  are vectors in  $\mathbb{R}^N$ ,  $t$  is the independent variable,  $\dot{y} = dy/dt$ , and initial values  $y_0$ ,  $\dot{y}_0$  are given. Often  $t$  is time, but it certainly need not be.

Additionally, if (2.1) depends on some parameters  $p \in \mathbb{R}^{N_p}$ , i.e.

$$\begin{aligned} F(t, y, \dot{y}, p) &= 0 \\ y(t_0) &= y_0(p), \quad \dot{y}(t_0) = \dot{y}_0(p), \end{aligned} \quad (2.2)$$

IDAS can also compute first order derivative information, performing either *forward sensitivity analysis* or *adjoint sensitivity analysis*. In the first case, IDAS computes the sensitivities of the solution with respect to the parameters  $p$ , while in the second case, IDAS computes the gradient of a *derived function* with respect to the parameters  $p$ .

### 2.1 Initial Condition

Prior to integrating a DAE initial-value problem, an important requirement is that the pair of vectors  $y_0$  and  $\dot{y}_0$  are both initialized to satisfy the DAE residual  $F(t_0, y_0, \dot{y}_0) = 0$ . For a class of problems that includes so-called semi-explicit index-one systems, IDAS provides a routine that computes consistent initial conditions from a user's initial guess [19]. For this, the user must identify sub-vectors of  $y$  (not necessarily contiguous), denoted  $y_d$  and  $y_a$ , which are its differential and algebraic parts, respectively, such that  $F$  depends on  $\dot{y}_d$  but not on any components of  $\dot{y}_a$ . The assumption that the system is “index one” means that for a given  $t$  and  $y_d$ , the system  $F(t, y, \dot{y}) = 0$  defines  $y_a$  uniquely. In this case, a solver within IDAS computes  $y_a$  and  $\dot{y}_d$  at  $t = t_0$ , given  $y_d$  and an initial guess for  $y_a$ . A second available option with this solver also computes all of  $y(t_0)$  given  $\dot{y}(t_0)$ ; this is intended mainly for quasi-steady-state problems, where  $\dot{y}(t_0) = 0$  is given. In both cases, IDAS solves the system  $F(t_0, y_0, \dot{y}_0) = 0$  for the unknown components of  $y_0$  and  $\dot{y}_0$ , using a Newton iteration augmented with a line search globalization strategy. In doing this, it makes use of the existing machinery that is to be used for solving the linear systems during the integration, in combination with certain tricks involving the step size (which is set artificially for this calculation). For problems that do not fall into either of these categories, the user is responsible for passing consistent values, or risks failure in the numerical integration.

## 2.2 IVP solution

The integration method used in IDAS is the variable-order, variable-coefficient BDF (Backward Differentiation Formula), in fixed-leading-coefficient form [14]. The method order ranges from 1 to 5, with the BDF of order  $q$  given by the multistep formula

$$\sum_{i=0}^q \alpha_{n,i} y_{n-i} = h_n \dot{y}_n, \quad (2.3)$$

where  $y_n$  and  $\dot{y}_n$  are the computed approximations to  $y(t_n)$  and  $\dot{y}(t_n)$ , respectively, and the step size is  $h_n = t_n - t_{n-1}$ . The coefficients  $\alpha_{n,i}$  are uniquely determined by the order  $q$ , and the history of the step sizes. The application of the BDF (2.3) to the DAE system (2.1) results in a nonlinear algebraic system to be solved at each step:

$$G(y_n) \equiv F \left( t_n, y_n, h_n^{-1} \sum_{i=0}^q \alpha_{n,i} y_{n-i} \right) = 0. \quad (2.4)$$

In the process of controlling errors at various levels, IDAS uses a weighted root-mean-square norm, denoted  $\|\cdot\|_{\text{WRMS}}$ , for all error-like quantities. The multiplicative weights used are based on the current solution and on the relative and absolute tolerances input by the user, namely

$$W_i = \frac{1}{\text{rtol} \cdot |y_i| + \text{atol}_i}. \quad (2.5)$$

Because  $1/W_i$  represents a tolerance in the component  $y_i$ , a vector whose norm is 1 is regarded as “small.” For brevity, we will usually drop the subscript WRMS on norms in what follows.

### 2.2.1 Nonlinear Solve

By default IDAS solves (2.4) with a Newton iteration but IDAS also allows for user-defined nonlinear solvers (see Chapter §9). Each Newton iteration requires the solution of a linear system of the form

$$J[y_{n(m+1)} - y_{n(m)}] = -G(y_{n(m)}), \quad (2.6)$$

where  $y_{n(m)}$  is the  $m$ -th approximation to  $y_n$ . Here  $J$  is some approximation to the system Jacobian

$$J = \frac{\partial G}{\partial y} = \frac{\partial F}{\partial y} + \alpha \frac{\partial F}{\partial \dot{y}}, \quad (2.7)$$

where  $\alpha = \alpha_{n,0}/h_n$ . The scalar  $\alpha$  changes whenever the step size or method order changes.

For the solution of the linear systems within the Newton iteration, IDAS provides several choices, including the option of a user-supplied linear solver (see Chapter §8). The linear solvers distributed with SUNDIALS are organized in two families, a *direct* family comprising direct linear solvers for dense, banded, or sparse matrices and a *spils* family comprising scaled preconditioned iterative (Krylov) linear solvers. The methods offered through these modules are as follows:

- dense direct solvers, including an internal implementation, an interface to BLAS/LAPACK, an interface to MAGMA [62] and an interface to the oneMKL library [3],
- band direct solvers, including an internal implementation or an interface to BLAS/LAPACK,
- sparse direct solver interfaces to various libraries, including KLU [4, 27], SuperLU\_MT [9, 29, 53], SuperLU\_Dist [8, 37, 54, 55], and cuSPARSE [7],
- SPGMR, a scaled preconditioned GMRES (Generalized Minimal Residual method) solver with or without restarts,



- SPFGMR, a scaled preconditioned FGMRES (Flexible Generalized Minimal Residual method) solver with or without restarts,
- SPBCG, a scaled preconditioned Bi-CGStab (Bi-Conjugate Gradient Stable method) solver,
- SPTFQMR, a scaled preconditioned TFQMR (Transpose-Free Quasi-Minimal Residual method) solver, or
- PCG, a scaled preconditioned CG (Conjugate Gradient method) solver.

For large stiff systems, where direct methods are not feasible, the combination of a BDF integrator and a preconditioned Krylov method yields a powerful tool because it combines established methods for stiff integration, nonlinear iteration, and Krylov (linear) iteration with a problem-specific treatment of the dominant source of stiffness, in the form of the user-supplied preconditioner matrix [17]. For the *spils* linear solvers with IDAS, preconditioning is allowed only on the left (see §2.3). Note that the dense, band, and sparse direct linear solvers can only be used with serial and threaded vector representations.

In the case of a matrix-based linear solver, the default Newton iteration is a Modified Newton iteration, in that the Jacobian  $J$  is fixed (and usually out of date) throughout the nonlinear iterations, with a coefficient  $\bar{\alpha}$  in place of  $\alpha$  in  $J$ . However, in the case that a matrix-free iterative linear solver is used, the default Newton iteration is an Inexact Newton iteration, in which  $J$  is applied in a matrix-free manner, with matrix-vector products  $Jv$  obtained by either difference quotients or a user-supplied routine. In this case, the linear residual  $J\Delta y + G$  is nonzero but controlled. With the default Newton iteration, the matrix  $J$  and preconditioner matrix  $P$  are updated as infrequently as possible to balance the high costs of matrix operations against other costs. Specifically, this matrix update occurs when:

- starting the problem,
- the value  $\bar{\alpha}$  at the last update is such that  $\alpha/\bar{\alpha} < 3/5$  or  $\alpha/\bar{\alpha} > 5/3$ , or
- a non-fatal convergence failure occurred with an out-of-date  $J$  or  $P$ .

The above strategy balances the high cost of frequent matrix evaluations and preprocessing with the slow convergence due to infrequent updates. To reduce storage costs on an update, Jacobian information is always reevaluated from scratch.

The default stopping test for nonlinear solver iterations in IDAS ensures that the iteration error  $y_n - y_{n(m)}$  is small relative to  $y$  itself. For this, we estimate the linear convergence rate at all iterations  $m > 1$  as

$$R = \left( \frac{\delta_m}{\delta_1} \right)^{\frac{1}{m-1}},$$

where the  $\delta_m = y_{n(m)} - y_{n(m-1)}$  is the correction at iteration  $m = 1, 2, \dots$ . The nonlinear solver iteration is halted if  $R > 0.9$ . The convergence test at the  $m$ -th iteration is then

$$S \|\delta_m\| < 0.33, \quad (2.8)$$

where  $S = R/(R - 1)$  whenever  $m > 1$  and  $R \leq 0.9$ . The user has the option of changing the constant in the convergence test from its default value of 0.33. The quantity  $S$  is set to  $S = 20$  initially and whenever  $J$  or  $P$  is updated, and it is reset to  $S = 100$  on a step with  $\alpha \neq \bar{\alpha}$ . Note that at  $m = 1$ , the convergence test (2.8) uses an old value for  $S$ . Therefore, at the first nonlinear solver iteration, we make an additional test and stop the iteration if  $\|\delta_1\| < 0.33 \cdot 10^{-4}$  (since such a  $\delta_1$  is probably just noise and therefore not appropriate for use in evaluating  $R$ ). We allow only a small number (default value 4) of nonlinear iterations. If convergence fails with  $J$  or  $P$  current, we are forced to reduce the step size  $h_n$ , and we replace  $h_n$  by  $h_n \eta_{cf}$  (by default  $\eta_{cf} = 0.25$ ). The integration is halted after a preset number (default value 10) of convergence failures. Both the maximum number of allowable nonlinear iterations and the maximum number of nonlinear convergence failures can be changed by the user from their default values.

When an iterative method is used to solve the linear system, to minimize the effect of linear iteration errors on the nonlinear and local integration error controls, we require the preconditioned linear residual to be small relative to the allowed error in the nonlinear iteration, i.e.,  $\|P^{-1}(Jx + G)\| < 0.05 \cdot 0.33$ . The safety factor 0.05 can be changed by the user.

When the Jacobian is stored using either the `SUNMATRIX_DENSE` or `SUNMATRIX_BAND` matrix objects, the Jacobian  $J$  defined in (2.7) can be either supplied by the user or IDAS can compute  $J$  internally by difference quotients. In the latter case, we use the approximation

$$J_{ij} = [F_i(t, y + \sigma_j e_j, \dot{y} + \alpha \sigma_j e_j) - F_i(t, y, \dot{y})] / \sigma_j, \text{ with} \\ \sigma_j = \sqrt{U} \max\{|y_j|, |h \dot{y}_j|, 1/W_j\} \text{sign}(h \dot{y}_j),$$

where  $U$  is the unit roundoff,  $h$  is the current step size, and  $W_j$  is the error weight for the component  $y_j$  defined by (2.5). We note that with sparse and user-supplied matrix objects, the Jacobian *must* be supplied by a user routine.

In the case of an iterative linear solver, if a routine for  $Jv$  is not supplied, such products are approximated by

$$Jv = [F(t, y + \sigma v, \dot{y} + \alpha \sigma v) - F(t, y, \dot{y})] / \sigma,$$

where the increment  $\sigma = 1/\|v\|$ . As an option, the user can specify a constant factor that is inserted into this expression for  $\sigma$ .

## 2.2.2 Local Error Test

During the course of integrating the system, IDAS computes an estimate of the local truncation error, LTE, at the  $n$ -th time step, and requires this to satisfy the inequality

$$\|\text{LTE}\|_{\text{WRMS}} \leq 1.$$

Asymptotically, LTE varies as  $h^{q+1}$  at step size  $h$  and order  $q$ , as does the predictor-corrector difference  $\Delta_n \equiv y_n - y_{n(0)}$ . Thus there is a constant  $C$  such that

$$\text{LTE} = C\Delta_n + O(h^{q+2}),$$

and so the norm of LTE is estimated as  $|C| \cdot \|\Delta_n\|$ . In addition, IDAS requires that the error in the associated polynomial interpolant over the current step be bounded by 1 in norm. The leading term of the norm of this error is bounded by  $\bar{C}\|\Delta_n\|$  for another constant  $\bar{C}$ . Thus the local error test in IDAS is

$$\max\{|C|, \bar{C}\} \|\Delta_n\| \leq 1. \quad (2.9)$$

A user option is available by which the algebraic components of the error vector are omitted from the test (2.9), if these have been so identified.

## 2.2.3 Step Size and Order Selection

In IDAS, the local error test is tightly coupled with the logic for selecting the step size and order. First, there is an initial phase that is treated specially; for the first few steps, the step size is doubled and the order raised (from its initial value of 1) on every step, until (a) the local error test (2.9) fails, (b) the order is reduced (by the rules given below), or (c) the order reaches 5 (the maximum). For step and order selection on the general step, IDAS uses a different set of local error estimates, based on the asymptotic behavior of the local error in the case of fixed step sizes. At each of the orders  $q'$  equal to  $q$ ,  $q-1$  (if  $q > 1$ ),  $q-2$  (if  $q > 2$ ), or  $q+1$  (if  $q < 5$ ), there are constants  $C(q')$  such that the norm of the local truncation error at order  $q'$  satisfies

$$\text{LTE}(q') = C(q') \|\phi(q'+1)\| + O(h^{q'+2}),$$

where  $\phi(k)$  is a modified divided difference of order  $k$  that is retained by IDAS (and behaves asymptotically as  $h^k$ ). Thus the local truncation errors are estimated as  $\text{ELTE}(q') = C(q') \|\phi(q'+1)\|$  to select step sizes. But the choice of

order in IDAS is based on the requirement that the scaled derivative norms,  $\|h^k y^{(k)}\|$ , are monotonically decreasing with  $k$ , for  $k$  near  $q$ . These norms are again estimated using the  $\phi(k)$ , and in fact

$$\|h^{q'+1} y^{(q'+1)}\| \approx T(q') \equiv (q' + 1) \text{ELTE}(q').$$

The step/order selection begins with a test for monotonicity that is made even *before* the local error test is performed. Namely, the order is reset to  $q' = q - 1$  if (a)  $q = 2$  and  $T(1) \leq T(2)/2$ , or (b)  $q > 2$  and  $\max\{T(q-1), T(q-2)\} \leq T(q)$ ; otherwise  $q' = q$ . Next the local error test (2.9) is performed, and if it fails, the step is redone at order  $q \leftarrow q'$  and a new step size  $h'$ . The latter is based on the  $h^{q+1}$  asymptotic behavior of  $\text{ELTE}(q)$ , and, with safety factors, is given by

$$\eta = h'/h = 0.9/[2 \text{ELTE}(q)]^{1/(q+1)}.$$

The value of  $\eta$  is adjusted so that  $\eta_{\min_{\text{ef}}} \leq \eta \leq \eta_{\text{low}}$  (by default  $\eta_{\min_{\text{ef}}} = 0.25$  and  $\eta_{\text{low}} = 0.9$ ) before setting  $h \leftarrow h' = \eta h$ . If the local error test fails a second time, IDA uses  $\eta = \eta_{\min_{\text{ef}}}$ , and on the third and subsequent failures it uses  $q = 1$  and  $\eta = \eta_{\min_{\text{ef}}}$ . After 10 failures, IDA returns with a give-up message.

As soon as the local error test has passed, the step and order for the next step may be adjusted. No order change is made if  $q' = q - 1$  from the prior test, if  $q = 5$ , or if  $q$  was increased on the previous step. Otherwise, if the last  $q + 1$  steps were taken at a constant order  $q < 5$  and a constant step size, IDAS considers raising the order to  $q + 1$ . The logic is as follows:

- a. If  $q = 1$ , then set  $q = 2$  if  $T(2) < T(1)/2$ .
- b. If  $q > 1$  then
  - set  $q \leftarrow q - 1$  if  $T(q-1) \leq \min\{T(q), T(q+1)\}$ , else
  - set  $q \leftarrow q + 1$  if  $T(q+1) < T(q)$ , otherwise
  - leave  $q$  unchanged, in this case  $T(q-1) > T(q) \leq T(q+1)$

In any case, the new step size  $h'$  is set much as before:

$$\eta = h'/h = 1/[2 \text{ELTE}(q)]^{1/(q+1)}.$$

The value of  $\eta$  is adjusted such that

- a. If  $\eta_{\min_{\text{fx}}} < \eta < \eta_{\max_{\text{fx}}}$ , set  $\eta = 1$ . The defaults are  $\eta_{\min_{\text{fx}}} = 1$  and  $\eta_{\max_{\text{fx}}} = 2$ .
- b. If  $\eta \geq \eta_{\max_{\text{fx}}}$ , the step size growth is restricted to  $\eta_{\max_{\text{fx}}} \leq \eta \leq \eta_{\max}$  with  $\eta_{\max} = 2$  by default.
- c. If  $\eta \leq \eta_{\min_{\text{fx}}}$ , the step size reduction is restricted to  $\eta_{\min} \leq \eta \leq \eta_{\text{low}}$  with  $\eta_{\min} = 0.5$  and  $\eta_{\text{low}} = 0.9$  by default.

Thus we do not increase the step size unless it can be doubled. If a step size reduction is called for, the step size will be cut by at least 10% and up to 50% for the next step. See [14] for details.

Finally  $h$  is set to  $h' = \eta h$  and  $|h|$  is restricted to  $h_{\min} \leq |h| \leq h_{\max}$  with the defaults  $h_{\min} = 0.0$  and  $h_{\max} = \infty$ .

Normally, IDAS takes steps until a user-defined output value  $t = t_{\text{out}}$  is overtaken, and then computes  $y(t_{\text{out}})$  by interpolation. However, a “one step” mode option is available, where control returns to the calling program after each step. There are also options to force IDAS not to integrate past a given stopping point  $t = t_{\text{stop}}$ .

## 2.2.4 Inequality Constraints

IDAS permits the user to impose optional inequality constraints on individual components of the solution vector  $y$ . Any of the following four constraints can be imposed:  $y_i > 0$ ,  $y_i < 0$ ,  $y_i \geq 0$ , or  $y_i \leq 0$ . The constraint satisfaction is tested after a successful nonlinear system solution. If any constraint fails, we declare a convergence failure of the nonlinear iteration and reduce the step size. Rather than cutting the step size by some arbitrary factor, IDAS estimates a new step size  $h'$  using a linear approximation of the components in  $y$  that failed the constraint test (including a safety factor of 0.9 to cover the strict inequality case). These additional constraints are also imposed during the calculation of consistent initial conditions. If a step fails to satisfy the constraints repeatedly within a step attempt then the integration is halted and an error is returned. In this case the user may need to employ other strategies as discussed in §5.1.3.2 to satisfy the inequality constraints.

## 2.3 Preconditioning

When using a nonlinear solver that requires the solution of a linear system of the form  $J\Delta y = -G$  (e.g., the default Newton iteration), IDAS makes repeated use of a linear solver. If this linear system solve is done with one of the scaled preconditioned iterative linear solvers supplied with SUNDIALS, these solvers are rarely successful if used without preconditioning; it is generally necessary to precondition the system in order to obtain acceptable efficiency. A system  $Ax = b$  can be preconditioned on the left, on the right, or on both sides. The Krylov method is then applied to a system with the matrix  $P^{-1}A$ , or  $AP^{-1}$ , or  $P_L^{-1}AP_R^{-1}$ , instead of  $A$ . However, within IDAS, preconditioning is allowed *only* on the left, so that the iterative method is applied to systems  $(P^{-1}J)\Delta y = -P^{-1}G$ . Left preconditioning is required to make the norm of the linear residual in the nonlinear iteration meaningful; in general,  $\|J\Delta y + G\|$  is meaningless, since the weights used in the WRMS-norm correspond to  $y$ .

In order to improve the convergence of the Krylov iteration, the preconditioner matrix  $P$  should in some sense approximate the system matrix  $A$ . Yet at the same time, in order to be cost-effective, the matrix  $P$  should be reasonably efficient to evaluate and solve. Finding a good point in this tradeoff between rapid convergence and low cost can be very difficult. Good choices are often problem-dependent (for example, see [17] for an extensive study of preconditioners for reaction-transport systems).

Typical preconditioners used with IDAS are based on approximations to the iteration matrix of the systems involved; in other words,  $P \approx \frac{\partial F}{\partial y} + \alpha \frac{\partial F}{\partial \dot{y}}$ , where  $\alpha$  is a scalar inversely proportional to the integration step size  $h$ . Because the Krylov iteration occurs within a nonlinear solver iteration and further also within a time integration, and since each of these iterations has its own test for convergence, the preconditioner may use a very crude approximation, as long as it captures the dominant numerical feature(s) of the system. We have found that the combination of a preconditioner with the Newton-Krylov iteration, using even a fairly poor approximation to the Jacobian, can be surprisingly superior to using the same matrix without Krylov acceleration (i.e., a modified Newton iteration), as well as to using the Newton-Krylov method with no preconditioning.

## 2.4 Rootfinding

The IDAS solver has been augmented to include a rootfinding feature. This means that, while integrating the Initial Value Problem (2.1), IDAS can also find the roots of a set of user-defined functions  $g_i(t, y, \dot{y})$  that depend on  $t$ , the solution vector  $y = y(t)$ , and its  $t$ -derivative  $\dot{y}(t)$ . The number of these root functions is arbitrary, and if more than one  $g_i$  is found to have a root in any given interval, the various root locations are found and reported in the order that they occur on the  $t$  axis, in the direction of integration.

Generally, this rootfinding feature finds only roots of odd multiplicity, corresponding to changes in sign of  $g_i(t, y(t), \dot{y}(t))$ , denoted  $g_i(t)$  for short. If a user root function has a root of even multiplicity (no sign change), it will probably be missed by IDAS. If such a root is desired, the user should reformulate the root function so that it changes sign at the desired root.

The basic scheme used is to check for sign changes of any  $g_i(t)$  over each time step taken, and then (when a sign change is found) to home in on the root (or roots) with a modified secant method [39]. In addition, each time  $g$  is computed, IDAS checks to see if  $g_i(t) = 0$  exactly, and if so it reports this as a root. However, if an exact zero of any  $g_i$  is found at a point  $t$ , IDAS computes  $g$  at  $t + \delta$  for a small increment  $\delta$ , slightly further in the direction of integration, and if any  $g_i(t + \delta) = 0$  also, IDAS stops and reports an error. This way, each time IDAS takes a time step, it is guaranteed that the values of all  $g_i$  are nonzero at some past value of  $t$ , beyond which a search for roots is to be done.

At any given time in the course of the time-stepping, after suitable checking and adjusting has been done, IDAS has an interval  $(t_{lo}, t_{hi}]$  in which roots of the  $g_i(t)$  are to be sought, such that  $t_{hi}$  is further ahead in the direction of integration, and all  $g_i(t_{lo}) \neq 0$ . The endpoint  $t_{hi}$  is either  $t_n$ , the end of the time step last taken, or the next requested output time  $t_{out}$  if this comes sooner. The endpoint  $t_{lo}$  is either  $t_{n-1}$ , or the last output time  $t_{out}$  (if this occurred within the last step), or the last root location (if a root was just located within this step), possibly adjusted slightly toward  $t_n$  if an exact zero was found. The algorithm checks  $g$  at  $t_{hi}$  for zeros and for sign changes in  $(t_{lo}, t_{hi})$ . If no sign changes are found, then either a root is reported (if some  $g_i(t_{hi}) = 0$ ) or we proceed to the next time interval (starting at  $t_{hi}$ ). If one or more sign changes were found, then a loop is entered to locate the root to within a rather tight tolerance, given by

$$\tau = 100 * U * (|t_n| + |h|) \quad (U = \text{unit roundoff}) .$$

Whenever sign changes are seen in two or more root functions, the one deemed most likely to have its root occur first is the one with the largest value of  $|g_i(t_{hi})|/|g_i(t_{hi}) - g_i(t_{lo})|$ , corresponding to the closest to  $t_{lo}$  of the secant method values. At each pass through the loop, a new value  $t_{mid}$  is set, strictly within the search interval, and the values of  $g_i(t_{mid})$  are checked. Then either  $t_{lo}$  or  $t_{hi}$  is reset to  $t_{mid}$  according to which subinterval is found to have the sign change. If there is none in  $(t_{lo}, t_{mid})$  but some  $g_i(t_{mid}) = 0$ , then that root is reported. The loop continues until  $|t_{hi} - t_{lo}| < \tau$ , and then the reported root location is  $t_{hi}$ .

In the loop to locate the root of  $g_i(t)$ , the formula for  $t_{mid}$  is

$$t_{mid} = t_{hi} - (t_{hi} - t_{lo})g_i(t_{hi})/[g_i(t_{hi}) - \alpha g_i(t_{lo})] ,$$

where  $\alpha$  a weight parameter. On the first two passes through the loop,  $\alpha$  is set to 1, making  $t_{mid}$  the secant method value. Thereafter,  $\alpha$  is reset according to the side of the subinterval (low vs high, i.e. toward  $t_{lo}$  vs toward  $t_{hi}$ ) in which the sign change was found in the previous two passes. If the two sides were opposite,  $\alpha$  is set to 1. If the two sides were the same,  $\alpha$  is halved (if on the low side) or doubled (if on the high side). The value of  $t_{mid}$  is closer to  $t_{lo}$  when  $\alpha < 1$  and closer to  $t_{hi}$  when  $\alpha > 1$ . If the above value of  $t_{mid}$  is within  $\tau/2$  of  $t_{lo}$  or  $t_{hi}$ , it is adjusted inward, such that its fractional distance from the endpoint (relative to the interval size) is between 0.1 and 0.5 (0.5 being the midpoint), and the actual distance from the endpoint is at least  $\tau/2$ .

## 2.5 Pure quadrature integration

In many applications, and most notably during the backward integration phase of an adjoint sensitivity analysis run §2.7 it is of interest to compute integral quantities of the form

$$z(t) = \int_{t_0}^t q(\tau, y(\tau), \dot{y}(\tau), p) d\tau . \quad (2.10)$$

The most effective approach to compute  $z(t)$  is to extend the original problem with the additional ODEs (obtained by applying Leibnitz's differentiation rule):

$$\dot{z} = q(t, y, \dot{y}, p) , \quad z(t_0) = 0 .$$

Note that this is equivalent to using a quadrature method based on the underlying linear multistep polynomial representation for  $y(t)$ .

This can be done at the “user level” by simply exposing to IDAS the extended DAE system (2.2) + (2.10). However, in the context of an implicit integration solver, this approach is not desirable since the nonlinear solver module will require

the Jacobian (or Jacobian-vector product) of this extended DAE. Moreover, since the additional states,  $z$ , do not enter the right-hand side of the ODE (2.10) and therefore the residual of the extended DAE system does not depend on  $z$ , it is much more efficient to treat the ODE system (2.10) separately from the original DAE system (2.2) by “taking out” the additional states  $z$  from the nonlinear system (2.4) that must be solved in the correction step of the LMM. Instead, “corrected” values  $z_n$  are computed explicitly as

$$z_n = \frac{1}{\alpha_{n,0}} \left( h_n q(t_n, y_n, \dot{y}_n, p) - \sum_{i=1}^q \alpha_{n,i} z_{n-i} \right),$$

once the new approximation  $y_n$  is available.

The quadrature variables  $z$  can be optionally included in the error test, in which case corresponding relative and absolute tolerances must be provided.

## 2.6 Forward sensitivity analysis

Typically, the governing equations of complex, large-scale models depend on various parameters, through the right-hand side vector and/or through the vector of initial conditions, as in (2.2). In addition to numerically solving the DAEs, it may be desirable to determine the sensitivity of the results with respect to the model parameters. Such sensitivity information can be used to estimate which parameters are most influential in affecting the behavior of the simulation or to evaluate optimization gradients (in the setting of dynamic optimization, parameter estimation, optimal control, etc.).

The *solution sensitivity* with respect to the model parameter  $p_i$  is defined as the vector  $s_i(t) = \partial y(t) / \partial p_i$  and satisfies the following *forward sensitivity equations* (or *sensitivity equations* for short):

$$\begin{aligned} \frac{\partial F}{\partial y} s_i + \frac{\partial F}{\partial \dot{y}} \dot{s}_i + \frac{\partial F}{\partial p_i} &= 0 \\ s_i(t_0) &= \frac{\partial y_0(p)}{\partial p_i}, \quad \dot{s}_i(t_0) = \frac{\partial \dot{y}_0(p)}{\partial p_i}, \end{aligned} \tag{2.11}$$

obtained by applying the chain rule of differentiation to the original DAEs (2.2).

When performing forward sensitivity analysis, IDAS carries out the time integration of the combined system, (2.2) and (2.11), by viewing it as a DAE system of size  $N(N_s + 1)$ , where  $N_s$  is the number of model parameters  $p_i$ , with respect to which sensitivities are desired ( $N_s \leq N_p$ ). However, major improvements in efficiency can be made by taking advantage of the special form of the sensitivity equations as linearizations of the original DAEs. In particular, the original DAE system and all sensitivity systems share the same Jacobian matrix  $J$  in (2.7).

The sensitivity equations are solved with the same linear multistep formula that was selected for the original DAEs and the same linear solver is used in the correction phase for both state and sensitivity variables. In addition, IDAS offers the option of including (*full error control*) or excluding (*partial error control*) the sensitivity variables from the local error test.

### 2.6.1 Forward sensitivity methods

In what follows we briefly describe three methods that have been proposed for the solution of the combined DAE and sensitivity system for the vector  $\hat{y} = [y, s_1, \dots, s_{N_s}]$ .

- *Staggered Direct* In this approach [25], the nonlinear system (2.4) is first solved and, once an acceptable numerical solution is obtained, the sensitivity variables at the new step are found by directly solving (2.11) after the BDF discretization is used to eliminate  $\dot{s}_i$ . Although the system matrix of the above linear system is based on exactly the same information as the matrix  $J$  in (2.7), it must be updated and factored at every step of the integration, in contrast to an evaluation of  $J$  which is updated only occasionally. For problems with many parameters (relative

to the problem size), the staggered direct method can outperform the methods described below [52]. However, the computational cost associated with matrix updates and factorizations makes this method unattractive for problems with many more states than parameters (such as those arising from semidiscretization of PDEs) and is therefore not implemented in IDAS.

- *Simultaneous Corrector* In this method [57], the discretization is applied simultaneously to both the original equations (2.2) and the sensitivity systems (2.11) resulting in an “extended” nonlinear system  $\hat{G}(\hat{y}_n) = 0$  where  $\hat{y}_n = [y_n, \dots, s_i, \dots]$ . This combined nonlinear system can be solved using a modified Newton method as in (2.6) by solving the corrector equation

$$\hat{J}[\hat{y}_{n(m+1)} - \hat{y}_{n(m)}] = -\hat{G}(\hat{y}_{n(m)}) \quad (2.12)$$

at each iteration, where

$$\hat{J} = \begin{bmatrix} J & & & & \\ J_1 & J & & & \\ J_2 & 0 & J & & \\ \vdots & \vdots & \ddots & \ddots & \\ J_{N_s} & 0 & \dots & 0 & J \end{bmatrix},$$

$J$  is defined as in (2.7), and  $J_i = (\partial/\partial y) [F_y s_i + F_{\dot{y}} \dot{s}_i + F_{p_i}]$ . It can be shown that 2-step quadratic convergence can be retained by using only the block-diagonal portion of  $\hat{J}$  in the corrector equation (2.12). This results in a decoupling that allows the reuse of  $J$  without additional matrix factorizations. However, the sum  $F_y s_i + F_{\dot{y}} \dot{s}_i + F_{p_i}$  must still be reevaluated at each step of the iterative process (2.12) to update the sensitivity portions of the residual  $\hat{G}$ .

- *Staggered corrector* In this approach [34], as in the staggered direct method, the nonlinear system (2.4) is solved first using the Newton iteration (2.6). Then, for each sensitivity vector  $\xi \equiv s_i$ , a separate Newton iteration is used to solve the sensitivity system (2.11):

$$J[\xi_{n(m+1)} - \xi_{n(m)}] = - \left[ F_y(t_n, y_n, \dot{y}_n) \xi_{n(m)} + F_{\dot{y}}(t_n, y_n, \dot{y}_n) \cdot h_n^{-1} \left( \alpha_{n,0} \xi_{n(m)} + \sum_{i=1}^q \alpha_{n,i} \xi_{n-i} \right) + F_{p_i}(t_n, y_n, \dot{y}_n) \right]. \quad (2.13)$$

In other words, a modified Newton iteration is used to solve a linear system. In this approach, the matrices  $\partial F/\partial y$ ,  $\partial F/\partial \dot{y}$  and vectors  $\partial f/\partial p_i$  need be updated only once per integration step, after the state correction phase (2.6) has converged.

IDAS implements both the simultaneous corrector method and the staggered corrector method.

An important observation is that the staggered corrector method, combined with a Krylov linear solver, effectively results in a staggered direct method. Indeed, the Krylov solver requires only the action of the matrix  $J$  on a vector, and this can be provided with the current Jacobian information. Therefore, the modified Newton procedure (2.13) will theoretically converge after one iteration.

## 2.6.2 Selection of the absolute tolerances for sensitivity variables

If the sensitivities are included in the error test, IDAS provides an automated estimation of absolute tolerances for the sensitivity variables based on the absolute tolerance for the corresponding state variable. The relative tolerance for sensitivity variables is set to be the same as for the state variables. The selection of absolute tolerances for the sensitivity variables is based on the observation that the sensitivity vector  $s_i$  will have units of  $[y]/[p_i]$ . With this, the absolute tolerance for the  $j$ -th component of the sensitivity vector  $s_i$  is set to  $\text{atol}_j/|\bar{p}_i|$ , where  $\text{atol}_j$  are the absolute tolerances for the state variables and  $\bar{p}$  is a vector of scaling factors that are dimensionally consistent with the model parameters  $p$  and give an indication of their order of magnitude. This choice of relative and absolute tolerances is equivalent to requiring that the weighted root-mean-square norm of the sensitivity vector  $s_i$  with weights based on  $s_i$



be the same as the weighted root-mean-square norm of the vector of scaled sensitivities  $\bar{s}_i = |\bar{p}_i|s_i$  with weights based on the state variables (the scaled sensitivities  $\bar{s}_i$  being dimensionally consistent with the state variables). However, this choice of tolerances for the  $s_i$  may be a poor one, and the user of IDAS can provide different values as an option.

### 2.6.3 Evaluation of the sensitivity right-hand side

There are several methods for evaluating the residual functions in the sensitivity systems (2.11): analytic evaluation, automatic differentiation, complex-step approximation, and finite differences (or directional derivatives). IDAS provides all the software hooks for implementing interfaces to automatic differentiation (AD) or complex-step approximation; future versions will include a generic interface to AD-generated functions. At the present time, besides the option for analytical sensitivity right-hand sides (user-provided), IDAS can evaluate these quantities using various finite difference-based approximations to evaluate the terms  $(\partial F/\partial y)s_i + (\partial F/\partial \dot{y})\dot{s}_i$  and  $(\partial f/\partial p_i)$ , or using directional derivatives to evaluate  $[(\partial F/\partial y)s_i + (\partial F/\partial \dot{y})\dot{s}_i + (\partial f/\partial p_i)]$ . As is typical for finite differences, the proper choice of perturbations is a delicate matter. IDAS takes into account several problem-related features: the relative DAE error tolerance  $\text{rtol}$ , the machine unit roundoff  $U$ , the scale factor  $\bar{p}_i$ , and the weighted root-mean-square norm of the sensitivity vector  $s_i$ .

Using central finite differences as an example, the two terms  $(\partial F/\partial y)s_i + (\partial F/\partial \dot{y})\dot{s}_i$  and  $\partial f/\partial p_i$  in (2.11) can be evaluated either separately:

$$\frac{\partial F}{\partial y}s_i + \frac{\partial F}{\partial \dot{y}}p\dot{s}_i \approx \frac{F(t, y + \sigma_y s_i, \dot{y} + \sigma_y \dot{s}_i, p) - F(t, y - \sigma_y s_i, \dot{y} - \sigma_y \dot{s}_i, p)}{2\sigma_y}, \quad (2.14)$$

$$\frac{\partial F}{\partial p_i} \approx \frac{F(t, y, \dot{y}, p + \sigma_i e_i) - F(t, y, \dot{y}, p - \sigma_i e_i)}{2\sigma_i}, \quad (2.15)$$

$$\sigma_i = |\bar{p}_i| \sqrt{\max(\text{rtol}, U)}, \quad \sigma_y = \frac{1}{\max(1/\sigma_i, \|s_i\|_{\text{WRMS}}/|\bar{p}_i|)}$$

or simultaneously:

$$\frac{\partial F}{\partial y}s_i + \frac{\partial F}{\partial \dot{y}}p\dot{s}_i + \frac{\partial F}{\partial p_i} \approx \frac{F(t, y + \sigma s_i, \dot{y} + \sigma \dot{s}_i, p + \sigma e_i) - F(t, y - \sigma s_i, \dot{y} - \sigma \dot{s}_i, p - \sigma e_i)}{2\sigma}, \quad (2.16)$$

$$\sigma = \min(\sigma_i, \sigma_y),$$

or by adaptively switching between (2.14) + (2.15) and (2.16), depending on the relative size of the two finite difference increments  $\sigma_i$  and  $\sigma_y$ . In the adaptive scheme, if  $\rho = \max(\sigma_i/\sigma_y, \sigma_y/\sigma_i)$ , we use separate evaluations if  $\rho > \rho_{\max}$  (an input value), and simultaneous evaluations otherwise.

These procedures for choosing the perturbations  $(\sigma_i, \sigma_y, \sigma)$  and switching between derivative formulas have also been implemented for one-sided difference formulas. Forward finite differences can be applied to  $(\partial F/\partial y)s_i + (\partial F/\partial \dot{y})\dot{s}_i$  and  $\partial F/\partial p_i$  separately, or the single directional derivative formula

$$\frac{\partial F}{\partial y}s_i + \frac{\partial F}{\partial \dot{y}}p\dot{s}_i + \frac{\partial F}{\partial p_i} \approx \frac{F(t, y + \sigma s_i, \dot{y} + \sigma \dot{s}_i, p + \sigma e_i) - F(t, y, \dot{y}, p)}{\sigma}$$

can be used. In IDAS, the default value of  $\rho_{\max} = 0$  indicates the use of the second-order centered directional derivative formula (2.16) exclusively. Otherwise, the magnitude of  $\rho_{\max}$  and its sign (positive or negative) indicates whether this switching is done with regard to (centered or forward) finite differences, respectively.



## 2.6.4 Quadratures depending on forward sensitivities

If pure quadrature variables are also included in the problem definition (see §2.5), IDAS does *not* carry their sensitivities automatically. Instead, we provide a more general feature through which integrals depending on both the states  $y$  of (2.2) and the state sensitivities  $s_i$  of (2.11) can be evaluated. In other words, IDAS provides support for computing integrals of the form:

$$\bar{z}(t) = \int_{t_0}^t \bar{q}(\tau, y(\tau), \dot{y}(\tau), s_1(\tau), \dots, s_{N_p}(\tau), p) d\tau.$$

If the sensitivities of the quadrature variables  $z$  of (2.10) are desired, these can then be computed by using:

$$\bar{q}_i = q_y s_i + q_{\dot{y}} \dot{s}_i + q_p, \quad i = 1, \dots, N_p,$$

as integrands for  $\bar{z}$ , where  $q_y$ ,  $q_{\dot{y}}$ , and  $q_p$  are the partial derivatives of the integrand function  $q$  of (2.10).

As with the quadrature variables  $z$ , the new variables  $\bar{z}$  are also excluded from any nonlinear solver phase and “corrected” values  $\bar{z}_n$  are obtained through explicit formulas.

## 2.7 Adjoint sensitivity analysis

In the *forward sensitivity approach* described in the previous section, obtaining sensitivities with respect to  $N_s$  parameters is roughly equivalent to solving an DAE system of size  $(1 + N_s)N$ . This can become prohibitively expensive, especially for large-scale problems, if sensitivities with respect to many parameters are desired. In this situation, the *adjoint sensitivity method* is a very attractive alternative, provided that we do not need the solution sensitivities  $s_i$ , but rather the gradients with respect to model parameters of a relatively few derived functionals of the solution. In other words, if  $y(t)$  is the solution of (2.2), we wish to evaluate the gradient  $dG/dp$  of

$$G(p) = \int_{t_0}^T g(t, y, p) dt, \quad (2.17)$$

or, alternatively, the gradient  $dg/dp$  of the function  $g(t, y, p)$  at the final time  $t = T$ . The function  $g$  must be smooth enough that  $\partial g/\partial y$  and  $\partial g/\partial p$  exist and are bounded.

In what follows, we only sketch the analysis for the sensitivity problem for both  $G$  and  $g$ . For details on the derivation see [24].

### 2.7.1 Sensitivity of $G(p)$

We focus first on solving the sensitivity problem for  $G(p)$  defined by (2.17). Introducing a Lagrange multiplier  $\lambda$ , we form the augmented objective function

$$I(p) = G(p) - \int_{t_0}^T \lambda^* F(t, y, \dot{y}, p) dt.$$

Since  $F(t, y, \dot{y}, p) = 0$ , the sensitivity of  $G$  with respect to  $p$  is

$$\frac{dG}{dp} = \frac{dI}{dp} = \int_{t_0}^T (g_p + g_y y_p) dt - \int_{t_0}^T \lambda^* (F_p + F_y y_p + F_{\dot{y}} \dot{y}_p) dt, \quad (2.18)$$

where subscripts on functions such as  $F$  or  $g$  are used to denote partial derivatives. By integration by parts, we have

$$\int_{t_0}^T \lambda^* F_{\dot{y}} \dot{y}_p dt = (\lambda^* F_{\dot{y}} y_p) \Big|_{t_0}^T - \int_{t_0}^T (\lambda^* F_{\dot{y}})' y_p dt,$$

where  $(\dots)'$  denotes the  $t$ -derivative. Thus equation (2.18) becomes

$$\frac{dG}{dp} = \int_{t_0}^T (g_p - \lambda^* F_p) dt - \int_{t_0}^T [-g_y + \lambda^* F_y - (\lambda^* F_{\dot{y}})'] y_p dt - (\lambda^* F_{\dot{y}} y_p) \Big|_{t_0}^T.$$

Now by requiring  $\lambda$  to satisfy

$$(\lambda^* F_{\dot{y}})' - \lambda^* F_y = -g_y, \quad (2.19)$$

we obtain

$$\frac{dG}{dp} = \int_{t_0}^T (g_p - \lambda^* F_p) dt - (\lambda^* F_{\dot{y}} y_p) \Big|_{t_0}^T. \quad (2.20)$$

Note that  $y_p$  at  $t = t_0$  is the sensitivity of the initial conditions with respect to  $p$ , which is easily obtained. To find the initial conditions (at  $t = T$ ) for the adjoint system, we must take into consideration the structure of the DAE system.

For index-0 and index-1 DAE systems, we can simply take

$$\lambda^* F_{\dot{y}} \Big|_{t=T} = 0, \quad (2.21)$$

yielding the sensitivity equation for  $dG/dp$

$$\frac{dG}{dp} = \int_{t_0}^T (g_p - \lambda^* F_p) dt + (\lambda^* F_{\dot{y}} y_p) \Big|_{t=t_0}^T. \quad (2.22)$$

This choice will not suffice for a Hessenberg index-2 DAE system. For a derivation of proper final conditions in such cases, see [24].

The first thing to notice about the adjoint system (2.19) is that there is no explicit specification of the parameters  $p$ ; this implies that, once the solution  $\lambda$  is found, the formula (2.20) can then be used to find the gradient of  $G$  with respect to any of the parameters  $p$ . The second important remark is that the adjoint system (2.19) is a terminal value problem which depends on the solution  $y(t)$  of the original IVP (2.2). Therefore, a procedure is needed for providing the states  $y$  obtained during a forward integration phase of (2.2) to IDAS during the backward integration phase of (2.19). The approach adopted in IDAS, based on *checkpointing*, is described in §2.7.3 below.

## 2.7.2 Sensitivity of $g(T, p)$

Now let us consider the computation of  $dg/dp(T)$ . From  $dg/dp(T) = (d/dT)(dG/dp)$  and equation (2.20), we have

$$\frac{dg}{dp} = (g_p - \lambda^* F_p)(T) - \int_{t_0}^T \lambda_T^* F_p dt + (\lambda_T^* F_{\dot{y}} y_p) \Big|_{t=t_0} - \frac{d(\lambda^* F_{\dot{y}} y_p)}{dT} \quad (2.23)$$

where  $\lambda_T$  denotes  $\partial\lambda/\partial T$ . For index-0 and index-1 DAEs, we obtain

$$\frac{d(\lambda^* F_{\dot{y}} y_p) \Big|_{t=T}}{dT} = 0,$$

while for a Hessenberg index-2 DAE system we have

$$\frac{d(\lambda^* F_{\dot{y}} y_p) \Big|_{t=T}}{dT} = - \frac{d(g_y^a (CB)^{-1} f_p^2)}{dt} \Big|_{t=T}.$$

The corresponding adjoint equations are

$$(\lambda_T^* F_{\dot{y}})' - \lambda_T^* F_y = 0. \quad (2.24)$$

For index-0 and index-1 DAEs (as shown above, the index-2 case is different), to find the boundary condition for this equation we write  $\lambda$  as  $\lambda(t, T)$  because it depends on both  $t$  and  $T$ . Then

$$\lambda^*(T, T)F_{\dot{y}} \Big|_{t=T} = 0.$$

Taking the total derivative, we obtain

$$(\lambda_t + \lambda_T)^*(T, T)F_{\dot{y}} \Big|_{t=T} + \lambda^*(T, T) \frac{dF_{\dot{y}}}{dt} \Big|_{t=T} = 0.$$

Since  $\lambda_t$  is just  $\dot{\lambda}$ , we have the boundary condition

$$(\lambda_T^* F_{\dot{y}}) \Big|_{t=T} = - \left[ \lambda^*(T, T) \frac{dF_{\dot{y}}}{dt} + \dot{\lambda}^* F_{\dot{y}} \right] \Big|_{t=T}.$$

For the index-one DAE case, the above relation and (2.19) yield

$$(\lambda_T^* F_{\dot{y}}) \Big|_{t=T} = [g_y - \lambda^* F_y] \Big|_{t=T}.$$

For the regular implicit ODE case,  $F_{\dot{y}}$  is invertible; thus we have  $\lambda(T, T) = 0$ , which leads to  $\lambda_T(T) = -\dot{\lambda}(T)$ . As with the final conditions for  $\lambda(T)$  in (2.19), the above selection for  $\lambda_T(T)$  is not sufficient for index-two Hessenberg DAEs (see [24] for details).

### 2.7.3 Checkpointing scheme

During the backward integration, the evaluation of the right-hand side of the adjoint system requires, at the current time, the states  $y$  which were computed during the forward integration phase. Since IDAS implements variable-step integration formulas, it is unlikely that the states will be available at the desired time and so some form of interpolation is needed. The IDAS implementation being also variable-order, it is possible that during the forward integration phase the order may be reduced as low as first order, which means that there may be points in time where only  $y$  and  $\dot{y}$  are available. These requirements therefore limit the choices for possible interpolation schemes. IDAS implements two interpolation methods: a cubic Hermite interpolation algorithm and a variable-degree polynomial interpolation method which attempts to mimic the BDF interpolant for the forward integration.

However, especially for large-scale problems and long integration intervals, the number and size of the vectors  $y$  and  $\dot{y}$  that would need to be stored make this approach computationally intractable. Thus, IDAS settles for a compromise between storage space and execution time by implementing a so-called *checkpointing scheme*. At the cost of at most one additional forward integration, this approach offers the best possible estimate of memory requirements for adjoint sensitivity analysis. To begin with, based on the problem size  $N$  and the available memory, the user decides on the number  $N_d$  of data pairs  $(y, \dot{y})$  if cubic Hermite interpolation is selected, or on the number  $N_d$  of  $y$  vectors in the case of variable-degree polynomial interpolation, that can be kept in memory for the purpose of interpolation. Then, during the first forward integration stage, after every  $N_d$  integration steps a checkpoint is formed by saving enough information (either in memory or on disk) to allow for a hot restart, that is a restart which will exactly reproduce the forward integration. In order to avoid storing Jacobian-related data at each checkpoint, a reevaluation of the iteration matrix is forced before each checkpoint. At the end of this stage, we are left with  $N_c$  checkpoints, including one at  $t_0$ . During the backward integration stage, the adjoint variables are integrated backwards from  $T$  to  $t_0$ , going from one checkpoint to the previous one. The backward integration from checkpoint  $i + 1$  to checkpoint  $i$  is preceded by a forward integration from  $i$  to  $i + 1$  during which the  $N_d$  vectors  $y$  (and, if necessary  $\dot{y}$ ) are generated and stored in memory for interpolation.

**Note:** The degree of the interpolation polynomial is always that of the current BDF order for the forward interpolation at the first point to the right of the time at which the interpolated value is sought (unless too close to the  $i$ -th checkpoint,

in which case it uses the BDF order at the right-most relevant point). However, because of the FLC BDF implementation (see §2.2), the resulting interpolation polynomial is only an approximation to the underlying BDF interpolant.

The Hermite cubic interpolation option is present because it was implemented chronologically first and it is also used by other adjoint solvers (e.g. DASPKADJOINT). The variable-degree polynomial is more memory-efficient (it requires only half of the memory storage of the cubic Hermite interpolation) and is more accurate.




Fig. 2.1: Illustration of the checkpointing algorithm for generation of the forward solution during the integration of the adjoint system.

This approach transfers the uncertainty in the number of integration steps in the forward integration phase to uncertainty in the final number of checkpoints. However,  $N_c$  is much smaller than the number of steps taken during the forward integration, and there is no major penalty for writing/reading the checkpoint data to/from a temporary file. Note that, at the end of the first forward integration stage, interpolation data are available from the last checkpoint to the end of the interval of integration. If no checkpoints are necessary ( $N_d$  is larger than the number of integration steps taken in the solution of (2.2)), the total cost of an adjoint sensitivity computation can be as low as one forward plus one backward integration. In addition, IDAS provides the capability of reusing a set of checkpoints for multiple backward integrations, thus allowing for efficient computation of gradients of several functionals (2.17).

Finally, we note that the adjoint sensitivity module in IDAS provides the necessary infrastructure to integrate backwards in time any DAE terminal value problem dependent on the solution of the IVP (2.2), including adjoint systems (2.19) or (2.24), as well as any other quadrature ODEs that may be needed in evaluating the integrals in (2.20). In particular, for DAE systems arising from semi-discretization of time-dependent PDEs, this feature allows for integration of either the discretized adjoint PDE system or the adjoint of the discretized PDE.

## 2.8 Second-order sensitivity analysis

In some applications (e.g., dynamically-constrained optimization) it may be desirable to compute second-order derivative information. Considering the DAE problem (2.2) and some model output functional  $g(y)$ , the Hessian  $d^2g/dp^2$  can be obtained in a forward sensitivity analysis setting as

$$\frac{d^2g}{dp^2} = (g_y \otimes I_{N_p}) y_{pp} + y_p^T g_{yy} y_p,$$

where  $\otimes$  is the Kronecker product. The second-order sensitivities are solution of the matrix DAE system:

$$(F_{\dot{y}} \otimes I_{N_p}) \cdot \dot{y}_{pp} + (F_y \otimes I_{N_p}) \cdot y_{pp} + (I_N \otimes \dot{y}_p^T) \cdot (F_{\dot{y}\dot{y}} \dot{y}_p + F_{y\dot{y}} y_p) + (I_N \otimes y_p^T) \cdot (F_{y\dot{y}} \dot{y}_p + F_{yy} y_p) = 0$$

$$y_{pp}(t_0) = \frac{\partial^2 y_0}{\partial p^2}, \quad \dot{y}_{pp}(t_0) = \frac{\partial^2 \dot{y}_0}{\partial p^2},$$

where  $y_p$  denotes the first-order sensitivity matrix, the solution of  $N_p$  systems (2.11), and  $y_{pp}$  is a third-order tensor. It is easy to see that, except for situations in which the number of parameters  $N_p$  is very small, the computational cost of this so-called *forward-over-forward* approach is exorbitant as it requires the solution of  $N_p + N_p^2$  additional DAE systems of the same dimension as (2.2).

---

**Note:** For the sake of simplicity in presentation, we do not include explicit dependencies of  $g$  on time  $t$  or parameters  $p$ . Moreover, we only consider the case in which the dependency of the original DAE (2.2) on the parameters  $p$  is through its initial conditions only. For details on the derivation in the general case, see [58].

---

A much more efficient alternative is to compute Hessian-vector products using a so-called *forward-over-adjoint* approach. This method is based on using the same “trick” as the one used in computing gradients of pointwise functionals with the adjoint method, namely applying a formal directional forward derivation to the gradient of (2.20) (or the equivalent one for a pointwise functional  $g(T, y(T))$ ). With that, the cost of computing a full Hessian is roughly equivalent to the cost of computing the gradient with forward sensitivity analysis. However, Hessian-vector products can be cheaply computed with one additional adjoint solve.

As an illustration, consider the ODE problem (the derivation for the general DAE case is too involved for the purposes of this discussion)

$$\dot{y} = f(t, y), \quad y(t_0) = y_0(p),$$

depending on some parameters  $p$  through the initial conditions only and consider the model functional output  $G(p) = \int_{t_0}^{t_f} g(t, y) dt$ . It can be shown that the product between the Hessian of  $G$  (with respect to the parameters  $p$ ) and some vector  $u$  can be computed as

$$\frac{\partial^2 G}{\partial p^2} u = [(\lambda^T \otimes I_{N_p}) y_{pp} u + y_p^T \mu]_{t=t_0},$$

where  $\lambda$  and  $\mu$  are solutions of

$$\begin{aligned} -\dot{\mu} &= f_y^T \mu + (\lambda^T \otimes I_n) f_{yy} s; & \mu(t_f) &= 0 \\ -\dot{\lambda} &= f_y^T \lambda + g_y^T; & \lambda(t_f) &= 0 \\ \dot{s} &= f_y s; & s(t_0) &= y_{0p} u. \end{aligned}$$

In the above equation,  $s = y_p u$  is a linear combination of the columns of the sensitivity matrix  $y_p$ . The *forward-over-adjoint* approach hinges crucially on the fact that  $s$  can be computed at the cost of a forward sensitivity analysis with respect to a single parameter (the last ODE problem above) which is possible due to the linearity of the forward sensitivity equations (2.11).

Therefore (and this is also valid for the DAE case), the cost of computing the Hessian-vector product is roughly that of two forward and two backward integrations of a system of DAEs of size  $N$ . For more details, including the corresponding formulas for a pointwise model functional output, see the work by Ozyurt and Barton [58] who discuss this problem for ODE initial value problems. As far as we know, there is no published equivalent work on DAE problems. However, the derivations given in [58] for ODE problems can be extended to DAEs with some careful consideration given to the derivation of proper final conditions on the adjoint systems, following the ideas presented in [24].

To allow the *forward-over-adjoint* approach described above, IDAS provides support for:

- the integration of multiple backward problems depending on the same underlying forward problem (2.2), and
- the integration of backward problems and computation of backward quadratures depending on both the states  $y$  and forward sensitivities (for this particular application,  $s$ ) of the original problem (2.2).



## Chapter 3

# Code Organization

The IDAS package is written in ANSI C. The following summarizes the basic structure of the package, although knowledge of this structure is not necessary for its use.

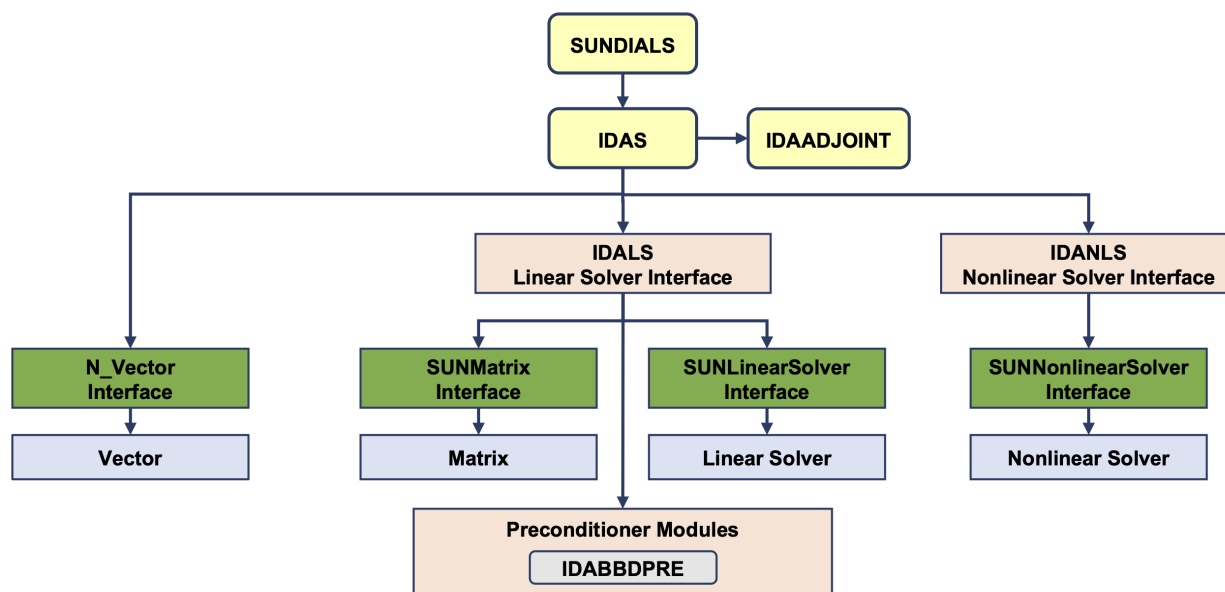


Fig. 3.1: Overall structure diagram of the IDAS package. Components specific to IDAS begin with “IDA” (IDALS, IDANLS, and IDABBDPRE), all other items correspond to generic SUNDIALS vector, matrix, and solver interfaces.

The overall organization of the IDAS package is shown in Fig. 3.1. IDAS utilizes generic linear and nonlinear solvers defined by the `SUNLinearSolver` (see §8) and `SUNNonlinearSolver` interfaces (see §9) respectively. As such, IDAS has no knowledge of the method being used to solve the linear and nonlinear systems that arise. For any given user problem, there exists a single nonlinear solver interface and, if necessary, one of the linear system solver interfaces is specified, and invoked as needed during the integration.

IDAS has a single unified linear solver interface, IDALS, supporting both direct and iterative linear solvers built using the generic `SUNLinearSolver` interface (see §8). These solvers may utilize a `SUNMatrix` object (see §7) for storing Jacobian information, or they may be matrix-free. Since IDAS can operate on any valid `SUNLinearSolver`, the set of linear solver modules available to IDAS will expand as new `SUNLinearSolver` implementations are developed.

For users employing `SUNMATRIX_DENSE` or `SUNMATRIX_BAND` Jacobian matrices, IDAS includes algorithms for their approximation through difference quotients, although the user also has the option of supplying a routine to compute

the Jacobian (or an approximation to it) directly. This user-supplied routine is required when using sparse or user-supplied Jacobian matrices.

For users employing matrix-free iterative linear solvers, IDAS includes an algorithm for the approximation by difference quotients of the product  $Jv$ . Again, the user has the option of providing routines for this operation, in two phases: setup (preprocessing of Jacobian data) and multiplication.

For preconditioned iterative methods, the preconditioning must be supplied by the user, again in two phases: setup and solve. While there is no default choice of preconditioner analogous to the difference-quotient approximation in the direct case, the references [17, 21], together with the example and demonstration programs included with IDAS, offer considerable assistance in building preconditioners.

IDA's linear solver interface consists of four primary phases, devoted to (1) memory allocation and initialization, (2) setup of the matrix data involved, (3) solution of the system, and (4) freeing of memory. The setup and solution phases are separate because the evaluation of Jacobians and preconditioners is done only periodically during the integration, and only as required to achieve convergence. The call list within the central IDAS module to each of the four associated functions is fixed, thus allowing the central module to be completely independent of the linear system method.

IDAS also provides a preconditioner module, for use with any of the Krylov iterative linear solvers. It works in conjunction with the `NVECTOR_PARALLEL` and generates a preconditioner that is a block-diagonal matrix with each block being a banded matrix.

All state information used by IDAS to solve a given problem is stored in `N_Vector` instances. There is no global data in the IDAS package, and so, in this respect, it is reentrant. State information specific to the linear and nonlinear solver are saved in the `SUNLinearSolver` and `SUNNonlinearSolver` instances respectively. The reentrancy of IDAS enables the setting where two or more problems are solved by intermixed or parallel calls to different instances of the package from within a single user program.



## Chapter 4

# Getting Started

The packages that make up SUNDIALS are built upon shared classes for vectors, matrices, and algebraic solvers. In addition, the packages all leverage some other common infrastructure, which we discuss in this section.

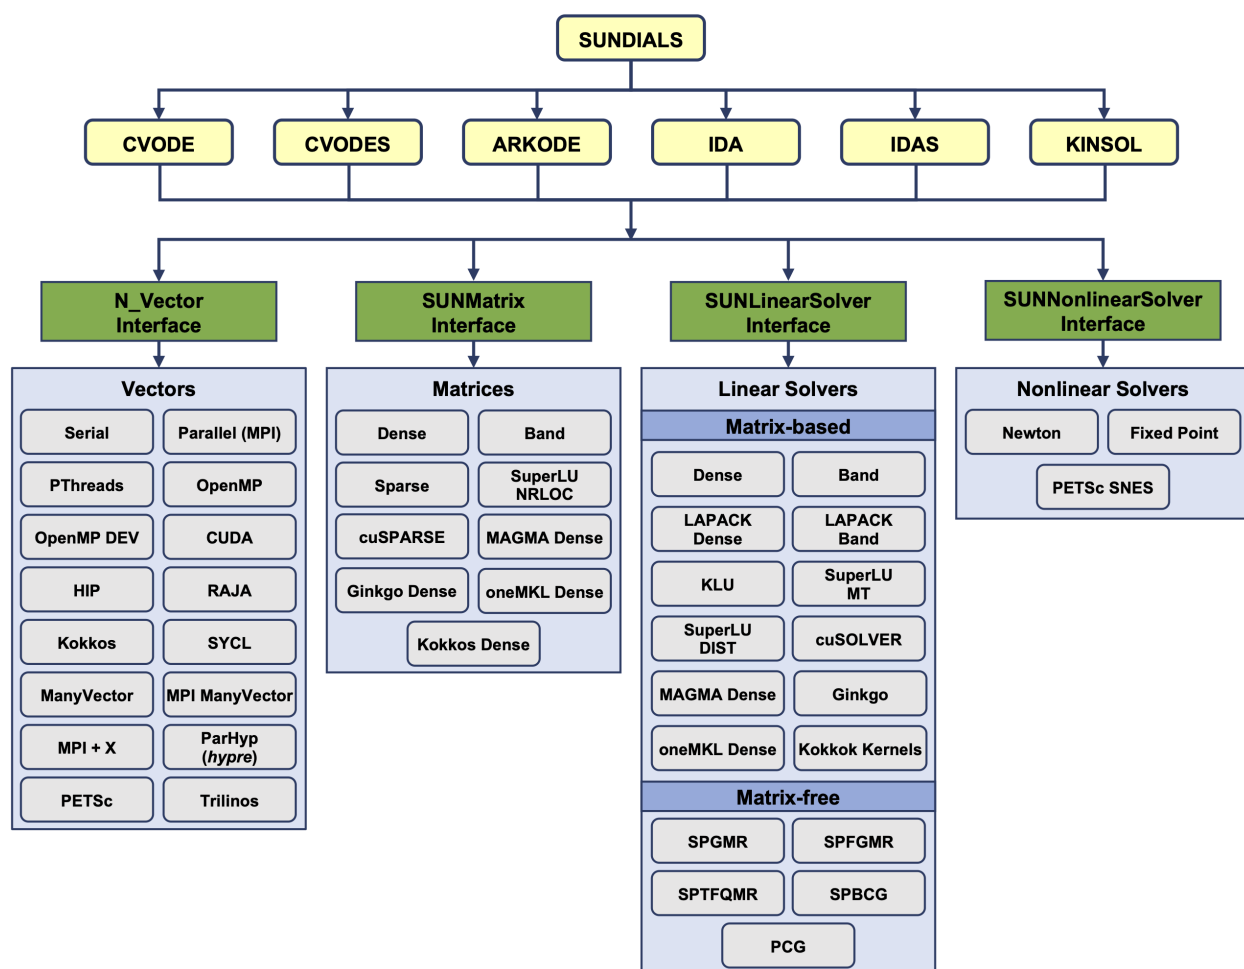


Fig. 4.1: High-level diagram of the SUNDIALS suite.

## 4.1 Data Types

SUNDIALS defines several data types in the header file `sundials_types.h`. These types are used in the SUNDIALS API and internally in SUNDIALS. It is not necessary to use these types in your application, but the type must be compatible with the SUNDIALS types in the API when calling SUNDIALS functions. The types that are defined are:

- *sunrealtype* – the floating-point type used by the SUNDIALS packages
- *sunindextype* – the integer type used for vector and matrix indices
- *sunbooleantype* – the type used for logic operations within SUNDIALS
- *SUNOutputFormat* – an enumerated type for SUNDIALS output formats
- *SUNComm* – a simple typedef to an *int* when SUNDIALS is built without MPI, or a *MPI\_Comm* when built with MPI.

### 4.1.1 Floating point types

type **sunrealtype**

The type `sunrealtype` can be `float`, `double`, or `long double`, with the default being `double`. The user can change the precision of the arithmetic used in the SUNDIALS solvers at the configuration stage (see [SUNDIALS - PRECISION](#)).

Additionally, based on the current precision, `sundials_types.h` defines `SUN_BIG_REAL` to be the largest value representable as a `sunrealtype`, `SUN_SMALL_REAL` to be the smallest value representable as a `sunrealtype`, and `SUN_UNIT_ROUNDOFF` to be the difference between 1.0 and the minimum `sunrealtype` greater than 1.0.

Within SUNDIALS, real constants are set by way of a macro called `SUN_RCONST`. It is this macro that needs the ability to branch on the definition of `sunrealtype`. In ANSI C, a floating-point constant with no suffix is stored as a `double`. Placing the suffix “F” at the end of a floating point constant makes it a `float`, whereas using the suffix “L” makes it a `long double`. For example,

```
#define A 1.0
#define B 1.0F
#define C 1.0L
```

defines A to be a `double` constant equal to 1.0, B to be a `float` constant equal to 1.0, and C to be a `long double` constant equal to 1.0. The macro call `SUN_RCONST(1.0)` automatically expands to `1.0` if `sunrealtype` is `double`, to `1.0F` if `sunrealtype` is `float`, or to `1.0L` if `sunrealtype` is `long double`. SUNDIALS uses the `SUN_RCONST` macro internally to declare all of its floating-point constants.

Additionally, SUNDIALS defines several macros for common mathematical functions *e.g.*, `fabs`, `sqrt`, `exp`, etc. in `sundials_math.h`. The macros are prefixed with `SUNR` and expand to the appropriate C function based on the `sunrealtype`. For example, the macro `SUNRabs` expands to the C function `fabs` when `sunrealtype` is `double`, `fabsf` when `sunrealtype` is `float`, and `fabsl` when `sunrealtype` is `long double`.

A user program which uses the type `sunrealtype`, the `SUN_RCONST` macro, and the `SUNR` mathematical function macros is precision-independent except for any calls to precision-specific library functions. Our example programs use `sunrealtype`, `SUN_RCONST`, and the `SUNR` macros. Users can, however, use the type `double`, `float`, or `long double` in their code (assuming that this usage is consistent with the typedef for `sunrealtype`) and call the appropriate math library functions directly. Thus, a previously existing piece of C or C++ code can use SUNDIALS without modifying the code to use `sunrealtype`, `SUN_RCONST`, or the `SUNR` macros so long as the SUNDIALS libraries are built to use the corresponding precision (see §12.2).

### 4.1.2 Integer types used for indexing

type **sunindextype**

The type `sunindextype` is used for indexing array entries in SUNDIALS modules as well as for storing the total problem size (*e.g.*, vector lengths and matrix sizes). During configuration `sunindextype` may be selected to be either a 32- or 64-bit *signed* integer with the default being 64-bit (see [SUNDIALS\\_INDEX\\_SIZE](#)).

When using a 32-bit integer the total problem size is limited to  $2^{31} - 1$  and with 64-bit integers the limit is  $2^{63} - 1$ . For users with problem sizes that exceed the 64-bit limit an advanced configuration option is available to specify the type used for `sunindextype` (see [SUNDIALS\\_INDEX\\_TYPE](#)).

A user program which uses `sunindextype` to handle indices will work with both index storage types except for any calls to index storage-specific external libraries. Our C and C++ example programs use `sunindextype`. Users can, however, use any compatible type (*e.g.*, `int`, `long int`, `int32_t`, `int64_t`, or `long long int`) in their code, assuming that this usage is consistent with the typedef for `sunindextype` on their architecture. Thus, a previously existing piece of C or C++ code can use SUNDIALS without modifying the code to use `sunindextype`, so long as the SUNDIALS libraries use the appropriate index storage type (for details see §12.2).

### 4.1.3 Boolean type

type **sunboolean**

As ANSI C89 (ISO C90) does not have a built-in boolean data type, SUNDIALS defines the type `sunboolean` as an `int`.

The advantage of using the name `sunboolean` (instead of `int`) is an increase in code readability. It also allows the programmer to make a distinction between `int` and boolean data. Variables of type `sunboolean` are intended to have only the two values: [SUNFALSE](#) or [SUNTRUE](#).

**SUNFALSE**

False (0)

**SUNTRUE**

True (1)

### 4.1.4 Output formatting type

enum **SUNOutputFormat**

The enumerated type [SUNOutputFormat](#) defines the enumeration constants for SUNDIALS output formats

enumerator **SUN\_OUTPUTFORMAT\_TABLE**

The output will be a table of values

enumerator **SUN\_OUTPUTFORMAT\_CSV**

The output will be a comma-separated list of key and value pairs *e.g.*, `key1,value1,key2,value2,...`

---

**Note:** The file `scripts/sundials_csv.py` provides python utility functions to read and output the data from a SUNDIALS CSV output file using the key and value pair format.

---

### 4.1.5 MPI types

type **SUNComm**

A simple typedef to an *int* when SUNDIALS is built without MPI, or a `MPI_Comm` when built with MPI. This type exists solely to ensure SUNDIALS can support MPI and non-MPI builds.

**SUN\_COMM\_NULL**

A macro defined as `0` when SUNDIALS is built without MPI, or as `MPI_COMM_NULL` when built with MPI.

## 4.2 The SUNContext Type

New in version 6.0.0.

All of the SUNDIALS objects (vectors, linear and nonlinear solvers, matrices, etc.) that collectively form a SUNDIALS simulation, hold a reference to a common simulation context object defined by the *SUNContext* class.

type **SUNContext**

An opaque pointer used by SUNDIALS objects for error handling, logging, profiling, etc.

Users should create a *SUNContext* object prior to any other calls to SUNDIALS library functions by calling:

*SUNErrCode* **SUNContext\_Create**(*SUNComm* comm, *SUNContext* \*sunctx)

Creates a *SUNContext* object associated with the thread of execution. The data of the *SUNContext* class is private.

#### Parameters

- **comm** – the MPI communicator or `SUN_COMM_NULL` if not using MPI.
- **sunctx** – [in,out] upon successful exit, a pointer to the newly created *SUNContext* object.

#### Returns

*SUNErrCode* indicating success or failure.

The created *SUNContext* object should be provided to the constructor routines for different SUNDIALS classes/modules e.g.,

```
SUNContext sunctx;
void* package_mem;
N_Vector x;

SUNContext_Create(SUN_COMM_NULL, &sunctx);

package_mem = CVodeCreate(..., sunctx);
package_mem = IDACreate(..., sunctx);
package_mem = KINCreate(..., sunctx);
package_mem = ARKStepCreate(..., sunctx);

x = N_VNew_<SomeVector>(..., sunctx);
```

After all other SUNDIALS code, the *SUNContext* object should be freed with a call to:

*SUNErrCode* **SUNContext\_Free**(*SUNContext* \*sunctx)

Frees the *SUNContext* object.

#### Parameters

- **sunctx** – pointer to a valid *SUNContext* object, NULL upon successful return.

**Returns**

*SUNErrCode* indicating success or failure.

**Warning:** When MPI is being used, the *SUNContext\_Free()* must be called prior to MPI\_Finalize.

The *SUNContext* API further consists of the following functions:

*SUNErrCode* **SUNContext\_GetLastError**(*SUNContext* sunctx)

Gets the last error code set by a SUNDIALS function call. The function then resets the last error code to *SUN\_SUCCESS*.

**Parameters**

- **sunctx** – a valid *SUNContext* object.

**Returns**

the last *SUNErrCode* recorded.

*SUNErrCode* **SUNContext\_PeekLastError**(*SUNContext* sunctx)

Gets the last error code set by a SUNDIALS function call. The function *does not* reset the last error code to *SUN\_SUCCESS*.

**Parameters**

- **sunctx** – a valid *SUNContext* object.

**Returns**

the last *SUNErrCode* recorded.

*SUNErrCode* **SUNContext\_PushErrorHandler**(*SUNContext* sunctx, *SUNErrorHandlerFn* err\_fn, void \*err\_user\_data)

Pushes a new *SUNErrorHandlerFn* onto the error handler stack so that it is called when an error occurs inside of SUNDIALS.

**Parameters**

- **sunctx** – a valid *SUNContext* object.
- **err\_fn** – a callback function of type *SUNErrorHandlerFn* to be pushed onto the error handler stack.
- **err\_user\_data** – a pointer that will be passed back to the callback function when it is called.

**Returns**

*SUNErrCode* indicating success or failure.

*SUNErrCode* **SUNContext\_PopErrorHandler**(*SUNContext* sunctx)

Pops the last *SUNErrorHandlerFn* off of the error handler stack.

**Parameters**

- **sunctx** – a valid *SUNContext* object.

**Returns**

*SUNErrCode* indicating success or failure.

*SUNErrCode* **SUNContext\_ClearErrorHandlers**(*SUNContext* sunctx)

Clears the entire error handler stack. After doing this it is important to push an error handler onto the stack with *SUNContext\_PushErrorHandler* otherwise errors will be ignored.

#### Parameters

- **sunctx** – a valid *SUNContext* object.

#### Returns

*SUNErrCode* indicating success or failure.

*SUNErrCode* **SUNContext\_GetProfiler**(*SUNContext* sunctx, *SUNProfiler* \*profiler)

Gets the *SUNProfiler* object associated with the *SUNContext* object.

#### Parameters

- **sunctx** – a valid *SUNContext* object.
- **profiler** – [in,out] a pointer to the *SUNProfiler* object associated with this context; will be NULL if profiling is not enabled.

#### Returns

*SUNErrCode* indicating success or failure.

*SUNErrCode* **SUNContext\_SetProfiler**(*SUNContext* sunctx, *SUNProfiler* profiler)

Sets the *SUNProfiler* object associated with the *SUNContext* object.

#### Parameters

- **sunctx** – a valid *SUNContext* object.
- **profiler** – a *SUNProfiler* object to associate with this context; this is ignored if profiling is not enabled.

#### Returns

*SUNErrCode* indicating success or failure.

*SUNErrCode* **SUNContext\_SetLogger**(*SUNContext* sunctx, *SUNLogger* logger)

Sets the *SUNLogger* object associated with the *SUNContext* object.

#### Parameters

- **sunctx** – a valid *SUNContext* object.
- **logger** – a *SUNLogger* object to associate with this context; this is ignored if logging is not enabled.

#### Returns

*SUNErrCode* indicating success or failure.

New in version 6.2.0.

*SUNErrCode* **SUNContext\_GetLogger**(*SUNContext* sunctx, *SUNLogger* \*logger)

Gets the *SUNLogger* object associated with the *SUNContext* object.

#### Parameters

- **sunctx** – a valid *SUNContext* object.
- **logger** – [in,out] a pointer to the *SUNLogger* object associated with this context; will be NULL if logging is not enabled.

#### Returns

*SUNErrCode* indicating success or failure.

New in version 6.2.0.

## 4.2.1 Implications for task-based programming and multi-threading

Applications that need to have *concurrently initialized* SUNDIALS simulations need to take care to understand the following:

#. A `SUNContext` object must only be associated with *one* SUNDIALS simulation (a solver object and its associated vectors etc.) at a time.

- Concurrently initialized is not the same as concurrently executing. Even if two SUNDIALS simulations execute sequentially, if both are initialized at the same time with the same `SUNContext`, behavior is undefined.
- It is OK to reuse a `SUNContext` object with another SUNDIALS simulation after the first simulation has completed and all of the simulation's associated objects (vectors, matrices, algebraic solvers, etc.) have been destroyed.

#. The creation and destruction of a `SUNContext` object is cheap, especially in comparison to the cost of creating/destroying a SUNDIALS solver object.

The following (incomplete) code examples demonstrate these points using CVODE as the example SUNDIALS package.

```
SUNContext sunctxs[num_threads];
int ccode_initialized[num_threads];
void* ccode_mem[num_threads];

// Create
for (int i = 0; i < num_threads; i++) {
    sunctxs[i] = SUNContext_Create(...);
    ccode_mem[i] = CCodeCreate(..., sunctxs[i]);
    ccode_initialized[i] = 0; // not yet initialized
    // set optional ccode inputs...
}

// Solve
#pragma omp parallel for
for (int i = 0; i < num_problems; i++) {
    int retval = 0;
    int tid = omp_get_thread_num();
    if (!ccode_initialized[tid]) {
        retval = CCodeInit(ccode_mem[tid], ...);
        ccode_initialized[tid] = 1;
    } else {
        retval = CCodeReInit(ccode_mem[tid], ...);
    }
    CCode(ccode_mem[i], ...);
}

// Destroy
for (int i = 0; i < num_threads; i++) {
    // get optional ccode outputs...
    CCodeFree(&ccode_mem[i]);
    SUNContext_Free(&sunctxs[i]);
}
```

Since each thread has its own unique CVODE and `SUNContext` object pair, there should be no thread-safety issues. Users should be sure that you apply the same idea to the other SUNDIALS objects needed as well (e.g. an `N_Vector`).

The variation of the above code example demonstrates another possible approach:

```
// Create, Solve, Destroy
#pragma omp parallel for
for (int i = 0; i < num_problems; i++) {
    int retval = 0;
    void* ccode_mem;
    SUNContext suncctx;

    suncctx = SUNContext_Create(...);
    ccode_mem = CCodeCreate(..., suncctx);
    retval = CCodeInit(ccode_mem, ...);

    // set optional ccode inputs...

    CCode(ccode_mem, ...);

    // get optional ccode outputs...

    CCodeFree(&ccode_mem);
    SUNContext_Free(&suncctx);
}
```

So long as the overhead of creating/destroying the CVODE object is small compared to the cost of solving the ODE, this approach is a fine alternative to the first approach since `SUNContext_Create()` and `SUNContext_Free()` are much cheaper than the CVODE create/free routines.

## 4.2.2 Convenience class for C++ Users

For C++ users a RAII safe class, `sundials::Context`, is provided:

```
namespace sundials {

class Context : public sundials::ConvertibleTo<SUNContext>
{
public:
    explicit Context(SUNComm comm = SUN_COMM_NULL)
    {
        suncctx_ = std::make_unique<SUNContext>();
        SUNContext_Create(comm, suncctx_.get());
    }

    /* disallow copy, but allow move construction */
    Context(const Context&) = delete;
    Context(Context&&)      = default;

    /* disallow copy, but allow move operators */
    Context& operator=(const Context&) = delete;
    Context& operator=(Context&&)     = default;

    SUNContext Convert() override
    {
        return *suncctx_.get();
    }
}
```

(continues on next page)



(continued from previous page)

```

}
SUNContext Convert() const override
{
    return *sunctx_.get();
}
operator SUNContext() override
{
    return *sunctx_.get();
}
operator SUNContext() const override
{
    return *sunctx_.get();
}

~Context()
{
    if (sunctx_) SUNContext_Free(sunctx_.get());
}

private:
std::unique_ptr<SUNContext> sunctx_;
};

} // namespace sundials

```

## 4.3 Error Checking

New in version 7.0.0.

Until version 7.0.0, error reporting and handling was inconsistent throughout SUNDIALS. Starting with version 7.0.0 all of SUNDIALS (the core, implementations of core modules, and packages) reports error messages through the *SUN-Logger* API. Furthermore, functions in the SUNDIALS core API (i.e., SUN or N\_V functions only) either return a *SUNErrCode*, or (if they don't return a *SUNErrCode*) they internally record an error code (if an error occurs) within the *SUNContext* for the execution stream. This “last error” is accessible via the *SUNContext\_GetLastError()* or *SUNContext\_PeekLastError()* functions.

typedef int **SUNErrCode**

Thus, in user code, SUNDIALS core API functions can be checked for errors in one of two ways:

```

SUNContext sunctx;
SUNErrCode sunerr;
N_Vector v;
int length;
sunrealtype dotprod;

// Every code that uses SUNDIALS must create a SUNContext.
sunctx = SUNContext_Create(...);

// Create a SUNDIALS serial vector.
// Some functions do not return an error code.

```

(continues on next page)

(continued from previous page)

```
// We have to check for errors in these functions using SUNContext_GetLastError.
length = 2;
v = N_VNew_Serial(length, sunctx);
sunerr = SUNContext_GetLastError(sunctx);
if (sunerr) { /* an error occurred, do something */ }

// If the function returns a SUNErrCode, we can check it directly
sunerr = N_VLinearCombination(...);
if (sunerr) { /* an error occurred, do something */ }

// Another function that does not return a SUNErrCode.
dotprod = N_VDotProd(...);
SUNContext_GetLastError(sunctx);
if (sunerr) {
    /* an error occurred, do something */
} else {
    print("dotprod = %.2f\n", dotprod);
}
```

The function `SUNGetErrMsg()` can be used to get a message describing the error code.

```
const char *SUNGetErrMsg(SUNErrCode code)
```

Returns a message describing the error code.

#### Parameters

- **code** – the error code

#### Returns

a message describing the error code.

---

**Note:** It is recommended in most cases that users check for an error after calling SUNDIALS functions. However, users concerned with getting the most performance might choose to exclude or limit these checks.

---

**Warning:** If a function returns a `SUNErrCode` then the return value is the only place the error is available i.e., these functions do not store their error code as the “last error” so it is invalid to use `SUNContext_GetLastError()` to check these functions for errors.

### 4.3.1 Error Handler Functions

When an error occurs in SUNDIALS, it calls error handler functions that have been pushed onto the error handler stack in last-in first-out order. Specific error handlers can be enabled by pushing them onto the error handler stack with the function `SUNContext_PushErrorHandler()`. They may disabled by calling `SUNContext_PopErrorHandler()` or `SUNContext_ClearErrHandlers()`. A SUNDIALS error handler function has the type

```
typedef void (*SUNErrorHandlerFn)(int line, const char *func, const char *file, const char *msg, SUNErrCode
err_code, void *err_user_data, SUNContext sunctx)
```

SUNDIALS provides a few different error handlers that can be used, or a custom one defined by the user can be provided (useful for linking SUNDIALS errors to your application’s error handling). The default error handler is `SUNLogErrorHandlerFn()` which logs an error to a specified file or `stderr` if no file is specified.

The error handlers provided in SUNDIALS are:

```
void SUNLogErrorHandlerFn(int line, const char *func, const char *file, const char *msg, SUNErrCode err_code,
                          void *err_user_data, SUNContext sunctx)
```

Logs the error that occurred using the *SUNLogger* from *sunctx*. This is the default error handler.

#### Parameters

- **line** – the line number at which the error occurred
- **func** – the function in which the error occurred
- **file** – the file in which the error occurred
- **msg** – the message to log, if this is NULL then the default error message for the error code will be used
- **err\_code** – the error code for the error that occurred
- **err\_user\_data** – the user pointer provided to *SUNContext\_PushErrorHandler()*
- **sunctx** – pointer to a valid *SUNContext* object

#### Returns

void

```
void SUNAbortErrorHandlerFn(int line, const char *func, const char *file, const char *msg, SUNErrCode err_code,
                             void *err_user_data, SUNContext sunctx)
```

Logs the error and aborts the program if an error occurred.

#### Parameters

- **line** – the line number at which the error occurred
- **func** – the function in which the error occurred
- **file** – the file in which the error occurred
- **msg** – this parameter is ignored
- **err\_code** – the error code for the error that occurred
- **err\_user\_data** – the user pointer provided to *SUNContext\_PushErrorHandler()*
- **sunctx** – pointer to a valid *SUNContext* object

#### Returns

void

```
void SUNMPIAbortErrorHandlerFn(int line, const char *func, const char *file, const char *msg, SUNErrCode
                                err_code, void *err_user_data, SUNContext sunctx)
```

Logs the error and calls `MPI_Abort` if an error occurred.

#### Parameters

- **line** – the line number at which the error occurred
- **func** – the function in which the error occurred
- **file** – the file in which the error occurred
- **msg** – this parameter is ignored
- **err\_code** – the error code for the error that occurred
- **err\_user\_data** – the user pointer provided to *SUNContext\_PushErrorHandler()*

- **sunctx** – pointer to a valid *SUNContext* object

**Returns**

void

## 4.4 Status and Error Logging

New in version 6.2.0.

SUNDIALS includes a built-in logging functionality which can be used to direct error messages, warning messages, informational output, and debugging output to specified files. This capability requires enabling both build-time and run-time options to ensure the best possible performance is achieved.

### 4.4.1 Enabling Logging

To enable logging, the CMake option *SUNDIALS\_LOGGING\_LEVEL* must be set to a value greater than 0 when configuring SUNDIALS. This option specifies the maximum desired output level. See the documentation entry for *SUNDIALS\_LOGGING\_LEVEL* for the numeric values correspond to errors, warnings, info output, and debug output where errors < warnings < info output < debug output < extra debug output. More details in regards to configuring SUNDIALS with CMake can be found in §11.

---

**Note:** As of version 7.0.0, enabling MPI in SUNDIALS enables MPI-aware logging.

---

When SUNDIALS is built with logging enabled, then the default logger (stored in the *SUNContext* object) may be configured through environment variables without any changes to user code. The available environment variables are:

```
SUNLOGGER_ERROR_FILENAME
SUNLOGGER_WARNING_FILENAME
SUNLOGGER_INFO_FILENAME
SUNLOGGER_DEBUG_FILENAME
```

These environment variables may be set to a filename string. There are two special filenames: `stdout` and `stderr`. These two filenames will result in output going to the standard output file and standard error file. The different variables may all be set to the same file, or to distinct files, or some combination there of. To disable output for one of the streams, then do not set the environment variable, or set it to an empty string.

**Warning:** A non-default logger should be created prior to any other SUNDIALS calls in order to capture all log events.

---

**Note:** If *SUNDIALS\_LOGGING\_LEVEL* was set to 1 (corresponding to error-level output) at build-time, then setting the environment variable *SUNLOGGER\_INFO\_FILENAME* will do nothing.

---

---

**Note:** Extra debugging output is turned on by setting *SUNDIALS\_LOGGING\_LEVEL* to 5. This extra output includes vector-values (so long as the *N\_Vector* used supports printing).

---

### 4.4.2 Logger API

The central piece of the Logger API is the *SUNLogger* type:

type **SUNLogger**

An opaque pointer containing logging information.

When SUNDIALS is built with logging enabled, a default logging object is stored in the *SUNContext* object and can be accessed with a call to *SUNContext\_GetLogger()*.

The enumerated type *SUNLogLevel* is used by some of the logging functions to identify the output level or file.

enum **SUNLogLevel**

The SUNDIALS logging level

enumerator **SUN\_LOGLEVEL\_ALL**

Represents all output levels

enumerator **SUN\_LOGLEVEL\_NONE**

Represents none of the output levels

enumerator **SUN\_LOGLEVEL\_ERROR**

Represents error-level logging messages

enumerator **SUN\_LOGLEVEL\_WARNING**

Represents warning-level logging messages

enumerator **SUN\_LOGLEVEL\_INFO**

Represents info-level logging messages

enumerator **SUN\_LOGLEVEL\_DEBUG**

Represents debug-level logging messages

The *SUNLogger* class provides the following methods.

int **SUNLogger\_Create**(*SUNComm* comm, int output\_rank, *SUNLogger* \*logger)

Creates a new *SUNLogger* object.

**Arguments:**

- *comm* – the MPI communicator to use, if MPI is enabled, otherwise can be *SUN\_COMM\_NULL*.
- *output\_rank* – the MPI rank used for output (can be -1 to print to all ranks).
- *logger* – [in,out] On input this is a pointer to a *SUNLogger*, on output it will point to a new *SUNLogger* instance.

**Returns:**

- Returns zero if successful, or non-zero if an error occurred.

int **SUNLogger\_CreateFromEnv**(*SUNComm* comm, *SUNLogger* \*logger)

Creates a new *SUNLogger* object and opens the output streams/files from the environment variables:

```
SUNLOGGER_ERROR_FILENAME
SUNLOGGER_WARNING_FILENAME
SUNLOGGER_INFO_FILENAME
SUNLOGGER_DEBUG_FILENAME
```

**Arguments:**

- `comm` – the MPI communicator to use, if MPI is enabled, otherwise can be `SUN_COMM_NULL`.
- **logger** – [in,out] On input this is a pointer to a *SUNLogger*, on output it will point to a new *SUNLogger* instance.

**Returns:**

- Returns zero if successful, or non-zero if an error occurred.

int **SUNLogger\_SetErrorFilename**(*SUNLogger* logger, const char \*error\_filename)

Sets the filename for error output.

**Arguments:**

- `logger` – a *SUNLogger* object.
- `error_filename` – the name of the file to use for error output.

**Returns:**

- Returns zero if successful, or non-zero if an error occurred.

int **SUNLogger\_SetWarningFilename**(*SUNLogger* logger, const char \*warning\_filename)

Sets the filename for warning output.

**Arguments:**

- `logger` – a *SUNLogger* object.
- `warning_filename` – the name of the file to use for warning output.

**Returns:**

- Returns zero if successful, or non-zero if an error occurred.

int **SUNLogger\_SetInfoFilename**(*SUNLogger* logger, const char \*info\_filename)

Sets the filename for info output.

**Arguments:**

- `logger` – a *SUNLogger* object.
- `info_filename` – the name of the file to use for info output.

**Returns:**

- Returns zero if successful, or non-zero if an error occurred.

int **SUNLogger\_SetDebugFilename**(*SUNLogger* logger, const char \*debug\_filename)

Sets the filename for debug output.

**Arguments:**

- `logger` – a *SUNLogger* object.
- `debug_filename` – the name of the file to use for debug output.

**Returns:**

- Returns zero if successful, or non-zero if an error occurred.

int **SUNLogger\_QueueMsg**(*SUNLogger* logger, *SUNLogLevel* lvl, const char \*scope, const char \*label, const char \*msg\_txt, ...)

Queues a message to the output log level.

**Arguments:**

- `logger` – a *SUNLogger* object.
- `lvl` – the message log level (i.e. error, warning, info, debug).
- `scope` – the message scope (e.g. the function name).
- `label` – the message label.
- `msg_txt` – the message text itself.
- `...` – the format string arguments

**Returns:**

- Returns zero if successful, or non-zero if an error occurred.

**Warning:** When compiling for ANSI C / C89 / C90 (and without compiler extensions), it is dangerous to pass any user input to this function because it falls back to using `sprintf` with a fixed buffer size.

It is **highly recommended** to compile with C99 or newer if your compiler does not support `snprintf` through extensions.

int **SUNLogger\_Flush**(*SUNLogger* logger, *SUNLogLevel* lvl)

Flush the message queue(s).

**Arguments:**

- `logger` – a *SUNLogger* object.
- `lvl` – the message log level (i.e. error, warning, info, debug or all).

**Returns:**

- Returns zero if successful, or non-zero if an error occurred.

int **SUNLogger\_GetOutputRank**(*SUNLogger* logger, int \*output\_rank)

Get the output MPI rank for the logger.

**Arguments:**

- `logger` – a *SUNLogger* object.
- `output_rank` – [in,out] On input this is a pointer to an int, on output it points to the int holding the output rank.

**Returns:**

- Returns zero if successful, or non-zero if an error occurred.

int **SUNLogger\_Destroy**(*SUNLogger* \*logger)

Free the memory for the *SUNLogger* object.

**Arguments:**

- `logger` – a pointer to the *SUNLogger* object.

**Returns:**

- Returns zero if successful, or non-zero if an error occur.

### 4.4.3 Example Usage

As previously mentioned, if it is enabled at build time, there is a default [SUNLogger](#) attached to a [SUNContext](#) instance when it is created. This logger can be configured using the environment variables, e.g.,

```
SUNDIALS_INFO_FILENAME=stdout ./examples/cvode/serial/cvKrylovDemo_ls
```

SUNDIALS also includes several example codes that demonstrate how to use the logging interface via the C API.

```
examples/arkode/CXX_serial/ark_analytic_sys.cpp
examples/cvode/serial/cvAdvDiff_bnd.c
examples/cvode/parallel/cvAdvDiff_diag_p.c
examples/kinsol/CXX_parallel/kin_em_p.cpp
examples/kinsol/CUDA_mpi/kin_em_mpicuda.cpp
```

## 4.5 Performance Profiling

New in version 6.0.0.

SUNDIALS includes a lightweight performance profiling layer that can be enabled at compile-time. Optionally, this profiling layer can leverage Caliper [13] for more advanced instrumentation and profiling. By default, only SUNDIALS library code is profiled. However, a public profiling API can be utilized to leverage the SUNDIALS profiler to time user code regions as well (see §4.5.2).

### 4.5.1 Enabling Profiling

To enable profiling, SUNDIALS must be built with the CMake option [SUNDIALS\\_BUILD\\_WITH\\_PROFILING](#) set to ON. To utilize Caliper support, the CMake option [ENABLE\\_CALIPER](#) must also be set to ON. More details in regards to configuring SUNDIALS with CMake can be found in §11.

When SUNDIALS is built with profiling enabled and **without Caliper**, then the environment variable `SUNPROFILER_PRINT` can be utilized to enable/disable the printing of profiler information. Setting `SUNPROFILER_PRINT=1` will cause the profiling information to be printed to stdout when the SUNDIALS simulation context is freed. Setting `SUNPROFILER_PRINT=0` will result in no profiling information being printed unless the [SUNProfiler\\_Print\(\)](#) function is called explicitly. By default, `SUNPROFILER_PRINT` is assumed to be 0. `SUNPROFILER_PRINT` can also be set to a file path where the output should be printed.

If Caliper is enabled, then users should refer to the [Caliper documentation](#) for information on getting profiler output. In most cases, this involves setting the `CALI_CONFIG` environment variable.

---

**Note:** The SUNDIALS profiler requires POSIX timers or the Windows `profileapi.h` timers.

---

**Warning:** While the SUNDIALS profiling scheme is relatively lightweight, enabling profiling can still negatively impact performance. As such, it is recommended that profiling is enabled judiciously.



## 4.5.2 Profiler API

The primary way of interacting with the SUNDIALS profiler is through the following macros:

```
SUNDIALS_MARK_FUNCTION_BEGIN(profobj)
SUNDIALS_MARK_FUNCTION_END(profobj)
SUNDIALS_WRAP_STATEMENT(profobj, name, stmt)
SUNDIALS_MARK_BEGIN(profobj, name)
SUNDIALS_MARK_END(profobj, name)
```

Additionally, in C++ applications, the follow macro is available:

```
SUNDIALS_CXX_MARK_FUNCTION(profobj)
```

These macros can be used to time specific functions or code regions. When using the \*\_BEGIN macros, it is important that a matching \*\_END macro is placed at all exit points for the scope/function. The SUNDIALS\_CXX\_MARK\_FUNCTION macro only needs to be placed at the beginning of a function, and leverages RAII to implicitly end the region.

The `profobj` argument to the macro should be a `SUNProfiler` object, i.e.

type **SUNProfiler**

An opaque pointer containing profiling information.

When SUNDIALS is built with profiling, a default profiling object is stored in the `SUNContext` object and can be accessed with a call to [`SUNContext\_GetProfiler\(\)`](#).

The `name` argument should be a unique string indicating the name of the region/function. It is important that the name given to the \*\_BEGIN macros matches the name given to the \*\_END macros.

In addition to the macros, the following methods of the `SUNProfiler` class are available.

```
int SUNProfiler_Create(SUNComm comm, const char *title, SUNProfiler *p)
```

Creates a new `SUNProfiler` object.

### Arguments:

- `comm` – the MPI communicator to use, if MPI is enabled, otherwise can be `SUN_COMM_NULL`.
- `title` – a title or description of the profiler
- `p` – [in,out] On input this is a pointer to a `SUNProfiler`, on output it will point to a new `SUNProfiler` instance

### Returns:

- Returns zero if successful, or non-zero if an error occurred

```
int SUNProfiler_Free(SUNProfiler *p)
```

Frees a `SUNProfiler` object.

### Arguments:

- `p` – [in,out] On input this is a pointer to a `SUNProfiler`, on output it will be `NULL`

### Returns:

- Returns zero if successful, or non-zero if an error occurred

```
int SUNProfiler_Begin(SUNProfiler p, const char *name)
```

Starts timing the region indicated by the `name`.

### Arguments:

- `p` – a `SUNProfiler` object
- `name` – a name for the profiling region

**Returns:**

- Returns zero if successful, or non-zero if an error occurred

int **SUNProfiler\_End**(*SUNProfiler* p, const char \*name)

Ends the timing of a region indicated by the `name`.

**Arguments:**

- `p` – a `SUNProfiler` object
- `name` – a name for the profiling region

**Returns:**

- Returns zero if successful, or non-zero if an error occurred

int **SUNProfiler\_GetElapsedTime**(*SUNProfiler* p, const char \*name, double \*time)

Get the elapsed time for the timer “`name`” in seconds.

**Arguments:**

- `p` – a `SUNProfiler` object
- `name` – the name for the profiling region of interest
- `time` – upon return, the elapsed time for the timer

**Returns:**

- Returns zero if successful, or non-zero if an error occurred

int **SUNProfiler\_GetTimerResolution**(*SUNProfiler* p, double \*resolution)

Get the timer resolution in seconds.

**Arguments:**

- `p` – a `SUNProfiler` object
- `resolution` – upon return, the resolution for the timer

**Returns:**

- Returns zero if successful, or non-zero if an error occurred

int **SUNProfiler\_Print**(*SUNProfiler* p, FILE \*fp)

Prints out a profiling summary. When constructed with an MPI comm the summary will include the average and maximum time per rank (in seconds) spent in each marked up region.

**Arguments:**

- `p` – a `SUNProfiler` object
- `fp` – the file handler to print to

**Returns:**

- Returns zero if successful, or non-zero if an error occurred

int **SUNProfiler\_Reset**(*SUNProfiler* p)

Resets the region timings and counters to zero.

**Arguments:**

- `p` – a `SUNProfiler` object

**Returns:**

- Returns zero if successful, or non-zero if an error occurred

### 4.5.3 Example Usage

The following is an excerpt from the CVODE example code `examples/cvode/serial/cvAdvDiff_bnd.c`. It is applicable to any of the SUNDIALS solver packages.

```
SUNContext ctx;
SUNProfiler profobj;

/* Create the SUNDIALS context */
retval = SUNContext_Create(SUN_COMM_NULL, &ctx);

/* Get a reference to the profiler */
retval = SUNContext_GetProfiler(ctx, &profobj);

/* ... */

SUNDIALS_MARK_BEGIN(profobj, "Integration loop");
umax = N_VMaxNorm(u);
PrintHeader(reltol, abstol, umax);
for(iout=1, tout=T1; iout <= NOUT; iout++, tout += DTOUT) {
    retval = CVode(cvode_mem, tout, u, &t, CV_NORMAL);
    umax = N_VMaxNorm(u);
    retval = CVodeGetNumSteps(cvode_mem, &nst);
    PrintOutput(t, umax, nst);
}
SUNDIALS_MARK_END(profobj, "Integration loop");
PrintFinalStats(cvode_mem); /* Print some final statistics */
```

### 4.5.4 Other Considerations

If many regions are being timed, it may be necessary to increase the maximum number of profiler entries (the default is 2560). This can be done by setting the environment variable `SUNPROFILER_MAX_ENTRIES`.

## 4.6 Getting Version Information

SUNDIALS provides additional utilities to all packages, that may be used to retrieve SUNDIALS version information at runtime.

int **SUNDIALSGetVersion**(char \*version, int len)

This routine fills a string with SUNDIALS version information.

**Arguments:**

- *version* – character array to hold the SUNDIALS version information.
- *len* – allocated length of the *version* character array.

**Return value:**

- 0 if successful
- -1 if the input string is too short to store the SUNDIALS version

**Notes:**

An array of 25 characters should be sufficient to hold the version information.

int **SUNDIALSGetVersionNumber**(int \*major, int \*minor, int \*patch, char \*label, int len)

This routine sets integers for the SUNDIALS major, minor, and patch release numbers and fills a string with the release label if applicable.

**Arguments:**

- *major* – SUNDIALS release major version number.
- *minor* – SUNDIALS release minor version number.
- *patch* – SUNDIALS release patch version number.
- *label* – string to hold the SUNDIALS release label.
- *len* – allocated length of the *label* character array.

**Return value:**

- 0 if successful
- -1 if the input string is too short to store the SUNDIALS label

**Notes:**

An array of 10 characters should be sufficient to hold the label information. If a label is not used in the release version, no information is copied to *label*.

## 4.7 Fortran Interface

SUNDIALS provides modern, Fortran 2003 based, interfaces as Fortran modules to most of the C API including:

- The SUNDIALS core types, utilities, and data structures via the `fsundials_core_mod` module.
- All of the time-stepping modules in ARKODE:
  - The `farkode_arkstep_mod`, `farkode_erkstep_mod`, `farkode_mristep_mod`, and `farkode_sprkstep_mod` modules provide interfaces to the ARKStep, ERKStep, MRISStep, and SPRKStep integrators respectively.
  - The `farkode_mod` module interfaces to the components of ARKODE which are shared by the time-stepping modules.
- CVODE via the `fcvode_mod` module.
- CVODES via the `fcvodes_mod` module.
- IDA via the `fida_mod` module.
- IDAS via the `fidas_mod` module.
- KINSOL via the `fkinsol_mod` module.

Additionally, all of the SUNDIALS base classes (*N\_Vector*, *SUNMatrix*, *SUNLinearSolver*, and *SUNNonlinearSolver*) include Fortran interface modules. A complete list of class implementations with Fortran 2003 interface modules is given in [Table 4.1](#).

An interface module can be accessed with the `use` statement, e.g.

```

use fsundials_core_mod      ! this is needed to access core SUNDIALS types, utilities, and data structures
use fcvode_mod              ! this is needed to access CVODE functions and types
use fnvector_openmp_mod     ! this is needed to access the OpenMP implementation of the N_Vector class

```

and by linking to the Fortran 2003 library in addition to the C library, e.g. `libsundials_fcore_mod.<so|a>`, `libsundials_core.<so|a>`, `libsundials_fnvecpenmp_mod.<so|a>`, `libsundials_nvecopenmp.<so|a>`, `libsundials_fcvode_mod.<so|a>` and `libsundials_cvode.<so|a>`. The use statements mirror the `#include` statements needed when using the C API.

The Fortran 2003 interfaces leverage the `iso_c_binding` module and the `bind(C)` attribute to closely follow the SUNDIALS C API (modulo language differences). The SUNDIALS classes, e.g. `N_Vector`, are interfaced as Fortran derived types, and function signatures are matched but with an `F` prepending the name, e.g. `FN_VConst` instead of `N_VConst()` or `FCvodeCreate` instead of `CVodeCreate`. Constants are named exactly as they are in the C API. Accordingly, using SUNDIALS via the Fortran 2003 interfaces looks just like using it in C. Some caveats stemming from the language differences are discussed in §4.7.2. A discussion on the topic of equivalent data types in C and Fortran 2003 is presented in §4.7.1.

Further information on the Fortran 2003 interfaces specific to the `N_Vector`, `SUNMatrix`, `SUNLinearSolver`, and `SUNNonlinearSolver` classes is given alongside the C documentation. For details on where the Fortran 2003 module (`.mod`) files and libraries are installed see §11.

The Fortran 2003 interface modules were generated with SWIG Fortran [47], a fork of SWIG. Users who are interested in the SWIG code used in the generation process should contact the SUNDIALS development team.

Table 4.1: List of SUNDIALS Fortran 2003 interface modules

Class/Module	Fortran 2003 Module Name
SUNDIALS core	fsundials_core_mod
ARKODE	farkode_mod
ARKODE::ARKSTEP	farkode_arkstep_mod
ARKODE::ERKSTEP	farkode_erkstep_mod
ARKODE::MRISTEP	farkode_mristep_mod
ARKODE::SPRKSTEP	farkode_sprkstep_mod
CVODE	fcvode_mod
CVODES	fcvodes_mod
IDA	fida_mod
IDAS	fidas_mod
KINSOL	fkinsol_mod
NVECTOR_SERIAL	fnvector_serial_mod
NVECTOR_OPENMP	fnvector_openmp_mod
NVECTOR_PTHREADS	fnvector_pthreads_mod
NVECTOR_PARALLEL	fnvector_parallel_mod
NVECTOR_PARHYP	Not interfaced
NVECTOR_PETSC	Not interfaced
NVECTOR_CUDA	Not interfaced
NVECTOR_RAJA	Not interfaced
NVECTOR_SYCL	Not interfaced
NVECTOR_MANVECTOR	fnvector_manyvector_mod
NVECTOR_MPIMANVECTOR	fnvector_mpimanyvector_mod
NVECTOR_MPIPLUSX	fnvector_mpiplusx_mod
SUNMATRIX_BAND	fsunmatrix_band_mod
SUNMATRIX_DENSE	fsunmatrix_dense_mod
SUNMATRIX_MAGMADENSE	Not interfaced
SUNMATRIX_ONEMKLDENSE	Not interfaced

continues on next page

Table 4.1 – continued from previous page

Class/Module	Fortran 2003 Module Name
SUNMATRIX_SPARSE	fsunmatrix_sparse_mod
SUNLINSOL_BAND	fsunlinsol_band_mod
SUNLINSOL_DENSE	fsunlinsol_dense_mod
SUNLINSOL_LAPACKBAND	Not interfaced
SUNLINSOL_LAPACKDENSE	Not interfaced
SUNLINSOL_MAGMADENSE	Not interfaced
SUNLINSOL_ONEMKLDENSE	Not interfaced
SUNLINSOL_KLU	fsunlinsol_klu_mod
SUNLINSOL_SLUMT	Not interfaced
SUNLINSOL_SLUDIST	Not interfaced
SUNLINSOL_SPGMR	fsunlinsol_spgmr_mod
SUNLINSOL_SPGMR	fsunlinsol_spgmr_mod
SUNLINSOL_SPBCGS	fsunlinsol_spgmr_mod
SUNLINSOL_SPTFQMR	fsunlinsol_sptfqr_mod
SUNLINSOL_PCG	fsunlinsol_pcg_mof
SUNNONLINSOL_NEWTON	fsunnonlinsol_newton_mod
SUNNONLINSOL_FIXEDPOINT	fsunnonlinsol_fixedpoint_mod
SUNNONLINSOL_PETSCSNES	Not interfaced

### 4.7.1 Data Types

Generally, the Fortran 2003 type that is equivalent to the C type is what one would expect. Primitive types map to the `iso_c_binding` type equivalent. SUNDIALS classes map to a Fortran derived type. However, the handling of pointer types is not always clear as they can depend on the parameter direction. Table 4.2 presents a summary of the type equivalencies with the parameter direction in mind.

**Warning:** Currently, the Fortran 2003 interfaces are only compatible with SUNDIALS builds where the `sunrealtype` is double-precision.

Changed in version 7.1.0: The Fortran interfaces can now be built with 32-bit `sunindextype` in addition to 64-bit `sunindextype`.

Table 4.2: C/Fortran-2003 Equivalent Types

C Type	Parameter Direction	Fortran 2003 type
SUNComm	in, inout, out, return	<code>integer(c_int)</code>
SUNErrCode	in, inout, out, return	<code>integer(c_int)</code>
double	in, inout, out, return	<code>real(c_double)</code>
int	in, inout, out, return	<code>integer(c_int)</code>
long	in, inout, out, return	<code>integer(c_long)</code>
sunboolean	in, inout, out, return	<code>integer(c_int)</code>
sunrealtype	in, inout, out, return	<code>real(c_double)</code>
sunindextype	in, inout, out, return	<code>integer(c_long)</code>
double*	in, inout, out	<code>real(c_double), dimension(*)</code>
double*	return	<code>real(c_double), pointer, dimension(:)</code>
int*	in, inout, out	<code>real(c_int), dimension(*)</code>
int*	return	<code>real(c_int), pointer, dimension(:)</code>
long*	in, inout, out	<code>real(c_long), dimension(*)</code>

continues on next page

Table 4.2 – continued from previous page

C Type	Parameter Direction	Fortran 2003 type
long*	return	real(c_long), pointer, dimension(:)
sunrealtype*	in, inout, out	real(c_double), dimension(*)
sunrealtype*	return	real(c_double), pointer, dimension(:)
sunindextype*	in, inout, out	real(c_long), dimension(*)
sunindextype*	return	real(c_long), pointer, dimension(:)
sunrealtype[]	in, inout, out	real(c_double), dimension(*)
sunindextype[]	in, inout, out	integer(c_long), dimension(*)
N_Vector	in, inout, out	type(N_Vector)
N_Vector	return	type(N_Vector), pointer
SUNMatrix	in, inout, out	type(SUNMatrix)
SUNMatrix	return	type(SUNMatrix), pointer
SUNLinearSolver	in, inout, out	type(SUNLinearSolver)
SUNLinearSolver	return	type(SUNLinearSolver), pointer
SUNNonlinearSolver	in, inout, out	type(SUNNonlinearSolver)
SUNNonlinearSolver	return	type(SUNNonlinearSolver), pointer
FILE*	in, inout, out, return	type(c_ptr)
void*	in, inout, out, return	type(c_ptr)
T**	in, inout, out, return	type(c_ptr)
T***	in, inout, out, return	type(c_ptr)
T****	in, inout, out, return	type(c_ptr)

## 4.7.2 Notable Fortran/C usage differences

While the Fortran 2003 interface to SUNDIALS closely follows the C API, some differences are inevitable due to the differences between Fortran and C. In this section, we note the most critical differences. Additionally, §4.7.1 discusses equivalencies of data types in the two languages.

### 4.7.2.1 Creating generic SUNDIALS objects

In the C API a SUNDIALS class, such as an `N_Vector`, is actually a pointer to an underlying C struct. However, in the Fortran 2003 interface, the derived type is bound to the C struct, not the pointer to the struct. For example, `type(N_Vector)` is bound to the C struct `_generic_N_Vector` not the `N_Vector` type. The consequence of this is that creating and declaring SUNDIALS objects in Fortran is nuanced. This is illustrated in the code snippets below:

C code:

```
N_Vector x;
x = N_VNew_Serial(N, sunctx);
```

Fortran code:

```
type(N_Vector), pointer :: x
x => FN_VNew_Serial(N, sunctx)
```

Note that in the Fortran declaration, the vector is a `type(N_Vector), pointer`, and that the pointer assignment operator is then used.

#### 4.7.2.2 Arrays and pointers

Unlike in the C API, in the Fortran 2003 interface, arrays and pointers are treated differently when they are return values versus arguments to a function. Additionally, pointers which are meant to be out parameters, not arrays, in the C API must still be declared as a rank-1 array in Fortran. The reason for this is partially due to the Fortran 2003 standard for C bindings, and partially due to the tool used to generate the interfaces. Regardless, the code snippets below illustrate the differences.

C code:

```
N_Vector x;
sunrealtype* xdata;
long int leniw, lenrw;

/* create a new serial vector */
x = N_VNew_Serial(N, sunctx);

/* capturing a returned array/pointer */
xdata = N_VGetArrayPointer(x)

/* passing array/pointer to a function */
N_VSetArrayPointer(xdata, x)

/* pointers that are out-parameters */
N_VSpace(x, &leniw, &lenrw);
```

Fortran code:

```
type(N_Vector), pointer :: x
real(c_double), pointer :: xdataptr(:)
real(c_double)          :: xdata(N)
integer(c_long)         :: leniw(1), lenrw(1)

! create a new serial vector
x => FN_VNew_Serial(x, sunctx)

! capturing a returned array/pointer
xdataptr => FN_VGetArrayPointer(x)

! passing array/pointer to a function
call FN_VSetArrayPointer(xdata, x)

! pointers that are out-parameters
call FN_VSpace(x, leniw, lenrw)
```



#### 4.7.2.3 Passing procedure pointers and user data

Since functions/subroutines passed to SUNDIALS will be called from within C code, the Fortran procedure must have the attribute `bind(C)`. Additionally, when providing them as arguments to a Fortran 2003 interface routine, it is required to convert a procedure's Fortran address to C with the Fortran intrinsic `c_funloc`.

Typically when passing user data to a SUNDIALS function, a user may simply cast some custom data structure as a `void*`. When using the Fortran 2003 interfaces, the same thing can be achieved. Note, the custom data structure *does not* have to be `bind(C)` since it is never accessed on the C side.

C code:

```
MyUserData *udata;
void *ccode_mem;

ierr = CCodeSetUserData(ccode_mem, udata);
```

Fortran code:

```
type(MyUserData) :: udata
type(c_ptr)      :: arkode_mem

ierr = FARKStepSetUserData(arkode_mem, c_loc(udata))
```

On the other hand, Fortran users may instead choose to store problem-specific data, e.g. problem parameters, within modules, and thus do not need the SUNDIALS-provided `user_data` pointers to pass such data back to user-supplied functions. These users should supply the `c_null_ptr` input for `user_data` arguments to the relevant SUNDIALS functions.

#### 4.7.2.4 Passing NULL to optional parameters

In the SUNDIALS C API some functions have optional parameters that a caller can pass as `NULL`. If the optional parameter is of a type that is equivalent to a Fortran `type(c_ptr)` (see §4.7.1), then a Fortran user can pass the intrinsic `c_null_ptr`. However, if the optional parameter is of a type that is not equivalent to `type(c_ptr)`, then a caller must provide a Fortran pointer that is dissociated. This is demonstrated in the code example below.

C code:

```
SUNLinearSolver LS;
N_Vector x, b;

/* SUNLinSolSolve expects a SUNMatrix or NULL as the second parameter. */
ierr = SUNLinSolSolve(LS, NULL, x, b);
```

Fortran code:

```
type(SUNLinearSolver), pointer :: LS
type(SUNMatrix), pointer      :: A
type(N_Vector), pointer       :: x, b

! Dissociate A
A => null()

! SUNLinSolSolve expects a type(SUNMatrix), pointer as the second parameter.
```

(continues on next page)

(continued from previous page)

```
! Therefore, we cannot pass a c_null_ptr, rather we pass a disassociated A.
ierr = FSUNLinSolSolve(LS, A, x, b)
```

#### 4.7.2.5 Working with N\_Vector arrays

Arrays of *N\_Vector* objects are interfaced to Fortran 2003 as an opaque type(*c\_ptr*). As such, it is not possible to directly index an array of *N\_Vector* objects returned by the *N\_Vector* “VectorArray” operations, or packages with sensitivity capabilities (CVODES and IDAS). Instead, SUNDIALS provides a utility function *FN\_VGetVecAtIndexVectorArray* wrapping *N\_VGetVecAtIndexVectorArray()*. The example below demonstrates accessing a vector in a vector array.

C code:

```
N_Vector x;
N_Vector* vecs;

/* Create an array of N_Vectors */
vecs = N_VCloneVectorArray(count, x);

/* Fill each array with ones */
for (int i = 0; i < count; ++i)
    N_VConst(vecs[i], 1.0);
```

Fortran code:

```
type(N_Vector), pointer :: x, xi
type(c_ptr)             :: vecs

! Create an array of N_Vectors
vecs = FN_VCloneVectorArray(count, x)

! Fill each array with ones
do index = 0, count-1
    xi => FN_VGetVecAtIndexVectorArray(vecs, index)
    call FN_VConst(xi, 1.d0)
enddo
```

SUNDIALS also provides the functions *N\_VSetVecAtIndexVectorArray()* and *N\_VNewVectorArray()* for working with *N\_Vector* arrays, that have corresponding Fortran interfaces *FN\_VSetVecAtIndexVectorArray* and *FN\_VNewVectorArray*, respectively. These functions are particularly useful for users of the Fortran interface to the *NVECTOR\_MANYVECTOR* or *NVECTOR\_MPIMANYVECTOR* when creating the subvector array. Both of these functions along with *N\_VGetVecAtIndexVectorArray()* (wrapped as *FN\_VGetVecAtIndexVectorArray*) are further described in §6.1.1.

#### 4.7.2.6 Providing file pointers

There are a few functions in the SUNDIALS C API which take a `FILE*` argument. Since there is no portable way to convert between a Fortran file descriptor and a C file pointer, SUNDIALS provides two utility functions for creating a `FILE*` and destroying it. These functions are defined in the module `fsundials_core_mod`.

*SUNErrCode* **SUNDIALSFileOpen**(const char \*filename, const char \*mode, FILE \*\*fp)

The function allocates a `FILE*` by calling the C function `fopen` with the provided filename and I/O mode.

##### Parameters

- **filename** – the path to the file, that should have Fortran type `character(kind=C_CHAR, len=*)`. There are two special filenames: `stdout` and `stderr` – these two filenames will result in output going to the standard output file and standard error file, respectively.
- **mode** – the I/O mode to use for the file. This should have the Fortran type `character(kind=C_CHAR, len=*)`. The string begins with one of the following characters:
  - `r` to open a text file for reading
  - `r+` to open a text file for reading/writing
  - `w` to truncate a text file to zero length or create it for writing
  - **w+ to open a text file for reading/writing or create it if it does not exist**
  - **a to open a text file for appending, see documentation of `fopen` for your system/compiler**
  - **a+ to open a text file for reading/appending, see documentation for `fopen` for your system/compiler**
- **fp** – The `FILE*` that will be open when the function returns. This should be a *type(c\_ptr)* in the Fortran.

##### Returns

A *SUNErrCode*

Usage example:

```
type(c_ptr) :: fp

! Open up the file output.log for writing
ierr = FSUNDIALSFileOpen("output.log", "w+", fp)

! The C function ARKStepPrintMem takes void* arkode_mem and FILE* fp as arguments
call FARKStepPrintMem(arkode_mem, fp)

! Close the file
ierr = FSUNDIALSFileClose(fp)
```

Changed in version 7.0.0: The function signature was updated to return a *SUNErrCode* and take a *FILE\*\** as the last input parameter rather than return a *FILE\**.

*SUNErrCode* **SUNDIALSFileClose**(FILE \*\*fp)

The function deallocates a C `FILE*` by calling the C function `fclose` with the provided pointer.

##### Parameters

- **fp** – the C FILE\* that was previously obtained from fopen. This should have the Fortran type type(c\_ptr). Note that if either stdout or stderr were opened using [SUNDIALS-FileOpen\(\)](#)

**Returns**A *SUNErrCode*

Changed in version 7.0.0: The function signature was updated to return a *SUNErrCode* and the *fp* parameter was changed from *FILE\** to *FILE\*\**.

### 4.7.3 Important notes on portability

The SUNDIALS Fortran 2003 interface *should* be compatible with any compiler supporting the Fortran 2003 ISO standard.

Upon compilation of SUNDIALS, Fortran module (.mod) files are generated for each Fortran 2003 interface. These files are highly compiler specific, and thus it is almost always necessary to compile a consuming application with the same compiler that was used to generate the modules.

### 4.7.4 Common Issues

In this subsection, we list some common issues users run into when using the Fortran interfaces.

#### Strange Segmentation Fault in User-Supplied Functions

One common issue we have seen trip up users (and even ourselves) has the symptom of segmentation fault in a user-supplied function (such as the RHS) when trying to use one of the callback arguments. For example, in the following RHS function, we will get a segfault on line 21:

```

1  integer(c_int) function ff(t, yvec, ydotvec, user_data) &
2      result(ierr) bind(C)
3
4      use, intrinsic :: iso_c_binding
5      use fsundials_nvector_mod
6      implicit none
7
8      real(c_double) :: t ! <===== Missing value attribute
9      type(N_Vector) :: yvec
10     type(N_Vector) :: ydotvec
11     type(c_ptr)     :: user_data
12
13     real(c_double) :: e
14     real(c_double) :: u, v
15     real(c_double) :: tmp1, tmp2
16     real(c_double), pointer :: yarr(:)
17     real(c_double), pointer :: ydotarr(:)
18
19     ! get N_Vector data arrays
20     yarr => FN_VGetArrayPointer(yvec)
21     ydotarr => FN_VGetArrayPointer(ydotvec) ! <===== SEGFAULTS HERE
22
23     ! extract variables
24     u = yarr(1)
25     v = yarr(2)

```

(continues on next page)

(continued from previous page)

```

26
27  ! fill in the RHS function:
28  !  [0  0]*[(-1+u^2-r(t))/(2*u)] + [          0          ]
29  !  [e -1]*[(-2+v^2-s(t))/(2*v)]   [sdot(t)/(2*vtrue(t))]
30  tmp1 = (-ONE+u*u-r(t))/(TWO*u)
31  tmp2 = (-TWO+v*v-s(t))/(TWO*v)
32  ydotarr(1) = ZERO
33  ydotarr(2) = e*tmp1 - tmp2 + sdot(t)/(TWO*vtrue(t))
34
35  ! return success
36  ierr = 0
37  return
38
39 end function

```

The subtle bug in the code causing the segfault is on line 8. It should read `real(c_double), value :: t` instead of `real(c_double) :: t` (notice the `value` attribute). Fundamental types that are passed by value in C need the `value` attribute.

## 4.8 Features for GPU Accelerated Computing

In this section, we introduce the SUNDIALS GPU programming model and highlight SUNDIALS GPU features. The model leverages the fact that all of the SUNDIALS packages interact with simulation data either through the shared vector, matrix, and solver APIs or through user-supplied callback functions. Thus, under the model, the overall structure of the user's calling program, and the way users interact with the SUNDIALS packages is similar to using SUNDIALS in CPU-only environments.

### 4.8.1 SUNDIALS GPU Programming Model

As described in [12], within the SUNDIALS GPU programming model, all control logic executes on the CPU, and all simulation data resides wherever the vector or matrix object dictates as long as SUNDIALS is in control of the program. That is, SUNDIALS will not migrate data (explicitly) from one memory space to another. Except in the most advanced use cases, it is safe to assume that data is kept resident in the GPU-device memory space. The consequence of this is that, when control is passed from the user's calling program to SUNDIALS, simulation data in vector or matrix objects must be up-to-date in the device memory space. Similarly, when control is passed from SUNDIALS to the user's calling program, the user should assume that any simulation data in vector and matrix objects are up-to-date in the device memory space. To put it succinctly, *it is the responsibility of the user's calling program to manage data coherency between the CPU and GPU-device memory spaces* unless unified virtual memory (UVM), also known as managed memory, is being utilized. Typically, the GPU-enabled SUNDIALS modules provide functions to copy data from the host to the device and vice-versa as well as support for unmanaged memory or UVM. In practical terms, the way SUNDIALS handles distinct host and device memory spaces means that *users need to ensure that the user-supplied functions, e.g. the right-hand side function, only operate on simulation data in the device memory space* otherwise extra memory transfers will be required and performance will suffer. The exception to this rule is if some form of hybrid data partitioning (achievable with the `NVECTOR_MANYVECTOR`, see §6.17) is utilized.

SUNDIALS provides many native shared features and modules that are GPU-enabled. Currently, these include the NVIDIA CUDA platform [5], AMD ROCm/HIP [2], and Intel oneAPI [3]. Table 4.3–Table 4.6 summarize the shared SUNDIALS modules that are GPU-enabled, what GPU programming environments they support, and what class of memory they support (unmanaged or UVM). Users may also supply their own GPU-enabled `N_Vector`, `SUNMatrix`, `SUNLinearSolver`, or `SUNNonlinearSolver` implementation, and the capabilities will be leveraged since SUNDIALS operates on data through these APIs.

In addition, SUNDIALS provides a memory management helper module (see §10) to support applications which implement their own memory management or memory pooling.

Table 4.3: List of SUNDIALS GPU-enabled N\_Vector Modules

Module	CUDA	ROCm/HIP	oneAPI	Unmanaged Memory	UVM
<i>NVECTOR_CUDA</i>	X			X	X
<i>NVECTOR_HIP</i>	X	X		X	X
<i>NVECTOR_SYCL</i>	X <sup>3</sup>	X <sup>3</sup>	X	X	X
<i>NVECTOR_RAJA</i>	X	X	X	X	X
<i>NVECTOR_KOKKOS</i>	X	X	X	X	X
<i>NVECTOR_OPENMPDEV</i>	X	X <sup>2</sup>	X <sup>2</sup>	X	

Table 4.4: List of SUNDIALS GPU-enabled SUNMatrix Modules

Module	CUDA	ROCm/HIP	oneAPI	Unmanaged Memory	UVM
<i>SUNMATRIX_CUSPARSE</i>	X			X	X
<i>SUNMATRIX_ONEMKLDENSE</i>	X <sup>3</sup>	X <sup>3</sup>	X	X	X
<i>SUNMATRIX_MAGMADENSE</i>	X	X		X	X
<i>SUNMATRIX_GINKGO</i>	X	X		X	X
<i>SUNMATRIX_KOKKOSDENSE</i>	X	X		X	X

Table 4.5: List of SUNDIALS GPU-enabled SUNLinearSolver Modules

Module	CUDA	ROCm/HIP	oneAPI	Unmanaged Memory	UVM
<i>SUNLINSOL_CUSOLVERSP</i>	X			X	X
<i>SUNLINSOL_ONEMKLDENSE</i>	X <sup>3</sup>	X <sup>3</sup>	X	X	X
<i>SUNLINSOL_MAGMADENSE</i>	X			X	X
<i>SUNLINSOL_GINKGO</i>	X	X		X	X
<i>SUNLINSOL_KOKKOSDENSE</i>	X	X		X	X
<i>SUNLINSOL_SPGMR</i>	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>
<i>SUNLINSOL_SPGFMR</i>	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>
<i>SUNLINSOL_SPTFQMR</i>	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>
<i>SUNLINSOL_SPBCGS</i>	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>
<i>SUNLINSOL_PCG</i>	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>

Table 4.6: List of SUNDIALS GPU-enabled SUNNonlinearSolver Modules

Module	CUDA	ROCm/HIP	oneAPI	Unmanaged Memory	UVM
<i>SUNNONLINSOL_NEWTON</i>	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>
<i>SUNNONLINSOL_FIXEDPOINT</i>	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>	X <sup>1</sup>

Notes regarding the above tables:

1. This module inherits support from the NVECTOR module used
2. Support for ROCm/HIP and oneAPI are currently untested.

3. Support for CUDA and ROCm/HIP are currently untested.

In addition, note that implicit UVM (i.e. `malloc` returning UVM) is not accounted for.

### 4.8.2 Steps for Using GPU Accelerated SUNDIALS

For any SUNDIALS package, the generalized steps a user needs to take to use GPU accelerated SUNDIALS are:

1. Utilize a GPU-enabled `N_Vector` implementation. Initial data can be loaded on the host, but must be in the device memory space prior to handing control to SUNDIALS.
2. Utilize a GPU-enabled `SUNLinearSolver` linear solver (if applicable).
3. Utilize a GPU-enabled `SUNMatrix` implementation (if using a matrix-based linear solver).
4. Utilize a GPU-enabled `SUNNonlinearSolver` nonlinear solver (if applicable).
5. Write user-supplied functions so that they use data only in the device memory space (again, unless an atypical data partitioning is used). A few examples of these functions are the right-hand side evaluation function, the Jacobian evaluation function, or the preconditioner evaluation function. In the context of CUDA and the right-hand side function, one way a user might ensure data is accessed on the device is, for example, calling a CUDA kernel, which does all of the computation, from a CPU function which simply extracts the underlying device data array from the `N_Vector` object that is passed from SUNDIALS to the user-supplied function.

Users should refer to the above tables for a complete list of GPU-enabled native SUNDIALS modules.





# Chapter 5

## Using IDAS

### 5.1 Using IDAS for IVP Solution

This chapter is concerned with the use of IDAS for the integration of DAEs.

The following sections treat the header files and the layout of the user's main program, and provide descriptions of the IDAS user-callable functions and user-supplied functions. The sample programs described in the companion document [46] may also be helpful. Those codes may be used as templates (with the removal of some lines used in testing) and are included in the IDAS package.

IDAS uses various constants for both input and output. These are defined as needed in this chapter, but for convenience are also listed separately in §13.

The user should be aware that not all `SUNLinearSolver` and `SUNMatrix` objects are compatible with all `N_Vector` implementations. Details on compatibility are given in the documentation for each `SUNMatrix` (Chapter §7) and `SUNLinearSolver` (Chapter §8) implementation. For example, `NVECTOR_PARALLEL` is not compatible with the dense, banded, or sparse `SUNMatrix` types, or with the corresponding dense, banded, or sparse `SUNLinearSolver` objects. Please check Chapters §7 and §8 to verify compatibility between these objects. In addition to that documentation, we note that the `IDABBDPRE` preconditioner can only be used with `NVECTOR_PARALLEL`. It is not recommended to use a threaded vector object with `SuperLU_MT` unless it is the `NVECTOR_OPENMP` module, and `SuperLU_MT` is also compiled with OpenMP.

#### 5.1.1 Access to library and header files

At this point, it is assumed that the installation of IDAS, following the procedure described in §11, has been completed successfully. In the proceeding text, the directories `libdir` and `incdir` are the installation library and include directories, respectively. For a default installation, these are `instdir/lib` and `instdir/include`, respectively, where `instdir` is the directory where SUNDIALS was installed.

Regardless of where the user's application program resides, its associated compilation and load commands must make reference to the appropriate locations for the library and header files required by IDAS. IDAS symbols are found in `libdir/libsundials_idas.lib`. Thus, in addition to linking to `libdir/libsundials_core.lib`, IDAS users need to link to the IDAS library. Symbols for additional SUNDIALS modules, vectors and algebraic solvers, are found in

```
<libdir>/libsundials_nvec*.lib  
<libdir>/libsundials_sunmat*.lib  
<libdir>/libsundials_sunlinsol*.lib
```

(continues on next page)

(continued from previous page)

```
<libdir>/libsundials_sunnonlinsol*.lib
<libdir>/libsundials_sunmem*.lib
```

The file extension `.lib` is typically `.so` for shared libraries and `.a` for static libraries.

The relevant header files for IDAS are located in the subdirectories `incdir/include/idas`. To use IDAS the application needs to include the header file for IDAS in addition to the SUNDIALS core header file:

```
#include <sundials/sundials_core.h> // Provides core SUNDIALS types
#include <idas/idas.h>              // IDAS provides methods for DAEs with sensitivity analysis
```

The calling program must also include an `N_Vector` implementation header file, of the form `nvector/nvector_*.h`. See §6 for the appropriate name.

If using a non-default nonlinear solver object, or when interacting with a `SUNNonlinearSolver` object directly, the calling program must also include a `SUNNonlinearSolver` implementation header file, of the form `sunnonlinsol/sunnonlinsol_*.h` where `*` is the name of the nonlinear solver (see Chapter §9 for more information).

If using a nonlinear solver that requires the solution of a linear system of the form (2.4) (e.g., the default Newton iteration), the calling program must also include a `SUNLinearSolver` implementation header file, of the form `sunlinsol/sunlinsol_*.h` where `*` is the name of the linear solver (see Chapter §8 for more information).

If the linear solver is matrix-based, the linear solver header will also include a header file of the form `sunmatrix/sunmatrix_*.h` where `*` is the name of the matrix implementation compatible with the linear solver. (see Chapter §7 for more information).

Other headers may be needed, according to the choice of preconditioner, etc. For example, in the example `idasFood-Web_kry_p` (see [46]), preconditioning is done with a block-diagonal matrix. For this, even though the `SUNLINSOL_SPGMR` linear solver is used, the header `sundials/sundials_dense.h` is included for access to the underlying generic dense matrix arithmetic routines.

## 5.1.2 A skeleton of the user's main program

The following is a skeleton of the user's main program (or calling program) for the integration of a DAE IVP. Most of the steps are independent of the `N_Vector`, `SUNMatrix`, `SUNLinearSolver`, and `SUNNonlinearSolver` implementations used. For the steps that are not, refer to Chapters §6, §7, §8, and §9 for the specific name of the function to be called or macro to be referenced.

### 1. Initialize parallel or multi-threaded environment (if appropriate)

For example, call `MPI_Init` to initialize MPI if used.

### 2. Create the SUNDIALS context object

Call `SUNContext_Create()` to allocate the `SUNContext` object.

### 3. Create the vector of initial values

Construct an `N_Vector` of initial values using the appropriate functions defined by the particular `N_Vector` implementation (see §6 for details).

For native SUNDIALS vector implementations, use a call of the form `y0 = N_VMake_***(..., ydata)` if the array containing the initial values of  $y$  already exists. Otherwise, create a new vector by making a call of the form `N_VNew_***(...)`, and then set its elements by accessing the underlying data with a call of the form `ydata = N_VGetArrayPointer(y0)`. Here, `***` is the name of the vector implementation.

For *hypre*, PETSc, and Trilinos vector wrappers, first create and initialize the underlying vector, and then create an `N_Vector` wrapper with a call of the form `y0 = N_VMake_***(yvec)`, where `yvec` is a *hypre*, PETSc, or

Trilinos vector. Note that calls like `N_VNew_***(...)` and `N_VGetArrayPointer(...)` are not available for these vector wrappers.

Set the vector `yp0` of initial conditions for  $y$  similarly.

#### 4. Create matrix object (if appropriate)

If a linear solver is required (e.g., when using the default Newton solver) and the linear solver will be a matrix-based linear solver, then a template Jacobian matrix must be created by calling the appropriate constructor defined by the particular `SUNMatrix` implementation.

For the native SUNDIALS `SUNMatrix` implementations, the matrix object may be created using a call of the form `SUN***Matrix(...)` where `***` is the name of the matrix (see §7 for details).

#### 5. Create linear solver object (if appropriate)

If a linear solver is required (e.g., when using the default Newton solver), then the desired linear solver object must be created by calling the appropriate constructor defined by the particular `SUNLinearSolver` implementation.

For any of the native SUNDIALS `SUNLinearSolver` implementations, the linear solver object may be created using a call of the form `SUNLinearSolver LS = SUNLinSol_***(...)`; where `***` is the name of the linear solver (see §8 for details).

#### 6. Create nonlinear solver object (if appropriate)

If using a non-default nonlinear solver, then the desired nonlinear solver object must be created by calling the appropriate constructor defined by the particular `SUNNonlinearSolver` implementation.

For any of the native SUNDIALS `SUNNonLinearSolver` implementations, the nonlinear solver object may be created using a call of the form `SUNNonlinearSolver NLS = SUNNonlinSol_***(...)`; where `***` is the name of the nonlinear solver (see §9 for details).

#### 7. Create IDAS object

Call `IDACreate()` to create the IDAS solver object.

#### 8. Initialize IDAS solver

Call `IDAInit()` to provide the initial condition vectors created above, set the DAE residual function, and initialize IDAS.

#### 9. Specify integration tolerances

Call one of the following functions to set the integration tolerances:

- `IDASTolerances()` to specify scalar relative and absolute tolerances.
- `IDASVTolerances()` to specify a scalar relative tolerance and a vector of absolute tolerances.
- `IDAWFTolerances()` to specify a function which sets directly the weights used in evaluating WRMS vector norms.

See §5.1.3.3 for general advice on selecting tolerances and §5.1.3.4 for advice on controlling unphysical values.

#### 10. Attach the linear solver (if appropriate)

If a linear solver was created above, initialize the IDALS linear solver interface by attaching the linear solver object (and matrix object, if applicable) with `IDASetLinearSolver()`.

#### 11. Set linear solver optional inputs (if appropriate)

See Table 5.2 for IDALS optional inputs and Chapter §8 for linear solver specific optional inputs.

#### 12. Attach nonlinear solver module (if appropriate)

If a nonlinear solver was created above, initialize the IDANLS nonlinear solver interface by attaching the nonlinear solver object with `IDASetNonlinearSolver()`.

**13. Set nonlinear solver optional inputs** (*if appropriate*)

See Table 5.3 for IDANLS optional inputs and Chapter §9 for nonlinear solver specific optional inputs. Note, solver specific optional inputs *must* be called after `IDASetNonlinearSolver()`, otherwise the optional inputs will be overridden by IDAS defaults.

**14. Specify rootfinding problem** (*optional*)

Call `IDARootInit()` to initialize a rootfinding problem to be solved during the integration of the ODE system. See Table 5.6 for relevant optional input calls.

**15. Set optional inputs**

Call `IDASet***` functions to change any optional inputs that control the behavior of IDAS from their default values. See §5.1.3.10 for details.

**16. Correct initial values** (*optional*)

Call `IDACalcIC()` to correct the initial values `y0` and `yp0` passed to `IDAInit()`. See Table 5.4 for relevant optional input calls.

**17. Advance solution in time**

For each point at which output is desired, call `ier = IDASolve(ida_mem, tout, &tret, yret, ypret, itask)`. Here `itask` specifies the return mode. The vector `yret` (which can be the same as the vector `y0` above) will contain  $y(t)$ , while the vector `ypret` (which can be the same as the vector `yp0` above) will contain  $\dot{y}(t)$ .

See `IDASolve()` for details.

**18. Get optional outputs**

Call `IDAGet***` functions to obtain optional output. See §5.1.3.12 for details.

**19. Destroy objects**

Upon completion of the integration call the following functions, as necessary, to destroy any objects created above:

- Call `N_VDestroy()` to free vector objects.
- Call `SUNMatDestroy()` to free matrix objects.
- Call `SUNLinSolFree()` to free linear solvers objects.
- Call `SUNNonlinSolFree()` to free nonlinear solvers objects.
- Call `IDAFree()` to free the memory allocated by IDAS.
- Call `SUNContext_Free()` to free the SUNDIALS context.

**20. Finalize MPI, if used**

Call `MPI_Finalize` to terminate MPI.

### 5.1.3 User-callable functions

This section describes the IDAS functions that are called by the user to setup and then solve an IVP. Some of these are required. However, starting with §5.1.3.10, the functions listed involve optional inputs/outputs or restarting, and those paragraphs may be skipped for a casual use of IDAS. In any case, refer to §5.1.2 for the correct order of these calls.

On an error, each user-callable function returns a negative value and sends an error message to the error handler routine, which prints the message on `stderr` by default. However, the user can set a file as error output or can provide his own error handler function (see §5.1.3.10).

### 5.1.3.1 IDAS initialization and deallocation functions

void \***IDACreate**(*SUNContext* sunctx)

The function *IDACreate()* instantiates an IDAS solver object.

**Arguments:**

- sunctx – the *SUNContext* object (see §4.2)

**Return value:**

- void\* pointer the IDAS solver object.

int **IDAInit**(void \*ida\_mem, *IDAResFn* res, *sunrealtype* t0, *N\_Vector* y0, *N\_Vector* yp0)

The function *IDAInit()* provides required problem and solution specifications, allocates internal memory, and initializes IDAS.

**Arguments:**

- ida\_mem – pointer to the IDAS solver object.
- res – is the function which computes the residual function  $F(t, y, \dot{y})$  for the DAE. For full details see *IDAResFn*.
- t0 – is the initial value of  $t$ .
- y0 – is the initial value of  $y$ .
- yp0 – is the initial value of  $\dot{y}$ .

**Return value:**

- IDA\_SUCCESS – The call was successful.
- IDA\_MEM\_NULL – The ida\_mem argument was NULL.
- IDA\_MEM\_FAIL – A memory allocation request has failed.
- IDA\_ILL\_INPUT – An input argument to *IDAInit()* has an illegal value.

**Notes:**

If an error occurred, *IDAInit()* also sends an error message to the error handler function.

void **IDAFree**(void \*\*ida\_mem)

The function *IDAFree()* frees the pointer allocated by a previous call to *IDACreate()*.

**Arguments:**

- ida\_mem – pointer to the IDAS solver object.

**Return value:**

- void

### 5.1.3.2 IDAS tolerance specification functions

One of the following three functions must be called to specify the integration tolerances (or directly specify the weights used in evaluating WRMS vector norms). Note that this call must be made after the call to [IDAInit\(\)](#).

int **IDASStolerances**(void \*ida\_mem, *sunrealtype* reltol, *sunrealtype* abstol)

The function [IDASStolerances\(\)](#) specifies scalar relative and absolute tolerances.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `reltol` – is the scalar relative error tolerance.
- `abstol` – is the scalar absolute error tolerance.

**Return value:**

- `IDA_SUCCESS` – The call was successful.
- `IDA_MEM_NULL` – The `ida_mem` argument was NULL.
- `IDA_NO_MALLOC` – The allocation function [IDAInit\(\)](#) has not been called.
- `IDA_ILL_INPUT` – One of the input tolerances was negative.

int **IDASVtolerances**(void \*ida\_mem, *sunrealtype* reltol, *N\_Vector* abstol)

The function [IDASVtolerances\(\)](#) specifies scalar relative tolerance and vector absolute tolerances.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `reltol` – is the scalar relative error tolerance.
- `abstol` – is the vector of absolute error tolerances.

**Return value:**

- `IDA_SUCCESS` – The call was successful.
- `IDA_MEM_NULL` – The `ida_mem` argument was NULL.
- `IDA_NO_MALLOC` – The allocation function [IDAInit\(\)](#) has not been called.
- `IDA_ILL_INPUT` – The relative error tolerance was negative or the absolute tolerance vector had a negative component.

**Notes:**

This choice of tolerances is important when the absolute error tolerance needs to be different for each component of the state vector  $y$ .

int **IDAWFtolerances**(void \*ida\_mem, *IDAEvtFn* efun)

The function [IDAWFtolerances\(\)](#) specifies a user-supplied function `efun` that sets the multiplicative error weights  $W_i$  for use in the weighted RMS norm, which are normally defined by (2.5).

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object. [IDACreate\(\)](#)
- `efun` – is the function which defines the ewt vector. For full details see [IDAEwtFn](#).

**Return value:**

- `IDA_SUCCESS` – The call was successful.
- `IDA_MEM_NULL` – The `ida_mem` argument was NULL.

- IDA\_NO\_MALLOC – The allocation function `IDAInit()` has not been called.

### 5.1.3.3 General advice on choice of tolerances

For many users, the appropriate choices for tolerance values in `reltol` and `abstol` are a concern. The following pieces of advice are relevant.

1. The scalar relative tolerance `reltol` is to be set to control relative errors. So `reltol` of  $10^{-4}$  means that errors are controlled to .01%. We do not recommend using `reltol` larger than  $10^{-3}$ . On the other hand, `reltol` should not be so small that it is comparable to the unit roundoff of the machine arithmetic (generally around  $10^{-15}$ ).
2. The absolute tolerances `abstol` (whether scalar or vector) need to be set to control absolute errors when any components of the solution vector `y` may be so small that pure relative error control is meaningless. For example, if `y[i]` starts at some nonzero value, but in time decays to zero, then pure relative error control on `y[i]` makes no sense (and is overly costly) after `y[i]` is below some noise level. Then `abstol` (if a scalar) or `abstol[i]` (if a vector) needs to be set to that noise level. If the different components have different noise levels, then `abstol` should be a vector. See the example `idaRoberts_dns` in the IDAS package, and the discussion of it in the IDAS Examples document [46]. In that problem, the three components vary between 0 and 1, and have different noise levels; hence the `abstol` vector. It is impossible to give any general advice on `abstol` values, because the appropriate noise levels are completely problem-dependent. The user or modeler hopefully has some idea as to what those noise levels are.
3. Finally, it is important to pick all the tolerance values conservatively, because they control the error committed on each individual time step. The final (global) errors are some sort of accumulation of those per-step errors. A good rule of thumb is to reduce the tolerances by a factor of .01 from the actual desired limits on errors. So if you want .01% accuracy (globally), a good choice is to is a `reltol` of  $10^{-6}$ . But in any case, it is a good idea to do a few experiments with the tolerances to see how the computed solution values vary as tolerances are reduced.

### 5.1.3.4 Advice on controlling unphysical negative values

In many applications, some components in the true solution are always positive or non-negative, though at times very small. In the numerical solution, however, small negative (hence unphysical) values can then occur. In most cases, these values are harmless, and simply need to be controlled, not eliminated. The following pieces of advice are relevant.

1. The way to control the size of unwanted negative computed values is with tighter absolute tolerances. Again this requires some knowledge of the noise level of these components, which may or may not be different for different components. Some experimentation may be needed.
2. If output plots or tables are being generated, and it is important to avoid having negative numbers appear there (for the sake of avoiding a long explanation of them, if nothing else), then eliminate them, but only in the context of the output medium. Then the internal values carried by the solver are unaffected. Remember that a small negative value in `yret` returned by IDAS, with magnitude comparable to `abstol` or less, is equivalent to zero as far as the computation is concerned.
3. The user's residual function `res` should never change a negative value in the solution vector `yy` to a non-negative value, as a "solution" to this problem. This can cause instability. If the `res` routine cannot tolerate a zero or negative value (e.g., because there is a square root or log of it), then the offending value should be changed to zero or a tiny positive number in a temporary variable (not in the input `yy` vector) for the purposes of computing  $F(t, y, \dot{y})$ .
4. IDAS provides the option of enforcing positivity or non-negativity on components. Also, such constraints can be enforced by use of the recoverable error return feature in the user-supplied residual function. However, because these options involve some extra overhead cost, they should only be exercised if the use of absolute tolerances to control the computed values is unsuccessful.



### 5.1.3.5 Linear solver interface functions

As previously explained, if the nonlinear solver requires the solution of linear systems of the form (2.6), e.g., the default Newton solver, then the solution of these linear systems is handled with the IDALS linear solver interface. This interface supports all valid `SUNLinearSolver` objects. Here, a matrix-based `SUNLinearSolver` utilizes `SUNMatrix` objects to store the Jacobian matrix  $J = \frac{\partial F}{\partial y} + \alpha \frac{\partial F}{\partial \dot{y}}$  and factorizations used throughout the solution process. Conversely, matrix-free `SUNLinearSolver` object instead use iterative methods to solve the linear systems of equations, and only require the *action* of the Jacobian on a vector,  $Jv$ .

With most iterative linear solvers, preconditioning can be done on the left only, on the right only, on both the left and the right, or not at all. The exceptions to this rule are SPFGMR that supports right preconditioning only and PCG that performs symmetric preconditioning. However, in IDAS only left preconditioning is supported. For the specification of a preconditioner, see the iterative linear solver sections in §5.1.3.10 and §5.1.4. A preconditioner matrix  $P$  must approximate the Jacobian  $J$ , at least crudely.

To attach a generic linear solver to IDAS, after the call to `IDACreate()` but before any calls to `IDASolve()`, the user's program must create the appropriate `SUNLinearSolver` object and call the function `IDASetLinearSolver()`. To create the `SUNLinearSolver` object, the user may call one of the SUNDIALS-packaged `SUNLinearSolver` constructors via a call of the form

```
SUNLinearSolver LS = SUNLinSol_*(...);
```

Alternately, a user-supplied `SUNLinearSolver` object may be created and used instead. The use of each of the generic linear solvers involves certain constants, functions and possibly some macros, that are likely to be needed in the user code. These are available in the corresponding header file associated with the specific `SUNMatrix` or `SUNLinearSolver` object in question, as described in Chapters §7 and §8.

Once this solver object has been constructed, the user should attach it to IDAS via a call to `IDASetLinearSolver()`. The first argument passed to this function is the IDAS memory pointer returned by `IDACreate()`; the second argument is the desired `SUNLinearSolver` object to use for solving systems. The third argument is an optional `SUNMatrix` object to accompany matrix-based `SUNLinearSolver` inputs (for matrix-free linear solvers, the third argument should be NULL). A call to this function initializes the IDALS linear solver interface, linking it to the main IDAS integrator, and allows the user to specify additional parameters and routines pertinent to their choice of linear solver.

```
int IDASetLinearSolver(void *ida_mem, SUNLinearSolver LS, SUNMatrix J)
```

The function `IDASetLinearSolver()` attaches a `SUNLinearSolver` object LS and corresponding template Jacobian `SUNMatrix` object J (if applicable) to IDAS, initializing the IDALS linear solver interface.

#### Arguments:

- `ida_mem` – pointer to the IDAS solver object.
- `LS` – `SUNLinearSolver` object to use for solving linear systems of the form (2.6).
- `J` – `SUNMatrix` object for used as a template for the Jacobian or NULL if not applicable.

#### Return value:

- `IDALS_SUCCESS` – The IDALS initialization was successful.
- `IDALS_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDALS_ILL_INPUT` – The IDALS interface is not compatible with the LS or J input objects or is incompatible with the `N_Vector` object passed to `IDAINIT()`.
- `IDALS_SUNLS_FAIL` – A call to the LS object failed.
- `IDALS_MEM_FAIL` – A memory allocation request failed.

#### Notes:

If LS is a matrix-based linear solver, then the template Jacobian matrix J will be used in the solve process,



so if additional storage is required within the `SUNMatrix` object (e.g., for factorization of a banded matrix), ensure that the input object is allocated with sufficient size (see the documentation of the particular `SUNMatrix` in Chapter §7 for further information).

New in version 3.0.0: Replaces the deprecated function `IDADlsSetLinearSolver`.

### 5.1.3.6 Nonlinear solver interface function

By default IDAS uses the `SUNNonlinearSolver` implementation of Newton's method (see §9.3). To attach a different nonlinear solver in IDAS, the user's program must create a `SUNNonlinearSolver` object by calling the appropriate constructor routine. The user must then attach the `SUNNonlinearSolver` object to IDAS by calling `IDASetNonlinearSolver()`.

When changing the nonlinear solver in IDAS, `IDASetNonlinearSolver()` must be called after `IDAInit()`. If any calls to `IDASolve()` have been made, then IDAS will need to be reinitialized by calling `IDAREInit()` to ensure that the nonlinear solver is initialized correctly before any subsequent calls to `IDASolve()`.

The first argument passed to `IDASetNonlinearSolver()` is the IDAS memory pointer returned by `IDACreate()` and the second argument is the `SUNNonlinearSolver` object to use for solving the nonlinear system (2.4). A call to this function attaches the nonlinear solver to the main IDAS integrator. We note that at present, the `SUNNonlinearSolver` object *must be of type* `SUNNONLINEARSOLVER_ROOTFIND`.

int `IDASetNonlinearSolver`(void \*ida\_mem, *SUNNonlinearSolver* NLS)

The function `IDASetNonlinearSolver()` attaches a `SUNNonlinearSolver` object (NLS) to IDAS.

#### Arguments:

- `ida_mem` – pointer to the IDAS solver object.
- `NLS` – `SUNNonlinearSolver` object to use for solving nonlinear systems.

#### Return value:

- `IDA_SUCCESS` – The nonlinear solver was successfully attached.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_ILL_INPUT` – The `SUNNonlinearSolver` object is NULL, does not implement the required nonlinear solver operations, is not of the correct type, or the residual function, convergence test function, or maximum number of nonlinear iterations could not be set.

#### Notes:

When forward sensitivity analysis capabilities are enabled and the `IDA_STAGGERED` corrector method is used this function sets the nonlinear solver method for correcting state variables (see §5.4.2.3 for more details).

### 5.1.3.7 Initial condition calculation function

`IDACalcIC()` calculates corrected initial conditions for the DAE system for certain index-one problems including a class of systems of semi-implicit form (see §2.2 and [19]). It uses a Newton iteration combined with a linesearch algorithm. Calling `IDACalcIC()` is optional. It is only necessary when the initial conditions do not satisfy the given system. Thus if  $y_0$  and  $yp_0$  are known to satisfy  $F(t_0, y_0, \dot{y}_0) = 0$ , then a call to `IDACalcIC()` is generally *not* necessary.

A call to the function `IDACalcIC()` must be preceded by successful calls to `IDACreate()` and `IDAInit()` (or `IDAREInit()`), and by a successful call to the linear system solver specification function. The call to `IDACalcIC()` should precede the call(s) to `IDASolve()` for the given problem.

int **IDACalcIC**(void \*ida\_mem, int icopt, *sunrealtype* tout1)

The function **IDACalcIC()** corrects the initial values  $y_0$  and  $yp_0$  at time  $t_0$ .

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `icopt` – is one of the following two options for the initial condition calculation.
  - `IDA_YA_YDP_INIT` directs **IDACalcIC()** to compute the algebraic components of  $y$  and differential components of  $\dot{y}$ , given the differential components of  $y$ . This option requires that the `N_Vector id` was set through **IDASetId()**, specifying the differential and algebraic components.
  - `IDA_Y_INIT` directs **IDACalcIC()** to compute all components of  $y$ , given  $\dot{y}$ . In this case, `id` is not required.
- `tout1` – is the first value of  $t$  at which a solution will be requested (from **IDASolve()**). This value is needed here only to determine the direction of integration and rough scale in the independent variable  $t$ .

**Return value:**

- `IDA_SUCCESS` – **IDACalcIC()** succeeded.
- `IDA_MEM_NULL` – The argument `ida_mem` was NULL.
- `IDA_NO_MALLOC` – The allocation function **IDAINIT()** has not been called.
- `IDA_ILL_INPUT` – One of the input arguments was illegal.
- `IDA_LSETUP_FAIL` – The linear solver's setup function failed in an unrecoverable manner.
- `IDA_LINIT_FAIL` – The linear solver's initialization function failed.
- `IDA_LSOLVE_FAIL` – The linear solver's solve function failed in an unrecoverable manner.
- `IDA_BAD_EWT` – Some component of the error weight vector is zero (illegal), either for the input value of  $y_0$  or a corrected value.
- `IDA_FIRST_RES_FAIL` – The user's residual function returned a recoverable error flag on the first call, but **IDACalcIC()** was unable to recover.
- `IDA_RES_FAIL` – The user's residual function returned a nonrecoverable error flag.
- `IDA_NO_RECOVERY` – The user's residual function, or the linear solver's setup or solve function had a recoverable error, but **IDACalcIC()** was unable to recover.
- `IDA_CONSTR_FAIL` – **IDACalcIC()** was unable to find a solution satisfying the inequality constraints.
- `IDA_LINESEARCH_FAIL` – The linesearch algorithm failed to find a solution with a step larger than `steptol` in weighted RMS norm, and within the allowed number of backtracks.
- `IDA_CONV_FAIL` – **IDACalcIC()** failed to get convergence of the Newton iterations.

**Notes:**

**IDACalcIC()** will correct the values of  $y(t_0)$  and  $\dot{y}(t_0)$  which were specified in the previous call to **IDAINIT()** or **IDAREINIT()**. To obtain the corrected values, call **IDAGETCONSISTENTIC()**.

### 5.1.3.8 Rootfinding initialization function

While solving the IVP, IDAS has the capability to find the roots of a set of user-defined functions. To activate the root finding algorithm, call the following function. This is normally called only once, prior to the first call to `IDASolve()`, but if the rootfinding problem is to be changed during the solution, `IDARootInit()` can also be called prior to a continuation call to `IDASolve()`.

int **IDARootInit**(void \*ida\_mem, int nrtfn, *IDARootFn* g)

The function `IDARootInit()` specifies that the roots of a set of functions  $g_i(t, y)$  are to be found while the IVP is being solved.

#### Arguments:

- `ida_mem` – pointer to the IDAS solver object.
- `nrtfn` – is the number of root functions.
- `g` – is the function which defines the `nrtfn` functions  $g_i(t, y, \dot{y})$  whose roots are sought. See *IDARootFn* for more details.

#### Return value:

- `IDA_SUCCESS` – The call was successful.
- `IDA_MEM_NULL` – The `ida_mem` argument was NULL.
- `IDA_MEM_FAIL` – A memory allocation failed.
- `IDA_ILL_INPUT` – The function `g` is NULL, but `nrtfn > 0`.

#### Notes:

If a new IVP is to be solved with a call to `IDAReInit()`, where the new IVP has no rootfinding problem but the prior one did, then call `IDARootInit()` with `nrtfn = 0`.

### 5.1.3.9 IDAS solver function

This is the central step in the solution process, the call to perform the integration of the DAE. The input arguments (`itask`) specifies one of two modes as to where IDAS is to return a solution. These modes are modified if the user has set a stop time (with `IDASetStopTime()`) or requested rootfinding (with `IDARootInit()`).

int **IDASolve**(void \*ida\_mem, *sunrealtype* tout, *sunrealtype* \*tret, *N\_Vector* yret, *N\_Vector* ypret, int itask)

The function `IDASolve()` integrates the DAE over an interval in `t`.

#### Arguments:

- `ida_mem` – pointer to the IDAS solver object.
- `tout` – the next time at which a computed solution is desired.
- `tret` – the time reached by the solver output.
- `yret` – the computed solution vector  $y$ .
- `ypret` – the computed solution vector  $\dot{y}$ .
- `itask` – a flag indicating the job of the solver for the next user step
  - `IDA_NORMAL` – the solver will take internal steps until it has reached or just passed the user specified `tout` parameter. The solver then interpolates in order to return approximate values of  $y(t_{out})$  and  $\dot{y}(t_{out})$ .
  - `IDA_ONE_STEP` – the solver will just take one internal step and return the solution at the point reached by that step.

**Return value:**

- IDA\_SUCCESS – The call was successful.
- IDA\_TSTOP\_RETURN – *IDASolve()* succeeded by reaching the stop point specified through the optional input function *IDASetStopTime()*.
- IDA\_ROOT\_RETURN – *IDASolve()* succeeded and found one or more roots. In this case, *tret* is the location of the root. If *nrtfn* > 1, call *IDAGetRootInfo()* to see which  $g_i$  were found to have a root.
- IDA\_MEM\_NULL – The *ida\_mem* argument was NULL.
- IDA\_ILL\_INPUT – One of the inputs to *IDASolve()* was illegal, or some other input to the solver was either illegal or missing. The latter category includes the following situations:
  - The tolerances have not been set.
  - A component of the error weight vector became zero during internal time-stepping.
  - The linear solver initialization function called by the user after calling *IDACreate()* failed to set the linear solver-specific *lsolve* field in *ida\_mem*.
  - A root of one of the root functions was found both at a point  $t$  and also very near  $t$ .

In any case, the user should see the printed error message for details.

- IDA\_TOO\_MUCH\_WORK – The solver took *mxstep* internal steps but could not reach *tout*. The default value for *mxstep* is *MXSTEP\_DEFAULT* = 500.
- IDA\_TOO\_MUCH\_ACC – The solver could not satisfy the accuracy demanded by the user for some internal step.
- IDA\_ERR\_FAIL – Error test failures occurred too many times (*MXNEF* = 10) during one internal time step or occurred with  $|h| = h_{\min}$ .
- IDA\_CONV\_FAIL – Convergence test failures occurred too many times (*MXNCF* = 10) during one internal time step or occurred with  $|h| = h_{\min}$ .
- IDA\_LINIT\_FAIL – The linear solver's initialization function failed.
- IDA\_LSETUP\_FAIL – The linear solver's setup function failed in an unrecoverable manner.
- IDA\_LSOLVE\_FAIL – The linear solver's solve function failed in an unrecoverable manner.
- IDA\_CONSTR\_FAIL – The inequality constraints were violated and the solver was unable to recover.
- IDA\_REP\_RES\_ERR – The user's residual function repeatedly returned a recoverable error flag, but the solver was unable to recover.
- IDA\_RES\_FAIL – The user's residual function returned a nonrecoverable error flag.
- IDA\_RTFUNC\_FAIL – The rootfinding function failed.

**Notes:**

The vectors *yret* and *ypret* can occupy the same space as the initial condition vectors *y0* and *yp0*, respectively, that were passed to *IDAInit()*.

In the IDA\_ONE\_STEP mode, *tout* is used on the first call only, and only to get the direction and rough scale of the independent variable.

If a stop time is enabled (through a call to *IDASetStopTime()*), then *IDASolve()* returns the solution at *tstop*. Once the integrator returns at a stop time, any future testing for *tstop* is disabled (and can be reenabled only through a new call to *IDASetStopTime()*).

All failure return values are negative and therefore a test *flag* < 0 will trap all *IDASolve()* failures.

On any error return in which one or more internal steps were taken by `IDASolve()`, the returned values of `tret`, `yret`, and `ypret` correspond to the farthest point reached in the integration. On all other error returns, these values are left unchanged from the previous `IDASolve()` return.

### 5.1.3.10 Optional input functions

There are numerous optional input parameters that control the behavior of the IDAS solver. IDAS provides functions that can be used to change these optional input parameters from their default values. The main inputs are divided in the following categories:

- Table 5.1 list the main IDAS optional input functions,
- Table 5.2 lists the IDALS linear solver interface optional input functions,
- Table 5.3 lists the IDANLS nonlinear solver interface optional input functions,
- Table 5.4 lists the initial condition calculation optional input functions,
- Table 5.5 lists the IDAS step size adaptivity optional input functions, and
- Table 5.6 lists the rootfinding optional input functions.

These optional inputs are described in detail in the remainder of this section. For the most casual use of IDAS, the reader can skip to §5.1.4.

We note that, on an error return, all of the optional input functions also send an error message to the error handler function. All error return values are negative, so the test `flag < 0` will catch all errors.

The optional input calls can, unless otherwise noted, be executed in any order. Finally, a call to an `IDASet***` function can, unless otherwise noted, be made at any time from the user's calling program and, if successful, takes effect immediately.

### Main solver optional input functions

Table 5.1: Optional inputs for IDAS

Optional input	Function name	Default
User data	<code>IDASetUserData()</code>	NULL
Maximum order for BDF method	<code>IDASetMaxOrd()</code>	5
Maximum no. of internal steps before $t_{out}$	<code>IDASetMaxNumSteps()</code>	500
Initial step size	<code>IDASetInitStep()</code>	estimated
Minimum absolute step size $h_{min}$	<code>IDASetMinStep()</code>	0
Maximum absolute step size $h_{max}$	<code>IDASetMaxStep()</code>	$\infty$
Value of $t_{stop}$	<code>IDASetStopTime()</code>	undefined
Disable the stop time	<code>IDAClearStopTime()</code>	N/A
Maximum no. of error test failures	<code>IDASetMaxErrTestFails()</code>	10
Suppress alg. vars. from error test	<code>IDASetSuppressAlg()</code>	SUNFALSE
Variable types (differential/algebraic)	<code>IDASetId()</code>	NULL
Inequality constraints on solution	<code>IDASetConstraints()</code>	NULL

int `IDASetUserData`(void \*ida\_mem, void \*user\_data)

The function `IDASetUserData()` attaches a user-defined data pointer to the main IDAS solver object.

#### Arguments:

- `ida_mem` – pointer to the IDAS solver object.

- `user_data` – pointer to the user data.

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

**Notes:**

If specified, the pointer to `user_data` is passed to all user-supplied functions that have it as an argument. Otherwise, a NULL pointer is passed.

**Warning:** If `user_data` is needed in user linear solver or preconditioner functions, the call to `IDASetUserData()` must be made before the call to specify the linear solver.

int **IDASetMaxOrd**(void \*ida\_mem, int maxord)

The function `IDASetMaxOrd()` specifies the maximum order of the linear multistep method.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `maxord` – value of the maximum method order. This must be positive.

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_ILL_INPUT` – The input value `maxord` is  $\leq 0$ , or larger than the max order value when `IDAINIT()` was called.

**Notes:**

The default value is 5. If the input value exceeds 5, the value 5 will be used. If called before `IDAINIT()`, `maxord` limits the memory requirements for the internal IDAS memory block and its value cannot be increased past the value set when `IDAINIT()` was called.

int **IDASetMaxNumSteps**(void \*ida\_mem, long int mxsteps)

The function `IDASetMaxNumSteps()` specifies the maximum number of steps to be taken by the solver in its attempt to reach the next output time.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `mxsteps` – maximum allowed number of steps.

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

**Notes:**

Passing `mxsteps = 0` results in IDAS using the default value (500). Passing `mxsteps < 0` disables the test (not recommended).

int **IDASetInitStep**(void \*ida\_mem, *sunrealtype* hin)

The function `IDASetInitStep()` specifies the initial step size.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `hin` – value of the initial step size to be attempted. Pass 0.0 to have IDAS use the default value.

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

**Notes:**

By default, IDAS estimates the initial step as the solution of  $\|h\dot{y}\|_{WRMS} = 1/2$ , with an added restriction that  $|h| \leq .001|t_{\text{out}} - t_0|$ .

int **IDASetMinStep**(void \*ida\_mem, *sunrealtype* hmin)

The function *IDASetMinStep()* specifies the minimum absolute value of the step size.

Pass `hmin = 0` to obtain the default value of 0.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `hmin` – minimum absolute value of the step size.

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_ILL_INPUT` – `hmin` is negative.

New in version 5.2.0.

int **IDASetMaxStep**(void \*ida\_mem, *sunrealtype* hmax)

The function *IDASetMaxStep()* specifies the maximum absolute value of the step size.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `hmax` – maximum absolute value of the step size.

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_ILL_INPUT` – Either `hmax` is not positive or it is smaller than the minimum allowable step.

**Notes:**

Pass `hmax = 0` to obtain the default value  $\infty$ .

int **IDASetStopTime**(void \*ida\_mem, *sunrealtype* tstop)

The function *IDASetStopTime()* specifies the value of the independent variable  $t$  past which the solution is not to proceed.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `tstop` – value of the independent variable past which the solution should not proceed.

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.

- IDA\_MEM\_NULL – The `ida_mem` pointer is NULL.
- IDA\_ILL\_INPUT – The value of `tstop` is not beyond the current  $t$  value,  $t_n$ .

**Notes:**

The default, if this routine is not called, is that no stop time is imposed. Once the integrator returns at a stop time, any future testing for `tstop` is disabled (and can be reenabled only through a new call to `IDASetStopTime()`).

A stop time not reached before a call to `IDAREInit()` will remain active but can be disabled by calling `IDAClearStopTime()`.

int **IDAClearStopTime**(void \*ida\_mem)

Disables the stop time set with `IDASetStopTime()`.

**Arguments:**

- `ida_mem` – pointer to the IDA memory block.

**Return value:**

- IDA\_SUCCESS if successful
- IDA\_MEM\_NULL if the IDA memory is NULL

**Notes:**

The stop time can be reenabled through a new call to `IDASetStopTime()`.

New in version 6.5.1.

int **IDASetMaxErrTestFails**(void \*ida\_mem, int maxnef)

The function `IDASetMaxErrTestFails()` specifies the maximum number of error test failures in attempting one step.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `maxnef` – maximum number of error test failures allowed on one step (>0).

**Return value:**

- IDA\_SUCCESS – The optional value has been successfully set.
- IDA\_MEM\_NULL – The `ida_mem` pointer is NULL.

**Notes:**

The default value is 10.

int **IDASetSuppressAlg**(void \*ida\_mem, *sunbooleantype* suppressalg)

The function `IDASetSuppressAlg()` indicates whether or not to suppress algebraic variables in the local error test.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `suppressalg` – indicates whether to suppress (SUNTRUE) or include (SUNFALSE) the algebraic variables in the local error test.

**Return value:**

- IDA\_SUCCESS – The optional value has been successfully set.
- IDA\_MEM\_NULL – The `ida_mem` pointer is NULL.



**Notes:**

The default value is `SUNFALSE`. If `suppressalg = SUNTRUE` is selected, then the `id` vector must be set (through `IDASetId()`) to specify the algebraic components. In general, the use of this option (with `suppressalg = SUNTRUE`) is *discouraged* when solving DAE systems of index 1, whereas it is generally *encouraged* for systems of index 2 or more. See pp. 146-147 of [14] for more on this issue.

int **IDASetId**(void \*ida\_mem, *N\_Vector* id)

The function `IDASetId()` specifies algebraic/differential components in the  $y$  vector.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `id` – a vector of values identifying the components of  $y$  as differential or algebraic variables. A value of 1.0 indicates a differential variable, while 0.0 indicates an algebraic variable.

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

**Notes:**

The vector `id` is required if the algebraic variables are to be suppressed from the local error test (see `IDASetSuppressAlg()`) or if `IDACalcIC()` is to be called with `icopt = IDA_YA_YDP_INIT`.

int **IDASetConstraints**(void \*ida\_mem, *N\_Vector* constraints)

The function `IDASetConstraints()` specifies a vector defining inequality constraints for each component of the solution vector  $y$ .

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `constraints` – vector of constraint flags.
  - If `constraints[i] = 0`, no constraint is imposed on  $y_i$ .
  - If `constraints[i] = 1`,  $y_i$  will be constrained to be  $y_i \geq 0.0$ .
  - If `constraints[i] = -1`,  $y_i$  will be constrained to be  $y_i \leq 0.0$ .
  - If `constraints[i] = 2`,  $y_i$  will be constrained to be  $y_i > 0.0$ .
  - If `constraints[i] = -2`,  $y_i$  will be constrained to be  $y_i < 0.0$ .

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_ILL_INPUT` – The constraints vector contains illegal values or the simultaneous corrector option has been selected when doing forward sensitivity analysis.

**Notes:**

The presence of a non-NULL constraints vector that is not 0.0 in all components will cause constraint checking to be performed. However, a call with 0.0 in all components of constraints vector will result in an illegal input return. A NULL input will disable constraint checking.

Constraint checking when doing forward sensitivity analysis with the simultaneous corrector option is currently disallowed and will result in an illegal input return.

## Linear solver interface optional input functions

Table 5.2: Optional inputs for the IDALS linear solver interface

Optional input	Function name	Default
Jacobian function	<i>IDASetJacFn()</i>	DQ
Set parameter determining if a $c_j$ change requires a linear solver setup call	<i>IDASetDeltaCjLSetup()</i>	0.25
Enable or disable linear solution scaling	<i>IDASetLinearSolutionScaling()</i>	on
Jacobian-times-vector function	<i>IDASetJacTimes()</i>	NULL, DQ
Preconditioner functions	<i>IDASetPreconditioner()</i>	NULL, NULL
Ratio between linear and nonlinear tolerances	<i>IDASetEpsLin()</i>	0.05
Increment factor used in DQ $J_v$ approx.	<i>IDASetIncrementFactor()</i>	1.0
Jacobian-times-vector DQ Res function	<i>IDASetJacTimesResFn()</i>	NULL
Newton linear solve tolerance conversion factor	<i>IDASetLSNormFactor()</i>	vector length

The mathematical explanation of the linear solver methods available to IDAS is provided in §2.2. We group the user-callable routines into four categories: general routines concerning the overall IDALS linear solver interface, optional inputs for matrix-based linear solvers, optional inputs for matrix-free linear solvers, and optional inputs for iterative linear solvers. We note that the matrix-based and matrix-free groups are mutually exclusive, whereas the “iterative” tag can apply to either case.

When using matrix-based linear solver modules, the IDALS solver interface needs a function to compute an approximation to the Jacobian matrix  $J(t, y, \dot{y})$ . This function must be of type *IDALSJacFn*. The user can supply a Jacobian function or, if using the *SUNMATRIX\_DENSE* or *SUNMATRIX\_BAND* modules for the matrix  $J$ , can use the default internal difference quotient approximation that comes with the IDALS interface. To specify a user-supplied Jacobian function *jac*, IDALS provides the function *IDASetJacFn()*. The IDALS interface passes the pointer *user\_data* to the Jacobian function. This allows the user to create an arbitrary structure with relevant problem data and access it during the execution of the user-supplied Jacobian function, without using global data in the program. The pointer *user\_data* may be specified through *IDASetUserData()*.

int *IDASetJacFn*(void \*ida\_mem, *IDALSJacFn* jac)

The function *IDASetJacFn()* specifies the Jacobian approximation function to be used for a matrix-based solver within the IDALS interface.

**Arguments:**

- *ida\_mem* – pointer to the IDAS solver object.
- *jac* – user-defined Jacobian approximation function. See *IDALSJacFn* for more details.

**Return value:**

- *IDALS\_SUCCESS* – The optional value has been successfully set.
- *IDALS\_MEM\_NULL* – The *ida\_mem* pointer is NULL.
- *IDALS\_LMEM\_NULL* – The IDALS linear solver interface has not been initialized.

**Notes:**

This function must be called after the IDALS linear solver interface has been initialized through a call to *IDASetLinearSolver()*. By default, IDALS uses an internal difference quotient function for the *SUNMATRIX\_DENSE* and *SUNMATRIX\_BAND* modules. If NULL is passed to *jac*, this default function is used. An error will occur if no *jac* is supplied when using other matrix types.

New in version 3.0.0: Replaces the deprecated function `IDADlsSetJacFn`.

When using a matrix-based linear solver the matrix information will be updated infrequently to reduce matrix construction and, with direct solvers, factorization costs. As a result the value of  $\alpha$  may not be current and a scaling factor is applied to the solution of the linear system to account for the lagged value of  $\alpha$ . See §8.2.1 for more details. The function `IDASetLinearSolutionScaling()` can be used to disable this scaling when necessary, e.g., when providing a custom linear solver that updates the matrix using the current  $\alpha$  as part of the solve.

int **IDASetDeltaCjLSetup**(void \*ida\_mem, *sunrealtype* dcj)

The function `IDASetDeltaCjLSetup` specifies the parameter that determines the bounds on a change in  $c_j$  that require a linear solver setup call. If  $c_{j\_current} / c_{j\_previous} < (1 - dcj) / (1 + dcj)$  or  $c_{j\_current} / c_{j\_previous} > (1 + dcj) / (1 - dcj)$ , the linear solver setup function is called.

If  $dcj$  is  $< 0$  or  $\geq 1$  then the default value (0.25) is used.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `dcj` – the  $c_j$  change threshold.

**Return value:**

- `IDA_SUCCESS` – The flag value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

New in version 5.2.0.

int **IDASetLinearSolutionScaling**(void \*ida\_mem, *sunbooleantype* onoff)

The function `IDASetLinearSolutionScaling()` enables or disables scaling the linear system solution to account for a change in  $\alpha$  in the linear system. For more details see §8.2.1.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `onoff` – flag to enable (`SUNTRUE`) or disable (`SUNFALSE`) scaling.

**Return value:**

- `IDALS_SUCCESS` – The flag value has been successfully set.
- `IDALS_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDALS_LMEM_NULL` – The IDALS linear solver interface has not been initialized.
- `IDALS_ILL_INPUT` – The attached linear solver is not matrix-based.

**Notes:**

This function must be called after the IDALS linear solver interface has been initialized through a call to `IDASetLinearSolver()`. By default scaling is enabled with matrix-based linear solvers.

When using matrix-free linear solver modules, the IDALS solver interface requires a function to compute an approximation to the product between the Jacobian matrix  $J(t, y, \dot{y})$  and a vector  $v$ . The user can supply a Jacobian-times-vector approximation function, or use the default internal difference quotient function that comes with the IDALS solver interface.

A user-defined Jacobian-vector product function must be of type `IDALSJacTimesVecFn` and can be specified through a call to `IDASetJacTimes()`. The evaluation and processing of any Jacobian-related data needed by the user's Jacobian-vector product function may be done in the optional user-supplied function `jtsetup` (see §5.1.4.6 for specification details). The pointer `user_data` received through `IDASetUserData()` (or a pointer to NULL if `user_data` was not specified) is passed to the Jacobian-vector product setup and product functions, `jtsetup` and `jt看times`, each time they are called. This allows the user to create an arbitrary structure with relevant problem data and access it during the execution of the user-supplied functions without using global data in the program.

int **IDASetJacTimes**(void \*ida\_mem, *IDALsJacTimesSetupFn* jsetup, *IDALsJacTimesVecFn* jtimes)

The function *IDASetJacTimes()* specifies the Jacobian-vector product setup and product functions.

**Arguments:**

- *ida\_mem* – pointer to the IDAS solver object.
- *jsetup* – user-defined function to set up the Jacobian-vector product. See *IDALsJacTimesSetupFn* for more details. Pass NULL if no setup is necessary.
- *jtimes* – user-defined Jacobian-vector product function. See *IDALsJacTimesVecFn* for more details.

**Return value:**

- IDALS\_SUCCESS – The optional value has been successfully set.
- IDALS\_MEM\_NULL – The *ida\_mem* pointer is NULL.
- IDALS\_LMEM\_NULL – The IDALS linear solver has not been initialized.
- IDALS\_SUNLS\_FAIL – An error occurred when setting up the system matrix-times-vector routines in the SUNLinearSolver object used by the IDALS interface.

**Notes:**

The default is to use an internal finite difference quotient for *jtimes* and to omit *jsetup*. If NULL is passed to *jtimes*, these defaults are used. A user may specify non-NULL *jtimes* and NULL *jsetup* inputs. This function must be called after the IDALS linear solver interface has been initialized through a call to *IDASetLinearSolver()*.

New in version 3.0.0: Replaces the deprecated function *IDASpilsSetJacTimes*.

When using the default difference-quotient approximation to the Jacobian-vector product, the user may specify the factor to use in setting increments for the finite-difference approximation, via a call to *IDASetIncrementFactor()*.

int **IDASetIncrementFactor**(void \*ida\_mem, *sunrealtype* dqincfac)

The function *IDASetIncrementFactor()* specifies the increment factor to be used in the difference-quotient approximation to the product  $Jv$ . Specifically,  $Jv$  is approximated via the formula

$$Jv = \frac{1}{\sigma} [F(t, \tilde{y}, \tilde{y}) - F(t, y, \dot{y})],$$

where  $\tilde{y} = y + \sigma v$ ,  $\tilde{y} = \dot{y} + c_j \sigma v$ ,  $c_j$  is a BDF parameter proportional to the step size,  $\sigma = dqincfac\sqrt{N}$ , and  $N$  is the number of equations in the DAE system.

**Arguments:**

- *ida\_mem* – pointer to the IDAS solver object.
- *dqincfac* – user-specified increment factor positive.

**Return value:**

- IDALS\_SUCCESS – The optional value has been successfully set.
- IDALS\_MEM\_NULL – The *ida\_mem* pointer is NULL.
- IDALS\_LMEM\_NULL – The IDALS linear solver has not been initialized.
- IDALS\_ILL\_INPUT – The specified value of *dqincfac* is  $\leq 0$ .

**Notes:**

The default value is 1.0. This function must be called after the IDALS linear solver interface has been initialized through a call to *IDASetLinearSolver()*.

New in version 3.0.0: Replaces the deprecated function *IDASpilsSetIncrementFactor*.

Additionally, when using the internal difference quotient, the user may also optionally supply an alternative residual function for use in the Jacobian-vector product approximation by calling [IDASetJacTimesResFn\(\)](#). The alternative residual function should compute a suitable (and differentiable) approximation to the residual function provided to [IDAInit\(\)](#). For example, as done in [31] for an ODE in explicit form, the alternative function may use lagged values when evaluating a nonlinearity to avoid differencing a potentially non-differentiable factor.

```
int IDASetJacTimesResFn(void *ida_mem, IDAResFn jtimesResFn)
```

The function [IDASetJacTimesResFn\(\)](#) specifies an alternative DAE residual function for use in the internal Jacobian-vector product difference quotient approximation.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `jtimesResFn` – is the function which computes the alternative DAE residual function to use in Jacobian-vector product difference quotient approximations. See [IDAResFn](#) for more details.

**Return value:**

- `IDALS_SUCCESS` – The optional value has been successfully set.
- `IDALS_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDALS_LMEM_NULL` – The IDALS linear solver has not been initialized.
- `IDALS_ILL_INPUT` – The internal difference quotient approximation is disabled.

**Notes:**

The default is to use the residual function provided to [IDAInit\(\)](#) in the internal difference quotient. If the input residual function is NULL, the default is used. This function must be called after the IDALS linear solver interface has been initialized through a call to [IDASetLinearSolver\(\)](#).

When using an iterative linear solver, the user may supply a preconditioning operator to aid in solution of the system. This operator consists of two user-supplied functions, `psetup` and `psolve`, that are supplied to IDAS using the function [IDASetPreconditioner\(\)](#). The `psetup` function supplied to this routine should handle evaluation and preprocessing of any Jacobian data needed by the user's preconditioner solve function, `psolve`. Both of these functions are fully specified in §5.1.4.7 and §5.1.4.8). The user data pointer received through [IDASetUserData\(\)](#) (or NULL if a user data pointer was not specified) is passed to the `psetup` and `psolve` functions. This allows the user to create an arbitrary structure with relevant problem data and access it during the execution of the user-supplied preconditioner functions without using global data in the program.

```
int IDASetPreconditioner(void *ida_mem, IDALSPrecSetupFn psetup, IDALSPrecSolveFn psolve)
```

The function [IDASetPreconditioner\(\)](#) specifies the preconditioner setup and solve functions.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `psetup` – user-defined function to set up the preconditioner. See [IDALSPrecSetupFn](#) for more details. Pass NULL if no setup is necessary.
- `psolve` – user-defined preconditioner solve function. See [IDALSPrecSolveFn](#) for more details.

**Return value:**

- `IDALS_SUCCESS` – The optional values have been successfully set.
- `IDALS_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDALS_LMEM_NULL` – The IDALS linear solver has not been initialized.
- `IDALS_SUNLS_FAIL` – An error occurred when setting up preconditioning in the SUNLinearSolver object used by the IDALS interface.

**Notes:**

The default is NULL for both arguments (i.e., no preconditioning). This function must be called after the IDALS linear solver interface has been initialized through a call to [IDASetLinearSolver\(\)](#).

New in version 3.0.0: Replaces the deprecated function IDASpilsSetPreconditioner.

Also, as described in §2.2, the IDALS interface requires that iterative linear solvers stop when the norm of the preconditioned residual satisfies

$$\|r\| \leq \frac{\epsilon_L \epsilon}{10}$$

where  $\epsilon$  is the nonlinear solver tolerance, and the default  $\epsilon_L = 0.05$ ; this value may be modified by the user through the [IDASetEpsLin\(\)](#) function.

int [IDASetEpsLin](#)(void \*ida\_mem, *sunrealtype* eplifac)

The function [IDASetEpsLin\(\)](#) specifies the factor by which the Krylov linear solver's convergence test constant is reduced from the nonlinear iteration test constant.

**Arguments:**

- ida\_mem – pointer to the IDAS solver object.
- eplifac – linear convergence safety factor  $\geq 0.0$ .

**Return value:**

- IDALS\_SUCCESS – The optional value has been successfully set.
- IDALS\_MEM\_NULL – The ida\_mem pointer is NULL.
- IDALS\_LMEM\_NULL – The IDALS linear solver has not been initialized.
- IDALS\_ILL\_INPUT – The factor eplifac is negative.

**Notes:**

The default value is 0.05. This function must be called after the IDALS linear solver interface has been initialized through a call to [IDASetLinearSolver\(\)](#). If eplifac = 0.0 is passed, the default value is used.

New in version 3.0.0: Replaces the deprecated function IDASpilsSetEpsLin.

int [IDASetLSNormFactor](#)(void \*ida\_mem, *sunrealtype* nrmfac)

The function [IDASetLSNormFactor\(\)](#) specifies the factor to use when converting from the integrator tolerance (WRMS norm) to the linear solver tolerance (L2 norm) for Newton linear system solves e.g., `tol_L2 = fac * tol_WRMS`.

**Arguments:**

- ida\_mem – pointer to the IDAS solver object.
- nrmfac – the norm conversion factor.
  - If nrmfac > 0, the provided value is used.
  - If nrmfac = 0 then the conversion factor is computed using the vector length i.e., `nrmfac = N_VGetLength(y)` (default).
  - If nrmfac < 0 then the conversion factor is computed using the vector dot product `nrmfac = N_VDotProd(v,v)` where all the entries of v are one.

**Return value:**

- IDA\_SUCCESS – The optional value has been successfully set.
- IDA\_MEM\_NULL – The ida\_mem pointer is NULL.

**Notes:**

This function must be called after the IDALS linear solver interface has been initialized through a call to [IDASetLinearSolver\(\)](#). Prior to the introduction of [N\\_VGetLength\(\)](#) in SUNDIALS v5.0.0 (IDAS v4.0.0) the value of `nrmfac` was computed using [N\\_VDotProd\(\)](#) i.e., the `nrmfac < 0` case.

**Nonlinear solver interface optional input functions**

Table 5.3: Optional inputs for the IDANLS nonlinear solver interface

Optional input	Function name	Default
Maximum no. of nonlinear iterations	<a href="#">IDASetMaxNonlinIters()</a>	4
Maximum no. of convergence failures	<a href="#">IDASetMaxConvFails()</a>	10
Coeff. in the nonlinear convergence test	<a href="#">IDASetNonlinConvCoef()</a>	0.33
Residual function for nonlinear system evaluations	<a href="#">IDASetNlsResFn()</a>	NULL

The following functions can be called to set optional inputs controlling the nonlinear solver.

int [IDASetMaxNonlinIters](#)(void \*ida\_mem, int maxcor)

The function [IDASetMaxNonlinIters\(\)](#) specifies the maximum number of nonlinear solver iterations in one solve attempt.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `maxcor` – maximum number of nonlinear solver iterations allowed in one solve attempt (>0).

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_MEM_FAIL` – The `SUNNonlinearSolver` object is NULL.

**Notes:**

The default value is 4.

int [IDASetMaxConvFails](#)(void \*ida\_mem, int maxncf)

The function [IDASetMaxConvFails\(\)](#) specifies the maximum number of nonlinear solver convergence failures in one step.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `maxncf` – maximum number of allowable nonlinear solver convergence failures in one step (>0).

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

**Notes:**

The default value is 10.

int [IDASetNonlinConvCoef](#)(void \*ida\_mem, *sunrealtype* nlscoef)

The function [IDASetNonlinConvCoef\(\)](#) specifies the safety factor in the nonlinear convergence test; see (2.8).

**Arguments:**



- `ida_mem` – pointer to the IDAS solver object.
- `nlscoef` – coefficient in nonlinear convergence test ( $>0.0$ ).

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_ILL_INPUT` – The value of `nlscoef` is  $\leq 0.0$ .

**Notes:**

The default value is 0.33.

int **IDASetNlsResFn**(void \*ida\_mem, *IDAResFn* res)

The function *IDASetNlsResFn()* specifies an alternative residual function for use in nonlinear system function evaluations.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `res` – the alternative function which computes the DAE residual function  $F(t, y, \dot{y})$ . See *IDAResFn* for more details.

**Return value:**

- `IDA_SUCCESS` – The optional function has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

**Notes:**

The default is to use the residual function provided to *IDAINit()* in nonlinear system function evaluations. If the input residual function is NULL, the default is used.

When using a non-default nonlinear solver, this function must be called after *IDASetNonlinearSolver()*.

When doing forward sensitivity analysis with the simultaneous solver strategy is function must be called after *IDASetNonlinearSolverSensSim()*.

**Initial condition calculation optional input functions**

Table 5.4: Optional inputs for IDAS initial condition calculation

Optional input	Function name	Default
Coeff. in the nonlinear convergence test	<i>IDASetNonlinConvCoefIC()</i>	0.0033
Maximum no. of steps	<i>IDASetMaxNumStepsIC()</i>	5
Maximum no. of Jacobian/precond. evals.	<i>IDASetMaxNumJacsIC()</i>	4
Maximum no. of Newton iterations	<i>IDASetMaxNumItersIC()</i>	10
Max. linesearch backtracks per Newton iter.	<i>IDASetMaxBacksIC()</i>	100
Turn off linesearch	<i>IDASetLineSearchOffIC()</i>	SUNFALSE
Lower bound on Newton step	<i>IDASetStepToleranceIC()</i>	around <sup>2/3</sup>

The following functions can be called just prior to calling *IDACalcIC()* to set optional inputs controlling the initial condition calculation.

int **IDASetNonlinConvCoefIC**(void \*ida\_mem, *sunrealtype* epiccon)

The function *IDASetNonlinConvCoefIC()* specifies the positive constant in the Newton iteration convergence test within the initial condition calculation.



**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `epiccon` – coefficient in the Newton convergence test ( $> 0$ ).

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_ILL_INPUT` – The `epiccon` factor is  $\leq 0.0$ .

**Notes:**

The default value is  $0.01 \cdot 0.33$ . This test uses a weighted RMS norm (with weights defined by the tolerances). For new initial value vectors  $y$  and  $\dot{y}$  to be accepted, the norm of  $J^{-1}F(t_0, y, \dot{y})$  must be  $\leq \text{epiccon}$ , where  $J$  is the system Jacobian.

int **IDASetMaxNumStepsIC**(void \*ida\_mem, int maxnh)

The function *IDASetMaxNumStepsIC()* specifies the maximum number of steps allowed when `icopt = IDA_YA_YDP_INIT` in *IDACalcIC()*, where  $h$  appears in the system Jacobian,  $J = \frac{\partial F}{\partial y} + \left(\frac{1}{h}\right) \frac{\partial F}{\partial \dot{y}}$ .

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `maxnh` – maximum allowed number of values for  $h$ .

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_ILL_INPUT` – `maxnh` is non-positive.

**Notes:**

The default value is 5.

int **IDASetMaxNumJacsIC**(void \*ida\_mem, int maxnj)

The function *IDASetMaxNumJacsIC()* specifies the maximum number of the approximate Jacobian or preconditioner evaluations allowed when the Newton iteration appears to be slowly converging.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `maxnj` – maximum allowed number of Jacobian or preconditioner evaluations.

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_ILL_INPUT` – `maxnj` is non-positive.

**Notes:**

The default value is 4.

int **IDASetMaxNumItersIC**(void \*ida\_mem, int maxnit)

The function *IDASetMaxNumItersIC()* specifies the maximum number of Newton iterations allowed in any one attempt to solve the initial conditions calculation problem.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `maxnit` – maximum number of Newton iterations.

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_ILL_INPUT` – `maxnit` is non-positive.

**Notes:**

The default value is 10.

int **IDASetMaxBacksIC**(void \*ida\_mem, int maxbacks)

The function *IDASetMaxBacksIC()* specifies the maximum number of linesearch backtracks allowed in any Newton iteration, when solving the initial conditions calculation problem.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `maxbacks` – maximum number of linesearch backtracks per Newton step.

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_ILL_INPUT` – `maxbacks` is non-positive.

**Notes:**

The default value is 100.

If *IDASetMaxBacksIC()* is called in a Forward Sensitivity Analysis, the the limit `maxbacks` applies in the calculation of both the initial state values and the initial sensitivities.

int **IDASetLineSearchOffIC**(void \*ida\_mem, *sunbooleantype* lsoff)

The function *IDASetLineSearchOffIC()* specifies whether to turn on or off the linesearch algorithm.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `lsoff` – a flag to turn off (`SUNTRUE`) or keep (`SUNFALSE`) the linesearch algorithm.

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

**Notes:**

The default value is `SUNFALSE`.

int **IDASetStepToleranceIC**(void \*ida\_mem, int steptol)

The function *IDASetStepToleranceIC()* specifies a positive lower bound on the Newton step.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `steptol` – Minimum allowed WRMS-norm of the Newton step ( $> 0.0$ ).

**Return value:**

- IDA\_SUCCESS – The optional value has been successfully set.
- IDA\_MEM\_NULL – The `ida_mem` pointer is NULL.
- IDA\_ILL\_INPUT – The `steptol` tolerance is  $\leq 0.0$ .

**Notes:**

The default value is  $(\text{unit roundoff})^{2/3}$ .

**Time step adaptivity optional input functions**

Table 5.5: Optional inputs for IDAS time step adaptivity

Optional input	Function name	Default
Fixed step size bounds $\eta_{\min\_fx}$ and $\eta_{\max\_fx}$	<i>IDASetEtaFixedStep-</i> <i>Bounds()</i>	1.0 and 2.0
Maximum step size growth factor $\eta_{\max}$	<i>IDASetEtaMax()</i>	2.0
Minimum step size reduction factor $\eta_{\min}$	<i>IDASetEtaMin()</i>	0.5
Maximum step size reduction factor $\eta_{\text{low}}$	<i>IDASetEtaLow()</i>	0.9
Minimum step size reduction factor after an error test failure $\eta_{\min\_ef}$	<i>IDASetEtaMinErrFail()</i>	0.25
Step size reduction factor after a nonlinear solver convergence failure $\eta_{cf}$	<i>IDASetEtaConvFail()</i>	0.25

The following functions can be called to set optional inputs to control the step size adaptivity.

**Note:** The default values for the step size adaptivity tuning parameters have a long history of success and changing the values is generally discouraged. However, users that wish to experiment with alternative values should be careful to make changes gradually and with testing to determine their effectiveness.

int **IDASetEtaFixedStepBounds**(void \*ida\_mem, *sunrealtype* eta\_min\_fx, *sunrealtype* eta\_max\_fx)

The function `IDASetEtaFixedStepBounds` specifies the bounds  $\eta_{\min\_fx}$  and  $\eta_{\max\_fx}$ . If step size change factor  $\eta$  satisfies  $\eta_{\min\_fx} < \eta < \eta_{\max\_fx}$  the current step size is retained.

The default values are  $\eta_{fxmin} = 1$  and  $\eta_{fxmax} = 2$ .

`eta_fxmin` should satisfy  $0 < \eta_{fxmin} \leq 1$ , otherwise the default value is used. `eta_fxmax` should satisfy  $\eta_{fxmin} \geq 1$ , otherwise the default value is used.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `eta_min_fx` – value of the fixed step size lower bound.
- `eta_max_fx` – value of the fixed step size upper bound.

**Return value:**

- IDA\_SUCCESS – The optional value has been successfully set.
- IDA\_MEM\_NULL – The `ida_mem` pointer is NULL.

New in version 5.2.0.

int **IDASetEtaMax**(void \*ida\_mem, *sunrealtype* eta\_max)

The function IDASetEtaMax specifies the maximum step size growth factor  $\eta_{\max}$ .

The default value is  $\eta_{\max} = 2$ .

**Arguments:**

- ida\_mem – pointer to the IDAS solver object.
- eta\_max – maximum step size growth factor.

**Return value:**

- IDA\_SUCCESS – The optional value has been successfully set.
- IDA\_MEM\_NULL – The ida\_mem pointer is NULL.

New in version 5.2.0.

int **IDASetEtaMin**(void \*ida\_mem, *sunrealtype* eta\_min)

The function IDASetEtaMin specifies the minimum step size reduction factor  $\eta_{\min}$ .

The default value is  $\eta_{\min} = 0.5$ .

eta\_min should satisfy  $0 < \eta_{\min} < 1$ , otherwise the default value is used.

**Arguments:**

- ida\_mem – pointer to the IDAS solver object.
- eta\_min – value of the minimum step size reduction factor.

**Return value:**

- IDA\_SUCCESS – The optional value has been successfully set.
- IDA\_MEM\_NULL – The ida\_mem pointer is NULL.

New in version 5.2.0.

int **IDASetEtaLow**(void \*ida\_mem, *sunrealtype* eta\_low)

The function IDASetEtaLow specifies the maximum step size reduction factor  $\eta_{\text{low}}$ .

The default value is  $\eta_{\text{low}} = 0.9$ .

eta\_low should satisfy  $0 < \eta_{\text{low}} \leq 1$ , otherwise the default value is used.

**Arguments:**

- ida\_mem – pointer to the IDAS solver object.
- eta\_low – value of the maximum step size reduction factor.

**Return value:**

- IDA\_SUCCESS – The optional value has been successfully set.
- IDA\_MEM\_NULL – The ida\_mem pointer is NULL.

New in version 5.2.0.

int **IDASetEtaMinErrFail**(void \*ida\_mem, *sunrealtype* eta\_min\_ef)

The function IDASetEtaMinErrFail specifies the minimum step size reduction factor  $\eta_{\text{min\_ef}}$  after an error test failure.

The default value is  $\eta_{\text{min\_ef}} = 0.25$ .

If eta\_min\_ef is  $\leq 0$  or  $\geq 1$ , the default value is used.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `eta_low` – value of the minimum step size reduction factor.

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

New in version 5.2.0.

int **IDASetEtaConvFail**(void \*ida\_mem, *sunrealtype* eta\_cf)

The function `IDASetEtaConvFail` specifies the step size reduction factor  $\eta_{cf}$  after a nonlinear solver convergence failure.

The default value is  $\eta_{cf} = 0.25$ .

If `eta_cf` is  $\leq 0$  or  $\geq 1$ , the default value is used.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `eta_low` – value of the step size reduction factor.

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

New in version 5.2.0.

**Rootfinding optional input functions**

Table 5.6: Optional inputs for IDAS rootfinding

Optional input	Function name	Default
Direction of zero-crossing	<code>IDASetRootDirection()</code>	both
Disable rootfinding warnings	<code>IDASetNoInactiveRootWarn()</code>	none

The following functions can be called to set optional inputs to control the rootfinding algorithm.

int **IDASetRootDirection**(void \*ida\_mem, int \*rootdir)

The function `IDASetRootDirection()` specifies the direction of zero-crossings to be located and returned to the user.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `rootdir` – state array of length `nrtfn`, the number of root functions  $g_i$ , as specified in the call to the function `IDARootInit()`.
  - A value of 0 for `rootdir[i]` indicates that crossing in either direction should be reported for  $g_i$ .
  - A value of +1 or -1 for `rootdir[i]` indicates that the solver should report only zero-crossings where  $g_i$  is increasing or decreasing, respectively.

**Return value:**

- IDA\_SUCCESS – The optional value has been successfully set.
- IDA\_MEM\_NULL – The `ida_mem` pointer is NULL.
- IDA\_ILL\_INPUT – rootfinding has not been activated through a call to `IDARootInit()`.

**Notes:**

The default behavior is to locate both zero-crossing directions.

int **IDASetNoInactiveRootWarn**(void \*ida\_mem)

The function `IDASetNoInactiveRootWarn()` disables issuing a warning if some root function appears to be identically zero at the beginning of the integration.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.

**Return value:**

- IDA\_SUCCESS – The optional value has been successfully set.
- IDA\_MEM\_NULL – The `ida_mem` pointer is NULL.

**Notes:**

IDAS will not report the initial conditions as a possible zero-crossing (assuming that one or more components  $g_i$  are zero at the initial time). However, if it appears that some  $g_i$  is identically zero at the initial time (i.e.,  $g_i$  is zero at the initial time and after the first step), IDAS will issue a warning which can be disabled with this optional input function.

### 5.1.3.11 Interpolated output function

An optional function `IDAGetDky()` is available to obtain additional output values. This function must be called after a successful return from `IDASolve()` and provides interpolated values of  $y$  or its derivatives of order up to the last internal order used for any value of  $t$  in the last internal step taken by IDAS.

int **IDAGetDky**(void \*ida\_mem, *sunrealtype* t, int k, *N\_Vector* dky)

The function `IDAGetDky()` computes the interpolated values of the  $k^{th}$  derivative of  $y$  for any value of  $t$  in the last internal step taken by IDAS. The value of  $k$  must be non-negative and smaller than the last internal order used. A value of 0 for  $k$  means that the  $y$  is interpolated. The value of  $t$  must satisfy  $t_n - h_u \leq t \leq t_n$ , where  $t_n$  denotes the current internal time reached, and  $h_u$  is the last internal step size used successfully.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `t` – time at which to interpolate.
- `k` – integer specifying the order of the derivative of  $y$  wanted.
- `dky` – vector containing the interpolated  $k^{th}$  derivative of  $y(t)$ .

**Return value:**

- IDA\_SUCCESS – `IDAGetDky()` succeeded.
- IDA\_MEM\_NULL – The `ida_mem` argument was NULL.
- IDA\_BAD\_T – `t` is not in the interval  $[t_n - h_u, t_n]$ .
- IDA\_BAD\_K – `k` is not one of  $0, 1, \dots, k_{last}$ .
- IDA\_BAD\_DKY – `dky` is NULL.

**Notes:**

It is only legal to call the function *IDAGetDky()* after a successful return from *IDASolve()*. Functions *IDAGetCurrentTime()*, *IDAGetLastStep()* and *IDAGetLastOrder()* can be used to access  $t_n$ ,  $h_u$ , and  $k_{\text{last}}$ .

**5.1.3.12 Optional output functions**

IDAS provides an extensive list of functions that can be used to obtain solver performance information. Table 5.7 lists all optional output functions in IDAS, which are then described in detail in the remainder of this section.

Some of the optional outputs, especially the various counters, can be very useful in determining how successful the IDAS solver is in doing its job. For example, the counters *nsteps* and *nrevals* provide a rough measure of the overall cost of a given run, and can be compared among runs with differing input options to suggest which set of options is most efficient. The ratio *nniters/nsteps* measures the performance of the nonlinear solver in solving the nonlinear systems at each time step; typical values for this range from 1.1 to 1.8. The ratio *njevals/nniters* (in the case of a matrix-based linear solver), and the ratio *npevals/nniters* (in the case of an iterative linear solver) measure the overall degree of nonlinearity in these systems, and also the quality of the approximate Jacobian or preconditioner being used. Thus, for example, *njevals/nniters* can indicate if a user-supplied Jacobian is inaccurate, if this ratio is larger than for the case of the corresponding internal Jacobian. The ratio *nliters/nniters* measures the performance of the Krylov iterative linear solver, and thus (indirectly) the quality of the preconditioner.

Table 5.7: Optional outputs for IDAS, IDALS, and IDANLS

Optional output	Function name
Size of IDAS real and integer workspace	<i>IDAGetWorkSpa</i>
Cumulative number of internal steps	<i>IDAGetNumStep</i>
No. of calls to residual function	<i>IDAGetNumResE</i>
No. of calls to linear solver setup function	<i>IDAGetNumLinS</i>
No. of local error test failures that have occurred   <i>IDAGetNumErrTestFails()</i>	
No. of failed steps due to a nonlinear solver failure   <i>IDAGetNumStepSolveFails()</i>	
Order used during the last step	<i>IDAGetLastOrd</i>
Order to be attempted on the next step	<i>IDAGetCurrent</i>
Actual initial step size used	<i>IDAGetActualI</i>
Step size used for the last step	<i>IDAGetLastSte</i>
Step size to be attempted on the next step	<i>IDAGetCurrent</i>
Current internal time reached by the solver	<i>IDAGetCurrent</i>
Suggested factor for tolerance scaling	<i>IDAGetTolScal</i>
Error weight vector for state variables	<i>IDAGetErrWeig</i>
Estimated local errors	<i>IDAGetEstLoca</i>
All IDA integrator statistics	<i>IDAGetIntegra</i>
No. of nonlinear solver iterations	<i>IDAGetNumNonI</i>
No. of nonlinear convergence failures	<i>IDAGetNumNonI</i>
IDA nonlinear solver statistics	<i>IDAGetNonlinS</i>
User data pointer	<i>IDAGetUserDat</i>
Array showing roots found	<i>IDAGetRootIn</i>
No. of calls to user root function	<i>IDAGetNumGEv</i>
Print all statistics	<i>IDAPrintAllSt</i>
Name of constant associated with a return flag	<i>IDAGetReturnF</i>
Number of backtrack operations	<i>IDAGetNumBack</i>
Corrected initial conditions	<i>IDAGetConsist</i>
Stored Jacobian of the DAE residual function	<i>IDAGetJac()</i>
$c_j$ value used in the Jacobian evaluation	<i>IDAGetJacCj()</i>

Table 5.7 – continued from previous page

Optional output	Function name
Time at which the Jacobian was evaluated	<i>IDAGetJacTime</i>
Step number at which the Jacobian was evaluated	<i>IDAGetJacNumS</i>
Size of real and integer workspace	<i>IDAGetLinWork</i>
No. of Jacobian evaluations	<i>IDAGetNumJacE</i>
No. of residual calls for finite diff. Jacobian-vector evals.	<i>IDAGetNumLinF</i>
No. of linear iterations	<i>IDAGetNumLinI</i>
No. of linear convergence failures	<i>IDAGetNumLinC</i>
No. of preconditioner evaluations	<i>IDAGetNumPrec</i>
No. of preconditioner solves	<i>IDAGetNumPrec</i>
No. of Jacobian-vector setup evaluations	<i>IDAGetNumJTSe</i>
No. of Jacobian-vector product evaluations	<i>IDAGetNumJtin</i>
Last return from a linear solver function	<i>IDAGetLastLin</i>
Name of constant associated with a return flag	<i>IDAGetLinRetu</i>

### Main solver optional output functions

IDAS provides several user-callable functions that can be used to obtain different quantities that may be of interest to the user, such as solver workspace requirements, solver performance statistics, as well as additional data from the IDAS solver object (a suggested tolerance scaling factor, the error weight vector, and the vector of estimated local errors). Also provided are functions to extract statistics related to the performance of the nonlinear solver being used. As a convenience, additional extraction functions provide the optional outputs in groups. These optional output functions are described next.

int **IDAGetWorkSpace**(void \*ida\_mem, long int \*lenrw, long int \*leniw)

The function *IDAGetWorkSpace()* returns the IDAS real and integer workspace sizes.

#### Arguments:

- *ida\_mem* – pointer to the IDAS solver object.
- *lenrw* – number of real values in the IDAS workspace.
- *leniw* – number of integer values in the IDAS workspace.

#### Return value:

- *IDA\_SUCCESS* – The optional output value has been successfully set.
- *IDA\_MEM\_NULL* – The *ida\_mem* pointer is NULL.

#### Notes:

In terms of the problem size  $N$ , the maximum method order *maxord*, and the number of root functions *nrtfn* (see §5.1.3.8), the actual size of the real workspace, in *sunrealtype* words, is given by the following:

- base value:  $\text{lenrw} = 55 + (m + 6) * N_r + 3 * \text{nrtfn}$ ;
- with *IDASvtolerances()*:  $\text{lenrw} = \text{lenrw} + N_r$ ;
- with constraint checking (see *IDASetConstraints()*):  $\text{lenrw} = \text{lenrw} + N_r$ ;
- with *id* specified (see *IDASetId()*):  $\text{lenrw} = \text{lenrw} + N_r$ ;

where  $m = \max(\text{maxord}, 3)$ , and  $N_r$  is the number of real words in one *N\_Vector* ( $\approx N$ ).

The size of the integer workspace (without distinction between *int* and *long int* words) is given by:



- base value:  $\text{leniw} = 38 + (m + 6) * N_i + \text{nrtfn}$ ;
- with `IDASVtolerances()`:  $\text{leniw} = \text{leniw} + N_i$ ;
- with constraint checking:  $\text{lenrw} = \text{lenrw} + N_i$ ;
- with id specified (see `IDASetId()`):  $\text{lenrw} = \text{lenrw} + N_i$ ;

where  $N_i$  is the number of integer words in one `N_Vector` (= 1 for the serial `N_Vector` and  $2 * \text{nps}$  for the parallel `N_Vector` on `nps` processors). For the default value of `maxord`, with no rootfinding, no `id`, no constraints, and with no call to `IDASVtolerances()`, these lengths are given roughly by  $\text{lenrw} = 55 + 11 * N$  and  $\text{leniw} = 49$ .

Note that additional memory is allocated if quadratures and/or forward sensitivity integration is enabled. See §5.2.1 and §5.4.2.1 for more details.

**int IDAGetNumSteps**(void \*ida\_mem, long int \*nsteps)

The function `IDAGetNumSteps()` returns the cumulative number of internal steps taken by the solver (total so far).

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `nsteps` – number of steps taken by IDAS.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

**int IDAGetNumResEvals**(void \*ida\_mem, long int \*nrevals)

The function `IDAGetNumResEvals()` returns the number of calls to the user's residual evaluation function.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `nrevals` – number of calls to the user's `res` function.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

**Notes:**

The `nrevals` value returned by `IDAGetNumResEvals()` does not account for calls made to `res` from a linear solver or preconditioner module.

**int IDAGetNumLinSolvSetups**(void \*ida\_mem, long int \*nlinsetups)

The function `IDAGetNumLinSolvSetups()` returns the cumulative number of calls made to the linear solver's setup function (total so far).

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `nlinsetups` – number of calls made to the linear solver setup function.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

int **IDAGetNumErrTestFails**(void \*ida\_mem, long int \*netfails)

The function *IDAGetNumErrTestFails()* returns the cumulative number of local error test failures that have occurred (total so far).

**Arguments:**

- *ida\_mem* – pointer to the IDAS solver object.
- *netfails* – number of error test failures.

**Return value:**

- *IDA\_SUCCESS* – The optional output value has been successfully set.
- *IDA\_MEM\_NULL* – The *ida\_mem* pointer is NULL.

int **IDAGetNumStepSolveFails**(void \*ida\_mem, long int \*ncnf)

Returns the number of failed steps due to a nonlinear solver failure.

**Arguments:**

- *ida\_mem* – pointer to the IDAS solver object.
- *ncnf* – number of step failures.

**Return value:**

- *IDA\_SUCCESS* – The optional output value has been successfully set.
- *IDA\_MEM\_NULL* – The *ida\_mem* pointer is NULL.

int **IDAGetLastOrder**(void \*ida\_mem, int \*klast)

The function *IDAGetLastOrder()* returns the integration method order used during the last internal step.

**Arguments:**

- *ida\_mem* – pointer to the IDAS solver object.
- *klast* – method order used on the last internal step.

**Return value:**

- *IDA\_SUCCESS* – The optional output value has been successfully set.
- *IDA\_MEM\_NULL* – The *ida\_mem* pointer is NULL.

int **IDAGetCurrentOrder**(void \*ida\_mem, int \*kcur)

The function *IDAGetCurrentOrder()* returns the integration method order to be used on the next internal step.

**Arguments:**

- *ida\_mem* – pointer to the IDAS solver object.
- *kcur* – method order to be used on the next internal step.

**Return value:**

- *IDA\_SUCCESS* – The optional output value has been successfully set.
- *IDA\_MEM\_NULL* – The *ida\_mem* pointer is NULL.

int **IDAGetLastStep**(void \*ida\_mem, *sunrealtype* \*hlast)

The function *IDAGetLastStep()* returns the integration step size taken on the last internal step (if from *IDA-Solve()*), or the last value of the artificial step size *h* (if from *IDACalcIC()*).

**Arguments:**

- *ida\_mem* – pointer to the IDAS solver object.

- `hlast` – step size taken on the last internal step by IDAS, or last artificial step size used in `IDACalcIC()`, whichever was called last.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

int **IDAGGetCurrentStep**(void \*ida\_mem, *sunrealtype* \*hcur)

The function `IDAGGetCurrentStep()` returns the integration step size to be attempted on the next internal step.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `hcur` – step size to be attempted on the next internal step.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

int **IDAGGetActualInitStep**(void \*ida\_mem, *sunrealtype* \*hinused)

The function `IDAGGetActualInitStep()` returns the value of the integration step size used on the first step.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `hinused` – actual value of initial step size.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

**Notes:**

Even if the value of the initial integration step size was specified by the user through a call to `IDASetInitStep()`, this value might have been changed by IDAS to ensure that the step size is within the prescribed bounds ( $h_{min} \leq h_0 \leq h_{max}$ ), or to meet the local error test.

int **IDAGGetCurrentTime**(void \*ida\_mem, *sunrealtype* \*tcur)

The function `IDAGGetCurrentTime()` returns the current internal time reached by the solver.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `tcur` – current internal time reached.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

int **IDAGGetTolScaleFactor**(void \*ida\_mem, *sunrealtype* \*tolsfac)

The function `IDAGGetTolScaleFactor()` returns a suggested factor by which the user's tolerances should be scaled when too much accuracy has been requested for some internal step.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.

- `tolsfac` – suggested scaling factor for user tolerances.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

int **IDAGetErrWeights**(void \*ida\_mem, *N\_Vector* eweight)

The function *IDAGetErrWeights()* returns the solution error weights at the current time. These are the  $W_i$  given by (2.5) (or by the user's *IDAEwtFn*).

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `eweight` – solution error weights at the current time.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

**Warning:** The user must allocate space for `eweight`.

int **IDAGetEstLocalErrors**(void \*ida\_mem, *N\_Vector* ele)

The function *IDAGetEstLocalErrors()* returns the estimated local errors.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `ele` – estimated local errors at the current time.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

**Warning:** The user must allocate space for `ele`. The values returned in `ele` are only valid if *IDASolve()* returned a non-negative value.

---

**Note:** The `ele` vector, together with the `eweight` vector from *IDAGetErrWeights()*, can be used to determine how the various components of the system contributed to the estimated local error test. Specifically, that error test uses the RMS norm of a vector whose components are the products of the components of these two vectors. Thus, for example, if there were recent error test failures, the components causing the failures are those with largest values for the products, denoted loosely as `eweight[i]*ele[i]`.

---

int **IDAGetIntegratorStats**(void \*ida\_mem, long int \*nsteps, long int \*nrevals, long int \*nlinsetups, long int \*netfails, int \*klast, int \*kcur, *sunrealtype* \*hinused, *sunrealtype* \*hlast, *sunrealtype* \*hcur, *sunrealtype* \*tcur)

The function *IDAGetIntegratorStats()* returns the IDAS integrator stats in one function call.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.

- `nsteps` – cumulative number of steps taken by IDAS.
- `nrevals` – cumulative number of calls to the user's `res` functions.
- `nlinsetups` – cumulative number of calls made to the linear solver setup function.
- `netfails` – cumulative number of error test failures.
- `klast` – method order used on the last internal step.
- `kcur` – method order to be used on the next internal step.
- `hinused` – actual value of initial step size.
- `hlast` – step sized taken on the last internal step.
- `hcur` – step size to be attempted on the next internal step.
- `tcur` – current internal time reached.

**Return value:**

- `IDA_SUCCESS` – The optional output values have been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

int **IDAGetNumNonlinSolvIters**(void \*ida\_mem, long int \*nniters)

The function *IDAGetNumNonlinSolvIters()* returns the cumulative number of nonlinear iterations performed.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `nniters` – number of nonlinear iterations performed.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_MEM_FAIL` – The `SUNNonlinearSolver` object is NULL.

int **IDAGetNumNonlinSolvConvFails**(void \*ida\_mem, long int \*nncfails)

The function *IDAGetNumNonlinSolvConvFails()* returns the cumulative number of nonlinear convergence failures that have occurred.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `nncfails` – number of nonlinear convergence failures.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

int **IDAGetNonlinSolvStats**(void \*ida\_mem, long int \*nniters, long int \*nncfails)

The function *IDAGetNonlinSolvStats()* returns the IDAS nonlinear solver statistics as a group.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `nniters` – cumulative number of nonlinear iterations performed.
- `nncfails` – cumulative number of nonlinear convergence failures.

**Return value:**

- IDA\_SUCCESS – The optional output value has been successfully set.
- IDA\_MEM\_NULL – The `ida_mem` pointer is NULL.
- IDA\_MEM\_FAIL – The `SUNNonlinearSolver` object is NULL.

int **IDAGetUserData**(void \*ida\_mem, void \*\*user\_data)

The function `IDAGetUserData` returns the user data pointer provided to `IDASetUserData()`.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `user_data` – memory reference to a user data pointer.

**Return value:**

- IDA\_SUCCESS – The optional output value has been successfully set.
- IDA\_MEM\_NULL – The `ida_mem` pointer is NULL.

New in version 5.3.0.

int **IDAPrintAllStats**(void \*ida\_mem, FILE \*outfile, *SUNOutputFormat* fmt)

The function `IDAPrintAllStats` outputs all of the integrator, nonlinear solver, linear solver, and other statistics.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `outfile` – pointer to output file.
- `fmt` – the output format:
  - *SUN\_OUTPUTFORMAT\_TABLE* – prints a table of values
  - *SUN\_OUTPUTFORMAT\_CSV* – prints a comma-separated list of key and value pairs e.g., `key1, value1, key2, value2, ...`

**Return value:**

- IDA\_SUCCESS – The output was successfully.
- IDA\_MEM\_NULL – The `ida_mem` pointer is NULL.
- IDA\_ILL\_INPUT – An invalid formatting option was provided.

---

**Note:** The file `scripts/sundials_csv.py` provides python utility functions to read and output the data from a SUNDIALS CSV output file using the key and value pair format.

---

New in version 5.2.0.

char \***IDAGetReturnFlagName**(long int flag)

The function `IDAGetReturnFlagName()` returns the name of the IDAS constant corresponding to `flag`.

**Arguments:**

- `flag` – the flag returned by a call to an IDAS function.

**Return value:**

- `char*` – the flag name string.

### Initial condition calculation optional output functions

int **IDAGetNumBacktrackOps**(void \*ida\_mem, long int \*nbacktr)

The function *IDAGetNumBacktrackOps()* returns the number of backtrack operations done in the linesearch algorithm in *IDACalcIC()*.

**Arguments:**

- *ida\_mem* – pointer to the IDAS solver object.
- *nbacktr* – the cumulative number of backtrack operations.

**Return value:**

- IDA\_SUCCESS – The optional output value has been successfully set.
- IDA\_MEM\_NULL – The *ida\_mem* pointer is NULL.

int **IDAGetConsistentIC**(void \*ida\_mem, *N\_Vector* yy0\_mod, *N\_Vector* yp0\_mod)

The function *IDAGetConsistentIC()* returns the corrected initial conditions calculated by *IDACalcIC()*.

**Arguments:**

- *ida\_mem* – pointer to the IDAS solver object.
- *yy0\_mod* – consistent solution vector.
- *yp0\_mod* – consistent derivative vector.

**Return value:**

- IDA\_SUCCESS – The optional output value has been successfully set.
- IDA\_ILL\_INPUT – The function was not called before the first call to *IDASolve()*.
- IDA\_MEM\_NULL – The *ida\_mem* pointer is NULL.

**Notes:**

If the consistent solution vector or consistent derivative vector is not desired, pass NULL for the corresponding argument.

**Warning:** The user must allocate space for *yy0\_mod* and *yp0\_mod* (if not NULL).

### Rootfinding optional output functions

There are two optional output functions associated with rootfinding.

int **IDAGetRootInfo**(void \*ida\_mem, int \*rootsfound)

The function *IDAGetRootInfo()* returns an array showing which functions were found to have a root.

**Arguments:**

- *ida\_mem* – pointer to the IDAS solver object.
- *rootsfound* – array of length *nrtfn* with the indices of the user functions  $g_i$  found to have a root. For  $i = 0, \dots, \text{nrtfn} - 1$ ,  $\text{rootsfound}[i] \neq 0$  if  $g_i$  has a root, and  $= 0$  if not.

**Return value:**

- IDA\_SUCCESS – The optional output values have been successfully set.
- IDA\_MEM\_NULL – The *ida\_mem* pointer is NULL.

**Notes:**

Note that, for the components  $g_i$  for which a root was found, the sign of `rootsfound[i]` indicates the direction of zero-crossing. A value of +1 indicates that  $g_i$  is increasing, while a value of -1 indicates a decreasing  $g_i$ .

**Warning:** The user must allocate memory for the vector `rootsfound`.

int **IDAGetNumGEvals**(void \*ida\_mem, long int \*ngevals)

The function `IDAGetNumGEvals()` returns the cumulative number of calls to the user root function  $g$ .

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `ngevals` – number of calls to the user's function  $g$  so far.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

**IDALS linear solver interface optional output functions**

The following optional outputs are available from the IDALS modules:

int **IDAGetJac**(void \*ida\_mem, *SUNMatrix* \*J)

Returns the internally stored copy of the Jacobian matrix of the DAE residual function.

**Parameters**

- `ida_mem` – the IDAS memory structure
- `J` – the Jacobian matrix

**Return values**

- `IDALS_SUCCESS` – the output value has been successfully set
- `IDALS_MEM_NULL` – `ida_mem` was NULL
- `IDALS_LMEM_NULL` – the linear solver interface has not been initialized

**Warning:** With linear solvers that overwrite the input Jacobian matrix as part of the linear solver setup (e.g., performing an in-place LU factorization) the matrix returned by `IDAGetJac()` may differ from the matrix returned by the last Jacobian evaluation.

**Warning:** This function is provided for debugging purposes and the values in the returned matrix should not be altered.

int **IDAGetJacCj**(void \*ida\_mem, *sunrealtype* \*cj\_J)

Returns the  $c_j$  value used to compute the internally stored copy of the Jacobian matrix of the DAE residual function.

**Parameters**



- **ida\_mem** – the IDAS memory structure
- **cj\_J** – the  $c_j$  value used in the Jacobian was evaluation

**Return values**

- **IDALS\_SUCCESS** – the output value has been successfully set
- **IDALS\_MEM\_NULL** – **ida\_mem** was NULL
- **IDALS\_LMEM\_NULL** – the linear solver interface has not been initialized

int **IDAGetJacTime**(void \*ida\_mem, *sunrealtype* \*t\_J)

Returns the time at which the internally stored copy of the Jacobian matrix of the DAE residual function was evaluated.

**Parameters**

- **ida\_mem** – the IDAS memory structure
- **t\_J** – the time at which the Jacobian was evaluated

**Return values**

- **IDALS\_SUCCESS** – the output value has been successfully set
- **IDALS\_MEM\_NULL** – **ida\_mem** was NULL
- **IDALS\_LMEM\_NULL** – the linear solver interface has not been initialized

int **IDAGetJacNumSteps**(void \*ida\_mem, long int \*nst\_J)

Returns the value of the internal step counter at which the internally stored copy of the Jacobian matrix of the DAE residual function was evaluated.

**Parameters**

- **ida\_mem** – the IDAS memory structure
- **nst\_J** – the value of the internal step counter at which the Jacobian was evaluated

**Return values**

- **IDALS\_SUCCESS** – the output value has been successfully set
- **IDALS\_MEM\_NULL** – **ida\_mem** was NULL
- **IDALS\_LMEM\_NULL** – the linear solver interface has not been initialized

int **IDAGetLinWorkSpace**(void \*ida\_mem, long int \*lenrwLS, long int \*leniwLS)

The function **IDAGetLinWorkSpace()** returns the sizes of the real and integer workspaces used by the IDALS linear solver interface.

**Arguments:**

- **ida\_mem** – pointer to the IDAS solver object.
- **lenrwLS** – the number of real values in the IDALS workspace.
- **leniwLS** – the number of integer values in the IDALS workspace.

**Return value:**

- **IDALS\_SUCCESS** – The optional output value has been successfully set.
- **IDALS\_MEM\_NULL** – The **ida\_mem** pointer is NULL.
- **IDALS\_LMEM\_NULL** – The IDALS linear solver has not been initialized.

**Notes:**

The workspace requirements reported by this routine correspond only to memory allocated within this interface and to memory allocated by the `SUNLinearSolver` object attached to it. The template Jacobian matrix allocated by the user outside of IDALS is not included in this report.

New in version 3.0.0: Replaces the deprecated functions `IDADlsGetWorkspace` and `IDASpilsGetWorkspace`.

int **IDAGetNumJacEvals**(void \*ida\_mem, long int \*njevals)

The function `IDAGetNumJacEvals()` returns the cumulative number of calls to the IDALS Jacobian approximation function.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `njevals` – the cumulative number of calls to the Jacobian function total so far.

**Return value:**

- `IDALS_SUCCESS` – The optional output value has been successfully set.
- `IDALS_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDALS_LMEM_NULL` – The IDALS linear solver has not been initialized.

New in version 3.0.0: Replaces the deprecated function `IDADlsGetNumJacEvals`.

int **IDAGetNumLinResEvals**(void \*ida\_mem, long int \*nrevalsLS)

The function `IDAGetNumLinResEvals()` returns the cumulative number of calls to the user residual function due to the finite difference Jacobian approximation or finite difference Jacobian-vector product approximation.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `nrevalsLS` – the cumulative number of calls to the user residual function.

**Return value:**

- `IDALS_SUCCESS` – The optional output value has been successfully set.
- `IDALS_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDALS_LMEM_NULL` – The IDALS linear solver has not been initialized.

**Notes:**

The value `nrevalsLS` is incremented only if one of the default internal difference quotient functions is used.

New in version 3.0.0: Replaces the deprecated functions `IDADlsGetNumRhsEvals` and `IDASpilsGetNumRhsEvals`.

int **IDAGetNumLinIters**(void \*ida\_mem, long int \*nliters)

The function `IDAGetNumLinIters()` returns the cumulative number of linear iterations.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `nliters` – the current number of linear iterations.

**Return value:**

- `IDALS_SUCCESS` – The optional output value has been successfully set.
- `IDALS_MEM_NULL` – The `ida_mem` pointer is NULL.

- IDALS\_LMEM\_NULL – The IDALS linear solver has not been initialized.

New in version 3.0.0: Replaces the deprecated function `IDASpilsGetNumLinIters`.

int **IDAGetNumLinConvFails**(void \*ida\_mem, long int \*nlcfails)

The function `IDAGetNumLinConvFails()` returns the cumulative number of linear convergence failures.

**Arguments:**

- ida\_mem – pointer to the IDAS solver object.
- nlcfails – the current number of linear convergence failures.

**Return value:**

- IDALS\_SUCCESS – The optional output value has been successfully set.
- IDALS\_MEM\_NULL – The ida\_mem pointer is NULL.
- IDALS\_LMEM\_NULL – The IDALS linear solver has not been initialized.

New in version 3.0.0: Replaces the deprecated function `IDASpilsGetNumConvFails`.

int **IDAGetNumPrecEvals**(void \*ida\_mem, long int \*npevals)

The function `IDAGetNumPrecEvals()` returns the cumulative number of preconditioner evaluations, i.e., the number of calls made to `psetup`.

**Arguments:**

- ida\_mem – pointer to the IDAS solver object.
- npevals – the cumulative number of calls to `psetup`.

**Return value:**

- IDALS\_SUCCESS – The optional output value has been successfully set.
- IDALS\_MEM\_NULL – The ida\_mem pointer is NULL.
- IDALS\_LMEM\_NULL – The IDALS linear solver has not been initialized.

New in version 3.0.0: Replaces the deprecated function `IDASpilsGetNumPrecEvals`.

int **IDAGetNumPrecSolves**(void \*ida\_mem, long int \*npsolves)

The function `IDAGetNumPrecSolves()` returns the cumulative number of calls made to the preconditioner solve function, `psolve`.

**Arguments:**

- ida\_mem – pointer to the IDAS solver object.
- npsolves – the cumulative number of calls to `psolve`.

**Return value:**

- IDALS\_SUCCESS – The optional output value has been successfully set.
- IDALS\_MEM\_NULL – The ida\_mem pointer is NULL.
- IDALS\_LMEM\_NULL – The IDALS linear solver has not been initialized.

New in version 3.0.0: Replaces the deprecated function `IDASpilsGetNumPrecSolves`.

int **IDAGetNumJTSetupEvals**(void \*ida\_mem, long int \*njtsetup)

The function `IDAGetNumJTSetupEvals()` returns the cumulative number of calls made to the Jacobian-vector product setup function `jtsetup`.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `njtsetup` – the current number of calls to `jntsetup`.

**Return value:**

- `IDALS_SUCCESS` – The optional output value has been successfully set.
- `IDALS_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDALS_LMEM_NULL` – The IDALS linear solver has not been initialized.

New in version 3.0.0: Replaces the deprecated function `IDASpilsGetNumJTSetupEvals`.

int **IDAGetNumJtimesEvals**(void \*ida\_mem, long int \*njvevals)

The function `IDAGetNumJtimesEvals()` returns the cumulative number of calls made to the Jacobian-vector product function, `jtimes`.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `njvevals` – the cumulative number of calls to `jtimes`.

**Return value:**

- `IDALS_SUCCESS` – The optional output value has been successfully set.
- `IDALS_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDALS_LMEM_NULL` – The IDALS linear solver has not been initialized.

New in version 3.0.0: Replaces the deprecated function `IDASpilsGetNumJtimesEvals`.

int **IDAGetLastLinFlag**(void \*ida\_mem, long int \*lsflag)

The function `IDAGetLastLinFlag()` returns the last return value from an IDALS routine.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `lsflag` – the value of the last return flag from an IDALS function.

**Return value:**

- `IDALS_SUCCESS` – The optional output value has been successfully set.
- `IDALS_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDALS_LMEM_NULL` – The IDALS linear solver has not been initialized.

**Notes:**

If the IDALS setup function failed (i.e., `IDASolve()` returned `IDA_LSETUP_FAIL`) when using the `SUNLINSOL_DENSE` or `SUNLINSOL_BAND` modules, then the value of `lsflag` is equal to the column index (numbered from one) at which a zero diagonal element was encountered during the LU factorization of the (dense or banded) Jacobian matrix. If the IDALS setup function failed when using another `SUNLinearSolver` object, then `lsflag` will be `SUNLS_PSET_FAIL_UNREC`, `SUNLS_ASET_FAIL_UNREC`, or `SUN_ERR_EXT_FAIL`. If the IDALS solve function failed (`IDASolve()` returned `IDA_LSOLVE_FAIL`), `lsflag` contains the error return flag from the `SUNLinearSolver` object, which will be one of: `SUN_ERR_ARG_CORRUPTRRUPT`, indicating that the `SUNLinearSolver` memory is NULL; `SUNLS_ATIMES_FAIL_UNREC`, indicating an unrecoverable failure in the  $J * v$  function; `SUNLS_PSOLVE_FAIL_UNREC`, indicating that the preconditioner solve function `psolve` failed unrecoverably; `SUNLS_GS_FAIL`, indicating a failure in the Gram-Schmidt procedure (generated only in SPGMR or SPFGMR); `SUNLS_QRSOL_FAIL`, indicating that the matrix  $R$  was found to be singular during the QR solve phase (SPGMR and SPFGMR only); or `SUN_ERR_EXT_FAIL`, indicating an unrecoverable failure in an external iterative linear solver package.

New in version 3.0.0: Replaces the deprecated functions `IDADlsGetLastFlag` and `IDASpilsGetLastFlag`.

char \***IDAGetLinReturnFlagName**(long int lsflag)

The function `IDAGetLinReturnFlagName()` returns the name of the IDALS constant corresponding to `lsflag`.

**Arguments:**

- `flag` – the flag returned by a call to an IDAS function.

**Return value:**

- `char*` – the flag name string or if  $1 \leq \text{lsflag} \leq N$  (LU factorization failed), this function returns “NONE”.

New in version 3.0.0: Replaces the deprecated functions `IDADlsGetReturnFlagName` and `IDASpilsGetReturnFlagName`.

### 5.1.3.13 IDAS reinitialization function

The function `IDAREInit()` reinitializes the main IDAS solver for the solution of a new problem, where a prior call to `IDAINit()` has been made. The new problem must have the same size as the previous one. `IDAREInit()` performs the same input checking and initializations that `IDAINit()` does, but does no memory allocation, as it assumes that the existing internal memory is sufficient for the new problem. A call to `IDAREInit()` deletes the solution history that was stored internally during the previous integration. Following a successful call to `IDAREInit()`, call `IDASolve()` again for the solution of the new problem.

The use of `IDAREInit()` requires that the maximum method order, `maxord`, is no larger for the new problem than for the problem specified in the last call to `IDAINit()`. In addition, the same `N_Vector` module set for the previous problem will be reused for the new problem.

If there are changes to the linear solver specifications, make the appropriate calls to either the linear solver objects themselves, or to the IDALS interface routines, as described in §5.1.3.5.

If there are changes to any optional inputs, make the appropriate `IDASet***` calls, as described in §5.1.3.10. Otherwise, all solver inputs set previously remain in effect.

One important use of the `IDAREInit()` function is in the treating of jump discontinuities in the residual function. Except in cases of fairly small jumps, it is usually more efficient to stop at each point of discontinuity and restart the integrator with a readjusted DAE model, using a call to `IDAREInit()`. To stop when the location of the discontinuity is known, simply make that location a value of  $t_{\text{out}}$ . To stop when the location of the discontinuity is determined by the solution, use the rootfinding feature. In either case, it is critical that the residual function *not* incorporate the discontinuity, but rather have a smooth extension over the discontinuity, so that the step across it (and subsequent rootfinding, if used) can be done efficiently. Then use a switch within the residual function (communicated through `user_data`) that can be flipped between the stopping of the integration and the restart, so that the restarted problem uses the new values (which have jumped). Similar comments apply if there is to be a jump in the dependent variable vector.

int **IDAREInit**(void \*ida\_mem, sunrealtype t0, N\_Vector y0, N\_Vector yp0)

The function `IDAREInit()` provides required problem specifications and reinitializes IDAS.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `t0` – is the initial value of  $t$ .
- `y0` – is the initial value of  $y$ .
- `yp0` – is the initial value of  $\dot{y}$ .

**Return value:**

- IDA\_SUCCESS – The call to was successful.
- IDA\_MEM\_NULL – The IDAS solver object was not initialized through a previous call to `IDACreate()`.
- IDA\_NO\_MALLOC – Memory space for the IDAS solver object was not allocated through a previous call to `IDAInit()`.
- IDA\_ILL\_INPUT – An input argument to `IDAReInit()` has an illegal value.

**Notes:**

All previously set options are retained but may be updated by calling the appropriate “Set” functions.

If an error occurred, `IDARInit()` also sends an error message to the error handler function.

### 5.1.4 User-supplied functions

The user-supplied functions consist of one function defining the DAE residual, (optionally) a function that handles error and warning messages, (optionally) a function that provides the error weight vector, (optionally) one or two functions that provide Jacobian-related information for the linear solver, and (optionally) one or two functions that define the preconditioner for use in any of the Krylov iteration algorithms.

#### 5.1.4.1 DAE residual function

The user must provide a function of type `IDAResFn` defined as follows:

```
typedef int (*IDAResFn)(sunrealtype tt, N_Vector yy, N_Vector yp, N_Vector rr, void *user_data)
```

This function computes the problem residual for given values of the independent variable  $t$ , state vector  $y$ , and derivative  $\dot{y}$ .

**Arguments:**

- `tt` – is the current value of the independent variable.
- `yy` – is the current value of the dependent variable vector,  $y(t)$ .
- `yp` – is the current value of  $\dot{y}(t)$ .
- `rr` – is the output residual vector  $F(t, y, \dot{y})$ .
- `user_data` – is a pointer to user data, the same as the `user_data` pointer parameter passed to `IDASectUserData()`.

**Return value:**

An `IDAResFn` function type should return a value of 0 if successful, a positive value if a recoverable error occurred (e.g., `yy` has an illegal value), or a negative value if a nonrecoverable error occurred. In the last case, the integrator halts. If a recoverable error occurred, the integrator will attempt to correct and retry.

**Notes:**

A recoverable failure error return from the `IDAResFn` is typically used to flag a value of the dependent variable  $y$  that is “illegal” in some way (e.g., negative where only a non-negative value is physically meaningful). If such a return is made, IDAS will attempt to recover (possibly repeating the nonlinear solve, or reducing the step size) in order to avoid this recoverable error return.

For efficiency reasons, the DAE residual function is not evaluated at the converged solution of the nonlinear solver. Therefore, in general, a recoverable error in that converged value cannot be corrected. (It may be detected when the residual function is called the first time during the following integration step, but a successful step cannot be undone.)

However, if the user program also includes quadrature integration, the state variables can be checked for legality in the call to `IDAQuadRhsFn`, which is called at the converged solution of the nonlinear system,

and therefore IDAS can be flagged to attempt to recover from such a situation. Also, if sensitivity analysis is performed with the staggered method, the DAE residual function is called at the converged solution of the nonlinear system, and a recoverable error at that point can be flagged, and IDAS will then try to correct it.

#### 5.1.4.2 Error weight function

```
typedef int (*IDAwtFn)(N_Vector y, N_Vector ewt, void *user_data)
```

This function computes the WRMS error weights for the vector  $y$ .

##### Arguments:

- $y$  – is the value of the dependent variable vector at which the weight vector is to be computed.
- $ewt$  – is the output vector containing the error weights.
- $user\_data$  – is a pointer to user data, the same as the  $user\_data$  parameter passed to `IDASetUserData()`.

##### Return value:

- 0 – if it the error weights were successfully set.
- -1 – if any error occurred.

##### Notes:

Allocation of memory for  $ewt$  is handled within IDAS.

**Warning:** The error weight vector must have all components positive. It is the user's responsibility to perform this test and return -1 if it is not satisfied.

#### 5.1.4.3 Rootfinding function

If a rootfinding problem is to be solved during the integration of the DAE system, the user must supply a function of type `IDARootFn`, defined as follows:

```
typedef int (*IDARootFn)(sunrealtype t, N_Vector y, N_Vector yp, sunrealtype *gout, void *user_data)
```

This function computes a vector-valued function  $g(t, y, \dot{y})$  such that the roots of the  $nrtfn$  components  $g_i(t, y, \dot{y})$  are to be found during the integration.

##### Arguments:

- $t$  – is the current value of the independent variable.
- $y$  – is the current value of the dependent variable vector,  $y(t)$ .
- $yp$  – is the current value of  $\dot{y}(t)$ , the  $t$  – derivative of  $y$ .
- $gout$  – is the output array, of length  $nrtfn$ , with components  $g_i(t, y, \dot{y})$ .
- $user\_data$  – is a pointer to user data, the same as the  $user\_data$  parameter passed to `IDASetUserData()`.

##### Return value:

0 if successful or non-zero if an error occurred (in which case the integration is halted and `IDASolve()` returns `IDA_RTFUNC_FAIL`).

##### Notes:

Allocation of memory for  $gout$  is handled within IDAS.

#### 5.1.4.4 Jacobian construction (matrix-based linear solvers)

If a matrix-based linear solver module is used (i.e. a non-NULL `SUNMatrix` object was supplied to `IDASSetLinearSolver()`), the user may provide a function of type `IDALsJacFn` defined as follows:

```
typedef int (*IDALsJacFn)(sunrealtype t, sunrealtype c_j, N_Vector y, N_Vector yp, N_Vector r, SUNMatrix Jac,
void *user_data, N_Vector tmp1, N_Vector tmp2, N_Vector tmp3)
```

This function computes the Jacobian matrix  $J$  of the DAE system (or an approximation to it), defined by (2.7).

**Arguments:**

- `tt` – is the current value of the independent variable  $t$ .
- `cj` – is the scalar in the system Jacobian, proportional to the inverse of the step size ( $\alpha$  in (2.7)).
- `yy` – is the current value of the dependent variable vector,  $y(t)$ .
- `yp` – is the current value of  $\dot{y}(t)$ .
- `rr` – is the current value of the residual vector  $F(t, y, \dot{y})$ .
- `Jac` – is the output (approximate) Jacobian matrix (of type `SUNMatrix`),  $J = \frac{\partial F}{\partial y} + cj \frac{\partial F}{\partial \dot{y}}$ .
- `user_data` - is a pointer to user data, the same as the `user_data` parameter passed to `IDASetUserData()`.
- `tmp1`, `tmp2`, and `tmp3` – are pointers to memory allocated for variables of type `N_Vector` which can be used by `IDALsJacFn()` function as temporary storage or work space.

**Return value:**

An `IDALsJacFn` should return 0 if successful, a positive value if a recoverable error occurred, or a negative value if a nonrecoverable error occurred.

In the case of a recoverable error return, the integrator will attempt to recover by reducing the stepsize, and hence changing  $\alpha$  in (2.7).

**Notes:**

Information regarding the structure of the specific `SUNMatrix` structure (e.g., number of rows, upper/lower bandwidth, sparsity type) may be obtained through using the implementation-specific `SUNMatrix` interface functions (see Chapter §7 for details).

With direct linear solvers (i.e., linear solvers with type `SUNLINEARSOLVER_DIRECT`), the Jacobian matrix  $J(t, y, \dot{y})$  is zeroed out prior to calling the user-supplied Jacobian function so only nonzero elements need to be loaded into `Jac`.

With the default nonlinear solver (the native SUNDIALS Newton method), each call to the user's `IDALsJacFn()` function is preceded by a call to the `IDAResFn()` user function with the same  $(t, y, \dot{y})$  arguments. Thus the Jacobian function can use any auxiliary data that is computed and saved during the evaluation of the DAE residual. In the case of a user-supplied or external nonlinear solver, this is also true if the residual function is evaluated prior to calling the linear solver setup function (see §9.1.4 for more information).

If the user's `IDALsJacFn` function uses difference quotient approximations, it may need to access quantities not in the call list. These quantities may include the current stepsize, the error weights, etc. To obtain these, the user will need to add a pointer to `ida_mem` to `user_data` and then use the `IDAGet*` functions described in §5.1.3.12. The unit roundoff can be accessed as `SUN_UNIT_ROUNDOFF` defined in `sundials_types.h`.

**dense:**

A user-supplied dense Jacobian function must load the  $\text{Neq} \times \text{Neq}$  dense matrix `Jac` with an approximation to the Jacobian matrix  $J(t, y, \dot{y})$  at the point  $(tt, yy, yp)$ . The accessor macros `SM_ELEMENT_D` and `SM_COLUMN_D` allow the user to read and write dense matrix elements without making explicit references to the underlying representation of the `SUNMATRIX_DENSE` type. `SM_ELEMENT_D(J, i, j)` references the  $(i, j)$



$j$ )-th element of the dense matrix  $Jac$  (with  $i, j = 0 \dots N - 1$ ). This macro is meant for small problems for which efficiency of access is not a major concern. Thus, in terms of the indices  $m$  and  $n$  ranging from 1 to  $N$ , the Jacobian element  $J_{m,n}$  can be set using the statement `SM_ELEMENT_D(J, m-1, n-1) = Jm,n`. Alternatively, `SM_COLUMN_D(J, j)` returns a pointer to the first element of the  $j$ -th column of  $Jac$  (with  $j = 0 \dots N - 1$ ), and the elements of the  $j$ -th column can then be accessed using ordinary array indexing. Consequently,  $J_{m,n}$  can be loaded using the statements `col_n = SM_COLUMN_D(J, n-1); col_n[m-1] = Jm,n`. For large problems, it is more efficient to use `SM_COLUMN_D` than to use `SM_ELEMENT_D`. Note that both of these macros number rows and columns starting from 0. The `SUNMATRIX_DENSE` type and accessor macros are documented in §7.3.

#### **banded:**

A user-supplied banded Jacobian function must load the  $Neq \times Neq$  banded matrix  $Jac$  with an approximation to the Jacobian matrix  $J(t, y, \dot{y})$  at the point  $(tt, yy, yp)$ . The accessor macros `SM_ELEMENT_B`, `SM_COLUMN_B`, and `SM_COLUMN_ELEMENT_B` allow the user to read and write banded matrix elements without making specific references to the underlying representation of the `SUNMATRIX_BAND` type. `SM_ELEMENT_B(J, i, j)` references the  $(i, j)$ -th element of the banded matrix  $Jac$ , counting from 0. This macro is meant for use in small problems for which efficiency of access is not a major concern. Thus, in terms of the indices  $m$  and  $n$  ranging from 1 to  $N$  with  $(m, n)$  within the band defined by `mupper` and `mlower`, the Jacobian element  $J_{m,n}$  can be loaded using the statement `SM_ELEMENT_B(J, m-1, n-1) = Jm,n`. The elements within the band are those with  $-mupper \leq m-n \leq mlower$ . Alternatively, `SM_COLUMN_B(J, j)` returns a pointer to the diagonal element of the  $j$ -th column of  $Jac$ , and if we assign this address to `sunrealtype *col_j`, then the  $i$ -th element of the  $j$ -th column is given by `SM_COLUMN_ELEMENT_B(col_j, i, j)`, counting from 0. Thus, for  $(m, n)$  within the band,  $J_{m,n}$  can be loaded by setting `col_n = SM_COLUMN_B(J, n-1);` and `SM_COLUMN_ELEMENT_B(col_n, m-1, n-1) = Jm,n`. The elements of the  $j$ -th column can also be accessed via ordinary array indexing, but this approach requires knowledge of the underlying storage for a band matrix of type `SUNMATRIX_BAND`. The array `col_n` can be indexed from  $-mupper$  to `mlower`. For large problems, it is more efficient to use `SM_COLUMN_B` and `SM_COLUMN_ELEMENT_B` than to use the `SM_ELEMENT_B` macro. As in the dense case, these macros all number rows and columns starting from 0. The `SUNMATRIX_BAND` type and accessor macros are documented in §7.6.

#### **sparse:**

A user-supplied sparse Jacobian function must load the  $Neq \times Neq$  compressed-sparse-column or compressed-sparse-row matrix  $Jac$  with an approximation to the Jacobian matrix  $J(t, y, \dot{y})$  at the point  $(tt, yy, yp)$ . Storage for  $Jac$  already exists on entry to this function, although the user should ensure that sufficient space is allocated in  $Jac$  to hold the nonzero values to be set; if the existing space is insufficient the user may reallocate the data and index arrays as needed. The amount of allocated space in a `SUNMATRIX_SPARSE` object may be accessed using the macro `SM_NNZ_S` or the routine `SUNSparseMatrix_NNZ`. The `SUNMATRIX_SPARSE` type and accessor macros are documented in §7.8.

New in version 3.0.0: Replaces the deprecated type `IDADlsJacFn`.

#### **5.1.4.5 Jacobian-vector product (matrix-free linear solvers)**

If a matrix-free linear solver is to be used (i.e., a `NULL`-valued `SUNMatrix` was supplied to `IDASSetLinearSolver()`), the user may provide a function of type `IDALsJacTimesVecFn` in the following form, to compute matrix-vector products  $Jv$ . If such a function is not supplied, the default is a difference quotient approximation to these products.

```
typedef int (*IDALsJacTimesVecFn)(sunrealtype tt, N_Vector yy, N_Vector yp, N_Vector rr, N_Vector v, N_Vector Jv, sunrealtype cj, void *user_data, N_Vector tmp1, N_Vector tmp2)
```

This function computes the product  $Jv$  of the DAE system Jacobian  $J$  (or an approximation to it) and a given vector  $v$ , where  $J$  is defined by (2.7).

#### **Arguments:**

- `tt` – is the current value of the independent variable.

- `yy` – is the current value of the dependent variable vector,  $y(t)$ .
- `yp` – is the current value of  $\dot{y}(t)$ .
- `rr` – is the current value of the residual vector  $F(t, y, \dot{y})$ .
- `v` – is the vector by which the Jacobian must be multiplied to the right.
- `Jv` – is the computed output vector.
- `cj` – is the scalar in the system Jacobian, proportional to the inverse of the step size ( $\alpha$  in (2.7)).
- `user_data` – is a pointer to user data, the same as the `user_data` parameter passed to `IDASetUserData()`.
- `tmp1` and `tmp2` – are pointers to memory allocated for variables of type `N_Vector` which can be used by `IDALsJacTimesVecFn` as temporary storage or work space.

**Return value:**

The value returned by the Jacobian-times-vector function should be 0 if successful. A nonzero value indicates that a nonrecoverable error occurred.

**Notes:**

This function must return a value of `Jv` that uses an approximation to the **current** value of  $J$ , i.e. as evaluated at the current  $(t, y, \dot{y})$ .

If the user's `IDALsJacTimesVecFn()` function uses difference quotient approximations, it may need to access quantities not in the call list. These include the current stepsize, the error weights, etc. To obtain these, the user will need to add a pointer to `ida_mem` to `user_data` and then use the `IDAGet*` functions described in §5.1.3.12. The unit roundoff can be accessed as `SUN_UNIT_ROUNDOFF` defined in `sundials_types.h`.

New in version 3.0.0: Replaces the deprecated type `IDASpilsJacTimesVecFn`.

#### 5.1.4.6 Jacobian-vector product setup (matrix-free linear solvers)

If the user's Jacobian-vector product function requires that any Jacobian-related data be preprocessed or evaluated, then this needs to be done in a user-supplied function of type `IDALsJacTimesSetupFn`, defined as follows:

```
typedef int (*IDALsJacTimesSetupFn)(sunrealtype tt, N_Vector yy, N_Vector yp, N_Vector rr, sunrealtype cj, void *user_data);
```

This function setups any data needed by `Jv` product function (see `IDALsJacTimesVecFn`).

**Arguments:**

- `tt` – is the current value of the independent variable.
- `yy` – is the current value of the dependent variable vector,  $y(t)$ .
- `yp` – is the current value of  $\dot{y}(t)$ .
- `rr` – is the current value of the residual vector  $F(t, y, \dot{y})$ .
- `cj` – is the scalar in the system Jacobian, proportional to the inverse of the step size ( $\alpha$  in (2.7)).
- `user_data` – is a pointer to user data, the same as the `user_data` parameter passed to `IDASetUserData()`.

**Return value:**

The value returned by the Jacobian-vector setup function should be 0 if successful, positive for a recoverable error (in which case the step will be retried), or negative for an unrecoverable error (in which case the integration is halted).

**Notes:**

Each call to the Jacobian-vector product setup function is preceded by a call to the `IDAResFn` user function with the same  $(t, y, \dot{y})$  arguments. Thus, the setup function can use any auxiliary data that is computed and saved during the evaluation of the DAE residual.

If the user's `IDALsJacTimesVecFn` function uses difference quotient approximations, it may need to access quantities not in the call list. These include the current stepsize, the error weights, etc. To obtain these, the user will need to add a pointer to `ida_mem` to `user_data` and then use the `IDAGet*` functions described in §5.1.3.12. The unit roundoff can be accessed as `SUN_UNIT_ROUNDOFF` defined in `sundials_types.h`.

New in version 3.0.0: Replaces the deprecated type `IDASpilsJacTimesSetupFn`

**5.1.4.7 Preconditioner solve (iterative linear solvers)**

If a user-supplied preconditioner is to be used with a `SUNLinearSolver` solver module, then the user must provide a function to solve the linear system  $Pz = r$  where  $P$  is a left preconditioner matrix which approximates (at least crudely) the Jacobian matrix  $J = \partial F / \partial y + cj \partial F / \partial \dot{y}$ . This function must be of type `IDALsPrecSolveFn`, defined as follows:

```
typedef int (*IDALsPrecSolveFn)(sunrealtype tt, N_Vector yy, N_Vector yp, N_Vector rr, N_Vector rvec, N_Vector
zvec, sunrealtype cj, sunrealtype delta, void *user_data)
```

This function solves the preconditioning system  $Pz = r$ .

**Arguments:**

- `tt` – is the current value of the independent variable.
- `yy` – is the current value of the dependent variable vector,  $y(t)$ .
- `yp` – is the current value of  $\dot{y}(t)$ .
- `rr` – is the current value of the residual vector  $F(t, y, \dot{y})$ .
- `rvec` – is the right-hand side vector  $r$  of the linear system to be solved.
- `zvec` – is the computed output vector.
- `cj` – is the scalar in the system Jacobian, proportional to the inverse of the step size ( $\alpha$  in (2.7)).
- `delta` – is an input tolerance to be used if an iterative method is employed in the solution. In that case, the residual vector  $Res = r - Pz$  of the system should be made less than `delta` in weighted  $l_2$  norm, i.e.,  $\sqrt{\sum_i (Res_i \cdot ewt_i)^2} < delta$ . To obtain the `N_Vector ewt`, call `IDAGetErrWeights()`.
- `user_data` – is a pointer to user data, the same as the `user_data` parameter passed to `IDASetUserData()`.

**Return value:**

The value returned by the preconditioner solve function should be 0 if successful, positive for a recoverable error (in which case the step will be retried), or negative for an unrecoverable error (in which case the integration is halted).

#### 5.1.4.8 Preconditioner setup (iterative linear solvers)

If the user's preconditioner requires that any Jacobian-related data be evaluated or preprocessed, then this needs to be done in a user-supplied function of type `IDALSPrecSetupFn`, defined as follows:

```
typedef int (*IDALSPrecSetupFn)(sunrealtype tt, N_Vector yy, N_Vector yp, N_Vector rr, sunrealtype cj, void *user_data)
```

This function solves the preconditioning system  $Pz = r$ .

**Arguments:**

- `tt` – is the current value of the independent variable.
- `yy` – is the current value of the dependent variable vector,  $y(t)$ .
- `yp` – is the current value of  $\dot{y}(t)$ .
- `rr` – is the current value of the residual vector  $F(t, y, \dot{y})$ .
- `cj` – is the scalar in the system Jacobian, proportional to the inverse of the step size ( $\alpha$  in (2.7)).
- `user_data` – is a pointer to user data, the same as the `user_data` parameter passed to `IDASetUserData()`.

**Return value:**

The value returned by the preconditioner setup function should be 0 if successful, positive for a recoverable error (in which case the step will be retried), or negative for an unrecoverable error (in which case the integration is halted).

**Notes:**

With the default nonlinear solver (the native SUNDIALS Newton method), each call to the preconditioner setup function is preceded by a call to the `IDAResFn` user function with the same  $(t, y, \dot{y})$  arguments. Thus the preconditioner setup function can use any auxiliary data that is computed and saved during the evaluation of the DAE residual. In the case of a user-supplied or external nonlinear solver, this is also true if the residual function is evaluated prior to calling the linear solver setup function (see §9.1.4 for more information).

This function is not called in advance of every call to the preconditioner solve function, but rather is called only as often as needed to achieve convergence in the nonlinear solver.

If the user's `IDALSPrecSetupFn` function uses difference quotient approximations, it may need to access quantities not in the call list. These include the current stepsize, the error weights, etc. To obtain these, the user will need to add a pointer to `ida_mem` to `user_data` and then use the `IDAGet*` functions described in §5.1.3.12. The unit roundoff can be accessed as `SUN_UNIT_ROUNDOFF` defined in `sundials_types.h`.

## 5.2 Integration of pure quadrature equations

IDAS allows the DAE system to include *pure quadratures*. In this case, it is more efficient to treat the quadratures separately by excluding them from the nonlinear solution stage. To do this, begin by excluding the quadrature variables from the vectors `yy` and `yp` and the quadrature equations from within `res`. Thus a separate vector `yQ` of quadrature variables is to satisfy  $(d/dt)yQ = f_Q(t, y, \dot{y})$ .

The following is an overview of the sequence of calls in a user's main program in this situation. Steps that are unchanged from the skeleton program presented in §5.1.2 are grayed out and new or modified steps are in bold.

1. Initialize parallel or multi-threaded environment, if appropriate
2. Create the SUNDIALS context object
3. Set vector of initial values

4. Create matrix object
5. Create linear solver object
6. Create nonlinear solver object
7. Create IDAS object
8. Initialize IDAS solver
9. Specify integration tolerances
10. Set linear solver optional inputs
11. Attach linear solver module
12. Attach nonlinear solver module
13. Set nonlinear solver optional inputs

14. **Set vector of initial values for quadrature variables**

Typically, the quadrature variables should be initialized to 0.

15. **Initialize quadrature integration**

Call [\*IDAQuadInit\(\)\*](#) to specify the quadrature equation right-hand side function and to allocate internal memory related to quadrature integration. See §5.2.1 for details.

16. **Set optional inputs for quadrature integration**

Call [\*IDASetQuadErrCon\(\)\*](#) to indicate whether or not quadrature variables should be used in the step size control mechanism, and to specify the integration tolerances for quadrature variables. See §5.2.4 for details.

17. Specify rootfinding problem
18. Set optional inputs
19. Correct initial values
20. Advance solution in time

21. **Extract quadrature variables**

Call [\*IDAGetQuad\(\)\*](#) to obtain the values of the quadrature variables at the current time.

22. Get optional outputs

23. **Get quadrature optional outputs**

Call [\*IDAGetQuad\\*\\*\*](#) functions to obtain optional output related to the integration of quadratures. See §5.2.5 for details.

24. Destroy objects
25. Finalize MPI, if used

### 5.2.1 Quadrature initialization and deallocation functions

The function `IDAQuadInit()` activates integration of quadrature equations and allocates internal memory related to these calculations. The form of the call to this function is as follows:

```
int IDAQuadInit(void *ida_mem, IDAQuadRhsFn rhsQ, N_Vector yQ0)
```

The function `IDAQuadInit()` provides required problem specifications, allocates internal memory, and initializes quadrature integration.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block returned by `IDACreate()`.
- `rhsQ` – is the C function which computes  $f_Q$ , the right-hand side of the quadrature equations. This function has the form  $f(Q_t, yy, yp, rhsQ, user\_data)$  for full details see §5.2.6.
- `yQ0` – is the initial value of  $y_Q$ .

**Return value:**

- `IDA_SUCCESS` – The call to `IDAQuadInit()` was successful.
- `IDA_MEM_NULL` – The IDAS memory was not initialized by a prior call to `IDACreate()`.
- `IDA_MEM_FAIL` – A memory allocation request failed.

**Notes:**

If an error occurred, `IDAQuadInit()` also sends an error message to the error handler function.

In terms of the number of quadrature variables,  $N_q$ , and maximum method order, `maxord`, the size of the real and integer workspaces are increased by  $(\text{maxord} + 5)N_q$ . If `IDAQuadSVtolerances()` is called, the workspaces are further increased by  $N_q$ .

The function `IDAQuadReInit()`, useful during the solution of a sequence of problems of same size, reinitializes the quadrature-related internal memory and must follow a call to `IDAQuadInit()` (and maybe a call to `IDAReInit()`). The number  $N_q$  of quadratures is assumed to be unchanged from the prior call to `IDAQuadInit()`. The call to the `IDAQuadReInit()` function has the following form:

```
int IDAQuadReInit(void *ida_mem, N_Vector yQ0)
```

The function `IDAQuadReInit()` provides required problem specifications and reinitializes the quadrature integration.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `yQ0` – is the initial value of  $y_Q$ .

**Return value:**

- `IDA_SUCCESS` – The call to `IDAReInit()` was successful.
- `IDA_MEM_NULL` – The IDAS memory was not initialized by a prior call to `IDACreate()`.
- `IDA_NO_QUAD` – Memory space for the quadrature integration was not allocated by a prior call to `IDAQuadInit()`.

**Notes:**

If an error occurred, `IDAQuadReInit()` also sends an error message to the error handler function.

```
void IDAQuadFree(void *ida_mem)
```

The function `IDAQuadFree()` frees the memory allocated for quadrature integration.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.

**Return value:**

- The function has no return value.

**Notes:**

In general, `IDAQuadFree()` need not be called by the user as it is invoked automatically by `IDAFree()`.

## 5.2.2 IDAS solver function

Even if quadrature integration was enabled, the call to the main solver function `IDASolve()` is exactly the same. However, in this case the return value `flag` can also be one of the following:

- `IDA_QRHS_FAIL` – The quadrature right-hand side function failed in an unrecoverable manner.
- `IDA_FIRST_QRHS_ERR` – The quadrature right-hand side function failed at the first call.
- `IDA_REP_QRHS_ERR` – Convergence test failures occurred too many times due to repeated recoverable errors in the quadrature right-hand side function. This value will also be returned if the quadrature right-hand side function had repeated recoverable errors during the estimation of an initial step size (assuming the quadrature variables are included in the error tests).

## 5.2.3 Quadrature extraction functions

If quadrature integration has been initialized by a call to `IDAQuadInit()`, or reinitialized by a call to `IDAQuadReInit()`, then IDAS computes both a solution and quadratures at time `t`. However, `IDASolve()` will still return only the solution `y` in `y`. Solution quadratures can be obtained using the following function:

```
int IDAGetQuad(void *ida_mem, sunrealtype *tret, N_Vector yQ)
```

The function `IDAGetQuad()` returns the quadrature solution vector after a successful return from `IDASolve()`.

**Arguments:**

- `ida_mem` – pointer to the memory previously allocated by `IDAInit()`.
- `tret` – the time reached by the solver output.
- `yQ` – the computed quadrature vector.

**Return value:**

- `IDA_SUCCESS` – `IDAGetQuad()` was successful.
- `IDA_MEM_NULL` – `ida_mem` was NULL.
- `IDA_NO_QUAD` – Quadrature integration was not initialized.
- `IDA_BAD_DKY` – `yQ` is NULL.

The function `IDAGetQuadDky()` computes the `k`-th derivatives of the interpolating polynomials for the quadrature variables at time `t`. This function is called by `IDAGetQuad()` with `k = 0` and with the current time at which `IDASolve()` has returned, but may also be called directly by the user.

```
int IDAGetQuadDky(void *ida_mem, sunrealtype t, int k, N_Vector dkyQ)
```

The function `IDAGetQuadDky()` returns derivatives of the quadrature solution vector after a successful return from IDA.

**Arguments:**

- `ida_mem` – pointer to the memory previously allocated by `IDAInit()`.

- $t$  – the time at which quadrature information is requested. The time  $t$  must fall within the interval defined by the last successful step taken by IDAS.
- $k$  – order of the requested derivative. This must be  $\leq k_{\text{last}}$ .
- $\text{dkyQ}$  – the vector containing the derivative. This vector must be allocated by the user.

**Return value:**

- `IDA_SUCCESS` – `IDAGetQuadDky()` succeeded.
- `IDA_MEM_NULL` – The pointer to `ida_mem` was NULL.
- `IDA_NO_QUAD` – Quadrature integration was not initialized.
- `IDA_BAD_DKY` – The vector `dkyQ` is NULL.
- `IDA_BAD_K` –  $k$  is not in the range  $0, 1, \dots, k_{\text{last}}$ .
- `IDA_BAD_T` – The time  $t$  is not in the allowed range.

**Notes:**

In case of an error return, an error message is also sent to the error handler function.

## 5.2.4 Optional inputs for quadrature integration

IDAS provides the following optional input functions to control the integration of quadrature equations.

int **IDASetQuadErrCon**(void \*ida\_mem, *sunbooleantype* errconQ)

The function `IDASetQuadErrCon()` specifies whether or not the quadrature variables are to be used in the step size control mechanism within IDAS. If they are, the user must call either `IDAQuadSStolerances()` or `IDAQuadSVtolerances()` to specify the integration tolerances for the quadrature variables.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `errconQ` – specifies whether quadrature variables are included `SUNTRUE` or not `SUNFALSE` in the error control mechanism.

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_NO_QUAD` – Quadrature integration has not been initialized.

**Notes:**

By default, `errconQ` is set to `SUNFALSE`.

**Warning:** It is illegal to call `IDASetQuadErrCon()` before a call to `IDAQuadInit()`.

If the quadrature variables are part of the step size control mechanism, one of the following functions must be called to specify the integration tolerances for quadrature variables.

int **IDAQuadSStolerances**(void \*ida\_mem, *sunrealtype* reltolQ, *sunrealtype* abstolQ)

The function `IDAQuadSStolerances()` specifies scalar relative and absolute tolerances.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.



- `reltolQ` – tolerances is the scalar relative error tolerance.
- `abstolQ` – is the scalar absolute error tolerance.

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_NO_QUAD` – Quadrature integration was not initialized.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_ILL_INPUT` – One of the input tolerances was negative.

int **IDAQuadSVtolerances**(void \*ida\_mem, *sunrealtype* reltolQ, *N\_Vector* abstolQ)

The function *IDAQuadSVtolerances()* specifies scalar relative and vector absolute tolerances.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `reltolQ` – tolerances is the scalar relative error tolerance.
- `abstolQ` – is the vector absolute error tolerance.

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_NO_QUAD` – Quadrature integration was not initialized.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_ILL_INPUT` – One of the input tolerances was negative.

## 5.2.5 Optional outputs for quadrature integration

IDAS provides the following functions that can be used to obtain solver performance information related to quadrature integration.

int **IDAGetQuadNumRhsEvals**(void \*ida\_mem, long int \*nrhsQevals)

The function *IDAGetQuadNumRhsEvals()* returns the number of calls made to the user's quadrature right-hand side function.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `nrhsQevals` – number of calls made to the user's `rhsQ` function.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_NO_QUAD` – Quadrature integration has not been initialized.

int **IDAGetQuadNumErrTestFails**(void \*ida\_mem, long int \*nQetfails)

The function *IDAGetQuadNumErrTestFails()* returns the number of local error test failures due to quadrature variables.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.

- `nQetfails` – number of error test failures due to quadrature variables.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_NO_QUAD` – Quadrature integration has not been initialized.

int **IDAGetQuadErrWeights**(void \*ida\_mem, *N\_Vector* eQweight)

The function *IDAGetQuadErrWeights()* returns the quadrature error weights at the current time.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `eQweight` – quadrature error weights at the current time.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_NO_QUAD` – Quadrature integration has not been initialized.

**Warning:** The user must allocate memory for `eQweight`. If quadratures were not included in the error control mechanism (through a call to *IDASetQuadErrCon()* with `errconQ = SUNTRUE`), *IDAGetQuadErrWeights()* does not set the `eQweight` vector.

int **IDAGetQuadStats**(void \*ida\_mem, long int \*nrhsQevals, long int \*nQetfails)

The function *IDAGetQuadStats()* returns the IDAS integrator statistics as a group.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `nrhsQevals` – number of calls to the user's `rhsQ` function.
- `nQetfails` – number of error test failures due to quadrature variables.

**Return value:**

- `IDA_SUCCESS` – the optional output values have been successfully set.
- `IDA_MEM_NULL` – the `ida_mem` pointer is NULL.
- `IDA_NO_QUAD` – Quadrature integration has not been initialized.

## 5.2.6 User-supplied function for quadrature integration

For integration of quadrature equations, the user must provide a function that defines the right-hand side of the quadrature equations (in other words, the integrand function of the integral that must be evaluated). This function must be of type *IDAQuadRhsFn()* defined as follows:

```
typedef int (*IDAQuadRhsFn)(sunrealtype tres, N_Vector yy, N_Vector yp, N_Vector rrQ, void *user_data)
```

This function computes the quadrature equation right-hand side for a given value of the independent variable  $t$  and state vectors  $y$  and  $\dot{y}$ .

**Arguments:**

- `t` – is the current value of the independent variable.

- `yy` – is the current value of the dependent variable vector,  $y(t)$  .
- `yp` – is the current value of the dependent variable derivative vector,  $\dot{y}(t)$  .
- `rrQ` – is the output vector  $f_Q(t, y, \dot{y})$  .
- `user_data` – is the `user_data` pointer passed to `IDASSetUserData()` .

**Return value:**

A `IDAQuadRhsFn()` should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and `IDA_QRHS_FAIL` is returned).

**Notes:**

Allocation of memory for `rhsQ` is automatically handled within IDAS.

Both `y` and `rhsQ` are of type `N_Vector`, but they typically have different internal representations. It is the user's responsibility to access the vector data consistently.

There is one situation in which recovery is not possible even if `IDAQuadRhsFn()` function returns a recoverable error flag. This is when this occurs at the very first call to the `IDAQuadRhsFn()` (in which case IDAS returns `IDA_FIRST_QRHS_ERR`).

## 5.3 Preconditioner modules

A principal reason for using a parallel DAE solver such as IDAS lies in the solution of partial differential equations (PDEs). Moreover, the use of a Krylov iterative method for the solution of many such problems is motivated by the nature of the underlying linear system of equations (2.6) that must be solved at each time step. The linear algebraic system is large, sparse, and structured. However, if a Krylov iterative method is to be effective in this setting, then a nontrivial preconditioner needs to be used. Otherwise, the rate of convergence of the Krylov iterative method is usually unacceptably slow. Unfortunately, an effective preconditioner tends to be problem-specific.

However, we have developed one type of preconditioner that treats a rather broad class of PDE-based problems. It has been successfully used for several realistic, large-scale problems [42] and is included in a software module within the IDAS package. This module works with the parallel vector module `NVECTOR_PARALLEL` and generates a preconditioner that is a block-diagonal matrix with each block being a band matrix. The blocks need not have the same number of super- and sub-diagonals, and these numbers may vary from block to block. This Band-Block-Diagonal Preconditioner module is called `IDABBDPRE`.

One way to envision these preconditioners is to think of the domain of the computational PDE problem as being subdivided into  $M$  non-overlapping sub-domains. Each of these sub-domains is then assigned to one of the  $M$  processors to be used to solve the DAE system. The basic idea is to isolate the preconditioning so that it is local to each processor, and also to use a (possibly cheaper) approximate residual function. This requires the definition of a new function  $G(t, y, \dot{y})$  which approximates the function  $F(t, y, \dot{y})$  in the definition of the DAE system (2.1). However, the user may set  $G = F$ . Corresponding to the domain decomposition, there is a decomposition of the solution vectors  $y$  and  $\dot{y}$  into  $M$  disjoint blocks  $y_m$  and  $\dot{y}_m$ , and a decomposition of  $G$  into blocks  $G_m$ . The block  $G_m$  depends on  $y_m$  and  $\dot{y}_m$ , and also on components of  $y_{m'}$  and  $\dot{y}_{m'}$  associated with neighboring sub-domains (so-called ghost-cell data). Let  $\bar{y}_m$  and  $\bar{\dot{y}}_m$  denote  $y_m$  and  $\dot{y}_m$  (respectively) augmented with those other components on which  $G_m$  depends. Then we have

$$G(t, y, \dot{y}) = [G_1(t, \bar{y}_1, \bar{\dot{y}}_1), G_2(t, \bar{y}_2, \bar{\dot{y}}_2), \dots, G_M(t, \bar{y}_M, \bar{\dot{y}}_M)]^T,$$

and each of the blocks  $G_m(t, \bar{y}_m, \bar{\dot{y}}_m)$  is uncoupled from the others.

The preconditioner associated with this decomposition has the form

$$P = \begin{bmatrix} P_1 & & & \\ & P_2 & & \\ & & \ddots & \\ & & & P_M \end{bmatrix}$$

where

$$P_m \approx \frac{\partial G_m}{\partial y_m} + \alpha \frac{\partial G_m}{\partial \dot{y}_m}$$

This matrix is taken to be banded, with upper and lower half-bandwidths `mudq` and `mldq` defined as the number of non-zero diagonals above and below the main diagonal, respectively. The difference quotient approximation is computed using `mudq + mldq + 2` evaluations of  $G_m$ , but only a matrix of bandwidth `mukeep + mlkeep + 1` is retained.

Neither pair of parameters need be the true half-bandwidths of the Jacobians of the local block of  $G$ , if smaller values provide a more efficient preconditioner. Such an efficiency gain may occur if the couplings in the DAE system outside a certain bandwidth are considerably weaker than those within the band. Reducing `mukeep` and `mlkeep` while keeping `mudq` and `mldq` at their true values, discards the elements outside the narrower band. Reducing both pairs has the additional effect of lumping the outer Jacobian elements into the computed elements within the band, and requires more caution and experimentation.

The solution of the complete linear system

$$Px = b$$

reduces to solving each of the equations

$$P_m x_m = b_m$$

and this is done by banded LU factorization of  $P_m$  followed by a banded backsolve.

Similar block-diagonal preconditioners could be considered with different treatment of the blocks  $P_m$ . For example, incomplete LU factorization or an iterative method could be used instead of banded LU factorization.

### 5.3.1 A parallel band-block-diagonal preconditioner module

The `IDABBDPRE` module calls two user-provided functions to construct  $P$ : a required function `Gres` (of type `IDABBDLocalFn`) which approximates the residual function  $G(t, y, \dot{y}) \approx F(t, y, \dot{y})$  and which is computed locally, and an optional function `Gcomm` (of type `IDABBDCommFn`) which performs all inter-process communication necessary to evaluate the approximate residual  $G$ . These are in addition to the user-supplied residual function `res`. Both functions take as input the same pointer `user_data` as passed by the user to `IDASSetUserData()` and passed to the user's function `res`. The user is responsible for providing space (presumably within `user_data`) for components of `yy` and `yp` that are communicated by `Gcomm` from the other processors, and that are then used by `Gres`, which should not do any communication.

```
typedef int (*IDABBDLocalFn)(sunindextype Nlocal, sunrealtype tt, N_Vector yy, N_Vector yp, N_Vector gval, void *user_data)
```

This `Gres` function computes  $G(t, y, \dot{y})$ . It loads the vector `gval` as a function of `tt`, `yy`, and `yp`.

#### Arguments:

- `Nlocal` – is the local vector length.
- `tt` – is the value of the independent variable.
- `yy` – is the dependent variable.

- `yp` – is the derivative of the dependent variable.
- `gval` – is the output vector.
- `user_data` – is a pointer to user data, the same as the `user_data` parameter passed to `IDASetUserData()`.

**Return value:**

An `IDABBDLocalFn` function type should return 0 to indicate success, 1 for a recoverable error, or -1 for a non-recoverable error.

**Notes:**

This function must assume that all inter-processor communication of data needed to calculate `gval` has already been done, and this data is accessible within `user_data`.

The case where  $G$  is mathematically identical to  $F$  is allowed.

```
typedef int (*IDABBDCommFn)(sunindextype Nlocal, sunrealtype tt, N_Vector yy, N_Vector yp, void *user_data)
```

This `Gcomm` function performs all inter-processor communications necessary for the execution of the `Gres` function above, using the input vectors `yy` and `yp`.

**Arguments:**

- `Nlocal` – is the local vector length.
- `tt` – is the value of the independent variable.
- `yy` – is the dependent variable.
- `yp` – is the derivative of the dependent variable.
- `gval` – is the output vector.
- `user_data` – is a pointer to user data, the same as the `user_data` parameter passed to `IDASetUserData()`.

**Return value:**

An `IDABBDCommFn` function type should return 0 to indicate success, 1 for a recoverable error, or -1 for a non-recoverable error.

**Notes:**

The `Gcomm` function is expected to save communicated data in space defined within the structure `user_data`.

Each call to the `Gcomm` function is preceded by a call to the residual function `res` with the same  $(t, y, \dot{y})$  arguments. Thus `Gcomm` can omit any communications done by `res` if relevant to the evaluation of `Gres`. If all necessary communication was done in `res`, then `Gcomm = NULL` can be passed in the call to `IDABBDPrecInit()`.

Besides the header files required for the integration of the DAE problem (see §5.1.1), to use the `IDABBDPRE` module, the main program must include the header file `ida_bbdpre.h` which declares the needed function prototypes.

The following is a summary of the usage of this module and describes the sequence of calls in the user main program. Steps that are unchanged from the user main program presented in §5.1.2 are grayed out and new or modified steps are in bold.

1. Initialize parallel or multi-threaded environment
2. Create the vector of initial values
3. Create matrix object
4. **Create linear solver object**

When creating the iterative linear solver object, specify the use of left preconditioning (`SUN_PREC_LEFT`) as IDAS only supports left preconditioning.

5. Create nonlinear solver object
6. Create IDAS object
7. Initialize IDAS solver
8. Specify integration tolerances
9. Attach the linear solver
10. **Set linear solver optional inputs**

**Warning:** The user should not overwrite the preconditioner setup function or solve function through calls to `IDASetPreconditioner()` optional input function.

#### 11. **Initialize the IDABBDPRE preconditioner module**

Call `IDABBDPrecInit()` to allocate memory and initialize the internal preconditioner data. The last two arguments of `IDABBDPrecInit()` are the two user-supplied functions described above.

12. Attach nonlinear solver module
13. Set nonlinear solver optional inputs
14. Specify rootfinding problem
15. Set optional inputs
16. Advance solution in time

#### 17. **Get optional outputs**

Additional optional outputs associated with IDABBDPRE are available by way of two routines described below, `IDABBDPrecGetWorkSpace()` and `IDABBDPrecGetNumGfnEvals()`.

18. Destroy objects
19. Finalize MPI, if used

The user-callable functions that initialize or re-initialize the IDABBDPRE preconditioner module are described next.

```
int IDABBDPrecInit(void *ida_mem, sunindextype Nlocal, sunindextype mudq, sunindextype mldq, sunindextype
    mukeep, sunindextype mlkeep, sunrealtype dq_rel_yy, IDABBDLocalFn Gres,
    IDABBDCommFn Gcomm);
```

The function `IDABBDPrecInit()` initializes and allocates (internal) memory for the IDABBDPRE preconditioner.

#### **Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `Nlocal` – local vector dimension.
- `mudq` – upper half-bandwidth to be used in the difference-quotient Jacobian approximation.
- `mldq` – lower half-bandwidth to be used in the difference-quotient Jacobian approximation.
- `mukeep` – upper half-bandwidth of the retained banded approximate Jacobian block.
- `mlkeep` – lower half-bandwidth of the retained banded approximate Jacobian block.
- `dq_rel_yy` – the relative increment in components of  $y$  used in the difference quotient approximations. The default is  $dq\_rel\_yy = \sqrt{\text{unit roundoff}}$ , which can be specified by passing `dq_rel_yy = 0.0`.
- `Gres` – the function which computes the local residual approximation  $G(t, y, \dot{y})$ .

- `Gcomm` – the optional function which performs all inter-process communication required for the computation of  $G(t, y, \dot{y})$ .

**Return value:**

- `IDALS_SUCCESS` – The call was successful.
- `IDALS_MEM_NULL` – The `ida_mem` pointer was NULL.
- `IDALS_MEM_FAIL` – A memory allocation request has failed.
- `IDALS_LMEM_NULL` – An IDALS linear solver memory was not attached.
- `IDALS_ILL_INPUT` – The supplied vector implementation was not compatible with the block band preconditioner.

**Notes:**

If one of the half-bandwidths `mudq` or `mldq` to be used in the difference-quotient calculation of the approximate Jacobian is negative or exceeds the value `Nlocal-1`, it is replaced by 0 or `Nlocal-1` accordingly.

The half-bandwidths `mudq` and `mldq` need not be the true half-bandwidths of the Jacobian of the local block of  $G$ , when smaller values may provide a greater efficiency.

Also, the half-bandwidths `mukeep` and `mlkeep` of the retained banded approximate Jacobian block may be even smaller, to reduce storage and computation costs further.

For all four half-bandwidths, the values need not be the same on every processor.

The `IDABBDPRE` module also provides a reinitialization function to allow for a sequence of problems of the same size, with the same linear solver choice, provided there is no change in `local_N`, `mukeep`, or `mlkeep`. After solving one problem, and after calling `IDAREInit()` to re-initialize IDAS for a subsequent problem, a call to `IDABBDPrecReInit()` can be made to change any of the following: the half-bandwidths `mudq` and `mldq` used in the difference-quotient Jacobian approximations, the relative increment `dq_rel_yy`, or one of the user-supplied functions `Gres` and `Gcomm`. If there is a change in any of the linear solver inputs, an additional call to the “Set” routines provided by the `SUNLinearSolver` object, and/or one or more of the corresponding `IDASet***` functions, must also be made (in the proper order).

`int IDABBDPrecReInit(void *ida_mem, sunindextype mudq, sunindextype mldq, sunrealtype dq_rel_yy)`

The function `IDABBDPrecReInit()` reinitializes the `IDABBDPRE` preconditioner.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `mudq` – upper half-bandwidth to be used in the difference-quotient Jacobian approximation.
- `Mldq` – lower half-bandwidth to be used in the difference-quotient Jacobian approximation.
- `dq_rel_yy` – the relative increment in components of  $y$  used in the difference quotient approximations. The default is  $dq\_rel\_yy = \sqrt{\text{unit roundoff}}$ , which can be specified by passing `dq_rel_yy = 0.0`.

**Return value:**

- `IDALS_SUCCESS` – The call was successful.
- `IDALS_MEM_NULL` – The `ida_mem` pointer was NULL.
- `IDALS_LMEM_NULL` – An IDALS linear solver memory was not attached.
- `IDALS_PMEM_NULL` – The function `IDABBDPrecInit()` was not previously called.

**Notes:**

If one of the half-bandwidths `mudq` or `mldq` is negative or exceeds the value `Nlocal - 1`, it is replaced by 0 or `Nlocal - 1`, accordingly.

The following two optional output functions are available for use with the `IDABBDPRE` module:

int **IDABBDP**precGetWorkSpace(void \*ida\_mem, long int \*lenrwBBDP, long int \*leniwBBDP)

The function `IDABBDP`precGetWorkSpace() returns the local sizes of the IDABBDPRE real and integer workspaces.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `lenrwBBDP` – local number of real values in the IDABBDPRE workspace.
- `leniwBBDP` – local number of integer values in the IDABBDPRE workspace.

**Return value:**

- `IDALS_SUCCESS` – The optional output value has been successfully set.
- `IDALS_MEM_NULL` – The `ida_mem` pointer was NULL.
- `IDALS_PMEM_NULL` – The IDABBDPRE preconditioner has not been initialized.

**Notes:**

The workspace requirements reported by this routine correspond only to memory allocated within the IDABBDPRE module (the banded matrix approximation, banded SUNLinearSolver object, temporary vectors). These values are local to each process. The workspaces referred to here exist in addition to those given by the corresponding function `IDAGetLinWorkSpace()`.

int **IDABBDP**precGetNumGfnEvals(void \*ida\_mem, long int \*ngevalsBBDP)

The function `IDABBDP`precGetNumGfnEvals() returns the cumulative number of calls to the user `Gres` function due to the finite difference approximation of the Jacobian blocks used within IDABBDPRE's preconditioner setup function.

**Arguments:**

- `ida_mem` – pointer to the IDAS solver object.
- `ngevalsBBDP` – the cumulative number of calls to the user `Gres` function.

**Return value:**

- `IDALS_SUCCESS` – The optional output value has been successfully set.
- `IDALS_MEM_NULL` – The `ida_mem` pointer was NULL.
- `IDALS_PMEM_NULL` – The IDABBDPRE preconditioner has not been initialized.

In addition to the `ngevalsBBDP` evaluations of `Gres`, the costs associated with IDABBDPRE also includes `nlinsetups` LU factorizations, `nlinsetups` calls to `Gcomm`, `npsolves` banded backsolve calls, and `nrevalsLS` residual function evaluations, where `nlinsetups` is an optional IDAS output (see §5.1.3.12), and `npsolves` and `nrevalsLS` are linear solver optional outputs (see §5.1.3.12).

## 5.4 Using IDAS for Forward Sensitivity Analysis

This chapter describes the use of IDAS to compute solution sensitivities using forward sensitivity analysis. One of our main guiding principles was to design the IDAS user interface for forward sensitivity analysis as an extension of that for IVP integration. Assuming a user main program and user-defined support routines for IVP integration have already been defined, in order to perform forward sensitivity analysis the user only has to insert a few more calls into the main program and (optionally) define an additional routine which computes the residual of the sensitivity systems (2.11). The only departure from this philosophy is due to the `IDAResFn` type definition. Without changing the definition of this type, the only way to pass values of the problem parameters to the ODE residual function is to require the user data structure `f_data` to contain a pointer to the array of real parameters `p`.



IDAS uses various constants for both input and output. These are defined as needed in this chapter, but for convenience are also listed separately in §13.

We begin with a brief overview, in the form of a skeleton user program. Following that are detailed descriptions of the interface to the various user-callable routines and of the user-supplied routines that were not already described in §5.1 or §5.2.

### 5.4.1 A skeleton of the user's main program

The following is a skeleton of the user's main program (or calling program) as an application of IDAS. The user program is to have these steps in the order indicated, unless otherwise noted. For the sake of brevity, we defer many of the details to the later sections. As in §5.1.2, most steps are independent of the `N_Vector`, `SUNMatrix`, `SUNLinearSolver`, and `SUNNonlinearSolver` implementations used. For the steps that are not, refer to Chapters §6, §7, §8, §9 for the specific name of the function to be called or macro to be referenced.

First, note that no additional header files need be included for forward sensitivity analysis beyond those for IVP solution §5.1.2.

Steps that are unchanged from the user main program skeleton in §5.1.2 are grayed out and new or modified steps are in bold.

1. Initialize parallel or multi-threaded environment
2. Create the SUNDIALS context object
3. Set the vector of initial values
4. Create matrix object
5. Create linear solver object
6. Create nonlinear solver object
7. Create IDAS object
8. Initialize IDAS solver
9. Specify integration tolerances
10. Attach linear solver
11. Set linear solver optional inputs
12. Attach nonlinear solver
13. Set nonlinear solver optional inputs
14. **Initialize quadrature integration**

If the quadrature is not sensitivity-dependent, initialize the quadrature integration as described in §5.2. For integrating a problem where the quadrature depends on the forward sensitivities see §5.4.4.

15. **Set the sensitivity initial values**

Call `N_VCloneVectorArray()` to create `N_Vector` arrays `yS0` and `ypS0` to hold the initial values for the sensitivity vectors of  $y$  and sensitivity derivative vectors of  $\dot{y}$ , respectively.

```
yS0 = N_VCloneVectorArray(Ns, y0);
ypS0 = N_VCloneVectorArray(Ns, y0);
```

where `Ns` is the number of parameters with respect to which sensitivities are to be computed and `y0` serves only to provide an `N_Vector` template for cloning.

Then, load initial values for each sensitivity vector `yS0[i]` and sensitivity derivative vector `ypS0[i]` for  $i = 0, \dots, N_s - 1$ .

#### 16. Activate sensitivity calculations

Call `IDASensInit()` to activate forward sensitivity computations and allocate internal memory for IDAS related to sensitivity calculations.

If a sensitivity residual function is *not* provided to `IDASensInit()`, then `IDASetSensParams()` *must* be called after `IDASensInit()` and before `IDASolve()` to provide the array of problem parameters with respect to which the sensitivities are computed. This array must also be attached to the “user data” pointer set with `IDASetUserData()`. Optionally, an array of scaling factors for difference-quotient residual computations and a mask array to select which parameters with respect to which the sensitivities are computed may also be provided to `IDASetSensParams()`.

#### 17. Set sensitivity integration tolerances (optional)

Call `IDASensSStolerances()` or `IDASensSVtolerances()` to set the sensitivity integration tolerances or `IDASensEETolerances()` to have IDAS estimate tolerances for sensitivity variables based on the tolerances supplied for states variables.

If sensitivity tolerances are estimated by IDAS, the results will be more accurate if order of magnitude is provided by setting the `pbar` input to `IDASetSensParams()`.

#### 18. Create sensitivity nonlinear solver

If using a non-default nonlinear solver (see §5.4.2.3), then create the desired nonlinear solver object by calling the appropriate constructor function defined by the particular `SUNNonlinearSolver` implementation e.g.,

```
NLSSens = SUNNonlinSol_***Sens(...);
```

for the `IDA_SIMULTANEOUS` or `IDA_STAGGERED` options `***` is the name of the nonlinear solver and `...` are constructor specific arguments (see §9 for details).

#### 19. Attach the sensitivity nonlinear solver

If using a non-default nonlinear solver, then initialize the nonlinear solver interface by attaching the nonlinear solver object by calling `IDASetNonlinearSolverSensSim()` when using the `IDA_SIMULTANEOUS` corrector method, `IDASetNonlinearSolverSensStg()` when using the `IDA_STAGGERED` corrector method (see §5.4.2.3 for details).

#### 20. Set sensitivity nonlinear solver optional inputs

Call the appropriate set functions for the selected nonlinear solver module to change optional inputs specific to that nonlinear solver. These *must* be called after `IDASensInit()` if using the default nonlinear solver or after attaching a new nonlinear solver to IDAS, otherwise the optional inputs will be overridden by IDAS defaults. See §9 for more information on optional inputs.

#### 21. Specify rootfinding problem

#### 22. Set optional inputs

Call `IDASens*` routines to change from their default values any optional inputs that control the behavior of IDAS in computing forward sensitivities. See §5.4.2.7 for details.

#### 23. Correct initial values

#### 24. Advance solution in time

#### 25. Extract sensitivity solution

After each successful return from `IDASolve()`, the solution of the original IVP is available in the `y` argument of `IDASolve()`, while the sensitivity solution can be extracted into `yS` and `ypS` (which can be the same as `yS0` and

ypS0) by calling one of the routines *IDAGetSens()*, *IDAGetSens1()*, *IDAGetSensDky()*, or *IDAGetSensDky1()*.

26. Get optional outputs

27. **Deallocate memory**

Upon completion of the integration, deallocate memory for the vectors yS0 and yps0 using *N\_VDestroyVectorArray()*.

28. Finalize MPI, if used

## 5.4.2 User-callable routines for forward sensitivity analysis

This section describes the IDAS functions, in addition to those presented in §5.1.3, that are called by the user to setup and solve a forward sensitivity problem.

### 5.4.2.1 Forward sensitivity initialization and deallocation functions

Activation of forward sensitivity computation is done by calling *IDASensInit()*. The form of the call is as follows:

```
int IDASensInit(void *ida_mem, int Ns, int ism, IDASensResFn fS, N_Vector *yS0, N_Vector *ypS0)
```

The routine *IDASensInit()* activates forward sensitivity computations and allocates internal memory related to sensitivity calculations.

#### Arguments:

- *ida\_mem* – pointer to the IDAS memory block returned by *IDACreate()*.
- *Ns* – the number of sensitivities to be computed.
- *ism* – forward sensitivity analysis!correction strategies a flag used to select the sensitivity solution method. Its value can be *IDA\_SIMULTANEOUS* or *IDA\_STAGGERED* :
  - In the *IDA\_SIMULTANEOUS* approach, the state and sensitivity variables are corrected at the same time. If the default Newton nonlinear solver is used, this amounts to performing a modified Newton iteration on the combined nonlinear system.
  - In the *IDA\_STAGGERED* approach, the correction step for the sensitivity variables takes place at the same time for all sensitivity equations, but only after the correction of the state variables has converged and the state variables have passed the local error test.
- *resS* – is the C function which computes all sensitivity ODE residuals at the same time. For full details see *IDASensResFn*.
- *yS0* – a pointer to an array of *Ns* vectors containing the initial values of the sensitivities of *y*.
- *ypS0* – a pointer to an array of *Ns* vectors containing the initial values of the sensitivities of *y*.

#### Return value:

- *IDA\_SUCCESS* – The call to *IDASensInit()* was successful.
- *IDA\_MEM\_NULL* – The IDAS memory block was not initialized through a previous call to *IDACreate()*.
- *IDA\_MEM\_FAIL* – A memory allocation request has failed.
- *IDA\_ILL\_INPUT* – An input argument to *IDASensInit()* has an illegal value.

**Notes:**

Passing `fs == NULL` indicates using the default internal difference quotient sensitivity residual routine and `IDASetsensParams()` must be called before `IDASolve()`.

If an error occurred, `IDASensInit()` also sends an error message to the error handler function.

In terms of the problem size  $N$ , number of sensitivity vectors  $N_s$ , and maximum method order `maxord`, the size of the real workspace is increased as follows:

- Base value:  $\text{lenrw} = \text{lenrw} + (\text{maxord} + 5)N_sN$
- With `IDASensSVtolerances()`:  $\text{texttlenrw} = \text{lenrw} + N_sN$

the size of the integer workspace is increased as follows:

- Base value:  $\text{leniw} = \text{leniw} + (\text{maxord} + 5)N_sN_i$
- With `IDASensSVtolerances()`:  $\text{leniw} = \text{leniw} + N_sN_i$

where  $N_i$  is the number of integers in one `N_Vector`.

The routine `IDASensReInit()`, useful during the solution of a sequence of problems of same size, reinitializes the sensitivity-related internal memory. The call to it must follow a call to `IDASensInit()` (and maybe a call to `IDAREInit()`). The number `Ns` of sensitivities is assumed to be unchanged since the call to the initialization function. The call to the `IDASensReInit()` function has the form:

```
int IDASensReInit(void *ida_mem, int ism, N_Vector *yS0, N_Vector *ypS0)
```

The routine `IDASensReInit()` reinitializes forward sensitivity computations.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block returned by `IDACreate()`.
- `ism` – forward sensitivity analysis!correction strategies a flag used to select the sensitivity solution method. Its value can be `IDA_SIMULTANEOUS`, `IDA_STAGGERED`, or `IDA_STAGGERED1`.
- `yS0` – a pointer to an array of `Ns` variables of type `N_Vector` containing the initial values of the sensitivities.
- `ypS0` – a pointer to an array of `Ns` variables of type `N_Vector` containing the initial values of the sensitivities of  $\dot{y}$ .

**Return value:**

- `IDA_SUCCESS` – The call to `IDASensReInit()` was successful.
- `IDA_MEM_NULL` – The IDAS memory block was not initialized through a previous call to `IDACreate()`.
- `IDA_NO_SENS` – Memory space for sensitivity integration was not allocated through a previous call to `IDASensInit()`.
- `IDA_ILL_INPUT` – An input argument to `IDASensReInit()` has an illegal value.
- `IDA_MEM_FAIL` – A memory allocation request has failed.

**Notes:**

All arguments of `IDASensReInit()` are the same as those of the functions `IDASensInit()`. If an error occurred, `IDASensReInit()` also sends a message to the error handler function.

To deallocate all forward sensitivity-related memory (allocated in a prior call to `IDASensInit()`), the user must call

void **IDASensFree**(void \*ida\_mem)

The function *IDASensFree()* frees the memory allocated for forward sensitivity computations by a previous call to *IDASensInit()*.

**Arguments:**

- *ida\_mem* – pointer to the IDAS memory block returned by *IDACreate()*.

**Return value:**

- The function has no return value.

**Notes:**

In general, *IDASensFree()* need not be called by the user, as it is invoked automatically by *IDAFree()*.

After a call to *IDASensFree()*, forward sensitivity computations can be reactivated only by calling *IDASensInit()*.

To activate and deactivate forward sensitivity calculations for successive IDAS runs, without having to allocate and deallocate memory, the following function is provided:

int **IDASensToggleOff**(void \*ida\_mem)

The function *IDASensToggleOff()* deactivates forward sensitivity calculations. It does not deallocate sensitivity-related memory.

**Arguments:**

- *ida\_mem* – pointer to the memory previously returned by *IDACreate()*.

**Return value:**

- *IDA\_SUCCESS* – *IDASensToggleOff()* was successful.
- *IDA\_MEM\_NULL* – *ida\_mem* was NULL.

**Notes:**

Since sensitivity-related memory is not deallocated, sensitivities can be reactivated at a later time (using *IDASensReInit()*).

#### 5.4.2.2 Forward sensitivity tolerance specification functions

One of the following three functions must be called to specify the integration tolerances for sensitivities. Note that this call must be made after the call to *IDASensInit()*.

int **IDASensSStolerances**(void \*ida\_mem, *sunrealtype* reltolS, *sunrealtype* \*abstolS)

The function *IDASensSStolerances()* specifies scalar relative and absolute tolerances.

**Arguments:**

- *ida\_mem* – pointer to the IDAS memory block returned by *IDACreate()*.
- *reltolS* – is the scalar relative error tolerance.
- *abstolS* – is a pointer to an array of length *Ns* containing the scalar absolute error tolerances, one for each parameter.

**Return value:**

- *IDA\_SUCCESS* – The call to *IDASStolerances()* was successful.
- *IDA\_MEM\_NULL* – The IDAS memory block was not initialized through a previous call to *IDACreate()*.
- *IDA\_NO\_SENS* – The sensitivity allocation function *IDASensInit()* has not been called.

- IDA\_ILL\_INPUT – One of the input tolerances was negative.

int **IDASensSVtolerances**(void \*ida\_mem, *sunrealtype* reltolS, *N\_Vector* \*abstolS)

The function *IDASensSVtolerances()* specifies scalar relative tolerance and vector absolute tolerances.

**Arguments:**

- ida\_mem – pointer to the IDAS memory block returned by *IDACreate()*.
- reltolS – is the scalar relative error tolerance.
- abstolS – is an array of Ns variables of type *N\_Vector*. The *N\_Vector* from abstolS[is] specifies the vector tolerances for is -th sensitivity.

**Return value:**

- IDA\_SUCCESS – The call to *IDASVtolerances()* was successful.
- IDA\_MEM\_NULL – The IDAS memory block was not initialized through a previous call to *IDACreate()*.
- IDA\_NO\_SENS – The allocation function for sensitivities has not been called.
- IDA\_ILL\_INPUT – The relative error tolerance was negative or an absolute tolerance vector had a negative component.

**Notes:**

This choice of tolerances is important when the absolute error tolerance needs to be different for each component of any vector *yS[i]*.

int **IDASenseEetolerances**(void \*ida\_mem)

When *IDASenseEetolerances()* is called, IDAS will estimate tolerances for sensitivity variables based on the tolerances supplied for states variables and the scaling factors  $\bar{p}$ .

**Arguments:**

- ida\_mem – pointer to the IDAS memory block returned by *IDACreate()*.

**Return value:**

- IDA\_SUCCESS – The call to *IDASenseEetolerances()* was successful.
- IDA\_MEM\_NULL – The IDAS memory block was not initialized through a previous call to *IDACreate()*.
- IDA\_NO\_SENS – The sensitivity allocation function has not been called.

### 5.4.2.3 Forward sensitivity nonlinear solver interface functions

As in the pure DAE case, when computing solution sensitivities using forward sensitivity analysis IDAS uses the SUNNonlinearSolver implementation of Newton's method defined by the SUNNONLINSOL\_NEWTON module (see §9.3) by default. To specify a different nonlinear solver in IDAS, the user's program must create a SUNNonlinearSolver object by calling the appropriate constructor routine. The user must then attach the SUNNonlinearSolver object to IDAS by calling *IDASetNonlinearSolverSensSim()* when using the IDA\_SIMULTANEOUS corrector option, or *IDASetNonlinearSolver()* and *IDASetNonlinearSolverSensStg()* when using the IDA\_STAGGERED corrector option as documented below.

When changing the nonlinear solver in IDAS, *IDASetNonlinearSolver()* must be called after *IDAINit()*; similarly *IDASetNonlinearSolverSensSim()*, *IDASetNonlinearSolverSensStg()*, must be called after *IDASensInit()*. If any calls to *IDASolve()* have been made, then IDAS will need to be reinitialized by calling *IDAREinit()* to ensure that the nonlinear solver is initialized correctly before any subsequent calls to *IDASolve()*.

The first argument passed to the routines *IDASetNonlinearSolverSensSim()*, and *IDASetNonlinearSolverSensStg()*, is the IDAS memory pointer returned by *IDACreate()* and the second argument is the *SUNNonlinearSolver* object to use for solving the nonlinear systems (2.4). A call to this function attaches the nonlinear solver to the main IDAS integrator.

int **IDASetNonlinearSolverSensSim**(void \*ida\_mem, *SUNNonlinearSolver* NLS)

The function *IDASetNonlinearSolverSensSim()* attaches a *SUNNonlinearSolver* object (NLS) to IDAS when using the IDA\_SIMULTANEOUS approach to correct the state and sensitivity variables at the same time.

**Arguments:**

- ida\_mem – pointer to the IDAS memory block.
- NLS – *SUNNonlinearSolver* object to use for solving nonlinear system (2.4).

**Return value:**

- IDA\_SUCCESS – The nonlinear solver was successfully attached.
- IDA\_MEM\_NULL – The ida\_mem pointer is NULL.
- IDA\_ILL\_INPUT – The *SUNNONLINSOL* object is NULL, does not implement the required nonlinear solver operations, is not of the correct type, or the residual function, convergence test function, or maximum number of nonlinear iterations could not be set.

int **IDASetNonlinearSolverSensStg**(void \*ida\_mem, *SUNNonlinearSolver* NLS)

The function *IDASetNonlinearSolverSensStg()* attaches a *SUNNonlinearSolver* object (NLS) to IDAS when using the IDA\_STAGGERED approach to correct all the sensitivity variables after the correction of the state variables.

**Arguments:**

- ida\_mem – pointer to the IDAS memory block.
- NLS – *SUNNONLINSOL* object to use for solving nonlinear systems.

**Return value:**

- IDA\_SUCCESS – The nonlinear solver was successfully attached.
- IDA\_MEM\_NULL – The ida\_mem pointer is NULL.
- IDA\_ILL\_INPUT – The *SUNNONLINSOL* object is NULL, does not implement the required nonlinear solver operations, is not of the correct type, or the residual function, convergence test function, or maximum number of nonlinear iterations could not be set.

**Notes:**

This function only attaches the *SUNNonlinearSolver* object for correcting the sensitivity variables. To attach a *SUNNonlinearSolver* object for the state variable correction use *IDASetNonlinearSolver()*.

#### 5.4.2.4 Forward sensitivity initial condition calculation function

*IDACalcIC()* also calculates corrected initial conditions for sensitivity variables of a DAE system. When used for initial conditions calculation of the forward sensitivities, *IDACalcIC()* must be preceded by successful calls to *IDASensInit()* (or *IDASensReInit()*) and should precede the call(s) to *IDASolve()*. For restrictions that apply for initial conditions calculation of the state variables, see §5.1.3.7.

Calling *IDACalcIC()* is optional. It is only necessary when the initial conditions do not satisfy the sensitivity systems. Even if forward sensitivity analysis was enabled, the call to the initial conditions calculation function *IDACalcIC()* is exactly the same as for state variables.



```
flag = IDACalcIC(ida_mem, icopt, tout1);
```

See [IDACalcIC\(\)](#) for a list of possible return values.

#### 5.4.2.5 IDAS solver function

Even if forward sensitivity analysis was enabled, the call to the main solver function [IDASolve\(\)](#) is exactly the same as in §5.1. However, in this case the return value `flag` can also be one of the following:

- `IDA_SRES_FAIL` – The sensitivity residual function failed in an unrecoverable manner.
- `IDA_REP_SRES_ERR` – The user’s residual function repeatedly returned a recoverable error flag, but the solver was unable to recover.

#### 5.4.2.6 Forward sensitivity extraction functions

If forward sensitivity computations have been initialized by a call to [IDASensInit\(\)](#), or reinitialized by a call to [IDASensReInit\(\)](#), then IDAS computes both a solution and sensitivities at time `t`. However, [IDASolve\(\)](#) will still return only the solution  $y$  in `yout`. Solution sensitivities can be obtained through one of the following functions:

int [IDAGetSens](#)(void \*ida\_mem, *sunrealtype* \*tret, *N\_Vector* \*yS)

The function [IDAGetSens\(\)](#) returns the sensitivity solution vectors after a successful return from [IDASolve\(\)](#).

**Arguments:**

- `ida_mem` – pointer to the memory previously allocated by [IDAInit\(\)](#).
- `tret` – the time reached by the solver output.
- `yS` – array of computed forward sensitivity vectors. This vector array must be allocated by the user.

**Return value:**

- `IDA_SUCCESS` – [IDAGetSens\(\)](#) was successful.
- `IDA_MEM_NULL` – `ida_mem` was NULL.
- `IDA_NO_SENS` – Forward sensitivity analysis was not initialized.
- `IDA_BAD_DKY` – `yS` is NULL.

**Notes:**

Note that the argument `tret` is an output for this function. Its value will be the same as that returned at the last [IDASolve\(\)](#) call.

The function [IDAGetSensDky\(\)](#) computes the  $k$ -th derivatives of the interpolating polynomials for the sensitivity variables at time `t`. This function is called by [IDAGetSens\(\)](#) with  $k = 0$ , but may also be called directly by the user.

int [IDAGetSensDky](#)(void \*ida\_mem, *sunrealtype* t, int k, *N\_Vector* \*dkyS)

The function [IDAGetSensDky\(\)](#) returns derivatives of the sensitivity solution vectors after a successful return from [IDASolve\(\)](#).

**Arguments:**

- `ida_mem` – pointer to the memory previously allocated by [IDAInit\(\)](#).
- `t` – specifies the time at which sensitivity information is requested. The time `t` must fall within the interval defined by the last successful step taken by IDAS.
- `k` – order of derivatives. `k` must be in the range  $0, 1, \dots, klast$  where `klast` is the method order of the last successful step.



- dkyS – array of  $N_s$  vectors containing the derivatives on output. The space for dkyS must be allocated by the user.

**Return value:**

- IDA\_SUCCESS – *IDAGetSensDky()* succeeded.
- IDA\_MEM\_NULL – ida\_mem was NULL.
- IDA\_NO\_SENS – Forward sensitivity analysis was not initialized.
- IDA\_BAD\_DKY – One of the vectors dkyS[i] is NULL.
- IDA\_BAD\_K – k is not in the range 0, 1, ..., qlast.
- IDA\_BAD\_T – The time t is not in the allowed range.

Forward sensitivity solution vectors can also be extracted separately for each parameter in turn through the functions *IDAGetSens1()* and *IDAGetSensDky1()*, defined as follows:

int **IDAGetSens1**(void \*ida\_mem, *sunrealtype* \*tret, int is, *N\_Vector* yS)

The function IDAGetSens1 returns the is-th sensitivity solution vector after a successful return from *IDASolve()*.

**Arguments:**

- ida\_mem – pointer to the memory previously allocated by *IDAINIT()*.
- tret – the time reached by the solver output.
- is – specifies which sensitivity vector is to be returned  $0 \leq is < N_s$ .
- yS – the computed forward sensitivity vector. This vector array must be allocated by the user.

**Return value:**

- IDA\_SUCCESS – IDAGetSens1 was successful.
- IDA\_MEM\_NULL – ida\_mem was NULL.
- IDA\_NO\_SENS – Forward sensitivity analysis was not initialized.
- IDA\_BAD\_IS – The index is is not in the allowed range.
- IDA\_BAD\_DKY – yS is NULL.
- IDA\_BAD\_T – The time t is not in the allowed range.

**Notes:**

Note that the argument tret is an output for this function. Its value will be the same as that returned at the last *IDASolve()* call.

int **IDAGetSensDky1**(void \*ida\_mem, *sunrealtype* t, int k, int is, *N\_Vector* dkyS)

The function IDAGetSensDky1 returns the k-th derivative of the is-th sensitivity solution vector after a successful return from *IDASolve()*.

**Arguments:**

- ida\_mem – pointer to the memory previously allocated by *IDAINIT()*.
- t – specifies the time at which sensitivity information is requested. The time t must fall within the interval defined by the last successful step taken by IDAS.
- k – order of derivative.
- is – specifies the sensitivity derivative vector to be returned  $0 \leq is < N_s$ .
- dkyS – the vector containing the derivative. The space for dkyS must be allocated by the user.

**Return value:**

- IDA\_SUCCESS – IDAGetQuadDky1 succeeded.
- IDA\_MEM\_NULL – The pointer to `ida_mem` was NULL.
- IDA\_NO\_SENS – Forward sensitivity analysis was not initialized.
- IDA\_BAD\_DKY – `dkyS` or one of the vectors `dkyS[i]` is NULL.
- IDA\_BAD\_IS – The index `is` is not in the allowed range.
- IDA\_BAD\_K – `k` is not in the range  $0, 1, \dots, qlast$ .
- IDA\_BAD\_T – The time `t` is not in the allowed range.

**5.4.2.7 Optional inputs for forward sensitivity analysis**

Optional input variables that control the computation of sensitivities can be changed from their default values through calls to `IDASetsens*` functions. Table 5.8 lists all forward sensitivity optional input functions in IDAS which are described in detail in the remainder of this section.

We note that, on an error return, all of the optional input functions send an error message to the error handler function. All error return values are negative, so the test `flag < 0` will catch all errors. Finally, a call to a `IDASetsens***` function can be made from the user's calling program at any time and, if successful, takes effect immediately.

Table 5.8: Forward sensitivity optional inputs :align: center

Optional input	Routine name	Default
Sensitivity scaling factors	<code>IDASetsensParams()</code>	NULL
DQ approximation method	<code>IDASetsensDQMethod()</code>	centered/0.0
Error control strategy	<code>IDASetsensErrCon()</code>	SUNFALSE
Maximum no. of nonlinear iterations	<code>IDASetsensMaxNonlinIters()</code>	4

int **IDASetsensParams**(void \*ida\_mem, *sunrealtype* \*p, *sunrealtype* \*pbar, int \*plist)

The function `IDASetsensParams()` specifies problem parameter information for sensitivity calculations.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `p` – a pointer to the array of real problem parameters used to evaluate  $F(t, y, \dot{y}, p)$ . If non- NULL, `p` must point to a field in the user's data structure `user_data` passed to the residual function.
- `pbar` – an array of `Ns` positive scaling factors. If non- NULL, `pbar` must have all its components  $> 0.0$ .
- `plist` – an array of `Ns` non-negative indices to specify which components `p[i]` to use in estimating the sensitivity equations. If non- NULL, `plist` must have all components  $\geq 0$ .

**Return value:**

- IDA\_SUCCESS – The optional value has been successfully set.
- IDA\_MEM\_NULL – The `ida_mem` pointer is NULL.
- IDA\_NO\_SENS – Forward sensitivity analysis was not initialized.
- IDA\_ILL\_INPUT – An argument has an illegal value.

**Note:** The array `p` only needs to include the parameters with respect to which sensitivities are (potentially) desired.

If the user provides a function to evaluate the sensitivity residuals, `p` need not be specified.

When computing the sensitivity residual via a difference-quotient or estimating sensitivity tolerances the results will be more accurate if order of magnitude information is provided with `pbar`. Typically, if `p[0] != 0`, the value `pbar[i] = abs(p[plist[i]])` can be used. By default IDAS uses `p[i] = 1.0`.

If the user provides a function to evaluate the sensitivity residual and specifies tolerances for the sensitivity variables, `pbar` need not be specified.

By default IDA computes sensitivities with respect to the first  $N_s$  parameters in `p` i.e., `plist[i] = i` for  $i = 0, \dots, N_s-1$ . If sensitivities with respect to the  $j$ -th parameter `p[j]` are desired, set `plist[i] = j` for some  $0 \leq i < N_s$  and  $0 \leq j < N_p$  where  $N_p$  is the number of element in `p`.

If the user provides a function to evaluate the sensitivity residuals, `plist` need not be specified.

**Warning:** This function must be preceded by a call to `IDASensInit()`.

The array `p` *must* also be attached to the user data structure. For example, `user_data->p = p;`

int `IDASetsensDQMethod`(void \*ida\_mem, int DQtype, *sunrealtype* DQrhomax)

The function `IDASetsensDQMethod()` specifies the difference quotient strategy in the case in which the residual of the sensitivity equations are to be computed by IDAS.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `DQtype` – specifies the difference quotient type. Its value can be `IDA_CENTERED` or `IDA_FORWARD`.
- `DQrhomax` – positive value of the selection parameter used in deciding switching between a simultaneous or separate approximation of the two terms in the sensitivity residual.

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_ILL_INPUT` – An argument has an illegal value.

**Notes:**

If `DQrhomax = 0.0`, then no switching is performed. The approximation is done simultaneously using either centered or forward finite differences, depending on the value of `DQtype`. For values of `DQrhomax`  $\geq 1.0$ , the simultaneous approximation is used whenever the estimated finite difference perturbations for states and parameters are within a factor of `DQrhomax`, and the separate approximation is used otherwise. Note that a value `DQrhomax < 1.0` will effectively disable switching. See §2.6 for more details.

The default value are `DQtype == IDA_CENTERED` and `DQrhomax = 0.0`.

int `IDASetsensErrCon`(void \*ida\_mem, *sunbooleantype* errconS)

The function `IDASetsensErrCon()` specifies the error control strategy for sensitivity variables.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.

- `errconS` – specifies whether sensitivity variables are to be included `SUNTRUE` or not `SUNFALSE` in the error control mechanism.

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is `NULL`.

**Notes:**

By default, `errconS` is set to `SUNFALSE`. If `errconS = SUNTRUE` then both state variables and sensitivity variables are included in the error tests. If `errconS = SUNFALSE` then the sensitivity variables are excluded from the error tests. Note that, in any event, all variables are considered in the convergence tests.

int **IDASetsensMaxNonlinIters**(void \*ida\_mem, int maxcorS)

The function *IDASetsensMaxNonlinIters()* specifies the maximum number of nonlinear solver iterations for sensitivity variables per step.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `maxcorS` – maximum number of nonlinear solver iterations allowed per step  $> 0$ .

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is `NULL`.
- `IDA_MEM_FAIL` – The `SUNNONLINSOL` module is `NULL`.

**Notes:**

The default value is 3.

#### 5.4.2.8 Optional outputs for forward sensitivity analysis

Optional output functions that return statistics and solver performance information related to forward sensitivity computations are listed in [Table 5.9](#) and described in detail in the remainder of this section.

Table 5.9: Forward sensitivity optional outputs

Optional output	Routine name
No. of calls to sensitivity residual function	<i>IDAGetsensNumResEvals()</i>
No. of calls to residual function for sensitivity	<i>IDAGetNumResEvalsSens()</i>
No. of sensitivity local error test failures	<i>IDAGetsensNumErrTestFails()</i>
No. of failed steps due to sensitivity nonlinear solver failures	<i>IDAGetNumStepSensSolveFails()</i>
No. of calls to lin. solv. setup routine for sens.	<i>IDAGetsensNumLinSolvSetups()</i>
Error weight vector for sensitivity variables	<i>IDAGetsensErrWeights()</i>
Sensitivity-related statistics as a group	<i>IDAGetsensStats()</i>
No. of sens. nonlinear solver iterations	<i>IDAGetsensNumNonlinSolvIters()</i>
No. of sens. convergence failures	<i>IDAGetsensNumNonlinSolvConvFails()</i>
Sens. nonlinear solver statistics as a group	<i>IDAGetsensNonlinSolvStats()</i>

int **IDAGetsensNumResEvals**(void \*ida\_mem, long int \*nfSevals)

The function *IDAGetsensNumResEvals()* returns the number of calls to the sensitivity residual function.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `nfSevals` – number of calls to the sensitivity residual function.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_NO_SENS` – Forward sensitivity analysis was not initialized.

int **IDAGetNumResEvalsSens**(void \*ida\_mem, long int \*nfevalsS)

The function *IDAGetNumResEvalsSens()* returns the number of calls to the user's residual function due to the internal finite difference approximation of the sensitivity residuals.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `nfevalsS` – number of calls to the user's DAE residual function for the evaluation of sensitivity residuals.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_NO_SENS` – Forward sensitivity analysis was not initialized.

**Notes:**

This counter is incremented only if the internal finite difference approximation routines are used for the evaluation of the sensitivity residuals.

int **IDAGetSensNumErrTestFails**(void \*ida\_mem, long int \*nSetfails)

The function *IDAGetSensNumErrTestFails()* returns the number of local error test failures for the sensitivity variables that have occurred.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `nSetfails` – number of error test failures.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_NO_SENS` – Forward sensitivity analysis was not initialized.

**Notes:**

This counter is incremented only if the sensitivity variables have been included in the error test (see *IDASetSensErrCon()*). Even in that case, this counter is not incremented if the `ism = IDA_SIMULTANEOUS` sensitivity solution method has been used.

int **IDAGetNumStepSensSolveFails**(void \*ida\_mem, long int \*nSncfails)

Returns the number of failed steps due to a sensitivity nonlinear solver failure.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `nSncfails` – number of step failures.

**Return value:**

- IDA\_SUCCESS – The optional output value has been successfully set.
- IDA\_MEM\_NULL – The `ida_mem` pointer is NULL.
- IDA\_NO\_SENS – Forward sensitivity analysis was not initialized.

int **IDAGetSensNumLinSolvSetups**(void \*ida\_mem, long int \*nlinsetupsS)

The function *IDAGetSensNumLinSolvSetups()* returns the number of calls to the linear solver setup function due to forward sensitivity calculations.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `nlinsetupsS` – number of calls to the linear solver setup function.

**Return value:**

- IDA\_SUCCESS – The optional output value has been successfully set.
- IDA\_MEM\_NULL – The `ida_mem` pointer is NULL.
- IDA\_NO\_SENS – Forward sensitivity analysis was not initialized.

**Notes:**

This counter is incremented only if a nonlinear solver requiring a linear solve has been used and the `ism = IDA_STAGGERED` sensitivity solution method has been specified (see §5.4.2.1).

int **IDAGetSensStats**(void \*ida\_mem, long int \*nresSevals, long int \*nresevalsS, long int \*nSetfails, long int \*nlinsetupsS)

The function *IDAGetSensStats()* returns all of the above sensitivity-related solver statistics as a group.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `nresSevals` – number of calls to the sensitivity residual function.
- `nresevalsS` – number of calls to the user-supplied DAE residual function for sensitivity evaluations.
- `nSetfails` – number of error test failures.
- `nlinsetupsS` – number of calls to the linear solver setup function.

**Return value:**

- IDA\_SUCCESS – The optional output values have been successfully set.
- IDA\_MEM\_NULL – The `ida_mem` pointer is NULL.
- IDA\_NO\_SENS – Forward sensitivity analysis was not initialized.

int **IDAGetSensErrWeights**(void \*ida\_mem, *N\_Vector* \*eSweight)

The function *IDAGetSensErrWeights()* returns the sensitivity error weight vectors at the current time. These are the reciprocals of the  $W_i$  of (2.5) for the sensitivity variables.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `eSweight` – pointer to the array of error weight vectors.

**Return value:**

- IDA\_SUCCESS – The optional output value has been successfully set.

- IDA\_MEM\_NULL – The `ida_mem` pointer is NULL.
- IDA\_NO\_SENS – Forward sensitivity analysis was not initialized.

**Notes:**

The user must allocate memory for `eweightS`.

int **IDAGetSensNumNonlinSolvIters**(void \*ida\_mem, long int \*nSniters)

The function *IDAGetSensNumNonlinSolvIters()* returns the number of nonlinear iterations performed for sensitivity calculations.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `nSniters` – number of nonlinear iterations performed.

**Return value:**

- IDA\_SUCCESS – The optional output value has been successfully set.
- IDA\_MEM\_NULL – The `ida_mem` pointer is NULL.
- IDA\_NO\_SENS – Forward sensitivity analysis was not initialized.
- IDA\_MEM\_FAIL – The SUNNONLINSOL module is NULL.

**Notes:**

This counter is incremented only if `ism` was IDA\_STAGGERED or in the call to *IDASensInit()*.

int **IDAGetSensNumNonlinSolvConvFails**(void \*ida\_mem, long int \*nSncfails)

The function *IDAGetSensNumNonlinSolvConvFails()* returns the number of nonlinear convergence failures that have occurred for sensitivity calculations.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `nSncfails` – number of nonlinear convergence failures.

**Return value:**

- IDA\_SUCCESS – The optional output value has been successfully set.
- IDA\_MEM\_NULL – The `ida_mem` pointer is NULL.
- IDA\_NO\_SENS – Forward sensitivity analysis was not initialized.

**Notes:**

This counter is incremented only if `ism` was IDA\_STAGGERED or in the call to *IDASensInit()*.

int **IDAGetSensNonlinSolvStats**(void \*ida\_mem, long int \*nSniters, long int \*nSncfails)

The function *IDAGetSensNonlinSolvStats()* returns the sensitivity-related nonlinear solver statistics as a group.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `nSniters` – number of nonlinear iterations performed.
- `nSncfails` – number of nonlinear convergence failures.

**Return value:**

- IDA\_SUCCESS – The optional output values have been successfully set.
- IDA\_MEM\_NULL – The `ida_mem` pointer is NULL.

- IDA\_NO\_SENS – Forward sensitivity analysis was not initialized.
- IDA\_MEM\_FAIL – The SUNNONLINSOL module is NULL.

#### 5.4.2.9 Initial condition calculation optional output functions

The sensitivity consistent initial conditions found by IDAS (after a successful call to *IDACalcIC()*) can be obtained by calling the following function:

int **IDAGetSensConsistentIC**(void \*ida\_mem, *N\_Vector* \*yyS0\_mod, *N\_Vector* \*ypS0\_mod)

The function *IDAGetSensConsistentIC()* returns the corrected initial conditions calculated by *IDACalcIC()* for sensitivities variables.

**Arguments:**

- ida\_mem – pointer to the IDAS memory block.
- yyS0\_mod – a pointer to an array of Ns vectors containing consistent sensitivity vectors.
- ypS0\_mod – a pointer to an array of Ns vectors containing consistent sensitivity derivative vectors.

**Return value:**

- IDA\_SUCCESS – *IDAGetSensConsistentIC()* succeeded.
- IDA\_MEM\_NULL – The ida\_mem pointer is NULL.
- IDA\_NO\_SENS – The function *IDASensInit()* has not been previously called.
- IDA\_ILL\_INPUT – *IDASolve()* has been already called.

**Notes:**

If the consistent sensitivity vectors or consistent derivative vectors are not desired, pass NULL for the corresponding argument.

**Warning:** The user must allocate space for yyS0\_mod and ypS0\_mod (if not NULL).

### 5.4.3 User-supplied routines for forward sensitivity analysis

In addition to the required and optional user-supplied routines described in §5.1.4, when using IDAS for forward sensitivity analysis, the user has the option of providing a routine that calculates the residual of the sensitivity equations (2.11).

By default, IDAS uses difference quotient approximation routines for the residual of the sensitivity equations. However, IDAS allows the option for user-defined sensitivity residual routines (which also provides a mechanism for interfacing IDAS to routines generated by automatic differentiation).

The user may provide the residuals of the sensitivity equations (2.11) for all sensitivity parameters at once, through a function of type *IDASensResFn* defined by:

```
typedef int (*IDASensResFn)(int Ns, sunrealtype t, N_Vector yy, N_Vector yp, N_Vector resval, N_Vector *yS,  
N_Vector *ypS, N_Vector *resvalS, void *user_data, N_Vector tmp1, N_Vector tmp2, N_Vector tmp3)
```

This function computes the sensitivity residual for all sensitivity equations. It must compute the vectors  $(\partial F / \partial y_i) s_i(t) + (\partial F / \partial \dot{y}) \dot{s}_i(t) + (\partial F / \partial p_i)$  and store them in resvalS[i].

**Arguments:**

- Ns – is the number of sensitivities.



- $t$  – is the current value of the independent variable.
- $yy$  – is the current value of the state vector,  $y(t)$ .
- $yp$  – is the current value of  $\dot{y}(t)$ .
- $resval$  – contains the current value  $F$  of the original DAE residual.
- $yS$  – contains the current values of the sensitivities  $s_i$ .
- $ypS$  – contains the current values of the sensitivity derivatives  $\dot{s}_i$ .
- $resvalS$  – contains the output sensitivity residual vectors. Memory allocation for  $resvalS$  is handled within IDAS.
- $user\_data$  – is a pointer to user data.
- $tmp1, tmp2, tmp3$  – are `N_Vector` s of length  $N$  which can be used as temporary storage.

**Return value:**

An `IDASensResFn()` should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and `IDA_SRES_FAIL` is returned).

**Notes:**

There is one situation in which recovery is not possible even if `IDASensResFn()` function returns a recoverable error flag. That is when this occurs at the very first call to the `IDASensResFn()`, in which case IDAS returns `IDA_FIRST_RES_FAIL`.

## 5.4.4 Integration of quadrature equations depending on forward sensitivities

IDAS provides support for integration of quadrature equations that depends not only on the state variables but also on forward sensitivities.

The following is an overview of the sequence of calls in a user's main program in this situation. Steps that are unchanged from the skeleton program presented in §5.1.2 are grayed out and new or modified steps are in bold. See also §5.2.

1. Initialize parallel or multi-threaded environment, if appropriate
2. Create the SUNDIALS context object
3. Set vector of initial values
4. Create matrix object
5. Create linear solver object
6. Set linear solver optional inputs
7. Create nonlinear solver object
8. Create IDAS object
9. Initialize IDAS solver
10. Specify integration tolerances
11. Attach linear solver
12. Set linear solver optional inputs
13. Attach nonlinear solver
14. Set nonlinear solver optional inputs
15. Set sensitivity initial values

16. Activate sensitivity calculations
17. Set sensitivity integration tolerances
18. Create sensitivity nonlinear solver
19. Attach the sensitivity nonlinear solver
20. Set sensitivity nonlinear solver optional inputs

21. **Set vector of initial values for quadrature variables**

Typically, the quadrature variables should be initialized to 0.

22. **Initialize sensitivity-dependent quadrature integration**

Call `IDAQuadSensInit()` to specify the quadrature equation right-hand side function and to allocate internal memory related to quadrature integration.

23. Specify rootfinding problem

24. **Set optional inputs**

Call `IDASetQuadSensErrCon()` to indicate whether or not quadrature variables should be used in the step size control mechanism. If so, one of the `IDAQuadSens*tolerances` functions must be called to specify the integration tolerances for quadrature variables. See §5.2.4 for details.

25. Correct initial values

26. Advance solution in time

27. Extract sensitivity solution

28. **Extract sensitivity-dependent quadrature variables**

Call `IDAGetQuadSens()`, `IDAGetQuadSens1()`, `IDAGetQuadSensDky()` or `IDAGetQuadSensDky1()` to obtain the values of the quadrature variables or their derivatives at the current time.

29. **Get optional outputs**

Call `IDAQuadSens*` functions to obtain optional output related to the integration of sensitivity-dependent quadratures. See §5.4.4.5 for details.

30. Destroy objects

31. Finalize MPI, if used

#### 5.4.4.1 Sensitivity-dependent quadrature initialization and deallocation

The function `IDAQuadSensInit()` activates integration of quadrature equations depending on sensitivities and allocates internal memory related to these calculations. If `rhsQS` is input as `NULL`, then IDAS uses an internal function that computes difference quotient approximations to the functions  $\bar{q}_i = (\partial q / \partial y) s_i + (\partial q / \partial \dot{y}) \dot{s}_i + \partial q / \partial p_i$ , in the notation of (2.10). The form of the call to this function is as follows:

```
int IDAQuadSensInit(void *ida_mem, IDAQuadSensRhsFn rhsQS, N_Vector *yQS0)
```

The function `IDAQuadSensInit()` provides required problem specifications, allocates internal memory, and initializes quadrature integration.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block returned by `IDACreate()`.
- `rhsQS` – is the `IDAQuadSensRhsFn` function which computes  $f_{QS}$ , the right-hand side of the sensitivity-dependent quadrature equations.

- `yQS0` – contains the initial values of sensitivity-dependent quadratures.

**Return value:**

- `IDA_SUCCESS` – The call to `IDAQuadSensInit()` was successful.
- `IDA_MEM_NULL` – The IDAS memory was not initialized by a prior call to `IDACreate()`.
- `IDA_MEM_FAIL` – A memory allocation request failed.
- `IDA_NO_SENS` – The sensitivities were not initialized by a prior call to `IDASensInit()`.
- `IDA_ILL_INPUT` – The parameter `yQS0` is NULL.

**Notes:**

**Warning:** Before calling `IDAQuadSensInit()`, the user must enable the sensitivities by calling `IDASensInit()`. If an error occurred, `IDAQuadSensInit()` also sends an error message to the error handler function.

In terms of the number of quadrature variables  $N_q$  and maximum method order `maxord`, the size of the real workspace is increased as follows:

- Base value:  $\text{lenrw} = \text{lenrw} + (\text{maxord} + 5)N_q$
- If `IDAQuadSensSVtolerances()` is called:  $\text{lenrw} = \text{lenrw} + N_q N_s$

and the size of the integer workspace is increased as follows:

- Base value:  $\text{leniw} = \text{leniw} + (\text{maxord} + 5)N_q$
- If `IDAQuadSensSVtolerances()` is called:  $\text{leniw} = \text{leniw} + N_q N_s$

The function `IDAQuadSensReInit()`, useful during the solution of a sequence of problems of same size, reinitializes the quadrature related internal memory and must follow a call to `IDAQuadSensInit()`. The number  $N_q$  of quadratures as well as the number  $N_s$  of sensitivities are assumed to be unchanged from the prior call to `IDAQuadSensInit()`. The call to the `IDAQuadSensReInit()` function has the form:

```
int IDAQuadSensReInit(void *ida_mem, N_Vector *yQS0)
```

The function `IDAQuadSensReInit()` provides required problem specifications and reinitializes the sensitivity-dependent quadrature integration.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `yQS0` – contains the initial values of sensitivity-dependent quadratures.

**Return value:**

- `IDA_SUCCESS` – The call to `IDAQuadSensReInit()` was successful.
- `IDA_MEM_NULL` – The IDAS memory was not initialized by a prior call to `IDACreate()`.
- `IDA_NO_SENS` – Memory space for the sensitivity calculation was not allocated by a prior call to `IDASensInit()`.
- `IDA_NO_QUADSENS` – Memory space for the sensitivity quadratures integration was not allocated by a prior call to `IDAQuadSensInit()`.
- `IDA_ILL_INPUT` – The parameter `yQS0` is NULL.

**Notes:**

If an error occurred, `IDAQuadSensReInit()` also sends an error message to the error handler function.

void **IDAQuadSensFree**(void \*ida\_mem);

The function *IDAQuadSensFree()* frees the memory allocated for sensitivity quadrature integration.

**Arguments:**

- *ida\_mem* – pointer to the IDAS memory block.

**Return value:**

There is no return value.

**Notes:**

In general, *IDAQuadSensFree()* need not be called by the user as it is called automatically by *IDAFree()*.

#### 5.4.4.2 IDAS solver function

Even if quadrature integration was enabled, the call to the main solver function *IDASolve()* is exactly the same as in §5.1. However, in this case the return value *flag* can also be one of the following:

- *IDA\_QSRHS\_FAIL* – the sensitivity quadrature right-hand side function failed in an unrecoverable manner.
- *IDA\_FIRST\_QSRHS\_ERR* – the sensitivity quadrature right-hand side function failed at the first call.
- *IDA\_REP\_QSRHS\_ERR* – convergence test failures occurred too many times due to repeated recoverable errors in the quadrature right-hand side function. The *IDA\_REP\_RES\_ERR* will also be returned if the quadrature right-hand side function had repeated recoverable errors during the estimation of an initial step size (assuming the sensitivity quadrature variables are included in the error tests).

#### 5.4.4.3 Sensitivity-dependent quadrature extraction functions

If sensitivity-dependent quadratures have been initialized by a call to *IDAQuadSensInit()*, or reinitialized by a call to *IDAQuadSensReInit()*, then IDAS computes a solution, sensitivities, and quadratures depending on sensitivities at time *t*. However, *IDASolve()* will still return only the solutions *y* and *y*. Sensitivity-dependent quadratures can be obtained using one of the following functions:

int **IDAGetQuadSens**(void \*ida\_mem, *sunrealtype* \*tret, *N\_Vector* \*yQS)

The function *IDAGetQuadSens()* returns the quadrature sensitivity solution vectors after a successful return from *IDASolve()*.

**Arguments:**

- *ida\_mem* – pointer to the memory previously allocated by *IDAInit()*.
- *tret* – the time reached by the solver output.
- *yQS* – array of *Ns* computed sensitivity-dependent quadrature vectors. This array of vectors must be allocated by the user.

**Return value:**

- *IDA\_SUCCESS* – *IDAGetQuadSens()* was successful.
- *IDA\_MEM\_NULL* – *ida\_mem* was NULL.
- *IDA\_NO\_SENS* – Sensitivities were not activated.
- *IDA\_NO\_QUADSENS* – Quadratures depending on the sensitivities were not activated.
- *IDA\_BAD\_DKY* – *yQS* or one of the *yQS[i]* is NULL.

The function *IDAGetQuadSensDky()* computes the *k*-th derivatives of the interpolating polynomials for the sensitivity-dependent quadrature variables at time *t*. This function is called by *IDAGetQuadSens()* with *k* = 0, but may also be called directly by the user.

int **IDAGetQuadSensDky**(void \*ida\_mem, *sunrealtype* t, int k, *N\_Vector* \*dkyQS)

The function *IDAGetQuadSensDky()* returns derivatives of the quadrature sensitivities solution vectors after a successful return from *IDASolve()*.

**Arguments:**

- *ida\_mem* – pointer to the memory previously allocated by *IDAInit()*.
- *t* – the time at which information is requested. The time *t* must fall within the interval defined by the last successful step taken by IDAS.
- *k* – order of the requested derivative. *k* must be in the range  $0, 1, \dots, klast$  where *klast* is the method order of the last successful step.
- *dkyQS* – array of *N<sub>s</sub>* vectors containing the derivatives. This vector array must be allocated by the user.

**Return value:**

- *IDA\_SUCCESS* – *IDAGetQuadSensDky()* succeeded.
- *IDA\_MEM\_NULL* – *ida\_mem* was NULL.
- *IDA\_NO\_SENS* – Sensitivities were not activated.
- *IDA\_NO\_QUADSENS* – Quadratures depending on the sensitivities were not activated.
- *IDA\_BAD\_DKY* – *dkyQS* or one of the vectors *dkyQS[i]* is NULL.
- *IDA\_BAD\_K* – *k* is not in the range  $0, 1, \dots, klast$ .
- *IDA\_BAD\_T* – The time *t* is not in the allowed range.

Quadrature sensitivity solution vectors can also be extracted separately for each parameter in turn through the functions *IDAGetQuadSens1* and *IDAGetQuadSensDky1*, defined as follows:

int **IDAGetQuadSens1**(void \*ida\_mem, *sunrealtype* \*tret, int is, *N\_Vector* yQS)

The function *IDAGetQuadSens1* returns the *is*-th sensitivity of quadratures after a successful return from *IDASolve()*.

**Arguments:**

- *ida\_mem* – pointer to the memory previously allocated by *IDAInit()*.
- *tret* – the time reached by the solver output.
- *is* – specifies which sensitivity vector is to be returned  $0 \leq is < N_s$ .
- *yQS* – the computed sensitivity-dependent quadrature vector. This vector must be allocated by the user.

**Return value:**

- *IDA\_SUCCESS* – *IDAGetQuadSens1* was successful.
- *IDA\_MEM\_NULL* – *ida\_mem* was NULL.
- *IDA\_NO\_SENS* – Forward sensitivity analysis was not initialized.
- *IDA\_NO\_QUADSENS* – Quadratures depending on the sensitivities were not activated.
- *IDA\_BAD\_IS* – The index *is* is not in the allowed range.
- *IDA\_BAD\_DKY* – *yQS* is NULL.

int **IDAGetQuadSensDky1**(void \*ida\_mem, *sunrealtype* t, int k, int is, *N\_Vector* dkyQS)

The function *IDAGetQuadSensDky1* returns the *k*-th derivative of the *is*-th sensitivity solution vector after a successful return from *IDASolve()*.

**Arguments:**

- `ida_mem` – pointer to the memory previously allocated by `IDAInit()`.
- `t` – specifies the time at which sensitivity information is requested. The time `t` must fall within the interval defined by the last successful step taken by IDAS.
- `k` – order of derivative. `k` must be in the range  $0, 1, \dots, klast$  where `klast` is the method order of the last successful step.
- `is` – specifies the sensitivity derivative vector to be returned  $0 \leq is < N_s$ .
- `dkyQS` – the vector containing the derivative. The space for `dkyQS` must be allocated by the user.

**Return value:**

- `IDA_SUCCESS` – `IDAGetQuadDky1` succeeded.
- `IDA_MEM_NULL` – `ida_mem` was NULL.
- `IDA_NO_SENS` – Forward sensitivity analysis was not initialized.
- `IDA_NO_QUADSENS` – Quadratures depending on the sensitivities were not activated.
- `IDA_BAD_DKY` – `dkyQS` is NULL.
- `IDA_BAD_IS` – The index `is` is not in the allowed range.
- `IDA_BAD_K` – `k` is not in the range  $0, 1, \dots, klast$ .
- `IDA_BAD_T` – The time `t` is not in the allowed range.

#### 5.4.4.4 Optional inputs for sensitivity-dependent quadrature integration

IDAS provides the following optional input functions to control the integration of sensitivity-dependent quadrature equations.

int **IDASetQuadSensErrCon**(void \*ida\_mem, *sunbooleantype* errconQS)

The function `IDASetQuadSensErrCon()` specifies whether or not the quadrature variables are to be used in the local error control mechanism. If they are, the user must specify the error tolerances for the quadrature variables by calling `IDAQuadSensStolerances()`, `IDAQuadSensSVtolerances()`, or `IDAQuadSensEETolerances()`.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `errconQS` – specifies whether sensitivity quadrature variables are included `SUNTRUE` or not `SUNFALSE` in the error control mechanism.

**Return value:**

- `IDA_SUCCESS` – The optional value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_NO_SENS` – Sensitivities were not activated.
- `IDA_NO_QUADSENS` – Quadratures depending on the sensitivities were not activated.

**Notes:**

By default, `errconQS` is set to `SUNFALSE`.

**Warning:** It is illegal to call `IDASetQuadSensErrCon()` before a call to `IDAQuadSensInit()`.

If the quadrature variables are part of the step size control mechanism, one of the following functions must be called to specify the integration tolerances for quadrature variables.

int **IDAQuadSensSStolerances**(void \*ida\_mem, *sunrealtype* reltolQS, *sunrealtype* \*abstolQS)

The function *IDAQuadSensSStolerances()* specifies scalar relative and absolute tolerances.

**Arguments:**

- *ida\_mem* – pointer to the IDAS memory block.
- *reltolQS* – tolerances is the scalar relative error tolerance.
- *abstolQS* – is a pointer to an array containing the *Ns* scalar absolute error tolerances.

**Return value:**

- IDA\_SUCCESS – The optional value has been successfully set.
- IDA\_MEM\_NULL – The *ida\_mem* pointer is NULL.
- IDA\_NO\_SENS – Sensitivities were not activated.
- IDA\_NO\_QUADSENS – Quadratures depending on the sensitivities were not activated.
- IDA\_ILL\_INPUT – One of the input tolerances was negative.

int **IDAQuadSensSVtolerances**(void \*ida\_mem, *sunrealtype* reltolQS, *N\_Vector* \*abstolQS)

The function *IDAQuadSensSVtolerances()* specifies scalar relative and vector absolute tolerances.

**Arguments:**

- *ida\_mem* – pointer to the IDAS memory block.
- *reltolQS* – tolerances is the scalar relative error tolerance.
- *abstolQS* – is an array of *Ns* variables of type *N\_Vector*. The *N\_Vector* from *abstols[is]* specifies the vector tolerances for *is*-th quadrature sensitivity.

**Return value:**

- IDA\_SUCCESS – The optional value has been successfully set.
- IDA\_NO\_QUAD – Quadrature integration was not initialized.
- IDA\_MEM\_NULL – The *ida\_mem* pointer is NULL.
- IDA\_NO\_SENS – Sensitivities were not activated.
- IDA\_NO\_QUADSENS – Quadratures depending on the sensitivities were not activated.
- IDA\_ILL\_INPUT – One of the input tolerances was negative.

int **IDAQuadSenseEtolerances**(void \*ida\_mem)

The function *IDAQuadSenseEtolerances()* specifies that the tolerances for the sensitivity-dependent quadratures should be estimated from those provided for the pure quadrature variables.

**Arguments:**

- *ida\_mem* – pointer to the IDAS memory block.

**Return value:**

- IDA\_SUCCESS – The optional value has been successfully set.
- IDA\_MEM\_NULL – The *ida\_mem* pointer is NULL.
- IDA\_NO\_SENS – Sensitivities were not activated.
- IDA\_NO\_QUADSENS – Quadratures depending on the sensitivities were not activated.

**Notes:**

When *IDAQuadSenseEtolerances()* is used, before calling *IDASolve()*, integration of pure quadratures must be initialized (see §5.2) and tolerances for pure quadratures must be also specified (see §5.2.4).

**5.4.4.5 Optional outputs for sensitivity-dependent quadrature integration**

IDAS provides the following functions that can be used to obtain solver performance information related to quadrature integration.

int **IDAGetQuadSensNumRhsEvals**(void \*ida\_mem, long int \*nrhsQSevals)

The function *IDAGetQuadSensNumRhsEvals()* returns the number of calls made to the user's quadrature right-hand side function.

**Arguments:**

- *ida\_mem* – pointer to the IDAS memory block.
- *nrhsQSevals* – number of calls made to the user's *rhsQS* function.

**Return value:**

- *IDA\_SUCCESS* – The optional output value has been successfully set.
- *IDA\_MEM\_NULL* – The *ida\_mem* pointer is NULL.
- *IDA\_NO\_QUADSENS* – Sensitivity-dependent quadrature integration has not been initialized.

int **IDAGetQuadSensNumErrTestFails**(void \*ida\_mem, long int \*nQSetfails)

The function *IDAGetQuadSensNumErrTestFails()* returns the number of local error test failures due to quadrature variables.

**Arguments:**

- *ida\_mem* – pointer to the IDAS memory block.
- *nQSetfails* – number of error test failures due to quadrature variables.

**Return value:**

- *IDA\_SUCCESS* – The optional output value has been successfully set.
- *IDA\_MEM\_NULL* – The *ida\_mem* pointer is NULL.
- *IDA\_NO\_QUADSENS* – Sensitivity-dependent quadrature integration has not been initialized.

int **IDAGetQuadSensErrWeights**(void \*ida\_mem, *N\_Vector* \*eQSweight)

The function *IDAGetQuadSensErrWeights()* returns the quadrature error weights at the current time.

**Arguments:**

- *ida\_mem* – pointer to the IDAS memory block.
- *eQSweight* – array of quadrature error weight vectors at the current time.

**Return value:**

- *IDA\_SUCCESS* – The optional output value has been successfully set.
- *IDA\_MEM\_NULL* – The *ida\_mem* pointer is NULL.
- *IDA\_NO\_QUADSENS* – Sensitivity-dependent quadrature integration has not been initialized.

**Notes:**



**Warning:** The user must allocate memory for `eQSwight`. If quadratures were not included in the error control mechanism (through a call to `IDASetQuadSensErrCon()` with `errconQS=SUNTRUE`), `IDAGetQuadSensErrWeights()` does not set the `eQSwight` vector.

int **IDAGetQuadSensStats**(void \*ida\_mem, long int \*nrhsQSevals, long int \*nQSetfails)

The function `IDAGetQuadSensStats()` returns the IDAS integrator statistics as a group.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `nrhsQSevals` – number of calls to the user's `rhsQS` function.
- `nQSetfails` – number of error test failures due to quadrature variables.

**Return value:**

- `IDA_SUCCESS` – the optional output values have been successfully set.
- `IDA_MEM_NULL` – the `ida_mem` pointer is NULL.
- `IDA_NO_QUADSENS` – Sensitivity-dependent quadrature integration has not been initialized.

#### 5.4.4.6 User-supplied function for sensitivity-dependent quadrature integration

For the integration of sensitivity-dependent quadrature equations, the user must provide a function that defines the residual of those quadrature equations. For the sensitivities of quadratures (2.10) with integrand  $q$ , the appropriate residual functions are given by  $\bar{q}_i = \partial q / \partial y s_i + \partial q / \partial \dot{y} \dot{s}_i + \partial q \partial p_i$ . This user function must be of type `IDAQuadSensRhsFn` defined as follows:

```
typedef int (*IDAQuadSensRhsFn)(int Ns, sunrealtype t, N_Vector yy, N_Vector yp, N_Vector *yyS, N_Vector *ypS,
N_Vector rrQ, N_Vector *rhsvalQS, void *user_data, N_Vector tmp1, N_Vector tmp2, N_Vector tmp3)
```

This function computes the sensitivity quadrature equation right-hand side for a given value of the independent variable  $t$  and state vector  $y$ .

**Arguments:**

- `Ns` – is the number of sensitivity vectors.
- `t` – is the current value of the independent variable.
- `yy` – is the current value of the dependent variable vector,  $y(t)$ .
- `yp` – is the current value of the dependent variable vector,  $\dot{y}(t)$ .
- `yyS` – is an array of `Ns` variables of type `N_Vector` containing the dependent sensitivity vectors  $s_i$ .
- `ypS` – is an array of `Ns` variables of type `N_Vector` containing the dependent sensitivity derivatives  $\dot{s}_i$ .
- `rrQ` – is the current value of the quadrature right-hand side  $q$ .
- `rhsvalQS` – contains the `Ns` output vectors.
- `user_data` – is the `user_data` pointer passed to `IDASetUserData()`.
- `tmp1`, `tmp2`, `tmp3` – are `N_Vector` s which can be used as temporary storage.

**Return value:**

An `IDAQuadSensRhsFn` should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and `IDA_QRHS_FAIL` is returned).

**Notes:**

Allocation of memory for `rhsvalQS` is automatically handled within IDAS.

Both `yy` and `yp` are of type `N_Vector` and both `yyS` and `ypS` are pointers to an array containing `Ns` vectors of type `N_Vector`. It is the user's responsibility to access the vector data consistently (including the use of the correct accessor macros from each `N_Vector` implementation).

There is one situation in which recovery is not possible even if `IDAQuadSensRhsFn` function returns a recoverable error flag. That is when this occurs at the very first call to the `IDAQuadSensRhsFn`, in which case IDAS returns `IDA_FIRST_QSRHS_ERR`.

### 5.4.5 Note on using partial error control

For some problems, when sensitivities are excluded from the error control test, the behavior of IDAS may appear at first glance to be erroneous. One would expect that, in such cases, the sensitivity variables would not influence in any way the step size selection.

The short explanation of this behavior is that the step size selection implemented by the error control mechanism in IDAS is based on the magnitude of the correction calculated by the nonlinear solver. As mentioned in §5.4.2.1, even with partial error control selected in the call to `IDASensInit()`, the sensitivity variables are included in the convergence tests of the nonlinear solver.

When using the simultaneous corrector method §2.6, the nonlinear system that is solved at each step involves both the state and sensitivity equations. In this case, it is easy to see how the sensitivity variables may affect the convergence rate of the nonlinear solver and therefore the step size selection. The case of the staggered corrector approach is more subtle. The sensitivity variables at a given step are computed only once the solver for the nonlinear state equations has converged. However, if the nonlinear system corresponding to the sensitivity equations has convergence problems, IDAS will attempt to improve the initial guess by reducing the step size in order to provide a better prediction of the sensitivity variables. Moreover, even if there are no convergence failures in the solution of the sensitivity system, IDAS may trigger a call to the linear solver's setup routine which typically involves reevaluation of Jacobian information (Jacobian approximation in the case of matrix-based linear solvers, or preconditioner data in the case of the Krylov solvers). The new Jacobian information will be used by subsequent calls to the nonlinear solver for the state equations and, in this way, potentially affect the step size selection.

When using the simultaneous corrector method it is not possible to decide whether nonlinear solver convergence failures or calls to the linear solver setup routine have been triggered by convergence problems due to the state or the sensitivity equations. When using one of the staggered corrector methods, however, these situations can be identified by carefully monitoring the diagnostic information provided through optional outputs. If there are no convergence failures in the sensitivity nonlinear solver, and none of the calls to the linear solver setup routine were made by the sensitivity nonlinear solver, then the step size selection is not affected by the sensitivity variables.

Finally, the user must be warned that the effect of appending sensitivity equations to a given system of DAEs on the step size selection (through the mechanisms described above) is problem-dependent and can therefore lead to either an increase or decrease of the total number of steps that IDAS takes to complete the simulation. At first glance, one would expect that the impact of the sensitivity variables, if any, would be in the direction of increasing the step size and therefore reducing the total number of steps. The argument for this is that the presence of the sensitivity variables in the convergence test of the nonlinear solver can only lead to additional iterations (and therefore a smaller iteration error), or to additional calls to the linear solver setup routine (and therefore more up-to-date Jacobian information), both of which will lead to larger steps being taken by IDAS. However, this is true only locally. Overall, a larger integration step taken at a given time may lead to step size reductions at later times, due to either nonlinear solver convergence failures or error test failures.

## 5.5 Using IDAS for Adjoint Sensitivity Analysis

This chapter describes the use of IDAS to compute sensitivities of derived functions using adjoint sensitivity analysis. As mentioned before, the adjoint sensitivity module of IDAS provides the infrastructure for integrating backward in time any system of DAEs that depends on the solution of the original IVP, by providing various interfaces to the main IDAS integrator, as well as several supporting user-callable functions. For this reason, in the following sections we refer to the *backward problem* and not to the *adjoint problem* when discussing details relevant to the DAEs that are integrated backward in time. The backward problem can be the adjoint problem (2.19) or (2.24), and can be augmented with some quadrature differential equations.

IDAS uses various constants for both input and output. These are defined as needed in this chapter, but for convenience are also listed separately in Appendix §13.

We begin with a brief overview, in the form of a skeleton user program. Following that are detailed descriptions of the interface to the various user-callable functions and of the user-supplied functions that were not already described in §5.1.

### 5.5.1 A skeleton of the user's main program

The following is a skeleton of the user's main program as an application of IDAS. The user program is to have these steps in the order indicated, unless otherwise noted. For the sake of brevity, we defer many of the details to the later sections. As in §5.1.2, most steps are independent of the `N_Vector`, `SUNMatrix`, `SUNLinearSolver`, and `SUNNon-linearSolver` implementations used. For the steps that are not, refer to Chapters §6, §7, §8, and §9 for the specific name of the function to be called or macro to be referenced.

Steps that are unchanged from the skeleton programs presented in §5.1.2, §5.4.1, and §5.4.4 are grayed out and new or modified steps are in bold.

1. Initialize parallel or multi-threaded environment
2. Create the SUNDIALS context object

#### Forward Problem

3. Set initial conditions for the forward problem
4. Create matrix object for the forward problem
5. Create linear solver object for the forward problem
6. Create nonlinear solver module for the forward problem
7. Create IDAS object for the forward problem
8. Initialize IDAS solver for the forward problem
9. Specify integration tolerances for forward problem
10. Attach linear solver module for the forward problem
11. Set linear solver optional inputs for the forward problem
12. Attach nonlinear solver module for the forward problem
13. Set nonlinear solver optional inputs for the forward problem
14. Initialize quadrature problem or problems for forward problems
15. Initialize forward sensitivity problem
16. Specify rootfinding
17. Set optional inputs for the forward problem

### 18. Allocate space for the adjoint computation

Call `IDAAdjInit()` to allocate memory for the combined forward-backward problem. This call requires `Nd`, the number of steps between two consecutive checkpoints. `IDAAdjInit()` also specifies the type of interpolation used (see §2.7.3).

### 19. Integrate forward problem

Call `IDASolveF()`, a wrapper for the IDAS main integration function `IDASolve()`, either in `IDA_NORMAL` mode to the time `tout` or in `IDA_ONE_STEP` mode inside a loop (if intermediate solutions of the forward problem are desired (see §5.5.2.3)). The final value of `tret` is then the maximum allowable value for the endpoint  $T$  of the backward problem.

## Backward Problem(s)

### 20. Create vectors of endpoint values for the backward problem

Create the vectors `yB0` and `ypB0` at the endpoint time `tB0 = T` at which the backward problem starts.

### 21. Create the backward problem

Call `IDACreateB()`, a wrapper for `IDACreate()`, to create the IDAS memory block for the new backward problem. Unlike `IDACreate()`, the function `IDACreateB()` does not return a pointer to the newly created memory block (see §5.5.2.4). Instead, this pointer is attached to the internal adjoint memory block (created by `IDAAdjInit()`) and returns an identifier called `which` that the user must later specify in any actions on the newly created backward problem.

### 22. Allocate memory for the backward problem

Call `IDAInitB()` (or `IDAInitBS()`, when the backward problem depends on the forward sensitivities). The two functions are actually wrappers for `IDAInit()` and allocate internal memory, specify problem data, and initialize IDAS at `tB0` for the backward problem (see §5.5.2.4).

### 23. Specify integration tolerances for backward problem

Call `IDASSStolerancesB()` or `IDASVStolerancesB()` to specify a scalar relative tolerance and scalar absolute tolerance, or a scalar relative tolerance and a vector of absolute tolerances, respectively. The functions are wrappers for `IDASSStolerances()` and `IDASVStolerances()` but they require an extra argument `which`, the identifier of the backward problem returned by `IDACreateB()`. See §5.5.2.5 for more information.

### 24. Set optional inputs for the backward problem

Call `IDASet*B` functions to change from their default values any optional inputs that control the behavior of IDAS. Unlike their counterparts for the forward problem, these functions take an extra argument `which`, the identifier of the backward problem returned by `IDACreateB()` (see §5.5.2.10).

### 25. Create matrix object for the backward problem

If a nonlinear solver requiring a linear solve will be used (e.g., the the default Newton iteration) and the linear solver will be a direct linear solver, then a template Jacobian matrix must be created by calling the appropriate constructor function defined by the particular `SUNMatrix` implementation.

---

**Note:** The dense, banded, and sparse matrix objects are usable only in a serial or threaded environment.

It is not required to use the same matrix type for both the forward and the backward problems.

---

### 26. Create linear solver object for the backward problem

If a nonlinear solver requiring a linear solver is chosen (e.g., the default Newton iteration), then the desired linear solver object for the backward problem must be created by calling the appropriate constructor function defined by the particular `SUNLinearSolver` implementation.

---

**Note:** It is not required to use the same linear solver module for both the forward and the backward problems; for example, the forward problem could be solved with the SUNLINSOL\_BAND linear solver module and the backward problem with SUNLINSOL\_SPGMR linear solver module.

---

#### 27. Set linear solver interface optional inputs for the backward problem

Call `IDASet*B` functions to change optional inputs specific to the linear solver interface. See §5.5.2.10 for details.

#### 28. Attach linear solver module for the backward problem

If a nonlinear solver requiring a linear solver is chosen for the backward problem (e.g., the default Newton iteration), then initialize the IDALS linear solver interface by attaching the linear solver object (and matrix object, if applicable) with `IDASetLinearSolverB()` (for additional details see §5.5.2.6).

#### 29. Create nonlinear solver object for the backward problem (optional)

If using a non-default nonlinear solver for the backward problem, then create the desired nonlinear solver object by calling the appropriate constructor function defined by the particular `SUNNonlinearSolver` implementation e.g., `NLSB = SUNNonlinSol_***(...)`; where `***` is the name of the nonlinear solver (see Chapter §9 for details).

#### 30. Attach nonlinear solver module for the backward problem (optional)

If using a non-default nonlinear solver for the backward problem, then initialize the nonlinear solver interface by attaching the nonlinear solver object by calling `IDASetNonlinearSolverB()`.

#### 31. Initialize quadrature calculation

If additional quadrature equations must be evaluated, call `IDAQuadInitB()` or `IDAQuadInitBS()` (if quadrature depends also on the forward sensitivities) as shown in §5.5.2.12. These functions are wrappers around `IDAQuadInit()` and can be used to initialize and allocate memory for quadrature integration. Optionally, call `IDASetQuad*B` functions to change from their default values optional inputs that control the integration of quadratures during the backward phase.

#### 32. Integrate backward problem

Call `IDASolveB()`, a second wrapper around the IDAS main integration function `IDASolve()`, to integrate the backward problem from `tB0`. This function can be called either in `IDA_NORMAL` or `IDA_ONE_STEP` mode. Typically, `IDASolveB()` will be called in `IDA_NORMAL` mode with an end time equal to the initial time  $t_0$  of the forward problem.

#### 33. Extract quadrature variables

If applicable, call `IDAGetQuadB()`, a wrapper around `IDAGetQuad()`, to extract the values of the quadrature variables at the time returned by the last call to `IDASolveB()`.

#### 34. Destroy objects

Upon completion of the backward integration, call all necessary deallocation functions. These include appropriate destructors for the vectors `y` and `yB`, a call to `IDAFree()` to free the IDAS memory block for the forward problem. If one or more additional adjoint sensitivity analyses are to be done for this problem, a call to `IDAAdjFree()` (see §5.5.2.1) may be made to free and deallocate the memory allocated for the backward problems, followed by a call to `IDAAdjInit()`.

#### 35. Finalize MPI, if used

The above user interface to the adjoint sensitivity module in IDAS was motivated by the desire to keep it as close as possible in look and feel to the one for DAE IVP integration. Note that if steps (18) - (31) are not present, a program with the above structure will have the same functionality as one described in §5.1.2 for integration of DAEs, albeit with some overhead due to the checkpointing scheme.

If there are multiple backward problems associated with the same forward problem, repeat steps (18) - (31) above for each successive backward problem. In the process, If there are multiple backward problems associated with the same forward each call to `IDACreateB()` creates a new value of the identifier which.

## 5.5.2 User-callable functions for adjoint sensitivity analysis

### 5.5.2.1 Adjoint sensitivity allocation and deallocation functions

After the setup phase for the forward problem, but before the call to `IDASolveF()`, memory for the combined forward-backward problem must be allocated by a call to the function `IDAAdjInit()`. The form of the call to this function is

int **IDAAdjInit**(void \*ida\_mem, long int Nd, int interpType)

The function `IDAAdjInit()` updates IDAS memory block by allocating the internal memory needed for backward integration. Space is allocated for the  $N_d = N_d$  interpolation data points, and a linked list of checkpoints is initialized.

#### Arguments:

- `ida_mem` – is the pointer to the IDAS memory block returned by a previous call to `IDACreate()`.
- `Nd` – is the number of integration steps between two consecutive checkpoints.
- `interpType` – specifies the type of interpolation used and can be `IDA_POLYNOMIAL` or `IDA_HERMITE`, indicating variable-degree polynomial and cubic Hermite interpolation, respectively see §2.7.3.

#### Return value:

- `IDA_SUCCESS` – `IDAAdjInit()` was successful.
- `IDA_MEM_FAIL` – A memory allocation request has failed.
- `IDA_MEM_NULL` – `ida_mem` was NULL.
- `IDA_ILL_INPUT` – One of the parameters was invalid: `Nd` was not positive or `interpType` is not one of the `IDA_POLYNOMIAL` or `IDA_HERMITE`.

#### Notes:

The user must set `Nd` so that all data needed for interpolation of the forward problem solution between two checkpoints fits in memory. `IDAAdjInit()` attempts to allocate space for  $(2N_d + 3)$  variables of type `N_Vector`.

If an error occurred, `IDAAdjInit()` also sends a message to the error handler function.

int **IDAAdjReInit**(void \*ida\_mem)

The function `IDAAdjReInit()` reinitializes the IDAS memory block for ASA, assuming that the number of steps between check points and the type of interpolation remain unchanged.

#### Arguments:

- `ida_mem` – is the pointer to the IDAS memory block returned by a previous call to `IDACreate()`.

#### Return value:

- `IDA_SUCCESS` – `IDAAdjReInit()` was successful.
- `IDA_MEM_NULL` – `ida_mem` was NULL.
- `IDA_NO_ADJ` – The function `IDAAdjInit()` was not previously called.

**Notes:**

The list of check points (and associated memory) is deleted.

The list of backward problems is kept. However, new backward problems can be added to this list by calling *IDACreateB()*. If a new list of backward problems is also needed, then free the adjoint memory (by calling *IDAAdjFree()*) and reinitialize ASA with *IDAAdjInit()*.

The IDAS memory for the forward and backward problems can be reinitialized separately by calling *IDAREInit()* and *IDAREInitB()*, respectively.

void **IDAAdjFree**(void \*ida\_mem)

The function *IDAAdjFree()* frees the memory related to backward integration allocated by a previous call to *IDAAdjInit()*.

**Arguments:**

The only argument is the IDAS memory block pointer returned by a previous call to *IDACreate()*.

**Return value:**

The function *IDAAdjFree()* has no return value.

**Notes:**

This function frees all memory allocated by *IDAAdjInit()*. This includes workspace memory, the linked list of checkpoints, memory for the interpolation data, as well as the IDAS memory for the backward integration phase.

Unless one or more further calls to *IDAAdjInit()* are to be made, *IDAAdjFree()* should not be called by the user, as it is invoked automatically by *IDAFree()*.

**5.5.2.2 Adjoint sensitivity optional input**

At any time during the integration of the forward problem, the user can disable the checkpointing of the forward sensitivities by calling the following function:

int **IDAAdjSetNoSensi**(void \*ida\_mem)

The function *IDAAdjSetNoSensi()* instructs *IDASolveF()* not to save checkpointing data for forward sensitivities any more.

**Arguments:**

- *ida\_mem* – pointer to the IDAS memory block.

**Return value:**

- *IDA\_SUCCESS* – The call to *IDACreateB()* was successful.
- *IDA\_MEM\_NULL* – The *ida\_mem* was NULL.
- *IDA\_NO\_ADJ* – The function *IDAAdjInit()* has not been previously called.

**5.5.2.3 Forward integration function**

The function *IDASolveF()* is very similar to the IDAS function *IDASolve()* in that it integrates the solution of the forward problem and returns the solution  $(y, \dot{y})$ . At the same time, however, *IDASolveF()* stores checkpoint data every *Nd* integration steps. *IDASolveF()* can be called repeatedly by the user. Note that *IDASolveF()* is used only for the forward integration pass within an Adjoint Sensitivity Analysis. It is not for use in Forward Sensitivity Analysis; for that, see §5.4. The call to this function has the form



```
int IDASolveF(void *ida_mem, sunrealtype tout, sunrealtype *tret, N_Vector yret, N_Vector ypret, int itask, int *ncheck)
```

The function `IDASolveF()` integrates the forward problem over an interval in  $t$  and saves checkpointing data.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `tout` – the next time at which a computed solution is desired.
- `tret` – the time reached by the solver output.
- `yret` – the computed solution vector  $y$ .
- `ypret` – the computed solution vector  $\dot{y}$ .
- `itask` – a flag indicating the job of the solver for the next step. The `IDA_NORMAL` task is to have the solver take internal steps until it has reached or just passed the user-specified `tout` parameter. The solver then interpolates in order to return an approximate value of  $y(tout)$  and  $\dot{y}(tout)$ . The `IDA_ONE_STEP` option tells the solver to take just one internal step and return the solution at the point reached by that step.
- `ncheck` – the number of internal checkpoints stored so far.

**Return value:**

On return, `IDASolveF()` returns vectors `yret`, `ypret` and a corresponding independent variable value `t = tret`, such that `yret` is the computed value of  $y(t)$  and `ypret` the value of  $\dot{y}(t)$ . Additionally, it returns in `ncheck` the number of internal checkpoints saved; the total number of checkpoint intervals is `ncheck+1`. The return value flag (of type `int`) will be one of the following. For more details see the documentation for `IDASolve()`.

- `IDA_SUCCESS` – `IDASolveF()` succeeded.
- `IDA_TSTOP_RETURN` – `IDASolveF()` succeeded by reaching the optional stopping point.
- `IDA_ROOT_RETURN` – `IDASolveF()` succeeded and found one or more roots. In this case, `tret` is the location of the root. If `nrtfn > 1`, call `IDAGetRootInfo()` to see which  $g_i$  were found to have a root.
- `IDA_NO_MALLOC` – The function `IDAInit()` has not been previously called.
- `IDA_ILL_INPUT` – One of the inputs to `IDASolveF()` is illegal.
- `IDA_TOO_MUCH_WORK` – The solver took `mxstep` internal steps but could not reach `tout`.
- `IDA_TOO_MUCH_ACC` – The solver could not satisfy the accuracy demanded by the user for some internal step.
- `IDA_ERR_FAILURE` – Error test failures occurred too many times during one internal time step or occurred with  $|h| = h_{min}$ .
- `IDA_CONV_FAILURE` – Convergence test failures occurred too many times during one internal time step or occurred with  $|h| = h_{min}$ .
- `IDA_LSETUP_FAIL` – The linear solver's setup function failed in an unrecoverable manner.
- `IDA_LSOLVE_FAIL` – The linear solver's solve function failed in an unrecoverable manner.
- `IDA_NO_ADJ` – The function `IDAAdjInit()` has not been previously called.
- `IDA_MEM_FAIL` – A memory allocation request has failed in an attempt to allocate space for a new checkpoint.

**Notes:**

All failure return values are negative and therefore a test `flag < 0` will trap all `IDASolveF()` failures.



At this time, `IDASolveF()` stores checkpoint information in memory only. Future versions will provide for a safeguard option of dumping checkpoint data into a temporary file as needed. The data stored at each checkpoint is basically a snapshot of the IDAS internal memory block and contains enough information to restart the integration from that time and to proceed with the same step size and method order sequence as during the forward integration.

In addition, `IDASolveF()` also stores interpolation data between consecutive checkpoints so that, at the end of this first forward integration phase, interpolation information is already available from the last checkpoint forward. In particular, if no checkpoints were necessary, there is no need for the second forward integration phase.

**Warning:** It is illegal to change the integration tolerances between consecutive calls to `IDASolveF()`, as this information is not captured in the checkpoint data.

#### 5.5.2.4 Backward problem initialization functions

The functions `IDACreateB()` and `IDAInitB()` (or `IDAInitBS()`) must be called in the order listed. They instantiate an IDAS solver object, provide problem and solution specifications, and allocate internal memory for the backward problem.

int `IDACreateB`(void \*ida\_mem, int \*which)

The function `IDACreateB()` instantiates an IDAS solver object for the backward problem.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block returned by `IDACreate()`.
- `which` – contains the identifier assigned by IDAS for the newly created backward problem. Any call to IDA\*B functions requires such an identifier.

**Return value:**

- `IDA_SUCCESS` – The call to `IDACreateB()` was successful.
- `IDA_MEM_NULL` – The `ida_mem` was NULL.
- `IDA_NO_ADJ` – The function `IDAAdjInit()` has not been previously called.
- `IDA_MEM_FAIL` – A memory allocation request has failed.

There are two initialization functions for the backward problem – one for the case when the backward problem does not depend on the forward sensitivities, and one for the case when it does. These two functions are described next.

The function `IDAInitB()` initializes the backward problem when it does not depend on the forward sensitivities. It is essentially wrapper for `IDAInit` with some particularization for backward integration, as described below.

int `IDAInitB`(void \*ida\_mem, int which, `IDAResFnB` resB, `sunrealtype` tB0, `N_Vector` yB0, `N_Vector` ypB0)

The function `IDAInitB()` provides problem specification, allocates internal memory, and initializes the backward problem.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block returned by `IDACreate()`.
- `which` – represents the identifier of the backward problem.
- `resB` – is the C function which computes  $f_B$ , the residual of the backward DAE problem. This function has the form `resB(t, y, yp, yB, ypB, resvalB, user_dataB)` for full details see §5.5.3.1.

- `tB0` – specifies the endpoint  $T$  where final conditions are provided for the backward problem, normally equal to the endpoint of the forward integration.
- `yB0` – is the initial value at  $t = tB0$  of the backward solution.
- `ypB0` – is the initial derivative value at  $t = tB0$  of the backward solution.

**Return value:**

- `IDA_SUCCESS` – The call to `IDAInitB()` was successful.
- `IDA_NO_MALLOC` – The function `IDAInit()` has not been previously called.
- `IDA_MEM_NULL` – The `ida_mem` was NULL.
- `IDA_NO_ADJ` – The function `IDAAdjInit()` has not been previously called.
- `IDA_BAD_TB0` – The final time `tB0` was outside the interval over which the forward problem was solved.
- `IDA_ILL_INPUT` – The parameter which represented an invalid identifier, or one of `yB0`, `ypB0`, `resB` was NULL.

**Notes:**

The memory allocated by `IDAInitB()` is deallocated by the function `IDAAdjFree()`.

For the case when backward problem also depends on the forward sensitivities, user must call `IDAInitBS()` instead of `IDAInitB()`. Only the third argument of each function differs between these functions.

int **IDAInitBS**(void \*ida\_mem, int which, *IDAResFnBS* resBS, *sunrealtype* tB0, *N\_Vector* yB0, *N\_Vector* ypB0)

The function `IDAInitBS()` provides problem specification, allocates internal memory, and initializes the backward problem.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block returned by `IDACreate()`.
- `which` – represents the identifier of the backward problem.
- `resBS` – is the C function which computes  $fB$ , the residual or the backward DAE problem. This function has the form `resBS(t, y, yp, yS, ypS, yB, ypB, resvalB, user_dataB)` for full details see §5.5.3.2.
- `tB0` – specifies the endpoint  $T$  where final conditions are provided for the backward problem.
- `yB0` – is the initial value at  $t = tB0$  of the backward solution.
- `ypB0` – is the initial derivative value at  $t = tB0$  of the backward solution.

**Return value:**

- `IDA_SUCCESS` – The call to `IDAInitB()` was successful.
- `IDA_NO_MALLOC` – The function `IDAInit()` has not been previously called.
- `IDA_MEM_NULL` – The `ida_mem` was NULL.
- `IDA_NO_ADJ` – The function `IDAAdjInit()` has not been previously called.
- `IDA_BAD_TB0` – The final time `tB0` was outside the interval over which the forward problem was solved.
- `IDA_ILL_INPUT` – The parameter which represented an invalid identifier, or one of `yB0`, `ypB0`, `resB` was NULL, or sensitivities were not active during the forward integration.

**Notes:**

The memory allocated by `IDAInitBS()` is deallocated by the function `IDAAdjFree()`.

The function `IDAreInitB()` reinitializes idas for the solution of a series of backward problems, each identified by a value of the parameter `which`. `IDAreInitB()` is essentially a wrapper for `IDAReInit()`, and so all details given for `IDAReInit()` apply here. Also, `IDAreInitB()` can be called to reinitialize a backward problem even if it has been initialized with the sensitivity-dependent version `IDAInitBS()`. Before calling `IDAreInitB()` for a new backward problem, call any desired solution extraction functions `IDAGet**` associated with the previous backward problem. The call to the `IDAreInitB()` function has the form

```
int IDAreInitB(void *ida_mem, int which, sunrealtype tB0, N_Vector yB0, N_Vector ypB0)
```

The function `IDAreInitB()` reinitializes an IDAS backward problem.

**Arguments:**

- `ida_mem` – pointer to IDAS memory block returned by `IDACreate()`.
- `which` – represents the identifier of the backward problem.
- `tB0` – specifies the endpoint  $T$  where final conditions are provided for the backward problem.
- `yB0` – is the initial value at  $t = tB0$  of the backward solution.
- `ypB0` – is the initial derivative value at  $t = tB0$  of the backward solution.

**Return value:**

- `IDA_SUCCESS` – The call to `IDAreInitB()` was successful.
- `IDA_NO_MALLOC` – The function `IDAInit()` has not been previously called.
- `IDA_MEM_NULL` – The `ida_mem` memory block pointer was NULL.
- `IDA_NO_ADJ` – The function `IDAAdjInit()` has not been previously called.
- `IDA_BAD_TB0` – The final time `tB0` is outside the interval over which the forward problem was solved.
- `IDA_ILL_INPUT` – The parameter `which` represented an invalid identifier, or one of `yB0`, `ypB0` was NULL.

### 5.5.2.5 Tolerance specification functions for backward problem

One of the following two functions must be called to specify the integration tolerances for the backward problem. Note that this call must be made after the call to `IDAInitB()` or `IDAInitBS()`.

```
int IDASStolerancesB(void *ida_mem, int which, sunrealtype reltolB, sunrealtype abstolB)
```

The function `IDASStolerancesB()` specifies scalar relative and absolute tolerances.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block returned by `IDACreate()`.
- `which` – represents the identifier of the backward problem.
- `reltolB` – is the scalar relative error tolerance.
- `abstolB` – is the scalar absolute error tolerance.

**Return value:**

- `IDA_SUCCESS` – The call to `IDASStolerancesB()` was successful.
- `IDA_MEM_NULL` – The IDAS memory block was not initialized through a previous call to `IDACreate()`.
- `IDA_NO_MALLOC` – The allocation function `IDAInit()` has not been called.
- `IDA_NO_ADJ` – The function `IDAAdjInit()` has not been previously called.

- IDA\_ILL\_INPUT – One of the input tolerances was negative.

int **IDASVtolerancesB**(void \*ida\_mem, int which, *sunrealtype* reltolB, *N\_Vector* abstolB)

The function *IDASVtolerancesB()* specifies scalar relative tolerance and vector absolute tolerances.

**Arguments:**

- ida\_mem – pointer to the IDAS memory block returned by *IDACreate()*.
- which – represents the identifier of the backward problem.
- reltolB – is the scalar relative error tolerance.
- abstolB – is the vector of absolute error tolerances.

**Return value:**

- IDA\_SUCCESS – The call to *IDASVtolerancesB()* was successful.
- IDA\_MEM\_NULL – The IDAS memory block was not initialized through a previous call to *IDACreate()*.
- IDA\_NO\_MALLOC – The allocation function *IDAINit()* has not been called.
- IDA\_NO\_ADJ – The function *IDAAdjInit()* has not been previously called.
- IDA\_ILL\_INPUT – The relative error tolerance was negative or the absolute tolerance had a negative component.

**Notes:**

This choice of tolerances is important when the absolute error tolerance needs to be different for each component of the DAE state vector  $y$ .

### 5.5.2.6 Linear solver initialization functions for backward problem

All IDAS linear solver modules available for forward problems are available for the backward problem. They should be created as for the forward problem then attached to the memory structure for the backward problem using the following function.

int **IDASetLinearSolverB**(void \*ida\_mem, int which, *SUNLinearSolver* LS, *SUNMatrix* A)

The function *IDASetLinearSolverB()* attaches a generic *SUNLinearSolver* object LS and corresponding template Jacobian *SUNMatrix* object A (if applicable) to IDAS, initializing the IDALS linear solver interface for solution of the backward problem.

**Arguments:**

- ida\_mem – pointer to the IDAS memory block.
- which – represents the identifier of the backward problem returned by *IDACreateB()*.
- LS – *SUNLinearSolver* object to use for solving linear systems for the backward problem.
- A – *SUNMatrix* object for used as a template for the Jacobian for the backward problem or NULL if not applicable.

**Return value:**

- IDALS\_SUCCESS – The IDALS initialization was successful.
- IDALS\_MEM\_NULL – The ida\_mem pointer is NULL.
- IDALS\_ILL\_INPUT – The parameter which represented an invalid identifier.
- IDALS\_MEM\_FAIL – A memory allocation request failed.

- IDALS\_NO\_ADJ – The function *IDAAadjInit()* has not been previously called.

**Notes:**

If LS is a matrix-based linear solver, then the template Jacobian matrix A will be used in the solve process, so if additional storage is required within the SUNMatrix object (e.g. for factorization of a banded matrix), ensure that the input object is allocated with sufficient size (see the documentation of the particular SUNMatrix type in Chapter §7 for further information).

New in version 3.0.0: Replaces the deprecated functions IDADlsSetLinearSolverB and IDASpilsSetLinearSolverB.

**5.5.2.7 Nonlinear solver initialization functions for backward problem**

As with the forward problem IDAS uses the SUNNonlinearSolver implementation of Newton's method defined by the SUNNONLINSOL\_NEWTON module (see §9.3) by default.

To specify a different nonlinear solver in IDAS for the backward problem, the user's program must create a SUNNonlinearSolver object by calling the appropriate constructor routine. The user must then attach the SUNNonlinearSolver object to IDAS by calling *IDASetNonlinearSolverB()*, as documented below.

When changing the nonlinear solver in IDAS, *IDASetNonlinearSolverB()* must be called after *IDAInitB()*. If any calls to *IDASolveB()* have been made, then IDAS will need to be reinitialized by calling *IDAREInitB()* to ensure that the nonlinear solver is initialized correctly before any subsequent calls to *IDASolveB()*.

int **IDASetNonlinearSolverB**(void \*ida\_mem, int which, *SUNNonlinearSolver* NLS)

The function *IDASetNonlinearSolverB()* attaches a SUNNonlinearSolver object (NLS) to IDAS for the solution of the backward problem.

**Arguments:**

- ida\_mem – pointer to the IDAS memory block.
- which – represents the identifier of the backward problem returned by *IDACreateB()*.
- NLS – SUNNonlinearSolver object to use for solving nonlinear systems for the backward problem.

**Return value:**

- IDA\_SUCCESS – The nonlinear solver was successfully attached.
- IDA\_MEM\_NULL – The ida\_mem pointer is NULL.
- IDALS\_NO\_ADJ – The function IDAAadjInit has not been previously called.
- IDA\_ILL\_INPUT – The parameter which represented an invalid identifier or the SUNNonlinearSolver object is NULL, does not implement the required nonlinear solver operations, is not of the correct type, or the residual function, convergence test function, or maximum number of nonlinear iterations could not be set.

**5.5.2.8 Initial condition calculation functions for backward problem**

IDAS provides support for calculation of consistent initial conditions for certain backward index-one problems of semi-implicit form through the functions *IDACalcICB()* and *IDACalcICBS()*. Calling them is optional. It is only necessary when the initial conditions do not satisfy the adjoint system.

The above functions provide the same functionality for backward problems as *IDACalcIC()* with parameter icopt = IDA\_YA\_YDP\_INIT provides for forward problems: compute the algebraic components of  $yB$  and differential components of  $\dot{y}B$ , given the differential components of  $yB$ . They require that the IDASetIdB was previously called to specify the differential and algebraic components.

Both functions require forward solutions at the final time  $t_{B0}$ . `IDACalcICBS()` also needs forward sensitivities at the final time  $t_{B0}$ .

int **IDACalcICB**(void \*ida\_mem, int which, *sunrealtype* tBout1, *N\_Vector* yfin, *N\_Vector* ypfin)

The function `IDACalcICB()` corrects the initial values  $y_{B0}$  and  $yp_{B0}$  at time  $t_{B0}$  for the backward problem.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `which` – is the identifier of the backward problem.
- `tBout1` – is the first value of  $t$  at which a solution will be requested from `IDASolveB()`. This value is needed here only to determine the direction of integration and rough scale in the independent variable  $t$ .
- `yfin` – the forward solution at the final time  $t_{B0}$ .
- `ypfin` – the forward solution derivative at the final time  $t_{B0}$ .

**Return value:**

- `IDA_NO_ADJ` – `IDAAdjInit()` has not been previously called.
- `IDA_ILL_INPUT` – Parameter `which` represented an invalid identifier.

**Notes:**

All failure return values are negative and therefore a test `flag < 0` will trap all `IDACalcICB()` failures. Note that `IDACalcICB()` will correct the values of  $y_B(t_{B0})$  and  $\dot{y}_B(t_{B0})$  which were specified in the previous call to `IDAInitB()` or `IDAReInitB()`. To obtain the corrected values, call `IDAGetConsistentICB()` (see §5.5.2.11).

`IDACalcICB()` will correct the values of  $y_B(t_{B0})$  and  $\dot{y}_B(t_{B0})$  which were specified in the previous call to `IDAInitB()` or `IDAReInitB()`. To obtain the corrected values, call `c:func:IDAGetConsistentICB` (see §5.5.2.11).

In the case where the backward problem also depends on the forward sensitivities, user must call the following function to correct the initial conditions:

int **IDACalcICBS**(void \*ida\_mem, int which, *sunrealtype* tBout1, *N\_Vector* yfin, *N\_Vector* ypfin, *N\_Vector* ySfin, *N\_Vector* ypSfin)

The function `IDACalcICBS()` corrects the initial values  $y_{B0}$  and  $yp_{B0}$  at time  $t_{B0}$  for the backward problem.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `which` – is the identifier of the backward problem.
- `tBout1` – is the first value of  $t$  at which a solution will be requested from `IDASolveB()`. This value is needed here only to determine the direction of integration and rough scale in the independent variable  $t$ .
- `yfin` – the forward solution at the final time  $t_{B0}$ .
- `ypfin` – the forward solution derivative at the final time  $t_{B0}$ .
- `ySfin` – a pointer to an array of  $N_s$  vectors containing the sensitivities of the forward solution at the final time  $t_{B0}$ .
- `ypSfin` – a pointer to an array of  $N_s$  vectors containing the derivatives of the forward solution sensitivities at the final time  $t_{B0}$ .

**Return value:**

- `IDA_NO_ADJ` – `IDAAdjInit()` has not been previously called.

- IDA\_ILL\_INPUT – Parameter which represented an invalid identifier, sensitivities were not active during forward integration, or `IDAINITBS()` or `IDAREINITB()` has not been previously called.

#### Notes:

All failure return values are negative and therefore a test `flag < 0` will trap all `IDACalcICBS()` failures. Note that `IDACalcICBS()` will correct the values of  $yB(tB_0)$  and  $\dot{y}B(tB_0)$  which were specified in the previous call to `IDAINITBS()` or `IDAREINITB()`. To obtain the corrected values, call `IDAGetConsistentICB()` (see §5.5.2.11).

`IDACalcICBS()` will correct the values of  $yB(tB_0)$  and  $\dot{y}B(tB_0)$  which were specified in the previous call to `IDAINITBS()` or `IDAREINITB()`. To obtain the corrected values, :call `IDAGetConsistentICB()`.

### 5.5.2.9 Backward integration function

The function `IDASolveB()` performs the integration of the backward problem. It is essentially a wrapper for the IDAS main integration function `IDASolve()` and, in the case in which checkpoints were needed, it evolves the solution of the backward problem through a sequence of forward-backward integration pairs between consecutive checkpoints. In each pair, the first run integrates the original IVP forward in time and stores interpolation data; the second run integrates the backward problem backward in time and performs the required interpolation to provide the solution of the IVP to the backward problem.

The function `IDASolveB()` does not return the solution  $yB$  itself. To obtain that, call the function `IDAGetB()`, which is also described below.

The `IDASolveB()` function does not support rootfinding, unlike `IDASolveF()`, which supports the finding of roots of functions of  $(t, y, \dot{y})$ . If rootfinding was performed by `IDASolveF()`, then for the sake of efficiency, it should be disabled for `IDASolveB()` by first calling `IDARootInit()` with `nrtfn = 0`.

The call to `IDASolveB()` has the form

```
int IDASolveB(void *ida_mem, sunrealtype tBout, int itaskB)
```

The function `IDASolveB()` integrates the backward DAE problem.

#### Arguments:

- `ida_mem` – pointer to the IDAS memory returned by `IDACreate()`.
- `tBout` – the next time at which a computed solution is desired.
- `itaskB` – output mode a flag indicating the job of the solver for the next step. The IDA\_NORMAL task is to have the solver take internal steps until it has reached or just passed the user-specified value `tBout`. The solver then interpolates in order to return an approximate value of  $yB(tBout)$ . The IDA\_ONE\_STEP option tells the solver to take just one internal step in the direction of `tBout` and return.

#### Return value:

- IDA\_SUCCESS – `IDASolveB()` succeeded.
- IDA\_MEM\_NULL – The `ida_mem` was NULL.
- IDA\_NO\_ADJ – The function `IDAAdjInit()` has not been previously called.
- IDA\_NO\_BCK – No backward problem has been added to the list of backward problems by a call to `IDACreateB()`.
- IDA\_NO\_FWD – The function `IDASolveF()` has not been previously called.
- IDA\_ILL\_INPUT – One of the inputs to `IDASolveB()` is illegal.
- IDA\_BAD\_ITASK – The `itaskB` argument has an illegal value.
- IDA\_TOO\_MUCH\_WORK – The solver took `mxstep` internal steps but could not reach `tBout`.



- IDA\_TOO\_MUCH\_ACC – The solver could not satisfy the accuracy demanded by the user for some internal step.
- IDA\_ERR\_FAILURE – Error test failures occurred too many times during one internal time step.
- IDA\_CONV\_FAILURE – Convergence test failures occurred too many times during one internal time step.
- IDA\_LSETUP\_FAIL – The linear solver’s setup function failed in an unrecoverable manner.
- IDA\_SOLVE\_FAIL – The linear solver’s solve function failed in an unrecoverable manner.
- IDA\_BCKMEM\_NULL – The IDAS memory for the backward problem was not created with a call to [IDACreateB\(\)](#).
- IDA\_BAD\_TBOUT – The desired output time `tBout` is outside the interval over which the forward problem was solved.
- IDA\_REIFWD\_FAIL – Reinitialization of the forward problem failed at the first checkpoint corresponding to the initial time of the forward problem.
- IDA\_FWD\_FAIL – An error occurred during the integration of the forward problem.

**Notes:**

All failure return values are negative and therefore a test `flag < 0` will trap all [IDASolveB\(\)](#) failures. In the case of multiple checkpoints and multiple backward problems, a given call to [IDASolveB\(\)](#) in `IDA_ONE_STEP` mode may not advance every problem one step, depending on the relative locations of the current times reached. But repeated calls will eventually advance all problems to `tBout`.

To obtain the solution `yB` to the backward problem, call the function [IDAGetB\(\)](#) as follows:

```
int IDAGetB(void *ida_mem, int which, sunrealtype *tret, N_Vector yB, N_Vector ypB)
```

The function [IDAGetB\(\)](#) provides the solution `yB` of the backward DAE problem.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory returned by [IDACreate\(\)](#).
- `which` – the identifier of the backward problem.
- `tret` – the time reached by the solver output.
- `yB` – the backward solution at time `tret`.
- `ypB` – the backward solution derivative at time `tret`.

**Return value:**

- IDA\_SUCCESS – [IDAGetB\(\)](#) was successful.
- IDA\_MEM\_NULL – `ida_mem` is NULL.
- IDA\_NO\_ADJ – The function [IDAAdjInit\(\)](#) has not been previously called.
- IDA\_ILL\_INPUT – The parameter `which` is an invalid identifier.

**Notes:**

To obtain the solution associated with a given backward problem at some other time within the last integration step, first obtain a pointer to the proper IDAS memory structure by calling [IDAGetAdjIDABmem\(\)](#) and then use it to call [IDAGetDky\(\)](#).

**Warning:** The user must allocate space for `yB` and `ypB`.



### 5.5.2.10 Optional input functions for the backward problem

As for the forward problem there are numerous optional input parameters that control the behavior of the IDAS solver for the backward problem. IDAS provides functions that can be used to change these optional input parameters from their default values which are then described in detail in the remainder of this section, beginning with those for the main IDAS solver and continuing with those for the linear solver interfaces. For the most casual use of IDAS, the reader can skip to §5.5.3.

We note that, on an error return, all of the optional input functions send an error message to the error handler function. All error return values are negative, so the test `flag < 0` will catch all errors. Finally, a call to a `IDASet***B` function can be made from the user's calling program at any time and, if successful, takes effect immediately.

#### Main solver optional input functions

The adjoint module in IDAS provides wrappers for most of the optional input functions defined in §5.1.3.10. The only difference is that the user must specify the identifier which of the backward problem within the list managed by IDAS.

The optional input functions defined for the backward problem are:

```
flag = IDASetUserDataB(ida_mem, which, user_dataB);
flag = IDASetMaxOrdB(ida_mem, which, maxordB);
flag = IDASetMaxNumStepsB(ida_mem, which, mxstepsB);
flag = IDASetInitStepB(ida_mem, which, hinB);
flag = IDASetMaxStepB(ida_mem, which, hmaxB);
flag = IDASetSuppressAlgB(ida_mem, which, suppressalgB);
flag = IDASetIdB(ida_mem, which, idB);
flag = IDASetConstraintsB(ida_mem, which, constraintsB);
```

Their return value `flag` (of type `int`) can have any of the return values of their counterparts, but it can also be `IDA_NO_ADJ` if `IDAAdjInit()` has not been called, or `IDA_ILL_INPUT` if `which` was an invalid identifier.

#### Linear solver interface optional input functions

When using matrix-based linear solver modules for the backward problem, i.e., a non-NULL `SUNMatrix` object `A` was passed to `IDASetLinearSolverB()`, the IDALS linear solver interface needs a function to compute an approximation to the Jacobian matrix. This can be attached through a call to either `IDASetJacFnB()` or `IDASetJacFnBS()`, with the second used when the backward problem depends on the forward sensitivities.

`int IDASetJacFnB(void *ida_mem, int which, IDALSJacFnB jacB)`

The function `IDASetJacFnB()` specifies the Jacobian approximation function to be used for the backward problem.

#### Arguments:

- `ida_mem` – pointer to the IDAS memory block.
- `which` – represents the identifier of the backward problem.
- `jacB` – user-defined Jacobian approximation function.

#### Return value:

- `IDALS_SUCCESS` – `IDASetJacFnB()` succeeded.
- `IDALS_MEM_NULL` – The `ida_mem` was `NULL`.
- `IDALS_NO_ADJ` – The function `IDAAdjInit()` has not been previously called.

- IDALS\_LMEM\_NULL – The linear solver has not been initialized with a call to [IDASetLinearSolverB\(\)](#).
- IDALS\_ILL\_INPUT – The parameter `which` represented an invalid identifier.

New in version 3.0.0: Replaces the deprecated function `IDADlsSetJacFnB`.

int **IDASetJacFnBS**(void \*ida\_mem, int which, [IDALSJacFnBS](#) jacBS)

The function [IDASetJacFnBS\(\)](#) specifies the Jacobian approximation function to be used for the backward problem in the case where the backward problem depends on the forward sensitivities.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `which` – represents the identifier of the backward problem.
- `jacBS` – user-defined Jacobian approximation function.

**Return value:**

- IDALS\_SUCCESS – [IDASetJacFnBS\(\)](#) succeeded.
- IDALS\_MEM\_NULL – The `ida_mem` was NULL.
- IDALS\_NO\_ADJ – The function [IDAAdjInit\(\)](#) has not been previously called.
- IDALS\_LMEM\_NULL – The linear solver has not been initialized with a call to [IDASetLinearSolverB\(\)](#).
- IDALS\_ILL\_INPUT – The parameter `which` represented an invalid identifier.

New in version 3.0.0: Replaces the deprecated function `IDADlsSetJacFnBS`.

The function [IDASetLinearSolutionScalingB\(\)](#) can be used to enable or disable solution scaling when using a matrix-based linear solver.

int **IDASetLinearSolutionScalingB**(void \*ida\_mem, int which, [sunbooleantype](#) onoffB)

The function [IDASetLinearSolutionScalingB\(\)](#) enables or disables scaling the linear system solution to account for a change in  $\alpha$  in the linear system in the backward problem. For more details see §8.2.1.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `which` – represents the identifier of the backward problem.
- `onoffB` – flag to enable SUNTRUE or disable SUNFALSE scaling.

**Return value:**

- IDALS\_SUCCESS – The flag value has been successfully set.
- IDALS\_MEM\_NULL – The `ida_mem` pointer is NULL.
- IDALS\_LMEM\_NULL – The IDALS linear solver interface has not been initialized.
- IDALS\_ILL\_INPUT – The attached linear solver is not matrix-based.

**Notes:**

By default scaling is enabled with matrix-based linear solvers when using BDF methods.

By default scaling is enabled with matrix-based linear solvers when using BDF methods.

When using a matrix-free linear solver module for the backward problem, the IDALS linear solver interface requires a function to compute an approximation to the product between the Jacobian matrix  $J(t, y)$  and a vector  $v$ . This may be performed internally using a difference-quotient approximation, or it may be supplied by the user by calling one of the following two functions:

```
int IDASetJacTimesB(void *ida_mem, int which, IDALsJacTimesSetupFnB jsetupB, IDALsJacTimesVecFnB
                    jtimesB)
```

The function *IDASetJacTimesB()* specifies the Jacobian-vector setup and product functions to be used.

**Arguments:**

- *ida\_mem* – pointer to the IDAS memory block.
- *which* – the identifier of the backward problem.
- *jsetupB* – user-defined function to set up the Jacobian-vector product. Pass NULL if no setup is necessary.
- *jtimesB* – user-defined Jacobian-vector product function.

**Return value:**

- IDALS\_SUCCESS – The optional value has been successfully set.
- IDALS\_MEM\_NULL – The *ida\_mem* memory block pointer was NULL.
- IDALS\_LMEM\_NULL – The IDALS linear solver has not been initialized.
- IDALS\_NO\_ADJ – The function *IDAAdjInit()* has not been previously called.
- IDALS\_ILL\_INPUT – The parameter *which* represented an invalid identifier.

New in version 3.0.0: Replaces the deprecated function *IDASpilsSetJacTimesB*.

```
int IDASetJacTimesBS(void *ida_mem, int which, IDALsJacTimesSetupFnBS jsetupBS, IDALsJacTimesVecFnBS
                    jtimesBS)
```

The function *IDASetJacTimesBS()* specifies the Jacobian-vector product setup and evaluation functions to be used, in the case where the backward problem depends on the forward sensitivities.

**Arguments:**

- *ida\_mem* – pointer to the IDAS memory block.
- *which* – the identifier of the backward problem.
- *jsetupBS* – user-defined function to set up the Jacobian-vector product. Pass NULL if no setup is necessary.
- *jtimesBS* – user-defined Jacobian-vector product function.

**Return value:**

- IDALS\_SUCCESS – The optional value has been successfully set.
- IDALS\_MEM\_NULL – The *ida\_mem* memory block pointer was NULL.
- IDALS\_LMEM\_NULL – The IDALS linear solver has not been initialized.
- IDALS\_NO\_ADJ – The function *IDAAdjInit()* has not been previously called.
- IDALS\_ILL\_INPUT – The parameter *which* represented an invalid identifier.

New in version 3.0.0: Replaces the deprecated function *IDASpilsSetJacTimesBS*.

When using the default difference-quotient approximation to the Jacobian-vector product for the backward problem, the user may specify the factor to use in setting increments for the finite-difference approximation, via a call to *IDASetIncrementFactorB()*.

int **IDASetIncrementFactorB**(void \*ida\_mem, int which, *sunrealtype* dqincfacB)

The function *IDASetIncrementFactorB()* specifies the factor in the increments used in the difference quotient approximations to matrix-vector products for the backward problem. This routine can be used in both the cases where the backward problem does and does not depend on the forward sensitivities.

**Arguments:**

- *ida\_mem* – pointer to the IDAS memory block.
- *which* – the identifier of the backward problem.
- *dqincfacB* – difference quotient approximation factor.

**Return value:**

- IDALS\_SUCCESS – The optional value has been successfully set.
- IDALS\_MEM\_NULL – The *ida\_mem* pointer is NULL.
- IDALS\_LMEM\_NULL – The IDALS linear solver has not been initialized.
- IDALS\_NO\_ADJ – The function *IDAAdjInit()* has not been previously called.
- IDALS\_ILL\_INPUT – The parameter *which* represented an invalid identifier.

**Notes:**

The default value is 1.0.

New in version 3.0.0: Replaces the deprecated function *IDASpilsSetIncrementFactorB*.

Additionally, When using the internal difference quotient for the backward problem, the user may also optionally supply an alternative residual function for use in the Jacobian-vector product approximation by calling *IDASetJacTimesResFnB()*. The alternative residual side function should compute a suitable (and differentiable) approximation to the residual function provided to *IDAInitB()* or *IDAInitBS()*. For example, as done in [31] for the forward integration of an ODE in explicit form without sensitivity analysis, the alternative function may use lagged values when evaluating a nonlinearity in the right-hand side to avoid differencing a potentially non-differentiable factor.

int **IDASetJacTimesResFnB**(void \*ida\_mem, int which, *IDAResFn* jtimesResFn)

The function *IDASetJacTimesResFnB()* specifies an alternative DAE residual function for use in the internal Jacobian-vector product difference quotient approximation for the backward problem.

**Arguments:**

- *ida\_mem* – pointer to the IDAS memory block.
- *which* – the identifier of the backward problem.
- *jtimesResFn* – is the C function which computes the alternative DAE residual function to use in Jacobian-vector product difference quotient approximations. This function has the form `res(t, yy, yp, resval, user_data)`. For full details see §5.1.4.1.

**Return value:**

- IDALS\_SUCCESS – The optional value has been successfully set.
- IDALS\_MEM\_NULL – The *ida\_mem* pointer is NULL.
- IDALS\_LMEM\_NULL – The IDALS linear solver has not been initialized.
- IDALS\_NO\_ADJ – The function *IDAAdjInit()* has not been previously called.
- IDALS\_ILL\_INPUT – The parameter *which* represented an invalid identifier or the internal difference quotient approximation is disabled.

**Notes:**

The default is to use the residual function provided to *IDAInit()* in the internal difference quotient. If the input residual function is NULL, the default is used.

This function must be called *after* the IDALS linear solver interface has been initialized through a call to *IDASetLinearSolverB()*.

When using an iterative linear solver for the backward problem, the user may supply a preconditioning operator to aid in solution of the system, or she/he may adjust the convergence tolerance factor for the iterative linear solver. These may be accomplished through calling the following functions:

```
int IDASetPreconditionerB(void *ida_mem, int which, IDALsPrecSetupFnB psetupB, IDALsPrecSolveFnB
                        psolveB)
```

The function *IDASetPreconditionerB()* specifies the preconditioner setup and solve functions for the backward integration.

**Arguments:**

- *ida\_mem* – pointer to the IDAS memory block.
- *which* – the identifier of the backward problem.
- *psetupB* – user-defined preconditioner setup function.
- *psolveB* – user-defined preconditioner solve function.

**Return value:**

- *IDALS\_SUCCESS* – The optional value has been successfully set.
- *IDALS\_MEM\_NULL* – The *ida\_mem* memory block pointer was NULL.
- *IDALS\_LMEM\_NULL* – The IDALS linear solver has not been initialized.
- *IDALS\_NO\_ADJ* – The function *IDAAdjInit()* has not been previously called.
- *IDALS\_ILL\_INPUT* – The parameter *which* represented an invalid identifier.

**Notes:**

The *psetupB* argument may be NULL if no setup operation is involved in the preconditioner.

New in version 3.0.0: Replaces the deprecated function *IDASpilsSetPreconditionerB*.

```
int IDASetPreconditionerBS(void *ida_mem, int which, IDALsPrecSetupFnBS psetupBS, IDALsPrecSolveFnBS
                        psolveBS)
```

The function *IDASetPreconditionerBS()* specifies the preconditioner setup and solve functions for the backward integration, in the case where the backward problem depends on the forward sensitivities.

**Arguments:**

- *ida\_mem* – pointer to the IDAS memory block.
- *which* – the identifier of the backward problem.
- *psetupBS* – user-defined preconditioner setup function.
- *psolveBS* – user-defined preconditioner solve function.

**Return value:**

- *IDALS\_SUCCESS* – The optional value has been successfully set.
- *IDALS\_MEM\_NULL* – The *ida\_mem* memory block pointer was NULL.
- *IDALS\_LMEM\_NULL* – The IDALS linear solver has not been initialized.
- *IDALS\_NO\_ADJ* – The function *IDAAdjInit()* has not been previously called.

- IDALS\_ILL\_INPUT – The parameter which represented an invalid identifier.

**Notes:**

The psetupBS argument may be NULL if no setup operation is involved in the preconditioner.

New in version 3.0.0: Replaces the deprecated function IDASpilsSetPreconditionerBS.

int **IDASetEpsLinB**(void \*ida\_mem, int which, *sunrealtype* eplifacB)

The function *IDASetEpsLinB()* specifies the factor by which the Krylov linear solver's convergence test constant is reduced from the nonlinear iteration test constant. (See §2.2). This routine can be used in both the cases where the backward problem does and does not depend on the forward sensitivities.

**Arguments:**

- ida\_mem – pointer to the IDAS memory block.
- which – the identifier of the backward problem.
- eplifacB – linear convergence safety factor  $\geq 0.0$ .

**Return value:**

- IDALS\_SUCCESS – The optional value has been successfully set.
- IDALS\_MEM\_NULL – The ida\_mem pointer is NULL.
- IDALS\_LMEM\_NULL – The IDALS linear solver has not been initialized.
- IDALS\_NO\_ADJ – The function *IDAAadjInit()* has not been previously called.
- IDALS\_ILL\_INPUT – The parameter which represented an invalid identifier.

**Notes:**

The default value is 0.05.

Passing a value eplifacB = 0.0 also indicates using the default value.

New in version 3.0.0: Replaces the deprecated function IDASpilsSetEpsLinB.

int **IDASetLSNormFactorB**(void \*ida\_mem, int which, *sunrealtype* nrmfac)

The function *IDASetLSNormFactorB()* specifies the factor to use when converting from the integrator tolerance (WRMS norm) to the linear solver tolerance (L2 norm) for Newton linear system solves e.g.,  $tol\_L2 = fac * tol\_WRMS$ . This routine can be used in both the cases where the backward problem does and does not depend on the forward sensitivities.

**Arguments:**

- ida\_mem – pointer to the IDAS memory block.
- which – the identifier of the backward problem.
- nrmfac – the norm conversion factor. If nrmfac is:
  - $> 0$  then the provided value is used.
  - $= 0$  then the conversion factor is computed using the vector length i.e.,  $nrmfac = N\_VGetLength(y)$  default.
  - $< 0$  then the conversion factor is computed using the vector dot product  $nrmfac = N\_VDotProd(v, v)$  where all the entries of  $v$  are one.

**Return value:**

- IDALS\_SUCCESS – The optional value has been successfully set.
- IDALS\_MEM\_NULL – The ida\_mem pointer is NULL.

- IDALS\_LMEM\_NULL – The IDALS linear solver has not been initialized.
- IDALS\_NO\_ADJ – The function `IDAAdjInit()` has not been previously called.
- IDALS\_ILL\_INPUT – The parameter which represented an invalid identifier.

**Notes:**

This function must be called after the IDALS linear solver interface has been initialized through a call to `IDASetLinearSolverB()`.

Prior to the introduction of `N_VGetLength` in SUNDIALS v5.0.0 (IDAS v4.0.0) the value of `nrmfac` was computed using the vector dot product i.e., the `nrmfac < 0` case.

### 5.5.2.11 Optional output functions for the backward problem

#### Main solver optional output functions

The user of the adjoint module in IDAS has access to any of the optional output functions described in §5.1.3.12, both for the main solver and for the linear solver modules. The first argument of these `IDAGet*` and `IDA*Get*` functions is the pointer to the IDAS memory block for the backward problem. In order to call any of these functions, the user must first call the following function to obtain this pointer:

```
void *IDAGetAdjIDABmem(void *ida_mem, int which)
```

The function `IDAGetAdjIDABmem()` returns a pointer to the IDAS memory block for the backward problem.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block created by `IDACreate()`.
- `which` – the identifier of the backward problem.

**Return value:**

- The return value, `ida_memB` (of type `void *`), is a pointer to the idas memory for the backward problem.

**Warning:** The user should not modify `ida_memB` in any way.

Optional output calls should pass `ida_memB` as the first argument; thus, for example, to get the number of integration steps: `flag = IDAGetNumSteps(idas_memB, &nsteps)`.

To get values of the *forward* solution during a backward integration, use the following function. The input value of `t` would typically be equal to that at which the backward solution has just been obtained with `IDAGetB()`. In any case, it must be within the last checkpoint interval used by `IDASolveB()`.

```
int IDAGetAdjY(void *ida_mem, sunrealtype t, N_Vector y, N_Vector yp)
```

The function `IDAGetAdjY()` returns the interpolated value of the forward solution  $y$  and its derivative during a backward integration.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block created by `IDACreate()`.
- `t` – value of the independent variable at which  $y$  is desired input.
- $y$  – forward solution  $y(t)$ .
- $yp$  – forward solution derivative  $\dot{y}(t)$ .

**Return value:**

- IDA\_SUCCESS – *IDAGetAdjY()* was successful.
- IDA\_MEM\_NULL – *ida\_mem* was NULL.
- IDA\_GETY\_BADT – The value of *t* was outside the current checkpoint interval.

**Warning:** The user must allocate space for *y* and *yp*.

int **IDAGetAdjCheckPointsInfo**(void \*ida\_mem, *IDAadjCheckpointRec* \*ckpnt)

The function *IDAGetAdjCheckPointsInfo()* loads an array of *ncheck* + 1 records of type *IDAadjCheckpointRec*. The user must allocate space for the array *ckpnt*.

**Arguments:**

- *ida\_mem* – pointer to the IDAS memory block created by *IDACreate()*.
- *ckpnt* – array of *ncheck*+1 checkpoint records, each of type *IDAadjCheckpointRec()*.

**Return value:**

- void

The checkpoint structure is defined as

struct **IDAadjCheckpointRec**

void \***my\_addr**

The address of current checkpoint in *ida\_mem->ida\_adj\_mem*

void \***next\_addr**

The address of next checkpoint.

*sunrealtype* **t0**

The start time of the checkpoint interval

*sunrealtype* **t1**

The end time of the checkpoint interval

long int **nstep**

The step counter at *t0*

int **order**

The method order at *t0*

*sunrealtype* **step**

The step size at *t0*

**Initial condition calculation optional output function**

int **IDAGetConsistentICB**(void \*ida\_mem, int which, *N\_Vector* yB0\_mod, *N\_Vector* ypB0\_mod)

The function *IDAGetConsistentICB()* returns the corrected initial conditions for backward problem calculated by *IDACalcICB()*.

**Arguments:**

- *ida\_mem* – pointer to the IDAS memory block.
- *which* – is the identifier of the backward problem.



- `yB0_mod` – consistent initial vector.
- `ypB0_mod` – consistent initial derivative vector.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.
- `IDA_NO_ADJ` – `IDAAdjInit()` has not been previously called.
- `IDA_ILL_INPUT` – Parameter `which` did not refer a valid backward problem identifier.

**Notes:**

If the consistent solution vector or consistent derivative vector is not desired, pass NULL for the corresponding argument.

**Warning:** The user must allocate space for `yB0_mod` and `ypB0_mod` (if not NULL).

### 5.5.2.12 Backward integration of quadrature equations

Not only the backward problem but also the backward quadrature equations may or may not depend on the forward sensitivities. Accordingly, one of the `IDAQuadInitB()` or `IDAQuadInitBS()` should be used to allocate internal memory and to initialize backward quadratures. For any other operation (extraction, optional input/output, reinitialization, deallocation), the same function is called regardless of whether or not the quadratures are sensitivity-dependent.

#### Backward quadrature initialization functions

The function `IDAQuadInitB()` initializes and allocates memory for the backward integration of quadrature equations that do not depend on forward sensitivities. It has the following form:

int `IDAQuadInitB`(void \*ida\_mem, int which, `IDAQuadRhsFnB` rhsQB, *N\_Vector* yQB0)

The function `IDAQuadInitB()` provides required problem specifications, allocates internal memory, and initializes backward quadrature integration.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `which` – the identifier of the backward problem.
- `rhsQB` – is the C function which computes  $f_{QB}$ , the residual of the backward quadrature equations. This function has the form `rhsQB(t, y, yp, yB, ypB, rhsvalBQ, user_dataB)` see §5.5.3.3.
- `yQB0` – is the value of the quadrature variables at `tB0`.

**Return value:**

- `IDA_SUCCESS` – The call to `IDAQuadInitB()` was successful.
- `IDA_MEM_NULL` – `ida_mem` was NULL.
- `IDA_NO_ADJ` – The function `IDAAdjInit()` has not been previously called.
- `IDA_MEM_FAIL` – A memory allocation request has failed.
- `IDA_ILL_INPUT` – The parameter `which` is an invalid identifier.

int **IDAQuadInitBS**(void \*ida\_mem, int which, *IDAQuadRhsFnBS* rhsQBS, *N\_Vector* yQBS0)

The function *IDAQuadInitBS()* provides required problem specifications, allocates internal memory, and initializes backward quadrature integration with sensitivities.

**Arguments:**

- *ida\_mem* – pointer to the IDAS memory block.
- *which* – the identifier of the backward problem.
- *rhsQBS* – is the C function which computes  $fQBS$ , the residual of the backward quadrature equations. This function has the form *rhsQBS*(*t*, *y*, *yp*, *yS*, *ypS*, *yB*, *ypB*, *rhsvalBQS*, *user\_dataB*) see §5.5.3.4.
- *yQBS0* – is the value of the sensitivity-dependent quadrature variables at *tB0*.

**Return value:**

- IDA\_SUCCESS – The call to *IDAQuadInitBS()* was successful.
- IDA\_MEM\_NULL – *ida\_mem* was NULL.
- IDA\_NO\_ADJ – The function *IDAAdjInit()* has not been previously called.
- IDA\_MEM\_FAIL – A memory allocation request has failed.
- IDA\_ILL\_INPUT – The parameter *which* is an invalid identifier.

The integration of quadrature equations during the backward phase can be re-initialized by calling the following function. Before calling *IDAQuadReInitB()* for a new backward problem, call any desired solution extraction functions *IDAGet\*\** associated with the previous backward problem.

int **IDAQuadReInitB**(void \*ida\_mem, int which, *N\_Vector* yQB0)

The function *IDAQuadReInitB()* re-initializes the backward quadrature integration.

**Arguments:**

- *ida\_mem* – pointer to the IDAS memory block.
- *which* – the identifier of the backward problem.
- *yQB0* – is the value of the quadrature variables at *tB0*.

**Return value:**

- IDA\_SUCCESS – The call to *IDAQuadReInitB()* was successful.
- IDA\_MEM\_NULL – *ida\_mem* was NULL.
- IDA\_NO\_ADJ – The function *IDAAdjInit()* has not been previously called.
- IDA\_MEM\_FAIL – A memory allocation request has failed.
- IDA\_NO\_QUAD – Quadrature integration was not activated through a previous call to *IDAQuadInitB()*.
- IDA\_ILL\_INPUT – The parameter *which* is an invalid identifier.

**Notes:**

*IDAQuadReInitB()* can be used after a call to either *IDAQuadInitB()* or *IDAQuadInitBS()*.

## Backward quadrature extraction function

To extract the values of the quadrature variables at the last return time of `IDASolveB()`, IDAS provides a wrapper for the function `IDAGetQuad()`. The call to this function has the form

```
int IDAGetQuadB(void *ida_mem, int which, sunrealtype *tret, N_Vector yQB)
```

The function `IDAGetQuadB()` returns the quadrature solution vector after a successful return from `IDASolveB()`.

### Arguments:

- `ida_mem` – pointer to the IDAS memory.
- `tret` – the time reached by the solver output.
- `which` – the identifier of the backward problem.
- `yQB` – the computed quadrature vector.

### Return value:

- `IDA_SUCCESS` – `IDAGetQuadB()` was successful.
- `IDA_MEM_NULL` – `ida_mem` is NULL.
- `IDA_NO_ADJ` – The function `IDAAdjInit()` has not been previously called.
- `IDA_NO_QUAD` – Quadrature integration was not initialized.
- `IDA_BAD_DKY` – `yQB` was NULL.
- `IDA_ILL_INPUT` – The parameter `which` is an invalid identifier.

### Notes:

To obtain the quadratures associated with a given backward problem at some other time within the last integration step, first obtain a pointer to the proper IDAS memory structure by calling `IDAGetAdjIDABmem()` and then use it to call `IDAGetQuadDky()`.

**Warning:** The user must allocate space for `yQB`.

## Optional input/output functions for backward quadrature integration

Optional values controlling the backward integration of quadrature equations can be changed from their default values through calls to one of the following functions which are wrappers for the corresponding optional input functions defined in §5.2.4. The user must specify the identifier `which` of the backward problem for which the optional values are specified.

```
flag = IDASetQuadErrConB(ida_mem, which, errconQ);
flag = IDAQuadSStolerancesB(ida_mem, which, reltolQ, abstolQ);
flag = IDAQuadSVtolerancesB(ida_mem, which, reltolQ, abstolQ);
```

Their return value `flag` (of type `int`) can have any of the return values of its counterparts, but it can also be `IDA_NO_ADJ` if the function `IDAAdjInit()` has not been previously called or `IDA_ILL_INPUT` if the parameter `which` was an invalid identifier.

Access to optional outputs related to backward quadrature integration can be obtained by calling the corresponding `IDAGetQuad*` functions (see §5.2.5). A pointer `ida_memB` to the IDAS memory block for the backward problem, required as the first argument of these functions, can be obtained through a call to the functions `IDAGetAdjIDABmem()`.

### 5.5.3 User-supplied functions for adjoint sensitivity analysis

In addition to the required DAE residual function and any optional functions for the forward problem, when using the adjoint sensitivity module in IDAS, the user must supply one function defining the backward problem DAE and, optionally, functions to supply Jacobian-related information and one or two functions that define the preconditioner (if applicable for the choice of `SUNLinearSolver` object) for the backward problem. Type definitions for all these user-supplied functions are given below.

#### 5.5.3.1 DAE residual for the backward problem

The user must provide a `resB` function of type `IDAResFnB` defined as follows:

```
typedef int (*IDAResFnB)(sunrealtype t, N_Vector y, N_Vector yp, N_Vector yB, N_Vector ypB, N_Vector resvalB,  
void *user_dataB)
```

This function evaluates the residual of the backward problem DAE system. This could be (2.19) or (2.24).

##### Arguments:

- `t` – is the current value of the independent variable.
- `y` – is the current value of the forward solution vector.
- `yp` – is the current value of the forward solution derivative vector.
- `yB` – is the current value of the backward dependent variable vector.
- `ypB` – is the current value of the backward dependent derivative vector.
- `resvalB` – is the output vector containing the residual for the backward DAE problem.
- `user_dataB` – is a pointer to user data, same as passed to `IDASetUserDataB`.

##### Return value:

An `IDAResFnB` should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct), or a negative value if an unrecoverable failure occurred (in which case the integration stops and `IDASolveB()` returns `IDA_RESFUNC_FAIL`).

##### Notes:

Allocation of memory for `resvalB` is handled within IDAS. The `y`, `yp`, `yB`, `ypB`, and `resvalB` arguments are all of type `N_Vector`, but `yB`, `ypB`, and `resvalB` typically have different internal representations from `y` and `yp`. It is the user's responsibility to access the vector data consistently (including the use of the correct accessor macros from each `N_Vector` implementation). The `user_dataB` pointer is passed to the user's `resB` function every time it is called and can be the same as the `user_data` pointer used for the forward problem.

**Warning:** Before calling the user's `resB` function, IDAS needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, IDAS triggers an unrecoverable failure in the residual function which will halt the integration and `IDASolveB()` will return `IDA_RESFUNC_FAIL`.

### 5.5.3.2 DAE residual for the backward problem depending on the forward sensitivities

The user must provide a `resBS` function of type `IDAResFnBS` defined as follows:

```
typedef int (*IDAResFnBS)(sunrealtype t, N_Vector y, N_Vector yp, N_Vector *yS, N_Vector *ypS, N_Vector yB,
N_Vector ypB, N_Vector resvalB, void *user_dataB)
```

This function evaluates the residual of the backward problem DAE system. This could be (2.19) or (2.24).

#### Arguments:

- `t` – is the current value of the independent variable.
- `y` – is the current value of the forward solution vector.
- `yp` – is the current value of the forward solution derivative vector.
- `yS` – a pointer to an array of `Ns` vectors containing the sensitivities of the forward solution.
- `ypS` – a pointer to an array of `Ns` vectors containing the derivatives of the forward sensitivities.
- `yB` – is the current value of the backward dependent variable vector.
- `ypB` – is the current value of the backward dependent derivative vector.
- `resvalB` – is the output vector containing the residual for the backward DAE problem.
- `user_dataB` – is a pointer to user data, same as passed to `IDASetUserDataB`.

#### Return value:

An `IDAResFnBS` should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct), or a negative value if an unrecoverable error occurred (in which case the integration stops and `IDASolveB()` returns `IDA_RESFUNC_FAIL`).

#### Notes:

Allocation of memory for `resvalB` is handled within IDAS. The `y`, `yp`, `yB`, `ypB`, and `resvalB` arguments are all of type `N_Vector`, but `yB`, `ypB`, and `resvalB` typically have different internal representations from `y` and `yp`. Likewise for each `yS[i]` and `ypS[i]`. It is the user's responsibility to access the vector data consistently (including the use of the correct accessor macros from each `N_Vector` implementation). The `user_dataB` pointer is passed to the user's `resBS` function every time it is called and can be the same as the `user_data` pointer used for the forward problem.

**Warning:** Before calling the user's `resBS` function, IDAS needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, IDAS triggers an unrecoverable failure in the residual function which will halt the integration and `IDASolveB()` will return `IDA_RESFUNC_FAIL`.

### 5.5.3.3 Quadrature right-hand side for the backward problem

The user must provide an `fQB` function of type `IDAQuadRhsFnB` defined by

```
typedef int (*IDAQuadRhsFnB)(sunrealtype t, N_Vector y, N_Vector yp, N_Vector yB, N_Vector ypB, N_Vector
rhsvalBQ, void *user_dataB)
```

This function computes the quadrature equation right-hand side for the backward problem.

#### Arguments:

- `t` – is the current value of the independent variable.
- `y` – is the current value of the forward solution vector.

- `yp` – is the current value of the forward solution derivative vector.
- `yB` – is the current value of the backward dependent variable vector.
- `ypB` – is the current value of the backward dependent derivative vector.
- `rhsvalBQ` – is the output vector containing the residual for the backward quadrature equations.
- `user_dataB` – is a pointer to user data, same as passed to `IDASetUserDataB`.

**Return value:**

An `IDAQuadRhsFnB` should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and `IDASolveB()` returns `IDA_QRHSFUNC_FAIL`).

**Notes:**

Allocation of memory for `rhsvalBQ` is handled within IDAS. The `y`, `yp`, `yB`, `ypB`, and `rhsvalBQ` arguments are all of type `N_Vector`, but they typically all have different internal representations. It is the user's responsibility to access the vector data consistently (including the use of the correct accessor macros from each `N_Vector` implementation). For the sake of computational efficiency, the vector functions in the two `N_Vector` implementations provided with IDAS do not perform any consistency checks with respect to their `N_Vector` arguments (see §6). The `user_dataB` pointer is passed to the user's `fQB` function every time it is called and can be the same as the `user_data` pointer used for the forward problem.

**Warning:** Before calling the user's `fQB` function, IDAS needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, IDAS triggers an unrecoverable failure in the quadrature right-hand side function which will halt the integration and `IDASolveB()` will return `IDA_QRHSFUNC_FAIL`.

#### 5.5.3.4 Sensitivity-dependent quadrature right-hand side for the backward problem

The user must provide an `fQBS` function of type `IDAQuadRhsFnBS` defined by

```
typedef int (*IDAQuadRhsFnBS)(sunrealtype t, N_Vector y, N_Vector yp, N_Vector *yS, N_Vector *ypS, N_Vector yB, N_Vector ypB, N_Vector rhsvalBQS, void *user_dataB)
```

This function computes the quadrature equation residual for the backward problem.

**Arguments:**

- `t` – is the current value of the independent variable.
- `y` – is the current value of the forward solution vector.
- `yp` – is the current value of the forward solution derivative vector.
- `yS` – a pointer to an array of `Ns` vectors containing the sensitivities of the forward solution.
- `ypS` – a pointer to an array of `Ns` vectors containing the derivatives of the forward sensitivities.
- `yB` – is the current value of the backward dependent variable vector.
- `ypB` – is the current value of the backward dependent derivative vector.
- `rhsvalBQS` – is the output vector containing the residual for the backward quadrature equations.
- `user_dataB` – is a pointer to user data, same as passed to `IDASetUserDataB`.

**Return value:**

An `IDAQuadRhsFnBS` should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and `IDASolveB()` returns `IDA_QRHSFUNC_FAIL`).

**Notes:**

Allocation of memory for `rhsvalBQS` is handled within IDAS. The `y`, `yp`, `yB`, `ypB`, and `rhsvalBQS` arguments are all of type `N_Vector`, but they typically do not all have the same internal representations. Likewise for each `yS[i]` and `ypS[i]`. It is the user's responsibility to access the vector data consistently (including the use of the correct accessor macros from each `N_Vector` implementation). The `user_dataB` pointer is passed to the user's `fQBS` function every time it is called and can be the same as the `user_data` pointer used for the forward problem.

**Warning:** Before calling the user's `fQBS` function, IDAS needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, IDAS triggers an unrecoverable failure in the quadrature right-hand side function which will halt the integration and `IDASolveB()` will return `IDA_QRHSFUNC_FAIL`.

### 5.5.3.5 Jacobian construction for the backward problem (matrix-based linear solvers)

If a matrix-based linear solver module is used for the backward problem (i.e., `IDASetLinearSolverB()` is called with non-NULL `SUNMatrix` argument in the step described in §5.5.1), the user may provide a function of type `IDALSJacFnB` or `IDALSJacFnBS`, defined as follows:

```
typedef int (*IDALSJacFnB)(sunrealtype tt, sunrealtype c_jB, N_Vector yy, N_Vector yp, N_Vector yyB, N_Vector ypB, N_Vector rrB, SUNMatrix JacB, void *user_dataB, N_Vector tmp1B, N_Vector tmp2B, N_Vector tmp3B)
```

This function computes the Jacobian of the backward problem (or an approximation to it).

**Arguments:**

- `tt` – is the current value of the independent variable.
- `c_jB` – is the scalar in the system Jacobian, proportional to the inverse of the step size ( $\alpha$  in (2.7)).
- `yy` – is the current value of the forward solution vector.
- `yp` – is the current value of the forward solution derivative vector.
- `yB` – is the current value of the backward dependent variable vector.
- `ypB` – is the current value of the backward dependent derivative vector.
- `rrB` – is the current value of the residual for the backward problem.
- `JacB` – is the output approximate Jacobian matrix.
- `user_dataB` – is a pointer to user data — the parameter passed to `IDASetUserDataB`.
- `tmp1B`, `tmp2B`, `tmp3B` – are pointers to memory allocated for variables of type `N_Vector` which can be used by the `IDALSJacFnB` function as temporary storage or work space.

**Return value:**

An `IDALSJacFnB` should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct, while IDALS sets `last_flag` to `IDALS_JACFUNC_RECVR`), or a negative value if it failed unrecoverably (in which case the integration is halted, `IDASolveB()` returns `IDA_LSETUP_FAIL` and IDALS sets `last_flag` to `IDALS_JACFUNC_UNRECVR`).

**Notes:**

A user-supplied Jacobian function must load the matrix `JacB` with an approximation to the Jacobian matrix at the point  $(tt, yy, yB)$ , where `yy` is the solution of the original IVP at time `tt`, and `yB` is the solution of the backward problem at the same time. Information regarding the structure of the specific `SUNMatrix` structure (e.g. number of rows, upper/lower bandwidth, sparsity type) may be obtained through using the implementation-specific `SUNMatrix` interface functions (see Chapter §7 for details). With direct linear



solvers (i.e., linear solvers with type `SUNLINEARSOLVER_DIRECT`), the Jacobian matrix  $J(t, y)$  is zeroed out prior to calling the user-supplied Jacobian function so only nonzero elements need to be loaded into `JacB`.

**Warning:** Before calling the user's `IDALSJacFnB`, IDAS needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, IDAS triggers an unrecoverable failure in the Jacobian function which will halt the integration (`IDASolveB()` returns `IDA_LSETUP_FAIL` and IDALS sets `last_flag` to `IDALS_JACFUNC_UNRECVR`).

New in version 3.0.0: Replaces the deprecated type `IDADlsJacFnB`.

```
typedef int (*IDALSJacFnBS)(sunrealtype tt, sunrealtype c_jB, N_Vector yy, N_Vector yp, N_Vector *yS, N_Vector *ypS, N_Vector yyB, N_Vector ypB, N_Vector rrB, SUNMatrix JacB, void *user_dataB, N_Vector tmp1B, N_Vector tmp2B, N_Vector tmp3B);
```

This function computes the Jacobian of the backward problem (or an approximation to it), in the case where the backward problem depends on the forward sensitivities.

#### Arguments:

- `tt` – is the current value of the independent variable.
- `c_jB` – is the scalar in the system Jacobian, proportional to the inverse of the step size ( $\alpha$  in (2.7)).
- `yy` – is the current value of the forward solution vector.
- `yp` – is the current value of the forward solution derivative vector.
- `yS` – a pointer to an array of `Ns` vectors containing the sensitivities of the forward solution.
- `ypS` – a pointer to an array of `Ns` vectors containing the derivatives of the forward solution sensitivities.
- `yB` – is the current value of the backward dependent variable vector.
- `ypB` – is the current value of the backward dependent derivative vector.
- `rrb` – is the current value of the residual for the backward problem.
- `JacB` – is the output approximate Jacobian matrix.
- `user_dataB` – is a pointer to user data — the parameter passed to `IDASetUserDataB`.
- `tmp1B`, `tmp2B`, `tmp3B` – are pointers to memory allocated for variables of type `N_Vector` which can be used by `IDALSJacFnBS` as temporary storage or work space.

#### Return value:

An `IDALSJacFnBS` should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct, while IDALS sets `last_flag` to `IDALS_JACFUNC_RECVR`), or a negative value if it failed unrecoverably (in which case the integration is halted, `IDASolveB()` returns `IDA_LSETUP_FAIL` and IDALS sets `last_flag` to `IDALS_JACFUNC_UNRECVR`).

#### Notes:

A user-supplied dense Jacobian function must load the matrix `JacB` with an approximation to the Jacobian matrix at the point  $(tt, yy, yS, yB)$ , where `yy` is the solution of the original IVP at time `tt`, `yS` is the array of forward sensitivities at time `tt`, and `yB` is the solution of the backward problem at the same time. Information regarding the structure of the specific `SUNMatrix` structure (e.g. number of rows, upper/lower bandwidth, sparsity type) may be obtained through using the implementation-specific `SUNMatrix` interface functions (see Chapter §7 for details). With direct linear solvers (i.e., linear solvers with type `SUNLINEARSOLVER_DIRECT`), the Jacobian matrix  $J(t, y)$  is zeroed out prior to calling the user-supplied Jacobian function so only nonzero elements need to be loaded into `JacB`.



**Warning:** Before calling the user's `IDALSJacFnBS`, IDAS needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, IDAS triggers an unrecoverable failure in the Jacobian function which will halt the integration (`IDASolveB()` returns `IDA_LSETUP_FAIL` and IDALS sets `last_flag` to `IDALS_JACFUNC_UNRECVR`).

New in version 3.0.0: Replaces the deprecated type `IDADlsJacFnBS`.

### 5.5.3.6 Jacobian-vector product for the backward problem (matrix-free linear solvers)

If a matrix-free linear solver is selected for the backward problem (i.e., `IDASetLinearSolverB()` is called with NULL-valued `SUNMatrix` argument in the steps described in §5.5.1), the user may provide a function of type `IDALSJacTimesVecFnB` or `IDALSJacTimesVecFnBS` in the following form, to compute matrix-vector products  $Jv$ . If such a function is not supplied, the default is a difference quotient approximation to these products.

```
typedef int (*IDALSJacTimesVecFnB)(sunrealtype t, N_Vector yy, N_Vector yp, N_Vector yB, N_Vector ypB,
N_Vector resvalB, N_Vector vB, N_Vector JvB, sunrealtype cjB, void *user_dataB, N_Vector tmp1B, N_Vector
tmp2B)
```

This function computes the action of the backward problem Jacobian  $JB$  on a given vector  $vB$ .

#### Arguments:

- $t$  – is the current value of the independent variable.
- $yy$  – is the current value of the forward solution vector.
- $yp$  – is the current value of the forward solution derivative vector.
- $yB$  – is the current value of the backward dependent variable vector.
- $ypB$  – is the current value of the backward dependent derivative vector.
- $resvalB$  – is the current value of the residual for the backward problem.
- $vB$  – is the vector by which the Jacobian must be multiplied.
- $JvB$  – is the computed output vector,  $JB*vB$ .
- $cjB$  – is the scalar in the system Jacobian, proportional to the inverse of the step size ( $\alpha$  in (2.7)).
- $user\_dataB$  – is a pointer to user data — the same as the `user_dataB` parameter passed to `IDASetUserDataB`.
- $tmp1B, tmp2B$  – are pointers to memory allocated for variables of type `N_Vector` which can be used by `IDALSJacTimesVecFnB` as temporary storage or work space.

#### Return value:

The return value of a function of type `IDALSJacTimesVecFnB` should be 0 if successful or nonzero if an error was encountered, in which case the integration is halted.

#### Notes:

A user-supplied Jacobian-vector product function must load the vector  $JvB$  with the product of the Jacobian of the backward problem at the point  $(t, y, yB)$  and the vector  $vB$ . Here,  $y$  is the solution of the original IVP at time  $t$  and  $yB$  is the solution of the backward problem at the same time. The rest of the arguments are equivalent to those passed to a function of type `IDALSJacTimesVecFn` (see §5.1.4.5). If the backward problem is the adjoint of  $\dot{y} = f(t, y)$ , then this function is to compute  $-(\partial f / \partial y_i)^T v_B$ .

New in version 3.0.0: Replaces the deprecated type `IDASpilsJacTimesVecFnB`.

```
typedef int (*IDALSJacTimesVecFnBS)(sunrealtype t, N_Vector yy, N_Vector yp, N_Vector *yyS, N_Vector *ypS,  
N_Vector yB, N_Vector ypB, N_Vector resvalB, N_Vector vB, N_Vector JvB, sunrealtype cjB, void *user_dataB,  
N_Vector tmp1B, N_Vector tmp2B)
```

This function computes the action of the backward problem Jacobian JB on a given vector vB, in the case where the backward problem depends on the forward sensitivities.

**Arguments:**

- t – is the current value of the independent variable.
- yy – is the current value of the forward solution vector.
- yp – is the current value of the forward solution derivative vector.
- yyS – a pointer to an array of Ns vectors containing the sensitivities of the forward solution.
- ypS – a pointer to an array of Ns vectors containing the derivatives of the forward sensitivities.
- yB – is the current value of the backward dependent variable vector.
- ypB – is the current value of the backward dependent derivative vector.
- resvalB – is the current value of the residual for the backward problem.
- vB – is the vector by which the Jacobian must be multiplied.
- JvB – is the computed output vector, JB\*vB.
- cjB – is the scalar in the system Jacobian, proportional to the inverse of the step size (  $\alpha$  in (2.7) ).
- user\_dataB – is a pointer to user data — the same as the user\_dataB parameter passed to IDASetsUserDataB.
- tmp1B, tmp2B – are pointers to memory allocated for variables of type N\_Vector which can be used by IDALSJacTimesVecFnBS as temporary storage or work space.

**Return value:**

The return value of a function of type IDALSJacTimesVecFnBS should be 0 if successful or nonzero if an error was encountered, in which case the integration is halted.

**Notes:**

A user-supplied Jacobian-vector product function must load the vector JvB with the product of the Jacobian of the backward problem at the point (t, y, yB) and the vector vB. Here, y is the solution of the original IVP at time t and yB is the solution of the backward problem at the same time. The rest of the arguments are equivalent to those passed to a function of type IDALSJacTimesVecFn (see §5.1.4.5).

New in version 3.0.0: Replaces the deprecated type IDASpilsJacTimesVecFnBS.

### 5.5.3.7 Jacobian-vector product setup for the backward problem (matrix-free linear solvers)

If the user's Jacobian-times-vector requires that any Jacobian-related data be preprocessed or evaluated, then this needs to be done in a user-supplied function of type *IDALSJacTimesSetupFnB* or *IDALSJacTimesSetupFnBS*, defined as follows:

```
typedef int (*IDALSJacTimesSetupFnB)(sunrealtype tt, N_Vector yy, N_Vector yp, N_Vector yB, N_Vector ypB,  
N_Vector resvalB, sunrealtype cjB, void *user_dataB)
```

This function preprocesses and/or evaluates Jacobian data needed by the Jacobian-times-vector routine for the backward problem.

**Arguments:**

- tt – is the current value of the independent variable.

- `yy` – is the current value of the dependent variable vector,  $y(t)$  .
- `yp` – is the current value of  $\dot{y}(t)$  .
- `yB` – is the current value of the backward dependent variable vector.
- `ypB` – is the current value of the backward dependent derivative vector.
- `resvalB` – is the current value of the residual for the backward problem.
- `cjB` – is the scalar in the system Jacobian, proportional to the inverse of the step size (  $\alpha$  in (2.7) ).
- `user_dataB` – is a pointer to user data — the same as the `user_dataB` parameter passed to `IDASetUserDataB`.

**Return value:**

The value returned by the Jacobian-vector setup function should be if successful, positive for a recoverable error (in which case the step will be retried), or negative for an unrecoverable error (in which case the integration is halted).

**Notes:**

Each call to the Jacobian-vector setup function is preceded by a call to the backward problem residual user function with the same (`t`, `y`, `yp`, `yB`, `ypB`) arguments. Thus, the setup function can use any auxiliary data that is computed and saved during the evaluation of the DAE residual. If the user's `IDALSJacTimesVecFnB` function uses difference quotient approximations, it may need to access quantities not in the call list. These include the current stepsize, the error weights, etc. To obtain these, the user will need to add a pointer to `ida_mem` to `user_dataB` and then use the `IDAGet*` functions described in §5.1.3.12. The unit roundoff can be accessed as `SUN_UNIT_ROUNDOFF` defined in `sundials_types.h`.

New in version 3.0.0: Replaces the deprecated type `IDASpilsJacTimesSetupFnB`.

```
typedef int (*IDALSJacTimesSetupFnBS)(sunrealtype tt, N_Vector yy, N_Vector yp, N_Vector *yyS, N_Vector
*ypS, N_Vector yB, N_Vector ypB, N_Vector resvalB, sunrealtype cjB, void *user_dataB)
```

This function preprocesses and/or evaluates Jacobian data needed by the Jacobian-times-vector routine for the backward problem, in the case that the backward problem depends on the forward sensitivities.

**Arguments:**

- `tt` – is the current value of the independent variable.
- `yy` – is the current value of the dependent variable vector,  $y(t)$  .
- `yp` – is the current value of  $\dot{y}(t)$  .
- `yyS` – a pointer to an array of `Ns` vectors containing the sensitivities of the forward solution.
- `ypS` – a pointer to an array of `Ns` vectors containing the derivatives of the forward sensitivities.
- `yB` – is the current value of the backward dependent variable vector.
- `ypB` – is the current value of the backward dependent derivative vector.
- `resvalB` – is the current value of the residual for the backward problem.
- `cjB` – is the scalar in the system Jacobian, proportional to the inverse of the step size (  $\alpha$  in (2.7) ).
- `user_dataB` – is a pointer to user data — the same as the `user_dataB` parameter passed to `IDASetUserDataB`.

**Return value:**

The value returned by the Jacobian-vector setup function should be if successful, positive for a recoverable error (in which case the step will be retried), or negative for an unrecoverable error (in which case the integration is halted).

**Notes:**

Each call to the Jacobian-vector setup function is preceded by a call to the backward problem residual user function with the same (*t*, *y*, *yp*, *yyS*, *ypS*, *yB*, *ypB*) arguments. Thus, the setup function can use any auxiliary data that is computed and saved during the evaluation of the DAE residual. If the user's `IDALsJacTimesVecFnB` function uses difference quotient approximations, it may need to access quantities not in the call list. These include the current stepsize, the error weights, etc. To obtain these, the user will need to add a pointer to `ida_mem` to `user_dataB` and then use the `IDAGet*` functions described in §5.5.2.11. The unit roundoff can be accessed as `SUN_UNIT_ROUNDOFF` defined in `sundials_types.h`. The previous function type `IDASpilsJacTimesSetupFnBS` is deprecated.

New in version 3.0.0: Replaces the deprecated type `IDASpilsJacTimesSetupFnBS`.

**5.5.3.8 Preconditioner solve for the backward problem (iterative linear solvers)**

If preconditioning is used during integration of the backward problem, then the user must provide a function to solve the linear system  $Pz = r$ , where  $P$  is a left preconditioner matrix. This function must have one of the following two forms:

```
typedef int (*IDALsPrecSolveFnB)(sunrealtype t, N_Vector yy, N_Vector yp, N_Vector yB, N_Vector ypB,  
N_Vector resvalB, N_Vector rvecB, N_Vector zvecB, sunrealtype cJB, sunrealtype deltaB, void *user_dataB)
```

This function solves the preconditioning system  $Pz = r$  for the backward problem.

**Arguments:**

- *t* – is the current value of the independent variable.
- *yy* – is the current value of the forward solution vector.
- *yp* – is the current value of the forward solution derivative vector.
- *yB* – is the current value of the backward dependent variable vector.
- *ypB* – is the current value of the backward dependent derivative vector.
- *resvalB* – is the current value of the residual for the backward problem.
- *rvecB* – is the right-hand side vector  $r$  of the linear system to be solved.
- *zvecB* – is the computed output vector.
- *cJB* – is the scalar in the system Jacobian, proportional to the inverse of the step size ( $\alpha$  in (2.7)).
- *deltaB* – is an input tolerance to be used if an iterative method is employed in the solution.
- *user\_dataB* – is a pointer to user data — the same as the `user_dataB` parameter passed to the function `IDASetUserDataB`.

**Return value:**

The return value of a preconditioner solve function for the backward problem should be if successful, positive for a recoverable error (in which case the step will be retried), or negative for an unrecoverable error (in which case the integration is halted).

New in version 3.0.0: Replaces the deprecated type `IDASpilsPrecSolveFnB`.

```
typedef int (*IDALsPrecSolveFnBS)(sunrealtype t, N_Vector yy, N_Vector yp, N_Vector *yyS, N_Vector *ypS,  
N_Vector yB, N_Vector ypB, N_Vector resvalB, N_Vector rvecB, N_Vector zvecB, sunrealtype cJB, sunrealtype  
deltaB, void *user_dataB)
```

This function solves the preconditioning system  $Pz = r$  for the backward problem, for the case in which the backward problem depends on the forward sensitivities.

**Arguments:**

- `t` – is the current value of the independent variable.
- `yy` – is the current value of the forward solution vector.
- `yp` – is the current value of the forward solution derivative vector.
- `yyS` – a pointer to an array of `Ns` vectors containing the sensitivities of the forward solution.
- `ypS` – a pointer to an array of `Ns` vectors containing the derivatives of the forward sensitivities.
- `yB` – is the current value of the backward dependent variable vector.
- `ypB` – is the current value of the backward dependent derivative vector.
- `resvalB` – is the current value of the residual for the backward problem.
- `rvecB` – is the right-hand side vector  $r$  of the linear system to be solved.
- `zvecB` – is the computed output vector.
- `cjB` – is the scalar in the system Jacobian, proportional to the inverse of the step size (  $\alpha$  in (2.7) ).
- `deltaB` – is an input tolerance to be used if an iterative method is employed in the solution.
- `user_dataB` – is a pointer to user data — the same as the `user_dataB` parameter passed to the function `IDASetUserDataB`.

**Return value:**

The return value of a preconditioner solve function for the backward problem should be if successful, positive for a recoverable error (in which case the step will be retried), or negative for an unrecoverable error (in which case the integration is halted).

New in version 3.0.0: Replaces the deprecated type `IDASpilsPrecSolveFnBS`.

**5.5.3.9 Preconditioner setup for the backward problem (iterative linear solvers)**

If the user's preconditioner requires that any Jacobian-related data be preprocessed or evaluated, then this needs to be done in a user-supplied function of one of the following two types:

```
typedef int (*IDALsPrecSetupFnB)(sunrealtype t, N_Vector yy, N_Vector yp, N_Vector yB, N_Vector ypB,
N_Vector resvalB, sunrealtype cjB, void *user_dataB)
```

This function preprocesses and/or evaluates Jacobian-related data needed by the preconditioner for the backward problem.

**Arguments:**

- `t` – is the current value of the independent variable.
- `yy` – is the current value of the forward solution vector.
- `yp` – is the current value of the forward solution vector.
- `yB` – is the current value of the backward dependent variable vector.
- `ypB` – is the current value of the backward dependent derivative vector.
- `resvalB` – is the current value of the residual for the backward problem.
- `cjB` – is the scalar in the system Jacobian, proportional to the inverse of the step size (  $\alpha$  in (2.7) ).
- `user_dataB` – is a pointer to user data — the same as the `user_dataB` parameter passed to the function `IDASetUserDataB`.

**Return value:**

The return value of a preconditioner setup function for the backward problem should be if successful, positive for a recoverable error (in which case the step will be retried), or negative for an unrecoverable error (in which case the integration is halted).

New in version 3.0.0: Replaces the deprecated type IDASpilsPrecSetupFnB.

```
typedef int (*IDALSprecSetupFnBS)(sunrealtype t, N_Vector yy, N_Vector yp, N_Vector *yyS, N_Vector *ypS,  
N_Vector yB, N_Vector ypB, N_Vector resvalB, sunrealtype cjB, void *user_dataB)
```

This function preprocesses and/or evaluates Jacobian-related data needed by the preconditioner for the backward problem, in the case where the backward problem depends on the forward sensitivities.

**Arguments:**

- *t* – is the current value of the independent variable.
- *yy* – is the current value of the forward solution vector.
- *yp* – is the current value of the forward solution vector.
- *yyS* – a pointer to an array of *Ns* vectors containing the sensitivities of the forward solution.
- *ypS* – a pointer to an array of *Ns* vectors containing the derivatives of the forward sensitivities.
- *yB* – is the current value of the backward dependent variable vector.
- *ypB* – is the current value of the backward dependent derivative vector.
- *resvalB* – is the current value of the residual for the backward problem.
- *cjB* – is the scalar in the system Jacobian, proportional to the inverse of the step size (  $\alpha$  in (2.7) ).
- *user\_dataB* – is a pointer to user data — the same as the *user\_dataB* parameter passed to the function IDASetUserDataB.

**Return value:**

The return value of a preconditioner setup function for the backward problem should be if successful, positive for a recoverable error (in which case the step will be retried), or negative for an unrecoverable error (in which case the integration is halted).

New in version 3.0.0: Replaces the deprecated type IDASpilsPrecSetupFnBS.

## 5.5.4 Using the band-block-diagonal preconditioner for backward problems

As on the forward integration phase, the efficiency of Krylov iterative methods for the solution of linear systems can be greatly enhanced through preconditioning. The band-block-diagonal preconditioner module IDABBDPRE, provides interface functions through which it can be used on the backward integration phase.

The adjoint module in IDAS offers an interface to the band-block-diagonal preconditioner module IDABBDPRE described in section §5.3.1. This generates a preconditioner that is a block-diagonal matrix with each block being a band matrix and can be used with one of the Krylov linear solvers and with the MPI-parallel vector module NVECTOR\_PARALLEL.

In order to use the IDABBDPRE module in the solution of the backward problem, the user must define one or two additional functions, described at the end of this section.

#### 5.5.4.1 Usage of IDABBDPRE for the backward problem

The IDABBDPRE module is initialized by calling the following function, *after* an iterative linear solver for the backward problem has been attached to IDAS by calling `IDASetLinearSolverB()` (see §5.5.2.6).

```
int IDABBDPrecInitB(void *ida_mem, int which, sunindextype NlocalB, sunindextype mudqB, sunindextype
                    mldqB, sunindextype mukeepB, sunindextype mlkeepB, sunrealtype dqrelyB,
                    IDABBDLocalFnB GresB, IDABBDCommFnB GcommB)
```

The function `IDABBDPrecInitB()` initializes and allocates memory for the IDABBDPRE preconditioner for the backward problem.

##### Arguments:

- `ida_mem` – pointer to the IDAS memory block.
- `which` – the identifier of the backward problem.
- `NlocalB` – local vector dimension for the backward problem.
- `mudqB` – upper half-bandwidth to be used in the difference-quotient Jacobian approximation.
- `mldqB` – lower half-bandwidth to be used in the difference-quotient Jacobian approximation.
- `mukeepB` – upper half-bandwidth of the retained banded approximate Jacobian block.
- `mlkeepB` – lower half-bandwidth of the retained banded approximate Jacobian block.
- `dqrelyB` – the relative increment in components of  $y_B$  used in the difference quotient approximations. The default is  $dqrelyB = \sqrt{\text{unit roundoff}}$ , which can be specified by passing `dqrely = 0.0`.
- `GresB` – the C function which computes  $G_B(t, y, \dot{y}, y_B, \dot{y}_B)$ , the function approximating the residual of the backward problem.
- `GcommB` – the optional C function which performs all interprocess communication required for the computation of  $G_B$ .

##### Return value:

- `IDALS_SUCCESS` – The call to `IDABBDPrecInitB()` was successful.
- `IDALS_MEM_FAIL` – A memory allocation request has failed.
- `IDALS_MEM_NULL` – The `ida_mem` argument was NULL.
- `IDALS_LMEM_NULL` – No linear solver has been attached.
- `IDALS_ILL_INPUT` – An invalid parameter has been passed.

To reinitialize the IDABBDPRE preconditioner module for the backward problem, possibly with a change in `mudqB`, `mldqB`, or `dqrelyB`, call the following function:

```
int IDABBDPrecReInitB(void *ida_mem, int which, sunindextype mudqB, sunindextype mldqB, sunrealtype
                    dqrelyB)
```

The function `IDABBDPrecReInitB()` reinitializes the IDABBDPRE preconditioner for the backward problem.

##### Arguments:

- `ida_mem` – pointer to the IDAS memory block returned by `IDACreate()`.
- `which` – the identifier of the backward problem.
- `mudqB` – upper half-bandwidth to be used in the difference-quotient Jacobian approximation.
- `mldqB` – lower half-bandwidth to be used in the difference-quotient Jacobian approximation.
- `dqrelyB` – the relative increment in components of  $y_B$  used in the difference quotient approximations.



**Return value:**

- IDALS\_SUCCESS – The call to `IDABBDPrecReInitB()` was successful.
- IDALS\_MEM\_FAIL – A memory allocation request has failed.
- IDALS\_MEM\_NULL – The `ida_mem` argument was NULL.
- IDALS\_PMEM\_NULL – The `IDABBDPrecInitB()` has not been previously called.
- IDALS\_LMEM\_NULL – No linear solver has been attached.
- IDALS\_ILL\_INPUT – An invalid parameter has been passed.

**5.5.4.2 User-supplied functions for IDABBDPRE**

To use the IDABBDPRE module, the user must supply one or two functions which the module calls to construct the preconditioner: a required function `GresB` (of type `IDABBDLocalFnB`) which approximates the residual of the backward problem and which is computed locally, and an optional function `GcommB` (of type `IDABBDCommFnB`) which performs all interprocess communication necessary to evaluate this approximate residual (see §5.3.1). The prototypes for these two functions are described below.

```
typedef int (*IDABBDLocalFnB)(sunindextype NlocalB, sunrealtype t, N_Vector y, N_Vector yp, N_Vector yB,  
N_Vector ypB, N_Vector gB, void *user_dataB)
```

This `GresB` function loads the vector `gB`, an approximation to the residual of the backward problem, as a function of `t`, `y`, `yp`, and `yB` and `ypB`.

**Arguments:**

- `NlocalB` – is the local vector length for the backward problem.
- `t` – is the value of the independent variable.
- `y` – is the current value of the forward solution vector.
- `yp` – is the current value of the forward solution derivative vector.
- `yB` – is the current value of the backward dependent variable vector.
- `ypB` – is the current value of the backward dependent derivative vector.
- `gB` – is the output vector,  $G_B(t, y, \dot{y}, y_B, \dot{y}_B)$ .
- `user_dataB` – is a pointer to user data — the same as the `user_dataB` parameter passed to `IDASectUserDataB`.

**Return value:**

An `IDABBDLocalFnB` should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and `IDASolveB()` returns `IDA_LSETUP_FAIL`).

**Notes:**

This routine must assume that all interprocess communication of data needed to calculate `gB` has already been done, and this data is accessible within `user_dataB`.

**Warning:** Before calling the user's `IDABBDLocalFnB`, IDAS needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, IDAS triggers an unrecoverable failure in the preconditioner setup function which will halt the integration (`IDASolveB()` returns `IDA_LSETUP_FAIL`).



```
typedef int (*IDABBDCommFnB)(sunindextype NlocalB, sunrealtype t, N_Vector y, N_Vector yp, N_Vector yB,
N_Vector ypB, void *user_dataB)
```

This GcommB function performs all interprocess communications necessary for the execution of the GresB function above, using the input vectors y, yp, yB and ypB.

**Arguments:**

- NlocalB – is the local vector length.
- t – is the value of the independent variable.
- y – is the current value of the forward solution vector.
- yp – is the current value of the forward solution derivative vector.
- yB – is the current value of the backward dependent variable vector.
- ypB – is the current value of the backward dependent derivative vector.
- user\_dataB – is a pointer to user data — the same as the user\_dataB parameter passed to IDASetsUserDataB.

**Return value:**

An IDABBDCommFnB should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and *IDASolveB()* returns IDA\_LSETUP\_FAIL).

**Notes:**

The GcommB function is expected to save communicated data in space defined within the structure user\_dataB.

Each call to the GcommB function is preceded by a call to the function that evaluates the residual of the backward problem with the same t, y, yp, yB and ypB arguments. If there is no additional communication needed, then pass GcommB = NULL to *IDABBDPrecInitB()*.



## Chapter 6

# Vector Data Structures

The SUNDIALS library comes packaged with a variety of NVECTOR implementations, designed for simulations in serial, shared-memory parallel, and distributed-memory parallel environments, as well as interfaces to vector data structures used within external linear solver libraries. All native implementations assume that the process-local data is stored contiguously, and they in turn provide a variety of standard vector algebra operations that may be performed on the data.

In addition, SUNDIALS provides a simple interface for generic vectors (akin to a C++ *abstract base class*). All of the major SUNDIALS solvers (CVODE(s), IDA(s), KINSOL, ARKODE) in turn are constructed to only depend on these generic vector operations, making them immediately extensible to new user-defined vector objects. The only exceptions to this rule relate to the dense, banded and sparse-direct linear system solvers, since they rely on particular data storage and access patterns in the NVECTORS used.

### 6.1 Description of the NVECTOR Modules

SUNDIALS solvers are written in a data-independent manner. They all operate on generic vectors (of type *N\_Vector*) through a set of operations defined by, and specific to, the particular vector implementation. Users can provide a custom vector implementation or use one provided with SUNDIALS. The generic operations are described below. In the sections following, the implementations provided with SUNDIALS are described.

An *N\_Vector* is a pointer to the *\_generic\_N\_Vector* structure:

```
typedef struct _generic_N_Vector *N_Vector
```

```
struct _generic_N_Vector
```

The structure defining the SUNDIALS vector class.

void \***content**

Pointer to vector-specific member data.

*N\_Vector\_Ops* **ops**

A virtual table of vector operations provided by a specific implementation.

*SUNContext* **sunctx**

The SUNDIALS simulation context

The virtual table structure is defined as

```
typedef _generic_N_Vector_Ops *N_Vector_Ops
```

**struct \_generic\_N\_Vector\_Ops**

The structure defining *N\_Vector* operations.

*N\_Vector\_ID* (\***nvgetvectorid**)(*N\_Vector*)

The function implementing *N\_VGetVectorID()*

*N\_Vector* (\***nvclone**)(*N\_Vector*)

The function implementing *N\_VClone()*

*N\_Vector* (\***nvcloneempty**)(*N\_Vector*)

The function implementing *N\_VCloneEmpty()*

void (\***nvdestroy**)(*N\_Vector*)

The function implementing *N\_VDestroy()*

void (\***nvspace**)(*N\_Vector*, *sunindextype\**, *sunindextype\**)

The function implementing *N\_VSpace()*

*sunrealtype* \*(\***nvgetarraypointer**)(*N\_Vector*)

The function implementing *N\_VGetArrayPointer()*

*sunrealtype* \*(\***nvgetdevicearraypointer**)(*N\_Vector*)

The function implementing *N\_VGetDeviceArrayPointer()*

void (\***nvsetarraypointer**)(*sunrealtype\**, *N\_Vector*)

The function implementing *N\_VSetArrayPointer()*

*SUNComm* (\***nvgetcommunicator**)(*N\_Vector*)

The function implementing *N\_VGetCommunicator()*

*sunindextype* (\***nvgetlength**)(*N\_Vector*)

The function implementing *N\_VGetLength()*

*sunindextype* (\***nvgetlocallength**)(*N\_Vector*)

The function implementing *N\_VGetLocalLength()*

void (\***nvlinearsum**)(*sunrealtype*, *N\_Vector*, *sunrealtype*, *N\_Vector*, *N\_Vector*)

The function implementing *N\_VLinearSum()*

void (\***nvconst**)(*sunrealtype*, *N\_Vector*)

The function implementing *N\_VConst()*

void (\***nvprod**)(*N\_Vector*, *N\_Vector*, *N\_Vector*)

The function implementing *N\_VProd()*

void (\***nvdiv**)(*N\_Vector*, *N\_Vector*, *N\_Vector*)

The function implementing *N\_VDiv()*

void (\***nvscale**)(*sunrealtype*, *N\_Vector*, *N\_Vector*)

The function implementing *N\_VScale()*

void (\***nvabs**)(*N\_Vector*, *N\_Vector*)

The function implementing *N\_VAbs()*

void (\***nvinv**)(*N\_Vector*, *N\_Vector*)

The function implementing *N\_VInv()*

`void (*nvaddconst)(N_Vector, sunrealtype, N_Vector)`  
 The function implementing `N_VAddConst()`

`sunrealtype (*nvdotprod)(N_Vector, N_Vector)`  
 The function implementing `N_VDotProd()`

`sunrealtype (*nvmaxnorm)(N_Vector)`  
 The function implementing `N_VMaxNorm()`

`sunrealtype (*nvwrmsnorm)(N_Vector, N_Vector)`  
 The function implementing `N_VWrmsNorm()`

`sunrealtype (*nvwrmsnormmask)(N_Vector, N_Vector, N_Vector)`  
 The function implementing `N_VWrmsNormMask()`

`sunrealtype (*nvmin)(N_Vector)`  
 The function implementing `N_VMin()`

`sunrealtype (*nvwl2norm)(N_Vector, N_Vector)`  
 The function implementing `N_VWL2Norm()`

`sunrealtype (*nv1lnorm)(N_Vector)`  
 The function implementing `N_VL1Norm()`

`void (*nvcompare)(sunrealtype, N_Vector, N_Vector)`  
 The function implementing `N_VCompare()`

`sunboolean_t (*nvinvtest)(N_Vector, N_Vector)`  
 The function implementing `N_VInvTest()`

`sunboolean_t (*nvconstrmask)(N_Vector, N_Vector, N_Vector)`  
 The function implementing `N_VConstrMask()`

`sunrealtype (*nvminquotient)(N_Vector, N_Vector)`  
 The function implementing `N_VMinQuotient()`

`SUNErrCode (*nvlinearcombination)(int, sunrealtype*, N_Vector*, N_Vector)`  
 The function implementing `N_VLinearCombination()`

`SUNErrCode (*nvscaleaddmulti)(int, sunrealtype*, N_Vector, N_Vector*, N_Vector*)`  
 The function implementing `N_VScaleAddMulti()`

`SUNErrCode (*nvdotprodmulti)(int, N_Vector, N_Vector*, sunrealtype*)`  
 The function implementing `N_VDotProdMulti()`

`SUNErrCode (*nvlinearsumvectorarray)(int, sunrealtype, N_Vector*, sunrealtype, N_Vector*, N_Vector*)`  
 The function implementing `N_VLinearSumVectorArray()`

`SUNErrCode (*nvscalevectorarray)(int, sunrealtype*, N_Vector*, N_Vector*)`  
 The function implementing `N_VScaleVectorArray()`

`SUNErrCode (*nvconstvectorarray)(int, sunrealtype, N_Vector*)`  
 The function implementing `N_VConstVectorArray()`

`SUNErrCode (*nvwrmsnormvectorarray)(int, N_Vector*, N_Vector*, sunrealtype*)`  
 The function implementing `N_VWrmsNormVectorArray()`

*SUNErrorCode* (\***nvwrmsnormmaskvectorarray**)(int, *N\_Vector*\*, *N\_Vector*\*, *N\_Vector*, *sunrealtype*\*)

The function implementing *N\_VWrmsNormMaskVectorArray()*

*SUNErrorCode* (\***nvscaleaddmultivectorarray**)(int, int, *sunrealtype*\*, *N\_Vector*\*, *N\_Vector*\*\*, *N\_Vector*\*\*)

The function implementing *N\_VScaleAddMultiVectorArray()*

*SUNErrorCode* (\***nvlinearcombinationvectorarray**)(int, int, *sunrealtype*\*, *N\_Vector*\*\*, *N\_Vector*\*)

The function implementing *N\_VLinearCombinationVectorArray()*

*sunrealtype* (\***nvdotprodlocal**)(*N\_Vector*, *N\_Vector*)

The function implementing *N\_VDotProdLocal()*

*sunrealtype* (\***nvmaxnormlocal**)(*N\_Vector*)

The function implementing *N\_VMaxNormLocal()*

*sunrealtype* (\***nvminlocal**)(*N\_Vector*)

The function implementing *N\_VMinLocal()*

*sunrealtype* (\***nv1lnormlocal**)(*N\_Vector*)

The function implementing *N\_VL1NormLocal()*

*sunboolean* (\***nvinvtestlocal**)(*N\_Vector*, *N\_Vector*)

The function implementing *N\_VInvTestLocal()*

*sunboolean* (\***nvconstrmasklocal**)(*N\_Vector*, *N\_Vector*, *N\_Vector*)

The function implementing *N\_VConstrMaskLocal()*

*sunrealtype* (\***nvminquotientlocal**)(*N\_Vector*, *N\_Vector*)

The function implementing *N\_VMinQuotientLocal()*

*sunrealtype* (\***nvwsqrsumlocal**)(*N\_Vector*, *N\_Vector*)

The function implementing *N\_VWSqrSumLocal()*

*sunrealtype* (\***nvwsqrsummasklocal**)(*N\_Vector*, *N\_Vector*, *N\_Vector*)

The function implementing *N\_VWSqrSumMaskLocal()*

*SUNErrorCode* (\***nvdotprodmultilocal**)(int, *N\_Vector*, *N\_Vector*\*, *sunrealtype*\*)

The function implementing *N\_VDotProdMultiLocal()*

*SUNErrorCode* (\***nvdotprodmultiallreduce**)(int, *N\_Vector*, *sunrealtype*\*)

The function implementing *N\_VDotProdMultiAllReduce()*

*SUNErrorCode* (\***nvbufsize**)(*N\_Vector*, *sunindex*\*)

The function implementing *N\_VBufSize()*

*SUNErrorCode* (\***nvbufpack**)(*N\_Vector*, void\*)

The function implementing *N\_VBufPack()*

*SUNErrorCode* (\***nvbufunpack**)(*N\_Vector*, void\*)

The function implementing *N\_VBufUnpack()*

void (\***nvprint**)(*N\_Vector*)

The function implementing *N\_VPrint()*

```
void (*nvprintf)(N_Vector, FILE*)
```

The function implementing `N_VPrintfFile()`

The generic NVECTOR module defines and implements the vector operations acting on a `N_Vector`. These routines are nothing but wrappers for the vector operations defined by a particular NVECTOR implementation, which are accessed through the `ops` field of the `N_Vector` structure. To illustrate this point we show below the implementation of a typical vector operation from the generic NVECTOR module, namely `N_VScale`, which performs the operation  $z \leftarrow cx$  for vectors  $x$  and  $z$  and a scalar  $c$ :

```
void N_VScale(sunrealtype c, N_Vector x, N_Vector z) {
    z->ops->nvscale(c, x, z);
}
```

§6.2 contains a complete list of all standard vector operations defined by the generic NVECTOR module. §6.2.2, §6.2.3, §6.2.4, §6.2.5, and §6.2.6 list *optional* fused, vector array, local reduction, single buffer reduction, and exchange operations, respectively.

Fused and vector array operations (see §6.2.2 and §6.2.3) are intended to increase data reuse, reduce parallel communication on distributed memory systems, and lower the number of kernel launches on systems with accelerators. If a particular NVECTOR implementation defines a fused or vector array operation as `NULL`, the generic NVECTOR module will automatically call standard vector operations as necessary to complete the desired operation. In all SUNDIALS-provided NVECTOR implementations, all fused and vector array operations are disabled by default. However, these implementations provide additional user-callable functions to enable/disable any or all of the fused and vector array operations. See the following sections for the implementation specific functions to enable/disable operations.

Local reduction operations (see §6.2.4) are similarly intended to reduce parallel communication on distributed memory systems, particularly when NVECTOR objects are combined together within an `NVECTOR_MANYVECTOR` object (see §6.17). If a particular NVECTOR implementation defines a local reduction operation as `NULL`, the `NVECTOR_MANYVECTOR` module will automatically call standard vector reduction operations as necessary to complete the desired operation. All SUNDIALS-provided NVECTOR implementations include these local reduction operations, which may be used as templates for user-defined implementations.

The single buffer reduction operations (§6.2.5) are used in low-synchronization methods to combine separate reductions into one `MPI_Allreduce` call.

The exchange operations (see §6.2.6) are intended only for use with the XBraid library for parallel-in-time integration (accessible from ARKODE) and are otherwise unused by SUNDIALS packages.

### 6.1.1 NVECTOR Utility Functions

The generic NVECTOR module also defines several utility functions to aid in creation and management of arrays of `N_Vector` objects – these functions are particularly useful for Fortran users to utilize the `NVECTOR_MANYVECTOR` or SUNDIALS’ sensitivity-enabled packages CVODES and IDAS.

The functions `N_VCloneVectorArray()` and `N_VCloneVectorArrayEmpty()` create (by cloning) an array of *count* variables of type `N_Vector`, each of the same type as an existing `N_Vector` input:

```
N_Vector *N_VCloneVectorArray(int count, N_Vector w)
```

Clones an array of *count* `N_Vector` objects, allocating their data arrays (similar to `N_VClone()`).

#### Arguments:

- *count* – number of `N_Vector` objects to create.
- *w* – template `N_Vector` to clone.

#### Return value:

- pointer to a new `N_Vector` array on success.

- NULL pointer on failure.

*N\_Vector* \***N\_VCloneVectorArrayEmpty**(int count, *N\_Vector* w)

Clones an array of count *N\_Vector* objects, leaving their data arrays unallocated (similar to *N\_VCloneEmpty()*).

**Arguments:**

- count – number of *N\_Vector* objects to create.
- w – template *N\_Vector* to clone.

**Return value:**

- pointer to a new *N\_Vector* array on success.
- NULL pointer on failure.

An array of variables of type *N\_Vector* can be destroyed by calling *N\_VDestroyVectorArray()*:

void **N\_VDestroyVectorArray**(*N\_Vector* \*vs, int count)

Destroys an array of count *N\_Vector* objects.

**Arguments:**

- vs – *N\_Vector* array to destroy.
- count – number of *N\_Vector* objects in vs array.

**Notes:**

This routine will internally call the *N\_Vector* implementation-specific *N\_VDestroy()* operation.

If vs was allocated using *N\_VCloneVectorArray()* then the data arrays for each *N\_Vector* object will be freed; if vs was allocated using *N\_VCloneVectorArrayEmpty()* then it is the user's responsibility to free the data for each *N\_Vector* object.

Finally, we note that users of the Fortran 2003 interface may be interested in the additional utility functions *N\_VNewVectorArray()*, *N\_VGetVecAtIndexVectorArray()*, and *N\_VSetVecAtIndexVectorArray()*, that are wrapped as *FN\_NewVectorArray*, *FN\_VGetVecAtIndexVectorArray*, and *FN\_VSetVecAtIndexVectorArray*, respectively. These functions allow a Fortran 2003 user to create an empty vector array, access a vector from this array, and set a vector within this array:

*N\_Vector* \***N\_VNewVectorArray**(int count, *SUNContext* sunctx)

Creates an array of count *N\_Vector* objects, the pointers to each are initialized as NULL.

**Arguments:**

- count – length of desired *N\_Vector* array.
- sunctx – a *SUNContext* object

**Return value:**

- pointer to a new *N\_Vector* array on success.
- NULL pointer on failure.

Changed in version 7.0.0: The function signature was updated to add the *SUNContext* argument.

*N\_Vector* \***N\_VGetVecAtIndexVectorArray**(*N\_Vector* \*vs, int index)

Accesses the *N\_Vector* at the location index within the *N\_Vector* array vs.

**Arguments:**

- vs – *N\_Vector* array.



- `index` – desired `N_Vector` to access from within `vs`.

**Return value:**

- pointer to the indexed `N_Vector` on success.
- NULL pointer on failure (`index < 0` or `vs == NULL`).

**Notes:**

This routine does not verify that `index` is within the extent of `vs`, since `vs` is a simple `N_Vector` array that does not internally store its allocated length.

void **N\_VSetVecAtIndexVectorArray**(*N\_Vector* \*vs, int index, *N\_Vector* w)

Sets a pointer to `w` at the location `index` within the vector array `vs`.

**Arguments:**

- `vs` – `N_Vector` array.
- `index` – desired location to place the pointer to `w` within `vs`.
- `w` – `N_Vector` to set within `vs`.

**Notes:**

This routine does not verify that `index` is within the extent of `vs`, since `vs` is a simple `N_Vector` array that does not internally store its allocated length.

## 6.1.2 Implementing a custom NVECTOR

A particular implementation of the NVECTOR module must:

- Specify the *content* field of the `N_Vector` structure.
- Define and implement the vector operations. Note that the names of these routines should be unique to that implementation in order to permit using more than one NVECTOR module (each with different `N_Vector` internal data representations) in the same code.
- Define and implement user-callable constructor and destructor routines to create and free an `N_Vector` with the new *content* field and with *ops* pointing to the new vector operations.
- Optionally, define and implement additional user-callable routines acting on the newly-defined `N_Vector` (e.g., a routine to print the content for debugging purposes).
- Optionally, provide accessor macros as needed for that particular implementation to be used to access different parts in the *content* field of the newly-defined `N_Vector`.

To aid in the creation of custom NVECTOR modules, the generic NVECTOR module provides two utility functions *N\_VNewEmpty()* and *N\_VCopyOps()*. When used in custom NVECTOR constructors and clone routines these functions will ease the introduction of any new optional vector operations to the NVECTOR API by ensuring that only required operations need to be set, and that all operations are copied when cloning a vector.

*N\_Vector* **N\_VNewEmpty**(*SUNContext* sunctx)

This allocates a new generic `N_Vector` object and initializes its content pointer and the function pointers in the operations structure to NULL.

**Return value:** If successful, this function returns an `N_Vector` object. If an error occurs when allocating the object, then this routine will return NULL.

void **N\_VFreeEmpty**(*N\_Vector* v)

This routine frees the generic `N_Vector` object, under the assumption that any implementation-specific data that was allocated within the underlying content structure has already been freed. It will additionally test whether the *ops* pointer is NULL, and, if it is not, it will free it as well.

**Arguments:**

- $v$  – an `N_Vector` object

*SUNErrCode* **N\_VCopyOps**(*N\_Vector*  $w$ , *N\_Vector*  $v$ )

This function copies the function pointers in the ops structure of  $w$  into the ops structure of  $v$ .

**Arguments:**

- $w$  – the vector to copy operations from
- $v$  – the vector to copy operations to

**Return value:** Returns a *SUNErrCode*.

enum **N\_Vector\_ID**

Each *N\_Vector* implementation included in SUNDIALS has a unique identifier specified in enumeration and shown in Table 6.1. It is recommended that a user supplied NVECTOR implementation use the SUNDIALS\_NVEC\_CUSTOM identifier.

Table 6.1: Vector Identifications associated with vector kernels supplied with SUNDIALS

Vector ID	Vector type	ID Value
SUNDIALS_NVEC_SERIAL	Serial	0
SUNDIALS_NVEC_PARALLEL	Distributed memory parallel (MPI)	1
SUNDIALS_NVEC_OPENMP	OpenMP shared memory parallel	2
SUNDIALS_NVEC_PTHREADS	PThreads shared memory parallel	3
SUNDIALS_NVEC_PARHYP	<i>hypre</i> ParHyp parallel vector	4
SUNDIALS_NVEC_PETSC	PETSc parallel vector	5
SUNDIALS_NVEC_CUDA	CUDA vector	6
SUNDIALS_NVEC_HIP	HIP vector	7
SUNDIALS_NVEC_SYCL	SYCL vector	8
SUNDIALS_NVEC_RAJA	RAJA vector	9
SUNDIALS_NVEC_OPENMPDEV	OpenMP vector with device offloading	10
SUNDIALS_NVEC_TRILINOS	Trilinos Tpetra vector	11
SUNDIALS_NVEC_MANYVECTOR	“ManyVector” vector	12
SUNDIALS_NVEC_MPIMANYVECTOR	MPI-enabled “ManyVector” vector	13
SUNDIALS_NVEC_MPIPLUSX	MPI+X vector	14
SUNDIALS_NVEC_CUSTOM	User-provided custom vector	15

### 6.1.3 Support for complex-valued vectors

While SUNDIALS itself is written under an assumption of real-valued data, it does provide limited support for complex-valued problems. However, since none of the built-in NVECTOR modules supports complex-valued data, users must provide a custom NVECTOR implementation for this task. Many of the NVECTOR routines described in the subsection §6.2 naturally extend to complex-valued vectors; however, some do not. To this end, we provide the following guidance:

- *N\_VMin()* and *N\_VMinLocal()* should return the minimum of all *real* components of the vector, i.e.,  $m = \min_{0 \leq i < n} \text{real}(x_i)$ .
- *N\_VConst()* (and similarly *N\_VConstVectorArray()*) should set the real components of the vector to the input constant, and set all imaginary components to zero, i.e.,  $z_i = c + 0j$  for  $0 \leq i < n$ .
- *N\_VAddConst()* should only update the real components of the vector with the input constant, leaving all imaginary components unchanged.

- `N_VWrmsNorm()`, `N_VWrmsNormMask()`, `N_VWSqrSumLocal()` and `N_VWSqrSumMaskLocal()` should assume that all entries of the weight vector `w` and the mask vector `id` are real-valued.
- `N_VDotProd()` should mathematically return a complex number for complex-valued vectors; as this is not possible with SUNDIALS' current `sunrealtype`, this routine should be set to NULL in the custom NVECTOR implementation.
- `N_VCompare()`, `N_VConstrMask()`, `N_VMinQuotient()`, `N_VConstrMaskLocal()` and `N_VMinQuotientLocal()` are ill-defined due to the lack of a clear ordering in the complex plane. These routines should be set to NULL in the custom NVECTOR implementation.

While many SUNDIALS solver modules may be utilized on complex-valued data, others cannot. Specifically, although each package's linear solver interface (e.g., ARKLS or CVLS) may be used on complex-valued problems, none of the built-in SUNMatrix or SUNLinearSolver modules will work (all of the direct linear solvers must store complex-valued data, and all of the iterative linear solvers require `N_VDotProd()`). Hence a complex-valued user must provide custom linear solver modules for their problem. At a minimum this will consist of a custom SUNLinearSolver implementation (see §8.1.8), and optionally a custom SUNMatrix as well. The user should then attach these modules as normal to the package's linear solver interface.

Similarly, although both the `SUNNonlinearSolver_Newton` and `SUNNonlinearSolver_FixedPoint` modules may be used with any of the IVP solvers (CVODE(S), IDA(S) and ARKODE) for complex-valued problems, the Anderson-acceleration option with `SUNNonlinearSolver_FixedPoint` cannot be used due to its reliance on `N_VDotProd()`. By this same logic, the Anderson acceleration feature within KINSOL will also not work with complex-valued vectors.

Finally, constraint-handling features of each package cannot be used for complex-valued data, due to the issue of ordering in the complex plane discussed above with `N_VCompare()`, `N_VConstrMask()`, `N_VMinQuotient()`, `N_VConstrMaskLocal()` and `N_VMinQuotientLocal()`.

We provide a simple example of a complex-valued example problem, including a custom complex-valued Fortran 2003 NVECTOR module, in the files `examples/arkode/F2003_custom/ark_analytic_complex_f2003.f90`, `examples/arkode/F2003_custom/fnvector_complex_mod.f90`, and `examples/arkode/F2003_custom/test_fnvector_complex_mod.f90`.

## 6.2 Description of the NVECTOR operations

### 6.2.1 Standard vector operations

The standard vector operations defined by the generic `N_Vector` module are defined as follows. For each of these operations, we give the name, usage of the function, and a description of its mathematical operations below.

`N_Vector` **N\_VGetVectorID**(`N_Vector` `w`)

Returns the vector type identifier for the vector `w`. It is used to determine the vector implementation type (e.g. serial, parallel, ...) from the abstract `N_Vector` interface. Returned values are given in Table 6.1.

Usage:

```
id = N_VGetVectorID(w);
```

`N_Vector` **N\_VClone**(`N_Vector` `w`)

Creates a new `N_Vector` of the same type as an existing vector `w` and sets the `ops` field. It does not copy the vector, but rather allocates storage for the new vector.

Usage:

```
v = N_VClone(w);
```

***N\_Vector* N\_VCloneEmpty(*N\_Vector* w)**

Creates a new *N\_Vector* of the same type as an existing vector *w* and sets the *ops* field. It does not allocate storage for the new vector's data.

Usage:

```
v = N_VCloneEmpty(w);
```

**void N\_VDestroy(*N\_Vector* v)**

Destroys the *N\_Vector* *v* and frees memory allocated for its internal data.

Usage:

```
N_VDestroy(v);
```

**void N\_VSpace(*N\_Vector* v, *sunindextype* \*lrw, *sunindextype* \*liw)**

Returns storage requirements for the *N\_Vector* *v*:

- *lrw* contains the number of *sunrealtype* words
- *liw* contains the number of integer words.

This function is advisory only, for use in determining a user's total space requirements; it could be a dummy function in a user-supplied NVECTOR module if that information is not of interest.

Usage:

```
N_VSpace(nvSpec, &lrw, &liw);
```

***sunrealtype* \*N\_VGetArrayPointer(*N\_Vector* v)**

Returns a pointer to a *sunrealtype* array from the *N\_Vector* *v*. Note that this assumes that the internal data in the *N\_Vector* is a contiguous array of *sunrealtype* and is accesible from the CPU.

This routine is only used in the solver-specific interfaces to the dense and banded (serial) linear solvers, and in the interfaces to the banded (serial) and band-block-diagonal (parallel) preconditioner modules provided with SUNDIALS.

Usage:

```
vdata = N_VGetArrayPointer(v);
```

***sunrealtype* \*N\_VGetDeviceArrayPointer(*N\_Vector* v)**

Returns a device pointer to a *sunrealtype* array from the *N\_Vector* *v*. Note that this assumes that the internal data in *N\_Vector* is a contiguous array of *sunrealtype* and is accessible from the device (e.g., GPU).

This operation is *optional* except when using the GPU-enabled direct linear solvers.

Usage:

```
vdata = N_VGetArrayPointer(v);
```

**void N\_VSetArrayPointer(*sunrealtype* \*vdata, *N\_Vector* v)**

Replaces the data array pointer in an *N\_Vector* with a given array of *sunrealtype*. Note that this assumes that the internal data in the *N\_Vector* is a contiguous array of *sunrealtype*. This routine is only used in the interfaces to the dense (serial) linear solver, hence need not exist in a user-supplied NVECTOR module.

Usage:

```
N_VSetArrayPointer(vdata,v);
```

*SUNComm* **N\_VGetCommunicator**(*N\_Vector* v)

Returns the *SUNComm* (which is just an *MPI\_Comm* when SUNDIALS is built with MPI, otherwise it is an *int*) associated with the vector (if applicable). For MPI-unaware vector implementations, this should return *SUN\_COMM\_NULL*.

Usage:

```
MPI_Comm comm = N_VGetCommunicator(v); // Works if MPI is enabled
int comm = N_VGetCommunicator(v);      // Works if MPI is disabled
SUNComm comm = N_VGetCommunicator(v);  // Works with or without MPI
```

*sunindextype* **N\_VGetLength**(*N\_Vector* v)

Returns the global length (number of “active” entries) in the NVECTOR *v*. This value should be cumulative across all processes if the vector is used in a parallel environment. If *v* contains additional storage, e.g., for parallel communication, those entries should not be included.

Usage:

```
global_length = N_VGetLength(v);
```

*sunindextype* **N\_VGetLocalLength**(*N\_Vector* v)

Returns the local length (number of “active” entries) in the NVECTOR *v*. This value should be the length of the array returned by *N\_VGetArrayPointer()* or *N\_VGetDeviceArrayPointer()*.

Usage:

```
local_length = N_VGetLocalLength(v);
```

void **N\_VLinearSum**(*sunrealtype* a, *N\_Vector* x, *sunrealtype* b, *N\_Vector* y, *N\_Vector* z)

Performs the operation  $z = ax + by$ , where *a* and *b* are *sunrealtype* scalars and *x* and *y* are of type *N\_Vector*:

$$z_i = ax_i + by_i, \quad i = 0, \dots, n-1.$$

The output vector *z* can be the same as either of the input vectors (*x* or *y*).

Usage:

```
N_VLinearSum(a, x, b, y, z);
```

void **N\_VConst**(*sunrealtype* c, *N\_Vector* z)

Sets all components of the *N\_Vector* *z* to *sunrealtype* *c*:

$$z_i = c, \quad i = 0, \dots, n-1.$$

Usage:

```
N_VConst(c, z);
```

void **N\_VProd**(*N\_Vector* x, *N\_Vector* y, *N\_Vector* z)

Sets the *N\_Vector* *z* to be the component-wise product of the *N\_Vector* inputs *x* and *y*:

$$z_i = x_i y_i, \quad i = 0, \dots, n-1.$$

Usage:

```
N_VProd(x, y, z);
```

void **N\_VDiv**(*N\_Vector* x, *N\_Vector* y, *N\_Vector* z)

Sets the *N\_Vector*  $z$  to be the component-wise ratio of the *N\_Vector* inputs  $x$  and  $y$ :

$$z_i = \frac{x_i}{y_i}, \quad i = 0, \dots, n-1.$$

The  $y_i$  may not be tested for 0 values. It should only be called with a  $y$  that is guaranteed to have all nonzero components.

Usage:

```
N_VDiv(x, y, z);
```

void **N\_VScale**(*sunrealtype* c, *N\_Vector* x, *N\_Vector* z)

Scales the *N\_Vector*  $x$  by the *sunrealtype* scalar  $c$  and returns the result in  $z$ :

$$z_i = cx_i, \quad i = 0, \dots, n-1.$$

Usage:

```
N_VScale(c, x, z);
```

void **N\_VAbs**(*N\_Vector* x, *N\_Vector* z)

Sets the components of the *N\_Vector*  $z$  to be the absolute values of the components of the *N\_Vector*  $x$ :

$$z_i = |x_i|, \quad i = 0, \dots, n-1.$$

Usage:

```
N_VAbs(x, z);
```

void **N\_VInv**(*N\_Vector* x, *N\_Vector* z)

Sets the components of the *N\_Vector*  $z$  to be the inverses of the components of the *N\_Vector*  $x$ :

$$z_i = \frac{1}{x_i}, \quad i = 0, \dots, n-1.$$

This routine may not check for division by 0. It should be called only with an  $x$  which is guaranteed to have all nonzero components.

Usage:

```
N_VInv(x, z);
```

void **N\_VAddConst**(*N\_Vector* x, *sunrealtype* b, *N\_Vector* z)

Adds the *sunrealtype* scalar  $b$  to all components of  $x$  and returns the result in the *N\_Vector*  $z$ :

$$z_i = x_i + b, \quad i = 0, \dots, n-1.$$

Usage:

```
N_VAddConst(x, b, z);
```

*sunrealtype* **N\_VDotProd**(*N\_Vector* x, *N\_Vector* z)

Returns the value of the dot-product of the vectors  $x$  and  $y$ :

$$d = \sum_{i=0}^{n-1} x_i y_i.$$

Usage:

```
d = N_VDotProd(x, y);
```

*sunrealtype* **N\_VMaxNorm**(*N\_Vector* x)

Returns the value of the  $l_\infty$  norm of the *N\_Vector*  $x$ :

$$m = \max_{0 \leq i < n} |x_i|.$$

Usage:

```
m = N_VMaxNorm(x);
```

*sunrealtype* **N\_VWrmsNorm**(*N\_Vector* x, *N\_Vector* w)

Returns the weighted root-mean-square norm of the *N\_Vector*  $x$  with (positive) *sunrealtype* weight vector  $w$ :

$$m = \sqrt{\left( \sum_{i=0}^{n-1} (x_i w_i)^2 \right) / n}$$

Usage:

```
m = N_VWrmsNorm(x, w);
```

*sunrealtype* **N\_VWrmsNormMask**(*N\_Vector* x, *N\_Vector* w, *N\_Vector* id)

Returns the weighted root mean square norm of the *N\_Vector*  $x$  with *sunrealtype* weight vector  $w$  built using only the elements of  $x$  corresponding to positive elements of the *N\_Vector*  $id$ :

$$m = \sqrt{\left( \sum_{i=0}^{n-1} (x_i w_i H(id_i))^2 \right) / n},$$

$$\text{where } H(\alpha) = \begin{cases} 1 & \alpha > 0 \\ 0 & \alpha \leq 0 \end{cases}.$$

Usage:

```
m = N_VWrmsNormMask(x, w, id);
```

*sunrealtype* **N\_VMin**(*N\_Vector* x)

Returns the smallest element of the *N\_Vector*  $x$ :

$$m = \min_{0 \leq i < n} x_i.$$

Usage:

```
m = N_VMin(x);
```

*sunrealtype* **N\_VWL2Norm**(*N\_Vector* x, *N\_Vector* w)Returns the weighted Euclidean  $l_2$  norm of the *N\_Vector*  $x$  with *sunrealtype* weight vector  $w$ :

$$m = \sqrt{\sum_{i=0}^{n-1} (x_i w_i)^2}.$$

Usage:

```
m = N_VWL2Norm(x, w);
```

*sunrealtype* **N\_VL1Norm**(*N\_Vector* x)Returns the  $l_1$  norm of the *N\_Vector*  $x$ :

$$m = \sum_{i=0}^{n-1} |x_i|.$$

Usage:

```
m = N_VL1Norm(x);
```

void **N\_VCompare**(*sunrealtype* c, *N\_Vector* x, *N\_Vector* z)Compares the components of the *N\_Vector*  $x$  to the *sunrealtype* scalar  $c$  and returns an *N\_Vector*  $z$  such that for all  $0 \leq i < n$ ,

$$z_i = \begin{cases} 1.0 & \text{if } |x_i| \geq c, \\ 0.0 & \text{otherwise} \end{cases}.$$

Usage:

```
N_VCompare(c, x, z);
```

*sunbooleantype* **N\_VInvTest**(*N\_Vector* x, *N\_Vector* z)Sets the components of the *N\_Vector*  $z$  to be the inverses of the components of the *N\_Vector*  $x$ , with prior testing for zero values:

$$z_i = \frac{1}{x_i}, \quad i = 0, \dots, n-1.$$

This routine returns a boolean assigned to SUNTRUE if all components of  $x$  are nonzero (successful inversion) and returns SUNFALSE otherwise.

Usage:

```
t = N_VInvTest(x, z);
```

*sunbooleantype* **N\_VConstrMask**(*N\_Vector* c, *N\_Vector* x, *N\_Vector* m)Performs the following constraint tests based on the values in  $c_i$ :

$$\begin{aligned} x_i &> 0 & \text{if } c_i = 2, \\ x_i &\geq 0 & \text{if } c_i = 1, \\ x_i &< 0 & \text{if } c_i = -2, \\ x_i &\leq 0 & \text{if } c_i = -1. \end{aligned}$$

There is no constraint on  $x_i$  if  $c_i = 0$ . This routine returns a boolean assigned to SUNFALSE if any element failed the constraint test and assigned to SUNTRUE if all passed. It also sets a mask vector  $m$ , with elements equal to 1.0 where the constraint test failed, and 0.0 where the test passed. This routine is used only for constraint checking.

Usage:



```
t = N_VConstrMask(c, x, m);
```

*sunrealtype* **N\_VMinQuotient**(*N\_Vector* num, *N\_Vector* denom)

This routine returns the minimum of the quotients obtained by termwise dividing the elements of  $n$  by the elements in  $d$ :

$$\min_{0 \leq i < n} \frac{\text{num}_i}{\text{denom}_i}.$$

A zero element in *denom* will be skipped. If no such quotients are found, then the large value `SUN_BIG_REAL` (defined in the header file `sundials_types.h`) is returned.

Usage:

```
minq = N_VMinQuotient(num, denom);
```

## 6.2.2 Fused operations

The following fused vector operations are *optional*. These operations are intended to increase data reuse, reduce parallel communication on distributed memory systems, and lower the number of kernel launches on systems with accelerators. If a particular NVECTOR implementation defines one of the fused vector operations as `NULL`, the NVECTOR interface will call one of the above standard vector operations as necessary. As above, for each operation, we give the name, usage of the function, and a description of its mathematical operations below.

*SUNErrCode* **N\_VLinearCombination**(int nv, *sunrealtype* \*c, *N\_Vector* \*X, *N\_Vector* z)

This routine computes the linear combination of  $nv$  vectors with  $n$  elements:

$$z_i = \sum_{j=0}^{nv-1} c_j x_{j,i}, \quad i = 0, \dots, n-1,$$

where  $c$  is an array of  $nv$  scalars,  $x_j$  is a vector in the vector array  $X$ , and  $z$  is the output vector. If the output vector  $z$  is one of the vectors in  $X$ , then it *must* be the first vector in the vector array. The operation returns a *SUNErrCode*.

Usage:

```
retval = N_VLinearCombination(nv, c, X, z);
```

*SUNErrCode* **N\_VScaleAddMulti**(int nv, *sunrealtype* \*c, *N\_Vector* x, *N\_Vector* \*Y, *N\_Vector* \*Z)

This routine scales and adds one vector to  $nv$  vectors with  $n$  elements:

$$z_{j,i} = c_j x_i + y_{j,i}, \quad j = 0, \dots, nv-1 \quad i = 0, \dots, n-1,$$

where  $c$  is an array of scalars,  $x$  is a vector,  $y_j$  is a vector in the vector array  $Y$ , and  $z_j$  is an output vector in the vector array  $Z$ . The operation returns a *SUNErrCode*.

Usage:

```
retval = N_VScaleAddMulti(nv, c, x, Y, Z);
```

*SUNErrCode* **N\_VDotProdMulti**(int nv, *N\_Vector* x, *N\_Vector* \*Y, *sunrealtype* \*d)

This routine computes the dot product of a vector with  $nv$  vectors having  $n$  elements:

$$d_j = \sum_{i=0}^{n-1} x_i y_{j,i}, \quad j = 0, \dots, nv-1,$$

where  $d$  is an array of scalars containing the computed dot products,  $x$  is a vector, and  $y_j$  is a vector the vector array  $Y$ . The operation returns a [SUNErrCode](#).

Usage:

```
retval = N_VDotProdMulti(nv, x, Y, d);
```

### 6.2.3 Vector array operations

The following vector array operations are also *optional*. As with the fused vector operations, these are intended to increase data reuse, reduce parallel communication on distributed memory systems, and lower the number of kernel launches on systems with accelerators. If a particular NVECTOR implementation defines one of the fused or vector array operations as NULL, the NVECTOR interface will call one of the above standard vector operations as necessary. As above, for each operation, we give the name, usage of the function, and a description of its mathematical operations below.

[SUNErrCode](#) **N\_VLinearSumVectorArray**(int nv, *sunrealtype* a, *N\_Vector* X, *sunrealtype* b, *N\_Vector* \*Y, *N\_Vector* \*Z)

This routine computes the linear sum of two vector arrays of  $nv$  vectors with  $n$  elements:

$$z_{j,i} = ax_{j,i} + by_{j,i}, \quad i = 0, \dots, n-1 \quad j = 0, \dots, nv-1,$$

where  $a$  and  $b$  are scalars,  $x_j$  and  $y_j$  are vectors in the vector arrays  $X$  and  $Y$  respectively, and  $z_j$  is a vector in the output vector array  $Z$ . The operation returns a [SUNErrCode](#).

Usage:

```
retval = N_VLinearSumVectorArray(nv, a, X, b, Y, Z);
```

[SUNErrCode](#) **N\_VScaleVectorArray**(int nv, *sunrealtype* \*c, *N\_Vector* \*X, *N\_Vector* \*Z)

This routine scales each element in a vector of  $n$  elements in a vector array of  $nv$  vectors by a potentially different constant:

$$z_{j,i} = c_j x_{j,i}, \quad i = 0, \dots, n-1 \quad j = 0, \dots, nv-1,$$

where  $c$  is an array of scalars,  $x_j$  is a vector in the vector array  $X$ , and  $z_j$  is a vector in the output vector array  $Z$ . The operation returns a [SUNErrCode](#).

Usage:

```
retval = N_VScaleVectorArray(nv, c, X, Z);
```

[SUNErrCode](#) **N\_VConstVectorArray**(int nv, *sunrealtype* c, *N\_Vector* \*Z)

This routine sets each element in a vector of  $n$  elements in a vector array of  $nv$  vectors to the same value:

$$z_{j,i} = c, \quad i = 0, \dots, n-1 \quad j = 0, \dots, nv-1,$$

where  $c$  is a scalar and  $z_j$  is a vector in the vector array  $Z$ . The operation returns a [SUNErrCode](#).

Usage:

```
retval = N_VConstVectorArray(nv, c, Z);
```

**SUNErrCode N\_VWrmsNormVectorArray**(int nv, *N\_Vector* \*X, *N\_Vector* \*W, *sunrealtype* \*m)

This routine computes the weighted root mean square norm of each vector in a vector array:

$$m_j = \left( \frac{1}{n} \sum_{i=0}^{n-1} (x_{j,i} w_{j,i})^2 \right)^{1/2}, \quad j = 0, \dots, nv - 1,$$

where  $x_j$  is a vector in the vector array  $X$ ,  $w_j$  is a weight vector in the vector array  $W$ , and  $m$  is the output array of scalars containing the computed norms. The operation returns a *SUNErrCode*.

Usage:

```
retval = N_VWrmsNormVectorArray(nv, X, W, m);
```

**SUNErrCode N\_VWrmsNormMaskVectorArray**(int nv, *N\_Vector* \*X, *N\_Vector* \*W, *N\_Vector* id, *sunrealtype* \*m)

This routine computes the masked weighted root mean square norm of each vector in a vector array:

$$m_j = \left( \frac{1}{n} \sum_{i=0}^{n-1} (x_{j,i} w_{j,i} H(id_i))^2 \right)^{1/2}, \quad j = 0, \dots, nv - 1,$$

where  $H(id_i) = 1$  if  $id_i > 0$  and is zero otherwise,  $x_j$  is a vector in the vector array  $X$ ,  $w_j$  is a weight vector in the vector array  $W$ ,  $id$  is the mask vector, and  $m$  is the output array of scalars containing the computed norms. The operation returns a *SUNErrCode*.

Usage:

```
retval = N_VWrmsNormMaskVectorArray(nv, X, W, id, m);
```

**SUNErrCode N\_VScaleAddMultiVectorArray**(int nv, int nsum, *sunrealtype* \*c, *N\_Vector* \*X, *N\_Vector* \*\*YY, *N\_Vector* \*\*ZZ)

This routine scales and adds a vector array of  $nv$  vectors to  $nsum$  other vector arrays:

$$z_{k,j,i} = c_k x_{j,i} + y_{k,j,i}, \quad i = 0, \dots, n - 1 \quad j = 0, \dots, nv - 1, \quad k = 0, \dots, nsum - 1$$

where  $c$  is an array of scalars,  $x_j$  is a vector in the vector array  $X$ ,  $y_{k,j}$  is a vector in the array of vector arrays  $YY$ , and  $z_{k,j}$  is an output vector in the array of vector arrays  $ZZ$ . The operation returns a *SUNErrCode*.

Usage:

```
retval = N_VScaleAddMultiVectorArray(nv, nsum, c, x, YY, ZZ);
```

**SUNErrCode N\_VLinearCombinationVectorArray**(int nv, int nsum, *sunrealtype* \*c, *N\_Vector* \*\*XX, *N\_Vector* \*Z)

This routine computes the linear combination of  $nsum$  vector arrays containing  $nv$  vectors:

$$z_{j,i} = \sum_{k=0}^{nsum-1} c_k x_{k,j,i}, \quad i = 0, \dots, n - 1 \quad j = 0, \dots, nv - 1,$$

where  $c$  is an array of scalars,  $x_{k,j}$  is a vector in array of vector arrays  $XX$ , and  $z_{j,i}$  is an output vector in the vector array  $Z$ . If the output vector array is one of the vector arrays in  $XX$ , it *must* be the first vector array in  $XX$ . The operation returns a *SUNErrCode*.

Usage:

```
retval = N_VLinearCombinationVectorArray(nv, nsum, c, XX, Z);
```

## 6.2.4 Local reduction operations

The following local reduction operations are also *optional*. As with the fused and vector array operations, these are intended to reduce parallel communication on distributed memory systems. If a particular NVECTOR implementation defines one of the local reduction operations as NULL, the NVECTOR interface will call one of the above standard vector operations as necessary. As above, for each operation, we give the name, usage of the function, and a description of its mathematical operations below.

*sunrealtype* **N\_VDotProdLocal**(*N\_Vector* x, *N\_Vector* y)

This routine computes the MPI task-local portion of the ordinary dot product of  $x$  and  $y$ :

$$d = \sum_{i=0}^{n_{local}-1} x_i y_i,$$

where  $n_{local}$  corresponds to the number of components in the vector on this MPI task (or  $n_{local} = n$  for MPI-unaware applications).

Usage:

```
d = N_VDotProdLocal(x, y);
```

*sunrealtype* **N\_VMaxNormLocal**(*N\_Vector* x)

This routine computes the MPI task-local portion of the maximum norm of the NVECTOR  $x$ :

$$m = \max_{0 \leq i < n_{local}} |x_i|,$$

where  $n_{local}$  corresponds to the number of components in the vector on this MPI task (or  $n_{local} = n$  for MPI-unaware applications).

Usage:

```
m = N_VMaxNormLocal(x);
```

*sunrealtype* **N\_VMinLocal**(*N\_Vector* x)

This routine computes the smallest element of the MPI task-local portion of the NVECTOR  $x$ :

$$m = \min_{0 \leq i < n_{local}} x_i,$$

where  $n_{local}$  corresponds to the number of components in the vector on this MPI task (or  $n_{local} = n$  for MPI-unaware applications).

Usage:

```
m = N_VMinLocal(x);
```

*sunrealtype* **N\_VL1NormLocal**(*N\_Vector* x)

This routine computes the MPI task-local portion of the  $l_1$  norm of the N\_Vector  $x$ :

$$n = \sum_{i=0}^{n_{local}-1} |x_i|,$$

where  $n_{local}$  corresponds to the number of components in the vector on this MPI task (or  $n_{local} = n$  for MPI-unaware applications).

Usage:

```
n = N_VL1NormLocal(x);
```

*sunrealtype* **N\_VWSqrSumLocal**(*N\_Vector* x, *N\_Vector* w)

This routine computes the MPI task-local portion of the weighted squared sum of the NVECTOR  $x$  with weight vector  $w$ :

$$s = \sum_{i=0}^{n_{local}-1} (x_i w_i)^2,$$

where  $n_{local}$  corresponds to the number of components in the vector on this MPI task (or  $n_{local} = n$  for MPI-unaware applications).

Usage:

```
s = N_VWSqrSumLocal(x, w);
```

*sunrealtype* **N\_VWSqrSumMaskLocal**(*N\_Vector* x, *N\_Vector* w, *N\_Vector* id)

This routine computes the MPI task-local portion of the weighted squared sum of the NVECTOR  $x$  with weight vector  $w$  built using only the elements of  $x$  corresponding to positive elements of the NVECTOR  $id$ :

$$m = \sum_{i=0}^{n_{local}-1} (x_i w_i H(id_i))^2,$$

where

$$H(\alpha) = \begin{cases} 1 & \alpha > 0 \\ 0 & \alpha \leq 0 \end{cases}$$

and  $n_{local}$  corresponds to the number of components in the vector on this MPI task (or  $n_{local} = n$  for MPI-unaware applications).

Usage:

```
s = N_VWSqrSumMaskLocal(x, w, id);
```

*sunbooleantype* **N\_VInvTestLocal**(*N\_Vector* x)

This routine sets the MPI task-local components of the NVECTOR  $z$  to be the inverses of the components of the NVECTOR  $x$ , with prior testing for zero values:

$$z_i = \frac{1}{x_i}, \quad i = 0, \dots, n_{local} - 1$$

where  $n_{local}$  corresponds to the number of components in the vector on this MPI task (or  $n_{local} = n$  for MPI-unaware applications). This routine returns a boolean assigned to SUNTRUE if all task-local components of  $x$  are nonzero (successful inversion) and returns SUNFALSE otherwise.

Usage:

```
t = N_VInvTestLocal(x);
```

*sunbooleantype* **N\_VConstrMaskLocal**(*N\_Vector* c, *N\_Vector* x, *N\_Vector* m)

Performs the following constraint tests based on the values in  $c_i$ :

$$\begin{aligned} x_i &> 0 & \text{if } c_i = 2, \\ x_i &\geq 0 & \text{if } c_i = 1, \\ x_i &< 0 & \text{if } c_i = -2, \\ x_i &\leq 0 & \text{if } c_i = -1. \end{aligned}$$

for all MPI task-local components of the vectors. This routine returns a boolean assigned to `SUNFALSE` if any task-local element failed the constraint test and assigned to `SUNTRUE` if all passed. It also sets a mask vector  $m$ , with elements equal to 1.0 where the constraint test failed, and 0.0 where the test passed. This routine is used only for constraint checking.

Usage:

```
t = N_VConstrMaskLocal(c, x, m);
```

*sunrealtype* **N\_VMinQuotientLocal**(*N\_Vector* num, *N\_Vector* denom)

This routine returns the minimum of the quotients obtained by term-wise dividing  $num_i$  by  $denom_i$ , for all MPI task-local components of the vectors. A zero element in  $denom$  will be skipped. If no such quotients are found, then the large value `SUN_BIG_REAL` (defined in the header file `sundials_types.h`) is returned.

Usage:

```
minq = N_VMinQuotientLocal(num, denom);
```

## 6.2.5 Single Buffer Reduction Operations

The following *optional* operations are used to combine separate reductions into a single MPI call by splitting the local computation and communication into separate functions. These operations are used in low-synchronization orthogonalization methods to reduce the number of MPI `Allreduce` calls. If a particular `NVECTOR` implementation does not define these operations additional communication will be required.

*SUNErrCode* **N\_VDotProdMultiLocal**(int nv, *N\_Vector* x, *N\_Vector* \*Y, *sunrealtype* \*d)

This routine computes the MPI task-local portion of the dot product of a vector  $x$  with  $nv$  vectors  $y_j$ :

$$d_j = \sum_{i=0}^{n_{local}-1} x_i y_{j,i}, \quad j = 0, \dots, nv - 1,$$

where  $d$  is an array of scalars containing the computed dot products,  $x$  is a vector,  $y_j$  is a vector in the vector array  $Y$ , and  $n_{local}$  corresponds to the number of components in the vector on this MPI task. The operation returns a *SUNErrCode*.

Usage:

```
retval = N_VDotProdMultiLocal(nv, x, Y, d);
```

*SUNErrCode* **N\_VDotProdMultiAllReduce**(int nv, *N\_Vector* x, *sunrealtype* \*d)

This routine combines the MPI task-local portions of the dot product of a vector  $x$  with  $nv$  vectors:

```
retval = MPI_Allreduce(MPI_IN_PLACE, d, nv, MPI_SUNREALTYPE, MPI_SUM, comm)
```

where  $d$  is an array of  $nv$  scalars containing the local contributions to the dot product and  $comm$  is the MPI communicator associated with the vector  $x$ . The operation returns a *SUNErrCode*.

Usage:

```
retval = N_VDotProdMultiAllReduce(nv, x, d);
```

## 6.2.6 Exchange operations

The following vector exchange operations are also *optional* and are intended only for use when interfacing with the XBraid library for parallel-in-time integration. In that setting these operations are required but are otherwise unused by SUNDIALS packages and may be set to NULL. For each operation, we give the function signature, a description of the expected behavior, and an example of the function usage.

*SUNErrCode* **N\_VBufSize**(*N\_Vector* x, *sunindextype* \*size)

This routine returns the buffer size need to exchange in the data in the vector *x* between computational nodes.

Usage:

```
flag = N_VBufSize(x, &buf_size)
```

*SUNErrCode* **N\_VBufPack**(*N\_Vector* x, void \*buf)

This routine fills the exchange buffer *buf* with the vector data in *x*.

Usage:

```
flag = N_VBufPack(x, &buf)
```

*SUNErrCode* **N\_VBufUnpack**(*N\_Vector* x, void \*buf)

This routine unpacks the data in the exchange buffer *buf* into the vector *x*.

Usage:

```
flag = N_VBufUnpack(x, buf)
```

## 6.2.7 Output operations

The following optional vector operations are for writing vector data either to stdout or to a given file.

void **N\_VPrint**(*N\_Vector* x)

This routine prints vector data to stdout

Usage:

```
N_VPrint(x);
```

void **N\_VPrintFile**(*N\_Vector* x, FILE \*file)

This routine writes vector data to the given file pointer.

Usage:

```
FILE* fp = fopen("vector_data.txt", "w");
N_VPrintFile(x, fp);
fclose(fp);
```

### 6.3 NVECTOR functions used by IDAS

In Table 6.2 below, we list the vector functions used in the `N_Vector` module used by the IDAS package. The table also shows, for each function, which of the code modules uses the function. The IDAS column shows function usage within the main integrator module, while the remaining columns show function usage within the IDALS linear solvers interface, and the IDABBDPRE preconditioner module.

At this point, we should emphasize that the IDAS user does not need to know anything about the usage of vector functions by the IDAS code modules in order to use IDAS. The information is presented as an implementation detail for the interested reader.

Table 6.2: List of vector functions usage by IDAS code modules

	IDAS	IDALS	IDABBDPRE	IDAA
<code>N_VGetVectorID()</code>				
<code>N_VGetLength()</code>	4			
<code>N_VClone()</code>	x	x	x	x
<code>N_VCloneEmpty()</code>	1			
<code>N_VDestroy()</code>	x	x	x	x
<code>N_VCloneVectorArray()</code>	x	x		
<code>N_VDestroyVectorArray()</code>	x	x		
<code>N_VSpace()</code>	x	2		
<code>N_VGetArrayPointer()</code>	1	x		
<code>N_VSetArrayPointer()</code>	1			
<code>N_VLinearSum()</code>	x	x	x	
<code>N_VConst()</code>	x	x	x	
<code>N_VProd()</code>	x			
<code>N_VDiv()</code>	x			
<code>N_VScale()</code>	x	x	x	x
<code>N_VAbs()</code>	x			
<code>N_VInv()</code>	x			
<code>N_VAddConst()</code>	x			
<code>N_VMaxNorm()</code>	x			
<code>N_VWrmsNorm()</code>	x	x		
<code>N_VMin()</code>	x			
<code>N_VMinQuotient()</code>	x			
<code>N_VConstrMask()</code>	x			
<code>N_VWrmsNormMask()</code>	x			
<code>N_VCompare()</code>	x			
<code>N_VLinearCombination()</code>	x			
<code>N_VScaleAddMulti()</code>	x			
<code>N_VDotProdMulti()</code>	3			
<code>N_VLinearSumVectorArray()</code>	x			
<code>N_VScaleVectorArray()</code>	x			
<code>N_VConstVectorArray()</code>	x			
<code>N_VWrmsNormVectorArray()</code>	x			
<code>N_VWrmsNormMaskVectorArray()</code>	x			
<code>N_VScaleAddMultiVectorArray()</code>	x			
<code>N_VLinearCombinationVectorArray()</code>	x			

Special cases (numbers match markings in table):

1. These routines are only required if an internal difference-quotient routine for constructing `SUNMATRIX_DENSE` or `SUNMATRIX_BAND` Jacobian matrices is used.



2. This routine is optional, and is only used in estimating space requirements for IDAS modules for user feedback.
3. The optional function `N_VDotProdMulti` is only used when Classical Gram-Schmidt is enabled with SPGMR or SPFGMR. The remaining operations from Tables §6.2.2 and §6.2.3 not listed above are unused and a user-supplied `N_Vector` module for IDAS could omit these operations.
4. This routine is only used when an iterative or matrix iterative `SUNLinearSolver` module is supplied to IDAS.

Of the functions listed in §6.2, `N_VDotProd()`, `N_VWL2Norm()`, `N_VL1Norm()`, `N_VInvTest()`, and `N_VGetCommunicator()` are *not* used by IDAS. Therefore a user-supplied `N_Vector` module for IDAS could omit these functions (although some may be needed by `SUNNonlinearSolver` or `SUNLinearSolver` modules).

## 6.4 The NVECTOR\_SERIAL Module

The serial implementation of the `NVECTOR` module provided with SUNDIALS, `NVECTOR_SERIAL`, defines the *content* field of an `N_Vector` to be a structure containing the length of the vector, a pointer to the beginning of a contiguous data array, and a boolean flag *own\_data* which specifies the ownership of data.

```
struct _N_VectorContent_Serial {
    sunindextype length;
    sunbooleantype own_data;
    sunrealtype *data;
};
```

The header file to be included when using this module is `nvector_serial.h`. The installed module library to link to is `libsundials_nvecserial.lib` where `.lib` is typically `.so` for shared libraries and `.a` for static libraries.

### 6.4.1 NVECTOR\_SERIAL accessor macros

The following five macros are provided to access the content of an `NVECTOR_SERIAL` vector. The suffix `_S` in the names denotes the serial version.

#### NV\_CONTENT\_S(v)

This macro gives access to the contents of the serial vector `N_Vector v`.

The assignment `v_cont = NV_CONTENT_S(v)` sets `v_cont` to be a pointer to the serial `N_Vector` *content* structure.

Implementation:

```
#define NV_CONTENT_S(v) ( (N_VectorContent_Serial)(v->content) )
```

#### NV\_OWN\_DATA\_S(v)

Access the *own\_data* component of the serial `N_Vector v`.

Implementation:

```
#define NV_OWN_DATA_S(v) ( NV_CONTENT_S(v)->own_data )
```

#### NV\_DATA\_S(v)

The assignment `v_data = NV_DATA_S(v)` sets `v_data` to be a pointer to the first component of the *data* for the `N_Vector v`.

Similarly, the assignment `NV_DATA_S(v) = v_data` sets the component array of `v` to be `v_data` by storing the pointer `v_data`.

Implementation:

```
#define NV_DATA_S(v) ( NV_CONTENT_S(v)->data )
```

#### **NV\_LENGTH\_S(v)**

Access the *length* component of the serial `N_Vector` `v`.

The assignment `v_len = NV_LENGTH_S(v)` sets `v_len` to be the *length* of `v`. On the other hand, the call `NV_LENGTH_S(v) = len_v` sets the *length* of `v` to be `len_v`.

Implementation:

```
#define NV_LENGTH_S(v) ( NV_CONTENT_S(v)->length )
```

#### **NV\_Ith\_S(v, i)**

This macro gives access to the individual components of the *data* array of an `N_Vector`, using standard 0-based C indexing.

The assignment `r = NV_Ith_S(v, i)` sets `r` to be the value of the *i*-th component of `v`.

The assignment `NV_Ith_S(v, i) = r` sets the value of the *i*-th component of `v` to be `r`.

Here *i* ranges from 0 to  $n - 1$  for a vector of length *n*.

Implementation:

```
#define NV_Ith_S(v,i) ( NV_DATA_S(v)[i] )
```

## **6.4.2 NVECTOR\_SERIAL functions**

The `NVECTOR_SERIAL` module defines serial implementations of all vector operations listed in §6.2.1, §6.2.2, §6.2.3, and §6.2.4. Their names are obtained from those in those sections by appending the suffix `_Serial` (e.g. `N_VDestroy_Serial`). All the standard vector operations listed in §6.2.1 with the suffix `_Serial` appended are callable via the Fortran 2003 interface by prepending an `F` (e.g. `FN_VDestroy_Serial`).

The module `NVECTOR_SERIAL` provides the following additional user-callable routines:

`N_Vector N_VNew_Serial`(*sunindextype* `vec_length`, *SUNContext* `sunctx`)

This function creates and allocates memory for a serial `N_Vector`. Its only argument is the vector length.

`N_Vector N_VNewEmpty_Serial`(*sunindextype* `vec_length`, *SUNContext* `sunctx`)

This function creates a new serial `N_Vector` with an empty (NULL) data array.

`N_Vector N_VMake_Serial`(*sunindextype* `vec_length`, *sunrealttype* `*v_data`, *SUNContext* `sunctx`)

This function creates and allocates memory for a serial vector with user-provided data array, `v_data`.

(This function does *not* allocate memory for `v_data` itself.)

`void N_VPrint_Serial`(*N\_Vector* `v`)

This function prints the content of a serial vector to `stdout`.

`void N_VPrintFile_Serial`(*N\_Vector* `v`, `FILE` `*outfile`)

This function prints the content of a serial vector to `outfile`.

By default all fused and vector array operations are disabled in the `NVECTOR_SERIAL` module. The following additional user-callable routines are provided to enable or disable fused and vector array operations for a specific vector. To ensure consistency across vectors it is recommended to first create a vector with `N_VNew_Serial()`, enable/disable the desired operations for that vector with the functions below, and create any additional vectors from that vector using `N_VClone()`. This guarantees that the new vectors will have the same operations enabled/disabled as cloned

vectors inherit the same enable/disable options as the vector they are cloned, from while vectors created with `N_VNew_Serial()` will have the default settings for the NVECTOR\_SERIAL module.

**`SUNErrCode N_VEnableFusedOps_Serial(N_Vector v, sunbooleantype tf)`**

This function enables (SUNTRUE) or disables (SUNFALSE) all fused and vector array operations in the serial vector. The return value is a `SUNErrCode`.

**`SUNErrCode N_VEnableLinearCombination_Serial(N_Vector v, sunbooleantype tf)`**

This function enables (SUNTRUE) or disables (SUNFALSE) the linear combination fused operation in the serial vector. The return value is a `SUNErrCode`.

**`SUNErrCode N_VEnableScaleAddMulti_Serial(N_Vector v, sunbooleantype tf)`**

This function enables (SUNTRUE) or disables (SUNFALSE) the scale and add a vector to multiple vectors fused operation in the serial vector. The return value is a `SUNErrCode`.

**`SUNErrCode N_VEnableDotProdMulti_Serial(N_Vector v, sunbooleantype tf)`**

This function enables (SUNTRUE) or disables (SUNFALSE) the multiple dot products fused operation in the serial vector. The return value is a `SUNErrCode`.

**`SUNErrCode N_VEnableLinearSumVectorArray_Serial(N_Vector v, sunbooleantype tf)`**

This function enables (SUNTRUE) or disables (SUNFALSE) the linear sum operation for vector arrays in the serial vector. The return value is a `SUNErrCode`.

**`SUNErrCode N_VEnableScaleVectorArray_Serial(N_Vector v, sunbooleantype tf)`**

This function enables (SUNTRUE) or disables (SUNFALSE) the scale operation for vector arrays in the serial vector. The return value is a `SUNErrCode`.

**`SUNErrCode N_VEnableConstVectorArray_Serial(N_Vector v, sunbooleantype tf)`**

This function enables (SUNTRUE) or disables (SUNFALSE) the const operation for vector arrays in the serial vector. The return value is a `SUNErrCode`.

**`SUNErrCode N_VEnableWrmsNormVectorArray_Serial(N_Vector v, sunbooleantype tf)`**

This function enables (SUNTRUE) or disables (SUNFALSE) the WRMS norm operation for vector arrays in the serial vector. The return value is a `SUNErrCode`.

**`SUNErrCode N_VEnableWrmsNormMaskVectorArray_Serial(N_Vector v, sunbooleantype tf)`**

This function enables (SUNTRUE) or disables (SUNFALSE) the masked WRMS norm operation for vector arrays in the serial vector. The return value is a `SUNErrCode`.

**`SUNErrCode N_VEnableScaleAddMultiVectorArray_Serial(N_Vector v, sunbooleantype tf)`**

This function enables (SUNTRUE) or disables (SUNFALSE) the scale and add a vector array to multiple vector arrays operation in the serial vector. The return value is a `SUNErrCode`.

**`SUNErrCode N_VEnableLinearCombinationVectorArray_Serial(N_Vector v, sunbooleantype tf)`**

This function enables (SUNTRUE) or disables (SUNFALSE) the linear combination operation for vector arrays in the serial vector. The return value is a `SUNErrCode`.

## Notes

- When looping over the components of an `N_Vector v`, it is more efficient to first obtain the component array via `v_data = NV_DATA_S(v)`, or equivalently `v_data = N_VGetArrayPointer(v)`, and then access `v_data[i]` within the loop than it is to use `NV_Ith_S(v, i)` within the loop.
- `N_VNewEmpty_Serial()` and `N_VMake_Serial()` set the field `own_data` to `SUNFALSE`. The implementation of `N_VDestroy()` will not attempt to free the pointer data for any `N_Vector` with `own_data` set to `SUNFALSE`. In such a case, it is the user's responsibility to deallocate the data pointer.

- To maximize efficiency, vector operations in the NVECTOR\_SERIAL implementation that have more than one `N_Vector` argument do not check for consistent internal representation of these vectors. It is the user's responsibility to ensure that such routines are called with `N_Vector` arguments that were all created with the same length.

### 6.4.3 NVECTOR\_SERIAL Fortran Interface

The NVECTOR\_SERIAL module provides a Fortran 2003 module for use from Fortran applications.

The `fnvector_serial_mod` Fortran module defines interfaces to all NVECTOR\_SERIAL C functions using the intrinsic `iso_c_binding` module which provides a standardized mechanism for interoperating with C. As noted in the C function descriptions above, the interface functions are named after the corresponding C function, but with a leading `F`. For example, the function `N_VNew_Serial` is interfaced as `FN_VNew_Serial`.

The Fortran 2003 NVECTOR\_SERIAL interface module can be accessed with the `use` statement, i.e. `use fnvector_serial_mod`, and linking to the library `libsundials_fnvectorserial_mod.lib` in addition to the C library. For details on where the library and module file `fnvector_serial_mod.mod` are installed see §11. We note that the module is accessible from the Fortran 2003 SUNDIALS integrators *without* separately linking to the `libsundials_fnvectorserial_mod` library.

## 6.5 The NVECTOR\_PARALLEL Module

The NVECTOR\_PARALLEL implementation of the NVECTOR module provided with SUNDIALS is based on MPI. It defines the *content* field of an `N_Vector` to be a structure containing the global and local lengths of the vector, a pointer to the beginning of a contiguous local data array, an MPI communicator, an a boolean flag *own\_data* indicating ownership of the data array *data*.

```
struct _N_VectorContent_Parallel {  
    sunindextype local_length;  
    sunindextype global_length;  
    sunboolean_t own_data;  
    sunrealtype *data;  
    MPI_Comm comm;  
};
```

The header file to be included when using this module is `nvector_parallel.h`. The installed module library to link to is `libsundials_nvecparallel.lib` where `.lib` is typically `.so` for shared libraries and `.a` for static libraries.

### 6.5.1 NVECTOR\_PARALLEL accessor macros

The following seven macros are provided to access the content of a NVECTOR\_PARALLEL vector. The suffix `_P` in the names denotes the distributed memory parallel version.

#### NV\_CONTENT\_P(v)

This macro gives access to the contents of the parallel `N_Vector` *v*.

The assignment `v_cont = NV_CONTENT_P(v)` sets `v_cont` to be a pointer to the `N_Vector` *content* structure of type `struct N_VectorContent_Parallel`.

Implementation:

```
#define NV_CONTENT_P(v) ( (N_VectorContent_Parallel)(v->content) )
```

**NV\_OWN\_DATA\_P(v)**

Access the *own\_data* component of the parallel *N\_Vector* *v*.

Implementation:

```
#define NV_OWN_DATA_P(v) ( NV_CONTENT_P(v)->own_data )
```

**NV\_DATA\_P(v)**

The assignment `v_data = NV_DATA_P(v)` sets *v\_data* to be a pointer to the first component of the *local\_data* for the *N\_Vector* *v*.

The assignment `NV_DATA_P(v) = v_data` sets the component array of *v* to be *v\_data* by storing the pointer *v\_data* into *data*.

Implementation:

```
#define NV_DATA_P(v) ( NV_CONTENT_P(v)->data )
```

**NV\_LOCLENGTH\_P(v)**

The assignment `v_llen = NV_LOCLENGTH_P(v)` sets *v\_llen* to be the length of the local part of *v*.

The call `NV_LOCLENGTH_P(v) = llen_v` sets the *local\_length* of *v* to be *llen\_v*.

Implementation:

```
#define NV_LOCLENGTH_P(v) ( NV_CONTENT_P(v)->local_length )
```

**NV\_GLOBLENGTH\_P(v)**

The assignment `v_glen = NV_GLOBLENGTH_P(v)` sets *v\_glen* to be the *global\_length* of the vector *v*.

The call `NV_GLOBLENGTH_P(v) = glen_v` sets the *global\_length* of *v* to be *glen\_v*.

Implementation:

```
#define NV_GLOBLENGTH_P(v) ( NV_CONTENT_P(v)->global_length )
```

**NV\_COMM\_P(v)**

This macro provides access to the MPI communicator used by the parallel *N\_Vector* *v*.

Implementation:

```
#define NV_COMM_P(v) ( NV_CONTENT_P(v)->comm )
```

**NV\_Ith\_P(v, i)**

This macro gives access to the individual components of the *local\_data* array of an *N\_Vector*.

The assignment `r = NV_Ith_P(v, i)` sets *r* to be the value of the *i*-th component of the local part of *v*.

The assignment `NV_Ith_P(v, i) = r` sets the value of the *i*-th component of the local part of *v* to be *r*.

Here *i* ranges from 0 to  $n - 1$ , where  $n$  is the *local\_length*.

Implementation:

```
#define NV_Ith_P(v, i) ( NV_DATA_P(v)[i] )
```

## 6.5.2 NVECTOR\_PARALLEL functions

The NVECTOR\_PARALLEL module defines parallel implementations of all vector operations listed in §6.2. Their names are obtained from the generic names by appending the suffix `_Parallel` (e.g. `N_VDestroy_Parallel`). The module NVECTOR\_PARALLEL provides the following additional user-callable routines:

*N\_Vector* **N\_VNew\_Parallel**(MPI\_Comm comm, *sunindextype* local\_length, *sunindextype* global\_length, *SUNContext* sunctx)

This function creates and allocates memory for a parallel vector having global length *global\_length*, having processor-local length *local\_length*, and using the MPI communicator *comm*.

*N\_Vector* **N\_VNewEmpty\_Parallel**(MPI\_Comm comm, *sunindextype* local\_length, *sunindextype* global\_length, *SUNContext* sunctx)

This function creates a new parallel *N\_Vector* with an empty (NULL) data array.

*N\_Vector* **N\_VMake\_Parallel**(MPI\_Comm comm, *sunindextype* local\_length, *sunindextype* global\_length, *sunrealtype* \*v\_data, *SUNContext* sunctx)

This function creates and allocates memory for a parallel vector with user-provided data array.

(This function does *not* allocate memory for *v\_data* itself.)

*sunindextype* **N\_VGetLocalLength\_Parallel**(*N\_Vector* v)

This function returns the local vector length.

void **N\_VPrint\_Parallel**(*N\_Vector* v)

This function prints the local content of a parallel vector to stdout.

void **N\_VPrintFile\_Parallel**(*N\_Vector* v, FILE \*outfile)

This function prints the local content of a parallel vector to outfile.

By default all fused and vector array operations are disabled in the NVECTOR\_PARALLEL module. The following additional user-callable routines are provided to enable or disable fused and vector array operations for a specific vector. To ensure consistency across vectors it is recommended to first create a vector with `N_VNew_Parallel()`, enable/disable the desired operations for that vector with the functions below, and create any additional vectors from that vector using `N_VClone()`. This guarantees that the new vectors will have the same operations enabled/disabled as cloned vectors inherit the same enable/disable options as the vector they are cloned from, while vectors created with `N_VNew_Parallel()` will have the default settings for the NVECTOR\_PARALLEL module.

*SUNErrCode* **N\_VEnableFusedOps\_Parallel**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) all fused and vector array operations in the parallel vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearCombination\_Parallel**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear combination fused operation in the parallel vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleAddMulti\_Parallel**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale and add a vector to multiple vectors fused operation in the parallel vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableDotProdMulti\_Parallel**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the multiple dot products fused operation in the parallel vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearSumVectorArray\_Parallel**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear sum operation for vector arrays in the parallel vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleVectorArray\_Parallel**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale operation for vector arrays in the parallel vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableConstVectorArray\_Parallel**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the const operation for vector arrays in the parallel vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableWrmsNormVectorArray\_Parallel**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the WRMS norm operation for vector arrays in the parallel vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableWrmsNormMaskVectorArray\_Parallel**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the masked WRMS norm operation for vector arrays in the parallel vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleAddMultiVectorArray\_Parallel**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale and add a vector array to multiple vector arrays operation in the parallel vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearCombinationVectorArray\_Parallel**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear combination operation for vector arrays in the parallel vector. The return value is a *SUNErrCode*.

#### Notes

- When looping over the components of an *N\_Vector* v, it is more efficient to first obtain the local component array via `v_data = N_VGetArrayPointer(v)`, or equivalently `v_data = NV_DATA_P(v)`, and then access `v_data[i]` within the loop than it is to use `NV_Ith_P(v, i)` within the loop.
- `N_VNewEmpty_Parallel()` and `N_VMake_Parallel()` set the field `own_data` to `SUNFALSE`. The implementation of `N_VDestroy()` will not attempt to free the pointer data for any *N\_Vector* with `own_data` set to `SUNFALSE`. In such a case, it is the user's responsibility to deallocate the data pointer.
- To maximize efficiency, vector operations in the `NVECTOR_PARALLEL` implementation that have more than one *N\_Vector* argument do not check for consistent internal representation of these vectors. It is the user's responsibility to ensure that such routines are called with *N\_Vector* arguments that were all created with the same internal representations.

### 6.5.3 NVECTOR\_PARALLEL Fortran Interface

The `NVECTOR_PARALLEL` module provides a Fortran 2003 module for use from Fortran applications.

The `fnvector_parallel_mod` Fortran module defines interfaces to all `NVECTOR_PARALLEL` C functions using the intrinsic `iso_c_binding` module which provides a standardized mechanism for interoperating with C. As noted in the C function descriptions above, the interface functions are named after the corresponding C function, but with a leading F. For example, the function `N_VNew_Parallel` is interfaced as `FN_VNew_Parallel`.

The Fortran 2003 `NVECTOR_PARALLEL` interface module can be accessed with the `use` statement, i.e. `use fnvector_parallel_mod`, and linking to the library `libsundials_fnvectorparallel_mod.lib` in addition to the C library. For details on where the library and module file `fnvector_parallel_mod.mod` are installed see §11. We note that the module is accessible from the Fortran 2003 SUNDIALS integrators *without* separately linking to the `libsundials_fnvectorparallel_mod` library.



## 6.6 The NVECTOR\_OPENMP Module

In situations where a user has a multi-core processing unit capable of running multiple parallel threads with shared memory, SUNDIALS provides an implementation of NVECTOR using OpenMP, called NVECTOR\_OPENMP, and an implementation using Pthreads, called NVECTOR\_PTHREADS. Testing has shown that vectors should be of length at least 100,000 before the overhead associated with creating and using the threads is made up by the parallelism in the vector calculations.

The OpenMP NVECTOR implementation provided with SUNDIALS, NVECTOR\_OPENMP, defines the *content* field of *N\_Vector* to be a structure containing the length of the vector, a pointer to the beginning of a contiguous data array, a boolean flag *own\_data* which specifies the ownership of *data*, and the number of threads. Operations on the vector are threaded using OpenMP, the number of threads used is based on the supplied argument in the vector constructor.

```
struct _N_VectorContent_OpenMP {  
    sunindextype length;  
    sunbooleantype own_data;  
    sunrealtype *data;  
    int num_threads;  
};
```

The header file to be included when using this module is `nvector_openmp.h`. The installed module library to link to is `libsundials_nvecopenmp.lib` where `.lib` is typically `.so` for shared libraries and `.a` for static libraries. The Fortran module file to use when using the Fortran 2003 interface to this module is `fnvector_openmp_mod.mod`.

### 6.6.1 NVECTOR\_OPENMP accessor macros

The following six macros are provided to access the content of an NVECTOR\_OPENMP vector. The suffix `_OMP` in the names denotes the OpenMP version.

#### NV\_CONTENT\_OMP(v)

This macro gives access to the contents of the OpenMP vector *N\_Vector* *v*.

The assignment `v_cont = NV_CONTENT_OMP(v)` sets *v\_cont* to be a pointer to the OpenMP *N\_Vector* content structure.

Implementation:

```
#define NV_CONTENT_OMP(v) ( (_N_VectorContent_OpenMP)(v->content) )
```

#### NV\_OWN\_DATA\_OMP(v)

Access the *own\_data* component of the OpenMP *N\_Vector* *v*.

Implementation:

```
#define NV_OWN_DATA_OMP(v) ( NV_CONTENT_OMP(v)->own_data )
```

#### NV\_DATA\_OMP(v)

The assignment `v_data = NV_DATA_OMP(v)` sets *v\_data* to be a pointer to the first component of the *data* for the *N\_Vector* *v*.

Similarly, the assignment `NV_DATA_OMP(v) = v_data` sets the component array of *v* to be *v\_data* by storing the pointer *v\_data*.

Implementation:



```
#define NV_DATA_OMP(v) ( NV_CONTENT_OMP(v)->data )
```

### NV\_LENGTH\_OMP(v)

Access the *length* component of the OpenMP N\_Vector *v*.

The assignment `v_len = NV_LENGTH_OMP(v)` sets `v_len` to be the *length* of *v*. On the other hand, the call `NV_LENGTH_OMP(v) = len_v` sets the *length* of *v* to be `len_v`.

Implementation:

```
#define NV_LENGTH_OMP(v) ( NV_CONTENT_OMP(v)->length )
```

### NV\_NUM\_THREADS\_OMP(v)

Access the *num\_threads* component of the OpenMP N\_Vector *v*.

The assignment `v_threads = NV_NUM_THREADS_OMP(v)` sets `v_threads` to be the *num\_threads* of *v*. On the other hand, the call `NV_NUM_THREADS_OMP(v) = num_threads_v` sets the *num\_threads* of *v* to be `num_threads_v`.

Implementation:

```
#define NV_NUM_THREADS_OMP(v) ( NV_CONTENT_OMP(v)->num_threads )
```

### NV\_Ith\_OMP(v, i)

This macro gives access to the individual components of the *data* array of an N\_Vector, using standard 0-based C indexing.

The assignment `r = NV_Ith_OMP(v, i)` sets `r` to be the value of the *i*-th component of *v*.

The assignment `NV_Ith_OMP(v, i) = r` sets the value of the *i*-th component of *v* to be `r`.

Here *i* ranges from 0 to  $n - 1$  for a vector of length *n*.

Implementation:

```
#define NV_Ith_OMP(v,i) ( NV_DATA_OMP(v)[i] )
```

## 6.6.2 NVECTOR\_OPENMP functions

The NVECTOR\_OPENMP module defines OpenMP implementations of all vector operations listed in §6.2, §6.2.2, §6.2.3, and §6.2.4. Their names are obtained from those in those sections by appending the suffix `_OpenMP` (e.g. `N_VDestroy_OpenMP`). All the standard vector operations listed in §6.2 with the suffix `_OpenMP` appended are callable via the Fortran 2003 interface by prepending an *F*' (e.g. `FN_VDestroy_OpenMP`).

The module NVECTOR\_OPENMP provides the following additional user-callable routines:

*N\_Vector* **N\_VNew\_OpenMP**(*sunindextype* vec\_length, int num\_threads, *SUNContext* sunctx)

This function creates and allocates memory for a OpenMP N\_Vector. Arguments are the vector length and number of threads.

*N\_Vector* **N\_VNewEmpty\_OpenMP**(*sunindextype* vec\_length, int num\_threads, *SUNContext* sunctx)

This function creates a new OpenMP N\_Vector with an empty (NULL) data array.

*N\_Vector* **N\_VMake\_OpenMP**(*sunindextype* vec\_length, *sunrealtype* \*v\_data, int num\_threads, *SUNContext* sunctx)

This function creates and allocates memory for a OpenMP vector with user-provided data array, *v\_data*.

(This function does *not* allocate memory for *v\_data* itself.)

void **N\_VPrint\_OpenMP**(*N\_Vector* v)

This function prints the content of an OpenMP vector to stdout.

void **N\_VPrintFile\_OpenMP**(*N\_Vector* v, FILE \*outfile)

This function prints the content of an OpenMP vector to outfile.

By default all fused and vector array operations are disabled in the NVECTOR\_OPENMP module. The following additional user-callable routines are provided to enable or disable fused and vector array operations for a specific vector. To ensure consistency across vectors it is recommended to first create a vector with [N\\_VNew\\_OpenMP\(\)](#), enable/disable the desired operations for that vector with the functions below, and create any additional vectors from that vector using [N\\_VClone\(\)](#). This guarantees the new vectors will have the same operations enabled/disabled as cloned vectors inherit the same enable/disable options as the vector they are cloned from while vectors created with [N\\_VNew\\_OpenMP\(\)](#) will have the default settings for the NVECTOR\_OPENMP module.

*SUNErrCode* **N\_VEnableFusedOps\_OpenMP**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) all fused and vector array operations in the OpenMP vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearCombination\_OpenMP**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear combination fused operation in the OpenMP vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleAddMulti\_OpenMP**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale and add a vector to multiple vectors fused operation in the OpenMP vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableDotProdMulti\_OpenMP**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the multiple dot products fused operation in the OpenMP vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearSumVectorArray\_OpenMP**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear sum operation for vector arrays in the OpenMP vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleVectorArray\_OpenMP**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale operation for vector arrays in the OpenMP vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableConstVectorArray\_OpenMP**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the const operation for vector arrays in the OpenMP vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableWrmsNormVectorArray\_OpenMP**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the WRMS norm operation for vector arrays in the OpenMP vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableWrmsNormMaskVectorArray\_OpenMP**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the masked WRMS norm operation for vector arrays in the OpenMP vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleAddMultiVectorArray\_OpenMP**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale and add a vector array to multiple vector arrays operation in the OpenMP vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearCombinationVectorArray\_OpenMP**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear combination operation for vector arrays in the OpenMP vector. The return value is a *SUNErrCode*.

## Notes

- When looping over the components of an `N_Vector v`, it is more efficient to first obtain the component array via `v_data = N_VGetArrayPointer(v)`, or equivalently `v_data = NV_DATA_OMP(v)` and then access `v_data[i]` within the loop than it is to use `NV_Ith_OMP(v,i)` within the loop.
- `N_VNewEmpty_OpenMP()` and `N_VMake_OpenMP()` set the field `own_data` to `SUNFALSE`. The implementation of `N_VDestroy()` will not attempt to free the pointer data for any `N_Vector` with `own_data` set to `SUNFALSE`. In such a case, it is the user's responsibility to deallocate the data pointer.
- To maximize efficiency, vector operations in the `NVECTOR_OPENMP` implementation that have more than one `N_Vector` argument do not check for consistent internal representation of these vectors. It is the user's responsibility to ensure that such routines are called with `N_Vector` arguments that were all created with the same internal representations.

### 6.6.3 NVECTOR\_OPENMP Fortran Interface

The `NVECTOR_OPENMP` module provides a Fortran 2003 module for use from Fortran applications.

The `fnvector_openmp_mod` Fortran module defines interfaces to all `NVECTOR_OPENMP` C functions using the intrinsic `iso_c_binding` module which provides a standardized mechanism for interoperating with C. As noted in the C function descriptions above, the interface functions are named after the corresponding C function, but with a leading `F`. For example, the function `N_VNew_OpenMP` is interfaced as `FN_VNew_OpenMP`.

The Fortran 2003 `NVECTOR_OPENMP` interface module can be accessed with the `use` statement, i.e. `use fnvector_openmp_mod`, and linking to the library `libsundials_fnvectoropenmp_mod.lib` in addition to the C library. For details on where the library and module file `fnvector_openmp_mod.mod` are installed see §11.

## 6.7 The NVECTOR\_PTHREADS Module

In situations where a user has a multi-core processing unit capable of running multiple parallel threads with shared memory, SUNDIALS provides an implementation of `NVECTOR` using OpenMP, called `NVECTOR_OPENMP`, and an implementation using Pthreads, called `NVECTOR_PTHREADS`. Testing has shown that vectors should be of length at least 100,000 before the overhead associated with creating and using the threads is made up by the parallelism in the vector calculations.

The Pthreads `NVECTOR` implementation provided with SUNDIALS, denoted `NVECTOR_PTHREADS`, defines the `content` field of `N_Vector` to be a structure containing the length of the vector, a pointer to the beginning of a contiguous data array, a boolean flag `own_data` which specifies the ownership of `data`, and the number of threads. Operations on the vector are threaded using POSIX threads (Pthreads).

```
struct _N_VectorContent_Pthreads {
    sunindextype length;
    sunbooleantype own_data;
    sunrealtype *data;
    int num_threads;
};
```

The header file to be included when using this module is `nvector_pthreads.h`. The installed module library to link to is `libsundials_nvecpthreads.lib` where `.lib` is typically `.so` for shared libraries and `.a` for static libraries.

### 6.7.1 NVECTOR\_PTHREADS accessor macros

The following six macros are provided to access the content of an NVECTOR\_PTHREADS vector. The suffix `_PT` in the names denotes the Pthreads version.

#### NV\_CONTENT\_PT(v)

This macro gives access to the contents of the Pthreads vector `N_Vector v`.

The assignment `v_cont = NV_CONTENT_PT(v)` sets `v_cont` to be a pointer to the Pthreads `N_Vector` content structure.

Implementation:

```
#define NV_CONTENT_PT(v) ( (N_VectorContent_Pthreads)(v->content) )
```

#### NV\_OWN\_DATA\_PT(v)

Access the *own\_data* component of the Pthreads `N_Vector v`.

Implementation:

```
#define NV_OWN_DATA_PT(v) ( NV_CONTENT_PT(v)->own_data )
```

#### NV\_DATA\_PT(v)

The assignment `v_data = NV_DATA_PT(v)` sets `v_data` to be a pointer to the first component of the *data* for the `N_Vector v`.

Similarly, the assignment `NV_DATA_PT(v) = v_data` sets the component array of `v` to be `v_data` by storing the pointer `v_data`.

Implementation:

```
#define NV_DATA_PT(v) ( NV_CONTENT_PT(v)->data )
```

#### NV\_LENGTH\_PT(v)

Access the *length* component of the Pthreads `N_Vector v`.

The assignment `v_len = NV_LENGTH_PT(v)` sets `v_len` to be the *length* of `v`. On the other hand, the call `NV_LENGTH_PT(v) = len_v` sets the *length* of `v` to be `len_v`.

Implementation:

```
#define NV_LENGTH_PT(v) ( NV_CONTENT_PT(v)->length )
```

#### NV\_NUM\_THREADS\_PT(v)

Access the *num\_threads* component of the Pthreads `N_Vector v`.

The assignment `v_threads = NV_NUM_THREADS_PT(v)` sets `v_threads` to be the *num\_threads* of `v`. On the other hand, the call `NV_NUM_THREADS_PT(v) = num_threads_v` sets the *num\_threads* of `v` to be `num_threads_v`.

Implementation:

```
#define NV_NUM_THREADS_PT(v) ( NV_CONTENT_PT(v)->num_threads )
```

#### NV\_Ith\_PT(v, i)

This macro gives access to the individual components of the *data* array of an `N_Vector`, using standard 0-based C indexing.

The assignment `r = NV_Ith_PT(v, i)` sets `r` to be the value of the *i*-th component of `v`.

The assignment `NV_Ith_PT(v,i) = r` sets the value of the  $i$ -th component of  $v$  to be  $r$ .

Here  $i$  ranges from 0 to  $n - 1$  for a vector of length  $n$ .

Implementation:

```
#define NV_Ith_PT(v,i) ( NV_DATA_PT(v)[i] )
```

## 6.7.2 NVECTOR\_PTHREADS functions

The NVECTOR\_PTHREADS module defines Pthreads implementations of all vector operations listed in §6.2, §6.2.2, §6.2.3, and §6.2.4. Their names are obtained from those in those sections by appending the suffix `_Pthreads` (e.g. `N_VDestroy_Pthreads`). All the standard vector operations listed in §6.2 are callable via the Fortran 2003 interface by prepending an *F* (e.g. `FN_VDestroy_Pthreads`). The module NVECTOR\_PTHREADS provides the following additional user-callable routines:

*N\_Vector* **N\_VNew\_Pthreads**(*sunindextype* vec\_length, int num\_threads, *SUNContext* sunctx)

This function creates and allocates memory for a Pthreads *N\_Vector*. Arguments are the vector length and number of threads.

*N\_Vector* **N\_VNewEmpty\_Pthreads**(*sunindextype* vec\_length, int num\_threads, *SUNContext* sunctx)

This function creates a new Pthreads *N\_Vector* with an empty (NULL) data array.

*N\_Vector* **N\_VMake\_Pthreads**(*sunindextype* vec\_length, *sunrealtype* \*v\_data, int num\_threads, *SUNContext* sunctx)

This function creates and allocates memory for a Pthreads vector with user-provided data array, *v\_data*.

(This function does *not* allocate memory for *v\_data* itself.)

void **N\_VPrint\_Pthreads**(*N\_Vector* v)

This function prints the content of a Pthreads vector to `stdout`.

void **N\_VPrintFile\_Pthreads**(*N\_Vector* v, FILE \*outfile)

This function prints the content of a Pthreads vector to `outfile`.

By default all fused and vector array operations are disabled in the NVECTOR\_PTHREADS module. The following additional user-callable routines are provided to enable or disable fused and vector array operations for a specific vector. To ensure consistency across vectors it is recommended to first create a vector with `N_VNew_Pthreads()`, enable/disable the desired operations for that vector with the functions below, and create any additional vectors from that vector using `N_VClone()`. This guarantees the new vectors will have the same operations enabled/disabled as cloned vectors inherit the same enable/disable options as the vector they are cloned from while vectors created with `N_VNew_Pthreads()` will have the default settings for the NVECTOR\_PTHREADS module.

*SUNErrCode* **N\_VEnableFusedOps\_Pthreads**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) all fused and vector array operations in the Pthreads vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearCombination\_Pthreads**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear combination fused operation in the Pthreads vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleAddMulti\_Pthreads**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale and add a vector to multiple vectors fused operation in the Pthreads vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableDotProdMulti\_Pthreads**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the multiple dot products fused operation in the Pthreads vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearSumVectorArray\_Pthreads**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear sum operation for vector arrays in the Pthreads vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleVectorArray\_Pthreads**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale operation for vector arrays in the Pthreads vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableConstVectorArray\_Pthreads**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the const operation for vector arrays in the Pthreads vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableWrmsNormVectorArray\_Pthreads**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the WRMS norm operation for vector arrays in the Pthreads vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableWrmsNormMaskVectorArray\_Pthreads**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the masked WRMS norm operation for vector arrays in the Pthreads vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleAddMultiVectorArray\_Pthreads**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale and add a vector array to multiple vector arrays operation in the Pthreads vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearCombinationVectorArray\_Pthreads**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear combination operation for vector arrays in the Pthreads vector. The return value is a *SUNErrCode*.

#### Notes

- When looping over the components of an *N\_Vector* v, it is more efficient to first obtain the component array via `v_data = N_VGetArrayPointer(v)`, or equivalently `v_data = NV_DATA_PT(v)` and then access `v_data[i]` within the loop than it is to use `NV_Ith_S(v, i)` within the loop.
- *N\_VNewEmpty\_Pthreads()* and *N\_VMake\_Pthreads()* set the field `own_data` to `SUNFALSE`. The implementation of *N\_VDestroy()* will not attempt to free the pointer data for any *N\_Vector* with `own_data` set to `SUNFALSE`. In such a case, it is the user's responsibility to deallocate the data pointer.
- To maximize efficiency, vector operations in the `NVECTOR_PTHREADS` implementation that have more than one *N\_Vector* argument do not check for consistent internal representation of these vectors. It is the user's responsibility to ensure that such routines are called with *N\_Vector* arguments that were all created with the same internal representations.

### 6.7.3 NVECTOR\_PTHREADS Fortran Interface

The `NVECTOR_PTHREADS` module provides a Fortran 2003 module for use from Fortran applications.

The `fnvector_pthreads_mod` Fortran module defines interfaces to all `NVECTOR_PTHREADS` C functions using the intrinsic `iso_c_binding` module which provides a standardized mechanism for interoperating with C. As noted in the C function descriptions above, the interface functions are named after the corresponding C function, but with a leading F. For example, the function `N_VNew_Pthreads` is interfaced as `FN_VNew_Pthreads`.



The Fortran 2003 NVECTOR\_PTHREADS interface module can be accessed with the use statement, i.e. `use fn-vector_pthreads_mod`, and linking to the library `libsundials_fnvectorpthreads_mod.lib` in addition to the C library. For details on where the library and module file `fnvector_pthreads_mod.mod` are installed see §11.

## 6.8 The NVECTOR\_PARHYP Module

The NVECTOR\_PARHYP implementation of the NVECTOR module provided with SUNDIALS is a wrapper around HYPRE's ParVector class. Most of the vector kernels simply call HYPRE vector operations. The implementation defines the *content* field of `N_Vector` to be a structure containing the global and local lengths of the vector, a pointer to an object of type `hypre_ParVector`, an MPI communicator, and a boolean flag *own\_parvector* indicating ownership of the HYPRE parallel vector object *x*.

```
struct _N_VectorContent_ParHyp {
    sunindextype local_length;
    sunindextype global_length;
    sunboolean_t own_data;
    sunboolean_t own_parvector;
    sunrealtype *data;
    MPI_Comm comm;
    hypre_ParVector *x;
};
```

The header file to be included when using this module is `nvector_parhyp.h`. The installed module library to link to is `libsundials_nvecparhyp.lib` where `.lib` is typically `.so` for shared libraries and `.a` for static libraries.

Unlike native SUNDIALS vector types, NVECTOR\_PARHYP does not provide macros to access its member variables. Note that NVECTOR\_PARHYP requires SUNDIALS to be built with MPI support.

### 6.8.1 NVECTOR\_PARHYP functions

The NVECTOR\_PARHYP module defines implementations of all vector operations listed in §6.2 except for `N_VSetArrayPointer()` and `N_VGetArrayPointer()` because accessing raw vector data is handled by low-level HYPRE functions. As such, this vector is not available for use with SUNDIALS Fortran interfaces. When access to raw vector data is needed, one should extract the HYPRE vector first, and then use HYPRE methods to access the data. Usage examples of NVECTOR\_PARHYP are provided in the `cvAdvDiff_non_ph.c` example programs for CVODE and the `ark_diurnal_kry_ph.c` example program for ARKODE.

The names of parhyp methods are obtained from those in §6.2, §6.2.2, §6.2.3, and §6.2.4 by appending the suffix `_ParHyp` (e.g. `N_VDestroy_ParHyp`). The module NVECTOR\_PARHYP provides the following additional user-callable routines:

`N_Vector N_VNewEmpty_ParHyp(MPI_Comm comm, sunindextype local_length, sunindextype global_length, SUNContext sunctx)`

This function creates a new parhyp `N_Vector` with the pointer to the HYPRE vector set to NULL.

`N_Vector N_VMake_ParHyp(hypre_ParVector *x, SUNContext sunctx)`

This function creates an `N_Vector` wrapper around an existing HYPRE parallel vector. It does *not* allocate memory for *x* itself.

`hypre_ParVector *N_VGetVector_ParHyp(N_Vector v)`

This function returns a pointer to the underlying HYPRE vector.

void **N\_VPrint\_ParHyp**(*N\_Vector* v)

This function prints the local content of a parhyp vector to `stdout`.

void **N\_VPrintFile\_ParHyp**(*N\_Vector* v, FILE \*outfile)

This function prints the local content of a parhyp vector to `outfile`.

By default all fused and vector array operations are disabled in the `NVECTOR_PARHYP` module. The following additional user-callable routines are provided to enable or disable fused and vector array operations for a specific vector. To ensure consistency across vectors it is recommended to first create a vector with `N_VMake_ParHyp()`, enable/disable the desired operations for that vector with the functions below, and create any additional vectors from that vector using `N_VClone()`. This guarantees the new vectors will have the same operations enabled/disabled as cloned vectors inherit the same enable/disable options as the vector they are cloned from while vectors created with `N_VMake_ParHyp()` will have the default settings for the `NVECTOR_PARHYP` module.

*SUNErrCode* **N\_VEnableFusedOps\_ParHyp**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) all fused and vector array operations in the parhyp vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearCombination\_ParHyp**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the linear combination fused operation in the parhyp vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleAddMulti\_ParHyp**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the scale and add a vector to multiple vectors fused operation in the parhyp vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableDotProdMulti\_ParHyp**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the multiple dot products fused operation in the parhyp vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearSumVectorArray\_ParHyp**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the linear sum operation for vector arrays in the parhyp vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleVectorArray\_ParHyp**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the scale operation for vector arrays in the parhyp vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableConstVectorArray\_ParHyp**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the const operation for vector arrays in the parhyp vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableWrmsNormVectorArray\_ParHyp**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the WRMS norm operation for vector arrays in the parhyp vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableWrmsNormMaskVectorArray\_ParHyp**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the masked WRMS norm operation for vector arrays in the parhyp vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleAddMultiVectorArray\_ParHyp**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the scale and add a vector array to multiple vector arrays operation in the parhyp vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearCombinationVectorArray\_ParHyp**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the linear combination operation for vector arrays in the parhyp vector. The return value is a *SUNErrCode*.



## Notes

- When there is a need to access components of an `N_Vector_ParHyp` `v`, it is recommended to extract the HYPRE vector via `x_vec = N_VGetVector_ParHyp(v)` and then access components using appropriate HYPRE functions.
- `N_VNewEmpty_ParHyp()`, and `N_VMake_ParHyp()` set the field `own_parvector` to `SUNFALSE`. The implementation of `N_VDestroy()` will not attempt to delete an underlying HYPRE vector for any `N_Vector` with `own_parvector` set to `SUNFALSE`. In such a case, it is the user's responsibility to delete the underlying vector.
- To maximize efficiency, vector operations in the `NVECTOR_PARHYP` implementation that have more than one `N_Vector` argument do not check for consistent internal representations of these vectors. It is the user's responsibility to ensure that such routines are called with `N_Vector` arguments that were all created with the same internal representations.

## 6.9 The NVECTOR\_PETSC Module

The `NVECTOR_PETSC` module is an `NVECTOR` wrapper around the PETSc vector. It defines the `content` field of a `N_Vector` to be a structure containing the global and local lengths of the vector, a pointer to the PETSc vector, an MPI communicator, and a boolean flag `own_data` indicating ownership of the wrapped PETSc vector.

```
struct _N_VectorContent_Petsc {
    sunindextype local_length;
    sunindextype global_length;
    sunbooleantype own_data;
    Vec *pvec;
    MPI_Comm comm;
};
```

The header file to be included when using this module is `nvector_petsc.h`. The installed module library to link to is `libsundials_nvecpetsc.lib` where `.lib` is typically `.so` for shared libraries and `.a` for static libraries.

Unlike native SUNDIALS vector types, `NVECTOR_PETSC` does not provide macros to access its member variables. Note that `NVECTOR_PETSC` requires SUNDIALS to be built with MPI support.

### 6.9.1 NVECTOR\_PETSC functions

The `NVECTOR_PETSC` module defines implementations of all vector operations listed in §6.2 except for `N_VGetArrayPointer()` and `N_VSetArrayPointer()`. As such, this vector cannot be used with SUNDIALS Fortran interfaces. When access to raw vector data is needed, it is recommended to extract the PETSc vector first, and then use PETSc methods to access the data. Usage examples of `NVECTOR_PETSC` is provided in example programs for IDA.

The names of vector operations are obtained from those in §6.2, §6.2.2, §6.2.3, and §6.2.4 by appending the suffix `_Petsc` (e.g. `N_VDestroy_Petsc`). The module `NVECTOR_PETSC` provides the following additional user-callable routines:

`N_Vector N_VNewEmpty_Petsc`(MPI\_Comm comm, *sunindextype* local\_length, *sunindextype* global\_length, *SUNContext* sunctx)

This function creates a new PETSC `N_Vector` with the pointer to the wrapped PETSc vector set to `NULL`. It is used by the `N_VMake_Petsc` and `N_VClone_Petsc` implementations. It should be used only with great caution.

`N_Vector N_VMake_Petsc`(Vec \*pvec, *SUNContext* sunctx)

This function creates and allocates memory for an `NVECTOR_PETSC` wrapper with a user-provided PETSc vector. It does *not* allocate memory for the vector `pvec` itself.

Vec \***N\_VGetVector\_Petsc**(*N\_Vector* v)

This function returns a pointer to the underlying PETSc vector.

void **N\_VPrint\_Petsc**(*N\_Vector* v)

This function prints the global content of a wrapped PETSc vector to stdout.

void **N\_VPrintFile\_Petsc**(*N\_Vector* v, const char fname[])

This function prints the global content of a wrapped PETSc vector to fname.

By default all fused and vector array operations are disabled in the NVECTOR\_PETSC module. The following additional user-callable routines are provided to enable or disable fused and vector array operations for a specific vector. To ensure consistency across vectors it is recommended to first create a vector with *N\_VMake\_Petsc()*, enable/disable the desired operations for that vector with the functions below, and create any additional vectors from that vector using *N\_VClone()*. This guarantees the new vectors will have the same operations enabled/disabled as cloned vectors inherit the same enable/disable options as the vector they are cloned from while vectors created with *N\_VMake\_Petsc()* will have the default settings for the NVECTOR\_PETSC module.

*SUNErrCode* **N\_VEnableFusedOps\_Petsc**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) all fused and vector array operations in the PETSc vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearCombination\_Petsc**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear combination fused operation in the PETSc vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleAddMulti\_Petsc**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale and add a vector to multiple vectors fused operation in the PETSc vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableDotProdMulti\_Petsc**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the multiple dot products fused operation in the PETSc vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearSumVectorArray\_Petsc**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear sum operation for vector arrays in the PETSc vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleVectorArray\_Petsc**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale operation for vector arrays in the PETSc vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableConstVectorArray\_Petsc**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the const operation for vector arrays in the PETSc vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableWrmsNormVectorArray\_Petsc**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the WRMS norm operation for vector arrays in the PETSc vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableWrmsNormMaskVectorArray\_Petsc**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the masked WRMS norm operation for vector arrays in the PETSc vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleAddMultiVectorArray\_Petsc**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale and add a vector array to multiple vector arrays operation in the PETSc vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearCombinationVectorArray\_Petsc**(*N\_Vector* v, *sunboolean*type tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear combination operation for vector arrays in the PETSc vector. The return value is a *SUNErrCode*.

#### Notes

- When there is a need to access components of an *N\_Vector\_Petsc* v, it is recommended to extract the PETSc vector via `x_vec = N_VGetVector_Petsc(v)`; and then access components using appropriate PETSc functions.
- The functions *N\_VNewEmpty\_Petsc()* and *N\_VMake\_Petsc()*, set the field *own\_data* to SUNFALSE. The implementation of *N\_VDestroy()* will not attempt to free the pointer pvec for any *N\_Vector* with *own\_data* set to SUNFALSE. In such a case, it is the user's responsibility to deallocate the pvec pointer.
- To maximize efficiency, vector operations in the NVECTOR\_PETSC implementation that have more than one *N\_Vector* argument do not check for consistent internal representations of these vectors. It is the user's responsibility to ensure that such routines are called with *N\_Vector* arguments that were all created with the same internal representations.

## 6.10 The NVECTOR\_CUDA Module

The NVECTOR\_CUDA module is an NVECTOR implementation in the CUDA language. The module allows for SUNDIALS vector kernels to run on NVIDIA GPU devices. It is intended for users who are already familiar with CUDA and GPU programming. Building this vector module requires a CUDA compiler and, by extension, a C++ compiler. The vector content layout is as follows:

```
struct _N_VectorContent_Cuda
{
    sunindextype      length;
    sunbooleantype    own_helper;
    SUNMemory         host_data;
    SUNMemory         device_data;
    SUNCudaExecPolicy* stream_exec_policy;
    SUNCudaExecPolicy* reduce_exec_policy;
    SUNMemoryHelper    mem_helper;
    void*             priv; /* 'private' data */
};

typedef struct _N_VectorContent_Cuda *N_VectorContent_Cuda;
```

The content members are the vector length (size), boolean flags that indicate if the vector owns the execution policies and memory helper objects (i.e., it is in charge of freeing the objects), *SUNMemory* objects for the vector data on the host and device, pointers to execution policies that control how streaming and reduction kernels are launched, a *SUNMemoryHelper* for performing memory operations, and a private data structure which holds additional members that should not be accessed directly.

When instantiated with *N\_VNew\_Cuda()*, the underlying data will be allocated on both the host and the device. Alternatively, a user can provide host and device data arrays by using the *N\_VMake\_Cuda()* constructor. To use CUDA managed memory, the constructors *N\_VNewManaged\_Cuda()* and *N\_VMakeManaged\_Cuda()* are provided. Additionally, a user-defined *SUNMemoryHelper* for allocating/freeing data can be provided with the constructor *N\_VNewWithMemHelp\_Cuda()*. Details on each of these constructors are provided below.

To use the NVECTOR\_CUDA module, include `nvector_cuda.h` and link to the library `libsundials_nveccuda.lib`. The extension, `.lib`, is typically `.so` for shared libraries and `.a` for static libraries.

### 6.10.1 NVECTOR\_CUDA functions

Unlike other native SUNDIALS vector types, the NVECTOR\_CUDA module does not provide macros to access its member variables. Instead, user should use the accessor functions:

*sunrealtype* \***N\_VGetHostArrayPointer\_Cuda**(*N\_Vector* v)

This function returns pointer to the vector data on the host.

*sunrealtype* \***N\_VGetDeviceArrayPointer\_Cuda**(*N\_Vector* v)

This function returns pointer to the vector data on the device.

*sunbooleantype* **N\_VIsManagedMemory\_Cuda**(*N\_Vector* v)

This function returns a boolean flag indicating if the vector data array is in managed memory or not.

The NVECTOR\_CUDA module defines implementations of all standard vector operations defined in §6.2, §6.2.2, §6.2.3, and §6.2.4, except for *N\_VSetArrayPointer()*, and, if using unmanaged memory, *N\_VGetArrayPointer()*. As such, this vector can only be used with SUNDIALS direct solvers and preconditioners when using managed memory. The NVECTOR\_CUDA module provides separate functions to access data on the host and on the device for the unmanaged memory use case. It also provides methods for copying from the host to the device and vice versa. Usage examples of NVECTOR\_CUDA are provided in example programs for CVODE [43].

The names of vector operations are obtained from those in §6.2, §6.2.2, §6.2.3, and §6.2.4 by appending the suffix *\_Cuda* (e.g. *N\_VDestroy\_Cuda*). The module NVECTOR\_CUDA provides the following additional user-callable routines:

*N\_Vector* **N\_VNew\_Cuda**(*sunindextype* length, *SUNContext* sunctx)

This function creates and allocates memory for a CUDA *N\_Vector*. The vector data array is allocated on both the host and device.

*N\_Vector* **N\_VNewManaged\_Cuda**(*sunindextype* vec\_length, *SUNContext* sunctx)

This function creates and allocates memory for a CUDA *N\_Vector*. The vector data array is allocated in managed memory.

*N\_Vector* **N\_VNewWithMemHelp\_Cuda**(*sunindextype* length, *sunbooleantype* use\_managed\_mem,  
*SUNMemoryHelper* helper, *SUNContext* sunctx)

This function creates a new CUDA *N\_Vector* with a user-supplied *SUNMemoryHelper* for allocating/freeing memory.

*N\_Vector* **N\_VNewEmpty\_Cuda**(*sunindextype* vec\_length, *SUNContext* sunctx)

This function creates a new CUDA *N\_Vector* where the members of the content structure have not been allocated. This utility function is used by the other constructors to create a new vector.

*N\_Vector* **N\_VMake\_Cuda**(*sunindextype* vec\_length, *sunrealtype* \*h\_vdata, *sunrealtype* \*d\_vdata, *SUNContext* sunctx)

This function creates a CUDA *N\_Vector* with user-supplied vector data arrays for the host and the device.

*N\_Vector* **N\_VMakeManaged\_Cuda**(*sunindextype* vec\_length, *sunrealtype* \*vdata, *SUNContext* sunctx)

This function creates a CUDA *N\_Vector* with a user-supplied managed memory data array.

*N\_Vector* **N\_VMakeWithManagedAllocator\_Cuda**(*sunindextype* length, void \*(\*allocfn)(size\_t size), void  
(\*freefn)(void \*ptr))

This function creates a CUDA *N\_Vector* with a user-supplied memory allocator. It requires the user to provide a corresponding free function as well. The memory allocated by the allocator function must behave like CUDA managed memory.

The module NVECTOR\_CUDA also provides the following user-callable routines:

---

```
void N_VSetKernelExecPolicy_Cuda(N_Vector v, SUNCudaExecPolicy *stream_exec_policy,
                                SUNCudaExecPolicy *reduce_exec_policy)
```

This function sets the execution policies which control the kernel parameters utilized when launching the streaming and reduction CUDA kernels. By default the vector is setup to use the `SUNCudaThreadDirectExecPolicy` and `SUNCudaBlockReduceAtomicExecPolicy`. Any custom execution policy for reductions must ensure that the grid dimensions (number of thread blocks) is a multiple of the CUDA warp size (32). See §6.10.2 below for more information about the `SUNCudaExecPolicy` class. Providing NULL for an argument will result in the default policy being restored.

The input execution policies are cloned and, as such, may be freed after being attached to the desired vectors. A NULL input policy will reset the execution policy to the default setting.

---

**Note:** Note: All vectors used in a single instance of a SUNDIALS package must use the same execution policy. It is **strongly recommended** that this function is called immediately after constructing the vector, and any subsequent vector be created by cloning to ensure consistent execution policies across vectors

---

```
sunrealtype *N_VCopyToDevice_Cuda(N_Vector v)
```

This function copies host vector data to the device.

```
sunrealtype *N_VCopyFromDevice_Cuda(N_Vector v)
```

This function copies vector data from the device to the host.

```
void N_VPrint_Cuda(N_Vector v)
```

This function prints the content of a CUDA vector to stdout.

```
void N_VPrintFile_Cuda(N_Vector v, FILE *outfile)
```

This function prints the content of a CUDA vector to outfile.

By default all fused and vector array operations are disabled in the `NVECTOR_CUDA` module. The following additional user-callable routines are provided to enable or disable fused and vector array operations for a specific vector. To ensure consistency across vectors it is recommended to first create a vector with `N_VNew_Cuda()`, enable/disable the desired operations for that vector with the functions below, and create any additional vectors from that vector using `N_VClone()`. This guarantees the new vectors will have the same operations enabled/disabled as cloned vectors inherit the same enable/disable options as the vector they are cloned from while vectors created with `N_VNew_Cuda()` will have the default settings for the `NVECTOR_CUDA` module.

```
SUNErrCode N_VEnableFusedOps_Cuda(N_Vector v, sunboolean_t tf)
```

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) all fused and vector array operations in the CUDA vector. The return value is a `SUNErrCode`.

```
SUNErrCode N_VEnableLinearCombination_Cuda(N_Vector v, sunboolean_t tf)
```

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the linear combination fused operation in the CUDA vector. The return value is a `SUNErrCode`.

```
SUNErrCode N_VEnableScaleAddMulti_Cuda(N_Vector v, sunboolean_t tf)
```

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the scale and add a vector to multiple vectors fused operation in the CUDA vector. The return value is a `SUNErrCode`.

```
SUNErrCode N_VEnableDotProdMulti_Cuda(N_Vector v, sunboolean_t tf)
```

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the multiple dot products fused operation in the CUDA vector. The return value is a `SUNErrCode`.

```
SUNErrCode N_VEnableLinearSumVectorArray_Cuda(N_Vector v, sunboolean_t tf)
```

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the linear sum operation for vector arrays in the CUDA vector. The return value is a `SUNErrCode`.

*SUNErrCode* **N\_VEnableScaleVectorArray\_Cuda**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale operation for vector arrays in the CUDA vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableConstVectorArray\_Cuda**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the const operation for vector arrays in the CUDA vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableWrmsNormVectorArray\_Cuda**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the WRMS norm operation for vector arrays in the CUDA vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableWrmsNormMaskVectorArray\_Cuda**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the masked WRMS norm operation for vector arrays in the CUDA vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleAddMultiVectorArray\_Cuda**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale and add a vector array to multiple vector arrays operation in the CUDA vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearCombinationVectorArray\_Cuda**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear combination operation for vector arrays in the CUDA vector. The return value is a *SUNErrCode*.

#### Notes

- When there is a need to access components of an *N\_Vector\_Cuda*, v, it is recommended to use functions *N\_VGetDeviceArrayPointer\_Cuda()* or *N\_VGetHostArrayPointer\_Cuda()*. However, when using managed memory, the function *N\_VGetArrayPointer()* may also be used.
- To maximize efficiency, vector operations in the NVECTOR\_CUDA implementation that have more than one *N\_Vector* argument do not check for consistent internal representations of these vectors. It is the user's responsibility to ensure that such routines are called with *N\_Vector* arguments that were all created with the same internal representations.

### 6.10.2 The SUNCudaExecPolicy Class

In order to provide maximum flexibility to users, the CUDA kernel execution parameters used by kernels within SUNDIALS are defined by objects of the `sundials::cuda::ExecPolicy` abstract class type (this class can be accessed in the global namespace as `SUNCudaExecPolicy`). Thus, users may provide custom execution policies that fit the needs of their problem. The `SUNCudaExecPolicy` class is defined as

```
typedef sundials::cuda::ExecPolicy SUNCudaExecPolicy
```

where the `sundials::cuda::ExecPolicy` class is defined in the header file `sundials_cuda_policies.hpp`, as follows:

```
class sundials::cuda::ExecPolicy
```

```
    ExecPolicy(cudaStream_t stream = 0)
```

```
    virtual size_t gridSize(size_t numWorkUnits = 0, size_t blockDim = 0)
```

```
    virtual size_t blockSize(size_t numWorkUnits = 0, size_t gridDim = 0)
```

```
    virtual const cudaStream_t *stream() const
```



```

virtual ExecPolicy *clone() const

ExecPolicy *clone_new_stream(cudaStream_t stream) const

virtual bool atomic() const

virtual ~ExecPolicy()

```

To define a custom execution policy, a user simply needs to create a class that inherits from the abstract class and implements the methods. The SUNDIALS provided `sundials::cuda::ThreadDirectExecPolicy` (aka in the global namespace as `SUNCudaThreadDirectExecPolicy`) class is a good example of a what a custom execution policy may look like:

```

class ThreadDirectExecPolicy : public ExecPolicy
{
public:
    ThreadDirectExecPolicy(const size_t blockDim, cudaStream_t stream = 0)
        : blockDim_(blockDim), ExecPolicy(stream)
    {}

    ThreadDirectExecPolicy(const ThreadDirectExecPolicy& ex)
        : blockDim_(ex.blockDim_), ExecPolicy(ex.stream_)
    {}

    virtual size_t gridSize(size_t numWorkUnits = 0, size_t /*blockDim*/ = 0) const
    {
        /* ceil(n/m) = floor((n + m - 1) / m) */
        return (numWorkUnits + blockSize() - 1) / blockSize();
    }

    virtual size_t blockSize(size_t /*numWorkUnits*/ = 0, size_t /*gridDim*/ = 0) const
    {
        return blockDim_;
    }

    virtual ExecPolicy* clone() const
    {
        return static_cast<ExecPolicy*>(new ThreadDirectExecPolicy(*this));
    }

private:
    const size_t blockDim_;
};

```

In total, SUNDIALS provides 3 execution policies:

**SUNCudaThreadDirectExecPolicy**(const size\_t blockDim, const cudaStream\_t stream = 0)

Maps each CUDA thread to a work unit. The number of threads per block (blockDim) can be set to anything. The grid size will be calculated so that there are enough threads for one thread per element. If a CUDA stream is provided, it will be used to execute the kernel.

**SUNCudaGridStrideExecPolicy**(const size\_t blockDim, const size\_t gridDim, const cudaStream\_t stream = 0)

Is for kernels that use grid stride loops. The number of threads per block (blockDim) can be set to anything. The number of blocks (gridDim) can be set to anything. If a CUDA stream is provided, it will be used to execute the kernel.

**SUNCudaBlockReduceExecPolicy**(const size\_t blockDim, const cudaStream\_t stream = 0)

Is for kernels performing a reduction across individual thread blocks. The number of threads per block (blockDim) can be set to any valid multiple of the CUDA warp size. The grid size (gridDim) can be set to any value greater than 0. If it is set to 0, then the grid size will be chosen so that there is enough threads for one thread per work unit. If a CUDA stream is provided, it will be used to execute the kernel.

**SUNCudaBlockReduceAtomicExecPolicy**(const size\_t blockDim, const cudaStream\_t stream = 0)

Is for kernels performing a reduction across individual thread blocks using atomic operations. The number of threads per block (blockDim) can be set to any valid multiple of the CUDA warp size. The grid size (gridDim) can be set to any value greater than 0. If it is set to 0, then the grid size will be chosen so that there is enough threads for one thread per work unit. If a CUDA stream is provided, it will be used to execute the kernel.

For example, a policy that uses 128 threads per block and a user provided stream can be created like so:

```
cudaStream_t stream;
cudaStreamCreate(&stream);
SUNCudaThreadDirectExecPolicy thread_direct(128, stream);
```

These default policy objects can be reused for multiple SUNDIALS data structures (e.g. a *SUNMatrix* and an *N\_Vector*) since they do not hold any modifiable state information.

## 6.11 The NVECTOR\_HIP Module

The NVECTOR\_HIP module is an NVECTOR implementation using the AMD ROCm HIP library [2]. The module allows for SUNDIALS vector kernels to run on AMD or NVIDIA GPU devices. It is intended for users who are already familiar with HIP and GPU programming. Building this vector module requires the HIP-clang compiler. The vector content layout is as follows:

```
struct _N_VectorContent_Hip
{
    sunindextype      length;
    sunboolean_t      own_helper;
    SUNMemory         host_data;
    SUNMemory         device_data;
    SUNHipExecPolicy* stream_exec_policy;
    SUNHipExecPolicy* reduce_exec_policy;
    SUNMemoryHelper    mem_helper;
    void*             priv; /* 'private' data */
};

typedef struct _N_VectorContent_Hip *N_VectorContent_Hip;
```

The content members are the vector length (size), a boolean flag that signals if the vector owns the data (i.e. it is in charge of freeing the data), pointers to vector data on the host and the device, pointers to *SUNHipExecPolicy* implementations that control how the HIP kernels are launched for streaming and reduction vector kernels, and a private data structure which holds additional members that should not be accessed directly.

When instantiated with *N\_VNew\_Hip()*, the underlying data will be allocated on both the host and the device. Alternatively, a user can provide host and device data arrays by using the *N\_VMake\_Hip()* constructor. To use managed memory, the constructors *N\_VNewManaged\_Hip()* and *N\_VMakeManaged\_Hip()* are provided. Additionally, a user-defined *SUNMemoryHelper* for allocating/freeing data can be provided with the constructor *N\_VNewWithMemHelp\_Hip()*. Details on each of these constructors are provided below.



To use the NVECTOR\_HIP module, include `nvector_hip.h` and link to the library `libsundials_nvechip.lib`. The extension, `.lib`, is typically `.so` for shared libraries and `.a` for static libraries.

### 6.11.1 NVECTOR\_HIP functions

Unlike other native SUNDIALS vector types, the NVECTOR\_HIP module does not provide macros to access its member variables. Instead, user should use the accessor functions:

*sunrealtype* \***N\_VGetHostArrayPointer\_Hip**(*N\_Vector* v)

This function returns pointer to the vector data on the host.

*sunrealtype* \***N\_VGetDeviceArrayPointer\_Hip**(*N\_Vector* v)

This function returns pointer to the vector data on the device.

*sunbooleantype* **N\_VIsManagedMemory\_Hip**(*N\_Vector* v)

This function returns a boolean flag indicating if the vector data array is in managed memory or not.

The NVECTOR\_HIP module defines implementations of all standard vector operations defined in §6.2, §6.2.2, §6.2.3, and §6.2.4, except for `N_VSetArrayPointer()`. The names of vector operations are obtained from those in §6.2, §6.2.2, §6.2.3, and §6.2.4 by appending the suffix `_Hip` (e.g. `N_VDestroy_Hip`). The module NVECTOR\_HIP provides the following additional user-callable routines:

*N\_Vector* **N\_VNew\_Hip**(*sunindextype* length, *SUNContext* sunctx)

This function creates and allocates memory for a HIP *N\_Vector*. The vector data array is allocated on both the host and device.

*N\_Vector* **N\_VNewManaged\_Hip**(*sunindextype* vec\_length, *SUNContext* sunctx)

This function creates and allocates memory for a HIP *N\_Vector*. The vector data array is allocated in managed memory.

*N\_Vector* **N\_VNewWithMemHelp\_Hip**(*sunindextype* length, *sunbooleantype* use\_managed\_mem, *SUNMemoryHelper* helper, *SUNContext* sunctx)

This function creates a new HIP *N\_Vector* with a user-supplied *SUNMemoryHelper* for allocating/freeing memory.

*N\_Vector* **N\_VNewEmpty\_Hip**(*sunindextype* vec\_length, *SUNContext* sunctx)

This function creates a new HIP *N\_Vector* where the members of the content structure have not been allocated. This utility function is used by the other constructors to create a new vector.

*N\_Vector* **N\_VMake\_Hip**(*sunindextype* vec\_length, *sunrealtype* \*h\_vdata, *sunrealtype* \*d\_vdata, *SUNContext* sunctx)

This function creates a HIP *N\_Vector* with user-supplied vector data arrays for the host and the device.

*N\_Vector* **N\_VMakeManaged\_Hip**(*sunindextype* vec\_length, *sunrealtype* \*vdata, *SUNContext* sunctx)

This function creates a HIP *N\_Vector* with a user-supplied managed memory data array.

The module NVECTOR\_HIP also provides the following user-callable routines:

void **N\_VSetKernelExecPolicy\_Hip**(*N\_Vector* v, *SUNHipExecPolicy* \*stream\_exec\_policy, *SUNHipExecPolicy* \*reduce\_exec\_policy)

This function sets the execution policies which control the kernel parameters utilized when launching the streaming and reduction HIP kernels. By default the vector is setup to use the `SUNHipThreadDirectExecPolicy()` and `SUNHipBlockReduceExecPolicy()`. Any custom execution policy for reductions must ensure that the grid dimensions (number of thread blocks) is a multiple of the HIP warp size (32 for NVIDIA GPUs, 64 for AMD GPUs). See §6.11.2 below for more information about the `SUNHipExecPolicy` class. Providing NULL for an argument will result in the default policy being restored.

The input execution policies are cloned and, as such, may be freed after being attached to the desired vectors. A NULL input policy will reset the execution policy to the default setting.

---

**Note:** Note: All vectors used in a single instance of a SUNDIALS package must use the same execution policy. It is **strongly recommended** that this function is called immediately after constructing the vector, and any subsequent vector be created by cloning to ensure consistent execution policies across vectors\*

---

*sunrealtype* \*N\_VCopyToDevice\_Hip(*N\_Vector* v)

This function copies host vector data to the device.

*sunrealtype* \*N\_VCopyFromDevice\_Hip(*N\_Vector* v)

This function copies vector data from the device to the host.

void N\_VPrint\_Hip(*N\_Vector* v)

This function prints the content of a HIP vector to stdout.

void N\_VPrintFile\_Hip(*N\_Vector* v, FILE \*outfile)

This function prints the content of a HIP vector to outfile.

By default all fused and vector array operations are disabled in the NVECTOR\_HIP module. The following additional user-callable routines are provided to enable or disable fused and vector array operations for a specific vector. To ensure consistency across vectors it is recommended to first create a vector with *N\_VNew\_Hip()*, enable/disable the desired operations for that vector with the functions below, and create any additional vectors from that vector using *N\_VClone()*. This guarantees the new vectors will have the same operations enabled/disabled as cloned vectors inherit the same enable/disable options as the vector they are cloned from while vectors created with *N\_VNew\_Hip()* will have the default settings for the NVECTOR\_HIP module.

*SUNErrCode* N\_VEnableFusedOps\_Hip(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) all fused and vector array operations in the HIP vector. The return value is a *SUNErrCode*.

*SUNErrCode* N\_VEnableLinearCombination\_Hip(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear combination fused operation in the HIP vector. The return value is a *SUNErrCode*.

*SUNErrCode* N\_VEnableScaleAddMulti\_Hip(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale and add a vector to multiple vectors fused operation in the HIP vector. The return value is a *SUNErrCode*.

*SUNErrCode* N\_VEnableDotProdMulti\_Hip(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the multiple dot products fused operation in the HIP vector. The return value is a *SUNErrCode*.

*SUNErrCode* N\_VEnableLinearSumVectorArray\_Hip(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear sum operation for vector arrays in the HIP vector. The return value is a *SUNErrCode*.

*SUNErrCode* N\_VEnableScaleVectorArray\_Hip(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale operation for vector arrays in the HIP vector. The return value is a *SUNErrCode*.

*SUNErrCode* N\_VEnableConstVectorArray\_Hip(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the const operation for vector arrays in the HIP vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableWrmsNormVectorArray\_Hip**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the WRMS norm operation for vector arrays in the HIP vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableWrmsNormMaskVectorArray\_Hip**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the masked WRMS norm operation for vector arrays in the HIP vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleAddMultiVectorArray\_Hip**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale and add a vector array to multiple vector arrays operation in the HIP vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearCombinationVectorArray\_Hip**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear combination operation for vector arrays in the HIP vector. The return value is a *SUNErrCode*.

#### Notes

- When there is a need to access components of an *N\_Vector\_Hip*, v, it is recommended to use functions *N\_VGetDeviceArrayPointer\_Hip()* or *N\_VGetHostArrayPointer\_Hip()*. However, when using managed memory, the function *N\_VGetArrayPointer()* may also be used.
- To maximize efficiency, vector operations in the NVECTOR\_HIP implementation that have more than one *N\_Vector* argument do not check for consistent internal representations of these vectors. It is the user's responsibility to ensure that such routines are called with *N\_Vector* arguments that were all created with the same internal representations.

### 6.11.2 The SUNHipExecPolicy Class

In order to provide maximum flexibility to users, the HIP kernel execution parameters used by kernels within SUNDIALS are defined by objects of the `sundials::hip::ExecPolicy` abstract class type (this class can be accessed in the global namespace as `SUNHipExecPolicy`). Thus, users may provide custom execution policies that fit the needs of their problem. The `SUNHipExecPolicy` class is defined as

```
typedef sundials::hip::ExecPolicy SUNHipExecPolicy
```

where the `sundials::hip::ExecPolicy` class is defined in the header file `sundials_hip_policies.hpp`, as follows:

```
class sundials::hip::ExecPolicy
```

```
    ExecPolicy(hipStream_t stream = 0)

    virtual size_t gridSize(size_t numWorkUnits = 0, size_t blockDim = 0)

    virtual size_t blockSize(size_t numWorkUnits = 0, size_t gridDim = 0)

    virtual const hipStream_t *stream() const

    virtual ExecPolicy *clone() const

    ExecPolicy *clone_new_stream(hipStream_t stream) const

    virtual bool atomic() const
```

virtual ~ExecPolicy()

To define a custom execution policy, a user simply needs to create a class that inherits from the abstract class and implements the methods. The SUNDIALS provided `sundials::hip::ThreadDirectExecPolicy` (aka in the global namespace as `SUNHipThreadDirectExecPolicy`) class is a good example of a what a custom execution policy may look like:

```
class ThreadDirectExecPolicy : public ExecPolicy
{
public:
    ThreadDirectExecPolicy(const size_t blockDim, hipStream_t stream = 0)
        : blockDim_(blockDim), ExecPolicy(stream)
    {}

    ThreadDirectExecPolicy(const ThreadDirectExecPolicy& ex)
        : blockDim_(ex.blockDim_), ExecPolicy(ex.stream_)
    {}

    virtual size_t gridSize(size_t numWorkUnits = 0, size_t /*blockDim*/ = 0) const
    {
        /* ceil(n/m) = floor((n + m - 1) / m) */
        return (numWorkUnits + blockSize() - 1) / blockSize();
    }

    virtual size_t blockSize(size_t /*numWorkUnits*/ = 0, size_t /*gridDim*/ = 0) const
    {
        return blockDim_;
    }

    virtual ExecPolicy* clone() const
    {
        return static_cast<ExecPolicy*>(new ThreadDirectExecPolicy(*this));
    }

private:
    const size_t blockDim_;
};
```

In total, SUNDIALS provides 4 execution policies:

**SUNHipThreadDirectExecPolicy**(const size\_t blockDim, const hipStream\_t stream = 0)

Maps each HIP thread to a work unit. The number of threads per block (blockDim) can be set to anything. The grid size will be calculated so that there are enough threads for one thread per element. If a HIP stream is provided, it will be used to execute the kernel.

**SUNHipGridStrideExecPolicy**(const size\_t blockDim, const size\_t gridDim, const hipStream\_t stream = 0)

Is for kernels that use grid stride loops. The number of threads per block (blockDim) can be set to anything. The number of blocks (gridDim) can be set to anything. If a HIP stream is provided, it will be used to execute the kernel.

**SUNHipBlockReduceExecPolicy**(const size\_t blockDim, const hipStream\_t stream = 0)

Is for kernels performing a reduction across individual thread blocks. The number of threads per block (blockDim) can be set to any valid multiple of the HIP warp size. The grid size (gridDim) can be set to any value greater than 0. If it is set to 0, then the grid size will be chosen so that there is

enough threads for one thread per work unit. If a HIP stream is provided, it will be used to execute the kernel.

**SUNHipBlockReduceAtomicExecPolicy**(const size\_t blockDim, const hipStream\_t stream = 0)

Is for kernels performing a reduction across individual thread blocks using atomic operations. The number of threads per block (blockDim) can be set to any valid multiple of the HIP warp size. The grid size (gridDim) can be set to any value greater than 0. If it is set to 0, then the grid size will be chosen so that there is enough threads for one thread per work unit. If a HIP stream is provided, it will be used to execute the kernel.

For example, a policy that uses 128 threads per block and a user provided stream can be created like so:

```
hipStream_t stream;
hipStreamCreate(&stream);
SUNHipThreadDirectExecPolicy thread_direct(128, stream);
```

These default policy objects can be reused for multiple SUNDIALS data structures (e.g. a *SUNMatrix* and an *N\_Vector*) since they do not hold any modifiable state information.

## 6.12 The NVECTOR\_SYCL Module

The NVECTOR\_SYCL module is an experimental NVECTOR implementation using the SYCL abstraction layer. At present the only supported SYCL compiler is the DPC++ (Intel oneAPI) compiler. This module allows for SUNDIALS vector kernels to run on Intel GPU devices. The module is intended for users who are already familiar with SYCL and GPU programming.

The vector content layout is as follows:

```
struct _N_VectorContent_Sycl
{
    sunindextype      length;
    sunboolean_t      own_helper;
    SUNMemory         host_data;
    SUNMemory         device_data;
    SUNSyclExecPolicy* stream_exec_policy;
    SUNSyclExecPolicy* reduce_exec_policy;
    SUNMemoryHelper    mem_helper;
    sycl::queue*       queue;
    void*              priv; /* 'private' data */
};

typedef struct _N_VectorContent_Sycl *N_VectorContent_Sycl;
```

The content members are the vector length (size), boolean flags that indicate if the vector owns the execution policies and memory helper objects (i.e., it is in charge of freeing the objects), *SUNMemory* objects for the vector data on the host and device, pointers to execution policies that control how streaming and reduction kernels are launched, a *SUNMemoryHelper* for performing memory operations, the SYCL queue, and a private data structure which holds additional members that should not be accessed directly.

When instantiated with *N\_VNew\_Sycl()*, the underlying data will be allocated on both the host and the device. Alternatively, a user can provide host and device data arrays by using the *N\_VMake\_Sycl()* constructor. To use managed (shared) memory, the constructors *N\_VNewManaged\_Sycl()* and *N\_VMakeManaged\_Sycl()* are provided. Additionally, a user-defined *SUNMemoryHelper* for allocating/freeing data can be provided with the constructor *N\_VNewWithMemHelp\_Sycl()*. Details on each of these constructors are provided below.

The header file to include when using this is `nvector_sycl.h`. The installed module library to link to is `libsundials_nvecscyl.lib`. The extension `.lib` is typically `.so` for shared libraries `.a` for static libraries.

### 6.12.1 NVECTOR\_SYCL functions

The NVECTOR\_SYCL module implementations of all vector operations listed in §6.2, §6.2.2, §6.2.3, and §6.2.4, except for `N_VDotProdMulti()`, `N_VWrmsNormVectorArray()`, `N_VWrmsNormMaskVectorArray()` as support for arrays of reduction vectors is not yet supported. These functions will be added to the NVECTOR\_SYCL implementation in the future. The names of vector operations are obtained from those in the aforementioned sections by appending the suffix `_Sycl` (e.g., `N_VDestroy_Sycl`).

Additionally, the NVECTOR\_SYCL module provides the following user-callable constructors for creating a new NVECTOR\_SYCL:

`N_Vector N_VNew_Sycl`(sunindextype vec\_length, `sycl::queue *Q`, `SUNContext sunctx`)

This function creates and allocates memory for an NVECTOR\_SYCL. Vector data arrays are allocated on both the host and the device associated with the input queue. All operation are launched in the provided queue.

`N_Vector N_VNewManaged_Sycl`(sunindextype vec\_length, `sycl::queue *Q`, `SUNContext sunctx`)

This function creates and allocates memory for a NVECTOR\_SYCL. The vector data array is allocated in managed (shared) memory using the input queue. All operation are launched in the provided queue.

`N_Vector N_VMake_Sycl`(sunindextype length, `sunrealtype *h_vdata`, `sunrealtype *d_vdata`, `sycl::queue *Q`, `SUNContext sunctx`)

This function creates an NVECTOR\_SYCL with user-supplied host and device data arrays. This function does not allocate memory for data itself. All operation are launched in the provided queue.

`N_Vector N_VMakeManaged_Sycl`(sunindextype length, `sunrealtype *vdata`, `sycl::queue *Q`, `SUNContext sunctx`)

This function creates an NVECTOR\_SYCL with a user-supplied managed (shared) data array. This function does not allocate memory for data itself. All operation are launched in the provided queue.

`N_Vector N_VNewWithMemHelp_Sycl`(sunindextype length, `sunbooleantype use_managed_mem`, `SUNMemoryHelper helper`, `sycl::queue *Q`, `SUNContext sunctx`)

This function creates an NVECTOR\_SYCL with a user-supplied `SUNMemoryHelper` for allocating/freeing memory. All operation are launched in the provided queue.

`N_Vector N_VNewEmpty_Sycl`()

This function creates a new `N_Vector` where the members of the content structure have not been allocated. This utility function is used by the other constructors to create a new vector.

The following user-callable functions are provided for accessing the vector data arrays on the host and device and copying data between the two memory spaces. Note the generic NVECTOR operations `N_VGetArrayPointer()` and `N_VSetArrayPointer()` are mapped to the corresponding `HostArray` functions given below. To ensure memory coherency, a user will need to call the `CopyTo` or `CopyFrom` functions as necessary to transfer data between the host and device, unless managed (shared) memory is used.

`sunrealtype *N_VGetHostArrayPointer_Sycl`(`N_Vector v`)

This function returns a pointer to the vector host data array.

`sunrealtype *N_VGetDeviceArrayPointer_Sycl`(`N_Vector v`)

This function returns a pointer to the vector device data array.

`void N_VSetHostArrayPointer_Sycl`(`sunrealtype *h_vdata`, `N_Vector v`)

This function sets the host array pointer in the vector `v`.



void **N\_VSetDeviceArrayPointer\_Sycl**(sunrealtype \*d\_vdata, N\_Vector v)

This function sets the device array pointer in the vector v.

void **N\_VCopyToDevice\_Sycl**(N\_Vector v)

This function copies host vector data to the device.

void **N\_VCopyFromDevice\_Sycl**(N\_Vector v)

This function copies vector data from the device to the host.

sunboolean\_t **N\_VIsManagedMemory\_Sycl**(N\_Vector v)

This function returns SUNTRUE if the vector data is allocated as managed (shared) memory otherwise it returns SUNFALSE.

The following user-callable function is provided to set the execution policies for how SYCL kernels are launched on a device.

SUNErrCode **N\_VSetKernelExecPolicy\_Sycl**(N\_Vector v, *SUNSyclExecPolicy* \*stream\_exec\_policy, *SUNSyclExecPolicy* \*reduce\_exec\_policy)

This function sets the execution policies which control the kernel parameters utilized when launching the streaming and reduction kernels. By default the vector is setup to use the *SUNSyclThreadDirectExecPolicy()* and *SUNSyclBlockReduceExecPolicy()*. See §6.12.2 below for more information about the *SUNSyclExecPolicy* class.

The input execution policies are cloned and, as such, may be freed after being attached to the desired vectors. A NULL input policy will reset the execution policy to the default setting.

---

**Note:** All vectors used in a single instance of a SUNDIALS package must use the same execution policy. It is **strongly recommended** that this function is called immediately after constructing the vector, and any subsequent vector be created by cloning to ensure consistent execution policies across vectors.

---

The following user-callable functions are provided to print the host vector data array. Unless managed memory is used, a user may need to call *N\_VCopyFromDevice\_Sycl()* to ensure consistency between the host and device array.

void **N\_VPrint\_Sycl**(N\_Vector v)

This function prints the host data array to stdout.

void **N\_VPrintFile\_Sycl**(N\_Vector v, FILE \*outfile)

This function prints the host data array to outfile.

By default all fused and vector array operations are disabled in the NVECTOR\_SYCL module. The following additional user-callable routines are provided to enable or disable fused and vector array operations for a specific vector. To ensure consistency across vectors it is recommended to first create a vector with one of the above constructors, enable/disable the desired operations on that vector with the functions below, and then use this vector in conjunction with *N\_VClone()* to create any additional vectors. This guarantees the new vectors will have the same operations enabled/disabled as cloned vectors inherit the same enable/disable options as the vector they are cloned from while vectors created by any of the constructors above will have the default settings for the NVECTOR\_SYCL module.

SUNErrCode **N\_VEnableFusedOps\_Sycl**(N\_Vector v, sunboolean\_t tf)

This function enables (SUNTRUE) or disables (SUNFALSE) all fused and vector array operations in the SYCL vector. The return value is a *SUNErrCode*.

SUNErrCode **N\_VEnableLinearCombination\_Sycl**(N\_Vector v, sunboolean\_t tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear combination fused operation in the SYCL vector. The return value is a *SUNErrCode*.

SUNErrCode **N\_VEnableScaleAddMulti\_Sycl**(N\_Vector v, sunbooleantype tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale and add a vector to multiple vectors fused operation in the SYCL vector. The return value is a [SUNErrCode](#).

SUNErrCode **N\_VEnableLinearSumVectorArray\_Sycl**(N\_Vector v, sunbooleantype tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear sum operation for vector arrays in the SYCL vector. The return value is a [SUNErrCode](#).

SUNErrCode **N\_VEnableScaleVectorArray\_Sycl**(N\_Vector v, sunbooleantype tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale operation for vector arrays in the SYCL vector. The return value is a [SUNErrCode](#).

SUNErrCode **N\_VEnableConstVectorArray\_Sycl**(N\_Vector v, sunbooleantype tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the const operation for vector arrays in the SYCL vector. The return value is a [SUNErrCode](#).

SUNErrCode **N\_VEnableScaleAddMultiVectorArray\_Sycl**(N\_Vector v, sunbooleantype tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale and add a vector array to multiple vector arrays operation in the SYCL vector. The return value is a [SUNErrCode](#).

SUNErrCode **N\_VEnableLinearCombinationVectorArray\_Sycl**(N\_Vector v, sunbooleantype tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear combination operation for vector arrays in the SYCL vector. The return value is a [SUNErrCode](#).

#### Notes

- When there is a need to access components of an NVECTOR\_SYCL, v, it is recommended to use [N\\_VGetDeviceArrayPointer\(\)](#) to access the device array or [N\\_VGetArrayPointer\(\)](#) for the host array. When using managed (shared) memory, either function may be used. To ensure memory coherency, a user may need to call the CopyTo or CopyFrom functions as necessary to transfer data between the host and device, unless managed (shared) memory is used.
- To maximize efficiency, vector operations in the NVECTOR\_SYCL implementation that have more than one N\_Vector argument do not check for consistent internal representations of these vectors. It is the user's responsibility to ensure that such routines are called with N\_Vector arguments that were all created with the same internal representations.

### 6.12.2 The SUNSyclExecPolicy Class

In order to provide maximum flexibility to users, the SYCL kernel execution parameters used by kernels within SUNDIALS are defined by objects of the `sundials::sycl::ExecPolicy` abstract class type (this class can be accessed in the global namespace as `SUNSyclExecPolicy`). Thus, users may provide custom execution policies that fit the needs of their problem. The `SUNSyclExecPolicy` class is defined as

```
typedef sundials::sycl::ExecPolicy SUNSyclExecPolicy
```

where the `sundials::sycl::ExecPolicy` class is defined in the header file `sundials_sycl_policies.hpp`, as follows:

```
class sundials::sycl::ExecPolicy
```

```
    virtual size_t gridSize(size_t numWorkUnits = 0, size_t blockDim = 0)
```

```
    virtual size_t blockSize(size_t numWorkUnits = 0, size_t gridDim = 0)
```

```
    virtual ExecPolicy *clone() const
```



virtual **~ExecPolicy()**

For consistency the function names and behavior mirror the execution policies for the CUDA and HIP vectors. In the SYCL case the `blockSize` is the local work-group range in a one-dimensional `nd_range` (threads per group). The `gridSize` is the number of local work groups so the global work-group range in a one-dimensional `nd_range` is `blockSize * gridSize` (total number of threads). All vector kernels are written with a many-to-one mapping where work units (vector elements) are mapped in a round-robin manner across the global range. As such, the `blockSize` and `gridSize` can be set to any positive value.

To define a custom execution policy, a user simply needs to create a class that inherits from the abstract class and implements the methods. The SUNDIALS provided `sundials::sycl::ThreadDirectExecPolicy` (aka in the global namespace as `SUNSYCLThreadDirectExecPolicy`) class is a good example of what a custom execution policy may look like:

```
class ThreadDirectExecPolicy : public ExecPolicy
{
public:
    ThreadDirectExecPolicy(const size_t blockDim)
        : blockDim_(blockDim)
    {}

    ThreadDirectExecPolicy(const ThreadDirectExecPolicy& ex)
        : blockDim_(ex.blockDim_)
    {}

    virtual size_t gridSize(size_t numWorkUnits = 0, size_t blockDim = 0) const
    {
        return (numWorkUnits + blockSize() - 1) / blockSize();
    }

    virtual size_t blockSize(size_t numWorkUnits = 0, size_t gridDim = 0) const
    {
        return blockDim_;
    }

    virtual ExecPolicy* clone() const
    {
        return static_cast<ExecPolicy*>(new ThreadDirectExecPolicy(*this));
    }

private:
    const size_t blockDim_;
};
```

SUNDIALS provides the following execution policies:

**SUNSYCLThreadDirectExecPolicy**(const size\_t blockDim)

Is for kernels performing streaming operations and maps each work unit (vector element) to a work-item (thread). Based on the local work-group range (number of threads per group, `blockSize`) the number of local work-groups (`gridSize`) is computed so there are enough work-items in the global work-group range ( total number of threads, `blockSize * gridSize`) for one work unit per work-item (thread).

**SUNSYCLGridStrideExecPolicy**(const size\_t blockDim, const size\_t gridDim)

Is for kernels performing streaming operations and maps each work unit (vector element) to a work-item (thread) in a round-robin manner so the local work-group range (number of threads per group,

blockSize) and the number of local work-groups (gridSize) can be set to any positive value. In this case the global work-group range (total number of threads, blockSize \* gridSize) may be less than the number of work units (vector elements).

#### **SUNSYCLBlockReduceExecPolicy**(const size\_t blockDim)

Is for kernels performing a reduction, the local work-group range (number of threads per group, blockSize) and the number of local work-groups (gridSize) can be set to any positive value or the gridSize may be set to 0 in which case the global range is chosen so that there are enough threads for at most two work units per work-item.

By default the NVECTOR\_SYCL module uses the SUNSYCLThreadDirectExecPolicy and SUNSYCLBlockReduceExecPolicy where the default blockDim is determined by querying the device for the max\_work\_group\_size. User may specify different policies by constructing a new SyclExecPolicy and attaching it with `N_VSetKernelExecPolicy_Sycl()`. For example, a policy that uses 128 work-items (threads) per group can be created and attached like so:

```
N_Vector v = N_VNew_Sycl(length, SUNContext sunctx);
SUNSYCLThreadDirectExecPolicy thread_direct(128);
SUNSYCLBlockReduceExecPolicy block_reduce(128);
flag = N_VSetKernelExecPolicy_Sycl(v, &thread_direct, &block_reduce);
```

These default policy objects can be reused for multiple SUNDIALS data structures (e.g. a *SUNMatrix* and an *N\_Vector*) since they do not hold any modifiable state information.

## 6.13 The NVECTOR\_RAJA Module

The NVECTOR\_RAJA module is an experimental NVECTOR implementation using the RAJA hardware abstraction layer. In this implementation, RAJA allows for SUNDIALS vector kernels to run on AMD, NVIDIA, or Intel GPU devices. The module is intended for users who are already familiar with RAJA and GPU programming. Building this vector module requires a C++11 compliant compiler and either the NVIDIA CUDA programming environment, the AMD ROCm HIP programming environment, or a compiler that supports the SYCL abstraction layer. When using the AMD ROCm HIP environment, the HIP-clang compiler must be utilized. Users can select which backend to compile with by setting the SUNDIALS\_RAJA\_BACKENDS CMake variable to either CUDA, HIP, or SYCL. Besides the CUDA, HIP, and SYCL backends, RAJA has other backends such as serial, OpenMP, and OpenACC. These backends are not used in this SUNDIALS release.

The vector content layout is as follows:

```
struct _N_VectorContent_Raja
{
    sunindextype length;
    sunbooleantype own_data;
    sunrealtype* host_data;
    sunrealtype* device_data;
    void* priv; /* 'private' data */
};
```

The content members are the vector length (size), a boolean flag that signals if the vector owns the data (i.e., it is in charge of freeing the data), pointers to vector data on the host and the device, and a private data structure which holds the memory management type, which should not be accessed directly.

When instantiated with `N_VNew_Raja()`, the underlying data will be allocated on both the host and the device. Alternatively, a user can provide host and device data arrays by using the `N_VMake_Raja()` constructor. To use managed memory, the constructors `N_VNewManaged_Raja()` and `N_VMakeManaged_Raja()` are provided. Details on each of these constructors are provided below.

The header file to include when using this is `nvector_raja.h`. The installed module library to link to is `libsundials_nveccudaraja.lib` when using the CUDA backend, `libsundials_nvechipraja.lib` when using the HIP backend, and `libsundials_nvecsyclraja.lib` when using the SYCL backend. The extension `.lib` is typically `.so` for shared libraries `.a` for static libraries.

### 6.13.1 NVECTOR\_RAJA functions

Unlike other native SUNDIALS vector types, the NVECTOR\_RAJA module does not provide macros to access its member variables. Instead, user should use the accessor functions:

*sunrealtype* \***N\_VGetHostArrayPointer\_Raja**(*N\_Vector* v)

This function returns pointer to the vector data on the host.

*sunrealtype* \***N\_VGetDeviceArrayPointer\_Raja**(*N\_Vector* v)

This function returns pointer to the vector data on the device.

*sunbooleantype* **N\_VIsManagedMemory\_Raja**(*N\_Vector* v)

This function returns a boolean flag indicating if the vector data is allocated in managed memory or not.

The NVECTOR\_RAJA module defines the implementations of all vector operations listed in §6.2, §6.2.2, §6.2.3, and §6.2.4, except for `N_VDotProdMulti()`, `N_VWrmsNormVectorArray()`, and `N_VWrmsNormMaskVectorArray()` as support for arrays of reduction vectors is not yet supported in RAJA. These functions will be added to the NVECTOR\_RAJA implementation in the future. Additionally, the operations `N_VGetArrayPointer()` and `N_VSetArrayPointer()` are not implemented by the RAJA vector. As such, this vector cannot be used with SUNDIALS direct solvers and preconditioners. The NVECTOR\_RAJA module provides separate functions to access data on the host and on the device. It also provides methods for copying from the host to the device and vice versa. Usage examples of NVECTOR\_RAJA are provided in some example programs for CVODE [43].

The names of vector operations are obtained from those in §6.2, §6.2.2, §6.2.3, and §6.2.4 by appending the suffix `_Raja` (e.g. `N_VDestroy_Raja`). The module NVECTOR\_RAJA provides the following additional user-callable routines:

*N\_Vector* **N\_VNew\_Raja**(*sunindextype* vec\_length, *SUNContext* sunctx)

This function creates and allocates memory for a RAJA *N\_Vector*. The memory is allocated on both the host and the device. Its only argument is the vector length.

*N\_Vector* **N\_VNewManaged\_Raja**(*sunindextype* vec\_length, *SUNContext* sunctx)

This function creates and allocates memory for a RAJA *N\_Vector*. The vector data array is allocated in managed memory.

*N\_Vector* **N\_VMake\_Raja**(*sunindextype* length, *sunrealtype* \*h\_data, *sunrealtype* \*v\_data, *SUNContext* sunctx)

This function creates an NVECTOR\_RAJA with user-supplied host and device data arrays. This function does not allocate memory for data itself.

*N\_Vector* **N\_VMakeManaged\_Raja**(*sunindextype* length, *sunrealtype* \*vdata, *SUNContext* sunctx)

This function creates an NVECTOR\_RAJA with a user-supplied managed memory data array. This function does not allocate memory for data itself.

*N\_Vector* **N\_VNewWithMemHelp\_Raja**(*sunindextype* length, *sunbooleantype* use\_managed\_mem, *SUNMemoryHelper* helper, *SUNContext* sunctx)

This function creates an NVECTOR\_RAJA with a user-supplied *SUNMemoryHelper* for allocating/freeing memory.

*N\_Vector* **N\_VNewEmpty\_Raja**()

This function creates a new *N\_Vector* where the members of the content structure have not been allocated. This utility function is used by the other constructors to create a new vector.

void **N\_VCopyToDevice\_Raja**(*N\_Vector* v)

This function copies host vector data to the device.

void **N\_VCopyFromDevice\_Raja**(*N\_Vector* v)

This function copies vector data from the device to the host.

void **N\_VPrint\_Raja**(*N\_Vector* v)

This function prints the content of a RAJA vector to stdout.

void **N\_VPrintFile\_Raja**(*N\_Vector* v, FILE \*outfile)

This function prints the content of a RAJA vector to outfile.

By default all fused and vector array operations are disabled in the NVECTOR\_RAJA module. The following additional user-callable routines are provided to enable or disable fused and vector array operations for a specific vector. To ensure consistency across vectors it is recommended to first create a vector with *N\_VNew\_Raja()*, enable/disable the desired operations for that vector with the functions below, and create any additional vectors from that vector using *N\_VClone()*. This guarantees the new vectors will have the same operations enabled/disabled as cloned vectors inherit the same enable/disable options as the vector they are cloned from while vectors created with *N\_VNew\_Raja()* will have the default settings for the NVECTOR\_RAJA module.

*SUNErrCode* **N\_VEnableFusedOps\_Raja**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) all fused and vector array operations in the RAJA vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearCombination\_Raja**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear combination fused operation in the RAJA vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleAddMulti\_Raja**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale and add a vector to multiple vectors fused operation in the RAJA vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearSumVectorArray\_Raja**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear sum operation for vector arrays in the RAJA vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleVectorArray\_Raja**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale operation for vector arrays in the RAJA vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableConstVectorArray\_Raja**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the const operation for vector arrays in the RAJA vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleAddMultiVectorArray\_Raja**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale and add a vector array to multiple vector arrays operation in the RAJA vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearCombinationVectorArray\_Raja**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear combination operation for vector arrays in the RAJA vector. The return value is a *SUNErrCode*.

## Notes

- When there is a need to access components of an NVECTOR\_RAJA vector, it is recommended to use functions *N\_VGetDeviceArrayPointer\_Raja()* or *N\_VGetHostArrayPointer\_Raja()*. However, when using managed memory, the function *N\_VGetArrayPointer()* may also be used.

- To maximize efficiency, vector operations in the NVECTOR\_RAJA implementation that have more than one `N_Vector` argument do not check for consistent internal representations of these vectors. It is the user's responsibility to ensure that such routines are called with `N_Vector` arguments that were all created with the same internal representations.

## 6.14 The NVECTOR\_KOKKOS Module

New in version 6.4.0.

The NVECTOR\_KOKKOS `N_Vector` implementation provides a vector data structure using Kokkos [32, 64] to support a variety of backends including serial, OpenMP, CUDA, HIP, and SYCL. Since Kokkos is a modern C++ library, the module is also written in modern C++ (it requires C++14) as a header only library. To utilize this `N_Vector` users will need to include `nvector/nvector_kokkos.hpp`. More instructions on building SUNDIALS with Kokkos enabled are given in §12.4. For instructions on building and using Kokkos, refer to the [Kokkos](#) documentation.

### 6.14.1 Using NVECTOR\_KOKKOS

The NVECTOR\_KOKKOS module is defined by the `Vector` templated class in the `sundials::kokkos` namespace:

```
template<class ExecutionSpace = Kokkos::DefaultExecutionSpace,
        class MemorySpace = typename ExecutionSpace::memory_space>
class Vector : public sundials::impl::BaseNVector,
               public sundials::ConvertibleTo<N_Vector>
```

To use the NVECTOR\_KOKKOS module, we construct an instance of the `Vector` class e.g.,

```
// Vector with extent length using the default execution space
sundials::kokkos::Vector<> x{length, sunctx};

// Vector with extent length using the Cuda execution space
sundials::kokkos::Vector<Kokkos::Cuda> x{length, sunctx};

// Vector based on an existing Kokkos::View
Kokkos::View<> view{"a view", length};
sundials::kokkos::Vector<> x{view, sunctx};

// Vector based on an existing Kokkos::View for device and host
Kokkos::View<Kokkos::Cuda> device_view{"a view", length};
Kokkos::View<Kokkos::HostMirror> host_view{Kokkos::create_mirror_view(device_view)};
sundials::kokkos::Vector<> x{device_view, host_view, sunctx};
```

Instances of the `Vector` class are implicitly or explicitly (using the `Convert()` method) convertible to a `N_Vector` e.g.,

```
sundials::kokkos::Vector<> x{length, sunctx};
N_Vector x2 = x;           // implicit conversion to N_Vector
N_Vector x3 = x.Convert(); // explicit conversion to N_Vector
```

No further interaction with a `Vector` is required from this point, and it is possible to use the `N_Vector` API to operate on `x2` or `x3`.

**Warning:** `N_VDestroy()` should never be called on a `N_Vector` that was created via conversion from a `sundials::kokkos::Vector`. Doing so may result in a double free.

The underlying `Vector` can be extracted from a `N_Vector` using `GetVec()` e.g.,

```
auto x_vec = GetVec<>(x3);
```

## 6.14.2 NVECTOR\_KOKKOS API

In this section we list the public API of the `sundials::kokkos::Vector` class.

```
template<class ExecutionSpace = Kokkos::DefaultExecutionSpace, class MemorySpace = class  
ExecutionSpace::memory_space>
```

```
class Vector : public sundials::impl::BaseNVector, public sundials::ConvertibleTo<N_Vector>
```

```
    using view_type = Kokkos::View<sunrealtype*, MemorySpace>;
```

```
    using size_type = typename view_type::size_type;
```

```
    using host_view_type = typename view_type::HostMirror;
```

```
    using memory_space = MemorySpace;
```

```
    using exec_space = typename MemorySpace::execution_space;
```

```
    using range_policy = Kokkos::RangePolicy<exec_space>;
```

```
Vector() = default
```

Default constructor – the vector must be copied or moved to.

```
Vector(size_type length, SUNContext sunctx)
```

Constructs a single `Vector` which is based on a 1D `Kokkos::View` with the `ExecutionSpace` and `MemorySpace` provided as template arguments.

### Parameters

- **length** – length of the vector (i.e., the extent of the View)
- **sunctx** – the SUNDIALS simulation context object (*SUNContext*)

```
Vector(view_type view, SUNContext sunctx)
```

Constructs a single `Vector` from an existing `Kokkos::View`. The View `ExecutionSpace` and `MemorySpace` must match the `ExecutionSpace` and `MemorySpace` provided as template arguments.

### Parameters

- **view** – A 1D `Kokkos::View`
- **sunctx** – the SUNDIALS simulation context object (*SUNContext*)

```
Vector(view_type view, host_view_type host_view, SUNContext sunctx)
```

Constructs a single `Vector` from an existing `Kokkos::View` for the device and the host. The `ExecutionSpace` and `MemorySpace` of the device View must match the `ExecutionSpace` and `MemorySpace` provided as template arguments.

### Parameters

- **view** – A 1D `Kokkos::View` for the device

- **host\_view** – A 1D Kokkos::View that is a Kokkos::HostMirror for the device view
- **sunctx** – the SUNDIALS simulation context object (*SUNContext*)

**Vector**(*Vector* &&that\_vector) noexcept

Move constructor.

**Vector**(const *Vector* &that\_vector)

Copy constructor. This creates a clone of the Vector, i.e., it creates a new Vector with the same properties, such as length, but it does not copy the data.

*Vector* &**operator=**(*Vector* &&rhs) noexcept

Move assignment.

*Vector* &**operator=**(const *Vector* &rhs)

Copy assignment. This creates a clone of the Vector, i.e., it creates a new Vector with the same properties, such as length, but it does not copy the data.

virtual ~**Vector**() = default;

Default destructor.

*size\_type* **Length**()

Get the vector length i.e., `extent(0)`.

*view\_type* **View**()

Get the underlying Kokkos::View for the device.

*host\_view\_type* **HostView**()

Get the underlying Kokkos::View for the host.

**operator N\_Vector**() override

Implicit conversion to a *N\_Vector*.

**operator N\_Vector**() const override

Implicit conversion to a *N\_Vector*.

*N\_Vector* **Convert**() override

Explicit conversion to a *N\_Vector*.

*N\_Vector* **Convert**() const override

Explicit conversion to a *N\_Vector*.

template<class **VectorType**>

inline *VectorType* \***GetVec**(*N\_Vector* v)

Get the *Vector* wrapped by a *N\_Vector*.

void **CopyToDevice**(*N\_Vector* v)

Copy the data from the host view to the device view with Kokkos::deep\_copy.

void **CopyFromDevice**(*N\_Vector* v)

Copy the data to the host view from the device view with Kokkos::deep\_copy.

template<class **VectorType**>

void **CopyToDevice**(*VectorType* &v)

Copy the data from the host view to the device view with Kokkos::deep\_copy.

template<class **VectorType**>

void **CopyFromDevice**(*VectorType* &v)

Copy the data to the host view from the device view with Kokkos::deep\_copy.



## 6.15 The NVECTOR\_OPENMPDEV Module

In situations where a user has access to a device such as a GPU for offloading computation, SUNDIALS provides an NVECTOR implementation using OpenMP device offloading, called NVECTOR\_OPENMPDEV.

The NVECTOR\_OPENMPDEV implementation defines the *content* field of the `N_Vector` to be a structure containing the length of the vector, a pointer to the beginning of a contiguous data array on the host, a pointer to the beginning of a contiguous data array on the device, and a boolean flag *own\_data* which specifies the ownership of host and device data arrays.

```
struct _N_VectorContent_OpenMPDEV
{
    sunindextype length;
    sunbooleantype own_data;
    sunrealtype    *host_data;
    sunrealtype    *dev_data;
};
```

The header file to include when using this module is `nvector_openmpdev.h`. The installed module library to link to is `libsundials_nvecopenmpdev.lib` where `.lib` is typically `.so` for shared libraries and `.a` for static libraries.

### 6.15.1 NVECTOR\_OPENMPDEV accessor macros

The following macros are provided to access the content of an NVECTOR\_OPENMPDEV vector.

#### NV\_CONTENT\_OMPDEV(v)

This macro gives access to the contents of the NVECTOR\_OPENMPDEV `N_Vector v`.

The assignment `v_cont = NV_CONTENT_S(v)` sets `v_cont` to be a pointer to the NVECTOR\_OPENMPDEV content structure.

Implementation:

```
#define NV_CONTENT_OMPDEV(v) ( (N_VectorContent_OpenMPDEV)(v->content) )
```

#### NV\_OWN\_DATA\_OMPDEV(v)

Access the *own\_data* component of the OpenMPDEV `N_Vector v`.

The assignment `v_data = NV_DATA_HOST_OMPDEV(v)` sets `v_data` to be a pointer to the first component of the data on the host for the `N_Vector v`.

Implementation:

```
#define NV_OWN_DATA_OMPDEV(v) ( NV_CONTENT_OMPDEV(v)->own_data )
```

#### NV\_DATA\_HOST\_OMPDEV(v)

The assignment `NV_DATA_HOST_OMPDEV(v) = v_data` sets the host component array of `v` to be `v_data` by storing the pointer `v_data`.

Implementation:

```
#define NV_DATA_HOST_OMPDEV(v) ( NV_CONTENT_OMPDEV(v)->host_data )
```



**NV\_DATA\_DEV\_OMPDEV(v)**

The assignment `v_dev_data = NV_DATA_DEV_OMPDEV(v)` sets `v_dev_data` to be a pointer to the first component of the data on the device for the `N_Vector v`. The assignment `NV_DATA_DEV_OMPDEV(v) = v_dev_data` sets the device component array of `v` to be `v_dev_data` by storing the pointer `v_dev_data`.

Implementation:

```
#define NV_DATA_DEV_OMPDEV(v) ( NV_CONTENT_OMPDEV(v)->dev_data )
```

**NV\_LENGTH\_OMPDEV(V)**

Access the *length* component of the OpenMPDEV `N_Vector v`.

The assignment `v_len = NV_LENGTH_OMPDEV(v)` sets `v_len` to be the length of `v`. On the other hand, the call `NV_LENGTH_OMPDEV(v) = len_v` sets the length of `v` to be `len_v`.

```
#define NV_LENGTH_OMPDEV(v) ( NV_CONTENT_OMPDEV(v)->length )
```

**6.15.2 NVECTOR\_OPENMPDEV functions**

The `NVECTOR_OPENMPDEV` module defines OpenMP device offloading implementations of all vector operations listed in §6.2, §6.2.2, §6.2.3, and §6.2.4, except for `N_VSetArrayPointer()`. As such, this vector cannot be used with the SUNDIALS direct solvers and preconditioners. It also provides methods for copying from the host to the device and vice versa.

The names of the vector operations are obtained from those in §6.2, §6.2.2, §6.2.3, and §6.2.4 by appending the suffix `_OpenMPDEV` (e.g. `N_VDestroy_OpenMPDEV`). The module `NVECTOR_OPENMPDEV` provides the following additional user-callable routines:

*N\_Vector* **N\_VNew\_OpenMPDEV**(*sunindextype* vec\_length, *SUNContext* sunctx)

This function creates and allocates memory for an `NVECTOR_OPENMPDEV` `N_Vector`.

*N\_Vector* **N\_VNewEmpty\_OpenMPDEV**(*sunindextype* vec\_length, *SUNContext* sunctx)

This function creates a new `NVECTOR_OPENMPDEV` `N_Vector` with an empty (NULL) data array.

*N\_Vector* **N\_VMake\_OpenMPDEV**(*sunindextype* vec\_length, *sunrealtype* \*h\_vdata, *sunrealtype* \*d\_vdata, *SUNContext* sunctx)

This function creates an `NVECTOR_OPENMPDEV` vector with user-supplied vector data arrays `h_vdata` and `d_vdata`. This function does not allocate memory for data itself.

*sunrealtype* \***N\_VGetHostArrayPointer\_OpenMPDEV**(*N\_Vector* v)

This function returns a pointer to the host data array.

*sunrealtype* \***N\_VGetDeviceArrayPointer\_OpenMPDEV**(*N\_Vector* v)

This function returns a pointer to the device data array.

void **N\_VPrint\_OpenMPDEV**(*N\_Vector* v)

This function prints the content of an `NVECTOR_OPENMPDEV` vector to `stdout`.

void **N\_VPrintFile\_OpenMPDEV**(*N\_Vector* v, FILE \*outfile)

This function prints the content of an `NVECTOR_OPENMPDEV` vector to `outfile`.

void **N\_VCopyToDevice\_OpenMPDEV**(*N\_Vector* v)

This function copies the content of an `NVECTOR_OPENMPDEV` vector's host data array to the device data array.

void **N\_VCopyFromDevice\_OpenMPDEV**(*N\_Vector* v)

This function copies the content of an NVECTOR\_OPENMPDEV vector's device data array to the host data array.

By default all fused and vector array operations are disabled in the NVECTOR\_OPENMPDEV module. The following additional user-callable routines are provided to enable or disable fused and vector array operations for a specific vector. To ensure consistency across vectors it is recommended to first create a vector with **N\_VNew\_OpenMPDEV**, enable/disable the desired operations for that vector with the functions below, and create any additional vectors from that vector using **N\_VClone**. This guarantees the new vectors will have the same operations enabled/disabled as cloned vectors inherit the same enable/disable options as the vector they are cloned from while vectors created with **N\_VNew\_OpenMPDEV** will have the default settings for the NVECTOR\_OPENMPDEV module.

*SUNErrCode* **N\_VEnableFusedOps\_OpenMPDEV**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) all fused and vector array operations in the NVECTOR\_OPENMPDEV vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearCombination\_OpenMPDEV**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear combination fused operation in the NVECTOR\_OPENMPDEV vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleAddMulti\_OpenMPDEV**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale and add a vector to multiple vectors fused operation in the NVECTOR\_OPENMPDEV vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableDotProdMulti\_OpenMPDEV**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the multiple dot products fused operation in the NVECTOR\_OPENMPDEV vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearSumVectorArray\_OpenMPDEV**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear sum operation for vector arrays in the NVECTOR\_OPENMPDEV vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleVectorArray\_OpenMPDEV**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale operation for vector arrays in the NVECTOR\_OPENMPDEV vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableConstVectorArray\_OpenMPDEV**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the const operation for vector arrays in the NVECTOR\_OPENMPDEV vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableWrmsNormVectorArray\_OpenMPDEV**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the WRMS norm operation for vector arrays in the NVECTOR\_OPENMPDEV vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableWrmsNormMaskVectorArray\_OpenMPDEV**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the masked WRMS norm operation for vector arrays in the NVECTOR\_OPENMPDEV vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleAddMultiVectorArray\_OpenMPDEV**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale and add a vector array to multiple vector arrays operation in the NVECTOR\_OPENMPDEV vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearCombinationVectorArray\_OpenMPDEV**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear combination operation for vector arrays in the NVECTOR\_OPENMPDEV vector. The return value is a *SUNErrCode*.

## Notes

- When looping over the components of an `N_Vector v`, it is most efficient to first obtain the component array via `h_data = N_VGetArrayPointer(v)` for the host array or `v_data = N_VGetDeviceArrayPointer(v)` for the device array, or equivalently to use the macros `h_data = NV_DATA_HOST_OMPDEV(v)` for the host array or `v_data = NV_DATA_DEV_OMPDEV(v)` for the device array, and then access `h_data[i]` or `v_data[i]` within the loop.
- When accessing individual components of an `N_Vector v` on the host remember to first copy the array back from the device with `N_VCopyFromDevice_OpenMPDEV(v)` to ensure the array is up to date.
- `N_VNewEmpty_OpenMPDEV()` and `N_VMake_OpenMPDEV()` set the field `own_data` to `SUNFALSE`. The implementation of `N_VDestroy()` will not attempt to free the pointer data for any `N_Vector` with `own_data` set to `SUNFALSE`. In such a case, it is the user's responsibility to deallocate the data pointers.
- To maximize efficiency, vector operations in the `NVECTOR_OPENMPDEV` implementation that have more than one `N_Vector` argument do not check for consistent internal representation of these vectors. It is the user's responsibility to ensure that such routines are called with `N_Vector` arguments that were all created with the same length.

## 6.16 The NVECTOR\_TRILINOS Module

The `NVECTOR_TRILINOS` module is an `NVECTOR` wrapper around the `Trilinos` Tpetra vector. The interface to Tpetra is implemented in the `sundials::trilinos::nvector_tpetra::TpetraVectorInterface` class. This class simply stores a reference counting pointer to a Tpetra vector and inherits from an empty structure

```
struct _N_VectorContent_Trilinos {};
```

to interface the C++ class with the `NVECTOR` C code. A pointer to an instance of this class is kept in the `content` field of the `N_Vector` object, to ensure that the Tpetra vector is not deleted for as long as the `N_Vector` object exists.

The Tpetra vector type in the `sundials::trilinos::nvector_tpetra::TpetraVectorInterface` class is defined as:

```
typedef Tpetra::Vector<sunrealtype, int, sunindextype> vector_type;
```

The Tpetra vector will use the SUNDIALS-specified `sunrealtype` as its scalar type, `int` as the local ordinal type, and `sunindextype` as the global ordinal type. This type definition will use Tpetra's default node type. Available Kokkos node types as of the Trilinos 12.14 release are serial (single thread), OpenMP, Pthread, and CUDA. The default node type is selected when building the Kokkos package. For example, the Tpetra vector will use a CUDA node if Tpetra was built with CUDA support and the CUDA node was selected as the default when Tpetra was built.

The header file to include when using this module is `nvector_trilinos.h`. The installed module library to link to is `libsundials_nvectrilinos.lib` where `.lib` is typically `.so` for shared libraries and `.a` for static libraries.

### 6.16.1 NVECTOR\_TRILINOS functions

The `NVECTOR_TRILINOS` module defines implementations of all vector operations listed in §6.2, §6.2.2, §6.2.3, and §6.2.4, except for `N_VGetArrayPointer()` and `N_VSetArrayPointer()`. As such, this vector cannot be used with the SUNDIALS direct solvers and preconditioners. When access to raw vector data is needed, it is recommended to extract the Trilinos Tpetra vector first, and then use Tpetra vector methods to access the data. Usage examples of `NVECTOR_TRILINOS` are provided in example programs for IDA.

The names of vector operations are obtained from those in §6.2 by appending the suffix `_Trilinos` (e.g. `N_VDestroy_Trilinos`). Vector operations call existing `Tpetra::Vector` methods when available. Vector operations specific to SUNDIALS are implemented as standalone functions in the namespace `sundials::trilinos::nvector_tpetra::TpetraVector`, located in the file `SundialsTpetraVectorKernels.hpp`. The module `NVECTOR_TRILINOS` provides the following additional user-callable routines:

Teuchos::RCP<*vector\_type*> **N\_VGetVector\_Trilinos**(N\_Vector v)

This C++ function takes an N\_Vector as the argument and returns a reference counting pointer to the underlying Tpetra vector. This is a standalone function defined in the global namespace.

N\_Vector **N\_VMake\_Trilinos**(Teuchos::RCP<*vector\_type*> v)

This C++ function creates and allocates memory for an NVECTOR\_TRILINOS wrapper around a user-provided Tpetra vector. This is a standalone function defined in the global namespace.

#### Notes

- The template parameter `vector_type` should be set as:

```
typedef sundials::trilinos::nvector_tpetra::TpetraVectorInterface::vector_type vector_type
```

This will ensure that data types used in Tpetra vector match those in SUNDIALS.

- When there is a need to access components of an N\_Vector\_Trilinos v, it is recommended to extract the Trilinos vector object via `x_vec = N_VGetVector_Trilinos(v)` and then access components using the appropriate Trilinos functions.
- The function `N_VDestroy_Trilinos` only deletes the N\_Vector wrapper. The underlying Tpetra vector object will exist for as long as there is at least one reference to it.

## 6.17 The NVECTOR\_MANYVECTOR Module

The NVECTOR\_MANYVECTOR module is designed to facilitate problems with an inherent data partitioning within a computational node for the solution vector. These data partitions are entirely user-defined, through construction of distinct NVECTOR modules for each component, that are then combined together to form the NVECTOR\_MANYVECTOR. Two potential use cases for this flexibility include:

- Heterogenous computational architectures:* for data partitioning between different computing resources on a node, architecture-specific subvectors may be created for each partition. For example, a user could create one GPU-accelerated component based on `NVECTOR_CUDA`, and another CPU threaded component based on `NVECTOR_OPENMP`.
- Structure of arrays (SOA) data layouts:* for problems that require separate subvectors for each solution component. For example, in an incompressible Navier-Stokes simulation, separate subvectors may be used for velocities and pressure, which are combined together into a single NVECTOR\_MANYVECTOR for the overall “solution”.

The above use cases are neither exhaustive nor mutually exclusive, and the NVECTOR\_MANYVECTOR implementation should support arbitrary combinations of these cases.

The NVECTOR\_MANYVECTOR implementation is designed to work with any NVECTOR subvectors that implement the minimum “standard” set of operations in §6.2.1. Additionally, NVECTOR\_MANYVECTOR sets no limit on the number of subvectors that may be attached (aside from the limitations of using `sunindextype` for indexing, and standard per-node memory limitations). However, while this ostensibly supports subvectors with one entry each (i.e., one subvector for each solution entry), we anticipate that this extreme situation will hinder performance due to non-stride-one memory accesses and increased function call overhead. We therefore recommend a relatively coarse partitioning of the problem, although actual performance will likely be problem-dependent.

As a final note, in the coming years we plan to introduce additional algebraic solvers and time integration modules that will leverage the problem partitioning enabled by NVECTOR\_MANYVECTOR. However, even at present we anticipate that users will be able to leverage such data partitioning in their problem-defining ODE right-hand side function, DAE or nonlinear solver residual function, preconditioners, or custom `SUNLinearSolver` or `SUNNonlinearSolver` modules.

### 6.17.1 NVECTOR\_MANYVECTOR structure

The NVECTOR\_MANYVECTOR implementation defines the *content* field of `N_Vector` to be a structure containing the number of subvectors comprising the ManyVector, the global length of the ManyVector (including all subvectors), a pointer to the beginning of the array of subvectors, and a boolean flag `own_data` indicating ownership of the subvectors that populate `subvec_array`.

```
struct _N_VectorContent_ManyVector {
    sunindextype  num_subvectors; /* number of vectors attached */
    sunindextype  global_length; /* overall manyvector length */
    N_Vector*     subvec_array;   /* pointer to N_Vector array */
    sunbooleantype own_data;      /* flag indicating data ownership */
};
```

The header file to include when using this module is `nvector_manyvector.h`. The installed module library to link against is `libsundials_nvecmanyvector.lib` where `.lib` is typically `.so` for shared libraries and `.a` for static libraries.

### 6.17.2 NVECTOR\_MANYVECTOR functions

The NVECTOR\_MANYVECTOR module implements all vector operations listed in §6.2 except for `N_VGetArrayPointer()`, `N_VSetArrayPointer()`, `N_VScaleAddMultiVectorArray()`, and `N_VLinearCombinationVecArray()`. As such, this vector cannot be used with the SUNDIALS direct solvers and preconditioners. Instead, the NVECTOR\_MANYVECTOR module provides functions to access subvectors, whose data may in turn be accessed according to their NVECTOR implementations.

The names of vector operations are obtained from those in §6.2 by appending the suffix `_ManyVector` (e.g. `N_VDestroy_ManyVector`). The module NVECTOR\_MANYVECTOR provides the following additional user-callable routines:

*N\_Vector* **N\_VNew\_ManyVector**(*sunindextype* num\_subvectors, *N\_Vector*\* vec\_array, *SUNContext* sunctx)

This function creates a ManyVector from a set of existing NVECTOR objects.

This routine will copy all `N_Vector` pointers from the input `vec_array`, so the user may modify/free that pointer array after calling this function. However, this routine does *not* allocate any new subvectors, so the underlying NVECTOR objects themselves should not be destroyed before the ManyVector that contains them.

Upon successful completion, the new ManyVector is returned; otherwise this routine returns NULL (e.g., a memory allocation failure occurred).

Users of the Fortran 2003 interface to this function will first need to use the generic `N_Vector` utility functions `N_VNewVectorArray()`, and `N_VSetVecAtIndexVectorArray()` to create the `N_Vector*` argument. This is further explained in §4.7.2.5, and the functions are documented in §6.1.1.

*N\_Vector* **N\_VGetSubvector\_ManyVector**(*N\_Vector* v, *sunindextype* vec\_num)

This function returns the `vec_num` subvector from the NVECTOR array.

*sunindextype* **N\_VGetSubvectorLocalLength\_ManyVector**(*N\_Vector* v, *sunindextype* vec\_num)

This function returns the local length of the `vec_num` subvector from the NVECTOR array.

Usage:

```
local_length = N_VGetSubvectorLocalLength_ManyVector(v, 0);
```

*sunrealtype*\* **N\_VGetSubvectorArrayPointer\_ManyVector**(*N\_Vector* v, *sunindextype* vec\_num)

This function returns the data array pointer for the `vec_num` subvector from the NVECTOR array.



If the input *vec\_num* is invalid, or if the subvector does not support the `N_VGetArrayPointer` operation, then `NULL` is returned.

***SUNErrCode* N\_VSetSubvectorArrayPointer\_ManyVector(*sunrealtype* \*v\_data, *N\_Vector* v, *sunindextype* vec\_num)**

This function sets the data array pointer for the *vec\_num* subvector from the `NVECTOR` array.

The function returns a *SUNErrCode*.

***sunindextype* N\_VGetNumSubvectors\_ManyVector(*N\_Vector* v)**

This function returns the overall number of subvectors in the `ManyVector` object.

By default all fused and vector array operations are disabled in the `NVECTOR_MANYVECTOR` module, except for `N_VWrmsNormVectorArray()` and `N_VWrmsNormMaskVectorArray()`, that are enabled by default. The following additional user-callable routines are provided to enable or disable fused and vector array operations for a specific vector. To ensure consistency across vectors it is recommended to first create a vector with `N_VNew_ManyVector()`, enable/disable the desired operations for that vector with the functions below, and create any additional vectors from that vector using `N_VClone()`. This guarantees that the new vectors will have the same operations enabled/disabled, since cloned vectors inherit those configuration options from the vector they are cloned from, while vectors created with `N_VNew_ManyVector()` will have the default settings for the `NVECTOR_MANYVECTOR` module. We note that these routines *do not* call the corresponding routines on subvectors, so those should be set up as desired *before* attaching them to the `ManyVector` in `N_VNew_ManyVector()`.

***SUNErrCode* N\_VEnableFusedOps\_ManyVector(*N\_Vector* v, *sunbooleantype* tf)**

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) all fused and vector array operations in the `manyvector` vector. The return value is a *SUNErrCode*.

***SUNErrCode* N\_VEnableLinearCombination\_ManyVector(*N\_Vector* v, *sunbooleantype* tf)**

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the linear combination fused operation in the `manyvector` vector. The return value is a *SUNErrCode*.

***SUNErrCode* N\_VEnableScaleAddMulti\_ManyVector(*N\_Vector* v, *sunbooleantype* tf)**

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the scale and add a vector to multiple vectors fused operation in the `manyvector` vector. The return value is a *SUNErrCode*.

***SUNErrCode* N\_VEnableDotProdMulti\_ManyVector(*N\_Vector* v, *sunbooleantype* tf)**

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the multiple dot products fused operation in the `manyvector` vector. The return value is a *SUNErrCode*.

***SUNErrCode* N\_VEnableLinearSumVectorArray\_ManyVector(*N\_Vector* v, *sunbooleantype* tf)**

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the linear sum operation for vector arrays in the `manyvector` vector. The return value is a *SUNErrCode*.

***SUNErrCode* N\_VEnableScaleVectorArray\_ManyVector(*N\_Vector* v, *sunbooleantype* tf)**

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the scale operation for vector arrays in the `manyvector` vector. The return value is a *SUNErrCode*.

***SUNErrCode* N\_VEnableConstVectorArray\_ManyVector(*N\_Vector* v, *sunbooleantype* tf)**

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the const operation for vector arrays in the `manyvector` vector. The return value is a *SUNErrCode*.

***SUNErrCode* N\_VEnableWrmsNormVectorArray\_ManyVector(*N\_Vector* v, *sunbooleantype* tf)**

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the WRMS norm operation for vector arrays in the `manyvector` vector. The return value is a *SUNErrCode*.

***SUNErrCode* N\_VEnableWrmsNormMaskVectorArray\_ManyVector(*N\_Vector* v, *sunbooleantype* tf)**

This function enables (`SUNTRUE`) or disables (`SUNFALSE`) the masked WRMS norm operation for vector arrays in the `manyvector` vector. The return value is a *SUNErrCode*.

## Notes

- `N_VNew_ManyVector()` sets the field `own_data = SUNFALSE`. The `ManyVector` implementation of `N_VDestroy()` will not attempt to call `N_VDestroy()` on any subvectors contained in the subvector array for any `N_Vector` with `own_data` set to `SUNFALSE`. In such a case, it is the user's responsibility to deallocate the subvectors.
- To maximize efficiency, arithmetic vector operations in the `NVECTOR_MANYVECTOR` implementation that have more than one `N_Vector` argument do not check for consistent internal representation of these vectors. It is the user's responsibility to ensure that such routines are called with `N_Vector` arguments that were all created with the same subvector representations.

## 6.18 The NVECTOR\_MPIMANYVECTOR Module

The `NVECTOR_MPIMANYVECTOR` module is designed to facilitate problems with an inherent data partitioning for the solution vector, and when using distributed-memory parallel architectures. As such, this implementation supports all use cases allowed by the MPI-unaware `NVECTOR_MANYVECTOR` implementation, as well as partitioning data between nodes in a parallel environment. These data partitions are entirely user-defined, through construction of distinct `NVECTOR` modules for each component, that are then combined together to form the `NVECTOR_MPIMANYVECTOR`. Three potential use cases for this module include:

- Heterogenous computational architectures (single-node or multi-node)*: for data partitioning between different computing resources on a node, architecture-specific subvectors may be created for each partition. For example, a user could create one MPI-parallel component based on `NVECTOR_PARALLEL`, another GPU-accelerated component based on `NVECTOR_CUDA`.
- Process-based multiphysics decompositions (multi-node)*: for computations that combine separate MPI-based simulations together, each subvector may reside on a different MPI communicator, and the `MPIManyVector` combines these via an MPI *intercommunicator* that connects these distinct simulations together.
- Structure of arrays (SOA) data layouts (single-node or multi-node)*: for problems that require separate subvectors for each solution component. For example, in an incompressible Navier-Stokes simulation, separate subvectors may be used for velocities and pressure, which are combined together into a single `MPIManyVector` for the overall "solution".

The above use cases are neither exhaustive nor mutually exclusive, and the `NVECTOR_MPIMANYVECTOR` implementation should support arbitrary combinations of these cases.

The `NVECTOR_MPIMANYVECTOR` implementation is designed to work with any `NVECTOR` subvectors that implement the minimum "standard" set of operations in §6.2.1, however significant performance benefits may be obtained when subvectors additionally implement the optional local reduction operations listed in §6.2.4.

Additionally, `NVECTOR_MPIMANYVECTOR` sets no limit on the number of subvectors that may be attached (aside from the limitations of using `sunindextype` for indexing, and standard per-node memory limitations). However, while this ostensibly supports subvectors with one entry each (i.e., one subvector for each solution entry), we anticipate that this extreme situation will hinder performance due to non-stride-one memory accesses and increased function call overhead. We therefore recommend a relatively coarse partitioning of the problem, although actual performance will likely be problem-dependent.

As a final note, in the coming years we plan to introduce additional algebraic solvers and time integration modules that will leverage the problem partitioning enabled by `NVECTOR_MPIMANYVECTOR`. However, even at present we anticipate that users will be able to leverage such data partitioning in their problem-defining ODE right-hand side function, DAE or nonlinear solver residual function, preconditioners, or custom `SUNLinearSolver` or `SUNNonlinearSolver` modules.

### 6.18.1 NVECTOR\_MPIMANYVECTOR structure

The NVECTOR\_MPIMANYVECTOR implementation defines the *content* field of *N\_Vector* to be a structure containing the MPI communicator (or *MPI\_COMM\_NULL* if running on a single-node), the number of subvectors comprising the *MPIManyVector*, the global length of the *MPIManyVector* (including all subvectors on all MPI ranks), a pointer to the beginning of the array of subvectors, and a boolean flag *own\_data* indicating ownership of the subvectors that populate *subvec\_array*.

```
struct _N_VectorContent_MPIManyVector {
    MPI_Comm      comm;           /* overall MPI communicator      */
    sunindextype  num_subvectors; /* number of vectors attached   */
    sunindextype  global_length;  /* overall mpimanyvector length */
    N_Vector*     subvec_array;   /* pointer to N_Vector array    */
    sunboolean_t  own_data;       /* flag indicating data ownership */
};
```

The header file to include when using this module is *nvector\_mpimanyvector.h*. The installed module library to link against is *libsundials\_nvecmpimanyvector.lib* where *.lib* is typically *.so* for shared libraries and *.a* for static libraries.

---

**Note:** If SUNDIALS is configured with MPI disabled, then the *MPIManyVector* library will not be built. Furthermore, any user codes that include *nvector\_mpimanyvector.h* *must* be compiled using an MPI-aware compiler (whether the specific user code utilizes MPI or not). We note that the NVECTOR\_MANYVECTOR implementation is designed for *ManyVector* use cases in an MPI-unaware environment.

---

### 6.18.2 NVECTOR\_MPIMANYVECTOR functions

The NVECTOR\_MPIMANYVECTOR module implements all vector operations listed in §6.2, except for *N\_VGetArrayPointer()*, *N\_VSetArrayPointer()*, *N\_VScaleAddMultiVectorArray()*, and *N\_VLinearCombinationVectorArray()*. As such, this vector cannot be used with the SUNDIALS direct solvers and preconditioners. Instead, the NVECTOR\_MPIMANYVECTOR module provides functions to access subvectors, whose data may in turn be accessed according to their NVECTOR implementations.

The names of vector operations are obtained from those in §6.2 by appending the suffix *\_MPIManyVector* (e.g. *N\_VDestroy\_MPIManyVector*). The module NVECTOR\_MPIMANYVECTOR provides the following additional user-callable routines:

*N\_Vector* **N\_VNew\_MPIManyVector**(*sunindextype* num\_subvectors, *N\_Vector*\* vec\_array, *SUNContext* sunctx)

This function creates a *MPIManyVector* from a set of existing NVECTOR objects, under the requirement that all MPI-aware subvectors use the same MPI communicator (this is checked internally). If none of the subvectors are MPI-aware, then this may equivalently be used to describe data partitioning within a single node. We note that this routine is designed to support use cases A and C above.

This routine will copy all *N\_Vector* pointers from the input *vec\_array*, so the user may modify/free that pointer array after calling this function. However, this routine does *not* allocate any new subvectors, so the underlying NVECTOR objects themselves should not be destroyed before the *MPIManyVector* that contains them.

Upon successful completion, the new *MPIManyVector* is returned; otherwise this routine returns *NULL* (e.g., if two MPI-aware subvectors use different MPI communicators).

Users of the Fortran 2003 interface to this function will first need to use the generic *N\_Vector* utility functions *N\_VNewVectorArray()*, and *N\_VSetVecAtIndexVectorArray()* to create the *N\_Vector\** argument. This is further explained in §4.7.2.5, and the functions are documented in §6.1.1.



*N\_Vector* **N\_VMake\_MPIManyVector**(MPI\_Comm comm, *sunindextype* num\_subvectors, *N\_Vector* \*vec\_array, *SUNContext* sunctx)

This function creates a MPIManyVector from a set of existing NVECTOR objects, and a user-created MPI communicator that “connects” these subvectors. Any MPI-aware subvectors may use different MPI communicators than the input *comm*. We note that this routine is designed to support any combination of the use cases above.

The input *comm* should be this user-created MPI communicator. This routine will internally call `MPI_Comm_dup` to create a copy of the input *comm*, so the user-supplied *comm* argument need not be retained after the call to `N_VMake_MPIManyVector()`.

If all subvectors are MPI-unaware, then the input *comm* argument should be `MPI_COMM_NULL`, although in this case, it would be simpler to call `N_VNew_MPIManyVector()` instead, or to just use the NVECTOR\_MANYVECTOR module.

This routine will copy all *N\_Vector* pointers from the input *vec\_array*, so the user may modify/free that pointer array after calling this function. However, this routine does *not* allocate any new subvectors, so the underlying NVECTOR objects themselves should not be destroyed before the MPIManyVector that contains them.

Upon successful completion, the new MPIManyVector is returned; otherwise this routine returns `NULL` (e.g., if the input *vec\_array* is `NULL`).

*N\_Vector* **N\_VGetSubvector\_MPIManyVector**(*N\_Vector* v, *sunindextype* vec\_num)

This function returns the *vec\_num* subvector from the NVECTOR array.

*sunindextype* **N\_VGetSubvectorLocalLength\_MPIManyVector**(*N\_Vector* v, *sunindextype* vec\_num)

This function returns the local length of the *vec\_num* subvector from the NVECTOR array.

Usage:

```
local_length = N_VGetSubvectorLocalLength_MPIManyVector(v, 0);
```

*sunrealtype* \***N\_VGetSubvectorArrayPointer\_MPIManyVector**(*N\_Vector* v, *sunindextype* vec\_num)

This function returns the data array pointer for the *vec\_num* subvector from the NVECTOR array.

If the input *vec\_num* is invalid, or if the subvector does not support the `N_VGetArrayPointer` operation, then `NULL` is returned.

*SUNErrCode* **N\_VSetSubvectorArrayPointer\_MPIManyVector**(*sunrealtype* \*v\_data, *N\_Vector* v, *sunindextype* vec\_num)

This function sets the data array pointer for the *vec\_num* subvector from the NVECTOR array.

The function returns a *SUNErrCode*.

*sunindextype* **N\_VGetNumSubvectors\_MPIManyVector**(*N\_Vector* v)

This function returns the overall number of subvectors in the MPIManyVector object.

By default all fused and vector array operations are disabled in the NVECTOR\_MPIMANYVECTOR module, except for `N_VWrmsNormVectorArray()` and `N_VWrmsNormMaskVectorArray()`, that are enabled by default. The following additional user-callable routines are provided to enable or disable fused and vector array operations for a specific vector. To ensure consistency across vectors it is recommended to first create a vector with `N_VNew_MPIManyVector()` or `N_VMake_MPIManyVector()`, enable/disable the desired operations for that vector with the functions below, and create any additional vectors from that vector using `N_VClone()`. This guarantees that the new vectors will have the same operations enabled/disabled, since cloned vectors inherit those configuration options from the vector they are cloned from, while vectors created with `N_VNew_MPIManyVector()` and `N_VMake_MPIManyVector()` will have the default settings for the NVECTOR\_MPIMANYVECTOR module. We note that these routines *do not* call the corresponding routines on subvectors, so those should be set up as desired *before* attaching them to the MPIManyVector in `N_VNew_MPIManyVector()` or `N_VMake_MPIManyVector()`.

*SUNErrCode* **N\_VEnableFusedOps\_MPIManyVector**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) all fused and vector array operations in the MPI-ManyVector vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearCombination\_MPIManyVector**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear combination fused operation in the MPI-ManyVector vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleAddMulti\_MPIManyVector**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale and add a vector to multiple vectors fused operation in the MPIManyVector vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableDotProdMulti\_MPIManyVector**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the multiple dot products fused operation in the MPI-ManyVector vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableLinearSumVectorArray\_MPIManyVector**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear sum operation for vector arrays in the MPI-ManyVector vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableScaleVectorArray\_MPIManyVector**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the scale operation for vector arrays in the MPI-ManyVector vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableConstVectorArray\_MPIManyVector**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the const operation for vector arrays in the MPI-ManyVector vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableWrmsNormVectorArray\_MPIManyVector**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the WRMS norm operation for vector arrays in the MPIManyVector vector. The return value is a *SUNErrCode*.

*SUNErrCode* **N\_VEnableWrmsNormMaskVectorArray\_MPIManyVector**(*N\_Vector* v, *sunbooleantype* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the masked WRMS norm operation for vector arrays in the MPIManyVector vector. The return value is a *SUNErrCode*.

## Notes

- *N\_VNew\_MPIManyVector()* and *N\_VMake\_MPIManyVector()* set the field *own\_data* = SUNFALSE. The MPI-ManyVector implementation of *N\_VDestroy()* will not attempt to call *N\_VDestroy()* on any subvectors contained in the subvector array for any *N\_Vector* with *own\_data* set to SUNFALSE. In such a case, it is the user's responsibility to deallocate the subvectors.
- To maximize efficiency, arithmetic vector operations in the NVECTOR\_MPIMANYVECTOR implementation that have more than one *N\_Vector* argument do not check for consistent internal representation of these vectors. It is the user's responsibility to ensure that such routines are called with *N\_Vector* arguments that were all created with the same subvector representations.

## 6.19 The NVECTOR\_MPIPLUSX Module

The NVECTOR\_MPIPLUSX module is designed to facilitate the MPI+X paradigm, where X is some form of on-node (local) parallelism (e.g. OpenMP, CUDA). This paradigm is becoming increasingly popular with the rise of heterogeneous computing architectures.

The NVECTOR\_MPIPLUSX implementation is designed to work with any NVECTOR that implements the minimum “standard” set of operations in §6.2.1. However, it is not recommended to use the NVECTOR\_PARALLEL, NVECTOR\_PARHYP, NVECTOR\_PETSC, or NVECTOR\_TRILINOS implementations underneath the NVECTOR\_MPIPLUSX module since they already provide MPI capabilities.

### 6.19.1 NVECTOR\_MPIPLUSX structure

The NVECTOR\_MPIPLUSX implementation is a thin wrapper around the NVECTOR\_MPIMANYVECTOR. Accordingly, it adopts the same content structure as defined in §6.18.1.

The header file to include when using this module is `nvector_mpiplusx.h`. The installed module library to link against is `libsundials_nvecmpiplusx.lib` where `.lib` is typically `.so` for shared libraries and `.a` for static libraries.

---

**Note:** If SUNDIALS is configured with MPI disabled, then the `mpiplusx` library will not be built. Furthermore, any user codes that include `nvector_mpiplusx.h` *must* be compiled using an MPI-aware compiler.

---

### 6.19.2 NVECTOR\_MPIPLUSX functions

The NVECTOR\_MPIPLUSX module adopts all vector operations listed in §6.2, from the NVECTOR\_MPIMANYVECTOR (see §6.18) except for `N_VGetArrayPointer()`, and `N_VSetArrayPointer()`; the module provides its own implementation of these functions that call the local vector implementations. Therefore, the NVECTOR\_MPIPLUSX module implements all of the operations listed in the referenced sections except for `N_VScaleAddMultiVectorArray()`, and `N_VLinearCombinationVectorArray()`. Accordingly, its compatibility with the SUNDIALS direct solvers and preconditioners depends on the local vector implementation.

The module NVECTOR\_MPIPLUSX provides the following additional user-callable routines:

`N_Vector` **N\_VMake\_MPIplusX**(MPI\_Comm comm, `N_Vector` \*local\_vector, `SUNContext` sunctx)

This function creates a MPIplusX vector from an existing local (i.e. on node) NVECTOR object, and a user-created MPI communicator.

The input *comm* should be this user-created MPI communicator. This routine will internally call `MPI_Comm_dup` to create a copy of the input *comm*, so the user-supplied *comm* argument need not be retained after the call to `N_VMake_MPIplusX()`.

This routine will copy the NVECTOR pointer to the input *local\_vector*, so the underlying local NVECTOR object should not be destroyed before the `mpiplusx` that contains it.

Upon successful completion, the new MPIplusX is returned; otherwise this routine returns NULL (e.g., if the input *local\_vector* is NULL).

`N_Vector` **N\_VGetLocalVector\_MPIplusX**(`N_Vector` v)

This function returns the local vector underneath the MPIplusX NVECTOR.

`sunindextype` **N\_VGetLocalLength\_MPIplusX**(`N_Vector` v)

This function returns the local length of the vector underneath the MPIplusX NVECTOR.

Usage:

```
local_length = N_VGetLocalLength_MPIPlusX(v);
```

*sunrealtype* \***N\_VGetArrayPointer\_MPIPlusX**(*N\_Vector* v)

This function returns the data array pointer for the local vector.

If the local vector does not support the *N\_VGetArrayPointer()* operation, then NULL is returned.

void **N\_VSetArrayPointer\_MPIPlusX**(*sunrealtype* \*v\_data, *N\_Vector* v)

This function sets the data array pointer for the local vector if the local vector implements the *N\_VSetArrayPointer()* operation.

The NVECTOR\_MPIPLUSX module does not implement any fused or vector array operations. Instead users should enable/disable fused operations on the local vector.

#### Notes

- *N\_VMake\_MPIPlusX()* sets the field `own_data = SUNFALSE` and the MPIPlusX implementation of *N\_VDestroy()* will not call *N\_VDestroy()* on the local vector. In this a case, it is the user's responsibility to deallocate the local vector.
- To maximize efficiency, arithmetic vector operations in the NVECTOR\_MPIPLUSX implementation that have more than one *N\_Vector* argument do not check for consistent internal representation of these vectors. It is the user's responsibility to ensure that such routines are called with *N\_Vector* arguments that were all created with the same subvector representations.

## 6.20 NVECTOR Examples

There are NVECTOR examples that may be installed for each implementation. Each implementation makes use of the functions in `test_nvector.c`. These example functions show simple usage of the NVECTOR family of functions. The input to the examples are the vector length, number of threads (if threaded implementation), and a print timing flag.

The following is a list of the example functions in `test_nvector.c`:

- `Test_N_VClone`: Creates clone of vector and checks validity of clone.
- `Test_N_VCloneEmpty`: Creates clone of empty vector and checks validity of clone.
- `Test_N_VCloneVectorArray`: Creates clone of vector array and checks validity of cloned array.
- `Test_N_VCloneVectorArray`: Creates clone of empty vector array and checks validity of cloned array.
- `Test_N_VGetArrayPointer`: Get array pointer.
- `Test_N_VSetArrayPointer`: Allocate new vector, set pointer to new vector array, and check values.
- `Test_N_VGetLength`: Compares self-reported length to calculated length.
- `Test_N_VGetCommunicator`: Compares self-reported communicator to the one used in constructor; or for MPI-unaware vectors it ensures that NULL is reported.
- `Test_N_VLinearSum Case 1a`: Test  $y = x + y$
- `Test_N_VLinearSum Case 1b`: Test  $y = -x + y$
- `Test_N_VLinearSum Case 1c`: Test  $y = ax + y$
- `Test_N_VLinearSum Case 2a`: Test  $x = x + y$
- `Test_N_VLinearSum Case 2b`: Test  $x = x - y$
- `Test_N_VLinearSum Case 2c`: Test  $x = x + by$

- Test\_N\_VLinearSum Case 3: Test  $z = x + y$
- Test\_N\_VLinearSum Case 4a: Test  $z = x - y$
- Test\_N\_VLinearSum Case 4b: Test  $z = -x + y$
- Test\_N\_VLinearSum Case 5a: Test  $z = x + by$
- Test\_N\_VLinearSum Case 5b: Test  $z = ax + y$
- Test\_N\_VLinearSum Case 6a: Test  $z = -x + by$
- Test\_N\_VLinearSum Case 6b: Test  $z = ax - y$
- Test\_N\_VLinearSum Case 7: Test  $z = a(x + y)$
- Test\_N\_VLinearSum Case 8: Test  $z = a(x - y)$
- Test\_N\_VLinearSum Case 9: Test  $z = ax + by$
- Test\_N\_VConst: Fill vector with constant and check result.
- Test\_N\_VProd: Test vector multiply:  $z = x * y$
- Test\_N\_VDiv: Test vector division:  $z = x / y$
- Test\_N\_VScale: Case 1: scale:  $x = cx$
- Test\_N\_VScale: Case 2: copy:  $z = x$
- Test\_N\_VScale: Case 3: negate:  $z = -x$
- Test\_N\_VScale: Case 4: combination:  $z = cx$
- Test\_N\_VAbs: Create absolute value of vector.
- Test\_N\_VInv: Compute  $z[i] = 1 / x[i]$

\*\* Test\_N\_VAddConst: add constant vector:  $z = c + x$

- Test\_N\_VDotProd: Calculate dot product of two vectors.
- Test\_N\_VMaxNorm: Create vector with known values, find and validate the max norm.
- Test\_N\_VWrmsNorm: Create vector of known values, find and validate the weighted root mean square.
- Test\_N\_VWrmsNormMask: Create vector of known values, find and validate the weighted root mean square using all elements except one.
- Test\_N\_VMin: Create vector, find and validate the min.
- Test\_N\_VWL2Norm: Create vector, find and validate the weighted Euclidean L2 norm.
- Test\_N\_VL1Norm: Create vector, find and validate the L1 norm.
- Test\_N\_VCompare: Compare vector with constant returning and validating comparison vector.
- Test\_N\_VInvTest: Test  $z[i] = 1 / x[i]$
- Test\_N\_VConstrMask: Test mask of vector  $x$  with vector  $c$ .
- Test\_N\_VMinQuotient: Fill two vectors with known values. Calculate and validate minimum quotient.
- Test\_N\_VLinearCombination: Case 1a: Test  $x = a x$
- Test\_N\_VLinearCombination: Case 1b: Test  $z = a x$
- Test\_N\_VLinearCombination: Case 2a: Test  $x = a x + b y$
- Test\_N\_VLinearCombination: Case 2b: Test  $z = a x + b y$

- Test\_N\_VLinearCombination: Case 3a: Test  $x = x + a y + b z$
- Test\_N\_VLinearCombination: Case 3b: Test  $x = a x + b y + c z$
- Test\_N\_VLinearCombination: Case 3c: Test  $w = a x + b y + c z$
- Test\_N\_VScaleAddMulti: Case 1a:  $y = a x + y$
- Test\_N\_VScaleAddMulti: Case 1b:  $z = a x + y$
- Test\_N\_VScaleAddMulti: Case 2a:  $Y[i] = c[i] x + Y[i]$ ,  $i = 1, 2, 3$
- Test\_N\_VScaleAddMulti: Case 2b:  $Z[i] = c[i] x + Y[i]$ ,  $i = 1, 2, 3$
- Test\_N\_VDotProdMulti: Case 1: Calculate the dot product of two vectors
- Test\_N\_VDotProdMulti: Case 2: Calculate the dot product of one vector with three other vectors in a vector array.
- Test\_N\_VLinearSumVectorArray: Case 1:  $z = a x + b y$
- Test\_N\_VLinearSumVectorArray: Case 2a:  $Z[i] = a X[i] + b Y[i]$
- Test\_N\_VLinearSumVectorArray: Case 2b:  $X[i] = a X[i] + b Y[i]$
- Test\_N\_VLinearSumVectorArray: Case 2c:  $Y[i] = a X[i] + b Y[i]$
- Test\_N\_VScaleVectorArray: Case 1a:  $y = c y$
- Test\_N\_VScaleVectorArray: Case 1b:  $z = c y$
- Test\_N\_VScaleVectorArray: Case 2a:  $Y[i] = c[i] Y[i]$
- Test\_N\_VScaleVectorArray: Case 2b:  $Z[i] = c[i] Y[i]$
- Test\_N\_VConstVectorArray: Case 1a:  $z = c$
- Test\_N\_VConstVectorArray: Case 1b:  $Z[i] = c$
- Test\_N\_VWrmsNormVectorArray: Case 1a: Create a vector of know values, find and validate the weighted root mean square norm.
- Test\_N\_VWrmsNormVectorArray: Case 1b: Create a vector array of three vectors of know values, find and validate the weighted root mean square norm of each.
- Test\_N\_VWrmsNormMaskVectorArray: Case 1a: Create a vector of know values, find and validate the weighted root mean square norm using all elements except one.
- Test\_N\_VWrmsNormMaskVectorArray: Case 1b: Create a vector array of three vectors of know values, find and validate the weighted root mean square norm of each using all elements except one.
- Test\_N\_VScaleAddMultiVectorArray: Case 1a:  $y = a x + y$
- Test\_N\_VScaleAddMultiVectorArray: Case 1b:  $z = a x + y$
- Test\_N\_VScaleAddMultiVectorArray: Case 2a:  $Y[j][0] = a[j] X[0] + Y[j][0]$
- Test\_N\_VScaleAddMultiVectorArray: Case 2b:  $Z[j][0] = a[j] X[0] + Y[j][0]$
- Test\_N\_VScaleAddMultiVectorArray: Case 3a:  $Y[0][i] = a[0] X[i] + Y[0][i]$
- Test\_N\_VScaleAddMultiVectorArray: Case 3b:  $Z[0][i] = a[0] X[i] + Y[0][i]$
- Test\_N\_VScaleAddMultiVectorArray: Case 4a:  $Y[j][i] = a[j] X[i] + Y[j][i]$
- Test\_N\_VScaleAddMultiVectorArray: Case 4b:  $Z[j][i] = a[j] X[i] + Y[j][i]$
- Test\_N\_VLinearCombinationVectorArray: Case 1a:  $x = a x$
- Test\_N\_VLinearCombinationVectorArray: Case 1b:  $z = a x$

- `Test_N_VLinearCombinationVectorArray`: Case 2a:  $x = a x + b y$
- `Test_N_VLinearCombinationVectorArray`: Case 2b:  $z = a x + b y$
- `Test_N_VLinearCombinationVectorArray`: Case 3a:  $x = a x + b y + c z$
- `Test_N_VLinearCombinationVectorArray`: Case 3b:  $w = a x + b y + c z$
- `Test_N_VLinearCombinationVectorArray`: Case 4a:  $X[0][i] = c[0] X[0][i]$
- `Test_N_VLinearCombinationVectorArray`: Case 4b:  $Z[i] = c[0] X[0][i]$
- `Test_N_VLinearCombinationVectorArray`: Case 5a:  $X[0][i] = c[0] X[0][i] + c[1] X[1][i]$
- `Test_N_VLinearCombinationVectorArray`: Case 5b:  $Z[i] = c[0] X[0][i] + c[1] X[1][i]$
- `Test_N_VLinearCombinationVectorArray`: Case 6a:  $X[0][i] = X[0][i] + c[1] X[1][i] + c[2] X[2][i]$
- `Test_N_VLinearCombinationVectorArray`: Case 6b:  $X[0][i] = c[0] X[0][i] + c[1] X[1][i] + c[2] X[2][i]$
- `Test_N_VLinearCombinationVectorArray`: Case 6c:  $Z[i] = c[0] X[0][i] + c[1] X[1][i] + c[2] X[2][i]$
- `Test_N_VDotProdLocal`: Calculate MPI task-local portion of the dot product of two vectors.
- `Test_N_VMaxNormLocal`: Create vector with known values, find and validate the MPI task-local portion of the max norm.
- `Test_N_VMinLocal`: Create vector, find and validate the MPI task-local min.
- `Test_N_VL1NormLocal`: Create vector, find and validate the MPI task-local portion of the L1 norm.
- `Test_N_VWSqrSumLocal`: Create vector of known values, find and validate the MPI task-local portion of the weighted squared sum of two vectors.
- `Test_N_VWSqrSumMaskLocal`: Create vector of known values, find and validate the MPI task-local portion of the weighted squared sum of two vectors, using all elements except one.
- `Test_N_VInvTestLocal`: Test the MPI task-local portion of  $z[i] = 1 / x[i]$
- `Test_N_VConstrMaskLocal`: Test the MPI task-local portion of the mask of vector  $x$  with vector  $c$ .
- `Test_N_VMinQuotientLocal`: Fill two vectors with known values. Calculate and validate the MPI task-local minimum quotient.
- `Test_N_VMBufSize`: Tests for accuracy in the reported buffer size.
- `Test_N_VMBufPack`: Tests for accuracy in the buffer packing routine.
- `Test_N_VMBufUnpack`: Tests for accuracy in the buffer unpacking routine.





## Chapter 7

# Matrix Data Structures

The SUNDIALS library comes packaged with a variety of `SUNMatrix` implementations, designed for simulations requiring direct linear solvers for problems in serial or shared-memory parallel environments. SUNDIALS additionally provides a simple interface for generic matrices (akin to a C++ *abstract base class*). All of the major SUNDIALS packages (CVODE(s), IDA(s), KINSOL, ARKODE), are constructed to only depend on these generic matrix operations, making them immediately extensible to new user-defined matrix objects. For each of the SUNDIALS-provided matrix types, SUNDIALS also provides at least two `SUNLinearSolver` implementations that factor these matrix objects and use them in the solution of linear systems.

### 7.1 Description of the SUNMATRIX Modules

For problems that involve direct methods for solving linear systems, the SUNDIALS packages not only operate on generic vectors, but also on generic matrices (of type `SUNMatrix`), through a set of operations defined by the particular SUNMATRIX implementation. Users can provide their own specific implementation of the SUNMATRIX module, particularly in cases where they provide their own `N_Vector` and/or linear solver modules, and require matrices that are compatible with those implementations. The generic `SUNMatrix` operations are described below, and descriptions of the SUNMATRIX implementations provided with SUNDIALS follow.

The generic `SUNMatrix` type has been modeled after the object-oriented style of the generic `N_Vector` type. Specifically, a generic `SUNMatrix` is a pointer to a structure that has an implementation-dependent *content* field containing the description and actual data of the matrix, and an *ops* field pointing to a structure with generic matrix operations.

A `SUNMatrix` is a pointer to the `_generic_SUNMatrix` structure:

```
typedef struct _generic_SUNMatrix *SUNMatrix
```

```
struct _generic_SUNMatrix
```

The structure defining the SUNDIALS matrix class.

```
void *content
```

Pointer to matrix-specific member data

```
struct _generic_SUNMatrix_Ops *ops
```

A virtual table of matrix operations provided by a specific implementation

```
SUNContext sunctx
```

The SUNDIALS simulation context

The virtual table structure is defined as

**struct \_generic\_SUNMatrix\_Ops**

The structure defining *SUNMatrix* operations.

*SUNMatrix\_ID* (\***getid**)(*SUNMatrix*)

The function implementing *SUNMatGetID()*

*SUNMatrix* (\***clone**)(*SUNMatrix*)

The function implementing *SUNMatClone()*

void (\***destroy**)(*SUNMatrix*)

The function implementing *SUNMatDestroy()*

*SUNErrCode* (\***zero**)(*SUNMatrix*)

The function implementing *SUNMatZero()*

*SUNErrCode* (\***copy**)(*SUNMatrix*, *SUNMatrix*)

The function implementing *SUNMatCopy()*

*SUNErrCode* (\***scaleadd**)(*sunrealtype*, *SUNMatrix*, *SUNMatrix*)

The function implementing *SUNMatScaleAdd()*

*SUNErrCode* (\***scaleaddi**)(*sunrealtype*, *SUNMatrix*)

The function implementing *SUNMatScaleAddI()*

*SUNErrCode* (\***matvecsetup**)(*SUNMatrix*)

The function implementing *SUNMatMatvecSetup()*

*SUNErrCode* (\***matvec**)(*SUNMatrix*, *N\_Vector*, *N\_Vector*)

The function implementing *SUNMatMatvec()*

*SUNErrCode* (\***space**)(*SUNMatrix*, long int\*, long int\*)

The function implementing *SUNMatSpace()*

The generic SUNMATRIX module defines and implements the matrix operations acting on a *SUNMatrix*. These routines are nothing but wrappers for the matrix operations defined by a particular SUNMATRIX implementation, which are accessed through the *ops* field of the *SUNMatrix* structure. To illustrate this point we show below the implementation of a typical matrix operation from the generic SUNMATRIX module, namely *SUNMatZero*, which sets all values of a matrix *A* to zero, returning a flag denoting a successful/failed operation:

```
SUNErrCode SUNMatZero(SUNMatrix A)
{
    return(A->ops->zero(A));
}
```

§7.2 contains a complete list of all matrix operations defined by the generic SUNMATRIX module. A particular implementation of the SUNMATRIX module must:

- Specify the *content* field of the *SUNMatrix* object.
- Define and implement a minimal subset of the matrix operations. See the documentation for each SUNDIALS package and/or linear solver to determine which SUNMATRIX operations they require.

Note that the names of these routines should be unique to that implementation in order to permit using more than one SUNMATRIX module (each with different *SUNMatrix* internal data representations) in the same code.

- Define and implement user-callable constructor and destructor routines to create and free a *SUNMatrix* with the new *content* field and with *ops* pointing to the new matrix operations.
- Optionally, define and implement additional user-callable routines acting on the newly defined *SUNMatrix* (e.g., a routine to print the *content* for debugging purposes).

- Optionally, provide accessor macros as needed for that particular implementation to be used to access different parts in the content field of the newly defined `SUNMatrix`.

To aid in the creation of custom SUNMATRIX modules the generic SUNMATRIX module provides three utility functions `SUNMatNewEmpty()`, `SUNMatCopyOps()`, and `SUNMatFreeEmpty()`. When used in custom SUNMATRIX constructors and clone routines these functions will ease the introduction of any new optional matrix operations to the SUNMATRIX API by ensuring only required operations need to be set and all operations are copied when cloning a matrix.

*SUNMatrix* **SUNMatNewEmpty**(*SUNContext* suncctx)

This function allocates a new generic `SUNMatrix` object and initializes its content pointer and the function pointers in the operations structure to NULL.

**Return value:**

If successful, this function returns a `SUNMatrix` object. If an error occurs when allocating the object, then this routine will return NULL.

*SUNErrCode* **SUNMatCopyOps**(*SUNMatrix* A, *SUNMatrix* B)

This function copies the function pointers in the ops structure of A into the ops structure of B.

**Arguments:**

- A – the matrix to copy operations from.
- B – the matrix to copy operations to.

**Return value:**

- A *SUNErrCode*

void **SUNMatFreeEmpty**(*SUNMatrix* A)

This routine frees the generic `SUNMatrix` object, under the assumption that any implementation-specific data that was allocated within the underlying content structure has already been freed. It will additionally test whether the ops pointer is NULL, and, if it is not, it will free it as well.

**Arguments:**

- A – the `SUNMatrix` object to free

type **SUNMatrix\_ID**

Each SUNMATRIX implementation included in SUNDIALS has a unique identifier specified in enumeration and shown in Table 7.1. It is recommended that a user-supplied SUNMATRIX implementation use the `SUNMATRIX_CUSTOM` identifier.

Table 7.1: Identifiers associated with matrix kernels supplied with SUNDIALS

Matrix ID	Matrix type
<code>SUNMATRIX_BAND</code>	Band $M \times M$ matrix
<code>SUNMATRIX_CUSPARSE</code>	CUDA sparse CSR matrix
<code>SUNMATRIX_CUSTOM</code>	User-provided custom matrix
<code>SUNMATRIX_DENSE</code>	Dense $M \times N$ matrix
<code>SUNMATRIX_GINKGO</code>	<code>SUNMatrix</code> wrapper for Ginkgo matrices
<code>SUNMATRIX_MAGMADENSE</code>	Dense $M \times N$ matrix
<code>SUNMATRIX_ONEMKLDENSE</code>	oneMKL dense $M \times N$ matrix
<code>SUNMATRIX_SLUNRLOC</code>	<code>SUNMatrix</code> wrapper for SuperLU_DIST SuperMatrix
<code>SUNMATRIX_SPARSE</code>	Sparse (CSR or CSC) $M \times N$ matrix

## 7.2 Description of the SUNMATRIX operations

For each of the SUNMatrix operations, we give the name, usage of the function, and a description of its mathematical operations below.

### *SUNMatrix\_ID* **SUNMatGetID**(*SUNMatrix* A)

Returns the type identifier for the matrix *A*. It is used to determine the matrix implementation type (e.g. dense, banded, sparse,...) from the abstract SUNMatrix interface. This is used to assess compatibility with SUNDIALS-provided linear solver implementations. Returned values are given in [Table 7.1](#)

Usage:

```
id = SUNMatGetID(A);
```

### *SUNMatrix* **SUNMatClone**(*SUNMatrix* A)

Creates a new SUNMatrix of the same type as an existing matrix *A* and sets the *ops* field. It does not copy the matrix values, but rather allocates storage for the new matrix.

Usage:

```
B = SUNMatClone(A);
```

### void **SUNMatDestroy**(*SUNMatrix* A)

Destroys the SUNMatrix *A* and frees memory allocated for its internal data.

Usage:

```
SUNMatDestroy(A);
```

### *SUNErrCode* **SUNMatSpace**(*SUNMatrix* A, long int \*lrw, long int \*liw)

Returns the storage requirements for the matrix *A*. *lrw* contains the number of sunrealtype words and *liw* contains the number of integer words. The return value denotes success/failure of the operation.

This function is advisory only, for use in determining a user's total space requirements; it could be a dummy function in a user-supplied SUNMatrix module if that information is not of interest.

Usage:

```
retval = SUNMatSpace(A, &lrw, &liw);
```

### *SUNErrCode* **SUNMatZero**(*SUNMatrix* A)

Zeros all entries of the SUNMatrix *A*. The return value denotes the success/failure of the operation:

$$A_{i,j} = 0, \quad i = 1, \dots, m, \quad j = 1, \dots, n.$$

Usage:

```
retval = SUNMatZero(A);
```

### *SUNErrCode* **SUNMatCopy**(*SUNMatrix* A, *SUNMatrix* B)

Performs the operation *B gets A* for all entries of the matrices *A* and *B*. The return value denotes the success/failure of the operation:

$$B_{i,j} = A_{i,j}, \quad i = 1, \dots, m, \quad j = 1, \dots, n.$$

Usage:

```
retval = SUNMatCopy(A,B);
```

**SUNErrCode SUNMatScaleAdd**(*sunrealtype* c, *SUNMatrix* A, *SUNMatrix* B)

Performs the operation *A* gets  $cA + B$ . The return value denotes the success/failure of the operation:

$$A_{i,j} = cA_{i,j} + B_{i,j}, \quad i = 1, \dots, m, \quad j = 1, \dots, n.$$

Usage:

```
retval = SUNMatScaleAdd(c, A, B);
```

**SUNErrCode SUNMatScaleAddI**(*sunrealtype* c, *SUNMatrix* A)

Performs the operation *A* gets  $cA + I$ . The return value denotes the success/failure of the operation:

$$A_{i,j} = cA_{i,j} + \delta_{i,j}, \quad i, j = 1, \dots, n.$$

Usage:

```
retval = SUNMatScaleAddI(c, A);
```

**SUNErrCode SUNMatMatvecSetup**(*SUNMatrix* A)

Performs any setup necessary to perform a matrix-vector product. The return value denotes the success/failure of the operation. It is useful for *SUNMatrix* implementations which need to prepare the matrix itself, or communication structures before performing the matrix-vector product.

Usage:

```
retval = SUNMatMatvecSetup(A);
```

**SUNErrCode SUNMatMatvec**(*SUNMatrix* A, *N\_Vector* x, *N\_Vector* y)

Performs the matrix-vector product *y* gets  $Ax$ . It should only be called with vectors *x* and *y* that are compatible with the matrix *A* – both in storage type and dimensions. The return value denotes the success/failure of the operation:

$$y_i = \sum_{j=1}^n A_{i,j} x_j, \quad i = 1, \dots, m.$$

Usage:

```
retval = SUNMatMatvec(A, x, y);
```

## 7.3 The SUNMATRIX\_DENSE Module

The dense implementation of the *SUNMatrix* module, *SUNMATRIX\_DENSE*, defines the *content* field of *SUNMatrix* to be the following structure:

```
struct _SUNMatrixContent_Dense {
    sunindextype M;
    sunindextype N;
    sunrealtype *data;
    sunindextype ldata;
    sunrealtype **cols;
};
```

These entries of the *content* field contain the following information:

- **M** - number of rows
- **N** - number of columns
- **data** - pointer to a contiguous block of `sunrealtype` variables. The elements of the dense matrix are stored columnwise, i.e. the  $(i, j)$  element of a dense `SUNMatrix` object (with  $0 \leq i < M$  and  $0 \leq j < N$ ) may be accessed via `data[j*M+i]`.
- **ldata** - length of the data array ( $= M N$ ).
- **cols** - array of pointers. `cols[j]` points to the first element of the  $j$ -th column of the matrix in the array `data`. The  $(i, j)$  element of a dense `SUNMatrix` (with  $0 \leq i < M$  and  $0 \leq j < N$ ) may be accessed may be accessed via `cols[j][i]`.

The header file to be included when using this module is `sunmatrix/sunmatrix_dense.h`.

The following macros are provided to access the content of a `SUNMATRIX_DENSE` matrix. The prefix `SM_` in the names denotes that these macros are for *SUNMatrix* implementations, and the suffix `_D` denotes that these are specific to the *dense* version.

#### **SM\_CONTENT\_D(A)**

This macro gives access to the contents of the dense `SUNMatrix` *A*.

The assignment `A_cont = SM_CONTENT_D(A)` sets `A_cont` to be a pointer to the dense `SUNMatrix` content structure.

Implementation:

```
#define SM_CONTENT_D(A)    ( (SUNMatrixContent_Dense)(A->content) )
```

#### **SM\_ROWS\_D(A)**

Access the number of rows in the dense `SUNMatrix` *A*.

This may be used either to retrieve or to set the value. For example, the assignment `A_rows = SM_ROWS_D(A)` sets `A_rows` to be the number of rows in the matrix *A*. Similarly, the assignment `SM_ROWS_D(A) = A_rows` sets the number of columns in *A* to equal `A_rows`.

Implementation:

```
#define SM_ROWS_D(A)      ( SM_CONTENT_D(A)->M )
```

#### **SM\_COLUMNS\_D(A)**

Access the number of columns in the dense `SUNMatrix` *A*.

This may be used either to retrieve or to set the value. For example, the assignment `A_columns = SM_COLUMNS_D(A)` sets `A_columns` to be the number of columns in the matrix *A*. Similarly, the assignment `SM_COLUMNS_D(A) = A_columns` sets the number of columns in *A* to equal `A_columns`.

Implementation:

```
#define SM_COLUMNS_D(A)   ( SM_CONTENT_D(A)->N )
```

#### **SM\_LDATA\_D(A)**

Access the total data length in the dense `SUNMatrix` *A*.

This may be used either to retrieve or to set the value. For example, the assignment `A_ldata = SM_LDATA_D(A)` sets `A_ldata` to be the length of the data array in the matrix *A*. Similarly, the assignment `SM_LDATA_D(A) = A_ldata` sets the parameter for the length of the data array in *A* to equal `A_ldata`.

Implementation:

```
#define SM_LDATA_D(A)    ( SM_CONTENT_D(A)->ldata )
```

### SM\_DATA\_D(A)

This macro gives access to the data pointer for the matrix entries.

The assignment `A_data = SM_DATA_D(A)` sets `A_data` to be a pointer to the first component of the data array for the dense `SUNMatrix` `A`. The assignment `SM_DATA_D(A) = A_data` sets the data array of `A` to be `A_data` by storing the pointer `A_data`.

Implementation:

```
#define SM_DATA_D(A)    ( SM_CONTENT_D(A)->data )
```

### SM\_COLS\_D(A)

This macro gives access to the cols pointer for the matrix entries.

The assignment `A_cols = SM_COLS_D(A)` sets `A_cols` to be a pointer to the array of column pointers for the dense `SUNMatrix` `A`. The assignment `SM_COLS_D(A) = A_cols` sets the column pointer array of `A` to be `A_cols` by storing the pointer `A_cols`.

Implementation:

```
#define SM_COLS_D(A)    ( SM_CONTENT_D(A)->cols )
```

### SM\_COLUMN\_D(A)

This macros gives access to the individual columns of the data array of a dense `SUNMatrix`.

The assignment `col_j = SM_COLUMN_D(A, j)` sets `col_j` to be a pointer to the first entry of the  $j$ -th column of the  $M \times N$  dense matrix `A` (with  $0 \leq j < N$ ). The type of the expression `SM_COLUMN_D(A, j)` is `sunrealtype *`. The pointer returned by the call `SM_COLUMN_D(A, j)` can be treated as an array which is indexed from 0 to  $M-1$ .

Implementation:

```
#define SM_COLUMN_D(A, j) ( (SM_CONTENT_D(A)->cols)[j] )
```

### SM\_ELEMENT\_D(A)

This macro gives access to the individual entries of the data array of a dense `SUNMatrix`.

The assignments `SM_ELEMENT_D(A, i, j) = a_ij` and `a_ij = SM_ELEMENT_D(A, i, j)` reference the  $A_{i,j}$  element of the  $M \times N$  dense matrix `A` (with  $0 \leq i < M$  and  $0 \leq j < N$ ).

Implementation:

```
#define SM_ELEMENT_D(A, i, j) ( (SM_CONTENT_D(A)->cols)[j][i] )
```

The `SUNMATRIX_DENSE` module defines dense implementations of all matrix operations listed in §7.2. Their names are obtained from those in that section by appending the suffix `_Dense` (e.g. `SUNMatCopy_Dense`). The module `SUNMATRIX_DENSE` provides the following additional user-callable routines:

*SUNMatrix* **SUNDenseMatrix**(*sunindextype* M, *sunindextype* N, *SUNContext* sunctx)

This constructor function creates and allocates memory for a dense `SUNMatrix`. Its arguments are the number of rows, `M`, and columns, `N`, for the dense matrix.

void **SUNDenseMatrix\_Print**(*SUNMatrix* A, FILE \*outfile)

This function prints the content of a dense `SUNMatrix` to the output stream specified by `outfile`. Note: `stdout` or `stderr` may be used as arguments for `outfile` to print directly to standard output or standard error, respectively.

*sunindextype* **SUNDenseMatrix\_Rows**(*SUNMatrix* A)

This function returns the number of rows in the dense *SUNMatrix*.

*sunindextype* **SUNDenseMatrix\_Columns**(*SUNMatrix* A)

This function returns the number of columns in the dense *SUNMatrix*.

*sunindextype* **SUNDenseMatrix\_LData**(*SUNMatrix* A)

This function returns the length of the data array for the dense *SUNMatrix*.

*sunrealtype* \***SUNDenseMatrix\_Data**(*SUNMatrix* A)

This function returns a pointer to the data array for the dense *SUNMatrix*.

*sunrealtype* \*\***SUNDenseMatrix\_Cols**(*SUNMatrix* A)

This function returns a pointer to the cols array for the dense *SUNMatrix*.

*sunrealtype* \***SUNDenseMatrix\_Column**(*SUNMatrix* A, *sunindextype* j)

This function returns a pointer to the first entry of the *j*th column of the dense *SUNMatrix*. The resulting pointer should be indexed over the range 0 to *M*-1.

### Notes

- When looping over the components of a dense *SUNMatrix* A, the most efficient approaches are to:
  - First obtain the component array via *A\_data* = **SUNDenseMatrix\_Data**(A), or equivalently *A\_data* = **SM\_DATA\_D**(A), and then access *A\_data*[*i*] within the loop.
  - First obtain the array of column pointers via *A\_cols* = **SUNDenseMatrix\_Cols**(A), or equivalently *A\_cols* = **SM\_COLS\_D**(A), and then access *A\_cols*[*j*][*i*] within the loop.
  - Within a loop over the columns, access the column pointer via *A\_colj* = **SUNDenseMatrix\_Column**(A, *j*) and then to access the entries within that column using *A\_colj*[*i*] within the loop.

All three of these are more efficient than using **SM\_ELEMENT\_D**(A, *i*, *j*) within a double loop.

- Within the **SUNMatMatvec\_Dense** routine, internal consistency checks are performed to ensure that the matrix is called with consistent *N\_Vector* implementations. These are currently limited to: **NVECTOR\_SERIAL**, **NVECTOR\_OPENMP**, and **NVECTOR\_PTHREADS**. As additional compatible vector implementations are added to **SUNDIALS**, these will be included within this compatibility check.

## 7.4 The SUNMATRIX\_MAGMADENSE Module

The **SUNMATRIX\_MAGMADENSE** module interfaces to the **MAGMA** linear algebra library and can target NVIDIA's CUDA programming model or AMD's HIP programming model [62]. All data stored by this matrix implementation resides on the GPU at all times. The implementation currently supports a standard LAPACK column-major storage format as well as a low-storage format for block-diagonal matrices

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_0 & 0 & \cdots & 0 \\ 0 & \mathbf{A}_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mathbf{A}_{n-1} \end{bmatrix}$$

This matrix implementation is best paired with the *SUNLinearSolver\_MagmaDense* **SUNLinearSolver**.

The header file to include when using this module is `sunmatrix/sunmatrix_magmadense.h`. The installed library to link to is `libsundials_sunmatrixmagmadense.lib` where `lib` is typically `.so` for shared libraries and `.a` for static libraries.



**Warning:** The SUNMATRIX\_MAGMADENSE module is experimental and subject to change.

### 7.4.1 SUNMATRIX\_MAGMADENSE Functions

The SUNMATRIX\_MAGMADENSE module defines GPU-enabled implementations of all matrix operations listed in §7.2.

- `SUNMatGetID_MagmaDense` – returns `SUNMATRIX_MAGMADENSE`
- `SUNMatClone_MagmaDense`
- `SUNMatDestroy_MagmaDense`
- `SUNMatZero_MagmaDense`
- `SUNMatCopy_MagmaDense`
- `SUNMatScaleAdd_MagmaDense`
- `SUNMatScaleAddI_MagmaDense`
- `SUNMatMatvecSetup_MagmaDense`
- `SUNMatMatvec_MagmaDense`
- `SUNMatSpace_MagmaDense`

In addition, the SUNMATRIX\_MAGMADENSE module defines the following implementation specific functions:

*SUNMatrix* **SUNMatrix\_MagmaDense**(*sunindextype* M, *sunindextype* N, *SUNMemoryType* memtype, *SUNMemoryHelper* memhelper, void \*queue, *SUNContext* sunctx)

This constructor function creates and allocates memory for an  $M \times N$  SUNMATRIX\_MAGMADENSE `SUNMatrix`.

**Arguments:**

- $M$  – the number of matrix rows.
- $N$  – the number of matrix columns.
- *memtype* – the type of memory to use for the matrix data; can be `SUNMEMTYPE_UVM` or `SUNMEMTYPE_DEVICE`.
- *memhelper* – the memory helper used for allocating data.
- *queue* – a `cudaStream_t` when using CUDA or a `hipStream_t` when using HIP.
- *sunctx* – the *SUNContext* object (see §4.2)

**Return value:**

If successful, a `SUNMatrix` object otherwise `NULL`.

*SUNMatrix* **SUNMatrix\_MagmaDenseBlock**(*sunindextype* nblocks, *sunindextype* M\_block, *sunindextype* N\_block, *SUNMemoryType* memtype, *SUNMemoryHelper* memhelper, void \*queue, *SUNContext* sunctx)

This constructor function creates and allocates memory for a block diagonal SUNMATRIX\_MAGMADENSE `SUNMatrix` with *nblocks* of size  $M \times N$ .

**Arguments:**

- *nblocks* – the number of matrix rows.
- *M\_block* – the number of matrix rows in each block.

- $N_{block}$  – the number of matrix columns in each block.
- *memtype* – the type of memory to use for the matrix data; can be `SUNMEMTYPE_UVM` or `SUNMEMTYPE_DEVICE`.
- *memhelper* – the memory helper used for allocating data.
- *queue* – a `cudaStream_t` when using CUDA or a `hipStream_t` when using HIP.
- *sunctx* – the `SUNContext` object (see §4.2)

**Return value:**

If successful, a `SUNMatrix` object otherwise `NULL`.

*sunindextype* **SUNMatrix\_MagmaDense\_Rows**(*SUNMatrix* A)

This function returns the number of rows in the `SUNMatrix` object. For block diagonal matrices, the number of rows is computed as  $M_{block} \times nblocks$ .

**Arguments:**

- A – a `SUNMatrix` object.

**Return value:**

If successful, the number of rows in the `SUNMatrix` object otherwise `SUNMATRIX_ILL_INPUT`.

*sunindextype* **SUNMatrix\_MagmaDense\_Columns**(*SUNMatrix* A)

This function returns the number of columns in the `SUNMatrix` object. For block diagonal matrices, the number of columns is computed as  $N_{block} \times nblocks$ .

**Arguments:**

- A – a `SUNMatrix` object.

**Return value:**

If successful, the number of columns in the `SUNMatrix` object otherwise `SUNMATRIX_ILL_INPUT`.

*sunindextype* **SUNMatrix\_MagmaDense\_BlockRows**(*SUNMatrix* A)

This function returns the number of rows in a block of the `SUNMatrix` object.

**Arguments:**

- A – a `SUNMatrix` object.

**Return value:**

If successful, the number of rows in a block of the `SUNMatrix` object otherwise `SUNMATRIX_ILL_INPUT`.

*sunindextype* **SUNMatrix\_MagmaDense\_BlockColumns**(*SUNMatrix* A)

This function returns the number of columns in a block of the `SUNMatrix` object.

**Arguments:**

- A – a `SUNMatrix` object.

**Return value:**

If successful, the number of columns in a block of the `SUNMatrix` object otherwise `SUNMATRIX_ILL_INPUT`.

*sunindextype* **SUNMatrix\_MagmaDense\_LData**(*SUNMatrix* A)

This function returns the length of the `SUNMatrix` data array.

**Arguments:**

- A – a `SUNMatrix` object.

**Return value:**

If successful, the length of the SUNMatrix data array otherwise SUNMATRIX\_ILL\_INPUT.

*sunindextype* **SUNMatrix\_MagmaDense\_NumBlocks**(*SUNMatrix* A)

This function returns the number of blocks in the SUNMatrix object.

**Arguments:**

- A – a SUNMatrix object.

**Return value:**

If successful, the number of blocks in the SUNMatrix object otherwise SUNMATRIX\_ILL\_INPUT.

*sunrealtype* \***SUNMatrix\_MagmaDense\_Data**(*SUNMatrix* A)

This function returns the SUNMatrix data array.

**Arguments:**

- A – a SUNMatrix object.

**Return value:**

If successful, the SUNMatrix data array otherwise NULL.

*sunrealtype* \*\***SUNMatrix\_MagmaDense\_BlockData**(*SUNMatrix* A)

This function returns an array of pointers that point to the start of the data array for each block in the SUNMatrix.

**Arguments:**

- A – a SUNMatrix object.

**Return value:**

If successful, an array of data pointers to each of the SUNMatrix blocks otherwise NULL.

*sunrealtype* \***SUNMatrix\_MagmaDense\_Block**(*SUNMatrix* A, *sunindextype* k)

This function returns a pointer to the data array for block  $k$  in the SUNMatrix.

**Arguments:**

- A – a SUNMatrix object.
- $k$  – the block index.

**Return value:**

If successful, a pointer to the data array for the SUNMatrix block otherwise NULL.

---

**Note:** No bounds-checking is performed by this function,  $j$  should be strictly less than  $nblocks$ .

---

*sunrealtype* \***SUNMatrix\_MagmaDense\_Column**(*SUNMatrix* A, *sunindextype* j)

This function returns a pointer to the data array for column  $j$  in the SUNMatrix.

**Arguments:**

- A – a SUNMatrix object.
- $j$  – the column index.

**Return value:**

If successful, a pointer to the data array for the SUNMatrix column otherwise NULL.

---

**Note:** No bounds-checking is performed by this function,  $j$  should be strictly less than  $nblocks * N_{block}$ .

---

*sunrealtype* \*SUNMatrix\_MagmaDense\_BlockColumn(*SUNMatrix* A, *sunindextype* k, *sunindextype* j)

This function returns a pointer to the data array for column  $j$  of block  $k$  in the *SUNMatrix*.

**Arguments:**

- $A$  – a *SUNMatrix* object.
- $k$  – the block index.
- $j$  – the column index.

**Return value:**

If successful, a pointer to the data array for the *SUNMatrix* column otherwise NULL.

---

**Note:** No bounds-checking is performed by this function,  $k$  should be strictly less than  $nblocks$  and  $j$  should be strictly less than  $N_{block}$ .

---

*SUNErrCode* SUNMatrix\_MagmaDense\_CopyToDevice(*SUNMatrix* A, *sunrealtype* \*h\_data)

This function copies the matrix data to the GPU device from the provided host array.

**Arguments:**

- $A$  – a *SUNMatrix* object
- $h\_data$  – a host array pointer to copy data from.

**Return value:**

- *SUN\_SUCCESS* – if the copy is successful.
- *SUN\_ERR\_ARG\_INCOMPATIBLE* – if the *SUNMatrix* is not a *SUNMATRIX\_MAGMADENSE* matrix.
- *SUN\_ERR\_MEM\_FAIL* – if the copy fails.

*SUNErrCode* SUNMatrix\_MagmaDense\_CopyFromDevice(*SUNMatrix* A, *sunrealtype* \*h\_data)

This function copies the matrix data from the GPU device to the provided host array.

**Arguments:**

- $A$  – a *SUNMatrix* object
- $h\_data$  – a host array pointer to copy data to.

**Return value:**

- *SUN\_SUCCESS* – if the copy is successful.
- *SUN\_ERR\_ARG\_INCOMPATIBLE* – if the *SUNMatrix* is not a *SUNMATRIX\_MAGMADENSE* matrix.
- *SUN\_ERR\_MEM\_FAIL* – if the copy fails.

## 7.4.2 SUNMATRIX\_MAGMADENSE Usage Notes

**Warning:** When using the *SUNMATRIX\_MAGMADENSE* module with a *SUNDIALS* package (e.g. *CVODE*), the stream given to matrix should be the same stream used for the *NVECTOR* object that is provided to the package, and the *NVECTOR* object given to the *SUNMatvec* operation. If different streams are utilized, synchronization issues may occur.

## 7.5 The SUNMATRIX\_ONEMKLDENSE Module

The SUNMATRIX\_ONEMKLDENSE module is intended for interfacing with direct linear solvers from the [Intel oneAPI Math Kernel Library \(oneMKL\)](#) using the SYCL (DPC++) programming model. The implementation currently supports a standard LAPACK column-major storage format as well as a low-storage format for block-diagonal matrices,

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_0 & 0 & \cdots & 0 \\ 0 & \mathbf{A}_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mathbf{A}_{n-1} \end{bmatrix}$$

This matrix implementation is best paired with the [SUNLinearSolver\\_OneMklDense](#) linear solver.

The header file to include when using this class is `sunmatrix/sunmatrix_onemkldense.h`. The installed library to link to is `libsundials_sunmatrixonemkldense.lib` where `lib` is typically `.so` for shared libraries and `.a` for static libraries.

**Warning:** The SUNMATRIX\_ONEMKLDENSE class is experimental and subject to change.

### 7.5.1 SUNMATRIX\_ONEMKLDENSE Functions

The SUNMATRIX\_ONEMKLDENSE class defines implementations of the following matrix operations listed in §7.2.

- `SUNMatGetID_OneMklDense` – returns `SUNMATRIX_ONEMKLDENSE`
- `SUNMatClone_OneMklDense`
- `SUNMatDestroy_OneMklDense`
- `SUNMatZero_OneMklDense`
- `SUNMatCopy_OneMklDense`
- `SUNMatScaleAdd_OneMklDense`
- `SUNMatScaleAddI_OneMklDense`
- `SUNMatMatvec_OneMklDense`
- `SUNMatSpace_OneMklDense`

In addition, the SUNMATRIX\_ONEMKLDENSE class defines the following implementation specific functions.

#### 7.5.1.1 Constructors

SUNMatrix **SUNMatrix\_OneMklDense**(sunindextype M, sunindextype N, SUNMemoryType memtype, SUNMemoryHelper memhelper, sycl::queue \*queue, SUNContext sunctx)

This constructor function creates and allocates memory for an  $M \times N$  SUNMATRIX\_ONEMKLDENSE SUNMatrix.

**Arguments:**

- $M$  – the number of matrix rows.
- $N$  – the number of matrix columns.
- *memtype* – the type of memory to use for the matrix data; can be `SUNMEMTYPE_UVM` or `SUNMEMTYPE_DEVICE`.

- *memhelper* – the memory helper used for allocating data.
- *queue* – the SYCL queue to which operations will be submitted.
- *sunctx* – the *SUNContext* object (see §4.2)

**Return value:**

If successful, a *SUNMatrix* object otherwise NULL.

*SUNMatrix* **SUNMatrix\_OneMklDenseBlock**(sunindextype nblocks, sunindextype M\_block, sunindextype N\_block, SUNMemoryType memtype, SUNMemoryHelper memhelper, sycl::queue \*queue, SUNContext sunctx)

This constructor function creates and allocates memory for a block diagonal *SUNMATRIX\_ONEMKLDENSE* *SUNMatrix* with *nblocks* of size  $M_{block} \times N_{block}$ .

**Arguments:**

- *nblocks* – the number of matrix rows.
- *M\_block* – the number of matrix rows in each block.
- *N\_block* – the number of matrix columns in each block.
- *memtype* – the type of memory to use for the matrix data; can be *SUNMEMTYPE\_UVM* or *SUNMEMTYPE\_DEVICE*.
- *memhelper* – the memory helper used for allocating data.
- *queue* – the SYCL queue to which operations will be submitted.
- *sunctx* – the *SUNContext* object (see §4.2)

**Return value:**

If successful, a *SUNMatrix* object otherwise NULL.

### 7.5.1.2 Access Matrix Dimensions

*sunindextype* **SUNMatrix\_OneMklDense\_Rows**(*SUNMatrix* A)

This function returns the number of rows in the *SUNMatrix* object. For block diagonal matrices, the number of rows is computed as  $M_{block} \times nblocks$ .

**Arguments:**

- *A* – a *SUNMatrix* object.

**Return value:**

If successful, the number of rows in the *SUNMatrix* object otherwise *SUNMATRIX\_ILL\_INPUT*.

*sunindextype* **SUNMatrix\_OneMklDense\_Columns**(*SUNMatrix* A)

This function returns the number of columns in the *SUNMatrix* object. For block diagonal matrices, the number of columns is computed as  $N_{block} \times nblocks$ .

**Arguments:**

- *A* – a *SUNMatrix* object.

**Return value:**

If successful, the number of columns in the *SUNMatrix* object otherwise *SUNMATRIX\_ILL\_INPUT*.

### 7.5.1.3 Access Matrix Block Dimensions

*sunindextype* **SUNMatrix\_OneMklDense\_NumBlocks**(*SUNMatrix* A)

This function returns the number of blocks in the *SUNMatrix* object.

**Arguments:**

- A – a *SUNMatrix* object.

**Return value:**

If successful, the number of blocks in the *SUNMatrix* object otherwise *SUNMATRIX\_ILL\_INPUT*.

*sunindextype* **SUNMatrix\_OneMklDense\_BlockRows**(*SUNMatrix* A)

This function returns the number of rows in a block of the *SUNMatrix* object.

**Arguments:**

- A – a *SUNMatrix* object.

**Return value:**

If successful, the number of rows in a block of the *SUNMatrix* object otherwise *SUNMATRIX\_ILL\_INPUT*.

*sunindextype* **SUNMatrix\_OneMklDense\_BlockColumns**(*SUNMatrix* A)

This function returns the number of columns in a block of the *SUNMatrix* object.

**Arguments:**

- A – a *SUNMatrix* object.

**Return value:**

If successful, the number of columns in a block of the *SUNMatrix* object otherwise *SUNMATRIX\_ILL\_INPUT*.

### 7.5.1.4 Access Matrix Data

*sunindextype* **SUNMatrix\_OneMklDense\_LData**(*SUNMatrix* A)

This function returns the length of the *SUNMatrix* data array.

**Arguments:**

- A – a *SUNMatrix* object.

**Return value:**

If successful, the length of the *SUNMatrix* data array otherwise *SUNMATRIX\_ILL\_INPUT*.

*sunrealtype* \***SUNMatrix\_OneMklDense\_Data**(*SUNMatrix* A)

This function returns the *SUNMatrix* data array.

**Arguments:**

- A – a *SUNMatrix* object.

**Return value:**

If successful, the *SUNMatrix* data array otherwise *NULL*.

*sunrealtype* \***SUNMatrix\_OneMklDense\_Column**(*SUNMatrix* A, *sunindextype* j)

This function returns a pointer to the data array for column *j* in the *SUNMatrix*.

**Arguments:**

- A – a *SUNMatrix* object.
- *j* – the column index.

**Return value:**

If successful, a pointer to the data array for the `SUNMatrix` column otherwise `NULL`.

---

**Note:** No bounds-checking is performed by this function,  $j$  should be strictly less than  $nblocks * N_{\text{block}}$ .

---

### 7.5.1.5 Access Matrix Block Data

*sunindextype* `SUNMatrix_OneMklDense_BlockData(SUNMatrix A)`

This function returns the length of the `SUNMatrix` data array for each block of the `SUNMatrix` object.

**Arguments:**

- $A$  – a `SUNMatrix` object.

**Return value:**

If successful, the length of the `SUNMatrix` data array for each block otherwise `SUNMATRIX_ILL_INPUT`.

*sunrealtype* `**SUNMatrix_OneMklDense_BlockData(SUNMatrix A)`

This function returns an array of pointers that point to the start of the data array for each block in the `SUNMatrix`.

**Arguments:**

- $A$  – a `SUNMatrix` object.

**Return value:**

If successful, an array of data pointers to each of the `SUNMatrix` blocks otherwise `NULL`.

*sunrealtype* `*SUNMatrix_OneMklDense_Block(SUNMatrix A, sunindextype k)`

This function returns a pointer to the data array for block  $k$  in the `SUNMatrix`.

**Arguments:**

- $A$  – a `SUNMatrix` object.
- $k$  – the block index.

**Return value:**

If successful, a pointer to the data array for the `SUNMatrix` block otherwise `NULL`.

---

**Note:** No bounds-checking is performed by this function,  $j$  should be strictly less than  $nblocks$ .

---

*sunrealtype* `*SUNMatrix_OneMklDense_BlockColumn(SUNMatrix A, sunindextype k, sunindextype j)`

This function returns a pointer to the data array for column  $j$  of block  $k$  in the `SUNMatrix`.

**Arguments:**

- $A$  – a `SUNMatrix` object.
- $k$  – the block index.
- $j$  – the column index.

**Return value:**

If successful, a pointer to the data array for the `SUNMatrix` column otherwise `NULL`.

---

**Note:** No bounds-checking is performed by this function,  $k$  should be strictly less than  $nblocks$  and  $j$  should be strictly less than  $N_{\text{block}}$ .

---



### 7.5.1.6 Copy Data

*SUNErrCode* **SUNMatrix\_OneMklDense\_CopyToDevice**(*SUNMatrix* A, *sunrealtype* \*h\_data)

This function copies the matrix data to the GPU device from the provided host array.

**Arguments:**

- A – a *SUNMatrix* object
- h\_data – a host array pointer to copy data from.

**Return value:**

- SUN\_SUCCESS – if the copy is successful.
- SUN\_ERR\_ARG\_INCOMPATIBLE – if the *SUNMatrix* is not a *SUNMATRIX\_ONEMKLDENSE* matrix.
- SUN\_ERR\_MEM\_FAIL – if the copy fails.

*SUNErrCode* **SUNMatrix\_OneMklDense\_CopyFromDevice**(*SUNMatrix* A, *sunrealtype* \*h\_data)

This function copies the matrix data from the GPU device to the provided host array.

**Arguments:**

- A – a *SUNMatrix* object
- h\_data – a host array pointer to copy data to.

**Return value:**

- SUN\_SUCCESS – if the copy is successful.
- SUN\_ERR\_ARG\_INCOMPATIBLE – if the *SUNMatrix* is not a *SUNMATRIX\_ONEMKLDENSE* matrix.
- SUN\_ERR\_MEM\_FAIL – if the copy fails.

## 7.5.2 SUNMATRIX\_ONEMKLDENSE Usage Notes

**Warning:** The *SUNMATRIX\_ONEMKLDENSE* class only supports 64-bit indexing, thus *SUNDIALS* must be built for 64-bit indexing to use this class.

When using the *SUNMATRIX\_ONEMKLDENSE* class with a *SUNDIALS* package (e.g. *CVODE*), the queue given to matrix should be the same stream used for the *NVECTOR* object that is provided to the package, and the *NVECTOR* object given to the *SUNMatMatvec()* operation. If different streams are utilized, synchronization issues may occur.

## 7.6 The SUNMATRIX\_BAND Module

The banded implementation of the *SUNMatrix* module, *SUNMATRIX\_BAND*, defines the *content* field of *SUNMatrix* to be the following structure:

```
struct _SUNMatrixContent_Band {
    sunindextype M;
    sunindextype N;
    sunindextype mu;
    sunindextype ml;
    sunindextype smu;
```

(continues on next page)

(continued from previous page)

```

sunindextype ldim;
sunrealtype *data;
sunindextype ldata;
sunrealtype **cols;
};

```

A diagram of the underlying data representation in a banded matrix is shown in Fig. 7.1. A more complete description of the parts of this *content* field is given below:

- **M** - number of rows
- **N** - number of columns ( $N = M$ )
- **mu** - upper half-bandwidth,  $0 \leq \text{mu} < N$
- **ml** - lower half-bandwidth,  $0 \leq \text{ml} < N$
- **smu** - storage upper bandwidth,  $\text{mu} \leq \text{smu} < N$ . The LU decomposition routines in the associated *SUNLINSOL\_BAND* and *SUNLINSOL\_LAPACKBAND* modules write the LU factors into the existing storage for the band matrix. The upper triangular factor *U*, however, may have an upper bandwidth as big as  $\min(N-1, \text{mu}+\text{ml})$  because of partial pivoting. The **smu** field holds the upper half-bandwidth allocated for the band matrix.
- **ldim** - leading dimension ( $\text{ldim} \geq \text{smu} + \text{ml} + 1$ )
- **data** - pointer to a contiguous block of *sunrealtype* variables. The elements of the banded matrix are stored columnwise (i.e. columns are stored one on top of the other in memory). Only elements within the specified half-bandwidths are stored. **data** is a pointer to **ldata** contiguous locations which hold the elements within the banded matrix.
- **ldata** - length of the data array ( $= \text{ldim} \cdot N$ )
- **cols** - array of pointers. **cols[j]** is a pointer to the uppermost element within the band in the *j*-th column. This pointer may be treated as an array indexed from **smu**-**mu** (to access the uppermost element within the band in the *j*-th column) to **smu**+**ml** (to access the lowest element within the band in the *j*-th column). Indices from 0 to **smu**-**mu**-1 give access to extra storage elements required by the LU decomposition function. Finally, **cols[j][i-j+smu]** is the  $(i, j)$ -th element with  $j - \text{mu} \leq i \leq j + \text{ml}$ .

The header file to be included when using this module is *sunmatrix/sunmatrix\_band.h*.

The following macros are provided to access the content of a *SUNMATRIX\_BAND* matrix. The prefix **SM\_** in the names denotes that these macros are for *SUNMatrix* implementations, and the suffix **\_B** denotes that these are specific to the *banded* version.

#### **SM\_CONTENT\_B(A)**

This macro gives access to the contents of the banded *SUNMatrix* **A**.

The assignment **A\_cont = SM\_CONTENT\_B(A)** sets **A\_cont** to be a pointer to the banded *SUNMatrix* content structure.

Implementation:

```
#define SM_CONTENT_B(A)    ( (SUNMatrixContent_Band)(A->content) )
```

#### **SM\_ROWS\_B(A)**

Access the number of rows in the banded *SUNMatrix* **A**.

This may be used either to retrieve or to set the value. For example, the assignment **A\_rows = SM\_ROWS\_B(A)** sets **A\_rows** to be the number of rows in the matrix **A**. Similarly, the assignment **SM\_ROWS\_B(A) = A\_rows** sets the number of columns in **A** to equal **A\_rows**.

Implementation:

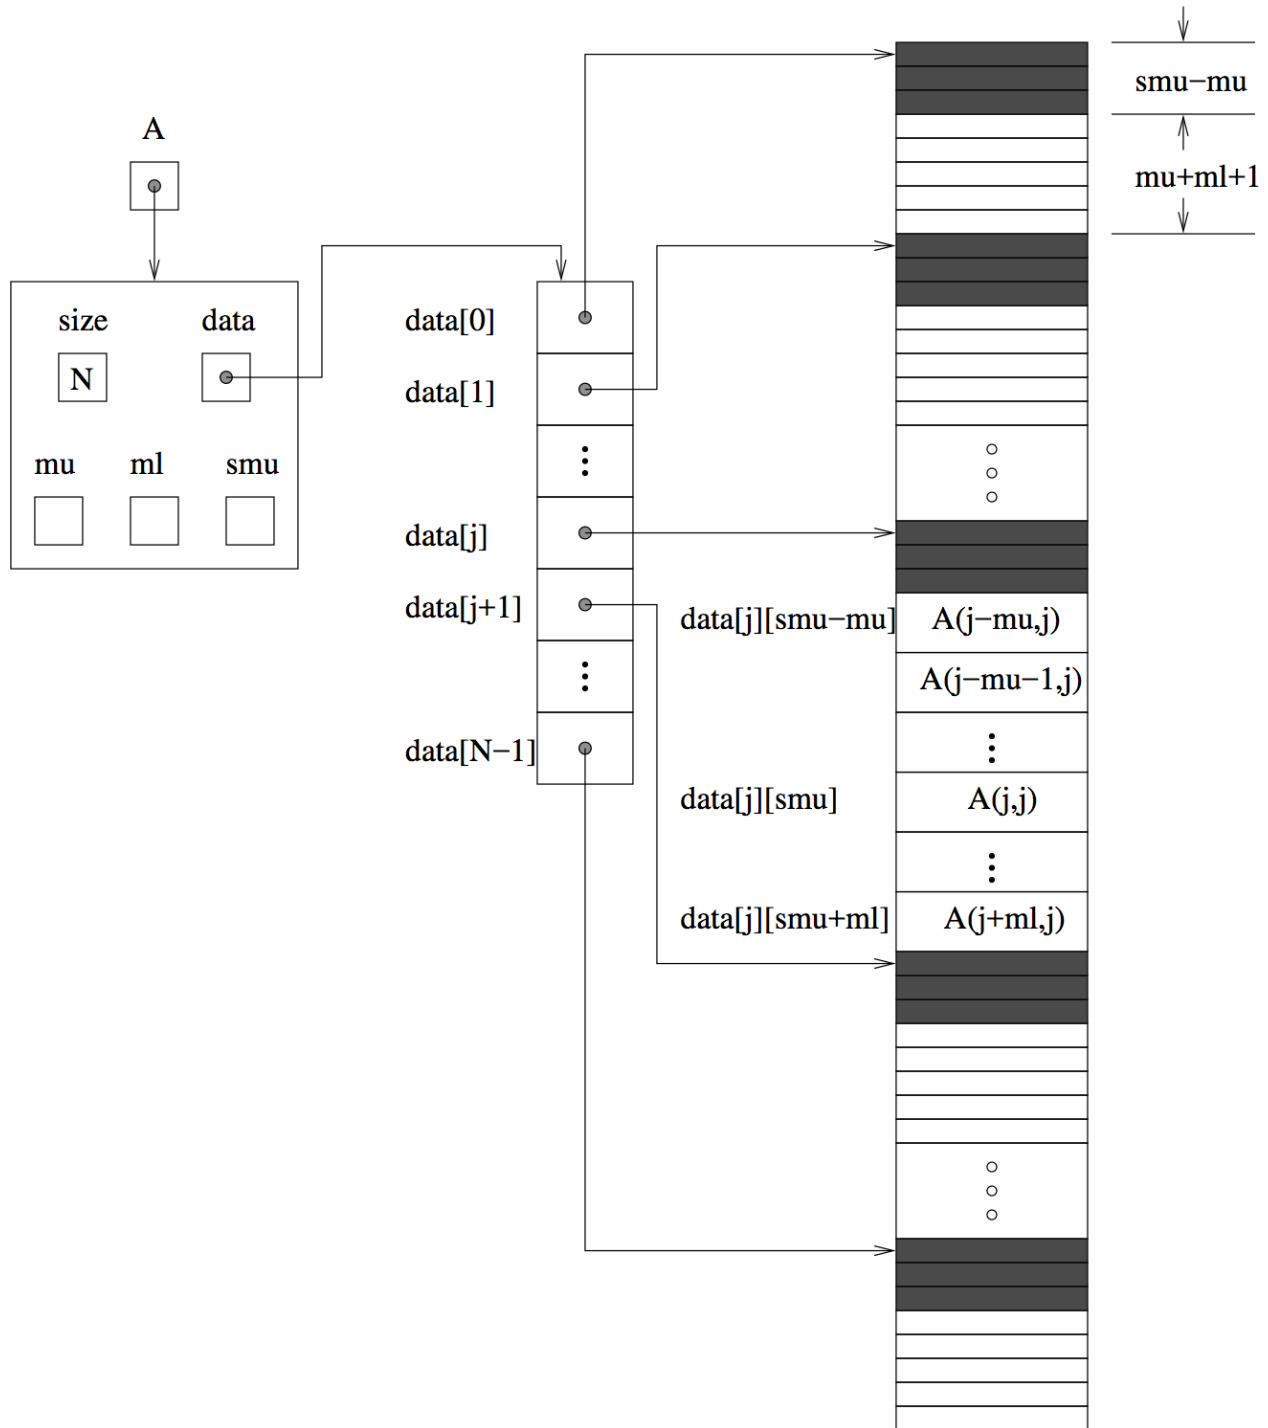


Fig. 7.1: Diagram of the storage for the `SUNMATRIX_BAND` module. Here  $A$  is an  $N \times N$  band matrix with upper and lower half-bandwidths  $\mu$  and  $m$ , respectively. The rows and columns of  $A$  are numbered from 0 to  $N-1$  and the  $(i, j)$ -th element of  $A$  is denoted  $A(i, j)$ . The greyed out areas of the underlying component storage are used by the associated `SUNLINSOL_BAND` or `SUNLINSOL_LAPACKBAND` linear solver.

```
#define SM_ROWS_B(A)    ( SM_CONTENT_B(A)->M )
```

**SM\_COLUMNS\_B(A)**

Access the number of columns in the banded SUNMatrix A. As with SM\_ROWS\_B, this may be used either to retrieve or to set the value.

Implementation:

```
#define SM_COLUMNS_B(A) ( SM_CONTENT_B(A)->N )
```

**SM\_UBAND\_B(A)**

Access the mu parameter in the banded SUNMatrix A. As with SM\_ROWS\_B, this may be used either to retrieve or to set the value.

Implementation:

```
#define SM_UBAND_B(A)   ( SM_CONTENT_B(A)->mu )
```

**SM\_LBAND\_B(A)**

Access the ml parameter in the banded SUNMatrix A. As with SM\_ROWS\_B, this may be used either to retrieve or to set the value.

Implementation:

```
#define SM_LBAND_B(A)   ( SM_CONTENT_B(A)->ml )
```

**SM\_SUBAND\_B(A)**

Access the smu parameter in the banded SUNMatrix A. As with SM\_ROWS\_B, this may be used either to retrieve or to set the value.

Implementation:

```
#define SM_SUBAND_B(A)  ( SM_CONTENT_B(A)->smu )
```

**SM\_LDIM\_B(A)**

Access the ldim parameter in the banded SUNMatrix A. As with SM\_ROWS\_B, this may be used either to retrieve or to set the value.

Implementation:

```
#define SM_LDIM_B(A)    ( SM_CONTENT_B(A)->ldim )
```

**SM\_LDATA\_B(A)**

Access the ldata parameter in the banded SUNMatrix A. As with SM\_ROWS\_B, this may be used either to retrieve or to set the value.

Implementation:

```
#define SM_LDATA_B(A)   ( SM_CONTENT_B(A)->ldata )
```

**SM\_DATA\_B(A)**

This macro gives access to the data pointer for the matrix entries.

The assignment `A_data = SM_DATA_B(A)` sets `A_data` to be a pointer to the first component of the data array for the banded SUNMatrix A. The assignment `SM_DATA_B(A) = A_data` sets the data array of A to be `A_data` by storing the pointer `A_data`.

Implementation:

```
#define SM_DATA_B(A)    ( SM_CONTENT_B(A)->data )
```

### SM\_COLS\_B(A)

This macro gives access to the cols pointer for the matrix entries.

The assignment `A_cols = SM_COLS_B(A)` sets `A_cols` to be a pointer to the array of column pointers for the banded SUNMatrix `A`. The assignment `SM_COLS_B(A) = A_cols` sets the column pointer array of `A` to be `A_cols` by storing the pointer `A_cols`.

Implementation:

```
#define SM_COLS_B(A)    ( SM_CONTENT_B(A)->cols )
```

### SM\_COLUMN\_B(A)

This macros gives access to the individual columns of the data array of a banded SUNMatrix.

The assignment `col_j = SM_COLUMN_B(A, j)` sets `col_j` to be a pointer to the diagonal element of the  $j$ -th column of the  $N \times N$  band matrix `A`,  $0 \leq j \leq N - 1$ . The type of the expression `SM_COLUMN_B(A, j)` is `sunrealtype *`. The pointer returned by the call `SM_COLUMN_B(A, j)` can be treated as an array which is indexed from `-mu` to `ml`.

Implementation:

```
#define SM_COLUMN_B(A, j)    ( ((SM_CONTENT_B(A)->cols)[j])+SM_SUBBAND_B(A) )
```

### SM\_ELEMENT\_B(A)

This macro gives access to the individual entries of the data array of a banded SUNMatrix.

The assignments `SM_ELEMENT_B(A, i, j) = a_ij` and `a_ij = SM_ELEMENT_B(A, i, j)` reference the  $(i, j)$ -th element of the  $N \times N$  band matrix `A`, where  $0 \leq i, j \leq N - 1$ . The location  $(i, j)$  should further satisfy  $j - \mu \leq i \leq j + \mu$ .

Implementation:

```
#define SM_ELEMENT_B(A, i, j)    ( (SM_CONTENT_B(A)->cols)[j][(i)-(j)+SM_SUBBAND_B(A)] )
```

### SM\_COLUMN\_ELEMENT\_B(A)

This macro gives access to the individual entries of the data array of a banded SUNMatrix.

The assignments `SM_COLUMN_ELEMENT_B(col_j, i, j) = a_ij` and `a_ij = SM_COLUMN_ELEMENT_B(col_j, i, j)` reference the  $(i, j)$ -th entry of the band matrix `A` when used in conjunction with `SM_COLUMN_B` to reference the  $j$ -th column through `col_j`. The index  $(i, j)$  should satisfy  $j - \mu \leq i \leq j + \mu$ .

Implementation:

```
#define SM_COLUMN_ELEMENT_B(col_j, i, j)    ( col_j[(i)-(j)] )
```

The `SUNMATRIX_BAND` module defines banded implementations of all matrix operations listed in §7.2. Their names are obtained from those in that section by appending the suffix `_Band` (e.g. `SUNMatCopy_Band`). The module `SUNMATRIX_BAND` provides the following additional user-callable routines:

*SUNMatrix* **SUNBandMatrix**(*sunindextype* N, *sunindextype* mu, *sunindextype* ml, *SUNContext* suncxt)

This constructor function creates and allocates memory for a banded SUNMatrix. Its arguments are the matrix size, `N`, and the upper and lower half-bandwidths of the matrix, `mu` and `ml`. The stored upper bandwidth is set to `mu+ml` to accommodate subsequent factorization in the `SUNLINSOL_BAND` and `SUNLINSOL_LAPACK-BAND` modules.

*SUNMatrix* **SUNBandMatrixStorage**(*sunindextype* N, *sunindextype* mu, *sunindextype* ml, *sunindextype* smu, *SUNContext* sunctx)

This constructor function creates and allocates memory for a banded *SUNMatrix*. Its arguments are the matrix size, N, the upper and lower half-bandwidths of the matrix, mu and ml, and the stored upper bandwidth, smu. When creating a band *SUNMatrix*, this value should be

- at least  $\min(N-1, \mu+ml)$  if the matrix will be used by the *SUNLinSol\_Band* module;
- exactly equal to  $\mu+ml$  if the matrix will be used by the *SUNLinSol\_LapackBand* module;
- at least  $\mu$  if used in some other manner.

---

**Note:** It is strongly recommended that users call the default constructor, *SUNBandMatrix()*, in all standard use cases. This advanced constructor is used internally within SUNDIALS solvers, and is provided to users who require banded matrices for non-default purposes.

---

void **SUNBandMatrix\_Print**(*SUNMatrix* A, FILE \*outfile)

This function prints the content of a banded *SUNMatrix* to the output stream specified by *outfile*. Note: *stdout* or *stderr* may be used as arguments for *outfile* to print directly to standard output or standard error, respectively.

*sunindextype* **SUNBandMatrix\_Rows**(*SUNMatrix* A)

This function returns the number of rows in the banded *SUNMatrix*.

*sunindextype* **SUNBandMatrix\_Columns**(*SUNMatrix* A)

This function returns the number of columns in the banded *SUNMatrix*.

*sunindextype* **SUNBandMatrix\_LowerBandwidth**(*SUNMatrix* A)

This function returns the lower half-bandwidth for the banded *SUNMatrix*.

*sunindextype* **SUNBandMatrix\_UpperBandwidth**(*SUNMatrix* A)

This function returns the upper half-bandwidth of the banded *SUNMatrix*.

*sunindextype* **SUNBandMatrix\_StoredUpperBandwidth**(*SUNMatrix* A)

This function returns the stored upper half-bandwidth of the banded *SUNMatrix*.

*sunindextype* **SUNBandMatrix\_LDim**(*SUNMatrix* A)

This function returns the length of the leading dimension of the banded *SUNMatrix*.

*sunindextype* **SUNBandMatrix\_LData**(*SUNMatrix* A)

This function returns the length of the data array for the banded *SUNMatrix*.

*sunrealtype* \***SUNBandMatrix\_Data**(*SUNMatrix* A)

This function returns a pointer to the data array for the banded *SUNMatrix*.

*sunrealtype* \*\***SUNBandMatrix\_Cols**(*SUNMatrix* A)

This function returns a pointer to the cols array for the band *SUNMatrix*.

*sunrealtype* \***SUNBandMatrix\_Column**(*SUNMatrix* A, *sunindextype* j)

This function returns a pointer to the diagonal entry of the j-th column of the banded *SUNMatrix*. The resulting pointer should be indexed over the range  $-\mu$  to  $ml$ .

**Warning:** When calling this function from the Fortran interfaces the shape of the array that is returned is [1], and the only element you can (legally) access is the diagonal element. Fortran users should instead work with the data array returned by *SUNBandMatrix\_Data()* directly.

## Notes

- When looping over the components of a banded SUNMatrix  $A$ , the most efficient approaches are to:
  - First obtain the component array via  $A\_data = \text{SUNBandMatrix\_Data}(A)$ , or equivalently  $A\_data = \text{SM\_DATA\_B}(A)$ , and then access  $A\_data[i]$  within the loop.
  - First obtain the array of column pointers via  $A\_cols = \text{SUNBandMatrix\_Cols}(A)$ , or equivalently  $A\_cols = \text{SM\_COLS\_B}(A)$ , and then access  $A\_cols[j][i]$  within the loop.
  - Within a loop over the columns, access the column pointer via  $A\_colj = \text{SUNBandMatrix\_Column}(A, j)$  and then to access the entries within that column using  $\text{SM\_COLUMN\_ELEMENT\_B}(A\_colj, i, j)$ .

All three of these are more efficient than using  $\text{SM\_ELEMENT\_B}(A, i, j)$  within a double loop.

- Within the  $\text{SUNMatMatvec\_Band}$  routine, internal consistency checks are performed to ensure that the matrix is called with consistent  $N\_Vector$  implementations. These are currently limited to:  $NVECTOR\_SERIAL$ ,  $NVECTOR\_OPENMP$ , and  $NVECTOR\_PTHREADS$ . As additional compatible vector implementations are added to SUNDIALS, these will be included within this compatibility check.

## 7.7 The SUNMATRIX\_CUSPARSE Module

The  $\text{SUNMATRIX\_CUSPARSE}$  module is an interface to the NVIDIA cuSPARSE matrix for use on NVIDIA GPUs [7]. All data stored by this matrix implementation resides on the GPU at all times.

The header file to be included when using this module is  $\text{sunmatrix/sunmatrix\_cusparses.h}$ . The installed library to link to is  $\text{libsundials\_sunmatrixcusparses.lib}$  where  $\text{.lib}$  is typically  $\text{.so}$  for shared libraries and  $\text{.a}$  for static libraries.

### 7.7.1 SUNMATRIX\_CUSPARSE Description

The implementation currently supports the cuSPARSE CSR matrix format described in the cuSPARSE documentation, as well as a unique low-storage format for block-diagonal matrices of the form

$$A = \begin{bmatrix} A_0 & 0 & \cdots & 0 \\ 0 & A_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_{n-1} \end{bmatrix},$$

where all the block matrices  $A_j$  share the same sparsity pattern. We will refer to this format as BCSR (not to be confused with the canonical BSR format where each block is stored as dense). In this format, the CSR column indices and row pointers are only stored for the first block and are computed only as necessary for other blocks. This can drastically reduce the amount of storage required compared to the regular CSR format when the number of blocks is large. This format is well-suited for, and intended to be used with, the  $\text{SUNLinearSolver\_cuSolverSp\_batchQR}$  linear solver (see §8.17).

**The  $\text{SUNMATRIX\_CUSPARSE}$  module is experimental and subject to change.**

## 7.7.2 SUNMATRIX\_CUSPARSE Functions

The SUNMATRIX\_CUSPARSE module defines GPU-enabled sparse implementations of all matrix operations listed in §7.2 except for the *SUNMatSpace()* and *SUNMatMatvecSetup()* operations:

- *SUNMatGetID\_cuSparse* – returns SUNMATRIX\_CUSPARSE
- *SUNMatClone\_cuSparse*
- *SUNMatDestroy\_cuSparse*
- *SUNMatZero\_cuSparse*
- *SUNMatCopy\_cuSparse*
- *SUNMatScaleAdd\_cuSparse* – performs  $A = cA + B$ , where  $A$  and  $B$  must have the same sparsity pattern
- *SUNMatScaleAddI\_cuSparse* – performs  $A = cA + I$ , where the diagonal of  $A$  must be present
- *SUNMatMatvec\_cuSparse*

In addition, the SUNMATRIX\_CUSPARSE module defines the following implementation specific functions:

*SUNMatrix* **SUNMatrix\_cuSparse\_NewCSR**(int M, int N, int NNZ, cusparseHandle\_t cusp, *SUNContext* sunctx)

This constructor function creates and allocates memory for a SUNMATRIX\_CUSPARSE *SUNMatrix* object that uses the CSR storage format. Its arguments are the number of rows and columns of the matrix, M and N, the number of nonzeros to be stored in the matrix, NNZ, and a valid *cusparseHandle\_t*.

*SUNMatrix* **SUNMatrix\_cuSparse\_NewBlockCSR**(int nblocks, int blockrows, int blockcols, int blocknnz, cusparseHandle\_t cusp, *SUNContext* sunctx)

This constructor function creates and allocates memory for a SUNMATRIX\_CUSPARSE *SUNMatrix* object that leverages the SUNMAT\_CUSPARSE\_BCSR storage format to store a block diagonal matrix where each block shares the same sparsity pattern. The blocks must be square. The function arguments are the number of blocks, nblocks, the number of rows, blockrows, the number of columns, blockcols, the number of nonzeros in each block, blocknnz, and a valid *cusparseHandle\_t*.

**Warning:** The SUNMAT\_CUSPARSE\_BCSR format currently only supports square matrices, i.e., blockrows == blockcols.

*SUNMatrix* **SUNMatrix\_cuSparse\_MakeCSR**(cusparseMatDescr\_t mat\_descr, int M, int N, int NNZ, int \*rowptrs, int \*colind, *sunrealtype* \*data, cusparseHandle\_t cusp, *SUNContext* sunctx)

This constructor function creates a SUNMATRIX\_CUSPARSE *SUNMatrix* object from user provided pointers. Its arguments are a *cusparseMatDescr\_t* that must have index base CUSPARSE\_INDEX\_BASE\_ZERO, the number of rows and columns of the matrix, M and N, the number of nonzeros to be stored in the matrix, NNZ, and a valid *cusparseHandle\_t*.

int **SUNMatrix\_cuSparse\_Rows**(*SUNMatrix* A)

This function returns the number of rows in the sparse *SUNMatrix*.

int **SUNMatrix\_cuSparse\_Columns**(*SUNMatrix* A)

This function returns the number of columns in the sparse *SUNMatrix*.

int **SUNMatrix\_cuSparse\_NNZ**(*SUNMatrix* A)

This function returns the number of entries allocated for nonzero storage for the sparse *SUNMatrix*.



int **SUNMatrix\_cuSparse\_SparseType**(*SUNMatrix* A)

This function returns the storage type (SUNMAT\_CUSPARSE\_CSR or SUNMAT\_CUSPARSE\_BCSR) for the sparse SUNMatrix.

*sunrealtype* \***SUNMatrix\_cuSparse\_Data**(*SUNMatrix* A)

This function returns a pointer to the data array for the sparse SUNMatrix.

int \***SUNMatrix\_cuSparse\_IndexValues**(*SUNMatrix* A)

This function returns a pointer to the index value array for the sparse SUNMatrix – for the CSR format this is an array of column indices for each nonzero entry. For the BCSR format this is an array of the column indices for each nonzero entry in the first block only.

int \***SUNMatrix\_cuSparse\_IndexPointers**(*SUNMatrix* A)

This function returns a pointer to the index pointer array for the sparse SUNMatrix – for the CSR format this is an array of the locations of the first entry of each row in the data and indexvalues arrays, for the BCSR format this is an array of the locations of each row in the data and indexvalues arrays in the first block only.

int **SUNMatrix\_cuSparse\_NumBlocks**(*SUNMatrix* A)

This function returns the number of matrix blocks.

int **SUNMatrix\_cuSparse\_BlockRows**(*SUNMatrix* A)

This function returns the number of rows in a matrix block.

int **SUNMatrix\_cuSparse\_BlockColumns**(*SUNMatrix* A)

This function returns the number of columns in a matrix block.

int **SUNMatrix\_cuSparse\_BlockNNZ**(*SUNMatrix* A)

This function returns the number of nonzeros in each matrix block.

*sunrealtype* \***SUNMatrix\_cuSparse\_BlockData**(*SUNMatrix* A, int blockidx)

This function returns a pointer to the location in the data array where the data for the block, blockidx, begins. Thus, blockidx must be less than SUNMatrix\_cuSparse\_NumBlocks(A). The first block in the SUNMatrix is index 0, the second block is index 1, and so on.

cusparseMatDescr\_t **SUNMatrix\_cuSparse\_MatDescr**(*SUNMatrix* A)

This function returns the cusparseMatDescr\_t object associated with the matrix.

*SUNErrCode* **SUNMatrix\_cuSparse\_CopyToDevice**(*SUNMatrix* A, *sunrealtype* \*h\_data, int \*h\_idxptrs, int \*h\_idxvals)

This functions copies the matrix information to the GPU device from the provided host arrays. A user may provide NULL for any of h\_data, h\_idxptrs, or h\_idxvals to avoid copying that information.

The function returns SUN\_SUCCESS if the copy operation(s) were successful, or a nonzero error code otherwise.

*SUNErrCode* **SUNMatrix\_cuSparse\_CopyFromDevice**(*SUNMatrix* A, *sunrealtype* \*h\_data, int \*h\_idxptrs, int \*h\_idxvals)

This functions copies the matrix information from the GPU device to the provided host arrays. A user may provide NULL for any of h\_data, h\_idxptrs, or h\_idxvals to avoid copying that information. Otherwise:

- The h\_data array must be at least SUNMatrix\_cuSparse\_NNZ(A)\*sizeof(sunrealtype) bytes.
- The h\_idxptrs array must be at least (SUNMatrix\_cuSparse\_BlockDim(A)+1)\*sizeof(int) bytes.
- The h\_idxvals array must be at least (SUNMatrix\_cuSparse\_BlockNNZ(A))\*sizeof(int) bytes.

The function returns SUN\_SUCCESS if the copy operation(s) were successful, or a nonzero error code otherwise.

*SUNErrCode* **SUNMatrix\_cuSparse\_SetFixedPattern**(*SUNMatrix* A, *sunbooleantype* yesno)

This function changes the behavior of the the `SUNMatZero` operation on the object A. By default the matrix sparsity pattern is not considered to be fixed, thus, the `SUNMatZero` operation zeros out all data array as well as the `indexvalues` and `indexpointers` arrays. Providing a value of 1 or `SUNTRUE` for the `yesno` argument changes the behavior of `SUNMatZero` on A so that only the data is zeroed out, but not the `indexvalues` or `indexpointers` arrays. Providing a value of 0 or `SUNFALSE` for the `yesno` argument is equivalent to the default behavior.

*SUNErrCode* **SUNMatrix\_cuSparse\_SetKernelExecPolicy**(*SUNMatrix* A, *SUNCudaExecPolicy* \*exec\_policy)

This function sets the execution policies which control the kernel parameters utilized when launching the CUDA kernels. By default the matrix is setup to use a policy which tries to leverage the structure of the matrix. See §6.10.2 for more information about the *SUNCudaExecPolicy* class.

### 7.7.3 SUNMATRIX\_CUSPARSE Usage Notes

The `SUNMATRIX_CUSPARSE` module only supports 32-bit indexing, thus `SUNDIALS` must be built for 32-bit indexing to use this module.

The `SUNMATRIX_CUSPARSE` module can be used with CUDA streams by calling the `cuSPARSE` function `cusparseSetStream` on the `cusparseHandle_t` that is provided to the `SUNMATRIX_CUSPARSE` constructor.

**Warning:** When using the `SUNMATRIX_CUSPARSE` module with a `SUNDIALS` package (e.g. `ARKODE`), the stream given to `cuSPARSE` should be the same stream used for the `NVECTOR` object that is provided to the package, and the `NVECTOR` object given to the `SUNMatvec` operation. If different streams are utilized, synchronization issues may occur.

## 7.8 The SUNMATRIX\_SPARSE Module

The sparse implementation of the `SUNMatrix` module, `SUNMATRIX_SPARSE`, is designed to work with either *compressed-sparse-column* (CSC) or *compressed-sparse-row* (CSR) sparse matrix formats. To this end, it defines the *content* field of `SUNMatrix` to be the following structure:

```
struct _SUNMatrixContent_Sparse {
    sunindextype M;
    sunindextype N;
    sunindextype NNZ;
    sunindextype NP;
    sunrealtype *data;
    int sparsetype;
    sunindextype *indexvals;
    sunindextype *indexptrs;
    /* CSC indices */
    sunindextype **rowvals;
    sunindextype **colptrs;
    /* CSR indices */
    sunindextype **colvals;
    sunindextype **rowptrs;
};
```

A diagram of the underlying data representation in a sparse matrix is shown in Fig. 7.2. A more complete description of the parts of this *content* field is given below:

- **M** - number of rows
- **N** - number of columns
- **NNZ** - maximum number of nonzero entries in the matrix (allocated length of **data** and **indexvals** arrays)
- **NP** - number of index pointers (e.g. number of column pointers for CSC matrix). For CSC matrices **NP=N**, and for CSR matrices **NP=M**. This value is set automatically at construction based the input choice for **sparsetype**.
- **data** - pointer to a contiguous block of **sunrealtype** variables (of length **NNZ**), containing the values of the nonzero entries in the matrix
- **sparsetype** - type of the sparse matrix (**CSC\_MAT** or **CSR\_MAT**)
- **indexvals** - pointer to a contiguous block of **int** variables (of length **NNZ**), containing the row indices (if **CSC**) or column indices (if **CSR**) of each nonzero matrix entry held in **data**
- **indexptrs** - pointer to a contiguous block of **int** variables (of length **NP+1**). For **CSC** matrices each entry provides the index of the first column entry into the **data** and **indexvals** arrays, e.g. if **indexptr[3]=7**, then the first nonzero entry in the fourth column of the matrix is located in **data[7]**, and is located in row **indexvals[7]** of the matrix. The last entry contains the total number of nonzero values in the matrix and hence points one past the end of the active data in the **data** and **indexvals** arrays. For **CSR** matrices, each entry provides the index of the first row entry into the **data** and **indexvals** arrays.

The following pointers are added to the **SUNMATRIX\_SPARSE** content structure for user convenience, to provide a more intuitive interface to the **CSC** and **CSR** sparse matrix data structures. They are set automatically when creating a sparse **SUNMatrix**, based on the sparse matrix storage type.

- **rowvals** - pointer to **indexvals** when **sparsetype** is **CSC\_MAT**, otherwise set to **NULL**.
- **colptrs** - pointer to **indexptrs** when **sparsetype** is **CSC\_MAT**, otherwise set to **NULL**.
- **colvals** - pointer to **indexvals** when **sparsetype** is **CSR\_MAT**, otherwise set to **NULL**.
- **rowptrs** - pointer to **indexptrs** when **sparsetype** is **CSR\_MAT**, otherwise set to **NULL**.

For example, the  $5 \times 4$  matrix

$$\begin{bmatrix} 0 & 3 & 1 & 0 \\ 3 & 0 & 0 & 2 \\ 0 & 7 & 0 & 0 \\ 1 & 0 & 0 & 9 \\ 0 & 0 & 0 & 5 \end{bmatrix}$$

could be stored as a **CSC** matrix in this structure as either

```
M = 5;
N = 4;
NNZ = 8;
NP = N;
data = {3.0, 1.0, 3.0, 7.0, 1.0, 2.0, 9.0, 5.0};
sparsetype = CSC_MAT;
indexvals = {1, 3, 0, 2, 0, 1, 3, 4};
indexptrs = {0, 2, 4, 5, 8};
```

or

```
M = 5;
N = 4;
NNZ = 10;
NP = N;
```

(continues on next page)

(continued from previous page)

```
data = {3.0, 1.0, 3.0, 7.0, 1.0, 2.0, 9.0, 5.0, *, *};
sparsetype = CSC_MAT;
indexvals = {1, 3, 0, 2, 0, 1, 3, 4, *, *};
indexptrs = {0, 2, 4, 5, 8};
```

where the first has no unused space, and the second has additional storage (the entries marked with \* may contain any values). Note in both cases that the final value in `indexptrs` is 8, indicating the total number of nonzero entries in the matrix.

Similarly, in CSR format, the same matrix could be stored as

```
M = 5;
N = 4;
NNZ = 8;
NP = M;
data = {3.0, 1.0, 3.0, 2.0, 7.0, 1.0, 9.0, 5.0};
sparsetype = CSR_MAT;
indexvals = {1, 2, 0, 3, 1, 0, 3, 3};
indexptrs = {0, 2, 4, 5, 7, 8};
```

The header file to be included when using this module is `sunmatrix/sunmatrix_sparse.h`.

The following macros are provided to access the content of a `SUNMATRIX_SPARSE` matrix. The prefix `SM_` in the names denotes that these macros are for *SUNMatrix* implementations, and the suffix `_S` denotes that these are specific to the *sparse* version.

#### **SM\_CONTENT\_S(A)**

This macro gives access to the contents of the sparse *SUNMatrix* *A*.

The assignment `A_cont = SM_CONTENT_S(A)` sets `A_cont` to be a pointer to the sparse *SUNMatrix* content structure.

Implementation:

```
#define SM_CONTENT_S(A)  ( (SUNMatrixContent_Sparse)(A->content) )
```

#### **SM\_ROWS\_S(A)**

Access the number of rows in the sparse *SUNMatrix* *A*.

This may be used either to retrieve or to set the value. For example, the assignment `A_rows = SM_ROWS_S(A)` sets `A_rows` to be the number of rows in the matrix *A*. Similarly, the assignment `SM_ROWS_S(A) = A_rows` sets the number of columns in *A* to equal `A_rows`.

Implementation:

```
#define SM_ROWS_S(A)    ( SM_CONTENT_S(A)->M )
```

#### **SM\_COLUMNS\_S(A)**

Access the number of columns in the sparse *SUNMatrix* *A*. As with `SM_ROWS_S`, this may be used either to retrieve or to set the value.

Implementation:

```
#define SM_COLUMNS_S(A) ( SM_CONTENT_S(A)->N )
```




Fig. 7.2: Diagram of the storage for a compressed-sparse-column matrix of type `SUNMATRIX_SPARSE`: Here  $A$  is an  $M \times N$  sparse CSC matrix with storage for up to `NNZ` nonzero entries (the allocated length of both `data` and `indexvals`). The entries in `indexvals` may assume values from 0 to  $M-1$ , corresponding to the row index (zero-based) of each nonzero value. The entries in `data` contain the values of the nonzero entries, with the row  $i$ , column  $j$  entry of  $A$  (again, zero-based) denoted as  $A(i, j)$ . The `indexptrs` array contains  $N+1$  entries; the first  $N$  denote the starting index of each column within the `indexvals` and `data` arrays, while the final entry points one past the final nonzero entry. Here, although `NNZ` values are allocated, only `nz` are actually filled in; the greyed-out portions of `data` and `indexvals` indicate extra allocated space.

**SM\_NNZ\_S(A)**

Access the allocated number of nonzeros in the sparse SUNMatrix *A*. As with `SM_ROWS_S`, this may be used either to retrieve or to set the value.

Implementation:

```
#define SM_NNZ_S(A)    ( SM_CONTENT_S(A)->NNZ )
```

**SM\_NP\_S(A)**

Access the number of index pointers *NP* in the sparse SUNMatrix *A*. As with `SM_ROWS_S`, this may be used either to retrieve or to set the value.

Implementation:

```
#define SM_NP_S(A)    ( SM_CONTENT_S(A)->NP )
```

**SM\_SPARSETYPE\_S(A)**

Access the sparsity type parameter in the sparse SUNMatrix *A*. As with `SM_ROWS_S`, this may be used either to retrieve or to set the value.

Implementation:

```
#define SM_SPARSETYPE_S(A)    ( SM_CONTENT_S(A)->sparsetype )
```

**SM\_DATA\_S(A)**

This macro gives access to the data pointer for the matrix entries.

The assignment `A_data = SM_DATA_S(A)` sets `A_data` to be a pointer to the first component of the data array for the sparse SUNMatrix *A*. The assignment `SM_DATA_S(A) = A_data` sets the data array of *A* to be `A_data` by storing the pointer `A_data`.

Implementation:

```
#define SM_DATA_S(A)    ( SM_CONTENT_S(A)->data )
```

**SM\_INDEXVALS\_S(A)**

This macro gives access to the `indexvals` pointer for the matrix entries.

The assignment `A_indexvals = SM_INDEXVALS_S(A)` sets `A_indexvals` to be a pointer to the array of index values (i.e. row indices for a CSC matrix, or column indices for a CSR matrix) for the sparse SUNMatrix *A*.

Implementation:

```
#define SM_INDEXVALS_S(A)    ( SM_CONTENT_S(A)->indexvals )
```

**SM\_INDEXPTRS\_S(A)**

This macro gives access to the `indexptrs` pointer for the matrix entries.

The assignment `A_indexptrs = SM_INDEXPTRS_S(A)` sets `A_indexptrs` to be a pointer to the array of index pointers (i.e. the starting indices in the data/indexvals arrays for each row or column in CSR or CSC formats, respectively).

Implementation:

```
#define SM_INDEXPTRS_S(A)    ( SM_CONTENT_S(A)->indexptrs )
```

The `SUNMATRIX_SPARSE` module defines sparse implementations of all matrix operations listed in §7.2. Their names are obtained from those in that section by appending the suffix `_Sparse` (e.g. `SUNMatCopy_Sparse`). The module `SUNMATRIX_SPARSE` provides the following additional user-callable routines:

*SUNMatrix* **SUNSparseMatrix**(*sunindextype* M, *sunindextype* N, *sunindextype* NNZ, int sparsetype, *SUNContext* sunctx)

This constructor function creates and allocates memory for a sparse *SUNMatrix*. Its arguments are the number of rows and columns of the matrix, *M* and *N*, the maximum number of nonzeros to be stored in the matrix, *NNZ*, and a flag *sparsetype* indicating whether to use CSR or CSC format (valid choices are *CSR\_MAT* or *CSC\_MAT*).

*SUNMatrix* **SUNSparseFromDenseMatrix**(*SUNMatrix* A, *sunrealtype* droptol, int sparsetype)

This constructor function creates a new sparse matrix from an existing *SUNMATRIX\_DENSE* object by copying all values with magnitude larger than *droptol* into the sparse matrix structure.

Requirements:

- A must have type *SUNMATRIX\_DENSE*
- *droptol* must be non-negative
- *sparsetype* must be either *CSC\_MAT* or *CSR\_MAT*

The function returns NULL if any requirements are violated, or if the matrix storage request cannot be satisfied.

*SUNMatrix* **SUNSparseFromBandMatrix**(*SUNMatrix* A, *sunrealtype* droptol, int sparsetype)

This constructor function creates a new sparse matrix from an existing *SUNMATRIX\_BAND* object by copying all values with magnitude larger than *droptol* into the sparse matrix structure.

Requirements:

- A must have type *SUNMATRIX\_BAND*
- *droptol* must be non-negative
- *sparsetype* must be either *CSC\_MAT* or *CSR\_MAT*.

The function returns NULL if any requirements are violated, or if the matrix storage request cannot be satisfied.

*SUNErrCode* **SUNSparseMatrix\_Realloc**(*SUNMatrix* A)

This function reallocates internal storage arrays in a sparse matrix so that the resulting sparse matrix has no wasted space (i.e. the space allocated for nonzero entries equals the actual number of nonzeros, *indexptrs[NP]*). Returns a *SUNErrCode*.

*SUNErrCode* **SUNSparseMatrix\_Reallocate**(*SUNMatrix* A, *sunindextype* NNZ)

Function to reallocate internal sparse matrix storage arrays so that the resulting sparse matrix has storage for a specified number of nonzeros. Returns a *SUNErrCode*.

void **SUNSparseMatrix\_Print**(*SUNMatrix* A, FILE \*outfile)

This function prints the content of a sparse *SUNMatrix* to the output stream specified by *outfile*. Note: *stdout* or *stderr* may be used as arguments for *outfile* to print directly to standard output or standard error, respectively.

*sunindextype* **SUNSparseMatrix\_Rows**(*SUNMatrix* A)

This function returns the number of rows in the sparse *SUNMatrix*.

*sunindextype* **SUNSparseMatrix\_Columns**(*SUNMatrix* A)

This function returns the number of columns in the sparse *SUNMatrix*.

*sunindextype* **SUNSparseMatrix\_NNZ**(*SUNMatrix* A)

This function returns the number of entries allocated for nonzero storage for the sparse *SUNMatrix*.

*sunindextype* **SUNSparseMatrix\_NP**(*SUNMatrix* A)

This function returns the number of index pointers for the sparse *SUNMatrix* (the *indexptrs* array has NP+1 entries).

int **SUNSparseMatrix\_SparseType**(*SUNMatrix* A)

This function returns the storage type (CSR\_MAT or CSC\_MAT) for the sparse *SUNMatrix*.

*sunrealtype* \***SUNSparseMatrix\_Data**(*SUNMatrix* A)

This function returns a pointer to the data array for the sparse *SUNMatrix*.

*sunindextype* \***SUNSparseMatrix\_IndexValues**(*SUNMatrix* A)

This function returns a pointer to index value array for the sparse *SUNMatrix* – for CSR format this is the column index for each nonzero entry, for CSC format this is the row index for each nonzero entry.

*sunindextype* \***SUNSparseMatrix\_IndexPointers**(*SUNMatrix* A)

This function returns a pointer to the index pointer array for the sparse *SUNMatrix* – for CSR format this is the location of the first entry of each row in the data and *indexvalues* arrays, for CSC format this is the location of the first entry of each column.

---

**Note:** Within the *SUNMatMatvec\_Sparse* routine, internal consistency checks are performed to ensure that the matrix is called with consistent *N\_Vector* implementations. These are currently limited to: *NVECTOR\_SERIAL*, *NVECTOR\_OPENMP*, *NVECTOR\_PTHREADS*, and *NVECTOR\_CUDA* when using managed memory. As additional compatible vector implementations are added to SUNDIALS, these will be included within this compatibility check.

---

## 7.9 The SUNMATRIX\_SLUNRLOC Module

The *SUNMATRIX\_SLUNRLOC* module is an interface to the *SuperMatrix* structure provided by the SuperLU-DIST sparse matrix factorization and solver library written by X. Sherry Li and collaborators [8, 37, 54, 55]. It is designed to be used with the SuperLU\_DIST *SUNLinearSolver* module discussed in §8.15. To this end, it defines the content field of *SUNMatrix* to be the following structure:

```
struct _SUNMatrixContent_SLUNRloc {
    sunboolean_t own_data;
    gridinfo_t *grid;
    sunindextype *row_to_proc;
    pdgsmv_comm_t *gsmv_comm;
    SuperMatrix *A_super;
    SuperMatrix *ACS_super;
};
```

A more complete description of the this content field is given below:

- *own\_data* – a flag which indicates if the *SUNMatrix* is responsible for freeing *A\_super*
- *grid* – pointer to the SuperLU\_DIST structure that stores the 2D process grid
- *row\_to\_proc* – a mapping between the rows in the matrix and the process it resides on; will be NULL until the *SUNMatMatvecSetup* routine is called
- *gsmv\_comm* – pointer to the SuperLU\_DIST structure that stores the communication information needed for matrix-vector multiplication; will be NULL until the *SUNMatMatvecSetup* routine is called
- *A\_super* – pointer to the underlying SuperLU\_DIST *SuperMatrix* with *Stype* = *SLU\_NR\_loc*, *Dtype* = *SLU\_D*, *Mtype* = *SLU\_GE*; must have the full diagonal present to be used with *SUNMatScaleAddI* routine
- *ACS\_super* – a column-sorted version of the matrix needed to perform matrix-vector multiplication; will be NULL until the routine *SUNMatMatvecSetup* routine is called



The header file to include when using this module is `sunmatrix/sunmatrix_slunrloc.h`. The installed module library to link to is `libsundials_sunmatrixslunrloc.lib` where `.lib` is typically `.so` for shared libraries and `.a` for static libraries.

## 7.9.1 SUNMATRIX\_SLUNRLOC Functions

The SUNMATRIX\_SLUNRLOC module provides the following user-callable routines:

*SUNMatrix* **SUNMatrix\_SLUNRloc**(SuperMatrix \*Asuper, gridinfo\_t \*grid, *SUNContext* sunctx)

This constructor function creates and allocates memory for a SUNMATRIX\_SLUNRLOC object. Its arguments are a fully-allocated SuperLU\_DIST SuperMatrix with `Stype = SLU_NR_loc`, `Dtype = SLU_D`, `Mtype = SLU_GE` and an initialized SuperLU\_DIST 2D process grid structure. It returns a SUNMatrix object if `Asuper` is compatible else it returns NULL.

void **SUNMatrix\_SLUNRloc\_Print**(*SUNMatrix* A, FILE \*fp)

This function prints the underlying SuperMatrix content. It is useful for debugging. Its arguments are the SUNMatrix object and a FILE pointer to print to. It returns void.

SuperMatrix \***SUNMatrix\_SLUNRloc\_SuperMatrix**(*SUNMatrix* A)

This function returns the underlying SuperMatrix of `A`. Its only argument is the SUNMatrix object to access.

gridinfo\_t \***SUNMatrix\_SLUNRloc\_ProcessGrid**(*SUNMatrix* A)

This function returns the SuperLU\_DIST 2D process grid associated with `A`. Its only argument is the SUNMatrix object to access.

*sunbooleantype* **SUNMatrix\_SLUNRloc\_OwnData**(*SUNMatrix* A)

This function returns true if the SUNMatrix object is responsible for freeing the underlying SuperMatrix, otherwise it returns false. Its only argument is the SUNMatrix object to access.

The SUNMATRIX\_SLUNRLOC module also defines implementations of all generic SUNMatrix operations listed in §7.2:

- `SUNMatGetID_SLUNRloc` – returns SUNMATRIX\_SLUNRLOC
- `SUNMatClone_SLUNRloc`
- `SUNMatDestroy_SLUNRloc`
- `SUNMatSpace_SLUNRloc` – this only returns information for the storage within the matrix interface, i.e. storage for `row_to_proc`
- `SUNMatZero_SLUNRloc`
- `SUNMatCopy_SLUNRloc`
- `SUNMatScaleAdd_SLUNRloc` – performs  $A = cA + B$ , where  $A$  and  $B$  must have the same sparsity pattern
- `SUNMatScaleAddI_SLUNRloc` – performs  $A = cA + I$ , where the diagonal of  $A$  must be present
- `SUNMatMatvecSetup_SLUNRloc` – initializes the SuperLU\_DIST parallel communication structures needed to perform a matrix-vector product; only needs to be called before the first call to `SUNMatMatvec()` or if the matrix changed since the last setup
- `SUNMatMatvec_SLUNRloc`

## 7.10 The SUNMATRIX\_GINKGO Module

New in version 6.4.0.

The SUNMATRIX\_GINKGO implementation of the SUNMatrix API provides an interface to the matrix data structure for the Ginkgo linear algebra library [11]. Ginkgo provides several different matrix formats and linear solvers which can run on a variety of hardware, such as NVIDIA, AMD, and Intel GPUs as well as multicore CPUs. Since Ginkgo is a modern C++ library, SUNMATRIX\_GINKGO is also written in modern C++ (it requires C++14). Unlike most other SUNDIALS modules, it is a header only library. To use the SUNMATRIX\_GINKGO SUNMatrix, users will need to include `sunmatrix/sunmatrix_ginkgo.hpp`. More instructions on building SUNDIALS with Ginkgo enabled are given in §12.4. For instructions on building and using Ginkgo itself, refer to the [Ginkgo website and documentation](#).

---

**Note:** It is assumed that users of this module are aware of how to use Ginkgo. This module does not try to encapsulate Ginkgo matrices, rather it provides a lightweight interoperability layer between Ginkgo and SUNDIALS.

---

The SUNMATRIX\_GINKGO module is defined by the `sundials::ginkgo::Matrix` templated class:

```
template<typename GkoMatType>
class Matrix : public sundials::impl::BaseMatrix, public sundials::ConvertibleTo<SUNMatrix>;
```

### 7.10.1 Compatible Vectors

The `N_Vector` to use with the SUNLINEARSOLVER\_GINKGO module depends on the `gko::Executor` utilized. That is, when using the `gko::CudaExecutor` you should use a CUDA capable `N_Vector` (e.g., §6.10), `gko::HipExecutor` goes with a HIP capable `N_Vector` (e.g., §6.11), `gko::DpcppExecutor` goes with a DPC++/SYCL capable `N_Vector` (e.g., §6.12), and `gko::OmpExecutor` goes with a CPU based `N_Vector` (e.g., §6.6). Specifically, what makes a `N_Vector` compatible with different Ginkgo executors is where they store the data. The GPU enabled Ginkgo executors need the data to reside on the GPU, so the `N_Vector` must implement `N_VGetDeviceArrayPointer()` and keep the data in GPU memory. The CPU-only enabled Ginkgo executors (e.g, `gko::OmpExecutor` and `gko::ReferenceExecutor`) need data to reside on the CPU and will use `N_VGetArrayPointer()` to access the `N_Vector` data.

### 7.10.2 Using SUNMATRIX\_GINKGO

To use the SUNMATRIX\_GINKGO module, we begin by creating an instance of a Ginkgo matrix using Ginkgo's API. For example, below we create a Ginkgo sparse matrix that uses the CSR storage format and then fill the diagonal of the matrix with ones to make an identity matrix:

```
auto gko_matrix{gko::matrix::Csr<sunrealtype, sunindextype>::create(gko_exec, matrix_dim)};
gko_matrix->read(gko::matrix_data<sunrealtype, sunindextype>::diag(matrix_dim, 1.0));
```

After we have a Ginkgo matrix object, we wrap it in an instance of the `sundials::ginkgo::Matrix` class. This object can be provided to other SUNDIALS functions that expect a `SUNMatrix` object via implicit conversion, or the `Convert()` method:

```
sundials::ginkgo::Matrix<gko::matrix::Csr> matrix{gko_matrix, sunctx};
SUNMatrix I1 = matrix.Convert(); // explicit conversion to SUNMatrix
SUNMatrix I2 = matrix;          // implicit conversion to SUNMatrix
```

No further interaction with `matrix` is required from this point, and it is possible to use the `SUNMatrix` API operating on `I1` or `I2` (or if needed, via Ginkgo operations on `gko_matrix`).

**Warning:** `SUNMatDestroy()` should never be called on a `SUNMatrix` that was created via conversion from a `sundials::ginkgo::Matrix`. Doing so may result in a double free.

### 7.10.3 SUNMATRIX\_GINKGO API

In this section we list the public API of the `sundials::ginkgo::Matrix` class.

```
template<typename GkoMatType>
```

```
class Matrix : public sundials::impl::BaseMatrix, public sundials::ConvertibleTo<SUNMatrix>
```

```
    Matrix() = default
```

Default constructor - means the matrix must be copied or moved to.

```
    Matrix(std::shared_ptr<GkoMatType> gko_mat, SUNContext sunctx)
```

Constructs a `Matrix` from an existing Ginkgo matrix object.

#### Parameters

- **gko\_mat** – A Ginkgo matrix object
- **sunctx** – The SUNDIALS simulation context object (*SUNContext*)

```
    Matrix(Matrix &&that_matrix) noexcept
```

Move constructor.

```
    Matrix(const Matrix &that_matrix)
```

Copy constructor (performs a deep copy).

```
    Matrix &operator=(Matrix &&rhs) noexcept
```

Move assignment.

```
    Matrix &operator=(const Matrix &rhs)
```

Copy assignment clones the `gko::matrix` and *SUNMatrix*. This is a deep copy (i.e. a new data array is created).

```
    virtual ~Matrix() = default;
```

Default destructor.

```
    std::shared_ptr<GkoMatType> GkoMtx() const
```

Get the underlying Ginkgo matrix object.

```
    std::shared_ptr<const gko::Executor> GkoExec() const
```

Get the `gko::Executor` associated with the Ginkgo matrix.

```
    const gko::dim<2> &GkoSize() const
```

Get the size, i.e. `gko::dim`, for the Ginkgo matrix.

```
    operator SUNMatrix() override
```

Implicit conversion to a *SUNMatrix*.

```
    operator SUNMatrix() const override
```

Implicit conversion to a *SUNMatrix*.

```
    SUNMatrix Convert() override
```

Explicit conversion to a *SUNMatrix*.

```
    SUNMatrix Convert() const override
```

Explicit conversion to a *SUNMatrix*.

## 7.11 The SUNMATRIX\_KOKKOSDENSE Module

New in version 6.4.0.

The SUNMATRIX\_KOKKOSDENSE *SUNMatrix* implementation provides a data structure for dense and dense batched (block-diagonal) matrices using Kokkos [32, 64] and KokkosKernels [63] to support a variety of backends including serial, OpenMP, CUDA, HIP, and SYCL. Since Kokkos is a modern C++ library, the module is also written in modern C++ (it requires C++14) as a header only library. To utilize this *SUNMatrix* users will need to include `sunmatrix/sunmatrix_kokkosdense.hpp`. More instructions on building SUNDIALS with Kokkos and KokkosKernels enabled are given in §12.4. For instructions on building and using Kokkos and KokkosKernels, refer to the [Kokkos](#) and [KokkosKernels](#) documentation.

### 7.11.1 Using SUNMATRIX\_KOKKOSDENSE

The SUNMATRIX\_KOKKOSDENSE module is defined by the `DenseMatrix` templated class in the `sundials::kokkos` namespace:

```
template<class ExecutionSpace = Kokkos::DefaultExecutionSpace,
         class MemorySpace = typename ExecutionSpace::memory_space>
class DenseMatrix : public sundials::impl::BaseMatrix,
                   public sundials::ConvertibleTo<SUNMatrix>
```

To use the SUNMATRIX\_KOKKOSDENSE module, we begin by constructing an instance of the Kokkos dense matrix e.g.,

```
// Single matrix using the default execution space
sundials::kokkos::DenseMatrix<> A{rows, cols, sunctx};

// Batched (block-diagonal) matrix using the default execution space
sundials::kokkos::DenseMatrix<> Abatch{blocks, rows, cols, sunctx};

// Batched (block-diagonal) matrix using the Cuda execution space
sundials::kokkos::DenseMatrix<Kokkos::Cuda> Abatch{blocks, rows, cols, sunctx};

// Batched (block-diagonal) matrix using the Cuda execution space and
// a non-default execution space instance
sundials::kokkos::DenseMatrix<Kokkos::Cuda> Abatch{blocks, rows, cols,
                                                    exec_space_instance,
                                                    sunctx};
```

Instances of the `DenseMatrix` class are implicitly or explicitly (using the `Convert()` method) convertible to a *SUNMatrix* e.g.,

```
sundials::kokkos::DenseMatrix<> A{rows, cols, sunctx};
SUNMatrix B = A;           // implicit conversion to SUNMatrix
SUNMatrix C = A.Convert(); // explicit conversion to SUNMatrix
```

No further interaction with a `DenseMatrix` is required from this point, and it is possible to use the *SUNMatrix* API to operate on B or C.

**Warning:** `SUNMatDestroy()` should never be called on a *SUNMatrix* that was created via conversion from a `sundials::kokkos::DenseMatrix`. Doing so may result in a double free.

The underlying `DenseMatrix` can be extracted from a `SUNMatrix` using `GetDenseMat()` e.g.,

```
auto A_dense_mat = GetDenseMat<>(A_sunmat);
```

The `SUNMATRIX_KOKKOSDENSE` module is compatible with the `NVECTOR_KOKKOS` vector module (see §6.14) and `SUNLINEARSOLVER_KOKKOSDENSE` linear solver module (see §8.19).

### 7.11.2 SUNMATRIX\_KOKKOSDENSE API

In this section we list the public API of the `sundials::kokkos::DenseMatrix` class.

```
template<class ExecutionSpace = Kokkos::DefaultExecutionSpace, class MemorySpace = typename
ExecutionSpace::memory_space>
```

```
class DenseMatrix : public sundials::impl::BaseMatrix, public sundials::ConvertibleTo<SUNMatrix>
```

```
using exec_space = ExecutionSpace;
```

```
using memory_space = MemorySpace;
```

```
using view_type = Kokkos::View<sunrealtype***, memory_space>;
```

```
using size_type = typename view_type::size_type;
```

```
using range_policy = Kokkos::MDRangePolicy<exec_space, Kokkos::Rank<3>>;
```

```
using team_policy = typename Kokkos::TeamPolicy<exec_space>;
```

```
using member_type = typename Kokkos::TeamPolicy<exec_space>::member_type;
```

```
DenseMatrix() = default
```

Default constructor – the matrix must be copied or moved to.

```
DenseMatrix(size_type rows, size_type cols, SUNContext sunctx)
```

Constructs a single `DenseMatrix` using the default execution space instance.

#### Parameters

- **rows** – number of matrix rows
- **cols** – number of matrix columns
- **sunctx** – the SUNDIALS simulation context object (*SUNContext*)

```
DenseMatrix(size_type rows, size_type cols, exec_space ex, SUNContext sunctx)
```

Constructs a single `DenseMatrix` using the provided execution space instance.

#### Parameters

- **rows** – number of matrix rows
- **cols** – number of matrix columns
- **ex** – an execution space
- **sunctx** – the SUNDIALS simulation context object (*SUNContext*)

```
DenseMatrix(size_type blocks, size_type block_rows, size_type block_cols, SUNContext sunctx)
```

Constructs a batched (block-diagonal) `DenseMatrix` using the default execution space instance.

#### Parameters

- **blocks** – number of matrix blocks

- **block\_rows** – number of rows in a block
- **block\_cols** – number of columns in a block
- **sunctx** – the SUNDIALS simulation context object (*SUNContext*)

**DenseMatrix**(*size\_type* blocks, *size\_type* block\_rows, *size\_type* block\_cols, *exec\_space* ex, *SUNContext* sunctx)

Constructs a batched (block-diagonal) DenseMatrix using the provided execution space instance.

**Parameters**

- **blocks** – number of matrix blocks
- **block\_rows** – number of rows in a block
- **block\_cols** – number of columns in a block
- **ex** – an execution space
- **sunctx** – the SUNDIALS simulation context object (*SUNContext*)

**DenseMatrix**(*DenseMatrix* &&that\_matrix) noexcept

Move constructor.

**DenseMatrix**(const *DenseMatrix* &that\_matrix)

Copy constructor. This creates a shallow clone of the Matrix, i.e., it creates a new Matrix with the same properties, such as size, but it does not copy the data.

*DenseMatrix* &**operator=**(*DenseMatrix* &&rhs) noexcept

Move assignment.

*DenseMatrix* &**operator=**(const *DenseMatrix* &rhs)

Copy assignment. This creates a shallow clone of the Matrix, i.e., it creates a new Matrix with the same properties, such as size, but it does not copy the data.

virtual ~**DenseMatrix**() = default;

Default destructor.

*exec\_space* **ExecSpace**()

Get the execution space instance used by the matrix.

*view\_type* **View**()

Get the underlying Kokkos view with extents {blocks, block\_rows, block\_cols}.

*size\_type* **Blocks**()

Get the number of blocks i.e., extent(0).

*size\_type* **BlockRows**()

Get the number of rows in a block i.e., extent(1).

*size\_type* **BlockCols**()

Get the number of columns in a block i.e., extent(2).

*size\_type* **Rows**()

Get the number of rows in the block-diagonal matrix i.e., extent(0) \* extent(1).

*size\_type* **Cols**()

Get the number of columns in the block-diagonal matrix i.e., extent(0) \* extent(2).

**operator SUNMatrix()** override

Implicit conversion to a *SUNMatrix*.

**operator SUNMatrix() const** override

Implicit conversion to a *SUNMatrix*.

**SUNMatrix Convert()** override

Explicit conversion to a *SUNMatrix*.

**SUNMatrix Convert() const** override

Explicit conversion to a *SUNMatrix*.

```
template<class ExecutionSpace = Kokkos::DefaultExecutionSpace, class MemorySpace = typename
ExecutionSpace::memory_space>
```

```
inline DenseMatrix<MatrixType> *GetDenseMat(SUNMatrix A)
```

Get the dense matrix wrapped by a SUNMatrix

## 7.12 SUNMATRIX Examples

There are *SUNMatrix* examples that may be installed for each implementation, that make use of the functions in `test_sunmatrix.c`. These example functions show simple usage of the *SUNMatrix* family of functions. The inputs to the examples depend on the matrix type, and are output to `stdout` if the example is run without the appropriate number of command-line arguments.

The following is a list of the example functions in `test_sunmatrix.c`:

- **Test\_SUNMatGetID**: Verifies the returned matrix ID against the value that should be returned.
- **Test\_SUNMatClone**: Creates clone of an existing matrix, copies the data, and checks that their values match.
- **Test\_SUNMatZero**: Zeros out an existing matrix and checks that each entry equals 0.0.
- **Test\_SUNMatCopy**: Clones an input matrix, copies its data to a clone, and verifies that all values match.
- **Test\_SUNMatScaleAdd**: Given an input matrix  $A$  and an input identity matrix  $I$ , this test clones and copies  $A$  to a new matrix  $B$ , computes  $B = -B + B$ , and verifies that the resulting matrix entries equal 0. Additionally, if the matrix is square, this test clones and copies  $A$  to a new matrix  $D$ , clones and copies  $I$  to a new matrix  $C$ , computes  $D = D + I$  and  $C = C + A$  using *SUNMatScaleAdd()*, and then verifies that  $C = D$ .
- **Test\_SUNMatScaleAddI**: Given an input matrix  $A$  and an input identity matrix  $I$ , this clones and copies  $I$  to a new matrix  $B$ , computes  $B = -B + I$  using *SUNMatScaleAddI()*, and verifies that the resulting matrix entries equal 0.
- **Test\_SUNMatMatvecSetup**: verifies that *SUNMatMatvecSetup()* can be called.
- **Test\_SUNMatMatvec**: Given an input matrix  $A$  and input vectors  $x$  and  $y$  such that  $y = Ax$ , this test has different behavior depending on whether  $A$  is square. If it is square, it clones and copies  $A$  to a new matrix  $B$ , computes  $B = 3B + I$  using *SUNMatScaleAddI()*, clones  $y$  to new vectors  $w$  and  $z$ , computes  $z = Bx$  using *SUNMatMatvec()*, computes  $w = 3y + x$  using `N_VLinearSum`, and verifies that  $w == z$ . If  $A$  is not square, it just clones  $y$  to a new vector  $z$ , computes  $z = Ax$  using *SUNMatMatvec()*, and verifies that  $y = z$ .
- **Test\_SUNMatSpace**: verifies that *SUNMatSpace()* can be called, and outputs the results to `stdout`.

## 7.13 SUNMatrix functions used by IDAS

In Table 7.2, we list the matrix functions in the `SUNMatrix` module used within the IDAS package. The table also shows, for each function, which of the code modules uses the function. The main IDAS integrator does not call any `SUNMatrix` functions directly, so the table columns are specific to the IDALS and IDABBDPRE preconditioner modules. We further note that the IDALS interface only utilizes these routines when supplied with a *matrix-based* linear solver, i.e., the `SUNMatrix` object passed to `IDASetLinearSolver()` was not NULL.

At this point, we should emphasize that the IDAS user does not need to know anything about the usage of matrix functions by the IDAS code modules in order to use IDAS. The information is presented as an implementation detail for the interested reader.

Table 7.2: List of matrix functions usage by IDAS code modules

	IDALS	IDABBDPRE
<code>SUNMatGetID()</code>	x	
<code>SUNMatDestroy()</code>		x
<code>SUNMatZero()</code>	x	x
<code>SUNMatSpace()</code>		†

The matrix functions listed with a † symbol are optionally used, in that these are only called if they are implemented in the `SUNMatrix` module that is being used (i.e. their function pointers are non-NULL). The matrix functions listed in §7.1 that are *not* used by IDAS are: `SUNMatCopy()`, `SUNMatClone()`, `SUNMatScaleAdd()`, `SUNMatScaleAddI()` and `SUNMatMatvec()`. Therefore a user-supplied `SUNMatrix` module for IDAS could omit these functions.

We note that the IDABBDPRE preconditioner module is hard-coded to use the SUNDIALS-supplied band `SUNMatrix` type, so the most useful information above for user-supplied `SUNMatrix` implementations is the column relating the IDALS requirements.



## Chapter 8

# Linear Algebraic Solvers

For problems that require the solution of linear systems of equations, the SUNDIALS packages operate using generic linear solver modules defined through the [SUNLinearSolver](#), or “SUNLinSol”, API. This allows SUNDIALS packages to utilize any valid SUNLinSol implementation that provides a set of required functions. These functions can be divided into three categories. The first are the core linear solver functions. The second group consists of “set” routines to supply the linear solver object with functions provided by the SUNDIALS package, or for modification of solver parameters. The last group consists of “get” routines for retrieving artifacts (statistics, residual vectors, etc.) from the linear solver. All of these functions are defined in the header file `sundials/sundials_linearsolver.h`.

The implementations provided with SUNDIALS work in coordination with the SUNDIALS [N\\_Vector](#), and optionally [SUNMatrix](#), modules to provide a set of compatible data structures and solvers for the solution of linear systems using direct or iterative (matrix-based or matrix-free) methods. Moreover, advanced users can provide a customized `SUNLinearSolver` implementation to any SUNDIALS package, particularly in cases where they provide their own `N_Vector` and/or `SUNMatrix` modules.

Historically, the SUNDIALS packages have been designed to specifically leverage the use of either *direct linear solvers* or matrix-free, *scaled, preconditioned, iterative linear solvers*. However, matrix-based iterative linear solvers are also supported.

The iterative linear solvers packaged with SUNDIALS leverage scaling and preconditioning, as applicable, to balance error between solution components and to accelerate convergence of the linear solver. To this end, instead of solving the linear system  $Ax = b$  directly, these apply the underlying iterative algorithm to the transformed system

$$\tilde{A}\tilde{x} = \tilde{b} \tag{8.1}$$

where

$$\begin{aligned} \tilde{A} &= S_1 P_1^{-1} A P_2^{-1} S_2^{-1}, \\ \tilde{b} &= S_1 P_1^{-1} b, \\ \tilde{x} &= S_2 P_2 x, \end{aligned} \tag{8.2}$$

and where

- $P_1$  is the left preconditioner,
- $P_2$  is the right preconditioner,
- $S_1$  is a diagonal matrix of scale factors for  $P_1^{-1}b$ ,
- $S_2$  is a diagonal matrix of scale factors for  $P_2x$ .

SUNDIALS solvers request that iterative linear solvers stop based on the 2-norm of the scaled preconditioned residual meeting a prescribed tolerance, i.e.,

$$\|\tilde{b} - \tilde{A}\tilde{x}\|_2 < \text{tol}.$$

When provided an iterative SUNLinSol implementation that does not support the scaling matrices  $S_1$  and  $S_2$ , the SUNDIALS packages will adjust the value of tol accordingly (see the iterative linear tolerance section that follows for more details). In this case, they instead request that iterative linear solvers stop based on the criterion

$$\|P_1^{-1}b - P_1^{-1}Ax\|_2 < \text{tol}.$$

We note that the corresponding adjustments to tol in this case may not be optimal, in that they cannot balance error between specific entries of the solution  $x$ , only the aggregate error in the overall solution vector.

We further note that not all of the SUNDIALS-provided iterative linear solvers support the full range of the above options (e.g., separate left/right preconditioning), and that some of the SUNDIALS packages only utilize a subset of these options. Further details on these exceptions are described in the documentation for each SUNLinearSolver implementation, or for each SUNDIALS package.

For users interested in providing their own SUNLinSol module, the following section presents the SUNLinSol API and its implementation beginning with the definition of SUNLinSol functions in §8.1.1 – §8.1.3. This is followed by the definition of functions supplied to a linear solver implementation in §8.1.4. The linear solver return codes are described in Table 8.1. The SUNLinearSolver type and the generic SUNLinSol module are defined in §8.1.6. §8.1.8 lists the requirements for supplying a custom SUNLinSol module and discusses some intended use cases. Users wishing to supply their own SUNLinSol module are encouraged to use the SUNLinSol implementations provided with SUNDIALS as a template for supplying custom linear solver modules. The section that then follows describes the SUNLinSol functions required by this SUNDIALS package, and provides additional package specific details. Then the remaining sections of this chapter present the SUNLinSol modules provided with SUNDIALS.

## 8.1 The SUNLinearSolver API

The SUNLinSol API defines several linear solver operations that enable SUNDIALS packages to utilize this API. These functions can be divided into three categories. The first are the core linear solver functions. The second consist of “set” routines to supply the linear solver with functions provided by the SUNDIALS packages and to modify solver parameters. The final group consists of “get” routines for retrieving linear solver statistics. All of these functions are defined in the header file `sundials/sundials_linear solver.h`.

### 8.1.1 SUNLinearSolver core functions

The core linear solver functions consist of two **required** functions: `SUNLinSolGetType()` returns the linear solver type, and `SUNLinSolSolve()` solves the linear system  $Ax = b$ .

The remaining **optional** functions return the solver ID (`SUNLinSolGetID()`), initialize the linear solver object once all solver-specific options have been set (`SUNLinSolInitialize()`), set up the linear solver object to utilize an updated matrix  $A$  (`SUNLinSolSetup()`), and destroy a linear solver object (`SUNLinSolFree()`).

enum **SUNLinearSolver\_Type**

An identifier indicating the type of linear solver.

---

**Note:** See §8.1.8.1 for more information on intended use cases corresponding to the linear solver type.

---

**enumerator `SUNLINEARSOLVER_DIRECT`**

The linear solver requires a matrix, and computes an “exact” solution to the linear system defined by that matrix.

**enumerator `SUNLINEARSOLVER_ITERATIVE`**

The linear solver does not require a matrix (though one may be provided), and computes an inexact solution to the linear system using a matrix-free iterative algorithm. That is it solves the linear system defined by the package-supplied `ATimes` routine (see `SUNLinSolSetATimes()` below), even if that linear system differs from the one encoded in the matrix object (if one is provided). As the solver computes the solution only inexactly (or may diverge), the linear solver should check for solution convergence/accuracy as appropriate.

**enumerator `SUNLINEARSOLVER_MATRIX_ITERATIVE`**

The linear solver module requires a matrix, and computes an inexact solution to the linear system defined by that matrix using an iterative algorithm. That is it solves the linear system defined by the matrix object even if that linear system differs from that encoded by the package-supplied `ATimes` routine. As the solver computes the solution only inexactly (or may diverge), the linear solver should check for solution convergence/accuracy as appropriate.

**enumerator `SUNLINEARSOLVER_MATRIX_EMBEDDED`**

The linear solver sets up and solves the specified linear system at each linear solve call. Any matrix-related data structures are held internally to the linear solver itself, and are not provided by the SUNDIALS package.

***SUNLinearSolver\_Type* `SUNLinSolGetType(SUNLinearSolver LS)`**

Returns the *SUNLinearSolver\_Type* type identifier for the linear solver.

**Usage:**

```
type = SUNLinSolGetType(LS);
```

***SUNLinearSolver\_ID* `SUNLinSolGetID(SUNLinearSolver LS)`**

Returns a non-negative linear solver identifier (of type `int`) for the linear solver *LS*.

**Return value:**

Non-negative linear solver identifier (of type `int`), defined by the enumeration `SUNLinearSolver_ID`, with values shown in Table 8.2 and defined in the `sundials_linearsolver.h` header file.

**Usage:**

```
id = SUNLinSolGetID(LS);
```

---

**Note:** It is recommended that a user-supplied `SUNLinearSolver` return the `SUNLINEARSOLVER_CUSTOM` identifier.

---

***SUNErrCode* `SUNLinSolInitialize(SUNLinearSolver LS)`**

Performs linear solver initialization (assuming that all solver-specific options have been set).

**Return value:**

A *SUNErrCode*.

**Usage:**

```
retval = SUNLinSolInitialize(LS);
```

int **SUNLinSolSetup**(*SUNLinearSolver* LS, *SUNMatrix* A)

Performs any linear solver setup needed, based on an updated system *SUNMatrix* A. This may be called frequently (e.g., with a full Newton method) or infrequently (for a modified Newton method), based on the type of integrator and/or nonlinear solver requesting the solves.

**Return value:**

Zero for a successful call, a positive value for a recoverable failure, and a negative value for an unrecoverable failure. Ideally this should return one of the generic error codes listed in [Table 8.1](#).

**Usage:**

```
retval = SUNLinSolSetup(LS, A);
```

int **SUNLinSolSolve**(*SUNLinearSolver* LS, *SUNMatrix* A, *N\_Vector* x, *N\_Vector* b, *sunrealtype* tol)

This *required* function solves a linear system  $Ax = b$ .

**Arguments:**

- *LS* – a *SUNLinSol* object.
- *A* – a *SUNMatrix* object.
- *x* – an *N\_Vector* object containing the initial guess for the solution of the linear system on input, and the solution to the linear system upon return.
- *b* – an *N\_Vector* object containing the linear system right-hand side.
- *tol* – the desired linear solver tolerance.

**Return value:**

Zero for a successful call, a positive value for a recoverable failure, and a negative value for an unrecoverable failure. Ideally this should return one of the generic error codes listed in [Table 8.1](#).

**Notes:**

**Direct solvers:** can ignore the *tol* argument.

**Matrix-free solvers:** (those that identify as *SUNLINEARSOLVER\_ITERATIVE*) can ignore the *SUNMatrix* input *A*, and should rely on the matrix-vector product function supplied through the routine *SUNLinSolSetATimes()*.

**Iterative solvers:** (those that identify as *SUNLINEARSOLVER\_ITERATIVE* or *SUNLINEARSOLVER\_MATRIX\_ITERATIVE*) should attempt to solve to the specified tolerance *tol* in a weighted 2-norm. If the solver does not support scaling then it should just use a 2-norm.

**Matrix-embedded solvers:** should ignore the *SUNMatrix* input *A* as this will be NULL. It is assumed that within this function, the solver will call interface routines from the relevant SUNDIALS package to directly form the linear system matrix *A*, and then solve  $Ax = b$  before returning with the solution *x*.

**Usage:**

```
retval = SUNLinSolSolve(LS, A, x, b, tol);
```

*SUNErrCode* **SUNLinSolFree**(*SUNLinearSolver* LS)

Frees memory allocated by the linear solver.

**Return value:**

A *SUNErrCode*.

**Usage:**

```
retval = SUNLinSolFree(LS);
```

### 8.1.2 SUNLinearSolver “set” functions

The following functions supply linear solver modules with functions defined by the SUNDIALS packages and modify solver parameters. Only the routine for setting the matrix-vector product routine is required, and even then is only required for matrix-free linear solver modules. Otherwise, all other set functions are optional. SUNLinSol implementations that do not provide the functionality for any optional routine should leave the corresponding function pointer NULL instead of supplying a dummy routine.

**SUNErrCode SUNLinSolSetATimes**(*SUNLinearSolver* LS, void \*A\_data, *SUNATimesFn* ATimes)

*Required for matrix-free linear solvers (otherwise optional).*

Provides a *SUNATimesFn* function pointer, as well as a void\* pointer to a data structure used by this routine, to the linear solver object LS. SUNDIALS packages call this function to set the matrix-vector product function to either a solver-provided difference-quotient via vector operations or a user-supplied solver-specific routine.

**Return value:**

A *SUNErrCode*.

**Usage:**

```
retval = SUNLinSolSetATimes(LS, A_data, ATimes);
```

**SUNErrCode SUNLinSolSetPreconditioner**(*SUNLinearSolver* LS, void \*P\_data, *SUNPSetupFn* Pset, *SUNPSolveFn* Psol)

This *optional* routine provides *SUNPSetupFn* and *SUNPSolveFn* function pointers that implement the preconditioner solves  $P_1^{-1}$  and  $P_2^{-1}$  from (8.2). This routine is called by a SUNDIALS package, which provides translation between the generic *Pset* and *Psol* calls and the package- or user-supplied routines.

**Return value:**

A *SUNErrCode*.

**Usage:**

```
retval = SUNLinSolSetPreconditioner(LS, Pdata, Pset, Psol);
```

**SUNErrCode SUNLinSolSetScalingVectors**(*SUNLinearSolver* LS, *N\_Vector* s1, *N\_Vector* s2)

This *optional* routine provides left/right scaling vectors for the linear system solve. Here, *s1* and *s2* are vectors of positive scale factors containing the diagonal of the matrices  $S_1$  and  $S_2$  from (8.2), respectively. Neither vector needs to be tested for positivity, and a NULL argument for either indicates that the corresponding scaling matrix is the identity.

**Return value:**

A *SUNErrCode*.

**Usage:**

```
retval = SUNLinSolSetScalingVectors(LS, s1, s2);
```

**SUNErrCode SUNLinSolSetZeroGuess**(*SUNLinearSolver* LS, *sunbooleantype* onoff)

This *optional* routine indicates if the upcoming *SUNLinSolSolve()* call will be made with a zero initial guess (SUNTRUE) or a non-zero initial guess (SUNFALSE).

**Return value:**

A *SUNErrCode*.

**Usage:**

```
retval = SUNLinSolSetZeroGuess(LS, onoff);
```

**Notes:**

It is assumed that the initial guess status is not retained across calls to *SUNLinSolSolve()*. As such, the linear solver interfaces in each of the SUNDIALS packages call *SUNLinSolSetZeroGuess()* prior to each call to *SUNLinSolSolve()*.

### 8.1.3 SUNLinearSolver “get” functions

The following functions allow SUNDIALS packages to retrieve results from a linear solve. *All routines are optional.*

int *SUNLinSolNumIters*(*SUNLinearSolver* LS)

This *optional* routine should return the number of linear iterations performed in the most-recent “solve” call.

**Usage:**

```
its = SUNLinSolNumIters(LS);
```

*sunrealtype* *SUNLinSolResNorm*(*SUNLinearSolver* LS)

This *optional* routine should return the final residual norm from the most-recent “solve” call.

**Usage:**

```
rnorm = SUNLinSolResNorm(LS);
```

*N\_Vector* *SUNLinSolResid*(*SUNLinearSolver* LS)

If an iterative method computes the preconditioned initial residual and returns with a successful solve without performing any iterations (i.e., either the initial guess or the preconditioner is sufficiently accurate), then this *optional* routine may be called by the SUNDIALS package. This routine should return the *N\_Vector* containing the preconditioned initial residual vector.

**Usage:**

```
rvec = SUNLinSolResid(LS);
```

**Notes:**

Since *N\_Vector* is actually a pointer, and the results are not modified, this routine should *not* require additional memory allocation. If the *SUNLinSol* object does not retain a vector for this purpose, then this function pointer should be set to NULL in the implementation.

*sunindextype* *SUNLinSolLastFlag*(*SUNLinearSolver* LS)

This *optional* routine should return the last error flag encountered within the linear solver. Although not called by the SUNDIALS packages directly, this may be called by the user to investigate linear solver issues after a failed solve.

**Usage:**

```
lflag = SUNLinSolLastFlag(LS);
```

*SUNErrCode* *SUNLinSolSpace*(*SUNLinearSolver* LS, long int \*lenrwLS, long int \*leniwLS)

This *optional* routine should return the storage requirements for the linear solver *LS*:

- *lrw* is a long int containing the number of sunrealtype words
- *liw* is a long int containing the number of integer words.

This function is advisory only, for use by users to help determine their total space requirements.

**Return value:**

A *SUNErrCode*.

**Usage:**

```
retval = SUNLinSolSpace(LS, &lrw, &liw);
```

## 8.1.4 Functions provided by SUNDIALS packages

To interface with SUNLinSol modules, the SUNDIALS packages supply a variety of routines for evaluating the matrix-vector product, and setting up and applying the preconditioner. These package-provided routines translate between the user-supplied ODE, DAE, or nonlinear systems and the generic linear solver API. The function types for these routines are defined in the header file `sundials/sundials_iterative.h`, and are described below.

typedef int (\***SUNTimesFn**)(void \*A\_data, *N\_Vector* v, *N\_Vector* z)

Computes the action of a matrix on a vector, performing the operation  $z \leftarrow Av$ . Memory for  $z$  will already be allocated prior to calling this function. The parameter *A\_data* is a pointer to any information about  $A$  which the function needs in order to do its job. The vector  $v$  should be left unchanged.

**Return value:**

Zero for a successful call, and non-zero upon failure.

typedef int (\***SUNPSetupFn**)(void \*P\_data)

Sets up any requisite problem data in preparation for calls to the corresponding *SUNPSolveFn*.

**Return value:**

Zero for a successful call, and non-zero upon failure.

typedef int (\***SUNPSolveFn**)(void \*P\_data, *N\_Vector* r, *N\_Vector* z, *sunrealtype* tol, int lr)

Solves the preconditioner equation  $Pz = r$  for the vector  $z$ . Memory for  $z$  will already be allocated prior to calling this function. The parameter *P\_data* is a pointer to any information about  $P$  which the function needs in order to do its job (set up by the corresponding *SUNPSetupFn*). The parameter *lr* is input, and indicates whether  $P$  is to be taken as the left or right preconditioner:  $lr = 1$  for left and  $lr = 2$  for right. If preconditioning is on one side only, *lr* can be ignored. If the preconditioner is iterative, then it should strive to solve the preconditioner equation so that

$$\|Pz - r\|_{\text{wrms}} < \text{tol}$$

where the error weight vector for the WRMS norm may be accessed from the main package memory structure. The vector  $r$  should not be modified by the *SUNPSolveFn*.

**Return value:**

Zero for a successful call, a negative value for an unrecoverable failure condition, or a positive value for a recoverable failure condition (thus the calling routine may reattempt the solution after updating preconditioner data).

### 8.1.5 SUNLinearSolver return codes

The functions provided to SUNLinSol modules by each SUNDIALS package, and functions within the SUNDIALS-provided SUNLinSol implementations, utilize a common set of return codes, listed in [Table 8.1](#). These adhere to a common pattern:

- 0 indicates success
- a positive value corresponds to a recoverable failure, and
- a negative value indicates a non-recoverable failure.

Aside from this pattern, the actual values of each error code provide additional information to the user in case of a linear solver failure.

Table 8.1: SUNLinSol error codes

Error code	Value	Meaning
SUN_SUCCESS	0	successful call or converged solve
SUNLS_ATIMES_NULL	-804	the <code>ATimes</code> function is NULL
SUNLS_ATIMES_FAIL_-UNREC	-805	an unrecoverable failure occurred in the <code>ATimes</code> routine
SUNLS_PSET_FAIL_-UNREC	-806	an unrecoverable failure occurred in the <code>Pset</code> routine
SUNLS_PSOLVE_NULL	-807	the preconditioner solve function is NULL
SUNLS_PSOLVE_FAIL_-UNREC	-808	an unrecoverable failure occurred in the <code>Psolve</code> routine
SUNLS_GS_FAIL	-810	a failure occurred during Gram-Schmidt orthogonalization (SPGMR/SPFGMR)
SUNLS_QRSOL_FAIL	-811	a singular $SR$ matrix was encountered in a QR factorization (SPGMR/SPFGMR)
SUNLS_RES_REDUCED	801	an iterative solver reduced the residual, but did not converge to the desired tolerance
SUNLS_CONV_FAIL	802	an iterative solver did not converge (and the residual was not reduced)
SUNLS_ATIMES_FAIL_-REC	803	a recoverable failure occurred in the <code>ATimes</code> routine
SUNLS_PSET_FAIL_REC	804	a recoverable failure occurred in the <code>Pset</code> routine
SUNLS_PSOLVE_FAIL_-REC	805	a recoverable failure occurred in the <code>Psolve</code> routine
SUNLS_PACKAGE_FAIL_-REC	806	a recoverable failure occurred in an external linear solver package
SUNLS_QRFACT_FAIL	807	a singular matrix was encountered during a QR factorization (SPGMR/SPFGMR)
SUNLS_LUFACT_FAIL	808	a singular matrix was encountered during a LU factorization



### 8.1.6 The generic SUNLinearSolver module

SUNDIALS packages interact with linear solver implementations through the *SUNLinearSolver* class. A *SUNLinearSolver* is a pointer to the *\_generic\_SUNLinearSolver* structure:

```
typedef struct _generic_SUNLinearSolver *SUNLinearSolver
```

```
struct _generic_SUNLinearSolver
```

The structure defining the SUNDIALS linear solver class.

void \***content**

Pointer to the linear solver-specific member data

*SUNLinearSolver\_Ops* **ops**

A virtual table of linear solver operations provided by a specific implementation

*SUNContext* **sunctx**

The SUNDIALS simulation context

The virtual table structure is defined as

```
typedef struct _generic_SUNLinearSolver_Ops *SUNLinearSolver_Ops
```

```
struct _generic_SUNLinearSolver_Ops
```

The structure defining *SUNLinearSolver* operations.

*SUNLinearSolver\_Type* (\***gettype**)(*SUNLinearSolver*)

The function implementing *SUNLinSolGetType()*

*SUNLinearSolver\_ID* (\***getid**)(*SUNLinearSolver*)

The function implementing *SUNLinSolGetID()*

*SUNErrCode* (\***setatimes**)(*SUNLinearSolver*, void\*, *SUNATimesFn*)

The function implementing *SUNLinSolSetATimes()*

*SUNErrCode* (\***setpreconditioner**)(*SUNLinearSolver*, void\*, *SUNPSSetupFn*, *SUNPSolveFn*)

The function implementing *SUNLinSolSetPreconditioner()*

*SUNErrCode* (\***setscalingvectors**)(*SUNLinearSolver*, *N\_Vector*, *N\_Vector*)

The function implementing *SUNLinSolSetScalingVectors()*

*SUNErrCode* (\***setzeroguess**)(*SUNLinearSolver*, *sunbooleantype*)

The function implementing *SUNLinSolSetZeroGuess()*

*SUNErrCode* (\***initialize**)(*SUNLinearSolver*)

The function implementing *SUNLinSolInitialize()*

int (\***setup**)(*SUNLinearSolver*, *SUNMatrix*)

The function implementing *SUNLinSolSetup()*

int (\***solve**)(*SUNLinearSolver*, *SUNMatrix*, *N\_Vector*, *N\_Vector*, *sunrealtype*)

The function implementing *SUNLinSolSolve()*

int (\***numiters**)(*SUNLinearSolver*)

The function implementing *SUNLinSolNumIters()*

*sunrealtype* (\***resnorm**)(*SUNLinearSolver*)

The function implementing *SUNLinSolResNorm()*

*sunindextype* (\***lastflag**)(*SUNLinearSolver*)

The function implementing *SUNLinSolLastFlag()*

*SUNErrCode* (\***space**)(*SUNLinearSolver*, long int\*, long int\*)

The function implementing *SUNLinSolSpace()*

*N\_Vector* (\***resid**)(*SUNLinearSolver*)

The function implementing *SUNLinSolResid()*

*SUNErrCode* (\***free**)(*SUNLinearSolver*)

The function implementing *SUNLinSolFree()*

The generic *SUNLinSol* class defines and implements the linear solver operations defined in §8.1.1 – §8.1.3. These routines are in fact only wrappers to the linear solver operations defined by a particular *SUNLinSol* implementation, which are accessed through the *ops* field of the *SUNLinearSolver* structure. To illustrate this point we show below the implementation of a typical linear solver operation from the *SUNLinearSolver* base class, namely *SUNLinSolInitialize()*, that initializes a *SUNLinearSolver* object for use after it has been created and configured, and returns a flag denoting a successful or failed operation:

```
int SUNLinSolInitialize(SUNLinearSolver S)
{
    return ((int) S->ops->initialize(S));
}
```

### 8.1.7 Compatibility of *SUNLinearSolver* modules

Not all *SUNLinearSolver* implementations are compatible with all *SUNMatrix* and *N\_Vector* implementations provided in SUNDIALS. More specifically, all of the SUNDIALS iterative linear solvers (*SPGMR*, *SPFGMR*, *SPBCGS*, *SPTFQMR*, and *PCG*) are compatible with all of the SUNDIALS *N\_Vector* modules, but the matrix-based direct *SUNLinSol* modules are specifically designed to work with distinct *SUNMatrix* and *N\_Vector* modules. In the list below, we summarize the compatibility of each matrix-based *SUNLinearSolver* module with the various *SUNMatrix* and *N\_Vector* modules. For a more thorough discussion of these compatibilities, we defer to the documentation for each individual *SUNLinSol* module in the sections that follow.

- *Dense*
  - *SUNMatrix*: *Dense* or user-supplied
  - *N\_Vector*: *Serial*, *OpenMP*, *Pthreads*, or user-supplied
- *LapackDense*
  - *SUNMatrix*: *Dense* or user-supplied
  - *N\_Vector*: *Serial*, *OpenMP*, *Pthreads*, or user-supplied
- *Band*
  - *SUNMatrix*: *Band* or user-supplied
  - *N\_Vector*: *Serial*, *OpenMP*, *Pthreads*, or user-supplied
- *LapackBand*
  - *SUNMatrix*: *Band* or user-supplied
  - *N\_Vector*: *Serial*, *OpenMP*, *Pthreads*, or user-supplied
- *KLU*
  - *SUNMatrix*: *Sparse* or user-supplied

- N\_Vector: *Serial*, *OpenMP*, *Pthreads*, or user-supplied
- *SuperLU\_MT*
  - SUNMatrix: *Sparse* or user-supplied
  - N\_Vector: *Serial*, *OpenMP*, *Pthreads*, or user-supplied
- *SuperLU\_Dist*
  - SUNMatrix: *SLUNRLOC* or user-supplied
  - N\_Vector: *Serial*, *OpenMP*, *Pthreads*, *Parallel*, *\*hypr\**, *PETSc*, or user-supplied
- *Magma Dense*
  - SUNMatrix: *Magma Dense* or user-supplied
  - N\_Vector: *HIP*, *RAJA*, or user-supplied
- *OneMKL Dense*
  - SUNMatrix: *One MKL Dense* or user-supplied
  - N\_Vector: *SYCL*, *RAJA*, or user-supplied
- *cuSolverSp batchQR*
  - SUNMatrix: *cuSparse* or user-supplied
  - N\_Vector: *CUDA*, *RAJA*, or user-supplied

### 8.1.8 Implementing a custom SUNLinearSolver module

A particular implementation of the SUNLinearSolver module must:

- Specify the *content* field of the SUNLinSol module.
- Define and implement the required linear solver operations.

---

**Note:** The names of these routines should be unique to that implementation in order to permit using more than one SUNLinSol module (each with different SUNLinearSolver internal data representations) in the same code.

---

- Define and implement user-callable constructor and destructor routines to create and free a SUNLinearSolver with the new *content* field and with *ops* pointing to the new linear solver operations.

We note that the function pointers for all unsupported optional routines should be set to NULL in the *ops* structure. This allows the SUNDIALS package that is using the SUNLinSol object to know whether the associated functionality is supported.

To aid in the creation of custom SUNLinearSolver modules the generic SUNLinearSolver module provides the utility function *SUNLinSolNewEmpty()*. When used in custom SUNLinearSolver constructors this function will ease the introduction of any new optional linear solver operations to the SUNLinearSolver API by ensuring that only required operations need to be set.

*SUNLinearSolver* **SUNLinSolNewEmpty**(*SUNContext* sunctx)

This function allocates a new generic SUNLinearSolver object and initializes its content pointer and the function pointers in the operations structure to NULL.

**Return value:**

If successful, this function returns a SUNLinearSolver object. If an error occurs when allocating the object, then this routine will return NULL.

void **SUNLinSolFreeEmpty**(*SUNLinearSolver* LS)

This routine frees the generic *SUNLinearSolver* object, under the assumption that any implementation-specific data that was allocated within the underlying content structure has already been freed. It will additionally test whether the ops pointer is NULL, and, if it is not, it will free it as well.

**Arguments:**

- *LS* – a *SUNLinearSolver* object

Additionally, a *SUNLinearSolver* implementation *may* do the following:

- Define and implement additional user-callable “set” routines acting on the *SUNLinearSolver*, e.g., for setting various configuration options to tune the linear solver for a particular problem.
- Provide additional user-callable “get” routines acting on the *SUNLinearSolver* object, e.g., for returning various solve statistics.

enum **SUNLinearSolver\_ID**

Each *SUNLinSol* implementation included in SUNDIALS has a unique identifier specified in enumeration and shown in [Table 8.2](#). It is recommended that a user-supplied *SUNLinSol* implementation use the *SUNLINEAR-SOLVER\_CUSTOM* identifier.

Table 8.2: Identifiers associated with *SUNLinearSolver* modules supplied with SUNDIALS

SUNLinSol ID	Linear solver type	ID Value
SUNLINEARSOLVER_BAND	Banded direct linear solver (internal)	0
SUNLINEARSOLVER_DENSE	Dense direct linear solver (internal)	1
SUNLINEARSOLVER_KLU	Sparse direct linear solver (KLU)	2
SUNLINEARSOLVER_LAPACKBAND	Banded direct linear solver (LAPACK)	3
SUNLINEARSOLVER_LAPACKDENSE	Dense direct linear solver (LAPACK)	4
SUNLINEARSOLVER_PCG	Preconditioned conjugate gradient iterative solver	5
SUNLINEARSOLVER_SPBCGS	Scaled-preconditioned BiCGStab iterative solver	6
SUNLINEARSOLVER_SPGMR	Scaled-preconditioned FGMRES iterative solver	7
SUNLINEARSOLVER_SPGMR	Scaled-preconditioned GMRES iterative solver	8
SUNLINEARSOLVER_SPTFQMR	Scaled-preconditioned TFQMR iterative solver	9
SUNLINEARSOLVER_SUPERLUDIST	Parallel sparse direct linear solver (SuperLU_Dist)	10
SUNLINEARSOLVER_SUPERLUMT	Threaded sparse direct linear solver (SuperLU_MT)	11
SUNLINEARSOLVER_CUSOLVERSP_BATCHQR	Sparse direct linear solver (CUDA)	12
SUNLINEARSOLVER_MAGMADENSE	Dense or block-dense direct linear solver (MAGMA)	13
SUNLINEARSOLVER_ONEMKLDENSE	Dense or block-dense direct linear solver (OneMKL)	14
SUNLINEARSOLVER_CUSTOM	User-provided custom linear solver	15

### 8.1.8.1 Intended use cases

The SUNLinSol and SUNMATRIX APIs are designed to require a minimal set of routines to ease interfacing with custom or third-party linear solver libraries. Many external solvers provide routines with similar functionality and thus may require minimal effort to wrap within custom SUNMATRIX and SUNLinSol implementations. As SUNDIALS packages utilize generic SUNLinSol modules they may naturally leverage user-supplied SUNLinearSolver implementations, thus there exist a wide range of possible linear solver combinations. Some intended use cases for both the SUNDIALS-provided and user-supplied SUNLinSol modules are discussed in the sections below.

#### Direct linear solvers

Direct linear solver modules require a matrix and compute an “exact” solution to the linear system *defined by the matrix*. SUNDIALS packages strive to amortize the high cost of matrix construction by reusing matrix information for multiple nonlinear iterations or time steps. As a result, each package’s linear solver interface recomputes matrix information as infrequently as possible.

Alternative matrix storage formats and compatible linear solvers that are not currently provided by, or interfaced with, SUNDIALS can leverage this infrastructure with minimal effort. To do so, a user must implement custom SUNMATRIX and SUNLinSol wrappers for the desired matrix format and/or linear solver following the APIs described in §7 and §8. *This user-supplied SUNLinSol module must then self-identify as having SUNLINEARSOLVER\_DIRECT type.*

#### Matrix-free iterative linear solvers

Matrix-free iterative linear solver modules do not require a matrix, and instead compute an inexact solution to the linear system *defined by the package-supplied ATimes routine*. SUNDIALS supplies multiple scaled, preconditioned iterative SUNLinSol modules that support scaling, allowing packages to handle non-dimensionalization, and users to define variables and equations as natural in their applications. However, for linear solvers that do not support left/right scaling, SUNDIALS packages must instead adjust the tolerance supplied to the linear solver to compensate (see the iterative linear tolerance section that follows for more details) – this strategy may be non-optimal since it cannot handle situations where the magnitudes of different solution components or equations vary dramatically within a single application.

To utilize alternative linear solvers that are not currently provided by, or interfaced with, SUNDIALS a user must implement a custom SUNLinSol wrapper for the linear solver following the API described in §8. *This user-supplied SUNLinSol module must then self-identify as having SUNLINEARSOLVER\_ITERATIVE type.*

#### Matrix-based iterative linear solvers (reusing A)

Matrix-based iterative linear solver modules require a matrix and compute an inexact solution to the linear system *defined by the matrix*. This matrix will be updated infrequently and reused across multiple solves to amortize the cost of matrix construction. As in the direct linear solver case, only thin SUNMATRIX and SUNLinSol wrappers for the underlying matrix and linear solver structures need to be created to utilize such a linear solver. *This user-supplied SUNLinSol module must then self-identify as having SUNLINEARSOLVER\_MATRIX\_ITERATIVE type.*

At present, SUNDIALS has one example problem that uses this approach for wrapping a structured-grid matrix, linear solver, and preconditioner from the *hypr* library; this may be used as a template for other customized implementations (see `examples/arkode/CXX_parhyp/ark_heat2D_hypr.cpp`).

### Matrix-based iterative linear solvers (current *A*)

For users who wish to utilize a matrix-based iterative linear solver where the matrix is *purely for preconditioning* and the linear system is *defined by the package-supplied `ATimes` routine*, we envision two current possibilities.

The preferred approach is for users to employ one of the SUNDIALS scaled, preconditioned iterative linear solver implementations (`SUNLinSol_SPGMR()`, `SUNLinSol_SPFGMR()`, `SUNLinSol_SPBCGS()`, `SUNLinSol_SPTFQMR()`, or `SUNLinSol_PCG()`) as the outer solver. The creation and storage of the preconditioner matrix, and interfacing with the corresponding matrix-based linear solver, can be handled through a package's preconditioner "setup" and "solve" functionality without creating SUNMATRIX and SUNLinSol implementations. This usage mode is recommended primarily because the SUNDIALS-provided modules support variable and equation scaling as described above.

A second approach supported by the linear solver APIs is as follows. If the SUNLinSol implementation is matrix-based, *self-identifies as having `SUNLINEARSOLVER_ITERATIVE` type*, and *also provides a non-NULL `SUNLinSolSetATimes()` routine*, then each SUNDIALS package will call that routine to attach its package-specific matrix-vector product routine to the SUNLinSol object. The SUNDIALS package will then call the SUNLinSol-provided `SUNLinSolSetup()` routine (infrequently) to update matrix information, but will provide current matrix-vector products to the SUNLinSol implementation through the package-supplied `SUNATimesFn` routine.

### Application-specific linear solvers with embedded matrix structure

Many applications can exploit additional linear system structure arising from the implicit couplings in their model equations. In certain circumstances, the linear solve  $Ax = b$  may be performed without the need for a global system matrix  $A$ , as the unformed  $A$  may be block diagonal or block triangular, and thus the overall linear solve may be performed through a sequence of smaller linear solves. In other circumstances, a linear system solve may be accomplished via specialized fast solvers, such as the fast Fourier transform, fast multipole method, or treecode, in which case no matrix structure may be explicitly necessary. In many of the above situations, construction and preprocessing of the linear system matrix  $A$  may be inexpensive, and thus increased performance may be possible if the current linear system information is used within every solve (instead of being lagged, as occurs with matrix-based solvers that reuse  $A$ ).

To support such application-specific situations, SUNDIALS supports user-provided linear solvers with the `SUNLINEARSOLVER_MATRIX_EMBEDDED` type. For an application to leverage this support, it should define a custom SUNLinSol implementation having this type, that only needs to implement the required `SUNLinSolGetType()` and `SUNLinSolSolve()` operations. Within `SUNLinSolSolve()`, the linear solver implementation should call package-specific interface routines (e.g., `ARKStepGetNonlinearSystemData`, `CVodeGetNonlinearSystemData`, `IDAGetNonlinearSystemData`, `ARKStepGetCurrentGamma`, `CVodeGetCurrentGamma`, `IDAGetCurrentCj`, or `MRISetGetCurrentGamma`) to construct the relevant system matrix  $A$  (or portions thereof), solve the linear system  $Ax = b$ , and return the solution vector  $x$ .

We note that when attaching this custom SUNLinearSolver object with the relevant SUNDIALS package `SetLinearSolver` routine, the input `SUNMatrix A` should be set to `NULL`.

For templates of such user-provided "matrix-embedded" SUNLinSol implementations, see the SUNDIALS examples `ark_analytic_mels.c`, `cvAnalytic_mels.c`, `cvsAnalytic_mels.c`, `idaAnalytic_mels.c`, and `idasAnalytic_mels.c`.

## 8.2 IDAS SUNLinearSolver interface

Table 8.3 below lists the SUNLinearSolver module linear solver functions used within the IDALS interface. As with the SUNMatrix module, we emphasize that the IDA user does not need to know detailed usage of linear solver functions by the IDA code modules in order to use IDA. The information is presented as an implementation detail for the interested reader.

The linear solver functions listed below are marked with ‘x’ to indicate that they are required, or with † to indicate that they are only called if they are non-NULL in the SUNLinearSolver implementation that is being used. Note:

1. Although IDALS does not call SUNLinSolLastFlag directly, this routine is available for users to query linear solver issues directly.
2. Although IDALS does not call SUNLinSolFree directly, this routine should be available for users to call when cleaning up from a simulation.

Table 8.3: List of linear solver function usage in the IDALS interface

	DIRECT	ITERATIVE	MATRIX_ITERATIVE
<i>SUNLinSolGetType()</i>	x	x	x
<i>SUNLinSolSetATimes()</i>	†	x	†
<i>SUNLinSolSetPreconditioner()</i>	†	†	†
<i>SUNLinSolSetScalingVectors()</i>	†	†	†
<i>SUNLinSolInitialize()</i>	x	x	x
<i>SUNLinSolSetup()</i>	x	x	x
<i>SUNLinSolSolve()</i>	x	x	x
<i>SUNLinSolNumIters()</i>		x	x
<i>SUNLinSolResid()</i>		x	x
<sup>1</sup> <i>SUNLinSolLastFlag()</i>			
<sup>2</sup> <i>SUNLinSolFree()</i>			
<i>SUNLinSolSpace()</i>	†	†	†

Since there are a wide range of potential SUNLinearSolver use cases, the following subsections describe some details of the IDALS interface, in the case that interested users wish to develop custom SUNLinearSolver modules.

### 8.2.1 Lagged matrix information

If the SUNLinearSolver object self-identifies as having type SUNLINEARSOLVER\_DIRECT or SUNLINEARSOLVER\_MATRIX\_ITERATIVE, then the SUNLinearSolver object solves a linear system *defined* by a SUNMatrix object. IDALS will update the matrix information infrequently according to the strategies outlined in §2. To this end, we differentiate between the *desired* linear system  $Jx = b$  with  $J = \left( \frac{\partial F}{\partial y} - c_j \frac{\partial F}{\partial \dot{y}} \right)$ , and the *actual* linear system  $\bar{J}\bar{x} = b$  with

$$\bar{J} = \frac{\partial \bar{F}}{\partial y} - \bar{c}_j \frac{\partial \bar{F}}{\partial \dot{y}},$$

where the overlines indicate the lagged versions of these numbers and matrices.

Since IDALS updates the SUNMatrix objects infrequently and it is likely that  $c_j \neq \bar{c}_j$ , then typically  $J \neq \bar{J}$ . Thus after calling the SUNLinearSolver-provided SUNLinSolSolve routine, we test whether  $\frac{c_j}{\bar{c}_j} \neq 1$ , and if this is the case we scale the solution  $\bar{x}$  to correct the linear system solution  $x$  via

$$x = \frac{2}{1 + c_j/\bar{c}_j} \bar{x}. \quad (8.3)$$



The motivation for this selection of the scaling factor  $c = 2/(1 + c_j/\bar{c}_j)$  is discussed in detail in [16, 40]. In short, if we consider a stationary iteration for the linear system as consisting of a solve with  $\bar{J}$  followed by scaling by  $c$ , then for a linear constant-coefficient problem, the error in the solution vector will be reduced at each iteration by the error matrix  $E = I - c\bar{J}^{-1}J$ , with a convergence rate given by the spectral radius of  $E$ . Assuming that stiff systems have a spectrum spread widely over the left half-plane,  $c$  is chosen to minimize the magnitude of the eigenvalues of  $E$ .

### 8.2.2 Iterative linear solver tolerance

If the `SUNLinearSolver` object self-identifies as having type `SUNLINEARSOLVER_ITERATIVE` or `SUNLINEAR-SOLVER_MATRIX_ITERATIVE` then IDALS will set the input tolerance `delta` as described in §2.2. However, if the iterative linear solver does not support scaling matrices (i.e., the `SUNLinSolSetScalingVectors` routine is `NULL`), then IDALS will attempt to adjust the linear solver tolerance to account for this lack of functionality. To this end, the following assumptions are made:

1. All solution components have similar magnitude; hence the error weight vector  $W$  used in the WRMS norm (see §2.2) should satisfy the assumption

$$W_i \approx W_{mean}, \quad \text{for } i = 0, \dots, n-1.$$

2. The `SUNLinearSolver` object uses a standard 2-norm to measure convergence.

Since IDA uses identical left and right scaling matrices,  $S_1 = S_2 = S = \text{diag}(W)$ , then the linear solver convergence requirement is converted as follows (using the notation from equations (8.1) – (8.2)):

$$\begin{aligned} & \|\tilde{b} - \tilde{A}\tilde{x}\|_2 < \text{tol} \\ \Leftrightarrow & \|SP_1^{-1}b - SP_1^{-1}Ax\|_2 < \text{tol} \\ \Leftrightarrow & \sum_{i=0}^{n-1} [W_i (P_1^{-1}(b - Ax))_i]^2 < \text{tol}^2 \\ \Leftrightarrow & W_{mean}^2 \sum_{i=0}^{n-1} [(P_1^{-1}(b - Ax))_i]^2 < \text{tol}^2 \\ \Leftrightarrow & \sum_{i=0}^{n-1} [(P_1^{-1}(b - Ax))_i]^2 < \left(\frac{\text{tol}}{W_{mean}}\right)^2 \\ \Leftrightarrow & \|P_1^{-1}(b - Ax)\|_2 < \frac{\text{tol}}{W_{mean}} \end{aligned}$$

Therefore the tolerance scaling factor

$$W_{mean} = \|W\|_2/\sqrt{n}$$

is computed and the scaled tolerance `delta = tol/Wmean` is supplied to the `SUNLinearSolver` object.

## 8.3 The SUNLinSol\_Band Module

The `SUNLinSol_Band` implementation of the `SUNLinearSolver` class is designed to be used with the corresponding `SUNMATRIX_BAND` matrix type, and one of the serial or shared-memory `N_Vector` implementations (`NVECTOR_SERIAL`, `NVECTOR_OPENMP` or `NVECTOR_PTHREADS`).



### 8.3.1 SUNLinSol\_Band Usage

The header file to be included when using this module is `sunlinsol/sunlinsol_band.h`. The `SUNLinSol_Band` module is accessible from all SUNDIALS packages *without* linking to the `libsundials_sunlinsolband` module library.

The `SUNLinSol_Band` module provides the following user-callable constructor routine:

*SUNLinearSolver* **SUNLinSol\_Band**(*N\_Vector* y, *SUNMatrix* A, *SUNContext* sunctx)

This function creates and allocates memory for a band `SUNLinearSolver`.

**Arguments:**

- y – vector used to determine the linear system size
- A – matrix used to assess compatibility
- sunctx – the *SUNContext* object (see §4.2)

**Return value:**

New `SUNLinSol_Band` object, or NULL if either A or y are incompatible.

**Notes:**

This routine will perform consistency checks to ensure that it is called with consistent *N\_Vector* and *SUNMatrix* implementations. These are currently limited to the `SUNMATRIX_BAND` matrix type and the `NVECTOR_SERIAL`, `NVECTOR_OPENMP`, and `NVECTOR_PTHREADS` vector types. As additional compatible matrix and vector implementations are added to SUNDIALS, these will be included within this compatibility check.

Additionally, this routine will verify that the input matrix A is allocated with appropriate upper bandwidth storage for the *LU* factorization.

### 8.3.2 SUNLinSol\_Band Description

The `SUNLinSol_Band` module defines the *content* field of a `SUNLinearSolver` to be the following structure:

```
struct _SUNLinearSolverContent_Band {
    sunindextype N;
    sunindextype *pivots;
    sunindextype last_flag;
};
```

These entries of the *content* field contain the following information:

- N - size of the linear system,
- pivots - index array for partial pivoting in LU factorization,
- last\_flag - last error return flag from internal function evaluations.

This solver is constructed to perform the following operations:

- The “setup” call performs an *LU* factorization with partial (row) pivoting,  $PA = LU$ , where *P* is a permutation matrix, *L* is a lower triangular matrix with 1’s on the diagonal, and *U* is an upper triangular matrix. This factorization is stored in-place on the input `SUNMATRIX_BAND` object *A*, with pivoting information encoding *P* stored in the *pivots* array.
- The “solve” call performs pivoting and forward and backward substitution using the stored *pivots* array and the *LU* factors held in the `SUNMATRIX_BAND` object.

- $A$  must be allocated to accommodate the increase in upper bandwidth that occurs during factorization. More precisely, if  $A$  is a band matrix with upper bandwidth  $\mu$  and lower bandwidth  $m_l$ , then the upper triangular factor  $U$  can have upper bandwidth as big as  $\text{smu} = \text{MIN}(N-1, \mu+m_l)$ . The lower triangular factor  $L$  has lower bandwidth  $m_l$ .

The `SUNLinSol_Band` module defines band implementations of all “direct” linear solver operations listed in §8.1:

- `SUNLinSolGetType_Band`
- `SUNLinSolInitialize_Band` – this does nothing, since all consistency checks are performed at solver creation.
- `SUNLinSolSetup_Band` – this performs the  $LU$  factorization.
- `SUNLinSolSolve_Band` – this uses the  $LU$  factors and `pivots` array to perform the solve.
- `SUNLinSolLastFlag_Band`
- `SUNLinSolSpace_Band` – this only returns information for the storage *within* the solver object, i.e. storage for `N`, `last_flag`, and `pivots`.
- `SUNLinSolFree_Band`

## 8.4 The `SUNLinSol_Dense` Module

The `SUNLinSol_Dense` implementation of the `SUNLinearSolver` class is designed to be used with the corresponding `SUNMATRIX_DENSE` matrix type, and one of the serial or shared-memory `N_Vector` implementations (`NVECTOR_SERIAL`, `NVECTOR_OPENMP` or `NVECTOR_PTHREADS`).

### 8.4.1 `SUNLinSol_Dense` Usage

The header file to be included when using this module is `sunlinsol/sunlinsol_dense.h`. The `SUNLinSol_Dense` module is accessible from all SUNDIALS solvers *without* linking to the `libsundials_sunlinsoldense` module library.

The module `SUNLinSol_Dense` provides the following user-callable constructor routine:

*SUNLinearSolver* **`SUNLinSol_Dense`**(*N\_Vector* `y`, *SUNMatrix* `A`, *SUNContext* `sunctx`)

This function creates and allocates memory for a dense `SUNLinearSolver`.

**Arguments:**

- `y` – vector used to determine the linear system size.
- `A` – matrix used to assess compatibility.
- `sunctx` – the *SUNContext* object (see §4.2)

**Return value:**

New `SUNLinSol_Dense` object, or `NULL` if either `A` or `y` are incompatible.

**Notes:**

This routine will perform consistency checks to ensure that it is called with consistent `N_Vector` and `SUNMatrix` implementations. These are currently limited to the `SUNMATRIX_DENSE` matrix type and the `NVECTOR_SERIAL`, `NVECTOR_OPENMP`, and `NVECTOR_PTHREADS` vector types. As additional compatible matrix and vector implementations are added to SUNDIALS, these will be included within this compatibility check.

## 8.4.2 SUNLinSol\_Dense Description

The SUNLinSol\_Dense module defines the *content* field of a SUNLinearSolver to be the following structure:

```
struct _SUNLinearSolverContent_Dense {
    sunindextype N;
    sunindextype *pivots;
    sunindextype last_flag;
};
```

These entries of the *content* field contain the following information:

- *N* - size of the linear system,
- *pivots* - index array for partial pivoting in LU factorization,
- *last\_flag* - last error return flag from internal function evaluations.

This solver is constructed to perform the following operations:

- The “setup” call performs an *LU* factorization with partial (row) pivoting ( $\mathcal{O}(N^3)$  cost),  $PA = LU$ , where *P* is a permutation matrix, *L* is a lower triangular matrix with 1’s on the diagonal, and *U* is an upper triangular matrix. This factorization is stored in-place on the input SUNMATRIX\_DENSE object *A*, with pivoting information encoding *P* stored in the *pivots* array.
- The “solve” call performs pivoting and forward and backward substitution using the stored *pivots* array and the *LU* factors held in the SUNMATRIX\_DENSE object ( $\mathcal{O}(N^2)$  cost).

The SUNLinSol\_Dense module defines dense implementations of all “direct” linear solver operations listed in §8.1:

- SUNLinSolGetType\_Dense
- SUNLinSolInitialize\_Dense – this does nothing, since all consistency checks are performed at solver creation.
- SUNLinSolSetup\_Dense – this performs the *LU* factorization.
- SUNLinSolSolve\_Dense – this uses the *LU* factors and *pivots* array to perform the solve.
- SUNLinSolLastFlag\_Dense
- SUNLinSolSpace\_Dense – this only returns information for the storage *within* the solver object, i.e. storage for *N*, *last\_flag*, and *pivots*.
- SUNLinSolFree\_Dense

## 8.5 The SUNLinSol\_KLU Module

The SUNLinSol\_KLU implementation of the SUNLinearSolver class is designed to be used with the corresponding SUNMATRIX\_SPARSE matrix type, and one of the serial or shared-memory N\_Vector implementations (NVECTOR\_SERIAL, NVECTOR\_OPENMP, or NVECTOR\_PTHREADS).

### 8.5.1 SUNLinSol\_KLU Usage

The header file to be included when using this module is `sunlinsol/sunlinsol_klu.h`. The installed module library to link to is `libsundials_sunlinsolklu.lib` where `.lib` is typically `.so` for shared libraries and `.a` for static libraries.

The module SUNLinSol\_KLU provides the following additional user-callable routines:

*SUNLinearSolver* **SUNLinSol\_KLU**(*N\_Vector* y, *SUNMatrix* A, *SUNContext* sunctx)

This constructor function creates and allocates memory for a SUNLinSol\_KLU object.

**Arguments:**

- y – vector used to determine the linear system size.
- A – matrix used to assess compatibility.
- sunctx – the *SUNContext* object (see §4.2)

**Return value:**

New SUNLinSol\_KLU object, or NULL if either A or y are incompatible.

**Notes:**

This routine will perform consistency checks to ensure that it is called with consistent *N\_Vector* and *SUNMatrix* implementations. These are currently limited to the *SUNMATRIX\_SPARSE* matrix type (using either CSR or CSC storage formats) and the *NVECTOR\_SERIAL*, *NVECTOR\_OPENMP*, and *NVECTOR\_PTHREADS* vector types. As additional compatible matrix and vector implementations are added to SUNDIALS, these will be included within this compatibility check.

*SUNErrCode* **SUNLinSol\_KLUReInit**(*SUNLinearSolver* S, *SUNMatrix* A, *sunindextype* nnz, int reinit\_type)

This function reinitializes memory and flags for a new factorization (symbolic and numeric) to be conducted at the next solver setup call. This routine is useful in the cases where the number of nonzeros has changed or if the structure of the linear system has changed which would require a new symbolic (and numeric factorization).

**Arguments:**

- S – existing SUNLinSol\_KLU object to reinitialize.
- A – sparse *SUNMatrix* matrix (with updated structure) to use for reinitialization.
- nnz – maximum number of nonzeros expected for Jacobian matrix.
- reinit\_type – governs the level of reinitialization. The allowed values are:
  1. The Jacobian matrix will be destroyed and a new one will be allocated based on the nnz value passed to this call. New symbolic and numeric factorizations will be completed at the next solver setup.
  2. Only symbolic and numeric factorizations will be completed. It is assumed that the Jacobian size has not exceeded the size of nnz given in the sparse matrix provided to the original constructor routine (or the previous *SUNKLUREInit* call).

**Return value:**

- A *SUNErrCode*

**Notes:**

This routine assumes no other changes to solver use are necessary.

*SUNErrCode* **SUNLinSol\_KLUSetOrdering**(*SUNLinearSolver* S, int ordering\_choice)

This function sets the ordering used by KLU for reducing fill in the linear solve.

**Arguments:**

- S – existing SUNLinSol\_KLU object to update.

- *ordering\_choice* – type of ordering to use, options are:

0. AMD,
1. COLAMD, and
2. the natural ordering.

The default is 1 for COLAMD.

**Return value:**

- A *SUNErrCode*

*sun\_klu\_symbolic* \***SUNLinSol\_KLUGetSymbolic**(*SUNLinearSolver* S)

This function returns a pointer to the KLU symbolic factorization stored in the SUNLinSol\_KLU content structure.

type **sun\_klu\_symbolic**

This type is an alias that depends on the SUNDIALS index size, see *sunindextype* and *SUNDIALS\_INDEX\_SIZE*. It is equivalent to:

- *klu\_symbolic* when SUNDIALS is compiled with 32-bit indices
- *klu\_l\_symbolic* when SUNDIALS is compiled with 64-bit indices

*sun\_klu\_numeric* \***SUNLinSol\_KLUGetNumeric**(*SUNLinearSolver* S)

This function returns a pointer to the KLU numeric factorization stored in the SUNLinSol\_KLU content structure.

type **sun\_klu\_numeric**

This type is an alias that depends on the SUNDIALS index size, see *sunindextype* and *SUNDIALS\_INDEX\_SIZE*. It is equivalent to:

- *klu\_numeric* when SUNDIALS is compiled with 32-bit indices
- *klu\_l\_numeric* when SUNDIALS is compiled with 64-bit indices

*sun\_klu\_common* \***SUNLinSol\_KLUGetCommon**(*SUNLinearSolver* S)

This function returns a pointer to the KLU common structure stored in the SUNLinSol\_KLU content structure.

type **sun\_klu\_common**

This type is an alias that depends on the SUNDIALS index size, see *sunindextype* and *SUNDIALS\_INDEX\_SIZE*. It is equivalent to:

- *klu\_common* when SUNDIALS is compiled with 32-bit indices
- *klu\_l\_common* when SUNDIALS is compiled with 64-bit indices

## 8.5.2 SUNLinSol\_KLU Description

The SUNLinSol\_KLU module defines the *content* field of a *SUNLinearSolver* to be the following structure:

```
struct _SUNLinearSolverContent_KLU {
    int      last_flag;
    int      first_factorize;
    sun_klu_symbolic *symbolic;
    sun_klu_numeric *numeric;
    sun_klu_common common;
```

(continues on next page)

(continued from previous page)

```

sunindextype      (*klu_solver)(sun_klu_symbolic*, sun_klu_numeric*,
                                sunindextype, sunindextype,
                                double*, sun_klu_common*);
};

```

These entries of the *content* field contain the following information:

- **last\_flag** - last error return flag from internal function evaluations,
- **first\_factorize** - flag indicating whether the factorization has ever been performed,
- **symbolic** - KLU storage structure for symbolic factorization components, with underlying type `klu_symbolic` or `klu_l_symbolic`, depending on whether SUNDIALS was installed with 32-bit versus 64-bit indices, respectively,
- **numeric** - KLU storage structure for numeric factorization components, with underlying type `klu_numeric` or `klu_l_numeric`, depending on whether SUNDIALS was installed with 32-bit versus 64-bit indices, respectively,
- **common** - storage structure for common KLU solver components, with underlying type `klu_common` or `klu_l_common`, depending on whether SUNDIALS was installed with 32-bit versus 64-bit indices, respectively,
- **klu\_solver** – pointer to the appropriate KLU solver function (depending on whether it is using a CSR or CSC sparse matrix, and on whether SUNDIALS was installed with 32-bit or 64-bit indices).

The `SUNLinSol_KLU` module is a `SUNLinearSolver` wrapper for the KLU sparse matrix factorization and solver library written by Tim Davis and collaborators ([4, 27]). In order to use the `SUNLinSol_KLU` interface to KLU, it is assumed that KLU has been installed on the system prior to installation of SUNDIALS, and that SUNDIALS has been configured appropriately to link with KLU (see §12.4 for details). Additionally, this wrapper only supports double-precision calculations, and therefore cannot be compiled if SUNDIALS is configured to have `sunrealtype` set to either `extended` or `single` (see §4.1 for details). Since the KLU library supports both 32-bit and 64-bit integers, this interface will be compiled for either of the available `sunindextype` options.

The KLU library has a symbolic factorization routine that computes the permutation of the linear system matrix to block triangular form and the permutations that will pre-order the diagonal blocks (the only ones that need to be factored) to reduce fill-in (using AMD, COLAMD, CHOLAMD, natural, or an ordering given by the user). Of these ordering choices, the default value in the `SUNLinSol_KLU` module is the COLAMD ordering.

KLU breaks the factorization into two separate parts. The first is a symbolic factorization and the second is a numeric factorization that returns the factored matrix along with final pivot information. KLU also has a refactor routine that can be called instead of the numeric factorization. This routine will reuse the pivot information. This routine also returns diagnostic information that a user can examine to determine if numerical stability is being lost and a full numerical factorization should be done instead of the refactor.

Since the linear systems that arise within the context of SUNDIALS calculations will typically have identical sparsity patterns, the `SUNLinSol_KLU` module is constructed to perform the following operations:

- The first time that the “setup” routine is called, it performs the symbolic factorization, followed by an initial numerical factorization.
- On subsequent calls to the “setup” routine, it calls the appropriate KLU “refactor” routine, followed by estimates of the numerical conditioning using the relevant “rcond”, and if necessary “condest”, routine(s). If these estimates of the condition number are larger than  $\varepsilon^{-2/3}$  (where  $\varepsilon$  is the double-precision unit roundoff), then a new factorization is performed.
- The module includes the routine `SUNKLUREInit`, that can be called by the user to force a full refactorization at the next “setup” call.
- The “solve” call performs pivoting and forward and backward substitution using the stored KLU data structures. We note that in this solve KLU operates on the native data arrays for the right-hand side and solution vectors, without requiring costly data copies.

The SUNLinSol\_KLU module defines implementations of all “direct” linear solver operations listed in §8.1:

- SUNLinSolGetType\_KLU
- SUNLinSolInitialize\_KLU – this sets the `first_factorize` flag to 1, forcing both symbolic and numerical factorizations on the subsequent “setup” call.
- SUNLinSolSetup\_KLU – this performs either a *LU* factorization or refactorization of the input matrix.
- SUNLinSolSolve\_KLU – this calls the appropriate KLU solve routine to utilize the *LU* factors to solve the linear system.
- SUNLinSolLastFlag\_KLU
- SUNLinSolSpace\_KLU – this only returns information for the storage within the solver *interface*, i.e. storage for the integers `last_flag` and `first_factorize`. For additional space requirements, see the KLU documentation.
- SUNLinSolFree\_KLU

## 8.6 The SUNLinSol\_LapackBand Module

The SUNLinSol\_LapackBand implementation of the SUNLinearSolver class is designed to be used with the corresponding SUNMATRIX\_BAND matrix type, and one of the serial or shared-memory N\_Vector implementations (NVECTOR\_SERIAL, NVECTOR\_OPENMP, or NVECTOR\_PTHREADS). The

### 8.6.1 SUNLinSol\_LapackBand Usage

The header file to be included when using this module is `sunlinsol/sunlinsol_lapackband.h`. The installed module library to link to is `libsundials_sunlinsollapackband.lib` where `.lib` is typically `.so` for shared libraries and `.a` for static libraries.

The module SUNLinSol\_LapackBand provides the following user-callable routine:

*SUNLinearSolver* **SUNLinSol\_LapackBand**(*N\_Vector* y, *SUNMatrix* A, *SUNContext* sunctx)

This function creates and allocates memory for a LAPACK band SUNLinearSolver.

#### Arguments:

- y – vector used to determine the linear system size.
- A – matrix used to assess compatibility.
- sunctx – the *SUNContext* object (see §4.2)

#### Return value:

New SUNLinSol\_LapackBand object, or NULL if either A or y are incompatible.

#### Notes:

This routine will perform consistency checks to ensure that it is called with consistent *N\_Vector* and *SUNMatrix* implementations. These are currently limited to the SUNMATRIX\_BAND matrix type and the NVECTOR\_SERIAL, NVECTOR\_OPENMP, and NVECTOR\_PTHREADS vector types. As additional compatible matrix and vector implementations are added to SUNDIALS, these will be included within this compatibility check.

Additionally, this routine will verify that the input matrix A is allocated with appropriate upper bandwidth storage for the *LU* factorization.

## 8.6.2 SUNLinSol\_LapackBand Description

SUNLinSol\_LapackBand module defines the *content* field of a SUNLinearSolver to be the following structure:

```
struct _SUNLinearSolverContent_Band {  
    sunindextype N;  
    sunindextype *pivots;  
    sunindextype last_flag;  
};
```

These entries of the *content* field contain the following information:

- *N* - size of the linear system,
- *pivots* - index array for partial pivoting in LU factorization,
- *last\_flag* - last error return flag from internal function evaluations.

The SUNLinSol\_LapackBand module is a SUNLinearSolver wrapper for the LAPACK band matrix factorization and solve routines, \*GBTRF and \*GBTRS, where \* is either D or S, depending on whether SUNDIALS was configured to have *sunrealtype* set to double or single, respectively (see §4.1 for details). In order to use the SUNLinSol\_LapackBand module it is assumed that LAPACK has been installed on the system prior to installation of SUNDIALS, and that SUNDIALS has been configured appropriately to link with LAPACK (see §12.4 for details). We note that since there do not exist 128-bit floating-point factorization and solve routines in LAPACK, this interface cannot be compiled when using extended precision for *sunrealtype*. Similarly, since there do not exist 64-bit integer LAPACK routines, the SUNLinSol\_LapackBand module also cannot be compiled when using *int64\_t* for the *sunindextype*.

This solver is constructed to perform the following operations:

- The “setup” call performs an *LU* factorization with partial (row) pivoting,  $PA = LU$ , where *P* is a permutation matrix, *L* is a lower triangular matrix with 1’s on the diagonal, and *U* is an upper triangular matrix. This factorization is stored in-place on the input SUNMATRIX\_BAND object *A*, with pivoting information encoding *P* stored in the *pivots* array.
- The “solve” call performs pivoting and forward and backward substitution using the stored *pivots* array and the *LU* factors held in the SUNMATRIX\_BAND object.
- *A* must be allocated to accommodate the increase in upper bandwidth that occurs during factorization. More precisely, if *A* is a band matrix with upper bandwidth *mu* and lower bandwidth *mL*, then the upper triangular factor *U* can have upper bandwidth as big as  $smu = \text{MIN}(N-1, mu+mL)$ . The lower triangular factor *L* has lower bandwidth *mL*.

The SUNLinSol\_LapackBand module defines band implementations of all “direct” linear solver operations listed in §8.1:

- SUNLinSolGetType\_LapackBand
- SUNLinSolInitialize\_LapackBand – this does nothing, since all consistency checks are performed at solver creation.
- SUNLinSolSetup\_LapackBand – this calls either DGBTRF or SGBTRF to perform the *LU* factorization.
- SUNLinSolSolve\_LapackBand – this calls either DGBTRS or SGBTRS to use the *LU* factors and *pivots* array to perform the solve.
- SUNLinSolLastFlag\_LapackBand
- SUNLinSolSpace\_LapackBand – this only returns information for the storage *within* the solver object, i.e. storage for *N*, *last\_flag*, and *pivots*.
- SUNLinSolFree\_LapackBand



## 8.7 The SUNLinSol\_LapackDense Module

The SUNLinSol\_LapackDense implementation of the SUNLinearSolver class is designed to be used with the corresponding SUNMATRIX\_DENSE matrix type, and one of the serial or shared-memory N\_Vector implementations (NVECTOR\_SERIAL, NVECTOR\_OPENMP, or NVECTOR\_PTHREADS).

### 8.7.1 SUNLinSol\_LapackDense Usage

The header file to be included when using this module is `sunlinsol/sunlinsol_lapackdense.h`. The installed module library to link to is `libsundials_sunlinsollapackdense.lib` where *.lib* is typically *.so* for shared libraries and *.a* for static libraries.

The module SUNLinSol\_LapackDense provides the following additional user-callable constructor routine:

*SUNLinearSolver* **SUNLinSol\_LapackDense**(*N\_Vector* y, *SUNMatrix* A, *SUNContext* sunctx)

This function creates and allocates memory for a LAPACK dense SUNLinearSolver.

#### Arguments:

- *y* – vector used to determine the linear system size.
- *A* – matrix used to assess compatibility.
- *sunctx* – the *SUNContext* object (see §4.2)

#### Return value:

New SUNLinSol\_LapackDense object, or NULL if either *A* or *y* are incompatible.

#### Notes:

This routine will perform consistency checks to ensure that it is called with consistent *N\_Vector* and *SUNMatrix* implementations. These are currently limited to the SUNMATRIX\_DENSE matrix type and the NVECTOR\_SERIAL, NVECTOR\_OPENMP, and NVECTOR\_PTHREADS vector types. As additional compatible matrix and vector implementations are added to SUNDIALS, these will be included within this compatibility check.

### 8.7.2 SUNLinSol\_LapackDense Description

The SUNLinSol\_LapackDense module defines the *content* field of a SUNLinearSolver to be the following structure:

```
struct _SUNLinearSolverContent_Dense {
    sunindextype N;
    sunindextype *pivots;
    sunindextype last_flag;
};
```

These entries of the *content* field contain the following information:

- *N* - size of the linear system,
- *pivots* - index array for partial pivoting in LU factorization,
- *last\_flag* - last error return flag from internal function evaluations.

The SUNLinSol\_LapackDense module is a SUNLinearSolver wrapper for the LAPACK dense matrix factorization and solve routines, \*GETRF and \*GETRS, where \* is either D or S, depending on whether SUNDIALS was configured to have *sunrealtype* set to double or single, respectively (see §4.1 for details). In order to use the SUNLinSol\_LapackDense module it is assumed that LAPACK has been installed on the system prior to installation of SUNDIALS, and that SUNDIALS has been configured appropriately to link with LAPACK (see §12.4 for details). We note that since

there do not exist 128-bit floating-point factorization and solve routines in LAPACK, this interface cannot be compiled when using extended precision for [sunrealtype](#). Similarly, since there do not exist 64-bit integer LAPACK routines, the SUNLinSol\_LapackDense module also cannot be compiled when using `int64_t` for the [sunindextype](#).

This solver is constructed to perform the following operations:

- The “setup” call performs an  $LU$  factorization with partial (row) pivoting ( $\mathcal{O}(N^3)$  cost),  $PA = LU$ , where  $P$  is a permutation matrix,  $L$  is a lower triangular matrix with 1’s on the diagonal, and  $U$  is an upper triangular matrix. This factorization is stored in-place on the input SUNMATRIX\_DENSE object  $A$ , with pivoting information encoding  $P$  stored in the `pivots` array.
- The “solve” call performs pivoting and forward and backward substitution using the stored `pivots` array and the  $LU$  factors held in the SUNMATRIX\_DENSE object ( $\mathcal{O}(N^2)$  cost).

The SUNLinSol\_LapackDense module defines dense implementations of all “direct” linear solver operations listed in §8.1:

- `SUNLinSolGetType_LapackDense`
- `SUNLinSolInitialize_LapackDense` – this does nothing, since all consistency checks are performed at solver creation.
- `SUNLinSolSetup_LapackDense` – this calls either DGETRF or SGETRF to perform the  $LU$  factorization.
- `SUNLinSolSolve_LapackDense` – this calls either DGETRS or SGETRS to use the  $LU$  factors and `pivots` array to perform the solve.
- `SUNLinSolLastFlag_LapackDense`
- `SUNLinSolSpace_LapackDense` – this only returns information for the storage *within* the solver object, i.e. storage for `N`, `last_flag`, and `pivots`.
- `SUNLinSolFree_LapackDense`

## 8.8 The SUNLinSol\_MagmaDense Module

The SUNLinearSolver\_MagmaDense implementation of the SUNLinearSolver class is designed to be used with the SUNMATRIX\_MAGMADENSE matrix, and a GPU-enabled vector. The header file to include when using this module is `sunlinsol/sunlinsol_magmadense.h`. The installed library to link to is `libsundials-sunlinsolmagmadense.lib` where `lib` is typically `.so` for shared libraries and `.a` for static libraries.

**Warning:** The SUNLinearSolver\_MagmaDense module is experimental and subject to change.

### 8.8.1 SUNLinearSolver\_MagmaDense Description

The SUNLinearSolver\_MagmaDense implementation provides an interface to the dense LU and dense batched LU methods in the [MAGMA](#) linear algebra library [62]. The batched LU methods are leveraged when solving block diagonal linear systems of the form

$$\begin{bmatrix} \mathbf{A}_0 & 0 & \cdots & 0 \\ 0 & \mathbf{A}_1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mathbf{A}_{n-1} \end{bmatrix} x_j = b_j.$$

## 8.8.2 SUNLinearSolver\_MagmaDense Functions

The SUNLinearSolver\_MagmaDense module defines implementations of all “direct” linear solver operations listed in §8.1:

- SUNLinSolGetType\_MagmaDense
- SUNLinSolInitialize\_MagmaDense
- SUNLinSolSetup\_MagmaDense
- SUNLinSolSolve\_MagmaDense
- SUNLinSolLastFlag\_MagmaDense
- SUNLinSolFree\_MagmaDense

In addition, the module provides the following user-callable routines:

*SUNLinearSolver* **SUNLinSol\_MagmaDense**(*N\_Vector* y, *SUNMatrix* A, *SUNContext* sunctx)

This constructor function creates and allocates memory for a SUNLinearSolver object.

### Arguments:

- y – a vector for checking compatibility with the solver.
- A – a SUNMATRIX\_MAGMADENSE matrix for checking compatibility with the solver.
- sunctx – the *SUNContext* object (see §4.2)

### Return value:

If successful, a SUNLinearSolver object. If either A or y are incompatible then this routine will return NULL. This routine analyzes the input matrix and vector to determine the linear system size and to assess compatibility with the solver.

*SUNErrCode* **SUNLinSol\_MagmaDense\_SetAsync**(*SUNLinearSolver* LS, *sunbooleantype* onoff)

This function can be used to toggle the linear solver between asynchronous and synchronous modes. In asynchronous mode (default), SUNLinearSolver operations are asynchronous with respect to the host. In synchronous mode, the host and GPU device are synchronized prior to the operation returning.

### Arguments:

- LS – a SUNLinSol\_MagmaDense object
- onoff – 0 for synchronous mode or 1 for asynchronous mode (default 1)

### Return value:

- A *SUNErrCode*

## 8.8.3 SUNLinearSolver\_MagmaDense Content

The SUNLinearSolver\_MagmaDense module defines the object *content* field of a SUNLinearSolver to be the following structure:

```
struct _SUNLinearSolverContent_MagmaDense {
    int last_flag;
    sunbooleantype async;
    sunindextype N;
    SUNMemory pivots;
    SUNMemory pivotsarr;
```

(continues on next page)

(continued from previous page)

```

SUNMemory      dpivotsarr;
SUNMemory      infoarr;
SUNMemory      rhsarr;
SUNMemoryHelper memhelp;
magma_queue_t   q;
};

```

## 8.9 The SUNLinSol\_OneMklDense Module

The `SUNLinearSolver_OneMklDense` implementation of the `SUNLinearSolver` class interfaces to the direct linear solvers from the [Intel oneAPI Math Kernel Library \(oneMKL\)](#) for solving dense systems or block-diagonal systems with dense blocks. This linear solver is best paired with the `SUNMatrix_OneMklDense` matrix.

The header file to include when using this class is `sunlinsol/sunlinsol_onemkldense.h`. The installed library to link to is `libsundials_sunlinsolonemkldense.lib` where `lib` is typically `.so` for shared libraries and `.a` for static libraries.

**Warning:** The `SUNLinearSolver_OneMklDense` class is experimental and subject to change.

### 8.9.1 SUNLinearSolver\_OneMklDense Functions

The `SUNLinearSolver_OneMklDense` class defines implementations of all “direct” linear solver operations listed in §8.1:

- `SUNLinSolGetType_OneMklDense` – returns `SUNLINEARSOLVER_ONEMKLDENSE`
- `SUNLinSolInitialize_OneMklDense`
- `SUNLinSolSetup_OneMklDense`
- `SUNLinSolSolve_OneMklDense`
- `SUNLinSolLastFlag_OneMklDense`
- `SUNLinSolFree_OneMklDense`

In addition, the class provides the following user-callable routines:

*SUNLinearSolver* **`SUNLinSol_OneMklDense`**(*N\_Vector* `y`, *SUNMatrix* `A`, *SUNContext* `sunctx`)

This constructor function creates and allocates memory for a `SUNLinearSolver` object.

**Arguments:**

- `y` – a vector for checking compatibility with the solver.
- `A` – a `SUNMatrix_OneMklDense` matrix for checking compatibility with the solver.
- `sunctx` – the *SUNContext* object (see §4.2)

**Return value:**

If successful, a `SUNLinearSolver` object. If either `A` or `y` are incompatible then this routine will return `NULL`. This routine analyzes the input matrix and vector to determine the linear system size and to assess compatibility with the solver.

### 8.9.2 SUNLinearSolver\_OneMklDense Usage Notes

**Warning:** The SUNLinearSolver\_OneMklDense class only supports 64-bit indexing, thus SUNDIALS must be built for 64-bit indexing to use this class.

When using the SUNLinearSolver\_OneMklDense class with a SUNDIALS package (e.g. CVODE), the queue given to the matrix is also used for the linear solver.

## 8.10 The SUNLinSol\_PCG Module

The SUNLinSol\_PCG implementation of the SUNLinearSolver class performs the PCG (Preconditioned Conjugate Gradient [38]) method; this is an iterative linear solver that is designed to be compatible with any `N_Vector` implementation that supports a minimal subset of operations (`N_VClone()`, `N_VDotProd()`, `N_VScale()`, `N_VLinearSum()`, `N_VProd()`, and `N_VDestroy()`). Unlike the SPGMR and SPFGMR algorithms, PCG requires a fixed amount of memory that does not increase with the number of allowed iterations.

Unlike all of the other iterative linear solvers supplied with SUNDIALS, PCG should only be used on *symmetric* linear systems (e.g. mass matrix linear systems encountered in ARKODE). As a result, the explanation of the role of scaling and preconditioning matrices given in general must be modified in this scenario. The PCG algorithm solves a linear system  $Ax = b$  where  $A$  is a symmetric ( $A^T = A$ ), real-valued matrix. Preconditioning is allowed, and is applied in a symmetric fashion on both the right and left. Scaling is also allowed and is applied symmetrically. We denote the preconditioner and scaling matrices as follows:

- $P$  is the preconditioner (assumed symmetric),
- $S$  is a diagonal matrix of scale factors.

The matrices  $A$  and  $P$  are not required explicitly; only routines that provide  $A$  and  $P^{-1}$  as operators are required. The diagonal of the matrix  $S$  is held in a single `N_Vector`, supplied by the user.

In this notation, PCG applies the underlying CG algorithm to the equivalent transformed system

$$\tilde{A}\tilde{x} = \tilde{b} \tag{8.4}$$

where

$$\begin{aligned} \tilde{A} &= SP^{-1}AP^{-1}S, \\ \tilde{b} &= SP^{-1}b, \\ \tilde{x} &= S^{-1}Px. \end{aligned} \tag{8.5}$$

The scaling matrix must be chosen so that the vectors  $SP^{-1}b$  and  $S^{-1}Px$  have dimensionless components.

The stopping test for the PCG iterations is on the L2 norm of the scaled preconditioned residual:

$$\begin{aligned} &\|\tilde{b} - \tilde{A}\tilde{x}\|_2 < \delta \\ \Leftrightarrow & \|SP^{-1}b - SP^{-1}Ax\|_2 < \delta \\ \Leftrightarrow & \|P^{-1}b - P^{-1}Ax\|_S < \delta \end{aligned}$$

where  $\|v\|_S = \sqrt{v^T S^T S v}$ , with an input tolerance  $\delta$ .

### 8.10.1 SUNLinSol\_PCG Usage

The header file to be included when using this module is `sunlinsol/sunlinsol_pcg.h`. The `SUNLinSol_PCG` module is accessible from all SUNDIALS solvers *without* linking to the `libsundials_sunlinsolpcg` module library.

The module `SUNLinSol_PCG` provides the following user-callable routines:

*SUNLinearSolver* **SUNLinSol\_PCG**(*N\_Vector* y, int pretype, int maxl, *SUNContext* suncx)

This constructor function creates and allocates memory for a PCG `SUNLinearSolver`.

**Arguments:**

- y – a template vector.
- pretype – a flag indicating the type of preconditioning to use:
  - SUN\_PREC\_NONE
  - SUN\_PREC\_LEFT
  - SUN\_PREC\_RIGHT
  - SUN\_PREC\_BOTH
- maxl – the maximum number of linear iterations to allow.
- suncx – the *SUNContext* object (see §4.2)

**Return value:**

If successful, a `SUNLinearSolver` object. If either y is incompatible then this routine will return NULL.

**Notes:**

This routine will perform consistency checks to ensure that it is called with a consistent `N_Vector` implementation (i.e. that it supplies the requisite vector operations).

A `maxl` argument that is  $\leq 0$  will result in the default value (5).

Since the PCG algorithm is designed to only support symmetric preconditioning, then any of the `pretype` inputs `SUN_PREC_LEFT`, `SUN_PREC_RIGHT`, or `SUN_PREC_BOTH` will result in use of the symmetric preconditioner; any other integer input will result in the default (no preconditioning). Although some SUNDIALS solvers are designed to only work with left preconditioning (IDA and IDAS) and others with only right preconditioning (KINSOL), PCG should *only* be used with these packages when the linear systems are known to be *symmetric*. Since the scaling of matrix rows and columns must be identical in a symmetric matrix, symmetric preconditioning should work appropriately even for packages designed with one-sided preconditioning in mind.

*SUNErrCode* **SUNLinSol\_PCGSetPrecType**(*SUNLinearSolver* S, int pretype)

This function updates the flag indicating use of preconditioning.

**Arguments:**

- S – `SUNLinSol_PCG` object to update.
- pretype – a flag indicating the type of preconditioning to use:
  - SUN\_PREC\_NONE
  - SUN\_PREC\_LEFT
  - SUN\_PREC\_RIGHT
  - SUN\_PREC\_BOTH

**Return value:**

- A *SUNErrCode*

**Notes:**

As above, any one of the input values, SUN\_PREC\_LEFT, SUN\_PREC\_RIGHT, or SUN\_PREC\_BOTH will enable preconditioning; SUN\_PREC\_NONE disables preconditioning.

*SUNErrCode* **SUNLinSol\_PCGSetMaxl** (*SUNLinearSolver* S, int maxl)

This function updates the number of linear solver iterations to allow.

**Arguments:**

- *S* – SUNLinSol\_PCG object to update.
- *maxl* – maximum number of linear iterations to allow. Any non-positive input will result in the default value (5).

**Return value:**

- A *SUNErrCode*

## 8.10.2 SUNLinSol\_PCG Description

The SUNLinSol\_PCG module defines the *content* field of a SUNLinearSolver to be the following structure:

```
struct _SUNLinearSolverContent_PCG {
    int maxl;
    int pretype;
    sunboolean_t zeroguess;
    int numiters;
    sunreal_t resnorm;
    int last_flag;
    SUNATimesFn ATimes;
    void* ATData;
    SUNPSetupFn Psetup;
    SUNPSolveFn Psolve;
    void* PData;
    N_Vector s;
    N_Vector r;
    N_Vector p;
    N_Vector z;
    N_Vector Ap;
};
```

These entries of the *content* field contain the following information:

- *maxl* - number of PCG iterations to allow (default is 5),
- *pretype* - flag for use of preconditioning (default is none),
- *numiters* - number of iterations from the most-recent solve,
- *resnorm* - final linear residual norm from the most-recent solve,
- *last\_flag* - last error return flag from an internal function,
- *ATimes* - function pointer to perform  $Av$  product,
- *ATData* - pointer to structure for *ATimes*,
- *Psetup* - function pointer to preconditioner setup routine,
- *Psolve* - function pointer to preconditioner solve routine,

- PData - pointer to structure for Psetup and Psolve,
- s - vector pointer for supplied scaling matrix (default is NULL),
- r - a N\_Vector which holds the preconditioned linear system residual,
- p, z, Ap - N\_Vector used for workspace by the PCG algorithm.

This solver is constructed to perform the following operations:

- During construction all N\_Vector solver data is allocated, with vectors cloned from a template N\_Vector that is input, and default solver parameters are set.
- User-facing “set” routines may be called to modify default solver parameters.
- Additional “set” routines are called by the SUNDIALS solver that interfaces with SUNLinSol\_PCG to supply the ATimes, PSetup, and Psolve function pointers and s scaling vector.
- In the “initialize” call, the solver parameters are checked for validity.
- In the “setup” call, any non-NULL PSetup function is called. Typically, this is provided by the SUNDIALS solver itself, that translates between the generic PSetup function and the solver-specific routine (solver-supplied or user-supplied).
- In the “solve” call the PCG iteration is performed. This will include scaling and preconditioning if those options have been supplied.

The SUNLinSol\_PCG module defines implementations of all “iterative” linear solver operations listed in §8.1:

- SUNLinSolGetType\_PCG
- SUNLinSolInitialize\_PCG
- SUNLinSolSetATimes\_PCG
- SUNLinSolSetPreconditioner\_PCG
- SUNLinSolSetScalingVectors\_PCG – since PCG only supports symmetric scaling, the second N\_Vector argument to this function is ignored.
- SUNLinSolSetZeroGuess\_PCG – note the solver assumes a non-zero guess by default and the zero guess flag is reset to SUNFALSE after each call to SUNLinSolSolve\_PCG.
- SUNLinSolSetup\_PCG
- SUNLinSolSolve\_PCG
- SUNLinSolNumIters\_PCG
- SUNLinSolResNorm\_PCG
- SUNLinSolResid\_PCG
- SUNLinSolLastFlag\_PCG
- SUNLinSolSpace\_PCG
- SUNLinSolFree\_PCG



## 8.11 The SUNLinSol\_SPBCGS Module

The SUNLinSol\_SPBCGS implementation of the SUNLinearSolver class performs a Scaled, Preconditioned, Bi-Conjugate Gradient, Stabilized [65] method; this is an iterative linear solver that is designed to be compatible with any N\_Vector implementation that supports a minimal subset of operations ([N\\_VClone\(\)](#), [N\\_VDotProd\(\)](#), [N\\_VScale\(\)](#), [N\\_VLinearSum\(\)](#), [N\\_VProd\(\)](#), [N\\_VDiv\(\)](#), and [N\\_VDestroy\(\)](#)). Unlike the SPGMR and SPFGMR algorithms, SPBCGS requires a fixed amount of memory that does not increase with the number of allowed iterations.

### 8.11.1 SUNLinSol\_SPBCGS Usage

The header file to be included when using this module is `sunlinsol/sunlinsol_spbcgs.h`. The SUNLinSol\_SPBCGS module is accessible from all SUNDIALS solvers *without* linking to the `libsundials_sunlinsolspbcgs` module library.

The module SUNLinSol\_SPBCGS provides the following user-callable routines:

*SUNLinearSolver* **SUNLinSol\_SPBCGS**(*N\_Vector* y, int pretype, int maxl, *SUNContext* sunctx)

This constructor function creates and allocates memory for a SPBCGS SUNLinearSolver.

**Arguments:**

- y – a template vector.
- pretype – a flag indicating the type of preconditioning to use:
  - SUN\_PREC\_NONE
  - SUN\_PREC\_LEFT
  - SUN\_PREC\_RIGHT
  - SUN\_PREC\_BOTH
- maxl – the maximum number of linear iterations to allow.
- sunctx – the *SUNContext* object (see §4.2)

**Return value:**

If successful, a SUNLinearSolver object. If either y is incompatible then this routine will return NULL.

**Notes:**

This routine will perform consistency checks to ensure that it is called with a consistent N\_Vector implementation (i.e. that it supplies the requisite vector operations).

A maxl argument that is  $\leq 0$  will result in the default value (5).

Some SUNDIALS solvers are designed to only work with left preconditioning (IDA and IDAS) and others with only right preconditioning (KINSOL). While it is possible to configure a SUNLinSol\_SPBCGS object to use any of the preconditioning options with these solvers, this use mode is not supported and may result in inferior performance.

---

**Note:** With SUN\_PREC\_RIGHT or SUN\_PREC\_BOTH the initial guess must be zero (use [SUNLinSolSetZeroGuess\(\)](#) to indicate the initial guess is zero).

---

*SUNErrCode* **SUNLinSol\_SPBCGSSetPrecType**(*SUNLinearSolver* S, int pretype)

This function updates the flag indicating use of preconditioning.

**Arguments:**

- *S* – SUNLinSol\_SPBCGS object to update.
- *pretype* – a flag indicating the type of preconditioning to use:
  - SUN\_PREC\_NONE
  - SUN\_PREC\_LEFT
  - SUN\_PREC\_RIGHT
  - SUN\_PREC\_BOTH

**Return value:**

- A *SUNErrCode*

*SUNErrCode* **SUNLinSol\_SPBCGSsetMaxl**(*SUNLinearSolver* S, int maxl)

This function updates the number of linear solver iterations to allow.

**Arguments:**

- *S* – SUNLinSol\_SPBCGS object to update.
- *maxl* – maximum number of linear iterations to allow. Any non-positive input will result in the default value (5).

**Return value:**

- A *SUNErrCode*

### 8.11.2 SUNLinSol\_SPBCGS Description

The SUNLinSol\_SPBCGS module defines the *content* field of a *SUNLinearSolver* to be the following structure:

```
struct _SUNLinearSolverContent_SPBCGS {  
    int maxl;  
    int pretype;  
    sunboolean_t zeroguess;  
    int numiters;  
    sunreal_t resnorm;  
    int last_flag;  
    SUNATimesFn ATimes;  
    void* ATData;  
    SUNPSetupFn Psetup;  
    SUNPSolveFn Psolve;  
    void* PData;  
    N_Vector s1;  
    N_Vector s2;  
    N_Vector r;  
    N_Vector r_star;  
    N_Vector p;  
    N_Vector q;  
    N_Vector u;  
    N_Vector Ap;  
    N_Vector vtemp;  
};
```

These entries of the *content* field contain the following information:

- *maxl* - number of SPBCGS iterations to allow (default is 5),

- `pretype` - flag for type of preconditioning to employ (default is none),
- `numiters` - number of iterations from the most-recent solve,
- `resnorm` - final linear residual norm from the most-recent solve,
- `last_flag` - last error return flag from an internal function,
- `ATimes` - function pointer to perform  $Av$  product,
- `ATData` - pointer to structure for `ATimes`,
- `Psetup` - function pointer to preconditioner setup routine,
- `Psolve` - function pointer to preconditioner solve routine,
- `PData` - pointer to structure for `Psetup` and `Psolve`,
- `s1`, `s2` - vector pointers for supplied scaling matrices (default is NULL),
- `r` - a `N_Vector` which holds the current scaled, preconditioned linear system residual,
- `r_star` - a `N_Vector` which holds the initial scaled, preconditioned linear system residual,
- `p`, `q`, `u`, `Ap`, `vtemp` - `N_Vector` used for workspace by the SPBCGS algorithm.

This solver is constructed to perform the following operations:

- During construction all `N_Vector` solver data is allocated, with vectors cloned from a template `N_Vector` that is input, and default solver parameters are set.
- User-facing “set” routines may be called to modify default solver parameters.
- Additional “set” routines are called by the SUNDIALS solver that interfaces with `SUNLinSol_SPBCGS` to supply the `ATimes`, `PSetup`, and `Psolve` function pointers and `s1` and `s2` scaling vectors.
- In the “initialize” call, the solver parameters are checked for validity.
- In the “setup” call, any non-NULL `PSetup` function is called. Typically, this is provided by the SUNDIALS solver itself, that translates between the generic `PSetup` function and the solver-specific routine (solver-supplied or user-supplied).
- In the “solve” call the SPBCGS iteration is performed. This will include scaling and preconditioning if those options have been supplied.

The `SUNLinSol_SPBCGS` module defines implementations of all “iterative” linear solver operations listed in §8.1:

- `SUNLinSolGetType_SPBCGS`
- `SUNLinSolInitialize_SPBCGS`
- `SUNLinSolSetATimes_SPBCGS`
- `SUNLinSolSetPreconditioner_SPBCGS`
- `SUNLinSolSetScalingVectors_SPBCGS`
- `SUNLinSolSetZeroGuess_SPBCGS` – note the solver assumes a non-zero guess by default and the zero guess flag is reset to `SUNFALSE` after each call to `SUNLinSolSolve_SPBCGS`.
- `SUNLinSolSetup_SPBCGS`
- `SUNLinSolSolve_SPBCGS`
- `SUNLinSolNumIters_SPBCGS`
- `SUNLinSolResNorm_SPBCGS`
- `SUNLinSolResid_SPBCGS`

- `SUNLinSolLastFlag_SPBCGS`
- `SUNLinSolSpace_SPBCGS`
- `SUNLinSolFree_SPBCGS`

## 8.12 The SUNLinSol\_SPFGMR Module

The `SUNLinSol_SPFGMR` implementation of the `SUNLinearSolver` class performs a Scaled, Preconditioned, Flexible, Generalized Minimum Residual [59] method; this is an iterative linear solver that is designed to be compatible with any `N_Vector` implementation that supports a minimal subset of operations (`N_VClone()`, `N_VDotProd()`, `N_VScale()`, `N_VLinearSum()`, `N_VProd()`, `N_VConst()`, `N_VDiv()`, and `N_VDestroy()`). Unlike the other Krylov iterative linear solvers supplied with SUNDIALS, FGMRES is specifically designed to work with a changing preconditioner (e.g. from an iterative method).

### 8.12.1 SUNLinSol\_SPFGMR Usage

The header file to be included when using this module is `sunlinsol/sunlinsol_spfgmr.h`. The `SUNLinSol_SPFGMR` module is accessible from all SUNDIALS solvers *without* linking to the `libsundials_sunlinsolspfgmr` module library.

The module `SUNLinSol_SPFGMR` provides the following user-callable routines:

*SUNLinearSolver* **`SUNLinSol_SPFGMR`**(*N\_Vector* y, int pretype, int maxl, *SUNContext* sunctx)

This constructor function creates and allocates memory for a SPFGMR `SUNLinearSolver`.

**Arguments:**

- y – a template vector.
- pretype – a flag indicating the type of preconditioning to use:
  - `SUN_PREC_NONE`
  - `SUN_PREC_LEFT`
  - `SUN_PREC_RIGHT`
  - `SUN_PREC_BOTH`
- maxl – the number of Krylov basis vectors to use.
- sunctx – the *SUNContext* object (see §4.2)

**Return value:**

If successful, a `SUNLinearSolver` object. If either y is incompatible then this routine will return `NULL`.

**Notes:**

This routine will perform consistency checks to ensure that it is called with a consistent `N_Vector` implementation (i.e. that it supplies the requisite vector operations).

A `maxl` argument that is  $\leq 0$  will result in the default value (5).

Since the FGMRES algorithm is designed to only support right preconditioning, then any of the `pretype` inputs `SUN_PREC_LEFT`, `SUN_PREC_RIGHT`, or `SUN_PREC_BOTH` will result in use of `SUN_PREC_RIGHT`; any other integer input will result in the default (no preconditioning). We note that some SUNDIALS solvers are designed to only work with left preconditioning (IDA and IDAS). While it is possible to use a right-preconditioned `SUNLinSol_SPFGMR` object for these packages, this use mode is not supported and may result in inferior performance.

*SUNErrCode* **SUNLinSol\_SPFGMRSetPrecType**(*SUNLinearSolver* S, int pretype)

This function updates the flag indicating use of preconditioning.

**Arguments:**

- *S* – SUNLinSol\_SPFGMR object to update.
- *pretype* – a flag indicating the type of preconditioning to use:
  - SUN\_PREC\_NONE
  - SUN\_PREC\_LEFT
  - SUN\_PREC\_RIGHT
  - SUN\_PREC\_BOTH

**Return value:**

- A *SUNErrCode*

**Notes:**

Since the FGMRES algorithm is designed to only support right preconditioning, then any of the *pretype* inputs SUN\_PREC\_LEFT, SUN\_PREC\_RIGHT, or SUN\_PREC\_BOTH will result in use of SUN\_PREC\_RIGHT; any other integer input will result in the default (no preconditioning).

*SUNErrCode* **SUNLinSol\_SPFGMRSetGSType**(*SUNLinearSolver* S, int gstype)

This function sets the type of Gram-Schmidt orthogonalization to use.

**Arguments:**

- *S* – SUNLinSol\_SPFGMR object to update.
- *gstype* – a flag indicating the type of orthogonalization to use:
  - SUN\_MODIFIED\_GS
  - SUN\_CLASSICAL\_GS

**Return value:**

- A *SUNErrCode*

*SUNErrCode* **SUNLinSol\_SPFGMRSetMaxRestarts**(*SUNLinearSolver* S, int maxrs)

This function sets the number of FGMRES restarts to allow.

**Arguments:**

- *S* – SUNLinSol\_SPFGMR object to update.
- *maxrs* – maximum number of restarts to allow. A negative input will result in the default of 0.

**Return value:**

- A *SUNErrCode*

### 8.12.2 SUNLinSol\_SPFGMR Description

The SUNLinSol\_SPFGMR module defines the *content* field of a SUNLinearSolver to be the following structure:

```
struct _SUNLinearSolverContent_SPFGMR {  
    int maxl;  
    int pretype;  
    int gstype;  
    int max_restarts;  
    sunbooleantype zeroguess;  
    int numiters;  
    sunrealtype resnorm;  
    int last_flag;  
    SUNATimesFn ATimes;  
    void* ATData;  
    SUNPSetupFn Psetup;  
    SUNPSolveFn Psolve;  
    void* PData;  
    N_Vector s1;  
    N_Vector s2;  
    N_Vector *V;  
    N_Vector *Z;  
    sunrealtype **Hes;  
    sunrealtype *givens;  
    N_Vector xcor;  
    sunrealtype *yg;  
    N_Vector vtemp;  
};
```

These entries of the *content* field contain the following information:

- `maxl` - number of FGMRES basis vectors to use (default is 5),
- `pretype` - flag for use of preconditioning (default is none),
- `gstype` - flag for type of Gram-Schmidt orthogonalization (default is modified Gram-Schmidt),
- `max_restarts` - number of FGMRES restarts to allow (default is 0),
- `numiters` - number of iterations from the most-recent solve,
- `resnorm` - final linear residual norm from the most-recent solve,
- `last_flag` - last error return flag from an internal function,
- `ATimes` - function pointer to perform  $Av$  product,
- `ATData` - pointer to structure for `ATimes`,
- `Psetup` - function pointer to preconditioner setup routine,
- `Psolve` - function pointer to preconditioner solve routine,
- `PData` - pointer to structure for `Psetup` and `Psolve`,
- `s1`, `s2` - vector pointers for supplied scaling matrices (default is NULL),
- `V` - the array of Krylov basis vectors  $v_1, \dots, v_{\text{maxl}+1}$ , stored in `V[0]`, `...`, `V[maxl]`. Each  $v_i$  is a vector of type `N_Vector`,
- `Z` - the array of preconditioned Krylov basis vectors  $z_1, \dots, z_{\text{maxl}+1}$ , stored in `Z[0]`, `...`, `Z[maxl]`. Each  $z_i$  is a vector of type `N_Vector`,

- **Hes** - the  $(\text{maxl} + 1) \times \text{maxl}$  Hessenberg matrix. It is stored row-wise so that the  $(i,j)$ th element is given by `Hes[i][j]`,
- **givens** - a length  $2 \text{maxl}$  array which represents the Givens rotation matrices that arise in the FGMRES algorithm. These matrices are  $F_0, F_1, \dots, F_j$ , where

$$F_i = \begin{bmatrix} 1 & & & & & & \\ & \ddots & & & & & \\ & & 1 & & & & \\ & & & c_i & -s_i & & \\ & & & s_i & c_i & & \\ & & & & & 1 & \\ & & & & & & \ddots & \\ & & & & & & & 1 \end{bmatrix},$$

are represented in the **givens** vector as `givens[0] = c0`, `givens[1] = s0`, `givens[2] = c1`, `givens[3] = s1`, ..., `givens[2j] = cj`, `givens[2j+1] = sj`,

- **xcor** - a vector which holds the scaled, preconditioned correction to the initial guess,
- **yg** - a length  $(\text{maxl} + 1)$  array of `sunrealtype` values used to hold “short” vectors (e.g.  $y$  and  $g$ ),
- **vtemp** - temporary vector storage.

This solver is constructed to perform the following operations:

- During construction, the **xcor** and **vtemp** arrays are cloned from a template `N_Vector` that is input, and default solver parameters are set.
- User-facing “set” routines may be called to modify default solver parameters.
- Additional “set” routines are called by the SUNDIALS solver that interfaces with `SUNLinSol_SPFGMR` to supply the `ATimes`, `PSetup`, and `Psolve` function pointers and `s1` and `s2` scaling vectors.
- In the “initialize” call, the remaining solver data is allocated (**V**, **Hes**, **givens**, and **yg**)
- In the “setup” call, any non-NULL `PSetup` function is called. Typically, this is provided by the SUNDIALS solver itself, that translates between the generic `PSetup` function and the solver-specific routine (solver-supplied or user-supplied).
- In the “solve” call, the FGMRES iteration is performed. This will include scaling, preconditioning, and restarts if those options have been supplied.

The `SUNLinSol_SPFGMR` module defines implementations of all “iterative” linear solver operations listed in §8.1:

- `SUNLinSolGetType_SPFGMR`
- `SUNLinSolInitialize_SPFGMR`
- `SUNLinSolSetATimes_SPFGMR`
- `SUNLinSolSetPreconditioner_SPFGMR`
- `SUNLinSolSetScalingVectors_SPFGMR`
- `SUNLinSolSetZeroGuess_SPFGMR` – note the solver assumes a non-zero guess by default and the zero guess flag is reset to `SUNFALSE` after each call to `SUNLinSolSolve_SPFGMR`.
- `SUNLinSolSetup_SPFGMR`
- `SUNLinSolSolve_SPFGMR`
- `SUNLinSolNumIters_SPFGMR`

- `SUNLinSolResNorm_SPFGMR`
- `SUNLinSolResid_SPFGMR`
- `SUNLinSolLastFlag_SPFGMR`
- `SUNLinSolSpace_SPFGMR`
- `SUNLinSolFree_SPFGMR`

## 8.13 The SUNLinSol\_SPGMR Module

The `SUNLinSol_SPGMR` implementation of the `SUNLinearSolver` class performs a Scaled, Preconditioned, Generalized Minimum Residual [60] method; this is an iterative linear solver that is designed to be compatible with any `N_Vector` implementation that supports a minimal subset of operations (`N_VClone()`, `N_VDotProd()`, `N_VScale()`, `N_VLinearSum()`, `N_VProd()`, `N_VConst()`, `N_VDiv()`, and `N_VDestroy()`).

### 8.13.1 SUNLinSol\_SPGMR Usage

The header file to be included when using this module is `sunlinsol/sunlinsol_spgmr.h`. The `SUNLinSol_SPGMR` module is accessible from all SUNDIALS solvers *without* linking to the `libsundials_sunlinsolspgmr` module library.

The module `SUNLinSol_SPGMR` provides the following user-callable routines:

*SUNLinearSolver* **SUNLinSol\_SPGMR**(*N\_Vector* y, int pretype, int maxl, *SUNContext* sunctx)

This constructor function creates and allocates memory for a SPGMR `SUNLinearSolver`.

**Arguments:**

- *y* – a template vector.
- *pretype* – a flag indicating the type of preconditioning to use:
  - `SUN_PREC_NONE`
  - `SUN_PREC_LEFT`
  - `SUN_PREC_RIGHT`
  - `SUN_PREC_BOTH`
- *maxl* – the number of Krylov basis vectors to use.

**Return value:**

If successful, a `SUNLinearSolver` object. If either *y* is incompatible then this routine will return `NULL`.

**Notes:**

This routine will perform consistency checks to ensure that it is called with a consistent `N_Vector` implementation (i.e. that it supplies the requisite vector operations).

A *maxl* argument that is  $\leq 0$  will result in the default value (5).

Some SUNDIALS solvers are designed to only work with left preconditioning (IDA and IDAS) and others with only right preconditioning (KINSOL). While it is possible to configure a `SUNLinSol_SPGMR` object to use any of the preconditioning options with these solvers, this use mode is not supported and may result in inferior performance.



*SUNErrCode* **SUNLinSol\_SPGMRSetPrecType**(*SUNLinearSolver* S, int pretype)

This function updates the flag indicating use of preconditioning.

**Arguments:**

- *S* – SUNLinSol\_SPGMR object to update.
- *pretype* – a flag indicating the type of preconditioning to use:
  - SUN\_PREC\_NONE
  - SUN\_PREC\_LEFT
  - SUN\_PREC\_RIGHT
  - SUN\_PREC\_BOTH

**Return value:**

- A *SUNErrCode*

*SUNErrCode* **SUNLinSol\_SPGMRSetGSType**(*SUNLinearSolver* S, int gstype)

This function sets the type of Gram-Schmidt orthogonalization to use.

**Arguments:**

- *S* – SUNLinSol\_SPGMR object to update.
- *gstype* – a flag indicating the type of orthogonalization to use:
  - SUN\_MODIFIED\_GS
  - SUN\_CLASSICAL\_GS

**Return value:**

- A *SUNErrCode*

*SUNErrCode* **SUNLinSol\_SPGMRSetMaxRestarts**(*SUNLinearSolver* S, int maxrs)

This function sets the number of GMRES restarts to allow.

**Arguments:**

- *S* – SUNLinSol\_SPGMR object to update.
- *maxrs* – maximum number of restarts to allow. A negative input will result in the default of 0.

**Return value:**

- A *SUNErrCode*

### 8.13.2 SUNLinSol\_SPGMR Description

The SUNLinSol\_SPGMR module defines the *content* field of a *SUNLinearSolver* to be the following structure:

```
struct _SUNLinearSolverContent_SPGMR {
    int maxl;
    int pretype;
    int gstype;
    int max_restarts;
    sunboolean_t zero_guess;
    int numiters;
    sunreal_t resnorm;
}
```

(continues on next page)

(continued from previous page)

```

int last_flag;
SUNATimesFn ATimes;
void* ATData;
SUNPSetupFn Psetup;
SUNPSolveFn Psolve;
void* PData;
N_Vector s1;
N_Vector s2;
N_Vector *V;
sunrealtype **Hes;
sunrealtype *givens;
N_Vector xcor;
sunrealtype *yg;
N_Vector vtemp;
};

```

These entries of the *content* field contain the following information:

- `maxl` - number of GMRES basis vectors to use (default is 5),
- `pretype` - flag for type of preconditioning to employ (default is none),
- `gstype` - flag for type of Gram-Schmidt orthogonalization (default is modified Gram-Schmidt),
- `max_restarts` - number of GMRES restarts to allow (default is 0),
- `numiters` - number of iterations from the most-recent solve,
- `resnorm` - final linear residual norm from the most-recent solve,
- `last_flag` - last error return flag from an internal function,
- `ATimes` - function pointer to perform  $Av$  product,
- `ATData` - pointer to structure for `ATimes`,
- `Psetup` - function pointer to preconditioner setup routine,
- `Psolve` - function pointer to preconditioner solve routine,
- `PData` - pointer to structure for `Psetup` and `Psolve`,
- `s1`, `s2` - vector pointers for supplied scaling matrices (default is NULL),
- `V` - the array of Krylov basis vectors  $v_1, \dots, v_{\text{maxl}+1}$ , stored in `V[0]`,  $\dots$  `V[maxl]`. Each  $v_i$  is a vector of type `N_Vector`,
- `Hes` - the  $(\text{maxl} + 1) \times \text{maxl}$  Hessenberg matrix. It is stored row-wise so that the  $(i,j)$ th element is given by `Hes[i][j]`,
- `givens` - a length  $2 \text{maxl}$  array which represents the Givens rotation matrices that arise in the GMRES algorithm.

These matrices are  $F_0, F_1, \dots, F_j$ , where

$$F_i = \begin{bmatrix} 1 & & & & & & \\ & \ddots & & & & & \\ & & 1 & & & & \\ & & & c_i & -s_i & & \\ & & & s_i & c_i & & \\ & & & & & 1 & \\ & & & & & & \ddots & \\ & & & & & & & 1 \end{bmatrix},$$

are represented in the `givens` vector as `givens[0] = c0`, `givens[1] = s0`, `givens[2] = c1`, `givens[3] = s1`, ..., `givens[2j] = cj`, `givens[2j+1] = sj`,

- `xcor` - a vector which holds the scaled, preconditioned correction to the initial guess,
- `yg` - a length  $(\text{maxl} + 1)$  array of `sunrealtype` values used to hold “short” vectors (e.g.  $y$  and  $g$ ),
- `vtemp` - temporary vector storage.

This solver is constructed to perform the following operations:

- During construction, the `xcor` and `vtemp` arrays are cloned from a template `N_Vector` that is input, and default solver parameters are set.
- User-facing “set” routines may be called to modify default solver parameters.
- Additional “set” routines are called by the SUNDIALS solver that interfaces with `SUNLinSol_SPGMR` to supply the `ATimes`, `PSetup`, and `Psolve` function pointers and `s1` and `s2` scaling vectors.
- In the “initialize” call, the remaining solver data is allocated (`V`, `Hes`, `givens`, and `yg` )
- In the “setup” call, any non-NULL `PSetup` function is called. Typically, this is provided by the SUNDIALS solver itself, that translates between the generic `PSetup` function and the solver-specific routine (solver-supplied or user-supplied).
- In the “solve” call, the GMRES iteration is performed. This will include scaling, preconditioning, and restarts if those options have been supplied.

The `SUNLinSol_SPGMR` module defines implementations of all “iterative” linear solver operations listed in §8.1:

- `SUNLinSolGetType_SPGMR`
- `SUNLinSolInitialize_SPGMR`
- `SUNLinSolSetATimes_SPGMR`
- `SUNLinSolSetPreconditioner_SPGMR`
- `SUNLinSolSetScalingVectors_SPGMR`
- `SUNLinSolSetZeroGuess_SPGMR` – note the solver assumes a non-zero guess by default and the zero guess flag is reset to `SUNFALSE` after each call to `SUNLinSolSolve_SPGMR`.
- `SUNLinSolSetup_SPGMR`
- `SUNLinSolSolve_SPGMR`
- `SUNLinSolNumIters_SPGMR`
- `SUNLinSolResNorm_SPGMR`
- `SUNLinSolResid_SPGMR`
- `SUNLinSolLastFlag_SPGMR`

- SUNLinSolSpace\_SPGMR
- SUNLinSolFree\_SPGMR

## 8.14 The SUNLinSol\_SPTFQMR Module

The SUNLinSol\_SPTFQMR implementation of the SUNLinearSolver class performs a Scaled, Preconditioned, Transpose-Free Quasi-Minimum Residual [35] method; this is an iterative linear solver that is designed to be compatible with any N\_Vector implementation that supports a minimal subset of operations (*N\_VClone()*, *N\_VDotProd()*, *N\_VScale()*, *N\_VLinearSum()*, *N\_VProd()*, *N\_VConst()*, *N\_VDiv()*, and *N\_VDestroy()*). Unlike the SPGMR and SPFGMR algorithms, SPTFQMR requires a fixed amount of memory that does not increase with the number of allowed iterations.

### 8.14.1 SUNLinSol\_SPTFQMR Usage

The header file to be included when using this module is `sunlinsol/sunlinsol_sptfqmr.h`. The SUNLinSol\_SPTFQMR module is accessible from all SUNDIALS solvers *without* linking to the `libsundials_sunlinsolsptfqmr` module library.

The module SUNLinSol\_SPTFQMR provides the following user-callable routines:

*SUNLinearSolver* **SUNLinSol\_SPTFQMR**(*N\_Vector* y, int pretype, int maxl, *SUNContext* sunctx)

This constructor function creates and allocates memory for a SPTFQMR SUNLinearSolver.

**Arguments:**

- y – a template vector.
- pretype – a flag indicating the type of preconditioning to use:
  - SUN\_PREC\_NONE
  - SUN\_PREC\_LEFT
  - SUN\_PREC\_RIGHT
  - SUN\_PREC\_BOTH
- maxl – the number of Krylov basis vectors to use.
- sunctx – the *SUNContext* object (see §4.2)

**Return value:**

If successful, a SUNLinearSolver object. If either y is incompatible then this routine will return NULL.

**Notes:**

This routine will perform consistency checks to ensure that it is called with a consistent N\_Vector implementation (i.e. that it supplies the requisite vector operations).

A maxl argument that is  $\leq 0$  will result in the default value (5).

Some SUNDIALS solvers are designed to only work with left preconditioning (IDA and IDAS) and others with only right preconditioning (KINSOL). While it is possible to configure a SUNLinSol\_SPTFQMR object to use any of the preconditioning options with these solvers, this use mode is not supported and may result in inferior performance.

---

**Note:** With SUN\_PREC\_RIGHT or SUN\_PREC\_BOTH the initial guess must be zero (use *SUNLinSolSetZeroGuess()* to indicate the initial guess is zero).

---

*SUNErrCode* **SUNLinSol\_SPTFQMRSetPrecType**(*SUNLinearSolver* S, int pretype)

This function updates the flag indicating use of preconditioning.

**Arguments:**

- *S* – SUNLinSol\_SPGMR object to update.
- *pretype* – a flag indicating the type of preconditioning to use:
  - SUN\_PREC\_NONE
  - SUN\_PREC\_LEFT
  - SUN\_PREC\_RIGHT
  - SUN\_PREC\_BOTH

**Return value:**

- A *SUNErrCode*

*SUNErrCode* **SUNLinSol\_SPTFQMRSetMaxl**(*SUNLinearSolver* S, int maxl)

This function updates the number of linear solver iterations to allow.

**Arguments:**

- *S* – SUNLinSol\_SPTFQMR object to update.
- *maxl* – maximum number of linear iterations to allow. Any non-positive input will result in the default value (5).

**Return value:**

- A *SUNErrCode*

## 8.14.2 SUNLinSol\_SPTFQMR Description

The SUNLinSol\_SPTFQMR module defines the *content* field of a *SUNLinearSolver* to be the following structure:

```
struct _SUNLinearSolverContent_SPTFQMR {
    int maxl;
    int pretype;
    sunboolean_t zeroguess;
    int numiters;
    sunreal_t resnorm;
    int last_flag;
    SUNATimesFn ATimes;
    void* ATData;
    SUNPSetupFn Psetup;
    SUNPSolveFn Psolve;
    void* PData;
    N_Vector s1;
    N_Vector s2;
    N_Vector r_star;
    N_Vector q;
    N_Vector d;
    N_Vector v;
    N_Vector p;
    N_Vector *r;
    N_Vector u;
```

(continues on next page)

(continued from previous page)

```
N_Vector vtemp1;  
N_Vector vtemp2;  
N_Vector vtemp3;  
};
```

These entries of the *content* field contain the following information:

- `maxl` - number of TFQMR iterations to allow (default is 5),
- `pretype` - flag for type of preconditioning to employ (default is none),
- `numiters` - number of iterations from the most-recent solve,
- `resnorm` - final linear residual norm from the most-recent solve,
- `last_flag` - last error return flag from an internal function,
- `ATimes` - function pointer to perform  $Av$  product,
- `ATData` - pointer to structure for `ATimes`,
- `Psetup` - function pointer to preconditioner setup routine,
- `Psolve` - function pointer to preconditioner solve routine,
- `PData` - pointer to structure for `Psetup` and `Psolve`,
- `s1`, `s2` - vector pointers for supplied scaling matrices (default is NULL),
- `r_star` - a `N_Vector` which holds the initial scaled, preconditioned linear system residual,
- `q`, `d`, `v`, `p`, `u` - `N_Vector` used for workspace by the SPTFQMR algorithm,
- `r` - array of two `N_Vector` used for workspace within the SPTFQMR algorithm,
- `vtemp1`, `vtemp2`, `vtemp3` - temporary vector storage.

This solver is constructed to perform the following operations:

- During construction all `N_Vector` solver data is allocated, with vectors cloned from a template `N_Vector` that is input, and default solver parameters are set.
- User-facing “set” routines may be called to modify default solver parameters.
- Additional “set” routines are called by the SUNDIALS solver that interfaces with `SUNLinSol_SPTFQMR` to supply the `ATimes`, `Psetup`, and `Psolve` function pointers and `s1` and `s2` scaling vectors.
- In the “initialize” call, the solver parameters are checked for validity.
- In the “setup” call, any non-NULL `Psetup` function is called. Typically, this is provided by the SUNDIALS solver itself, that translates between the generic `Psetup` function and the solver-specific routine (solver-supplied or user-supplied).
- In the “solve” call the TFQMR iteration is performed. This will include scaling and preconditioning if those options have been supplied.

The `SUNLinSol_SPTFQMR` module defines implementations of all “iterative” linear solver operations listed in §8.1:

- `SUNLinSolGetType_SPTFQMR`
- `SUNLinSolInitialize_SPTFQMR`
- `SUNLinSolSetATimes_SPTFQMR`
- `SUNLinSolSetPreconditioner_SPTFQMR`
- `SUNLinSolSetScalingVectors_SPTFQMR`

- `SUNLinSolSetZeroGuess_SPTFQMR` – note the solver assumes a non-zero guess by default and the zero guess flag is reset to `SUNFALSE` after each call to `SUNLinSolSolve_SPTFQMR`.
- `SUNLinSolSetup_SPTFQMR`
- `SUNLinSolSolve_SPTFQMR`
- `SUNLinSolNumIters_SPTFQMR`
- `SUNLinSolResNorm_SPTFQMR`
- `SUNLinSolResid_SPTFQMR`
- `SUNLinSolLastFlag_SPTFQMR`
- `SUNLinSolSpace_SPTFQMR`
- `SUNLinSolFree_SPTFQMR`

## 8.15 The SUNLinSol\_SuperLUDIST Module

The `SUNLinSol_SuperLUDIST` implementation of the `SUNLinearSolver` class interfaces with the `SuperLU_DIST` library. This is designed to be used with the `SUNMatrix_SLUNRloc` [SUNMatrix](#), and one of the serial, threaded or parallel `N_Vector` implementations (`NVECTOR_SERIAL`, `NVECTOR_OPENMP`, `NVECTOR_PTHREADS`, `NVECTOR_PARALLEL`, `NVECTOR_PARHYP`).

### 8.15.1 SUNLinSol\_SuperLUDIST Usage

The header file to be included when using this module is `sunlinsol/sunlinsol_superludist.h`. The installed module library to link to is `libsundials_sunlinsol_superludist.lib` where `.lib` is typically `.so` for shared libraries and `.a` for static libraries.

The module `SUNLinSol_SuperLUDIST` provides the following user-callable routines:

**Warning:** Starting with `SuperLU_DIST` version 6.3.0, some structures were renamed to have a prefix for the floating point type. The double precision API functions have the prefix ‘d’. To maintain backwards compatibility with the unprefixed types, SUNDIALS provides macros to these `SuperLU_DIST` types with an ‘x’ prefix that expand to the correct prefix. E.g., the SUNDIALS macro `xLUstruct_t` expands to `dLUstruct_t` or `LUstruct_t` based on the `SuperLU_DIST` version.

*SUNLinearSolver* **SUNLinSol\_SuperLUDIST**(*N\_Vector* y, SuperMatrix \*A, gridinfo\_t \*grid, xLUstruct\_t \*lu, xScalePermstruct\_t \*scaleperm, xSOLVEstruct\_t \*solve, SuperLUStat\_t \*stat, superlu\_dist\_options\_t \*options, *SUNContext* sunctx)

This constructor function creates and allocates memory for a `SUNLinSol_SuperLUDIST` object.

#### Arguments:

- y – a template vector.
- A – a template matrix
- grid, lu, scaleperm, solve, stat, options – `SuperLU_DIST` object pointers.
- sunctx – the [SUNContext](#) object (see §4.2)

**Return value:**

If successful, a `SUNLinearSolver` object; otherwise this routine will return `NULL`.

**Notes:**

This routine analyzes the input matrix and vector to determine the linear system size and to assess the compatibility with the SuperLU\_DIST library.

This routine will perform consistency checks to ensure that it is called with consistent `N_Vector` and `SUNMatrix` implementations. These are currently limited to the `SUNMatrix_SLUNRloc` matrix type and the `NVECTOR_SERIAL`, `NVECTOR_OPENMP`, `NVECTOR_PTHREADS`, `NVECTOR_PARALLEL`, and `NVECTOR_PARHYP` vector types. As additional compatible matrix and vector implementations are added to SUNDIALS, these will be included within this compatibility check.

The `grid`, `lu`, `scaleperm`, `solve`, and `options` arguments are not checked and are passed directly to SuperLU\_DIST routines.

Some struct members of the `options` argument are modified internally by the `SUNLinSol_SuperLUDIST` solver. Specifically, the member `Fact` is modified in the setup and solve routines.

*sunrealtype* **SUNLinSol\_SuperLUDIST\_GetBerr**(*SUNLinearSolver* LS)

This function returns the componentwise relative backward error of the computed solution. It takes one argument, the `SUNLinearSolver` object. The return type is *sunrealtype*.

*gridinfo\_t* \***SUNLinSol\_SuperLUDIST\_GetGridinfo**(*SUNLinearSolver* LS)

This function returns a pointer to the SuperLU\_DIST structure that contains the 2D process grid. It takes one argument, the `SUNLinearSolver` object.

*xLUstruct\_t* \***SUNLinSol\_SuperLUDIST\_GetLUstruct**(*SUNLinearSolver* LS)

This function returns a pointer to the SuperLU\_DIST structure that contains the distributed L and U structures. It takes one argument, the `SUNLinearSolver` object.

*superlu\_dist\_options\_t* \***SUNLinSol\_SuperLUDIST\_GetSuperLUOptions**(*SUNLinearSolver* LS)

This function returns a pointer to the SuperLU\_DIST structure that contains the options which control how the linear system is factorized and solved. It takes one argument, the `SUNLinearSolver` object.

*xScalePermstruct\_t* \***SUNLinSol\_SuperLUDIST\_GetScalePermstruct**(*SUNLinearSolver* LS)

This function returns a pointer to the SuperLU\_DIST structure that contains the vectors that describe the transformations done to the matrix A. It takes one argument, the `SUNLinearSolver` object.

*xSOLVEstruct\_t* \***SUNLinSol\_SuperLUDIST\_GetSOLVEstruct**(*SUNLinearSolver* LS)

This function returns a pointer to the SuperLU\_DIST structure that contains information for communication during the solution phase. It takes one argument the `SUNLinearSolver` object.

*SuperLUStat\_t* \***SUNLinSol\_SuperLUDIST\_GetSuperLUStat**(*SUNLinearSolver* LS)

This function returns a pointer to the SuperLU\_DIST structure that stores information about runtime and flop count. It takes one argument, the `SUNLinearSolver` object.

## 8.15.2 SUNLinSol\_SuperLUDIST Description

The `SUNLinSol_SuperLUDIST` module defines the *content* field of a `SUNLinearSolver` to be the following structure:

```
struct _SUNLinearSolverContent_SuperLUDIST {
    sunbooleantype    first_factorize;
    int               last_flag;
    sunrealtype       berr;
    gridinfo_t        *grid;
```

(continues on next page)



(continued from previous page)

```

xLUstruct_t      *lu;
superlu_dist_options_t *options;
xScalePermstruct_t *scaleperm;
xSOLVEstruct_t   *solve;
SuperLUStat_t     *stat;
sunindextype      N;
};

```

These entries of the *content* field contain the following information:

- `first_factorize` – flag indicating whether the factorization has ever been performed,
- `last_flag` – last error return flag from internal function evaluations,
- `berr` – the componentwise relative backward error of the computed solution,
- `grid` – pointer to the SuperLU\_DIST structure that stores the 2D process grid
- `lu` – pointer to the SuperLU\_DIST structure that stores the distributed L and U factors,
- `scaleperm` – pointer to the SuperLU\_DIST structure that stores vectors describing the transformations done to the matrix A,
- `options` – pointer to the SuperLU\_DIST structure which contains options that control how the linear system is factorized and solved,
- `solve` – pointer to the SuperLU\_DIST solve structure,
- `stat` – pointer to the SuperLU\_DIST structure that stores information about runtime and flop count,
- `N` – the number of equations in the system.

The `SUNLinSol_SuperLUDIST` module is a `SUNLinearSolver` adapter for the SuperLU\_DIST sparse matrix factorization and solver library written by X. Sherry Li and collaborators [8, 37, 54, 55]. The package uses a SPMD parallel programming model and multithreading to enhance efficiency in distributed-memory parallel environments with multi-core nodes and possibly GPU accelerators. It uses MPI for communication, OpenMP for threading, and CUDA for GPU support. In order to use the `SUNLinSol_SuperLUDIST` interface to SuperLU\_DIST, it is assumed that SuperLU\_DIST has been installed on the system prior to installation of SUNDIALS, and that SUNDIALS has been configured appropriately to link with SuperLU\_DIST (see §12.4 for details). Additionally, the wrapper only supports double-precision calculations, and therefore cannot be compiled if SUNDIALS is configured to use single or extended precision. Moreover, since the SuperLU\_DIST library may be installed to support either 32-bit or 64-bit integers, it is assumed that the SuperLU\_DIST library is installed using the same integer size as SUNDIALS.

The SuperLU\_DIST library provides many options to control how a linear system will be factorized and solved. These options may be set by a user on an instance of the `superlu_dist_options_t` struct, and then it may be provided as an argument to the `SUNLinSol_SuperLUDIST` constructor. The `SUNLinSol_SuperLUDIST` module will respect all options set except for `Fact` – this option is necessarily modified by the `SUNLinSol_SuperLUDIST` module in the setup and solve routines.

Since the linear systems that arise within the context of SUNDIALS calculations will typically have identical sparsity patterns, the `SUNLinSol_SuperLUDIST` module is constructed to perform the following operations:

- The first time that the “setup” routine is called, it sets the SuperLU\_DIST option `Fact` to `DOFACT` so that a subsequent call to the “solve” routine will perform a symbolic factorization, followed by an initial numerical factorization before continuing to solve the system.
- On subsequent calls to the “setup” routine, it sets the SuperLU\_DIST option `Fact` to `SamePattern` so that a subsequent call to “solve” will perform factorization assuming the same sparsity pattern as prior, i.e. it will reuse the column permutation vector.

- If “setup” is called prior to the “solve” routine, then the “solve” routine will perform a symbolic factorization, followed by an initial numerical factorization before continuing to the sparse triangular solves, and, potentially, iterative refinement. If “setup” is not called prior, “solve” will skip to the triangular solve step. We note that in this solve SuperLU\_DIST operates on the native data arrays for the right-hand side and solution vectors, without requiring costly data copies.

The SUNLinSol\_SuperLUDIST module defines implementations of all “direct” linear solver operations listed in §8.1:

- SUNLinSolGetType\_SuperLUDIST
- SUNLinSolInitialize\_SuperLUDIST – this sets the `first_factorize` flag to 1 and resets the internal SuperLU\_DIST statistics variables.
- SUNLinSolSetup\_SuperLUDIST – this sets the appropriate SuperLU\_DIST options so that a subsequent solve will perform a symbolic and numerical factorization before proceeding with the triangular solves
- SUNLinSolSolve\_SuperLUDIST – this calls the SuperLU\_DIST solve routine to perform factorization (if the setup routine was called prior) and then use the \$LU\$ factors to solve the linear system.
- SUNLinSolLastFlag\_SuperLUDIST
- SUNLinSolSpace\_SuperLUDIST – this only returns information for the storage within the solver *interface*, i.e. storage for the integers `last_flag` and `first_factorize`. For additional space requirements, see the SuperLU\_DIST documentation.
- SUNLinSolFree\_SuperLUDIST

## 8.16 The SUNLinSol\_SuperLUMT Module

The SUNLinSol\_SuperLUMT implementation of the SUNLinearSolver class interfaces with the SuperLU\_MT library. This is designed to be used with the corresponding SUNMATRIX\_SPARSE matrix type, and one of the serial or shared-memory N\_Vector implementations (NVECTOR\_SERIAL, NVECTOR\_OPENMP, or NVECTOR\_PTHREADS). While these are compatible, it is not recommended to use a threaded vector module with SUNLinSol\_SuperLUMT unless it is the NVECTOR\_OPENMP module and the SuperLU\_MT library has also been compiled with OpenMP.

### 8.16.1 SUNLinSol\_SuperLUMT Usage

The header file to be included when using this module is `sunlinsol/sunlinsol.SuperLUMT.h`. The installed module library to link to is `libsundials_sunlinsolsuperlumt.lib` where *.lib* is typically *.so* for shared libraries and *.a* for static libraries.

The module SUNLinSol\_SuperLUMT provides the following user-callable routines:

*SUNLinearSolver* **SUNLinSol\_SuperLUMT**(*N\_Vector* y, *SUNMatrix* A, int num\_threads, *SUNContext* sunctx)

This constructor function creates and allocates memory for a SUNLinSol\_SuperLUMT object.

#### Arguments:

- y – a template vector.
- A – a template matrix
- *num\_threads* – desired number of threads (OpenMP or Pthreads, depending on how SuperLU\_MT was installed) to use during the factorization steps.
- *sunctx* – the *SUNContext* object (see §4.2)

**Return value:**

If successful, a `SUNLinearSolver` object; otherwise this routine will return `NULL`.

**Notes:**

This routine analyzes the input matrix and vector to determine the linear system size and to assess compatibility with the SuperLU\_MT library.

This routine will perform consistency checks to ensure that it is called with consistent `N_Vector` and `SUN-Matrix` implementations. These are currently limited to the `SUNMATRIX_SPARSE` matrix type (using either CSR or CSC storage formats) and the `NVECTOR_SERIAL`, `NVECTOR_OPENMP`, and `NVECTOR_PTHREADS` vector types. As additional compatible matrix and vector implementations are added to SUNDIALS, these will be included within this compatibility check.

The `num_threads` argument is not checked and is passed directly to SuperLU\_MT routines.

*SUNErrCode* **SUNLinSol\_SuperLUMTSetOrdering**(*SUNLinearSolver* S, int ordering\_choice)

This function sets the ordering used by SuperLU\_MT for reducing fill in the linear solve.

**Arguments:**

- *S* – the `SUNLinSol_SuperLUMT` object to update.
- *ordering\_choice*:
  0. natural ordering
  1. minimal degree ordering on  $A^T A$
  2. minimal degree ordering on  $A^T + A$
  3. COLAMD ordering for unsymmetric matrices

The default is 3 for COLAMD.

**Return value:**

- A *SUNErrCode*

## 8.16.2 SUNLinSol\_SuperLUMT Description

The `SUNLinSol_SuperLUMT` module defines the *content* field of a `SUNLinearSolver` to be the following structure:

```
struct _SUNLinearSolverContent_SuperLUMT {
    int      last_flag;
    int      first_factorize;
    SuperMatrix *A, *AC, *L, *U, *B;
    Gstat_t  *Gstat;
    sunindextype *perm_r, *perm_c;
    sunindextype N;
    int      num_threads;
    sunrealtype diag_pivot_thresh;
    int      ordering;
    superlumt_options_t *options;
};
```

These entries of the *content* field contain the following information:

- `last_flag` - last error return flag from internal function evaluations,
- `first_factorize` - flag indicating whether the factorization has ever been performed,
- `A`, `AC`, `L`, `U`, `B` - `SuperMatrix` pointers used in solve,

- `Gstat` - `GStat_t` object used in solve,
- `perm_r`, `perm_c` - permutation arrays used in solve,
- `N` - size of the linear system,
- `num_threads` - number of OpenMP/Pthreads threads to use,
- `diag_pivot_thresh` - threshold on diagonal pivoting,
- `ordering` - flag for which reordering algorithm to use,
- `options` - pointer to SuperLU\_MT options structure.

The `SUNLinSol_SuperLUMT` module is a `SUNLinearSolver` wrapper for the SuperLU\_MT sparse matrix factorization and solver library written by X. Sherry Li and collaborators [9, 29, 53]. The package performs matrix factorization using threads to enhance efficiency in shared memory parallel environments. It should be noted that threads are only used in the factorization step. In order to use the `SUNLinSol_SuperLUMT` interface to SuperLU\_MT, it is assumed that SuperLU\_MT has been installed on the system prior to installation of SUNDIALS, and that SUNDIALS has been configured appropriately to link with SuperLU\_MT (see §12.4 for details). Additionally, this wrapper only supports single- and double-precision calculations, and therefore cannot be compiled if SUNDIALS is configured to have `sun-realttype` set to `extended` (see §4.1 for details). Moreover, since the SuperLU\_MT library may be installed to support either 32-bit or 64-bit integers, it is assumed that the SuperLU\_MT library is installed using the same integer precision as the SUNDIALS `sunindextype` option.

The SuperLU\_MT library has a symbolic factorization routine that computes the permutation of the linear system matrix to reduce fill-in on subsequent *LU* factorizations (using COLAMD, minimal degree ordering on  $A^T * A$ , minimal degree ordering on  $A^T + A$ , or natural ordering). Of these ordering choices, the default value in the `SUNLinSol_SuperLUMT` module is the COLAMD ordering.

Since the linear systems that arise within the context of SUNDIALS calculations will typically have identical sparsity patterns, the `SUNLinSol_SuperLUMT` module is constructed to perform the following operations:

- The first time that the “setup” routine is called, it performs the symbolic factorization, followed by an initial numerical factorization.
- On subsequent calls to the “setup” routine, it skips the symbolic factorization, and only refactors the input matrix.
- The “solve” call performs pivoting and forward and backward substitution using the stored SuperLU\_MT data structures. We note that in this solve SuperLU\_MT operates on the native data arrays for the right-hand side and solution vectors, without requiring costly data copies.

The `SUNLinSol_SuperLUMT` module defines implementations of all “direct” linear solver operations listed in §8.1:

- `SUNLinSolGetType_SuperLUMT`
- `SUNLinSolInitialize_SuperLUMT` – this sets the `first_factorize` flag to 1 and resets the internal SuperLU\_MT statistics variables.
- `SUNLinSolSetup_SuperLUMT` – this performs either a *LU* factorization or refactorization of the input matrix.
- `SUNLinSolSolve_SuperLUMT` – this calls the appropriate SuperLU\_MT solve routine to utilize the *LU* factors to solve the linear system.
- `SUNLinSolLastFlag_SuperLUMT`
- `SUNLinSolSpace_SuperLUMT` – this only returns information for the storage within the solver *interface*, i.e. storage for the integers `last_flag` and `first_factorize`. For additional space requirements, see the SuperLU\_MT documentation.
- `SUNLinSolFree_SuperLUMT`

## 8.17 The SUNLinSol\_cuSolverSp\_batchQR Module

The SUNLinSol\_cuSolverSp\_batchQR implementation of the SUNLinearSolver class is designed to be used with the SUNMATRIX\_CUSPARSE matrix, and the NVECTOR\_CUDA vector. The header file to include when using this module is sunlinsol/sunlinsol\_cusolversp\_batchqr.h. The installed library to link to is libsundials\_sunlinsolcusolversp.lib where .lib is typically .so for shared libraries and .a for static libraries.

**Warning:** The SUNLinearSolver\_cuSolverSp\_batchQR module is experimental and subject to change.

### 8.17.1 SUNLinSol\_cuSolverSp\_batchQR description

The SUNLinearSolver\_cuSolverSp\_batchQR implementation provides an interface to the batched sparse QR factorization method provided by the NVIDIA cuSOLVER library [6]. The module is designed for solving block diagonal linear systems of the form

$$\begin{bmatrix} \mathbf{A}_1 & 0 & \cdots & 0 \\ 0 & \mathbf{A}_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mathbf{A}_n \end{bmatrix} x_j = b_j$$

where all block matrices  $\mathbf{A}_j$  share the same sparsity pattern. The matrix must be the SUNMatrix.cuSparse.

### 8.17.2 SUNLinSol\_cuSolverSp\_batchQR functions

The SUNLinearSolver\_cuSolverSp\_batchQR module defines implementations of all “direct” linear solver operations listed in §8.1:

- SUNLinSolGetType\_cuSolverSp\_batchQR
- SUNLinSolInitialize\_cuSolverSp\_batchQR – this sets the first\_factorize flag to 1
- SUNLinSolSetup\_cuSolverSp\_batchQR – this always copies the relevant SUNMATRIX\_SPARSE data to the GPU; if this is the first setup it will perform symbolic analysis on the system
- SUNLinSolSolve\_cuSolverSp\_batchQR – this calls the cusolverSpXcsrqrsvBatched routine to perform factorization
- SUNLinSolLastFlag\_cuSolverSp\_batchQR
- SUNLinSolFree\_cuSolverSp\_batchQR

In addition, the module provides the following user-callable routines:

*SUNLinearSolver* **SUNLinSol\_cuSolverSp\_batchQR**(*N\_Vector* y, *SUNMatrix* A, cusolverHandle\_t cusol, *SUNContext* sunctx)

The function SUNLinSol\_cuSolverSp\_batchQR creates and allocates memory for a SUNLinearSolver object.

**Arguments:**

- y – a vector for checking compatibility with the solver.
- A – a SUNMATRIX\_cuSparse matrix for checking compatibility with the solver.
- cusol – cuSolverSp object to use.
- sunctx – the *SUNContext* object (see §4.2)

**Return value:**

If successful, a `SUNLinearSolver` object. If either  $A$  or  $y$  are incompatible then this routine will return `NULL`.

**Notes:**

This routine will perform consistency checks to ensure that it is called with consistent `N_Vector` and `SUNMatrix` implementations. These are currently limited to the `SUNMATRIX_CUSPARSE` matrix type and the `NVECTOR_CUDA` vector type. Since the `SUNMATRIX_CUSPARSE` matrix type is only compatible with the `NVECTOR_CUDA` the restriction is also in place for the linear solver. As additional compatible matrix and vector implementations are added to `SUNDIALS`, these will be included within this compatibility check.

void `SUNLinSol_cuSolverSp_batchQR_GetDescription`(*SUNLinearSolver* LS, char \*\*desc)

The function `SUNLinSol_cuSolverSp_batchQR_GetDescription` accesses the string description of the object (empty by default).

void `SUNLinSol_cuSolverSp_batchQR_SetDescription`(*SUNLinearSolver* LS, const char \*desc)

The function `SUNLinSol_cuSolverSp_batchQR_SetDescription` sets the string description of the object (empty by default).

void `SUNLinSol_cuSolverSp_batchQR_GetDeviceSpace`(*SUNLinearSolver* S, size\_t \*cuSolverInternal, size\_t \*cuSolverWorkspace)

The function `SUNLinSol_cuSolverSp_batchQR_GetDeviceSpace` returns the `cuSOLVER` batch QR method internal buffer size, in bytes, in the argument `cuSolverInternal` and the `cuSOLVER` batch QR workspace buffer size, in bytes, in the argument `cuSolverWorkspace`. The size of the internal buffer is proportional to the number of matrix blocks while the size of the workspace is almost independent of the number of blocks.

### 8.17.3 SUNLinSol\_cuSolverSp\_batchQR content

The `SUNLinSol_cuSolverSp_batchQR` module defines the *content* field of a `SUNLinearSolver` to be the following structure:

```
struct _SUNLinearSolverContent_cuSolverSp_batchQR {
    int                last_flag;           /* last return flag */
    sunboolean_t       first_factorize;     /* is this the first factorization? */
    size_t             internal_size;       /* size of cusolver buffer for Q and R */
    size_t             workspace_size;      /* size of cusolver memory for factorization */
    cusolverSpHandle_t cusolver_handle;     /* cuSolverSp context */
    csrqrInfo_t        info;               /* opaque cusolver data structure */
    void*              workspace;          /* memory block used by cusolver */
    const char*        desc;              /* description of this linear solver */
};
```

## 8.18 The SUNLINEARSOLVER\_GINKGO Module

New in version 6.4.0.

The `SUNLINEARSOLVER_GINKGO` implementation of the `SUNLinearSolver` API provides an interface to the linear solvers from the Ginkgo linear algebra library [11]. Since Ginkgo is a modern C++ library, `SUNLINEARSOLVER_GINKGO` is also written in modern C++ (specifically, C++14). Unlike most other `SUNDIALS` modules, it is a header only library. To use the `SUNLINEARSOLVER_GINKGO` `SUNLinearSolver`, users will need to include `sunlinsol/sunlinsol_ginkgo.hpp`. The module is meant to be used with the `SUNMATRIX_GINKGO` module

described in §7.10. Instructions on building SUNDIALS with Ginkgo enabled are given in §12.4. For instructions on building and using Ginkgo itself, refer to the [Ginkgo website and documentation](#).

**Note:** It is assumed that users of this module are aware of how to use Ginkgo. This module does not try to encapsulate Ginkgo linear solvers, rather it provides a lightweight interoperability layer between Ginkgo and SUNDIALS. Most, if not all, of the Ginkgo linear solver should work with this interface.

### 8.18.1 Using SUNLINEARSOLVER\_GINKGO

After choosing a compatible `N_Vector` (see §7.10.1) and creating a Ginkgo-enabled `SUNMatrix` (see §7.10) to use the `SUNLINEARSOLVER_GINKGO` module, we first create a Ginkgo stopping criteria object. Importantly, the `sundials::ginkgo::DefaultStop` class provided by SUNDIALS implements a stopping criterion that matches the default SUNDIALS stopping criterion. Namely, it checks if the max iterations (5 by default) were reached or if the absolute residual norm was below a specified tolerance. The criterion can be created just like any other Ginkgo stopping criteria:

```
auto crit{sundials::ginkgo::DefaultStop::build().with_max_iters(max_iters).on(gko_exec)};
```

**Warning:** It is *highly* recommended to employ this criterion when using Ginkgo solvers with SUNDIALS, but it is optional. However, to use the Ginkgo multigrid or cbgmres linear solvers, different Ginkgo criterion must be used.

Once we have created our stopping criterion, we create a Ginkgo solver factory object and wrap it in a `sundials::ginkgo::LinearSolver` object. In this example, we create a Ginkgo conjugate gradient solver:

```
using GkoMatrixType = gko::matrix::Csr<sunrealtype, sunindextype>;
using GkoSolverType = gko::solver::Cg<sunrealtype>;

auto gko_solver_factory = gko::share(
    GkoSolverType::build().with_criteria(std::move(crit)).on(gko_exec));

sundials::ginkgo::LinearSolver<GkoSolverType, GkoMatrixType> LS{
    gko_solver_factory, suncctx};
```

Finally, we can pass the instance of `sundials::ginkgo::LinearSolver` to any function expecting a `SUNLinearSolver` object through the implicit conversion operator or explicit conversion function.

```
// Attach linear solver and matrix to CVODE.
//
// Implicit conversion from sundials::ginkgo::LinearSolver<GkoSolverType, GkoMatrixType>
// to a SUNLinearSolver object is done.
//
// For details about creating A see the SUNMATRIX_GINKGO module.
CvodeSetLinearSolver(cvode_mem, LS, A);

// Alternatively with explicit conversion of LS to a SUNLinearSolver
// and A to a SUNMatrix:
CvodeSetLinearSolver(cvode_mem, LS->Convert(), A->Convert());
```



**Warning:** `SUNLinSolFree()` should never be called on a `SUNLinearSolver` that was created via conversion from a `sundials::ginkgo::LinearSolver`. Doing so may result in a double free.

## 8.18.2 SUNLINEARSOLVER\_GINKGO API

In this section we list the public API of the `sundials::ginkgo::LinearSolver` class.

```
template<class GkoSolverType, class GkoMatrixType>
```

```
class sundials::ginkgo::LinearSolver : public sundials::ConvertibleTo<SUNLinearSolver>
```

```
LinearSolver() = default;
```

Default constructor - means the solver must be moved to.

```
LinearSolver(std::shared_ptr<typename GkoSolverType::Factory> gko_solver_factory, SUNContext suncctx)
```

Constructs a new `LinearSolver` from a Ginkgo solver factory.

### Parameters

- **gko\_solver\_factory** – The Ginkgo solver factory (typically `gko::matrix::<type>::Factory`)
- **suncctx** – The SUNDIALS simulation context (`SUNContext`)

```
LinearSolver(LinearSolver &&that_solver) noexcept
```

Move constructor.

```
LinearSolver &operator=(LinearSolver &&rhs)
```

Move assignment.

```
~LinearSolver() override = default
```

Default destructor.

```
operator SUNLinearSolver() override
```

Implicit conversion to a `SUNLinearSolver`.

```
operator SUNLinearSolver() const override
```

Implicit conversion to a `SUNLinearSolver`.

```
SUNLinearSolver Convert() override
```

Explicit conversion to a `SUNLinearSolver`.

```
SUNLinearSolver Convert() const override
```

Explicit conversion to a `SUNLinearSolver`.

```
std::shared_ptr<const gko::Executor> GkoExec() const
```

Get the `gko::Executor` associated with the Ginkgo solver.

```
std::shared_ptr<typename GkoSolverType::Factory> GkoFactory()
```

Get the underlying Ginkgo solver factory.

```
GkoSolverType *GkoSolver()
```

Get the underlying Ginkgo solver.

---

**Note:** This will be `nullptr` until the linear solver setup phase.

---



int **NumIters**() const

Get the number of linear solver iterations in the most recent solve.

sunrealtype **ResNorm**() const

Get the residual norm of the solution at the end of the last solve.

The type of residual norm depends on the Ginkgo stopping criteria used with the solver. With the `DefaultStop` criteria this would be the absolute residual 2-norm.

*GkoSolverType* **\*Setup**(*Matrix*<*GkoMatrixType*> \*A)

Setup the linear system.

**Parameters**

**A** – the linear system matrix

**Returns**

Pointer to the Ginkgo solver generated from the factory

gko::LinOp **\*Solve**(N\_Vector b, N\_Vector x, sunrealtype tol)

Solve the linear system  $Ax = b$  to the specified tolerance.

**Parameters**

- **b** – the right-hand side vector
- **x** – the solution vector
- **tol** – the tolerance to solve the system to

**Returns**

gko::LinOp\* the solution

## 8.19 The SUNLINEARSOLVER\_KOKKOSDENSE Module

New in version 6.4.0.

The `SUNLINEARSOLVER_KOKKOSDENSE` *SUNLinearSolver* implementation provides an interface to KokkosKernels [63] linear solvers for dense and batched dense (block-diagonal) systems. Since Kokkos is a modern C++ library, the module is also written in modern C++ (it requires C++14) as a header only library. To utilize this *SUNLinearSolver* user will need to include `sunlinsol/sunlinsol_kokkosdense.hpp`. More instructions on building SUNDIALS with Kokkos and KokkosKernels enabled are given in §12.4. For instructions on building and using Kokkos and KokkosKernels, refer to the [Kokkos](#) and [KokkosKernels](#) documentation.

### 8.19.1 Using SUNLINEARSOLVER\_KOKKOSDENSE

The `SUNLINEARSOLVER_KOKKOSDENSE` module is defined by the `DenseLinearSolver` templated class in the `sundials::kokkos` namespace:

```
template<class ExecSpace = Kokkos::DefaultExecutionSpace,
        class MemSpace = typename ExecSpace::memory_space>
class DenseLinearSolver : public sundials::impl::BaseLinearSolver,
                        public sundials::ConvertibleTo<SUNLinearSolver>
```

To use the `SUNLINEARSOLVER_KOKKOSDENSE` module, we begin by constructing an instance of a dense linear solver e.g.,

```
// Create a dense linear solver
sundials::kokkos::DenseLinearSolver<> LS{sunctx};
```

Instances of the `DenseLinearSolver` class are implicitly or explicitly (using the `Convert()` method) convertible to a `SUNLinearSolver` e.g.,

```
sundials::kokkos::DenseLinearSolver<> LS{sunctx};
SUNLinearSolver LSA = LS;           // implicit conversion to SUNLinearSolver
SUNLinearSolver LSB = LS.Convert(); // explicit conversion to SUNLinearSolver
```

**Warning:** `SUNLinSolFree()` should never be called on a `SUNLinearSolver` that was created via conversion from a `sundials::kokkos::DenseLinearSolver`. Doing so may result in a double free.

The `SUNLINEARSOLVER_KOKKOSDENSE` module is compatible with the `NVECTOR_KOKKOS` vector module (see §6.14) and `SUNMATRIX_KOKKOSDENSE` matrix module (see §7.11).

## 8.19.2 SUNLINEARSOLVER\_KOKKOSDENSE API

In this section we list the public API of the `sundials::kokkos::DenseLinearSolver` class.

```
template<class ExecSpace = Kokkos::DefaultExecutionSpace, class MemSpace = typename
ExecSpace::memory_space>
class DenseLinearSolver : public sundials::impl::BaseLinearSolver, public
sundials::ConvertibleTo<SUNLinearSolver>
```

**DenseLinearSolver()** = default;

Default constructor - means the solver must be moved to.

**DenseLinearSolver(SUNContext sunctx)**

Constructs a new `DenseLinearSolver`.

### Parameters

**sunctx** – The SUNDIALS simulation context (*SUNContext*)

**DenseLinearSolver(DenseLinearSolver &&that\_solver)** noexcept

Move constructor.

*DenseLinearSolver* &**operator=**(*DenseLinearSolver* &&rhs)

Move assignment.

**~DenseLinearSolver()** override = default

Default destructor.

**operator SUNLinearSolver()** override

Implicit conversion to a *SUNLinearSolver*.

**operator SUNLinearSolver()** const override

Implicit conversion to a *SUNLinearSolver*.

*SUNLinearSolver* **Convert()** override

Explicit conversion to a *SUNLinearSolver*.

*SUNLinearSolver* **Convert()** const override

Explicit conversion to a *SUNLinearSolver*.

## 8.20 SUNLinearSolver Examples

There are SUNLinearSolver examples that may be installed for each implementation; these make use of the functions in `test_sunlinsol.c`. These example functions show simple usage of the SUNLinearSolver family of modules. The inputs to the examples depend on the linear solver type, and are output to `stdout` if the example is run without the appropriate number of command-line arguments.

The following is a list of the example functions in `test_sunlinsol.c`:

- `Test_SUNLinSolGetType`: Verifies the returned solver type against the value that should be returned.
- `Test_SUNLinSolGetID`: Verifies the returned solver identifier against the value that should be returned.
- `Test_SUNLinSolInitialize`: Verifies that `SUNLinSolInitialize` can be called and returns successfully.
- `Test_SUNLinSolSetup`: Verifies that `SUNLinSolSetup` can be called and returns successfully.
- `Test_SUNLinSolSolve`: Given a `SUNMatrix` object  $A$ , `N_Vector` objects  $x$  and  $b$  (where  $Ax = b$ ) and a desired solution tolerance `tol`, this routine clones  $x$  into a new vector  $y$ , calls `SUNLinSolSolve` to fill  $y$  as the solution to  $Ay = b$  (to the input tolerance), verifies that each entry in  $x$  and  $y$  match to within  $10 * \text{tol}$ , and overwrites  $x$  with  $y$  prior to returning (in case the calling routine would like to investigate further).
- `Test_SUNLinSolSetATimes` (iterative solvers only): Verifies that `SUNLinSolSetATimes` can be called and returns successfully.
- `Test_SUNLinSolSetPreconditioner` (iterative solvers only): Verifies that `SUNLinSolSetPreconditioner` can be called and returns successfully.
- `Test_SUNLinSolSetScalingVectors` (iterative solvers only): Verifies that `SUNLinSolSetScalingVectors` can be called and returns successfully.
- `Test_SUNLinSolSetZeroGuess` (iterative solvers only): Verifies that `SUNLinSolSetZeroGuess` can be called and returns successfully.
- `Test_SUNLinSolLastFlag`: Verifies that `SUNLinSolLastFlag` can be called, and outputs the result to `stdout`.
- `Test_SUNLinSolNumIters` (iterative solvers only): Verifies that `SUNLinSolNumIters` can be called, and outputs the result to `stdout`.
- `Test_SUNLinSolResNorm` (iterative solvers only): Verifies that `SUNLinSolResNorm` can be called, and that the result is non-negative.
- `Test_SUNLinSolResid` (iterative solvers only): Verifies that `SUNLinSolResid` can be called.
- `Test_SUNLinSolSpace` verifies that `SUNLinSolSpace` can be called, and outputs the results to `stdout`.

We'll note that these tests should be performed in a particular order. For either direct or iterative linear solvers, `Test_SUNLinSolInitialize` must be called before `Test_SUNLinSolSetup`, which must be called before `Test_SUNLinSolSolve`. Additionally, for iterative linear solvers `Test_SUNLinSolSetATimes`, `Test_SUNLinSolSetPreconditioner` and `Test_SUNLinSolSetScalingVectors` should be called before `Test_SUNLinSolInitialize`; similarly `Test_SUNLinSolNumIters`, `Test_SUNLinSolResNorm` and `Test_SUNLinSolResid` should be called after `Test_SUNLinSolSolve`. These are called in the appropriate order in all of the example problems.



## Chapter 9

# Nonlinear Algebraic Solvers

SUNDIALS time integration packages are written in terms of generic nonlinear solver operations defined by the SUNNonlinSol API and implemented by a particular SUNNonlinSol module of type `SUNNonlinearSolver`. Users can supply their own SUNNonlinSol module, or use one of the modules provided with SUNDIALS. Depending on the package, nonlinear solver modules can either target system presented in a rootfinding ( $F(y) = 0$ ) or fixed-point ( $G(y) = y$ ) formulation. For more information on the formulation of the nonlinear system(s) see the §9.2 section.

The time integrators in SUNDIALS specify a default nonlinear solver module and as such this chapter is intended for users that wish to use a non-default nonlinear solver module or would like to provide their own nonlinear solver implementation. Users interested in using a non-default solver module may skip the description of the SUNNonlinSol API in section §9.1 and proceeded to the subsequent sections in this chapter that describe the SUNNonlinSol modules provided with SUNDIALS.

For users interested in providing their own SUNNonlinSol module, the following section presents the SUNNonlinSol API and its implementation beginning with the definition of SUNNonlinSol functions in the sections §9.1.1, §9.1.2 and §9.1.3. This is followed by the definition of functions supplied to a nonlinear solver implementation in the section §9.1.4. The nonlinear solver return codes are given in the section §9.1.5. The `SUNNonlinearSolver` type and the generic SUNNonlinSol module are defined in the section §9.1.6. Finally, the section §9.1.7 lists the requirements for supplying a custom SUNNonlinSol module. Users wishing to supply their own SUNNonlinSol module are encouraged to use the SUNNonlinSol implementations provided with SUNDIALS as a template for supplying custom nonlinear solver modules.

## 9.1 The SUNNonlinearSolver API

The SUNNonlinSol API defines several nonlinear solver operations that enable SUNDIALS integrators to utilize any SUNNonlinSol implementation that provides the required functions. These functions can be divided into three categories. The first are the core nonlinear solver functions. The second consists of “set” routines to supply the nonlinear solver with functions provided by the SUNDIALS time integrators and to modify solver parameters. The final group consists of “get” routines for retrieving nonlinear solver statistics. All of these functions are defined in the header file `sundials/sundials_nonlinearsolver.h`.

### 9.1.1 SUNNonlinearSolver core functions

The core nonlinear solver functions consist of two required functions to get the nonlinear solver type (*SUNNonlinSolGetType()*) and solve the nonlinear system (*SUNNonlinSolSolve()*). The remaining three functions for nonlinear solver initialization (*SUNNonlinSolInitialize()*), setup (*SUNNonlinSolSetup()*), and destruction (*SUNNonlinSolFree()*) are optional.

enum **SUNNonlinearSolver\_Type**

An identifier indicating the form of the nonlinear system expected by the nonlinear solver.

enumerator **SUNNONLINEARSOLVER\_ROOTFIND**

The nonlinear solver expects systems in rootfinding form  $F(y) = 0$

enumerator **SUNNONLINEARSOLVER\_FIXEDPOINT**

The nonlinear solver expects systems in fixed-point form  $G(y) = y$ .

*SUNNonlinearSolver\_Type* **SUNNonlinSolGetType**(*SUNNonlinearSolver* NLS)

This *required* function returns the nonlinear solver type.

**Arguments:**

- *NLS* – a SUNNonlinSol object.

**Return value:**

The *SUNNonlinearSolver\_Type* type identifier for the nonlinear solver.

*SUNErrCode* **SUNNonlinSolInitialize**(*SUNNonlinearSolver* NLS)

This *optional* function handles nonlinear solver initialization and may perform any necessary memory allocations.

**Arguments:**

- *NLS* – a SUNNonlinSol object.

**Return value:**

A *SUNErrCode*.

**Notes:**

It is assumed all solver-specific options have been set prior to calling *SUNNonlinSolInitialize()*. SUNNonlinSol implementations that do not require initialization may set this operation to NULL.

*SUNErrCode* **SUNNonlinSolSetup**(*SUNNonlinearSolver* NLS, *N\_Vector* y, void \*mem)

This *optional* function performs any solver setup needed for a nonlinear solve.

**Arguments:**

- *NLS* – a SUNNonlinSol object.
- *y* – the initial guess passed to the nonlinear solver.
- *mem* – the SUNDIALS integrator memory structure.

**Return value:**

A *SUNErrCode*.

**Notes:**

SUNDIALS integrators call *SUNNonlinSolSetup()* before each step attempt. SUNNonlinSol implementations that do not require setup may set this operation to NULL.

int **SUNNonlinSolSolve**(*SUNNonlinearSolver* NLS, *N\_Vector* y0, *N\_Vector* ycor, *N\_Vector* w, *sunrealtype* tol, *sunbooleantype* callLSetup, void \*mem)

This *required* function solves the nonlinear system  $F(y) = 0$  or  $G(y) = y$ .

**Arguments:**

- *NLS* – a SUNNonlinSol object.
- *y0* – the predicted value for the new solution state. This *must* remain unchanged throughout the solution process.
- *ycor* – on input the initial guess for the correction to the predicted state (zero) and on output the final correction to the predicted state.
- *w* – the solution error weight vector used for computing weighted error norms.
- *tol* – the requested solution tolerance in the weighted root-mean-squared norm.
- *callLSetup* – a flag indicating that the integrator recommends for the linear solver setup function to be called.
- *mem* – the SUNDIALS integrator memory structure.

**Return value:**

The return value is zero for a successful solve, a positive value for a recoverable error (i.e., the solve failed and the integrator should reduce the step size and reattempt the step), and a negative value for an unrecoverable error (i.e., the solve failed and the integrator should halt and return an error to the user).

*SUNErrCode* **SUNNonlinSolFree**(*SUNNonlinearSolver* NLS)

This *optional* function frees any memory allocated by the nonlinear solver.

**Arguments:**

- *NLS* – a SUNNonlinSol object.

**Return value:**

- A *SUNErrCode*

### 9.1.2 SUNNonlinearSolver “set” functions

The following functions are used to supply nonlinear solver modules with functions defined by the SUNDIALS integrators and to modify solver parameters. Only the routine for setting the nonlinear system defining function (*SUNNonlinSolSetSysFn()*) is required. All other set functions are optional.

*SUNErrCode* **SUNNonlinSolSetSysFn**(*SUNNonlinearSolver* NLS, *SUNNonlinSolSysFn* SysFn)

This *required* function is used to provide the nonlinear solver with the function defining the nonlinear system. This is the function  $F(y)$  in  $F(y) = 0$  for SUNNONLINEARSOLVER\_ROOTFIND modules or  $G(y)$  in  $G(y) = y$  for SUNNONLINEARSOLVER\_FIXEDPOINT modules.

**Arguments:**

- *NLS* – a SUNNonlinSol object.
- *SysFn* – the function defining the nonlinear system. See §9.1.4 for the definition of *SUNNonlinSolSysFn*.

**Return value:**

- A *SUNErrCode*

*SUNErrCode* **SUNNonlinSolSetLSetupFn**(*SUNNonlinearSolver* NLS, *SUNNonlinSolLSetupFn* SetupFn)

This *optional* function is called by SUNDIALS integrators to provide the nonlinear solver with access to its linear solver setup function.

**Arguments:**

- *NLS* – a `SUNNonlinSol` object.
- *SetupFn* – a wrapper function to the SUNDIALS integrator’s linear solver setup function. See §9.1.4 for the definition of `SUNNonlinSolSetupFn`.

**Return value:**

- A `SUNErrCode`

**Notes:**

The `SUNNonlinSolSetupFn` function sets up the linear system  $Ax = b$  where  $A = \frac{\partial F}{\partial y}$  is the linearization of the nonlinear residual function  $F(y) = 0$  (when using `SUNLinSol` direct linear solvers) or calls the user-defined preconditioner setup function (when using `SUNLinSol` iterative linear solvers). `SUNNonlinSol` implementations that do not require solving this system, do not utilize `SUNLinSol` linear solvers, or use `SUNLinSol` linear solvers that do not require setup may set this operation to `NULL`.

`SUNErrCode` `SUNNonlinSolSetLSolveFn`(`SUNNonlinearSolver` *NLS*, `SUNNonlinSolSolveFn` *SolveFn*)

This *optional* function is called by SUNDIALS integrators to provide the nonlinear solver with access to its linear solver solve function.

**Arguments:**

- *NLS* – a `SUNNonlinSol` object.
- *SolveFn* – a wrapper function to the SUNDIALS integrator’s linear solver solve function. See §9.1.4 for the definition of `SUNNonlinSolSolveFn`.

**Return value:**

- A `SUNErrCode`

**Notes:**

The `SUNNonlinSolSolveFn` function solves the linear system  $Ax = b$  where  $A = \frac{\partial F}{\partial y}$  is the linearization of the nonlinear residual function  $F(y) = 0$ . `SUNNonlinSol` implementations that do not require solving this system or do not use `SUNLinSol` linear solvers may set this operation to `NULL`.

`SUNErrCode` `SUNNonlinSolSetConvTestFn`(`SUNNonlinearSolver` *NLS*, `SUNNonlinSolConvTestFn` *CTestFn*, void \*ctest\_data)

This *optional* function is used to provide the nonlinear solver with a function for determining if the nonlinear solver iteration has converged. This is typically called by SUNDIALS integrators to define their nonlinear convergence criteria, but may be replaced by the user.

**Arguments:**

- *NLS* – a `SUNNonlinSol` object.
- *CTestFn* – a SUNDIALS integrator’s nonlinear solver convergence test function. See §9.1.4 for the definition of `SUNNonlinSolConvTestFn`.
- *ctest\_data* – is a data pointer passed to *CTestFn* every time it is called.

**Return value:**

- A `SUNErrCode`

**Notes:**

`SUNNonlinSol` implementations utilizing their own convergence test criteria may set this function to `NULL`.

`SUNErrCode` `SUNNonlinSolSetMaxIters`(`SUNNonlinearSolver` *NLS*, int maxiters)

This *optional* function sets the maximum number of nonlinear solver iterations. This is typically called by SUNDIALS integrators to define their default iteration limit, but may be adjusted by the user.

**Arguments:**



- *NLS* – a SUNNonlinSol object.
- *maxiters* – the maximum number of nonlinear iterations.

**Return value:**

- A *SUNErrCode*

**9.1.3 SUNNonlinearSolver “get” functions**

The following functions allow SUNDIALS integrators to retrieve nonlinear solver statistics. The routines to get the number of iterations in the most recent solve (*SUNNonlinSolGetNumIters()*) and number of convergence failures are optional. The routine to get the current nonlinear solver iteration (*SUNNonlinSolGetCurIter()*) is required when using the convergence test provided by the SUNDIALS integrator or when using an iterative SUNLinSol linear solver module; otherwise *SUNNonlinSolGetCurIter()* is optional.

*SUNErrCode* **SUNNonlinSolGetNumIters**(*SUNNonlinearSolver* NLS, long int \*niters)

This *optional* function returns the number of nonlinear solver iterations in the most recent solve. This is typically called by the SUNDIALS integrator to store the nonlinear solver statistics, but may also be called by the user.

**Arguments:**

- *NLS* – a SUNNonlinSol object.
- *niters* – the total number of nonlinear solver iterations.

**Return value:**

- A *SUNErrCode*

*SUNErrCode* **SUNNonlinSolGetCurIter**(*SUNNonlinearSolver* NLS, int \*iter)

This function returns the iteration index of the current nonlinear solve. This function is *required* when using SUNDIALS integrator-provided convergence tests or when using an iterative SUNLinSol linear solver module; otherwise it is *optional*.

**Arguments:**

- *NLS* – a SUNNonlinSol object.
- *iter* – the nonlinear solver iteration in the current solve starting from zero.

**Return value:**

- A *SUNErrCode*

*SUNErrCode* **SUNNonlinSolGetNumConvFails**(*SUNNonlinearSolver* NLS, long int \*nconvfails)

This *optional* function returns the number of nonlinear solver convergence failures in the most recent solve. This is typically called by the SUNDIALS integrator to store the nonlinear solver statistics, but may also be called by the user.

**Arguments:**

- *NLS* – a SUNNonlinSol object.
- *nconvfails* – the total number of nonlinear solver convergence failures.

**Return value:**

- A *SUNErrCode*

### 9.1.4 Functions provided by SUNDIALS integrators

To interface with SUNNonlinSol modules, the SUNDIALS integrators supply a variety of routines for evaluating the nonlinear system, calling the SUNLinSol setup and solve functions, and testing the nonlinear iteration for convergence. These integrator-provided routines translate between the user-supplied ODE or DAE systems and the generic interfaces to the nonlinear or linear systems of equations that result in their solution. The functions provided to a SUNNonlinSol module have types defined in the header file `sundials/sundials_nonlinearsolver.h`; these are also described below.

```
typedef int (*SUNNonlinSolSysFn)(N_Vector ycor, N_Vector F, void *mem)
```

These functions evaluate the nonlinear system  $F(y)$  for SUNNONLINEARSOLVER\_ROOTFIND type modules or  $G(y)$  for SUNNONLINEARSOLVER\_FIXEDPOINT type modules. Memory for  $F$  must be allocated prior to calling this function. The vector  $ycor$  will be left unchanged.

**Arguments:**

- $ycor$  – is the current correction to the predicted state at which the nonlinear system should be evaluated.
- $F$  – is the output vector containing  $F(y)$  or  $G(y)$ , depending on the solver type.
- $mem$  – is the SUNDIALS integrator memory structure.

**Return value:**

The return value is zero for a successful solve, a positive value for a recoverable error, and a negative value for an unrecoverable error.

**Notes:**

SUNDIALS integrators formulate nonlinear systems as a function of the correction to the predicted solution. On each call to the nonlinear system function the integrator will compute and store the current solution based on the input correction. Additionally, the residual will store the value of the ODE right-hand side function or DAE residual used in computing the nonlinear system. These stored values are then directly used in the integrator-supplied linear solver setup and solve functions as applicable.

```
typedef int (*SUNNonlinSolLSetupFn)(sunbooleantype jbad, sunbooleantype *jcur, void *mem)
```

These functions are wrappers to the SUNDIALS integrator's function for setting up linear solves with SUNLinSol modules.

**Arguments:**

- $jbad$  – is an input indicating whether the nonlinear solver believes that  $A$  has gone stale (SUNTRUE) or not (SUNFALSE).
- $jcur$  – is an output indicating whether the routine has updated the Jacobian  $A$  (SUNTRUE) or not (SUNFALSE).
- $mem$  – is the SUNDIALS integrator memory structure.

**Return value:**

The return value is zero for a successful solve, a positive value for a recoverable error, and a negative value for an unrecoverable error.

**Notes:**

The `SUNNonlinSolLSetupFn` function sets up the linear system  $Ax = b$  where  $A = \frac{\partial F}{\partial y}$  is the linearization of the nonlinear residual function  $F(y) = 0$  (when using SUNLinSol direct linear solvers) or calls the user-defined preconditioner setup function (when using SUNLinSol iterative linear solvers). SUNNonlinSol implementations that do not require solving this system, do not utilize SUNLinSol linear solvers, or use SUNLinSol linear solvers that do not require setup may ignore these functions.

As discussed in the description of `SUNNonlinSolSysFn`, the linear solver setup function assumes that the nonlinear system function has been called prior to the linear solver setup function as the setup will utilize saved values from the nonlinear system evaluation (e.g., the updated solution).

```
typedef int (*SUNNonlinSolLSolveFn)(N_Vector b, void *mem)
```

These functions are wrappers to the SUNDIALS integrator's function for solving linear systems with SUNLinSol modules.

**Arguments:**

- *b* – contains the right-hand side vector for the linear solve on input and the solution to the linear system on output.
- *mem* – is the SUNDIALS integrator memory structure.

**Return value:**

The return value is zero for a successful solve, a positive value for a recoverable error, and a negative value for an unrecoverable error.

**Notes:**

The *SUNNonlinSolLSolveFn* function solves the linear system  $Ax = b$  where  $A = \frac{\partial F}{\partial y}$  is the linearization of the nonlinear residual function  $F(y) = 0$ . SUNNonlinSol implementations that do not require solving this system or do not use SUNLinSol linear solvers may ignore these functions.

As discussed in the description of *SUNNonlinSolSysFn*, the linear solver solve function assumes that the nonlinear system function has been called prior to the linear solver solve function as the setup may utilize saved values from the nonlinear system evaluation (e.g., the updated solution).

```
typedef int (*SUNNonlinSolConvTestFn)(SUNNonlinearSolver NLS, N_Vector ycor, N_Vector del, sunrealtype tol, N_Vector ewt, void *ctest_data)
```

These functions are SUNDIALS integrator-specific convergence tests for nonlinear solvers and are typically supplied by each SUNDIALS integrator, but users may supply custom problem-specific versions as desired.

**Arguments:**

- *NLS* – is the SUNNonlinSol object.
- *ycor* – is the current correction (nonlinear iterate).
- *del* – is the difference between the current and prior nonlinear iterates.
- *tol* – is the nonlinear solver tolerance.
- *ewt* – is the weight vector used in computing weighted norms.
- *ctest\_data* – is the data pointer provided to *SUNNonlinSolSetConvTestFn()*.

**Return value:**

The return value of this routine will be a negative value if an unrecoverable error occurred or one of the following:

- SUN\_SUCCESS – the iteration is converged.
- SUN\_NLS\_CONTINUE – the iteration has not converged, keep iterating.
- SUN\_NLS\_CONV\_RECVR – the iteration appears to be diverging, try to recover.

**Notes:**

The tolerance passed to this routine by SUNDIALS integrators is the tolerance in a weighted root-mean-squared norm with error weight vector *ewt*. SUNNonlinSol modules utilizing their own convergence criteria may ignore these functions.

### 9.1.5 SUNNonlinearSolver return codes

The functions provided to SUNNonlinSol modules by each SUNDIALS integrator, and functions within the SUNDIALS-provided SUNNonlinSol implementations, utilize a common set of return codes shown in Table 9.1. Here, negative values correspond to non-recoverable failures, positive values to recoverable failures, and zero to a successful call.

Table 9.1: Description of the SUNNonlinearSolver return codes.

Name	Value	Description
SUN_SUCCESS	0	successful call or converged solve
SUN_NLS_CONTINUE	901	the nonlinear solver is not converged, keep iterating
SUN_NLS_CONV_RECDVR	902	the nonlinear solver appears to be diverging, try to recover

### 9.1.6 The generic SUNNonlinearSolver module

SUNDIALS integrators interact with specific SUNNonlinSol implementations through the generic SUNNonlinSol module on which all other SUNNonlinSol implementations are built. The SUNNonlinearSolver type is a pointer to a structure containing an implementation-dependent *content* field and an *ops* field.

A *SUNNonlinearSolver* is a pointer to the *\_generic\_SUNNonlinearSolver* structure:

```
typedef struct _generic_SUNNonlinearSolver *SUNNonlinearSolver
```

```
struct _generic_SUNNonlinearSolver
```

The structure defining the SUNDIALS nonlinear solver class.

void \***content**

Pointer to nonlinear solver-specific member data

*SUNNonlinearSolver\_Ops* **ops**

A virtual table of nonlinear solver operations provided by a specific implementation

*SUNContext* **sunctx**

The SUNDIALS simulation context

The virtual table structure is defined as

```
typedef struct _generic_SUNNonlinearSolver_Ops *SUNNonlinearSolver_Ops
```

```
struct _generic_SUNNonlinearSolver_Ops
```

The structure defining *SUNNonlinearSolver* operations.

*SUNNonlinearSolver\_Type* (\***gettype**)(*SUNNonlinearSolver*)

The function implementing *SUNNonlinSolGetType()*

int (\***initialize**)(*SUNNonlinearSolver*)

The function implementing *SUNNonlinSolInitialize()*

int (\***setup**)(*SUNNonlinearSolver*, *N\_Vector*, void\*)

The function implementing *SUNNonlinSolSetup()*

int (\***solve**)(*SUNNonlinearSolver*, *N\_Vector*, *N\_Vector*, *N\_Vector*, *sunrealtype*, *sunbooleantype*, void\*)

The function implementing *SUNNonlinSolSolve()*

```
int (*free)(SUNNonlinearSolver)
    The function implementing SUNNonlinSolFree()

int (*setsysfn)(SUNNonlinearSolver, SUNNonlinSolSysFn)
    The function implementing SUNNonlinSolSetSysFn()

int (*setlsetupfn)(SUNNonlinearSolver, SUNNonlinSolLSetupFn)
    The function implementing SUNNonlinSolSetLSetupFn()

int (*setlsolvefn)(SUNNonlinearSolver, SUNNonlinSolLSolveFn)
    The function implementing SUNNonlinSolSetLSolveFn()

int (*setctestfn)(SUNNonlinearSolver, SUNNonlinSolConvTestFn, void*)
    The function implementing SUNNonlinSolSetConvTestFn()

int (*setmaxiters)(SUNNonlinearSolver, int)
    The function implementing SUNNonlinSolSetMaxIters()

int (*getnumiters)(SUNNonlinearSolver, long int*)
    The function implementing SUNNonlinSolGetNumIters()

int (*getcuriter)(SUNNonlinearSolver, int*)
    The function implementing SUNNonlinSolGetCurIter()

int (*getnumconvfails)(SUNNonlinearSolver, long int*)
    The function implementing SUNNonlinSolGetNumConvFails()
```

The generic `SUNNonlinSol` module defines and implements the nonlinear solver operations defined in §9.1.1–§9.1.3. These routines are in fact only wrappers to the nonlinear solver operations provided by a particular `SUNNonlinSol` implementation, which are accessed through the `ops` field of the `SUNNonlinearSolver` structure. To illustrate this point we show below the implementation of a typical nonlinear solver operation from the generic `SUNNonlinSol` module, namely `SUNNonlinSolSolve()`, which solves the nonlinear system and returns a flag denoting a successful or failed solve:

```
int SUNNonlinSolSolve(SUNNonlinearSolver NLS,
                     N_Vector y0, N_Vector y,
                     N_Vector w, sunrealtype tol,
                     sunbooleantype callLSetup, void* mem)
{
    return((int) NLS->ops->solve(NLS, y0, y, w, tol, callLSetup, mem));
}
```

### 9.1.7 Implementing a Custom `SUNNonlinearSolver` Module

A `SUNNonlinSol` implementation *must* do the following:

- Specify the content of the `SUNNonlinSol` module.
- Define and implement the required nonlinear solver operations defined in §9.1.1–§9.1.3. Note that the names of the module routines should be unique to that implementation in order to permit using more than one `SUNNonlinSol` module (each with different `SUNNonlinearSolver` internal data representations) in the same code.
- Define and implement a user-callable constructor to create a `SUNNonlinearSolver` object.

To aid in the creation of custom `SUNNonlinearSolver` modules, the generic `SUNNonlinearSolver` module provides the utility functions `SUNNonlinSolNewEmpty()` and `SUNNonlinSolFreeEmpty()`. When used in custom `SUNNonlinearSolver` constructors these functions will ease the introduction of any new optional nonlinear solver operations to the `SUNNonlinearSolver` API by ensuring that only required operations need to be set.

**SUNNonlinearSolver SUNNonlinSolNewEmpty(SUNContext sunctx)**

This function allocates a new generic SUNNonlinearSolver object and initializes its content pointer and the function pointers in the operations structure to NULL.

**Return value:**

If successful, this function returns a SUNNonlinearSolver object. If an error occurs when allocating the object, then this routine will return NULL.

**void SUNNonlinSolFreeEmpty(SUNNonlinearSolver NLS)**

This routine frees the generic SUNNonlinearSolver object, under the assumption that any implementation-specific data that was allocated within the underlying content structure has already been freed. It will additionally test whether the ops pointer is NULL, and, if it is not, it will free it as well.

**Arguments:**

- *NLS* – a SUNNonlinearSolver object

Additionally, a SUNNonlinearSolver implementation *may* do the following:

- Define and implement additional user-callable “set” routines acting on the SUNNonlinearSolver object, e.g., for setting various configuration options to tune the performance of the nonlinear solve algorithm.
- Provide additional user-callable “get” routines acting on the SUNNonlinearSolver object, e.g., for returning various solve statistics.

## 9.2 IDAS SUNNonlinearSolver interface

As discussed in Chapter §2 each integration step requires the (approximate) solution of the nonlinear system

$$G(y_n) = F \left( t_n, y_n, h_n^{-1} \sum_{i=0}^q \alpha_{n,i} y_{n-i} \right) = 0.$$

Rather than solving this system for the new state  $y_n$  IDAS reformulates the system to solve for the correction  $y_{cor}$  to the predicted new state  $y_{pred}$  and its derivative  $\dot{y}_{pred}$  so that  $y_n = y_{pred} + y_{cor}$  and  $\dot{y}_n = \dot{y}_{pred} + h_n^{-1} \alpha_{n,0} y_{cor}$ . The nonlinear system rewritten in terms of  $y_{cor}$  is

$$G(y_{cor}) = F(t_n, y_{pred} + y_{cor}, \dot{y}_{pred} + \alpha y_{cor}) = 0. \quad (9.1)$$

where  $\alpha = h_n^{-1} \alpha_{n,0}$ .

Similarly in the forward sensitivity analysis case the nonlinear system is also reformulated in terms of the correction to the predicted sensitivities.

The nonlinear system function provided by IDAS to the nonlinear solver module internally updates the current value of the new state and its derivative based on the current correction passed to the function (as well as the sensitivities). These values are used when calling the DAE residual function and when setting up linear solves (e.g., for updating the Jacobian or preconditioner).

IDAS provides several advanced functions that will not be needed by most users, but might be useful for users who choose to provide their own implementation of the SUNNonlinearSolver API. For example, such a user might need access to the current  $y$  and  $\dot{y}$  vectors to compute Jacobian data.

**int IDAGetCurrentCj(void \*ida\_mem, sunrealtype \*cj)**

The function *IDAGetCurrentCj()* returns the scalar  $c_j$  which is proportional to the inverse of the step size ( $\alpha$  in (2.7)).

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `cj` – the value of  $c_j$ .

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The IDAS memory block is `NULL`.

int **IDAGetCurrentY**(void \*ida\_mem, *N\_Vector* \*ycur)

The function *IDAGetCurrentY()* returns the current  $y$  vector.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `y` – the current  $y$  vector.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The IDAS memory block is `NULL`.

int **IDAGetCurrentYp**(void \*ida\_mem, *N\_Vector* \*ypcur)

The function *IDAGetCurrentYp()* returns the current  $\dot{y}$  vector.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `yp` – the current  $\dot{y}$  vector.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The IDAS memory block is `NULL`.

int **IDAGetCurrentYSens**(void \*ida\_mem, *N\_Vector* \*\*yyS)

The function *IDAGetCurrentYSens()* returns the current sensitivity vector array.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `yyS` – pointer to the vector array that is set to the array of sensitivity vectors.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is `NULL`.

int **IDAGetCurrentYpSens**(void \*ida\_mem, *N\_Vector* \*\*ypS)

The function *IDAGetCurrentYpSens()* returns the derivative the current sensitivity vector array.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `ypS` – pointer to the vector array that is set to the array of sensitivity vector derivatives.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is `NULL`.

```
int IDAGetNonlinearSystemData(void *ida_mem, sunrealtype *tcur, N_Vector *yypred, N_Vector *yypred,
                             N_Vector *yyn, N_Vector *ypn, N_Vector *res, sunrealtype *cj, void
                             **user_data)
```

The function `IDAGetNonlinearSystemData()` returns all internal data required to construct the current nonlinear system (9.1).

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `tcur` – current value of the independent variable  $t_n$ .
- `yypred` – predicted value of  $y_{pred}$  at  $t_n$ .
- `yypred` – predicted value of  $\dot{y}_{pred}$  at  $t_n$ .
- `yyn` – the vector  $y_n$ . This vector may not be current and may need to be filled (see the note below).
- `ypn` – the vector  $\dot{y}_n$ . This vector may not be current and may need to be filled (see the note below).
- `res` – the residual function evaluated at the current time and state,  $F(t_n, y_n, \dot{y}_n)$ . This vector may not be current and may need to be filled (see the note below).
- `cj` – the scalar  $c_j$  which is proportional to the inverse of the step size ( $\alpha$  in (9.1)).
- `user_data` – pointer to the user-defined data structures.

**Return value:**

- `IDA_SUCCESS` – The optional output values have been successfully set.
- `IDA_MEM_NULL` – The IDAS memory block is NULL.

**Notes:**

This routine is intended for users who wish to attach a custom `SUNNonlinSolSysFn` to an existing `SUNNonlinearSolver` object (through a call to `SUNNonlinSolSetSysFn()`) or who need access to nonlinear system data to compute the nonlinear system function as part of a custom `SUNNonlinearSolver` object.

When supplying a custom `SUNNonlinSolSysFn` to an existing `SUNNonlinearSolver` object, the user should call `IDAGetNonlinearSystemData()` inside the nonlinear system function to access the requisite data for evaluating the nonlinear system function of their choosing. Additionally, if the `SUNNonlinearSolver` object (existing or custom) leverages the `SUNNonlinSolLSetupFn` and/or `SUNNonlinSolLSolveFn` functions supplied by IDAS (through calls to `SUNNonlinSolSetLSetupFn()` and `SUNNonlinSolSetLSolveFn()` respectively) the vectors `yyn` and `ypn`, and `res` must be filled in by the user's `SUNNonlinSolSysFn` with the current state and corresponding evaluation of the right-hand side function respectively i.e.,

$$\begin{aligned} yyn &= y_{pred} + y_{cor}, \\ ypn &= \dot{y}_{pred} + \alpha \dot{y}_{cor}, \\ res &= F(t_n, y_n, \dot{y}_n), \end{aligned}$$

where  $y_{cor}$  was the first argument supplied to the `SUNNonlinSolSysFn`. If this function is called as part of a custom linear solver (i.e., the default `SUNNonlinSolSysFn` is used) then the vectors `yn`, `ypn` and `res` are only current when `IDAGetNonlinearSystemData()` is called after an evaluation of the nonlinear system function.

```
int IDAGetNonlinearSystemDataSens(void *ida_mem, sunrealtype *tcur, N_Vector **yySpred, N_Vector
                                 **ypSpred, N_Vector **yySn, N_Vector **ypSn, sunrealtype *cj, void
                                 **user_data)
```

The function `IDAGetNonlinearSystemDataSens()` returns all internal sensitivity data required to construct the current nonlinear system (9.1).



**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `tcur` – current value of the independent variable  $t_n$ .
- `yySpred` – predicted value of  $y_{S_{i,pred}}$  at  $t_n$  for  $i = 0 \dots N_s - 1$ .
- `ypSpred` – predicted value of  $\dot{y}_{S_{i,pred}}$  at  $t_n$  for  $i = 0 \dots N_s - 1$ .
- `yySn` – the vectors  $y_{S_{i,n}}$ . These vectors may be not current see the note below.
- `ypSn` – the vectors  $\dot{y}_{S_{i,n}}$ . These vectors may be not current see the note below.
- `cj` – the scalar  $c_j$  which is proportional to the inverse of the step size  $\alpha$  in (2.7).
- `user_data` – pointer to the user-defined data structures

**Return value:**

- `IDA_SUCCESS` – The optional output values have been successfully set.
- `IDA_MEM_NULL` – The `ida_mem` pointer is NULL.

**Notes:**

This routine is intended for users who wish to attach a custom [SUNNonlinSolSysFn](#) to an existing `SUNNonlinearSolver` object (through a call to [SUNNonlinSolSetSysFn\(\)](#)) or who need access to nonlinear system data to compute the nonlinear system function as part of a custom `SUNNonlinearSolver` object. When supplying a custom [SUNNonlinSolSysFn](#) to an existing `SUNNonlinearSolver` object, the user should call [IDAGetNonlinearSystemDataSens\(\)](#) inside the nonlinear system function to access the requisite data for evaluating the nonlinear system function of their choosing. Additionally, if the the vectors `yySn` and `ypSn` are provided as additional workspace and do not need to be filled in by the user's [SUNNonlinSolSysFn](#). If this function is called as part of a custom linear solver (i.e., the default [SUNNonlinSolSysFn](#) is used) then the vectors `yySn` and `ypSn` are only current when [IDAGetNonlinearSystemDataSens\(\)](#) is called after an evaluation of the nonlinear system function.

int **IDAComputeY**(void \*ida\_mem, [N\\_Vector](#) ycor, [N\\_Vector](#) y)

The function computes the current  $y(t)$  vector based on the given correction vector from the nonlinear solver.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `ycor` – the correction.
- `y` – the output vector.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.
- `IDA_MEM_NULL` – The IDAS memory block is NULL.

int **IDAComputeYp**(void \*ida\_mem, [N\\_Vector](#) ycor, [N\\_Vector](#) yp)

The function computes  $\dot{y}(t)$  based on the given correction vector from the nonlinear solver.

**Arguments:**

- `ida_mem` – pointer to the IDAS memory block.
- `ycor` – the correction.
- `yp` – the output vector array.

**Return value:**

- `IDA_SUCCESS` – The optional output value has been successfully set.

- IDA\_MEM\_NULL – The IDAS memory block is NULL.

int **IDAComputeYSens**(void \*ida\_mem, *N\_Vector* \*ycorS, *N\_Vector* \*yys)

The function computes the sensitivities based on the given correction vector from the nonlinear solver.

**Arguments:**

- ida\_mem – pointer to the IDAS memory block.
- ycorS – the correction.
- yyS – the output vector array.

**Return value:**

- IDA\_SUCCESS – The optional output value has been successfully set.
- IDA\_MEM\_NULL – The ida\_mem pointer is NULL.

int **IDAComputeYpSens**(void \*ida\_mem, *N\_Vector* \*ycorS, *N\_Vector* \*ypS)

The function computes the sensitivity derivatives based on the given correction vector from the nonlinear solver.

**Arguments:**

- ida\_mem – pointer to the IDAS memory block.
- ycorS – the correction.
- ypS – the output vector array.

**Return value:**

- IDA\_SUCCESS – The optional output value has been successfully set.
- IDA\_MEM\_NULL – The ida\_mem pointer is NULL.

## 9.3 The SUNNonlinSol\_Newton implementation

This section describes the SUNNonlinSol implementation of Newton's method. To access the SUNNonlinSol\_Newton module, include the header file `sunnonlinSol/sunnonlinSol_newton.h`. We note that the SUNNonlinSol\_Newton module is accessible from SUNDIALS integrators *without* separately linking to the `libsundials_sunnonlinSol-newton` module library.

### 9.3.1 SUNNonlinSol\_Newton description

To find the solution to

$$F(y) = 0 \tag{9.2}$$

given an initial guess  $y^{(0)}$ , Newton's method computes a series of approximate solutions

$$y^{(m+1)} = y^{(m)} + \delta^{(m+1)}$$

where  $m$  is the Newton iteration index, and the Newton update  $\delta^{(m+1)}$  is the solution of the linear system

$$A(y^{(m)})\delta^{(m+1)} = -F(y^{(m)}), \tag{9.3}$$

in which  $A$  is the Jacobian matrix

$$A \equiv \partial F / \partial y. \tag{9.4}$$

Depending on the linear solver used, the `SUNNonlinSol_Newton` module will employ either a Modified Newton method or an Inexact Newton method [15, 20, 28, 30, 48]. When used with a direct linear solver, the Jacobian matrix  $A$  is held constant during the Newton iteration, resulting in a Modified Newton method. With a matrix-free iterative linear solver, the iteration is an Inexact Newton method.

In both cases, calls to the integrator-supplied `SUNNonlinSolSetupFn` function are made infrequently to amortize the increased cost of matrix operations (updating  $A$  and its factorization within direct linear solvers, or updating the preconditioner within iterative linear solvers). Specifically, `SUNNonlinSol_Newton` will call the `SUNNonlinSolSetupFn` function in two instances:

- (a) when requested by the integrator (the input `callSetSetup` is `SUNTRUE`) before attempting the Newton iteration, or
- (b) when reattempting the nonlinear solve after a recoverable failure occurs in the Newton iteration with stale Jacobian information (`jcur` is `SUNFALSE`). In this case, `SUNNonlinSol_Newton` will set `jbad` to `SUNTRUE` before calling the `SUNNonlinSolSetupFn()` function.

Whether the Jacobian matrix  $A$  is fully or partially updated depends on logic unique to each integrator-supplied `SUNNonlinSolSetupFn` routine. We refer to the discussion of nonlinear solver strategies provided in the package-specific Mathematics section of the documentation for details.

The default maximum number of iterations and the stopping criteria for the Newton iteration are supplied by the `SUNDIALS` integrator when `SUNNonlinSol_Newton` is attached to it. Both the maximum number of iterations and the convergence test function may be modified by the user by calling the `SUNNonlinSolSetMaxIters()` and/or `SUNNonlinSolSetConvTestFn()` functions after attaching the `SUNNonlinSol_Newton` object to the integrator.

### 9.3.2 SUNNonlinSol\_Newton functions

The `SUNNonlinSol_Newton` module provides the following constructor for creating the `SUNNonlinearSolver` object.

*SUNNonlinearSolver* **SUNNonlinSol\_Newton**(*N\_Vector* y, *SUNContext* sunctx)

This creates a `SUNNonlinearSolver` object for use with `SUNDIALS` integrators to solve nonlinear systems of the form  $F(y) = 0$  using Newton's method.

**Arguments:**

- *y* – a template for cloning vectors needed within the solver.
- *sunctx* – the *SUNContext* object (see §4.2)

**Return value:**

A `SUNNonlinSol` object if the constructor exits successfully, otherwise it will be `NULL`.

The `SUNNonlinSol_Newton` module implements all of the functions defined in §9.1.1–§9.1.3 except for `SUNNonlinSolSetup()`. The `SUNNonlinSol_Newton` functions have the same names as those defined by the generic `SUNNonlinSol` API with `_Newton` appended to the function name. Unless using the `SUNNonlinSol_Newton` module as a standalone nonlinear solver the generic functions defined in §9.1.1–§9.1.3 should be called in favor of the `SUNNonlinSol_Newton`-specific implementations.

The `SUNNonlinSol_Newton` module also defines the following user-callable function.

*SUNErrCode* **SUNNonlinSolGetSysFn\_Newton**(*SUNNonlinearSolver* NLS, *SUNNonlinSolSysFn* \*SysFn)

This returns the residual function that defines the nonlinear system.

**Arguments:**

- *NLS* – a `SUNNonlinSol` object.
- *SysFn* – the function defining the nonlinear system.

**Return value:**

- A [SUNErrCode](#)

**Notes:**

This function is intended for users that wish to evaluate the nonlinear residual in a custom convergence test function for the SUNNonlinSol\_Newton module. We note that SUNNonlinSol\_Newton will not leverage the results from any user calls to *SysFn*.

### 9.3.3 SUNNonlinSol\_Newton content

The *content* field of the SUNNonlinSol\_Newton module is the following structure.

```
struct _SUNNonlinearSolverContent_Newton {  
  
    SUNNonlinSolSysFn      Sys;  
    SUNNonlinSolLSetupFn   LSetup;  
    SUNNonlinSolLSolveFn   LSolve;  
    SUNNonlinSolConvTestFn CTest;  
  
    N_Vector              delta;  
    sunbooleantype         jcur;  
    int                    curiter;  
    int                    maxiters;  
    long int               niters;  
    long int               nconvfails;  
    void*                  ctest_data;  
};
```

These entries of the *content* field contain the following information:

- Sys – the function for evaluating the nonlinear system,
- LSetup – the package-supplied function for setting up the linear solver,
- LSolve – the package-supplied function for performing a linear solve,
- CTest – the function for checking convergence of the Newton iteration,
- delta – the Newton iteration update vector,
- jcur – the Jacobian status (SUNTRUE = current, SUNFALSE = stale),
- curiter – the current number of iterations in the solve attempt,
- maxiters – the maximum number of Newton iterations allowed in a solve,
- niters – the total number of nonlinear iterations across all solves,
- nconvfails – the total number of nonlinear convergence failures across all solves,
- ctest\_data – the data pointer passed to the convergence test function,

## 9.4 The SUNNonlinSol\_FixedPoint implementation

This section describes the SUNNonlinSol implementation of a fixed point (functional) iteration with optional Anderson acceleration. To access the SUNNonlinSol\_FixedPoint module, include the header file `sunnonlinSol/sunnonlinSol_fixedpoint.h`. We note that the SUNNonlinSol\_FixedPoint module is accessible from SUNDIALS integrators *without* separately linking to the `libsundials_sunnonlinSolfixedpoint` module library.

### 9.4.1 SUNNonlinSol\_FixedPoint description

To find the solution to

$$G(y) = y \quad (9.5)$$

given an initial guess  $y^{(0)}$ , the fixed point iteration computes a series of approximate solutions

$$y^{(n+1)} = G(y^{(n)}) \quad (9.6)$$

where  $n$  is the iteration index. The convergence of this iteration may be accelerated using Anderson's method [10, 33, 56, 66]. With Anderson acceleration using subspace size  $m$ , the series of approximate solutions can be formulated as the linear combination

$$y^{(n+1)} = \beta \sum_{i=0}^{m_n} \alpha_i^{(n)} G(y^{(n-m_n+i)}) + (1 - \beta) \sum_{i=0}^{m_n} \alpha_i^{(n)} y_{n-m_n+i} \quad (9.7)$$

where  $m_n = \min \{m, n\}$  and the factors

$$\alpha^{(n)} = (\alpha_0^{(n)}, \dots, \alpha_{m_n}^{(n)})$$

solve the minimization problem  $\min_{\alpha} \|F_n \alpha^T\|_2$  under the constraint that  $\sum_{i=0}^{m_n} \alpha_i = 1$  where

$$F_n = (f_{n-m_n}, \dots, f_n)$$

with  $f_i = G(y^{(i)}) - y^{(i)}$ . Due to this constraint, in the limit of  $m = 0$  the accelerated fixed point iteration formula (9.7) simplifies to the standard fixed point iteration (9.6).

Following the recommendations made in [66], the SUNNonlinSol\_FixedPoint implementation computes the series of approximate solutions as

$$y^{(n+1)} = G(y^{(n)}) - \sum_{i=0}^{m_n-1} \gamma_i^{(n)} \Delta g_{n-m_n+i} - (1 - \beta)(f(y^{(n)})) - \sum_{i=0}^{m_n-1} \gamma_i^{(n)} \Delta f_{n-m_n+i} \quad (9.8)$$

with  $\Delta g_i = G(y^{(i+1)}) - G(y^{(i)})$  and where the factors

$$\gamma^{(n)} = (\gamma_0^{(n)}, \dots, \gamma_{m_n-1}^{(n)})$$

solve the unconstrained minimization problem  $\min_{\gamma} \|f_n - \Delta F_n \gamma^T\|_2$  where

$$\Delta F_n = (\Delta f_{n-m_n}, \dots, \Delta f_{n-1}),$$

with  $\Delta f_i = f_{i+1} - f_i$ . The least-squares problem is solved by applying a QR factorization to  $\Delta F_n = Q_n R_n$  and solving  $R_n \gamma = Q_n^T f_n$ .

The acceleration subspace size  $m$  is required when constructing the SUNNonlinSol\_FixedPoint object. The default maximum number of iterations and the stopping criteria for the fixed point iteration are supplied by the SUNDIALS integrator when SUNNonlinSol\_FixedPoint is attached to it. Both the maximum number of iterations and the convergence test function may be modified by the user by calling `SUNNonlinSolSetMaxIters()` and `SUNNonlinSolSetConvTestFn()` after attaching the SUNNonlinSol\_FixedPoint object to the integrator.

## 9.4.2 SUNNonlinSol\_FixedPoint functions

The SUNNonlinSol\_FixedPoint module provides the following constructor for creating the SUNNonlinearSolver object.

*SUNNonlinearSolver* **SUNNonlinSol\_FixedPoint**(*N\_Vector* y, int m, *SUNContext* sunctx)

This creates a SUNNonlinearSolver object for use with SUNDIALS integrators to solve nonlinear systems of the form  $G(y) = y$ .

**Arguments:**

- *y* – a template for cloning vectors needed within the solver.
- *m* – the number of acceleration vectors to use.
- *sunctx* – the *SUNContext* object (see §4.2)

**Return value:**

A SUNNonlinSol object if the constructor exits successfully, otherwise it will be NULL.

Since the accelerated fixed point iteration (9.6) does not require the setup or solution of any linear systems, the SUNNonlinSol\_FixedPoint module implements all of the functions defined in §9.1.1–§9.1.3 except for the *SUNNonlinSolSetup()*, *SUNNonlinSolSetLSetupFn()*, and *SUNNonlinSolSetLSolveFn()* functions, that are set to NULL. The SUNNonlinSol\_FixedPoint functions have the same names as those defined by the generic SUNNonlinSol API with *\_FixedPoint* appended to the function name. Unless using the SUNNonlinSol\_FixedPoint module as a standalone nonlinear solver the generic functions defined in §9.1.1–§9.1.3 should be called in favor of the SUNNonlinSol\_FixedPoint-specific implementations.

The SUNNonlinSol\_FixedPoint module also defines the following user-callable functions.

*SUNErrCode* **SUNNonlinSolGetSysFn\_FixedPoint**(*SUNNonlinearSolver* NLS, *SUNNonlinSolSysFn* \*SysFn)

This returns the fixed-point function that defines the nonlinear system.

**Arguments:**

- *NLS* – a SUNNonlinSol object.
- *SysFn* – the function defining the nonlinear system.

**Return value:**

- A *SUNErrCode*

**Notes:**

This function is intended for users that wish to evaluate the fixed-point function in a custom convergence test function for the SUNNonlinSol\_FixedPoint module. We note that SUNNonlinSol\_FixedPoint will not leverage the results from any user calls to *SysFn*.

*SUNErrCode* **SUNNonlinSolSetDamping\_FixedPoint**(*SUNNonlinearSolver* NLS, *sunrealtype* beta)

This sets the damping parameter  $\beta$  to use with Anderson acceleration. By default damping is disabled i.e.,  $\beta = 1.0$ .

**Arguments:**

- *NLS* – a SUNNonlinSol object.
- *beta* – the damping parameter  $0 < \beta \leq 1$ .

**Return value:**

- A *SUNErrCode*

**Notes:**

A beta value should satisfy  $0 < \beta < 1$  if damping is to be used. A value of one or more will disable damping.

### 9.4.3 SUNNonlinSol\_FixedPoint content

The *content* field of the SUNNonlinSol\_FixedPoint module is the following structure.

```
struct _SUNNonlinearSolverContent_FixedPoint {
    SUNNonlinSolSysFn      Sys;
    SUNNonlinSolConvTestFn CTest;

    int                    m;
    int                    *imap;
    sunrealtype            *R;
    sunbooleantype         damping;
    sunrealtype            beta;
    sunrealtype            *gamma;
    sunrealtype            *cvals;
    N_Vector               *df;
    N_Vector               *dg;
    N_Vector               *q;
    N_Vector               *Xvecs;
    N_Vector               yprev;
    N_Vector               gy;
    N_Vector               fold;
    N_Vector               gold;
    N_Vector               delta;
    int                    curiter;
    int                    maxiters;
    long int               niters;
    long int               nconvfails;
    void                   *ctest_data;
};
```

The following entries of the *content* field are always allocated:

- Sys – function for evaluating the nonlinear system,
- CTest – function for checking convergence of the fixed point iteration,
- yprev – N\_Vector used to store previous fixed-point iterate,
- gy – N\_Vector used to store  $G(y)$  in fixed-point algorithm,
- delta – N\_Vector used to store difference between successive fixed-point iterates,
- curiter – the current number of iterations in the solve attempt,
- maxiters – the maximum number of fixed-point iterations allowed in a solve,
- niters – the total number of nonlinear iterations across all solves,
- nconvfails – the total number of nonlinear convergence failures across all solves,
- ctest\_data – the data pointer passed to the convergence test function,
- m – number of acceleration vectors,

If Anderson acceleration is requested (i.e.,  $m > 0$  in the call to `SUNNonlinSol_FixedPoint()`), then the following items are also allocated within the *content* field:

- `imap` – index array used in acceleration algorithm (length  $m$ ),
- `damping` – a flag indicating if damping is enabled,
- `beta` – the damping parameter,
- `R` – small matrix used in acceleration algorithm (length  $m*m$ ),
- `gamma` – small vector used in acceleration algorithm (length  $m$ ),
- `cvals` – small vector used in acceleration algorithm (length  $m+1$ ),
- `df` – array of vectors used in acceleration algorithm (length  $m$ ),
- `dg` – array of vectors used in acceleration algorithm (length  $m$ ),
- `q` – array of vectors used in acceleration algorithm (length  $m$ ),
- `Xvecs` – vector pointer array used in acceleration algorithm (length  $m+1$ ),
- `fold` – vector used in acceleration algorithm, and
- `gold` – vector used in acceleration algorithm.

## 9.5 The SUNNonlinSol\_PetscSNES implementation

This section describes the SUNNonlinSol interface to the [PETSc SNES nonlinear solver\(s\)](#). To enable the SUNNonlinSol\_PetscSNES module, SUNDIALS must be configured to use PETSc. Instructions on how to do this are given in §12.4.7. To access the SUNNonlinSol\_PetscSNES module, include the header file `sunnonlinSol/sunnonlinSol_petscsnes.h`. The library to link to is `libsundials_sunnonlinSolpetsc.lib` where `.lib` is typically `.so` for shared libraries and `.a` for static libraries. Users of the SUNNonlinSol\_PetscSNES module should also see §6.9 which discusses the NVECTOR interface to the PETSc Vec API.

### 9.5.1 SUNNonlinSol\_PetscSNES description

The SUNNonlinSol\_PetscSNES implementation allows users to utilize a PETSc SNES nonlinear solver to solve the nonlinear systems that arise in the SUNDIALS integrators. Since SNES uses the KSP linear solver interface underneath it, the SUNNonlinSol\_PetscSNES implementation does not interface with SUNDIALS linear solvers. Instead, users should set nonlinear solver options, linear solver options, and preconditioner options through the PETSc SNES, KSP, and PC APIs.

*Important usage notes for the SUNNonlinSol\_PetscSNES implementation:*

- The SUNNonlinSol\_PetscSNES implementation handles calling `SNESSetFunction` at construction. The actual residual function  $F(y)$  is set by the SUNDIALS integrator when the SUNNonlinSol\_PetscSNES object is attached to it. Therefore, a user should not call `SNESSetFunction` on a SNES object that is being used with SUNNonlinSol\_PetscSNES. For these reasons it is recommended, although not always necessary, that the user calls `SUNNonlinSol_PetscSNES()` with the new SNES object immediately after calling `SNESCreate`.
- The number of nonlinear iterations is tracked by SUNDIALS separately from the count kept by SNES. As such, the function `SUNNonlinSolGetNumIters()` reports the cumulative number of iterations across the lifetime of the `SUNNonlinearSolver` object.



- Some “converged” and “diverged” convergence reasons returned by SNES are treated as recoverable convergence failures by SUNDIALS. Therefore, the count of convergence failures returned by *SUNNonlinSolGetNumConvFails()* will reflect the number of recoverable convergence failures as determined by SUNDIALS, and may differ from the count returned by *SNESGetNonlinearStepFailures*.
- The *SUNNonlinSol\_PetscSNES* module is not currently compatible with the CVODES or IDAS staggered or simultaneous sensitivity strategies.

### 9.5.2 SUNNonlinearSolver\_PetscSNES functions

The *SUNNonlinSol\_PetscSNES* module provides the following constructor for creating a *SUNNonlinearSolver* object.

*SUNNonlinearSolver* **SUNNonlinSol\_PetscSNES**(*N\_Vector* y, SNES snes, *SUNContext* sunctx)

This creates a *SUNNonlinSol* object that wraps a PETSc SNES object for use with SUNDIALS. This will call *SNESSetFunction* on the provided SNES object.

**Arguments:**

- *snes* – a PETSc SNES object.
- *y* – a *N\_Vector* object of type *NVECTOR\_PETSC* that is used as a template for the residual vector.
- *sunctx* – the *SUNContext* object (see §4.2)

**Return value:**

A *SUNNonlinSol* object if the constructor exits successfully, otherwise it will be NULL.

**Warning:** This function calls *SNESSetFunction* and will overwrite whatever function was previously set. Users should not call *SNESSetFunction* on the SNES object provided to the constructor.

The *SUNNonlinSol\_PetscSNES* module implements all of the functions defined in §9.1.1–§9.1.3 except for *SUNNonlinSolSetup()*, *SUNNonlinSolSetLSetupFn()*, *SUNNonlinSolSetLSolveFn()*, *SUNNonlinSolSetConvTestFn()*, and *SUNNonlinSolSetMaxIters()*.

The *SUNNonlinSol\_PetscSNES* functions have the same names as those defined by the generic *SUNNonlinSol* API with *\_PetscSNES* appended to the function name. Unless using the *SUNNonlinSol\_PetscSNES* module as a standalone nonlinear solver the generic functions defined in §9.1.1–§9.1.3 should be called in favor of the *SUNNonlinSol\_PetscSNES* specific implementations.

The *SUNNonlinSol\_PetscSNES* module also defines the following user-callable functions.

*SUNErrCode* **SUNNonlinSolGetSNES\_PetscSNES**(*SUNNonlinearSolver* NLS, SNES \*snes)

This gets the SNES object that was wrapped.

**Arguments:**

- *NLS* – a *SUNNonlinSol* object.
- *snes* – a pointer to a PETSc SNES object that will be set upon return.

**Return value:**

A *SUNErrCode*

*SUNErrCode* **SUNNonlinSolGetPetscError\_PetscSNES**(*SUNNonlinearSolver* NLS, PetscErrorCode \*error)

This gets the last error code returned by the last internal call to a PETSc API function.

**Arguments:**

- *NLS* – a *SUNNonlinSol* object.

- *error* – a pointer to a PETSc error integer that will be set upon return.

**Return value:**

A *SUNErrCode*

*SUNErrCode* **SUNNonlinSolGetSysFn\_PetscSNES**(*SUNNonlinearSolver* NLS, *SUNNonlinSolSysFn* \*SysFn)

This returns the residual function that defines the nonlinear system.

**Arguments:**

- *NLS* – a *SUNNonlinSol* object.
- *SysFn* – the function defining the nonlinear system.

**Return value:**

A *SUNErrCode*

### 9.5.3 SUNNonlinearSolver\_PetscSNES content

The *content* field of the *SUNNonlinSol\_PetscSNES* module is the following structure.

```
struct _SUNNonlinearSolverContent_PetscSNES {  
    int sysfn_last_err;  
    PetscErrorCode petsc_last_err;  
    long int nconvfails;  
    long int nni;  
    void *imem;  
    SNES snes;  
    Vec r;  
    N_Vector y, f;  
    SUNNonlinSolSysFn Sys;  
};
```

These entries of the *content* field contain the following information:

- *sysfn\_last\_err* – last error returned by the system defining function,
- *petsc\_last\_err* – last error returned by PETSc,
- *nconvfails* – number of nonlinear converge failures (recoverable or not),
- *nni* – number of nonlinear iterations,
- *imem* – SUNDIALS integrator memory,
- *snes* – PETSc SNES object,
- *r* – the nonlinear residual,
- *y* – wrapper for PETSc vectors used in the system function,
- *f* – wrapper for PETSc vectors used in the system function,
- *Sys* – nonlinear system defining function.

## Chapter 10

# Tools for Memory Management

To support applications which leverage memory pools, or utilize a memory abstraction layer, SUNDIALS provides a set of utilities that we collectively refer to as the SUNMemoryHelper API. The goal of this API is to allow users to leverage operations defined by native SUNDIALS data structures while allowing the user to have finer-grained control of the memory management.

### 10.1 The SUNMemoryHelper API

This API consists of three new SUNDIALS types: *SUNMemoryType*, *SUNMemory*, and *SUNMemoryHelper*:

typedef struct *SUNMemory\_* \***SUNMemory**

The *SUNMemory* type is a pointer to the structure

struct **SUNMemory\_**

void \***ptr**;

The actual data.

*SUNMemoryType* **type**;

The data memory type.

*sunbooleantype* **own**;

A flag indicating ownership.

size\_t **bytes**;

The size of the data allocated.

*SUNMemory* **SUNMemoryNewEmpty**(*SUNContext* sunctx)

This function returns an empty *SUNMemory* object.

#### Arguments:

- *sunctx* – the *SUNContext* object.

#### Returns:

- an uninitialized *SUNMemory* object

Changed in version 7.0.0: The function signature was updated to add the *SUNContext* argument.

**enum *SUNMemoryType***

The *SUNMemoryType* type is an enumeration that defines the supported memory types:

**enumerator *SUNMEMTYPE\_HOST***

Pageable memory accessible on the host

**enumerator *SUNMEMTYPE\_PINNED***

Page-locked memory accesible on the host

**enumerator *SUNMEMTYPE\_DEVICE***

Memory accessible from the device

**enumerator *SUNMEMTYPE\_UVM***

Memory accessible from the host or device

**typedef struct *SUNMemoryHelper\_* \**SUNMemoryHelper***

The *SUNMemoryHelper* type is a pointer to the structure

**struct *SUNMemoryHelper\_***

void \***content**;

Pointer to the implementation-specific member data

*SUNMemoryHelper\_Ops* **ops**;

A virtual method table of member functions

*SUNContext* **sunctx**;

The SUNDIALS simulation context

**typedef struct *SUNMemoryHelper\_Ops\_* \**SUNMemoryHelper\_Ops***

The *SUNMemoryHelper\_Ops* type is defined as a pointer to the structure containing the function pointers to the member function implementations. This structure is define as

**struct *SUNMemoryHelper\_Ops\_***

*SUNErrCode* (\***alloc**)(*SUNMemoryHelper*, *SUNMemory* \*memptr, size\_t mem\_size, *SUNMemoryType* mem\_type, void \*queue)

The function implementing *SUNMemoryHelper\_Alloc()*

*SUNErrCode* (\***dealloc**)(*SUNMemoryHelper*, *SUNMemory* mem, void \*queue)

The function implementing *SUNMemoryHelper\_Dealloc()*

*SUNErrCode* (\***copy**)(*SUNMemoryHelper*, *SUNMemory* dst, *SUNMemory* src, size\_t mem\_size, void \*queue)

The function implementing *SUNMemoryHelper\_Copy()*

*SUNErrCode* (\***copyasync**)(*SUNMemoryHelper*, *SUNMemory* dst, *SUNMemory* src, size\_t mem\_size, void \*queue)

The function implementing *SUNMemoryHelper\_CopyAsync()*

*SUNErrCode* (\***getallocstats**)(*SUNMemoryHelper*, *SUNMemoryType* mem\_type, unsigned long \*num\_allocations, unsigned long \*num\_deallocations, size\_t \*bytes\_allocated, size\_t \*bytes\_high\_watermark)

The function implementing *SUNMemoryHelper\_GetAllocStats()*

*SUNMemoryHelper* (\***clone**)(*SUNMemoryHelper*)

The function implementing *SUNMemoryHelper\_Clone()*

*SUNErrCode* (\*destroy)(*SUNMemoryHelper*)

The function implementing *SUNMemoryHelper\_Destroy()*

### 10.1.1 Implementation defined operations

The *SUNMemory* API defines the following operations that an implementation to must define:

*SUNMemory* **SUNMemoryHelper\_Alloc**(*SUNMemoryHelper* helper, *SUNMemory* \*memptr, size\_t mem\_size, *SUNMemoryType* mem\_type, void \*queue)

Allocates a *SUNMemory* object whose ptr field is allocated for mem\_size bytes and is of type mem\_type. The new object will have ownership of ptr and will be deallocated when *SUNMemoryHelper\_Dealloc()* is called.

#### Arguments:

- helper – the *SUNMemoryHelper* object.
- memptr – pointer to the allocated *SUNMemory*.
- mem\_size – the size in bytes of the ptr.
- mem\_type – the *SUNMemoryType* of the ptr.
- queue – typically a handle for an object representing an alternate execution stream (e.g., a CUDA/HIP stream or SYCL queue), but it can also be any implementation specific data.

#### Returns:

- A new *SUNMemory* object.

*SUNErrCode* **SUNMemoryHelper\_Dealloc**(*SUNMemoryHelper* helper, *SUNMemory* mem, void \*queue)

Deallocates the mem->ptr field if it is owned by mem, and then deallocates the mem object.

#### Arguments:

- helper – the *SUNMemoryHelper* object.
- mem – the *SUNMemory* object.
- queue – typically a handle for an object representing an alternate execution stream (e.g., a CUDA/HIP stream or SYCL queue), but it can also be any implementation specific data.

#### Returns:

- A *SUNErrCode* indicating success or failure.

*SUNErrCode* **SUNMemoryHelper\_Copy**(*SUNMemoryHelper* helper, *SUNMemory* dst, *SUNMemory* src, size\_t mem\_size, void \*queue)

Synchronously copies mem\_size bytes from the the source memory to the destination memory. The copy can be across memory spaces, e.g. host to device, or within a memory space, e.g. host to host. The helper object should use the memory types of dst and src to determine the appropriate transfer type necessary.

#### Arguments:

- helper – the *SUNMemoryHelper* object.
- dst – the destination memory to copy to.
- src – the source memory to copy from.
- mem\_size – the number of bytes to copy.
- queue – typically a handle for an object representing an alternate execution stream (e.g., a CUDA/HIP stream or SYCL queue), but it can also be any implementation specific data.

**Returns:**

- A *SUNErrCode* indicating success or failure.

## 10.1.2 Utility Functions

The SUNMemoryHelper API defines the following functions which do not require a SUNMemoryHelper instance:

*SUNMemory* **SUNMemoryHelper\_Alias**(*SUNMemoryHelper* helper, *SUNMemory* mem1)

Returns a SUNMemory object whose ptr field points to the same address as mem1. The new object *will not* have ownership of ptr, therefore, it will not free ptr when *SUNMemoryHelper\_Dealloc()* is called.

**Arguments:**

- helper – a SUNMemoryHelper object.
- mem1 – a SUNMemory object.

**Returns:**

- A SUNMemory object or NULL if an error occurs.

Changed in version 7.0.0: The SUNMemoryHelper argument was added to the function signature.

*SUNMemory* **SUNMemoryHelper\_Wrap**(*SUNMemoryHelper* helper, void \*ptr, *SUNMemoryType* mem\_type)

Returns a SUNMemory object whose ptr field points to the ptr argument passed to the function. The new object *will not* have ownership of ptr, therefore, it will not free ptr when *SUNMemoryHelper\_Dealloc()* is called.

**Arguments:**

- helper – a SUNMemoryHelper object.
- ptr – the data pointer to wrap in a SUNMemory object.
- mem\_type – the SUNMemoryType of the ptr.

**Returns:**

- A SUNMemory object or NULL if an error occurs.

Changed in version 7.0.0: The SUNMemoryHelper argument was added to the function signature.

*SUNMemoryHelper* **SUNMemoryHelper\_NewEmpty**(*SUNContext* sunctx)

Returns an empty SUNMemoryHelper. This is useful for building custom SUNMemoryHelper implementations.

**Arguments:**

- helper – a SUNMemoryHelper object.

**Returns:**

- A SUNMemoryHelper object or NULL if an error occurs.

Changed in version 7.0.0: The SUNMemoryHelper argument was added to the function signature.

*SUNErrCode* **SUNMemoryHelper\_CopyOps**(*SUNMemoryHelper* src, *SUNMemoryHelper* dst)

Copies the ops field of src to the ops field of dst. This is useful for building custom SUNMemoryHelper implementations.

**Arguments:**

- src – the object to copy from.
- dst – the object to copy to.

**Returns:**

- A *SUNErrCode* indicating success or failure.

*SUNErrCode* **SUNMemoryHelper\_GetAllocStats**(*SUNMemoryHelper* helper, *SUNMemoryType* mem\_type, unsigned long \*num\_allocations, unsigned long \*num\_deallocations, size\_t \*bytes\_allocated, size\_t \*bytes\_high\_watermark)

Returns statistics about the allocations performed with the helper.

**Arguments:**

- *helper* – the *SUNMemoryHelper* object.
- *mem\_type* – the *SUNMemoryType* to get stats for.
- *num\_allocations* – (output argument) number of allocations done through the helper.
- *num\_deallocations* – (output argument) number of deallocations done through the helper.
- *bytes\_allocated* – (output argument) total number of bytes allocated through the helper at the moment this function is called.
- *bytes\_high\_watermark* – (output argument) max number of bytes allocated through the helper at any moment in the lifetime of the helper.

**Returns:**

- A *SUNErrCode* indicating success or failure.

### 10.1.3 Implementation overridable operations with defaults

In addition, the *SUNMemoryHelper* API defines the following *optionally overridable* operations which an implementation may define:

*SUNErrCode* **SUNMemoryHelper\_CopyAsync**(*SUNMemoryHelper* helper, *SUNMemory* dst, *SUNMemory* src, size\_t mem\_size, void \*queue)

Asynchronously copies *mem\_size* bytes from the the source memory to the destination memory. The copy can be across memory spaces, e.g. host to device, or within a memory space, e.g. host to host. The *helper* object should use the memory types of *dst* and *src* to determine the appropriate transfer type necessary. The *ctx* argument is used when a different execution stream needs to be provided to perform the copy in, e.g. with CUDA this would be a *cudaStream\_t*.

**Arguments:**

- *helper* – the *SUNMemoryHelper* object.
- *dst* – the destination memory to copy to.
- *src* – the source memory to copy from.
- *mem\_size* – the number of bytes to copy.
- *queue* – typically a handle for an object representing an alternate execution stream (e.g., a CUDA/HIP stream or SYCL queue), but it can also be any implementation specific data.

**Returns:**

An int flag indicating success (zero) or failure (non-zero).

---

**Note:** If this operation is not defined by the implementation, then *SUNMemoryHelper\_Copy()* will be used.

---

*SUNMemoryHelper* **SUNMemoryHelper\_Clone**(*SUNMemoryHelper* helper)

Clones the SUNMemoryHelper object itself.

**Arguments:**

- helper – the SUNMemoryHelper object to clone.

**Returns:**

- A SUNMemoryHelper object.

---

**Note:** If this operation is not defined by the implementation, then the default clone will only copy the SUNMemoryHelper\_Ops structure stored in helper->ops, and not the helper->content field.

---

*SUNErrCode* **SUNMemoryHelper\_Destroy**(*SUNMemoryHelper* helper)

Destroys (frees) the SUNMemoryHelper object itself.

**Arguments:**

- helper – the SUNMemoryHelper object to destroy.

**Returns:**

- A *SUNErrCode* indicating success or failure.

---

**Note:** If this operation is not defined by the implementation, then the default destroy will only free the helper->ops field and the helper itself. The helper->content field will not be freed.

---

### 10.1.4 Implementing a custom SUNMemoryHelper

A particular implementation of the SUNMemoryHelper API must:

- Define and implement the required operations. Note that the names of these routines should be unique to that implementation in order to permit using more than one SUNMemoryHelper module in the same code.
- Optionally, specify the *content* field of SUNMemoryHelper.
- Optionally, define and implement additional user-callable routines acting on the newly defined SUNMemoryHelper.

An example of a custom SUNMemoryHelper is given in `examples/utilities/custom_memory_helper.h`.

## 10.2 The SUNMemoryHelper\_Cuda Implementation

The SUNMemoryHelper\_Cuda module is an implementation of the SUNMemoryHelper API that interfaces to the NVIDIA [5] library. The implementation defines the constructor

*SUNMemoryHelper* **SUNMemoryHelper\_Cuda**(*SUNContext* sunctx)

Allocates and returns a SUNMemoryHelper object for handling CUDA memory if successful. Otherwise it returns NULL.



### 10.2.1 SUNMemoryHelper\_Cuda API Functions

The implementation provides the following operations defined by the SUNMemoryHelper API:

*SUNMemory* **SUNMemoryHelper\_Alloc\_Cuda**(*SUNMemoryHelper* helper, *SUNMemory* memptr, size\_t mem\_size, *SUNMemoryType* mem\_type, void \*queue)

Allocates a SUNMemory object whose ptr field is allocated for mem\_size bytes and is of type mem\_type. The new object will have ownership of ptr and will be deallocated when *SUNMemoryHelper\_Dealloc()* is called.

**Arguments:**

- helper – the SUNMemoryHelper object.
- memptr – pointer to the allocated SUNMemory.
- mem\_size – the size in bytes of the ptr.
- mem\_type – the SUNMemoryType of the ptr. Supported values are:
  - SUNMEMTYPE\_HOST – memory is allocated with a call to malloc.
  - SUNMEMTYPE\_PINNED – memory is allocated with a call to cudaMallocHost.
  - SUNMEMTYPE\_DEVICE – memory is allocated with a call to cudaMalloc.
  - SUNMEMTYPE\_UVM – memory is allocated with a call to cudaMallocManaged.
- queue – currently unused.

**Returns:**

- A new *SUNMemory* object.

*SUNErrCode* **SUNMemoryHelper\_Dealloc\_Cuda**(*SUNMemoryHelper* helper, *SUNMemory* mem, void \*queue)

Deallocates the mem->ptr field if it is owned by mem, and then deallocates the mem object.

**Arguments:**

- helper – the SUNMemoryHelper object.
- mem – the SUNMemory object.
- queue – currently unused.

**Returns:**

- A *SUNErrCode* indicating success or failure.

*SUNErrCode* **SUNMemoryHelper\_Copy\_Cuda**(*SUNMemoryHelper* helper, *SUNMemory* dst, *SUNMemory* src, size\_t mem\_size, void \*queue)

Synchronously copies mem\_size bytes from the the source memory to the destination memory. The copy can be across memory spaces, e.g. host to device, or within a memory space, e.g. host to host. The helper object will use the memory types of dst and src to determine the appropriate transfer type necessary.

**Arguments:**

- helper – the SUNMemoryHelper object.
- dst – the destination memory to copy to.
- src – the source memory to copy from.
- mem\_size – the number of bytes to copy.
- queue – currently unused.

**Returns:**

- A *SUNErrCode* indicating success or failure.

*SUNErrCode* **SUNMemoryHelper\_CopyAsync\_Cuda**(*SUNMemoryHelper* helper, *SUNMemory* dst, *SUNMemory* src, size\_t mem\_size, void \*queue)

Asynchronously copies mem\_size bytes from the the source memory to the destination memory. The copy can be across memory spaces, e.g. host to device, or within a memory space, e.g. host to host. The helper object will use the memory types of dst and src to determine the appropriate transfer type necessary.

**Arguments:**

- helper – the SUNMemoryHelper object.
- dst – the destination memory to copy to.
- src – the source memory to copy from.
- mem\_size – the number of bytes to copy.
- queue – the cudaStream\_t handle for the stream that the copy will be performed on.

**Returns:**

- A *SUNErrCode* indicating success or failure.

*SUNErrCode* **SUNMemoryHelper\_GetAllocStats\_Cuda**(*SUNMemoryHelper* helper, *SUNMemoryType* mem\_type, unsigned long \*num\_allocations, unsigned long \*num\_deallocations, size\_t \*bytes\_allocated, size\_t \*bytes\_high\_watermark)

Returns statistics about memory allocations performed with the helper.

**Arguments:**

- helper – the SUNMemoryHelper object.
- mem\_type – the SUNMemoryType to get stats for.
- num\_allocations – (output argument) number of memory allocations done through the helper.
- num\_deallocations – (output argument) number of memory deallocations done through the helper.
- bytes\_allocated – (output argument) total number of bytes allocated through the helper at the moment this function is called.
- bytes\_high\_watermark – (output argument) max number of bytes allocated through the helper at any moment in the lifetime of the helper.

**Returns:**

- A *SUNErrCode* indicating success or failure.

## 10.3 The SUNMemoryHelper\_Hip Implementation

The SUNMemoryHelper\_Hip module is an implementation of the SUNMemoryHelper API that interfaces to the AMD ROCm HIP library [2]. The implementation defines the constructor

*SUNMemoryHelper* **SUNMemoryHelper\_Hip**(*SUNContext* sunctx)

Allocates and returns a SUNMemoryHelper object for handling HIP memory if successful. Otherwise it returns NULL.

### 10.3.1 SUNMemoryHelper\_Hip API Functions

The implementation provides the following operations defined by the SUNMemoryHelper API:

*SUNMemory* **SUNMemoryHelper\_Alloc\_Hip**(*SUNMemoryHelper* helper, *SUNMemory* memptr, size\_t mem\_size, *SUNMemoryType* mem\_type, void \*queue)

Allocates a SUNMemory object whose ptr field is allocated for mem\_size bytes and is of type mem\_type. The new object will have ownership of ptr and will be deallocated when *SUNMemoryHelper\_Dealloc()* is called.

**Arguments:**

- helper – the SUNMemoryHelper object.
- memptr – pointer to the allocated SUNMemory.
- mem\_size – the size in bytes of the ptr.
- mem\_type – the SUNMemoryType of the ptr. Supported values are:
  - SUNMEMTYPE\_HOST – memory is allocated with a call to malloc.
  - SUNMEMTYPE\_PINNED – memory is allocated with a call to hipMallocHost.
  - SUNMEMTYPE\_DEVICE – memory is allocated with a call to hipMalloc.
  - SUNMEMTYPE\_UVM – memory is allocated with a call to hipMallocManaged.
- queue – currently unused.

**Returns:**

- An int flag indicating success (zero) or failure (non-zero).

*SUNErrCode* **SUNMemoryHelper\_Dealloc\_Hip**(*SUNMemoryHelper* helper, *SUNMemory* mem, void \*queue)

Deallocates the mem->ptr field if it is owned by mem, and then deallocates the mem object.

**Arguments:**

- helper – the SUNMemoryHelper object.
- mem – the SUNMemory object.

**Returns:**

- An int flag indicating success (zero) or failure (non-zero).

*SUNErrCode* **SUNMemoryHelper\_Copy\_Hip**(*SUNMemoryHelper* helper, *SUNMemory* dst, *SUNMemory* src, size\_t mem\_size, void \*queue)

Synchronously copies mem\_size bytes from the the source memory to the destination memory. The copy can be across memory spaces, e.g. host to device, or within a memory space, e.g. host to host. The helper object will use the memory types of dst and src to determine the appropriate transfer type necessary.

**Arguments:**

- helper – the SUNMemoryHelper object.
- dst – the destination memory to copy to.
- src – the source memory to copy from.
- mem\_size – the number of bytes to copy.

**Returns:**

- An int flag indicating success (zero) or failure (non-zero).

*SUNErrCode* **SUNMemoryHelper\_CopyAsync\_Hip**(*SUNMemoryHelper* helper, *SUNMemory* dst, *SUNMemory* src, size\_t mem\_size, void \*queue)

Asynchronously copies mem\_size bytes from the the source memory to the destination memory. The copy can be across memory spaces, e.g. host to device, or within a memory space, e.g. host to host. The helper object will use the memory types of dst and src to determine the appropriate transfer type necessary.

**Arguments:**

- helper – the SUNMemoryHelper object.
- dst – the destination memory to copy to.
- src – the source memory to copy from.
- mem\_size – the number of bytes to copy.
- queue – the hipStream\_t handle for the stream that the copy will be performed on.

**Returns:**

- An int flag indicating success (zero) or failure (non-zero).

*SUNErrCode* **SUNMemoryHelper\_GetAllocStats\_Hip**(*SUNMemoryHelper* helper, *SUNMemoryType* mem\_type, unsigned long \*num\_allocations, unsigned long \*num\_deallocations, size\_t \*bytes\_allocated, size\_t \*bytes\_high\_watermark)

Returns statistics about memory allocations performed with the helper.

**Arguments:**

- helper – the SUNMemoryHelper object.
- mem\_type – the SUNMemoryType to get stats for.
- num\_allocations – (output argument) number of memory allocations done through the helper.
- num\_deallocations – (output argument) number of memory deallocations done through the helper.
- bytes\_allocated – (output argument) total number of bytes allocated through the helper at the moment this function is called.
- bytes\_high\_watermark – (output argument) max number of bytes allocated through the helper at any moment in the lifetime of the helper.

**Returns:**

- An int flag indicating success (zero) or failure (non-zero).

## 10.4 The SUNMemoryHelper\_Sycl Implementation

The SUNMemoryHelper\_Sycl module is an implementation of the SUNMemoryHelper API that interfaces to the SYCL abstraction layer. The implementation defines the constructor

*SUNMemoryHelper* **SUNMemoryHelper\_Sycl**(*SUNContext* sunctx)

Allocates and returns a SUNMemoryHelper object for handling SYCL memory using the provided queue. Otherwise it returns NULL.

### 10.4.1 SUNMemoryHelper\_Sycl API Functions

The implementation provides the following operations defined by the SUNMemoryHelper API:

*SUNMemory* **SUNMemoryHelper\_Alloc\_Sycl**(*SUNMemoryHelper* helper, *SUNMemory* memptr, size\_t mem\_size, *SUNMemoryType* mem\_type, void \*queue)

Allocates a SUNMemory object whose ptr field is allocated for mem\_size bytes and is of type mem\_type. The new object will have ownership of ptr and will be deallocated when *SUNMemoryHelper\_Dealloc()* is called.

**Arguments:**

- helper – the SUNMemoryHelper object.
- memptr – pointer to the allocated SUNMemory.
- mem\_size – the size in bytes of the ptr.
- mem\_type – the SUNMemoryType of the ptr. Supported values are:
  - SUNMEMTYPE\_HOST – memory is allocated with a call to malloc.
  - SUNMEMTYPE\_PINNED – memory is allocated with a call to `sycl::malloc_host`.
  - SUNMEMTYPE\_DEVICE – memory is allocated with a call to `sycl::malloc_device`.
  - SUNMEMTYPE\_UVM – memory is allocated with a call to `sycl::malloc_shared`.
- queue – the `sycl::queue` handle for the stream that the allocation will be performed on.

**Returns:**

- A new *SUNMemory* object.

*SUNErrCode* **SUNMemoryHelper\_Dealloc\_Sycl**(*SUNMemoryHelper* helper, *SUNMemory* mem, void \*queue)

Deallocates the mem->ptr field if it is owned by mem, and then deallocates the mem object.

**Arguments:**

- helper – the SUNMemoryHelper object.
- mem – the SUNMemory object.
- queue – the `sycl::queue` handle for the queue that the deallocation will be performed on.

**Returns:**

- A *SUNErrCode* indicating success or failure.

*SUNErrCode* **SUNMemoryHelper\_Copy\_Sycl**(*SUNMemoryHelper* helper, *SUNMemory* dst, *SUNMemory* src, size\_t mem\_size, void \*queue)

Synchronously copies mem\_size bytes from the the source memory to the destination memory. The copy can be across memory spaces, e.g. host to device, or within a memory space, e.g. host to host. The helper object will use the memory types of dst and src to determine the appropriate transfer type necessary.

**Arguments:**

- helper – the SUNMemoryHelper object.
- dst – the destination memory to copy to.
- src – the source memory to copy from.
- mem\_size – the number of bytes to copy.
- queue – the `sycl::queue` handle for the queue that the copy will be performed on.

**Returns:**

- A *SUNErrCode* indicating success or failure.

*SUNErrCode* **SUNMemoryHelper\_CopyAsync\_Sycl**(*SUNMemoryHelper* helper, *SUNMemory* dst, *SUNMemory* src, size\_t mem\_size, void \*queue)

Asynchronously copies mem\_size bytes from the the source memory to the destination memory. The copy can be across memory spaces, e.g. host to device, or within a memory space, e.g. host to host. The helper object will use the memory types of dst and src to determine the appropriate transfer type necessary.

**Arguments:**

- helper – the SUNMemoryHelper object.
- dst – the destination memory to copy to.
- src – the source memory to copy from.
- mem\_size – the number of bytes to copy.
- queue – the `sycl::queue` handle for the queue that the copy will be performed on.

**Returns:**

- A *SUNErrCode* indicating success or failure.

*SUNErrCode* **SUNMemoryHelper\_GetAllocStats\_Sycl**(*SUNMemoryHelper* helper, *SUNMemoryType* mem\_type, unsigned long \*num\_allocations, unsigned long \*num\_deallocations, size\_t \*bytes\_allocated, size\_t \*bytes\_high\_watermark)

Returns statistics about memory allocations performed with the helper.

**Arguments:**

- helper – the SUNMemoryHelper object.
- mem\_type – the SUNMemoryType to get stats for.
- num\_allocations – (output argument) number of memory allocations done through the helper.
- num\_deallocations – (output argument) number of memory deallocations done through the helper.
- bytes\_allocated – (output argument) total number of bytes allocated through the helper at the moment this function is called.
- bytes\_high\_watermark – (output argument) max number of bytes allocated through the helper at any moment in the lifetime of the helper.

**Returns:**

- A *SUNErrCode* indicating success or failure.

## Chapter 11

# Acquiring SUNDIALS

There are two supported ways for building and installing SUNDIALS from source. One option is to use the [Spack HPC package manager](#):

```
spack install sundials
```

The second supported option for building and installing SUNDIALS is with CMake. Before proceeding with CMake, the source code must be downloaded. This can be done by cloning the [SUNDIALS GitHub repository](https://github.com/LLNL/sundials) (run `git clone https://github.com/LLNL/sundials`), or by downloading the SUNDIALS release compressed archives (`.tar.gz`) from the [SUNDIALS website](#).

The compressed archives allow for downloading of individual SUNDIALS packages. The name of the distribution archive is of the form `SOLVER-7.1.0.tar.gz`, where `SOLVER` is one of: `sundials`, `cvode`, `cvodes`, `arkode`, `ida`, `idas`, or `kinsol`, and `7.1.0` represents the version number (of the SUNDIALS suite or of the individual solver). After downloading the relevant archives, uncompress and expand the sources, by running

```
% tar -zxvf SOLVER-7.1.0.tar.gz
```

This will extract source files under a directory `SOLVER-7.1.0`.

Starting with version 2.6.0 of SUNDIALS, CMake is the only supported method of installation. The explanations of the installation procedure begin with a few common observations:

1. The remainder of this chapter will follow these conventions:

`SOLVERDIR` is the directory `SOLVER-7.1.0` created above; i.e. the directory containing the SUNDIALS sources.

`BUILDDIR` is the (temporary) directory under which SUNDIALS is built.

`INSTDIR` is the directory under which the SUNDIALS exported header files and libraries will be installed. Typically, header files are exported under a directory `INSTDIR/include` while libraries are installed under `INSTDIR/lib`, with `INSTDIR` specified at configuration time.

2. For SUNDIALS' CMake-based installation, in-source builds are prohibited; in other words, the build directory `BUILDDIR` can **not** be the same as `SOLVERDIR` and such an attempt will lead to an error. This prevents “polluting” the source tree and allows efficient builds for different configurations and/or options.
3. The installation directory `INSTDIR` can not be the same as the source directory `SOLVERDIR`.
4. By default, only the libraries and header files are exported to the installation directory `INSTDIR`. If enabled by the user (with the appropriate toggle for CMake), the examples distributed with SUNDIALS will be built together with the solver libraries but the installation step will result in exporting (by default in a subdirectory of the installation directory) the example sources and sample outputs together with automatically generated configuration files that reference the *installed* SUNDIALS headers and libraries. As such, these configuration files for

the SUNDIALS examples can be used as “templates” for your own problems. CMake installs `CMakeLists.txt` files and also (as an option available only under Unix/Linux) `Makefile` files. Note this installation approach also allows the option of building the SUNDIALS examples without having to install them. (This can be used as a sanity check for the freshly built libraries.)

Further details on the CMake-based installation procedures, instructions for manual compilation, and a roadmap of the resulting installed libraries and exported header files, are provided in §12 and §12.8.



## Chapter 12

# Building and Installing with CMake

CMake-based installation provides a platform-independent build system. CMake can generate Unix and Linux Makefiles, as well as KDevelop, Visual Studio, and (Apple) XCode project files from the same configuration file. In addition, CMake also provides a GUI front end and which allows an interactive build and installation process.

The SUNDIALS build process requires CMake version 3.18.0 or higher and a working C compiler. On Unix-like operating systems, it also requires Make (and `curses`, including its development libraries, for the GUI front end to CMake, `ccmake` or `cmake-gui`), while on Windows it requires Visual Studio. While many Linux distributions offer CMake, the version included may be out of date. CMake adds new features regularly, and you should download the latest version from <http://www.cmake.org>. Build instructions for CMake (only necessary for Unix-like systems) can be found on the CMake website. Once CMake is installed, Linux/Unix users will be able to use `ccmake` or `cmake-gui` (depending on the version of CMake), while Windows users will be able to use `CMakeSetup`.

As previously noted, when using CMake to configure, build and install SUNDIALS, it is always required to use a separate build directory. While in-source builds are possible, they are explicitly prohibited by the SUNDIALS CMake scripts (one of the reasons being that, unlike autotools, CMake does not provide a `make distclean` procedure and it is therefore difficult to clean-up the source tree after an in-source build). By ensuring a separate build directory, it is an easy task for the user to clean-up all traces of the build by simply removing the build directory. CMake does generate a `make clean` which will remove files generated by the compiler and linker.

### 12.1 Configuring, building, and installing on Unix-like systems

The default CMake configuration will build all included solvers and associated examples and will build static and shared libraries. The `INSTDIR` defaults to `/usr/local` and can be changed by setting the `CMAKE_INSTALL_PREFIX` variable. Support for FORTRAN and all other options are disabled.

CMake can be used from the command line with the `cmake` command, or from a `curses`-based GUI by using the `ccmake` command, or from a `wxWidgets` or `QT` based GUI by using the `cmake-gui` command. Examples for using both text and graphical methods will be presented. For the examples shown it is assumed that there is a top level SUNDIALS directory with appropriate source, build and install directories:

```
$ mkdir (...)/INSTDIR
$ mkdir (...)/BUILDDIR
$ cd (...)/BUILDDIR
```

### 12.1.1 Building with the GUI

Using CMake with the `ccmake` GUI follows the general process:

1. Select and modify values, run configure (c key)
2. New values are denoted with an asterisk
3. To set a variable, move the cursor to the variable and press enter
  - If it is a boolean (ON/OFF) it will toggle the value
  - If it is string or file, it will allow editing of the string
  - For file and directories, the <tab> key can be used to complete
4. Repeat until all values are set as desired and the generate option is available (g key)
5. Some variables (advanced variables) are not visible right away; to see advanced variables, toggle to advanced mode (t key)
6. To search for a variable press the / key, and to repeat the search, press the n key

Using CMake with the `cmake-gui` GUI follows a similar process:

1. Select and modify values, click **Configure**
2. The first time you click **Configure**, make sure to pick the appropriate generator (the following will assume generation of Unix Makefiles).
3. New values are highlighted in red
4. To set a variable, click on or move the cursor to the variable and press enter
  - If it is a boolean (ON/OFF) it will check/uncheck the box
  - If it is string or file, it will allow editing of the string. Additionally, an ellipsis button will appear ... on the far right of the entry. Clicking this button will bring up the file or directory selection dialog.
  - For files and directories, the <tab> key can be used to complete
5. Repeat until all values are set as desired and click the **Generate** button
6. Some variables (advanced variables) are not visible right away; to see advanced variables, click the advanced button

To build the default configuration using the curses GUI, from the `BUILDDIR` enter the `ccmake` command and point to the `SOLVERDIR`:

```
$ ccmake (...) /SOLVERDIR
```

Similarly, to build the default configuration using the wxWidgets GUI, from the `BUILDDIR` enter the `cmake-gui` command and point to the `SOLVERDIR`:

```
$ cmake-gui (...) /SOLVERDIR
```

The default curses configuration screen is shown in the following figure.

The default `INSTDIR` for both `SUNDIALS` and the corresponding examples can be changed by setting the `CMAKE_INSTALL_PREFIX` and the `EXAMPLES_INSTALL_PATH` as shown in the following figure.

Pressing the g key or clicking **generate** will generate Makefiles including all dependencies and all rules to build `SUNDIALS` on this system. Back at the command prompt, you can now run:

```
$ make
```

```

Page 1 of 1
BUILD_ARKODE          *ON
BUILD_CVODE           *ON
BUILD_CVODES          *ON
BUILD_EXAMPLES        *ON
BUILD_IDA              *ON
BUILD_IDAS             *ON
BUILD_KINSOL           *ON
BUILD_SHARED_LIBS      *ON
BUILD_STATIC_LIBS      *ON
BUILD_TESTING          *ON
CMAKE_BUILD_TYPE       *
CMAKE_CXX_COMPILER      */usr/bin/c++
CMAKE_CXX_FLAGS         *
CMAKE_C_COMPILER        */usr/bin/cc
CMAKE_C_FLAGS           *
CMAKE_INSTALL_LIBDIR    */lib64
CMAKE_INSTALL_PREFIX    */usr/local
ENABLE_CUDA             *OFF
ENABLE_FORTRAN          *OFF
ENABLE_HYPRE            *OFF
ENABLE_KLU              *OFF
ENABLE_LAPACK           *OFF
ENABLE_MPI              *OFF
ENABLE_OPENMP           *OFF
ENABLE_OPENMP_DEVICE    *OFF
ENABLE_PETSC            *OFF
ENABLE_PTHREAD          *OFF
ENABLE_RAJA             *OFF
ENABLE_SUPERLUDIST      *OFF
ENABLE_SUPERLUMT        *OFF
ENABLE_TRILINOS         *OFF
EXAMPLES_ENABLE_C        *ON
EXAMPLES_ENABLE_CXX      *ON
EXAMPLES_INSTALL        *ON
EXAMPLES_INSTALL_PATH   */usr/local/examples
SUNDIALS_BUILD_WITH_MONITORING *OFF
SUNDIALS_INDEX_SIZE     *64
SUNDIALS_PRECISION      *DOUBLE
USE_GENERIC_MATH         *ON
USE_XSDK_DEFAULTS       *OFF

BUILD_ARKODE: Build the ARKODE library
Press [enter] to edit option Press [d] to delete an entry
Press [c] to configure
Press [h] for help          Press [q] to quit without generating
Press [t] to toggle advanced mode (Currently Off)
CMake Version 3.12.1

```

Fig. 12.1: Default configuration screen. Note: Initial screen is empty. To get this default configuration, press ‘c’ repeatedly (accepting default values denoted with asterisk) until the ‘g’ option is available.

```

Page 1 of 1
BUILD_ARKODE          *ON
BUILD_CVODE           *ON
BUILD_CVODES          *ON
BUILD_EXAMPLES        *ON
BUILD_IDA              *ON
BUILD_IDAS             *ON
BUILD_KINSOL           *ON
BUILD_SHARED_LIBS      *ON
BUILD_STATIC_LIBS      *ON
BUILD_TESTING          *ON
CMAKE_BUILD_TYPE       *
CMAKE_CXX_COMPILER      */usr/bin/c++
CMAKE_CXX_FLAGS         *
CMAKE_C_COMPILER        */usr/bin/cc
CMAKE_C_FLAGS           *
CMAKE_INSTALL_LIBDIR    *lib64
CMAKE_INSTALL_PREFIX    */usr/casc/sundials/instdir
ENABLE_CUDA             *OFF
ENABLE_FORTRAN          *OFF
ENABLE_HYPRE            *OFF
ENABLE_KLU              *OFF
ENABLE_LAPACK           *OFF
ENABLE_MPI              *OFF
ENABLE_OPENMP           *OFF
ENABLE_OPENMP_DEVICE    *OFF
ENABLE_PETSC            *OFF
ENABLE_PTHREAD          *OFF
ENABLE_RAJA             *OFF
ENABLE_SUPERLUDIST      *OFF
ENABLE_SUPERLUMT        *OFF
ENABLE_TRILINOS         *OFF
EXAMPLES_ENABLE_C       *ON
EXAMPLES_ENABLE_CXX     *ON
EXAMPLES_INSTALL        *ON
EXAMPLES_INSTALL_PATH   */usr/casc/sundials/instdir/examples
SUNDIALS_BUILD_WITH_MONITORING *OFF
SUNDIALS_INDEX_SIZE     *64
SUNDIALS_PRECISION      *DOUBLE
USE_GENERIC_MATH        *ON
USE_XSDK_DEFAULTS       *OFF

EXAMPLES_INSTALL_PATH: Output directory for installing example files
Press [enter] to edit option Press [d] to delete an entry
Press [c] to configure
Press [h] for help          Press [q] to quit without generating
Press [t] to toggle advanced mode (Currently Off)
CMake Version 3.12.1

```

Fig. 12.2: Changing the INSTDIR for SUNDIALS and corresponding EXAMPLES.

or for a faster parallel build (e.g. using 4 threads), you can run

```
$ make -j 4
```

To install SUNDIALS in the installation directory specified in the configuration, simply run:

```
$ make install
```

### 12.1.2 Building from the command line

Using CMake from the command line is simply a matter of specifying CMake variable settings with the `cmake` command. The following will build the default configuration:

```
$ cmake -DCMAKE_INSTALL_PREFIX=/home/myname/sundials/instdir \
> -DEXAMPLES_INSTALL_PATH=/home/myname/sundials/instdir/examples \
> ../srcdir
$ make
$ make install
```

## 12.2 Configuration options

A complete list of all available options for a CMake-based SUNDIALS configuration is provide below. Note that the default values shown are for a typical configuration on a Linux system and are provided as illustration only.

### **BUILD\_ARKODE**

Build the ARKODE library

Default: ON

### **BUILD\_CVODE**

Build the CVODE library

Default: ON

### **BUILD\_CVODES**

Build the CVODES library

Default: ON

### **BUILD\_IDA**

Build the IDA library

Default: ON

### **BUILD\_IDAS**

Build the IDAS library

Default: ON

### **BUILD\_KINSOL**

Build the KINSOL library

Default: ON

#### **BUILD\_SHARED\_LIBS**

Build shared libraries

Default: ON

#### **BUILD\_STATIC\_LIBS**

Build static libraries

Default: ON

#### **CMAKE\_BUILD\_TYPE**

Choose the type of build, options are: Debug, Release, RelWithDebInfo, and MinSizeRel

Default: RelWithDebInfo

---

**Note:** Specifying a build type will trigger the corresponding build type specific compiler flag options below which will be appended to the flags set by CMAKE\_<language>\_FLAGS.

---

#### **CMAKE\_C\_COMPILER**

C compiler

Default: /usr/bin/cc

#### **CMAKE\_C\_FLAGS**

Flags for C compiler

Default:

#### **CMAKE\_C\_FLAGS\_DEBUG**

Flags used by the C compiler during debug builds

Default: -g

#### **CMAKE\_C\_FLAGS\_MINSIZEREL**

Flags used by the C compiler during release minsize builds

Default: -Os -DNDEBUG

#### **CMAKE\_C\_FLAGS\_RELEASE**

Flags used by the C compiler during release builds

Default: -O3 -DNDEBUG

#### **CMAKE\_C\_STANDARD**

The C standard to build C parts of SUNDIALS with.

Default: 99

Options: 99, 11, 17.

#### **CMAKE\_C\_EXTENSIONS**

Enable compiler specific C extensions.

Default: OFF

#### **CMAKE\_CXX\_COMPILER**

C++ compiler

Default: /usr/bin/c++

---

**Note:** A C++ compiler is only required when a feature requiring C++ is enabled (e.g., CUDA, HIP, SYCL, RAJA, etc.) or the C++ examples are enabled.

All SUNDIALS solvers can be used from C++ applications without setting any additional configuration options.

---

**CMAKE\_CXX\_FLAGS**

Flags for C++ compiler

Default:

**CMAKE\_CXX\_FLAGS\_DEBUG**

Flags used by the C++ compiler during debug builds

Default: -g

**CMAKE\_CXX\_FLAGS\_MINSIZEREL**

Flags used by the C++ compiler during release minsize builds

Default: -Os -DNDEBUG

**CMAKE\_CXX\_FLAGS\_RELEASE**

Flags used by the C++ compiler during release builds

Default: -O3 -DNDEBUG

**CMAKE\_CXX\_STANDARD**

The C++ standard to build C++ parts of SUNDIALS with.

Default: 11

Options: 98, 11, 14, 17, 20.

**CMAKE\_CXX\_EXTENSIONS**

Enable compiler specific C++ extensions.

Default: OFF

**CMAKE\_Fortran\_COMPILER**

Fortran compiler

Default: /usr/bin/gfortran

---

**Note:** Fortran support (and all related options) are triggered only if either Fortran-C support (BUILD\_FORTRAN\_MODULE\_INTERFACE) or LAPACK (ENABLE\_LAPACK) support is enabled.

---

**CMAKE\_Fortran\_FLAGS**

Flags for Fortran compiler

Default:

**CMAKE\_Fortran\_FLAGS\_DEBUG**

Flags used by the Fortran compiler during debug builds

Default: -g

**CMAKE\_Fortran\_FLAGS\_MINSIZEREL**

Flags used by the Fortran compiler during release minsize builds

Default: -Os

#### **CMAKE\_Fortran\_FLAGS\_RELEASE**

Flags used by the Fortran compiler during release builds

Default: -O3

#### **CMAKE\_INSTALL\_LIBDIR**

The directory under which libraries will be installed.

Default: Set based on the system: lib, lib64, or lib/<multiarch-tuple>

#### **CMAKE\_INSTALL\_PREFIX**

Install path prefix, prepended onto install directories

Default: /usr/local

---

**Note:** The user must have write access to the location specified through this option. Exported SUNDIALS header files and libraries will be installed under subdirectories `include` and `lib` of `CMAKE_INSTALL_PREFIX`, respectively.

---

#### **ENABLE\_CUDA**

Build the SUNDIALS CUDA modules.

Default: OFF

#### **CMAKE\_CUDA\_ARCHITECTURES**

Specifies the CUDA architecture to compile for.

Default: sm\_30

#### **EXAMPLES\_ENABLE\_C**

Build the SUNDIALS C examples

Default: ON

#### **EXAMPLES\_ENABLE\_CXX**

Build the SUNDIALS C++ examples

Default: OFF

#### **EXAMPLES\_ENABLE\_CUDA**

Build the SUNDIALS CUDA examples

Default: OFF

---

**Note:** You need to enable CUDA support to build these examples.

---

#### **EXAMPLES\_ENABLE\_F2003**

Build the SUNDIALS Fortran2003 examples

Default: ON (if `BUILD_FORTRAN_MODULE_INTERFACE` is ON)

#### **EXAMPLES\_INSTALL**

Install example files

Default: ON

---

**Note:** This option is triggered when any of the SUNDIALS example programs are enabled (`EXAMPLES_ENABLE_<language>` is ON). If the user requires installation of example programs then the sources and sample

---



output files for all SUNDIALS modules that are currently enabled will be exported to the directory specified by `EXAMPLES_INSTALL_PATH`. A CMake configuration script will also be automatically generated and exported to the same directory. Additionally, if the configuration is done under a Unix-like system, makefiles for the compilation of the example programs (using the installed SUNDIALS libraries) will be automatically generated and exported to the directory specified by `EXAMPLES_INSTALL_PATH`.

---

#### **EXAMPLES\_INSTALL\_PATH**

Output directory for installing example files

Default: `/usr/local/examples`

---

**Note:** The actual default value for this option will be an `examples` subdirectory created under `CMAKE_INSTALL_PREFIX`.

---

#### **BUILD\_FORTRAN\_MODULE\_INTERFACE**

Enable Fortran 2003 interface

Default: OFF

**Warning:** There is a known issue with MSYS/gfortran and SUNDIALS shared libraries that causes linking the Fortran interfaces to fail when building SUNDIALS. For now the work around is to only build with static libraries when using MSYS with gfortran on Windows.

#### **SUNDIALS\_LOGGING\_LEVEL**

Set the maximum logging level for the SUNLogger runtime API. The higher this is set, the more output that may be logged, and the more performance may degrade. The options are:

- 0 – no logging
- 1 – log errors
- 2 – log errors + warnings
- 3 – log errors + warnings + informational output
- 4 – log errors + warnings + informational output + debug output
- 5 – log all of the above and even more (e.g. vector valued variables may be logged)

Default: 2

#### **SUNDIALS\_BUILD\_WITH\_MONITORING**

Build SUNDIALS with capabilities for fine-grained monitoring of solver progress and statistics. This is primarily useful for debugging.

Default: OFF

**Warning:** Building with monitoring may result in minor performance degradation even if monitoring is not utilized.

#### **SUNDIALS\_BUILD\_WITH\_PROFILING**

Build SUNDIALS with capabilities for fine-grained profiling. This requires POSIX timers or the Windows `profileapi.h` timers.

Default: OFF

**Warning:** Profiling will impact performance, and should be enabled judiciously.

#### **SUNDIALS\_ENABLE\_ERROR\_CHECKS**

Build SUNDIALS with more extensive checks for unrecoverable errors.

Default: OFF when CMAKE\_BUILD\_TYPE=Release|RelWithDebInfo `` and ``ON otherwise.

**Warning:** Error checks will impact performance, but can be helpful for debugging.

#### **SUNDIALS\_ENABLE\_EXTERNAL\_ADDONS**

Build SUNDIALS with any external addons that you have put in sundials/external.

Default: OFF

**Warning:** Addons are not maintained by the SUNDIALS team. Use at your own risk.

#### **ENABLE\_GINKGO**

Enable interfaces to the Ginkgo linear algebra library.

Default: OFF

#### **Ginkgo\_DIR**

Path to the Ginkgo installation.

Default: None

#### **SUNDIALS\_GINKGO\_BACKENDS**

Semi-colon separated list of Ginkgo target architectures/executors to build for. Options currently supported are REF (the Ginkgo reference executor), OMP, CUDA, HIP, and SYCL.

Default: "REF;OMP"

#### **ENABLE\_KOKKOS**

Enable the Kokkos based vector.

Default: OFF

#### **Kokkos\_DIR**

Path to the Kokkos installation.

Default: None

#### **ENABLE\_KOKKOS\_KERNELS**

Enable the Kokkos based dense matrix and linear solver.

Default: OFF

#### **KokkosKernels\_DIR**

Path to the Kokkos-Kernels installation.

Default: None

#### **ENABLE\_HIP**

Enable HIP Support

Default: OFF

**AMDGPU\_TARGETS**

Specify which AMDGPU processor(s) to target.

Default: None

**ENABLE\_HYPRE**

Flag to enable *hypre* support

Default: OFF

---

**Note:** See additional information on building with *hypre* enabled in §12.4.

---

**HYPRE\_INCLUDE\_DIR**

Path to *hypre* header files

Default: none

**HYPRE\_LIBRARY**

Path to *hypre* installed library files

Default: none

**ENABLE\_KLU**

Enable KLU support

Default: OFF

---

**Note:** See additional information on building with KLU enabled in §12.4.

---

**KLU\_INCLUDE\_DIR**

Path to SuiteSparse header files

Default: none

**KLU\_LIBRARY\_DIR**

Path to SuiteSparse installed library files

Default: none

**ENABLE\_LAPACK**

Enable LAPACK support

Default: OFF

---

**Note:** Setting this option to ON will trigger additional CMake options. See additional information on building with LAPACK enabled in §12.4.

---

**LAPACK\_LIBRARIES**

LAPACK (and BLAS) libraries

Default: /usr/lib/liblapack.so;/usr/lib/libblas.so

---

**Note:** CMake will search for libraries in your LD\_LIBRARY\_PATH prior to searching default system paths.

---

#### **ENABLE\_MAGMA**

Enable MAGMA support.

Default: OFF

---

**Note:** Setting this option to ON will trigger additional options related to MAGMA.

---

#### **MAGMA\_DIR**

Path to the root of a MAGMA installation.

Default: none

#### **SUNDIALS\_MAGMA\_BACKENDS**

Which MAGMA backend to use under the SUNDIALS MAGMA interface.

Default: CUDA

#### **ENABLE\_MPI**

Enable MPI support. This will build the parallel nvector and the MPI-aware version of the ManyVector library.

Default: OFF

---

**Note:** Setting this option to ON will trigger several additional options related to MPI.

---

#### **MPI\_C\_COMPILER**

mpicc program

Default:

#### **MPI\_CXX\_COMPILER**

mpicxx program

Default:

---

**Note:** This option is triggered only if MPI is enabled (ENABLE\_MPI is ON) and C++ examples are enabled (EXAMPLES\_ENABLE\_CXX is ON). All SUNDIALS solvers can be used from C++ MPI applications by default without setting any additional configuration options other than ENABLE\_MPI.

---

#### **MPI\_Fortran\_COMPILER**

mpi f90 program

Default:

---

**Note:** This option is triggered only if MPI is enabled (ENABLE\_MPI is ON) and Fortran-C support is enabled (EXAMPLES\_ENABLE\_F2003 is ON).

---

#### **MPIEXEC\_EXECUTABLE**

Specify the executable for running MPI programs

Default: mpirun

---

**Note:** This option is triggered only if MPI is enabled (ENABLE\_MPI is ON).

---

**ENABLE\_ONEMKL**

Enable oneMKL support.

Default: OFF

**ONEMKL\_DIR**

Path to oneMKL installation.

Default: none

**SUNDIALS\_ONEMKL\_USE\_GETRF\_LOOP**

This advanced debugging option replaces the batched LU factorization with a loop over each system in the batch and a non-batched LU factorization.

Default: OFF

**SUNDIALS\_ONEMKL\_USE\_GETRS\_LOOP**

This advanced debugging option replaces the batched LU solve with a loop over each system in the batch and a non-batched solve.

Default: OFF

**ENABLE\_OPENMP**

Enable OpenMP support (build the OpenMP NVector)

Default: OFF

**ENABLE\_PETSC**

Enable PETSc support

Default: OFF

---

**Note:** See additional information on building with PETSc enabled in §12.4.

---

**PETSC\_DIR**

Path to PETSc installation

Default: none

**PETSC\_LIBRARIES**

Semi-colon separated list of PETSc link libraries. Unless provided by the user, this is autopopulated based on the PETSc installation found in PETSC\_DIR.

Default: none

**PETSC\_INCLUDES**

Semi-colon separated list of PETSc include directories. Unless provided by the user, this is autopopulated based on the PETSc installation found in PETSC\_DIR.

Default: none

**ENABLE\_PTHREAD**

Enable Pthreads support (build the Pthreads NVector)

Default: OFF

**ENABLE\_RAJA**

Enable RAJA support.

Default: OFF

---

**Note:** You need to enable CUDA or HIP in order to build the RAJA vector module.

---

#### **SUNDIALS\_RAJA\_BACKENDS**

If building SUNDIALS with RAJA support, this sets the RAJA backend to target. Values supported are CUDA, HIP, or SYCL.

Default: CUDA

#### **ENABLE\_SUPERLUDIST**

Enable SuperLU\_DIST support

Default: OFF

---

**Note:** See additional information on building with SuperLU\_DIST enabled in §12.4.

---

#### **SUPERLUDIST\_DIR**

Path to SuperLU\_DIST installation.

Default: none

#### **SUPERLUDIST\_OpenMP**

Enable SUNDIALS support for SuperLU\_DIST built with OpenMP

Default: none

Note: SuperLU\_DIST must be built with OpenMP support for this option to function. Additionally the environment variable OMP\_NUM\_THREADS must be set to the desired number of threads.

#### **SUPERLUDIST\_INCLUDE\_DIRS**

List of include paths for SuperLU\_DIST (under a typical SuperLU\_DIST install, this is typically the SuperLU\_DIST SRC directory)

Default: none

---

**Note:** This is an advanced option. Prefer to use [SUPERLUDIST\\_DIR](#).

---

#### **SUPERLUDIST\_LIBRARIES**

Semi-colon separated list of libraries needed for SuperLU\_DIST

Default: none

---

**Note:** This is an advanced option. Prefer to use [SUPERLUDIST\\_DIR](#).

---

#### **SUPERLUDIST\_INCLUDE\_DIR**

Path to SuperLU\_DIST header files (under a typical SuperLU\_DIST install, this is typically the SuperLU\_DIST SRC directory)

Default: none

---

**Note:** This is an advanced option. This option is deprecated. Use [SUPERLUDIST\\_INCLUDE\\_DIRS](#).

---

**SUPERLUDIST\_LIBRARY\_DIR**

Path to SuperLU\_DIST installed library files

Default: none

---

**Note:** This option is deprecated. Use [SUPERLUDIST\\_DIR](#).

---

**ENABLE\_SUPERLUMT**

Enable SuperLU\_MT support

Default: OFF

---

**Note:** See additional information on building with SuperLU\_MT enabled in [§12.4](#).

---

**SUPERLUMT\_INCLUDE\_DIR**

Path to SuperLU\_MT header files (under a typical SuperLU\_MT install, this is typically the SuperLU\_MT SRC directory)

Default: none

**SUPERLUMT\_LIBRARY\_DIR**

Path to SuperLU\_MT installed library files

Default: none

**SUPERLUMT\_THREAD\_TYPE**

Must be set to Pthread or OpenMP, depending on how SuperLU\_MT was compiled.

Default: Pthread

**ENABLE\_SYCL**

Enable SYCL support.

Default: OFF

---

**Note:** Building with SYCL enabled requires a compiler that supports a subset of the of SYCL 2020 specification (specifically `sycl/sycl.hpp` must be available).

CMake does not currently support autodetection of SYCL compilers and `CMAKE_CXX_COMPILER` must be set to a valid SYCL compiler. At present the only supported SYCL compilers are the Intel oneAPI compilers i.e., `dpcpp` and `icpx`. When using `icpx` the `-fsycl` flag and any ahead of time compilation flags must be added to `CMAKE_CXX_FLAGS`.

---

**SUNDIALS\_SYCL\_2020\_UNSUPPORTED**

This advanced option disables the use of *some* features from the SYCL 2020 standard in SUNDIALS libraries and examples. This can be used to work around some cases of incomplete compiler support for SYCL 2020.

Default: OFF

**ENABLE\_CALIPER**

Enable CALIPER support

Default: OFF

---

**Note:** Using Caliper requires setting [SUNDIALS\\_BUILD\\_WITH\\_PROFILING](#) to ON.

---

#### **CALIPER\_DIR**

Path to the root of a Caliper installation

Default: None

#### **ENABLE\_ADIK**

Enable Adiak support

Default: OFF

#### **adiak\_DIR**

Path to the root of an Adiak installation

Default: None

#### **SUNDIALS\_LAPACK\_CASE**

Specify the case to use in the Fortran name-mangling scheme, options are: **lower** or **upper**

Default:

---

**Note:** The build system will attempt to infer the Fortran name-mangling scheme using the Fortran compiler. This option should only be used if a Fortran compiler is not available or to override the inferred or default (**lower**) scheme if one can not be determined. If used, **SUNDIALS\_LAPACK\_UNDERSCORES** must also be set.

---

#### **SUNDIALS\_LAPACK\_UNDERSCORES**

Specify the number of underscores to append in the Fortran name-mangling scheme, options are: **none**, **one**, or **two**

Default:

---

**Note:** The build system will attempt to infer the Fortran name-mangling scheme using the Fortran compiler. This option should only be used if a Fortran compiler is not available or to override the inferred or default (**one**) scheme if one can not be determined. If used, **SUNDIALS\_LAPACK\_CASE** must also be set.

---

#### **SUNDIALS\_INDEX\_TYPE**

Integer type used for SUNDIALS indices. The size must match the size provided for the **SUNDIALS\_INDEX\_SIZE** option.

Default: Automatically determined based on [\*SUNDIALS\\_INDEX\\_SIZE\*](#)

---

**Note:** In past SUNDIALS versions, a user could set this option to **INT64\_T** to use 64-bit integers, or **INT32\_T** to use 32-bit integers. Starting in SUNDIALS 3.2.0, these special values are deprecated. For SUNDIALS 3.2.0 and up, a user will only need to use the [\*SUNDIALS\\_INDEX\\_SIZE\*](#) option in most cases.

---

#### **SUNDIALS\_INDEX\_SIZE**

Integer size (in bits) used for indices in SUNDIALS, options are: 32 or 64

Default: 64

---

**Note:** The build system tries to find an integer type of appropriate size. Candidate 64-bit integer types are (in order of preference): **int64\_t**, **\_\_int64**, **long long**, and **long**. Candidate 32-bit integers are (in order of preference): **int32\_t**, **int**, and **long**. The advanced option, [\*SUNDIALS\\_INDEX\\_TYPE\*](#) can be used to provide a type not listed here.

---



**SUNDIALS\_PRECISION**

The floating-point precision used in SUNDIALS packages and class implementations, options are: double, single, or extended

Default: double

**SUNDIALS\_MATH\_LIBRARY**

The standard C math library (e.g., `libm`) to link with.

Default: `-lm` on Unix systems, none otherwise

**SUNDIALS\_INSTALL\_CMAKEDIR**

Installation directory for the SUNDIALS cmake files (relative to [\*CMAKE\\_INSTALL\\_PREFIX\*](#)).

Default: `CMAKE_INSTALL_PREFIX/cmake/sundials`

**ENABLE\_XBRAID**

Enable or disable the ARKStep + XBraid interface.

Default: OFF

---

**Note:** See additional information on building with *XBraid* enabled in §12.4.

---

**XBRAID\_DIR**

The root directory of the XBraid installation.

Default: OFF

**XBRAID\_INCLUDES**

Semi-colon separated list of XBraid include directories. Unless provided by the user, this is autopopulated based on the XBraid installation found in `XBRAID_DIR`.

Default: none

**XBRAID\_LIBRARIES**

Semi-colon separated list of XBraid link libraries. Unless provided by the user, this is autopopulated based on the XBraid installation found in `XBRAID_DIR`.

Default: none

**USE\_XSDK\_DEFAULTS**

Enable xSDK (see <https://xsdk.info> for more information) default configuration settings. This sets `CMAKE_BUILD_TYPE` to Debug, `SUNDIALS_INDEX_SIZE` to 32 and `SUNDIALS_PRECISION` to double.

Default: OFF

## 12.3 Configuration examples

The following examples will help demonstrate usage of the CMake configure options.

To configure SUNDIALS using the default C and Fortran compilers, and default `mpicc` and `mpif90` parallel compilers, enable compilation of examples, and install libraries, headers, and example sources under subdirectories of `/home/myname/sundials/`, use:

```
% cmake \  
> -DCMAKE_INSTALL_PREFIX=/home/myname/sundials/instdir \  
> -DEXAMPLES_INSTALL_PATH=/home/myname/sundials/instdir/examples \  
> -DENABLE_MPI=ON \  
> /home/myname/sundials/srcdir  
  
% make install
```

To disable installation of the examples, use:

```
% cmake \  
> -DCMAKE_INSTALL_PREFIX=/home/myname/sundials/instdir \  
> -DEXAMPLES_INSTALL_PATH=/home/myname/sundials/instdir/examples \  
> -DENABLE_MPI=ON \  
> -DEXAMPLES_INSTALL=OFF \  
> /home/myname/sundials/srcdir  
  
% make install
```

## 12.4 Working with external Libraries

The SUNDIALS suite contains many options to enable implementation flexibility when developing solutions. The following are some notes addressing specific configurations when using the supported third party libraries.

### 12.4.1 Building with Ginkgo

[Ginkgo](#) is a high-performance linear algebra library for manycore systems, with a focus on solving sparse linear systems. It is implemented using modern C++ (you will need at least a C++14 compliant compiler to build it), with GPU kernels implemented in CUDA (for NVIDIA devices), HIP (for AMD devices) and SYCL/DPC++ (for Intel devices and other supported hardware). To enable Ginkgo in SUNDIALS, set the [ENABLE\\_GINKGO](#) to ON and provide the path to the root of the Ginkgo installation in [Ginkgo\\_DIR](#). Additionally, [SUNDIALS\\_GINKGO\\_BACKENDS](#) must be set to a list of Ginkgo target architectures/executors. E.g.,

```
% cmake \  
> -DENABLE_GINKGO=ON \  
> -DGinkgo_DIR=/path/to/ginkgo/installation \  
> -DSUNDIALS_GINKGO_BACKENDS="REF;OMP;CUDA" \  
> /home/myname/sundials/srcdir
```

The SUNDIALS interfaces to Ginkgo are not compatible with [SUNDIALS\\_PRECISION](#) set to `extended`.

### 12.4.2 Building with Kokkos

**Kokkos** is a modern C++ (requires at least C++14) programming model for writing performance portable code for multi-core CPU and GPU-based systems including NVIDIA, AMD, and Intel accelerators. To enable Kokkos in SUNDIALS, set the `ENABLE_KOKKOS` to ON and provide the path to the root of the Kokkos installation in `Kokkos_DIR`. Additionally, the **Kokkos-Kernels** library provides common computational kernels for linear algebra. To enable Kokkos-Kernels in SUNDIALS, set the `ENABLE_KOKKOS_KERNELS` to ON and provide the path to the root of the Kokkos-Kernels installation in `KokkosKernels_DIR` e.g.,

```
% cmake \
> -DENABLE_KOKKOS=ON \
> -DKokkos_DIR=/path/to/kokkos/installation \
> -DENABLE_KOKKOS_KERNELS=ON \
> -DKokkosKernels_DIR=/path/to/kokkoskernels/installation \
> /home/myname/sundials/srcdir
```

---

**Note:** The minimum supported version of Kokkos-Kernels 3.7.00.

---

### 12.4.3 Building with LAPACK

To enable LAPACK, set the `ENABLE_LAPACK` option to ON. If the directory containing the LAPACK library is in the `LD_LIBRARY_PATH` environment variable, CMake will set the `LAPACK_LIBRARIES` variable accordingly, otherwise CMake will attempt to find the LAPACK library in standard system locations. To explicitly tell CMake what library to use, the `LAPACK_LIBRARIES` variable can be set to the desired libraries required for LAPACK.

```
% cmake \
> -DCMAKE_INSTALL_PREFIX=/home/myname/sundials/instdir \
> -DEXAMPLES_INSTALL_PATH=/home/myname/sundials/instdir/examples \
> -DENABLE_LAPACK=ON \
> -DLAPACK_LIBRARIES=/mylapackpath/lib/libblas.so;/mylapackpath/lib/liblapack.so \
> /home/myname/sundials/srcdir

% make install
```

---

**Note:** If a working Fortran compiler is not available to infer the Fortran name-mangling scheme, the options `SUNDIALS_F77_FUNC_CASE` and `SUNDIALS_F77_FUNC_UNDERSCORES` *must* be set in order to bypass the check for a Fortran compiler and define the name-mangling scheme. The defaults for these options in earlier versions of SUNDIALS were lower and one, respectively.

---

SUNDIALS has been tested with OpenBLAS 0.3.18.

### 12.4.4 Building with KLU

KLU is a software package for the direct solution of sparse nonsymmetric linear systems of equations that arise in circuit simulation and is part of SuiteSparse, a suite of sparse matrix software. The library is developed by Texas A&M University and is available from the [SuiteSparse GitHub repository](#).

To enable KLU, set `ENABLE_KLU` to ON, set `KLU_INCLUDE_DIR` to the `include` path of the KLU installation and set `KLU_LIBRARY_DIR` to the `lib` path of the KLU installation. In that case, the CMake configure will result in populating the following variables: `AMD_LIBRARY`, `AMD_LIBRARY_DIR`, `BTF_LIBRARY`, `BTF_LIBRARY_DIR`, `COLAMD_LIBRARY`, `COLAMD_LIBRARY_DIR`, and `KLU_LIBRARY`.

For SuiteSparse 7.4.0 and newer, the necessary information can also be gathered from a CMake import target. If SuiteSparse is installed in a non-default prefix, the path to the CMake Config file can be set using `CMAKE_PREFIX_PATH`. In that case, the CMake configure step won't populate the previously mentioned variables. It is still possible to set `KLU_INCLUDE_DIR` and `KLU_LIBRARY_DIR` which take precedence over a potentially installed CMake import target file.

In either case, a CMake target `SUNDIALS::KLU` will be created if the KLU library could be found. Dependent targets should link to that target.

SUNDIALS has been tested with SuiteSparse version 5.10.1.

### 12.4.5 Building with SuperLU\_DIST

SuperLU\_DIST is a general purpose library for the direct solution of large, sparse, nonsymmetric systems of linear equations in a distributed memory setting. The library is developed by Lawrence Berkeley National Laboratory and is available from the [SuperLU\\_DIST GitHub repository](#).

To enable SuperLU\_DIST, set `ENABLE_SUPERLUDIST` to ON, set `SUPERLUDIST_DIR` to the path where SuperLU\_DIST is installed. If SuperLU\_DIST was built with OpenMP then the option `SUPERLUDIST_OpenMP` and `ENABLE_OPENMP` should be set to ON.

SUNDIALS supports SuperLU\_DIST v7.0.0 – v8.x.x and has been tested with v7.2.0 and v8.1.0.

### 12.4.6 Building with SuperLU\_MT

SuperLU\_MT is a general purpose library for the direct solution of large, sparse, nonsymmetric systems of linear equations on shared memory parallel machines. The library is developed by Lawrence Berkeley National Laboratory and is available from the [SuperLU\\_MT GitHub repository](#).

To enable SuperLU\_MT, set `ENABLE_SUPERLUMT` to ON, set `SUPERLUMT_INCLUDE_DIR` to the `SRC` path of the SuperLU\_MT installation, and set the variable `SUPERLUMT_LIBRARY_DIR` to the `lib` path of the SuperLU\_MT installation. At the same time, the variable `SUPERLUMT_LIBRARIES` must be set to a semi-colon separated list of other libraries SuperLU\_MT depends on. For example, if SuperLU\_MT was build with an external blas library, then include the full path to the blas library in this list. Additionally, the variable `SUPERLUMT_THREAD_TYPE` must be set to either `Pthread` or `OpenMP`.

Do not mix thread types when building SUNDIALS solvers. If threading is enabled for SUNDIALS by having either `ENABLE_OPENMP` or `ENABLE_PTHREAD` set to ON then SuperLU\_MT should be set to use the same threading type.

SUNDIALS has been tested with SuperLU\_MT version 3.1.

### 12.4.7 Building with PETSc

The Portable, Extensible Toolkit for Scientific Computation (PETSc) is a suite of data structures and routines for simulating applications modeled by partial differential equations. The library is developed by Argonne National Laboratory and is available from the [PETSc GitLab repository](#).

To enable PETSc, set `ENABLE_PETSC` to ON, and set `PETSC_DIR` to the path of the PETSc installation. Alternatively, a user can provide a list of include paths in `PETSC_INCLUDES` and a list of complete paths to the PETSc libraries in `PETSC_LIBRARIES`.

SUNDIALS is regularly tested with the latest PETSc versions, specifically up to version 3.18.1 as of SUNDIALS version v6.1.0. SUNDIALS requires PETSc 3.5.0 or newer.

### 12.4.8 Building with hypre

*hypre* is a library of high performance preconditioners and solvers featuring multigrid methods for the solution of large, sparse linear systems of equations on massively parallel computers. The library is developed by Lawrence Livermore National Laboratory and is available from the [hypre GitHub repository](#).

To enable *hypre*, set `ENABLE_HYPRE` to ON, set `HYPRE_INCLUDE_DIR` to the include path of the *hypre* installation, and set the variable `HYPRE_LIBRARY_DIR` to the lib path of the *hypre* installation.

---

**Note:** SUNDIALS must be configured so that `SUNDIALS_INDEX_SIZE` is compatible with `HYPRE_BigInt` in the *hypre* installation.

---

SUNDIALS is regularly tested with the latest versions of *hypre*, specifically up to version 2.26.0 as of SUNDIALS version v6.1.0.

### 12.4.9 Building with MAGMA

The Matrix Algebra on GPU and Multicore Architectures (MAGMA) project provides a dense linear algebra library similar to LAPACK but targeting heterogeneous architectures. The library is developed by the University of Tennessee and is available from the [UTK webpage](#).

To enable the SUNDIALS MAGMA interface set `ENABLE_MAGMA` to ON, `MAGMA_DIR` to the MAGMA installation path, and `SUNDIALS_MAGMA_BACKENDS` to the desired MAGMA backend to use with SUNDIALS e.g., CUDA or HIP.

SUNDIALS has been tested with MAGMA version v2.6.1 and v2.6.2.

### 12.4.10 Building with oneMKL for SYCL

The Intel [oneAPI Math Kernel Library \(oneMKL\)](#) includes CPU and SYCL/DPC++ interfaces for LAPACK dense linear algebra routines. The SUNDIALS oneMKL interface targets the SYCL/DPC++ routines, to utilize the CPU routine see §12.4.3.

To enable the SUNDIALS oneMKL interface set `ENABLE_ONEMKL` to ON and `ONEMKL_DIR` to the oneMKL installation path.

SUNDIALS has been tested with oneMKL version 2021.4.

### 12.4.11 Building with CUDA

The NVIDIA CUDA Toolkit provides a development environment for GPU-accelerated computing with NVIDIA GPUs. The CUDA Toolkit and compatible NVIDIA drivers are available from the [NVIDIA developer website](#).

To enable CUDA, set `ENABLE_CUDA` to `ON`. If CUDA is installed in a nonstandard location, you may be prompted to set the variable `CUDA_TOOLKIT_ROOT_DIR` with your CUDA Toolkit installation path. To enable CUDA examples, set `EXAMPLES_ENABLE_CUDA` to `ON`.

SUNDIALS has been tested with the CUDA toolkit versions 10 and 11.

### 12.4.12 Building with HIP

HIP(heterogeneous-compute interface for portability) allows developers to create portable applications for AMD and NVIDIA GPUs. HIP can be obtained from [HIP GitHub repository](#).

To enable HIP, set `ENABLE_HIP` to `ON` and set `AMDGPU_TARGETS` to the desired target(ex. `gfx705`). In addition, set `CMAKE_C_COMPILER` and `CMAKE_CXX_COMPILER` to point to an installation of `hipcc`.

SUNDIALS has been tested with HIP versions between 5.0.0 to 5.4.3.

### 12.4.13 Building with RAJA

RAJA is a performance portability layer developed by Lawrence Livermore National Laboratory and can be obtained from the [RAJA GitHub repository](#).

Building SUNDIALS RAJA modules requires a CUDA, HIP, or SYCL enabled RAJA installation. To enable RAJA, set `ENABLE_RAJA` to `ON`, set `SUNDIALS_RAJA_BACKENDS` to the desired backend (CUDA, HIP, or SYCL), and set `ENABLE_CUDA`, `ENABLE_HIP`, or `ENABLE_SYCL` to `ON` depending on the selected backend. If RAJA is installed in a nonstandard location you will be prompted to set the variable `RAJA_DIR` with the path to the RAJA CMake configuration file. To enable building the RAJA examples set `EXAMPLES_ENABLE_CXX` to `ON`.

SUNDIALS has been tested with RAJA version 0.14.0.

### 12.4.14 Building with XBraid

XBraid is parallel-in-time library implementing an optimal-scaling multigrid reduction in time (MGRIT) solver. The library is developed by Lawrence Livermore National Laboratory and is available from the [XBraid GitHub repository](#).

To enable XBraid support, set `ENABLE_XBRAID` to `ON`, set `XBRAID_DIR` to the root install location of XBraid or the location of the clone of the XBraid repository.

---

**Note:** At this time the XBraid types `braid_Int` and `braid_Real` are hard-coded to `int` and `double` respectively. As such SUNDIALS must be configured with `SUNDIALS_INDEX_SIZE` set to 32 and `SUNDIALS_PRECISION` set to `double`. Additionally, SUNDIALS must be configured with `ENABLE_MPI` set to `ON`.

---

SUNDIALS has been tested with XBraid version 3.0.0.

## 12.5 Testing the build and installation

If SUNDIALS was configured with `EXAMPLES_ENABLE_<language>` options to ON, then a set of regression tests can be run after building with the `make` command by running:

```
% make test
```

Additionally, if `EXAMPLES_INSTALL` was also set to ON, then a set of smoke tests can be run after installing with the `make install` command by running:

```
% make test_install
```

## 12.6 Building and Running Examples

Each of the SUNDIALS solvers is distributed with a set of examples demonstrating basic usage. To build and install the examples, set at least of the `EXAMPLES_ENABLE_<language>` options to ON, and set `EXAMPLES_INSTALL` to ON. Specify the installation path for the examples with the variable `EXAMPLES_INSTALL_PATH`. CMake will generate `CMakeLists.txt` configuration files (and Makefile files if on Linux/Unix) that reference the *installed* SUNDIALS headers and libraries.

Either the `CMakeLists.txt` file or the traditional Makefile may be used to build the examples as well as serve as a template for creating user developed solutions. To use the supplied Makefile simply run `make` to compile and generate the executables. To use CMake from within the installed example directory, run `cmake` (or `ccmake` or `cmake-gui` to use the GUI) followed by `make` to compile the example code. Note that if CMake is used, it will overwrite the traditional Makefile with a new CMake-generated Makefile.

The resulting output from running the examples can be compared with example output bundled in the SUNDIALS distribution.

---

**Note:** There will potentially be differences in the output due to machine architecture, compiler versions, use of third party libraries etc.

---

## 12.7 Configuring, building, and installing on Windows

CMake can also be used to build SUNDIALS on Windows. To build SUNDIALS for use with Visual Studio the following steps should be performed:

1. Unzip the downloaded tar file(s) into a directory. This will be the `SOLVERDIR`
2. Create a separate `BUILDDIR`
3. Open a Visual Studio Command Prompt and `cd` to `BUILDDIR`
4. Run `cmake-gui ../SOLVERDIR`
  - a. Hit Configure
  - b. Check/Uncheck solvers to be built
  - c. Change `CMAKE_INSTALL_PREFIX` to `INSTDIR`
  - d. Set other options as desired
  - e. Hit Generate

5. Back in the VS Command Window:
  - a. Run `msbuild ALL_BUILD.vcxproj`
  - b. Run `msbuild INSTALL.vcxproj`

The resulting libraries will be in the `INSTDIR`.

The SUNDIALS project can also now be opened in Visual Studio. Double click on the `ALL_BUILD.vcxproj` file to open the project. Build the whole *solution* to create the SUNDIALS libraries. To use the SUNDIALS libraries in your own projects, you must set the include directories for your project, add the SUNDIALS libraries to your project solution, and set the SUNDIALS libraries as dependencies for your project.

## 12.8 Installed libraries and exported header files

Using the CMake SUNDIALS build system, the command

```
$ make install
```

will install the libraries under `LIBDIR` and the public header files under `INCLUDEDIR`. The values for these directories are `INSTDIR/lib` and `INSTDIR/include`, respectively. The location can be changed by setting the CMake variable `CMAKE_INSTALL_PREFIX`. Although all installed libraries reside under `LIBDIR/lib`, the public header files are further organized into subdirectories under `INCLUDEDIR/include`.

The installed libraries and exported header files are listed for reference in the table below. The file extension `.LIB` is typically `.so` for shared libraries and `.a` for static libraries. Note that, in this table names are relative to `LIBDIR` for libraries and to `INCLUDEDIR` for header files.

**Warning:** SUNDIALS installs some header files to `INSTDIR/include/sundials/priv`. All of the header files in this directory are private and **should not be included in user code**. The private headers are subject to change without any notice and relying on them may break your code.

## 12.9 Using SUNDIALS in your project

After building and installing SUNDIALS, using SUNDIALS in your application involves two steps: including the right header files and linking to the right libraries.

Depending on what features of SUNDIALS that your application uses, the header files needed will vary. For example, if you want to use CVODE for serial computations you need the following includes:

```
#include <cvode/cvode.h>
#include <nvector/nvector_serial.h>
```

If you wanted to use CVODE with the GMRES linear solver and our CUDA enabled vector:

```
#include <cvode/cvode.h>
#include <nvector/nvector_cuda.h>
#include <sunlinsol/sunlinsol_spgmr.h>
```

The story is similar for linking to SUNDIALS. Starting in v7.0.0, all applications will need to link to `libsundials_core`. Furthermore, depending on the packages and modules of SUNDIALS of interest an application will need to



link to a few more libraries. Using the same examples as for the includes, we would need to also link to `libsundials_cvode`, `libsundials_nvecserial` for the first example and `libsundials_cvode`, `libsundials_nveccuda`, `libsundials_sunlinsolspgmr` for the second.

Refer to the documentations sections for the individual packages and modules of SUNDIALS that interest you for the proper includes and libraries to link to.

## 12.10 Using SUNDIALS as a Third Party Library in other CMake Projects

The `make install` command will also install a [CMake package configuration file](#) that other CMake projects can load to get all the information needed to build against SUNDIALS. In the consuming project's CMake code, the `find_package` command may be used to search for the configuration file, which will be installed to `instdir/SUNDIALS_INSTALL_CMAKEDIR/SUNDIALSConfig.cmake` alongside a package version file `instdir/SUNDIALS_INSTALL_CMAKEDIR/SUNDIALSConfigVersion.cmake`. Together these files contain all the information the consuming project needs to use SUNDIALS, including exported CMake targets. The SUNDIALS exported CMake targets follow the same naming convention as the generated library binaries, e.g. the exported target for CVODE is `SUNDIALS::cvode`. The CMake code snipped below shows how a consuming project might leverage the SUNDIALS package configuration file to build against SUNDIALS in their own CMake project.

```
project(MyProject)

# Set the variable SUNDIALS_DIR to the SUNDIALS instdir.
# When using the cmake CLI command, this can be done like so:
#   cmake -D SUNDIALS_DIR=/path/to/sundials/installation

# Find any SUNDIALS version...
find_package(SUNDIALS REQUIRED)

# ... or find any version newer than some minimum...
find_package(SUNDIALS 7.1.0 REQUIRED)

# ... or find a version in a range
find_package(SUNDIALS 7.0.0...7.1.0 REQUIRED)

add_executable(myexec main.c)

# Link to SUNDIALS libraries through the exported targets.
# This is just an example, users should link to the targets appropriate
# for their use case.
target_link_libraries(myexec PUBLIC SUNDIALS::cvode SUNDIALS::nvecpetsc)
```

**Note:** Changed in version 7.1.0: A single version provided to `find_package` denotes the minimum version of SUNDIALS to look for, and any version equal or newer than what is specified will match. In prior versions `SUNDIALSConfig.cmake` required the version found to have the same major version number as the single version provided to `find_package`.

## 12.11 Table of SUNDIALS libraries and header files

Table 12.1: SUNDIALS shared libraries and header files

Core	Libraries	libsundials_core.LIB
	Headers	sundials/sundials_band.h
		sundials/sundials_config.h
		sundials/sundials_context.h
		sundials/sundials_cuda_policies.hpp
		sundials/sundials_dense.h
		sundials/sundials_direct.h
		sundials/sundials_hip_policies.hpp
		sundials/sundials_iterative.h
		sundials/sundials_linearsolver.h
		sundials/sundials_math.h
		sundials/sundials_matrix.h
		sundials/sundials_memory.h
		sundials/sundials_mpi_types.h
		sundials/sundials_nonlinearsolver.h
		sundials/sundials_nvector.h
		sundials/sundials_types.h
		sundials/sundials_version.h
		sundials/sundials_xbraid.h
	NVECTOR Modules	
SERIAL	Libraries	libsundials_nvecserial.LIB
	Headers	nvector/nvector_serial.h
PARALLEL	Libraries	libsundials_nvecparallel.LIB
	Headers	nvector/nvector_parallel.h
OPENMP	Libraries	libsundials_nvecopenmp.LIB
	Headers	nvector/nvector_openmp.h
PTHREADS	Libraries	libsundials_nvecpthreads.LIB
	Headers	nvector/nvector_pthreads.h
PARHYP	Libraries	libsundials_nvecparhyp.LIB
	Headers	nvector/nvector_parhyp.h
PETSC	Libraries	libsundials_nvecpetsc.LIB
	Headers	nvector/nvector_petsc.h
CUDA	Libraries	libsundials_nveccuda.LIB
	Headers	nvector/nvector_cuda.h
HIP	Libraries	libsundials_nvechip.LIB
	Headers	nvector/nvector_hip.h
RAJA	Libraries	libsundials_nveccudaraja.LIB
		libsundials_nvechipraja.LIB
	Headers	nvector/nvector_rajah.h
SYCL	Libraries	libsundials_nvecsycl.LIB
	Headers	nvector/nvector_sycl.h
MANYVECTOR	Libraries	libsundials_nvecmanyvector.LIB
	Headers	nvector/nvector_manyvector.h
MPIMANYVECTOR	Libraries	libsundials_nvecmpimanyvector.LIB
	Headers	nvector/nvector_mpimanyvector.h
MPIPLUSX	Libraries	libsundials_nvecmpiplusx.LIB
	Headers	nvector/nvector_mpiplusx.h

continues on next page

Table 12.1 – continued from previous page

<b>SUNMATRIX Modules</b>		
BAND	Libraries	libsundials_sunmatrixband.LIB
	Headers	sunmatrix/sunmatrix_band.h
CUSPARSE	Libraries	libsundials_sunmatrixcusparse.LIB
	Headers	sunmatrix/sunmatrix_cusparse.h
DENSE	Libraries	libsundials_sunmatrixdense.LIB
	Headers	sunmatrix/sunmatrix_dense.h
Ginkgo	Headers	sunmatrix/sunmatrix_ginkgo.hpp
MAGMADENSE	Libraries	libsundials_sunmatrixmagmadense.LIB
	Headers	sunmatrix/sunmatrix_magmadense.h
ONEMKLDENSE	Libraries	libsundials_sunmatrixonemkldense.LIB
	Headers	sunmatrix/sunmatrix_onemkldense.h
SPARSE	Libraries	libsundials_sunmatrixsparse.LIB
	Headers	sunmatrix/sunmatrix_sparse.h
SLUNRLOC	Libraries	libsundials_sunmatrixslunrloc.LIB
	Headers	sunmatrix/sunmatrix_slunrloc.h
<b>SUNLINSOL Modules</b>		
BAND	Libraries	libsundials_sunlinsolband.LIB
	Headers	sunlinsol/sunlinsol_band.h
CUSOLVERS_BATCHQR	Libraries	libsundials_sunlinsolcusolversp.LIB
	Headers	sunlinsol/sunlinsol_cusolversp_batchqr.h
DENSE	Libraries	libsundials_sunlinsoldense.LIB
	Headers	sunlinsol/sunlinsol_dense.h
Ginkgo	Headers	sunlinsol/sunlinsol_ginkgo.hpp
KLU	Libraries	libsundials_sunlinsolklu.LIB
	Headers	sunlinsol/sunlinsol_klu.h
LAPACKBAND	Libraries	libsundials_sunlinsollapackband.LIB
	Headers	sunlinsol/sunlinsol_lapackband.h
LAPACKDENSE	Libraries	libsundials_sunlinsollapackdense.LIB
	Headers	sunlinsol/sunlinsol_lapackdense.h
MAGMADENSE	Libraries	libsundials_sunlinsolmagmadense.LIB
	Headers	sunlinsol/sunlinsol_magmadense.h
ONEMKLDENSE	Libraries	libsundials_sunlinsolonemkldense.LIB
	Headers	sunlinsol/sunlinsol_onemkldense.h
PCG	Libraries	libsundials_sunlinsolpcg.LIB
	Headers	sunlinsol/sunlinsol_pcg.h
SPBCGS	Libraries	libsundials_sunlinsolspbcgs.LIB
	Headers	sunlinsol/sunlinsol_spbcgs.h
SPFGMR	Libraries	libsundials_sunlinsolspfgmr.LIB
	Headers	sunlinsol/sunlinsol_spfgmr.h
SPGMR	Libraries	libsundials_sunlinsolspgmr.LIB
	Headers	sunlinsol/sunlinsol_spgmr.h
SPTFQMR	Libraries	libsundials_sunlinsolsptfqmr.LIB
	Headers	sunlinsol/sunlinsol_sptfqmr.h
SUPERLUDIST	Libraries	libsundials_sunlinsolsuperludist.LIB
	Headers	sunlinsol/sunlinsol_superludist.h
SUPERLUMT	Libraries	libsundials_sunlinsolsuperlumt.LIB
	Headers	sunlinsol/sunlinsol_superlumt.h
<b>SUNNONLINSOL Modules</b>		
NEWTON	Libraries	libsundials_sunnonlinsolnewton.LIB
	Headers	sunnonlinsol/sunnonlinsol_newton.h

continues on next page

Table 12.1 – continued from previous page

FIXEDPOINT	Libraries	libsundials_sunnonlinolfixedpoint.LIB
	Headers	sunnonlinol/sunnonlinol_fixedpoint.h
PETSCSNES	Libraries	libsundials_sunnonlinolpetscsnes.LIB
	Headers	sunnonlinol/sunnonlinol_petscsnes.h
<b>SUNMEMORY Modules</b>		
SYSTEM	Libraries	libsundials_sunmemsys.LIB
	Headers	sunmemory/sunmemory_system.h
CUDA	Libraries	libsundials_sunmemcuda.LIB
	Headers	sunmemory/sunmemory_cuda.h
HIP	Libraries	libsundials_sunmemhip.LIB
	Headers	sunmemory/sunmemory_hip.h
SYCL	Libraries	libsundials_sunmemsycl.LIB
	Headers	sunmemory/sunmemory_sycl.h
<b>SUNDIALS Packages</b>		
CVODE	Libraries	libsundials_cvode.LIB
	Headers	cvode/cvode.h
		cvode/cvode_bandpre.h
		cvode/cvode_bbdpre.h
		cvode/cvode_diag.h
		cvode/cvode_impl.h
		cvode/cvode_ls.h
CVODES	Libraries	libsundials_cvodes.LIB
	Headers	cvodes/cvodes.h
		cvodes/cvodes_bandpre.h
		cvodes/cvodes_bbdpre.h
		cvodes/cvodes_diag.h
		cvodes/cvodes_impl.h
		cvodes/cvodes_ls.h
ARKODE	Libraries	libsundials_arkode.LIB
	Headers	libsundials_xbraid.LIB
		arkode/arkode.h
		arkode/arkode_arkstep.h
		arkode/arkode_bandpre.h
		arkode/arkode_bbdpre.h
		arkode/arkode_butcher.h
		arkode/arkode_butcher_dirk.h
		arkode/arkode_butcher_erk.h
		arkode/arkode_erkstep.h
		arkode/arkode_impl.h
		arkode/arkode_ls.h
		arkode/arkode_mristep.h
		arkode/arkode_xbraid.h
IDA	Libraries	libsundials_ida.LIB
	Headers	ida/ida.h
		ida/ida_bbdpre.h
		ida/ida_impl.h
IDAS	Headers	ida/ida_ls.h
		libsundials_idas.LIB
		idas/idas.h
		idas/idas_bbdpre.h

continues on next page

Table 12.1 – continued from previous page

KINSOL		idas/idas_impl.h
	Libraries	libsundials_kinsol.LIB
	Headers	kinsol/kinsol.h
		kinsol/kinsol_bbdpre.h
		kinsol/kinsol_impl.h
		kinsol/kinsol_ls.h

## 12.12 Installing SUNDIALS on HPC Clusters

This section is a guide for installing SUNDIALS on specific HPC clusters. In general, the procedure is the same as described previously for Linux machines. The main differences are in the modules and environment variables that are specific to different HPC clusters. We aim to keep this section as up to date as possible, but it may lag the latest software updates to each cluster.

### 12.12.1 Frontier

[Frontier](#) is an Exascale supercomputer at the Oak Ridge Leadership Computing Facility. If you are new to this system, then we recommend that you review the [Frontier user guide](#).

#### A Standard Installation

Clone SUNDIALS:

```
git clone https://github.com/LLNL/sundials.git && cd sundials
```

Next we load the modules and set the environment variables needed to build SUNDIALS. This configuration enables both MPI and HIP support for distributed and GPU parallelism. It uses the HIP compiler for C and C++ and the Cray Fortran compiler. Other configurations are possible.

```
# required dependencies
module load PrgEnv-cray-amd/8.5.0
module load craype-accel-amd-gfx90a
module load rocm/5.3.0
module load cmake/3.23.2

# GPU-aware MPI
export MPICH_GPU_SUPPORT_ENABLED=1

# compiler environment hints
export CC=$(which hipcc)
export CXX=$(which hipcc)
export FC=$(which ftn)
export CFLAGS="-I${ROCM_PATH}/include"
export CXXFLAGS="-I${ROCM_PATH}/include -Wno-pass-failed"
export LDFLAGS="-L${ROCM_PATH}/lib -lamdhip64 ${PE_MPICH_GTL_DIR_amd_gfx90a} -lmpi_gtl_hsa"
```

Now we can build SUNDIALS. In general, this is the same procedure described in the previous sections. The following command builds and installs SUNDIALS with MPI, HIP, and the Fortran interface enabled, where *<install path>* is your desired installation location, and *<account>* is your allocation account on Frontier:

```
cmake -S . -B builddir -DCMAKE_INSTALL_PREFIX=<install path> -DAMDGPU_TARGETS=gfx90a \
-DENABLE_HIP=ON -DENABLE_MPI=ON -DBUILD_FORTRAN_MODULE_INTERFACE=ON
cd builddir
make -j8 install
# Need an allocation to run the tests:
salloc -A <account> -t 10 -N 1 -p batch
make test
make test_install_all
```

## 12.13 Building with SUNDIALS Addons

SUNDIALS “addons” are community developed code additions for SUNDIALS that can be subsumed by the SUNDIALS build system so that they have full access to all internal SUNDIALS symbols. The intent is for SUNDIALS addons to function as if they are part of the SUNDIALS library, while allowing them to potentially have different licenses (although we encourage BSD-3-Clause still), code style (although we encourage them to follow the SUNDIALS style outlined [here](#)).

**Warning:** SUNDIALS addons are not maintained by the SUNDIALS team and may come with different licenses. Use them at your own risk.

To build with SUNDIALS addons,

1. Clone/copy the addon(s) into <sundials root>/external/
2. Copy the sundials-addon-example block in the <sundials root>/external/CMakeLists.txt, paste it below the example block, and modify the path listed for your own external addon(s).
3. When building SUNDIALS, set the CMake option `SUNDIALS_ENABLE_EXTERNAL_ADDONS` to `ON`
4. Build SUNDIALS as usual.

# Chapter 13

## IDAS Constants

Below we list all input and output constants used by the main solver and linear solver modules, together with their numerical values and a short description of their meaning.

### 13.1 IDAS input constants

Table 13.1: IDAS Input Constants

IDAS main solver module		
IDA_NORMAL	1	Solver returns at specified output time.
IDA_ONE_STEP	2	Solver returns after each successful step.
IDA_SIMULTANEOUS	1	Simultaneous corrector forward sensitivity method.
IDA_STAGGERED	2	Staggered corrector forward sensitivity method.
IDA_CENTERED	1	Central difference quotient approximation ( $2^{nd}$ order) of the sensitivity RHS.
IDA_FORWARD	2	Forward difference quotient approximation ( $1^{st}$ order) of the sensitivity RHS.
IDA_YA_YDP_INIT	1	Compute $y_a$ and $\dot{y}_d$ , given $y_d$ .
IDA_Y_INIT	2	Compute $y$ , given $\dot{y}$ .
IDAS adjoint solver module		
IDA_HERMITE	1	Use Hermite interpolation.
IDA_POLYNOMIAL	2	Use variable-degree polynomial interpolation.
Iterative linear solver module		
SUN_PREC_NONE	0	No preconditioning
SUN_PREC_LEFT	1	Preconditioning on the left.
SUN_MODIFIED_GS	1	Use modified Gram-Schmidt procedure.
SUN_CLASSICAL_GS	2	Use classical Gram-Schmidt procedure.

## 13.2 IDAS output constants

Table 13.2: IDAS Output Constants

IDAS main solver module		
IDA_SUCCESS	0	Successful function return.
IDA_TSTOP_RETURN	1	IDASolve succeeded by reaching the specified stopping point.
IDA_ROOT_RETURN	2	IDASolve succeeded and found one or more roots.
IDA_WARNING	99	IDASolve succeeded but an unusual situation occurred.
IDA_TOO_MUCH_WORK	-1	The solver took <code>mxstep</code> internal steps but could not reach tout.
IDA_TOO_MUCH_ACC	-2	The solver could not satisfy the accuracy demanded by the user for some internal step.
IDA_ERR_FAIL	-3	Error test failures occurred too many times during one internal time step or minimum step size was reached.
IDA_CONV_FAIL	-4	Convergence test failures occurred too many times during one internal time step or minimum step size was reached.
IDA_LINIT_FAIL	-5	The linear solver's initialization function failed.
IDA_LSETUP_FAIL	-6	The linear solver's setup function failed in an unrecoverable manner.
IDA_LSOLVE_FAIL	-7	The linear solver's solve function failed in an unrecoverable manner.
IDA_RES_FAIL	-8	The user-provided residual function failed in an unrecoverable manner.
IDA_REP_RES_FAIL	-9	The user-provided residual function repeatedly returned a recoverable error flag, but the solver was unable to recover.
IDA_RTFUNC_FAIL	-10	The rootfinding function failed in an unrecoverable manner.
IDA_CONSTR_FAIL	-11	The inequality constraints were violated and the solver was unable to recover.
IDA_FIRST_RES_FAIL	-12	The user-provided residual function failed recoverably on the first call.
IDA_LINESEARCH_FAIL	-13	The line search failed.
IDA_NO_RECOVERY	-14	The residual function, linear solver setup function, or linear solver solve function had a recoverable failure, but IDACalcIC could not recover.
IDA_NLS_INIT_FAIL	-15	The nonlinear solver's init routine failed.
IDA_NLS_SETUP_FAIL	-16	The nonlinear solver's setup routine failed.
IDA_MEM_NULL	-20	The <code>ida_mem</code> argument was NULL.
IDA_MEM_FAIL	-21	A memory allocation failed.
IDA_ILL_INPUT	-22	One of the function inputs is illegal.
IDA_NO_MALLOC	-23	The IDAS memory was not allocated by a call to <code>IDAINIT</code> .
IDA_BAD_EWT	-24	Zero value of some error weight component.
IDA_BAD_K	-25	The $k$ -th derivative is not available.
IDA_BAD_T	-26	The time $t$ is outside the last step taken.
IDA_BAD_DKY	-27	The vector argument where derivative should be stored is NULL.
IDA_NO_QUAD	-30	Quadratures were not initialized.
IDA_QRHS_FAIL	-31	The user-provided right-hand side function for quadratures failed in an unrecoverable manner.
IDA_FIRST_QRHS_ERR	-32	The user-provided right-hand side function for quadratures failed in an unrecoverable manner on the first call.
IDA_REP_QRHS_ERR	-33	The user-provided right-hand side repeatedly returned a recoverable error flag, but the solver was unable to recover.
IDA_NO_SENS	-40	Sensitivities were not initialized.

continues on next page



Table 13.2 – continued from previous page

IDAS main solver module		
IDA_SRES_FAIL	-41	The user-provided sensitivity residual function failed in an unrecoverable manner.
IDA_REP_SRES_ERR	-42	The user-provided sensitivity residual function repeatedly returned a recoverable error flag, but the solver was unable to recover.
IDA_BAD_IS	-43	The sensitivity identifier is not valid.
IDA_NO_QUADSENS	-50	Sensitivity-dependent quadratures were not initialized.
IDA_QSRHS_FAIL	-51	The user-provided sensitivity-dependent quadrature right-hand side function failed in an unrecoverable manner.
IDA_FIRST_QSRHS_ERR	-52	The user-provided sensitivity-dependent quadrature right-hand side function failed in an unrecoverable manner on the first call.
IDA_REP_QSRHS_ERR	-53	The user-provided sensitivity-dependent quadrature right-hand side repeatedly returned a recoverable error flag, but the solver was unable to recover.
IDAS adjoint solver module		
IDA_NO_ADJ	-101	The combined forward-backward problem has not been initialized.
IDA_NO_FWD	-102	IDASolveF has not been previously called.
IDA_NO_BCK	-103	No backward problem was specified.
IDA_BAD_TB0	-104	The desired output for backward problem is outside the interval over which the forward problem was solved.
IDA_REIFWD_FAIL	-105	No checkpoint is available for this hot start.
IDA_FWD_FAIL	-106	IDASolveB failed because IDASolve was unable to store data between two consecutive checkpoints.
IDA_GETY_BADT	-107	Wrong time in interpolation function.
IDALS linear solver interface		
IDALS_SUCCESS	0	Successful function return.
IDALS_MEM_NULL	-1	The ida_mem argument was NULL.
IDALS_LMEM_NULL	-2	The IDALS linear solver has not been initialized.
IDALS_ILL_INPUT	-3	The IDALS solver is not compatible with the current N_Vector module, or an input value was illegal.
IDALS_MEM_FAIL	-4	A memory allocation request failed.
IDALS_PMEM_NULL	-5	The preconditioner module has not been initialized.
IDALS_JACFUNC_UNRECV	-6	The Jacobian function failed in an unrecoverable manner.
IDALS_JACFUNC_RECVR	-7	The Jacobian function had a recoverable error.
IDALS_SUNMAT_FAIL	-8	An error occurred with the current SUNMatrix module.
IDALS_SUNLS_FAIL	-9	An error occurred with the current SUNLinearSolver module.
IDALS_NO_ADJ	-101	The combined forward-backward problem has not been initialized.
IDALS_LMEMB_NULL	-102	The linear solver was not initialized for the backward phase.



## Chapter 14

# Release History

Date	SUNDIALS	ARKODE	CVODE	CVODES	IDA	IDAS	KINSOL
Jun 2024	7.1.0	6.1.0	7.1.0	7.1.0	7.1.0	6.1.0	7.1.0
Feb 2024	7.0.0	6.0.0	7.0.0	7.0.0	7.0.0	6.0.0	7.0.0
Dec 2023	6.7.0	5.7.0	6.7.0	6.7.0	6.7.0	5.7.0	6.7.0
Nov 2023	6.6.2	5.6.2	6.6.2	6.6.2	6.6.2	5.6.2	6.6.2
Sep 2023	6.6.1	5.6.1	6.6.1	6.6.1	6.6.1	5.6.1	6.6.1
Jul 2023	6.6.0	5.6.0	6.6.0	6.6.0	6.6.0	5.6.0	6.6.0
Mar 2023	6.5.1	5.5.1	6.5.1	6.5.1	6.5.1	5.5.1	6.5.1
Dec 2022	6.5.0	5.5.0	6.5.0	6.5.0	6.5.0	5.5.0	6.5.0
Oct 2022	6.4.1	5.4.1	6.4.1	6.4.1	6.4.1	5.4.1	6.4.1
Oct 2022	6.4.0	5.4.0	6.4.0	6.4.0	6.4.0	5.4.0	6.4.0
Aug 2022	6.3.0	5.3.0	6.3.0	6.3.0	6.3.0	5.3.0	6.3.0
Apr 2022	6.2.0	5.2.0	6.2.0	6.2.0	6.2.0	5.2.0	6.2.0
Feb 2022	6.1.1	5.1.1	6.1.1	6.1.1	6.1.1	5.1.1	6.1.1
Jan 2022	6.1.0	5.1.0	6.1.0	6.1.0	6.1.0	5.1.0	6.1.0
Dec 2021	6.0.0	5.0.0	6.0.0	6.0.0	6.0.0	5.0.0	6.0.0
Sep 2021	5.8.0	4.8.0	5.8.0	5.8.0	5.8.0	4.8.0	5.8.0
Jan 2021	5.7.0	4.7.0	5.7.0	5.7.0	5.7.0	4.7.0	5.7.0
Dec 2020	5.6.1	4.6.1	5.6.1	5.6.1	5.6.1	4.6.1	5.6.1
Dec 2020	5.6.0	4.6.0	5.6.0	5.6.0	5.6.0	4.6.0	5.6.0
Oct 2020	5.5.0	4.5.0	5.5.0	5.5.0	5.5.0	4.5.0	5.5.0
Sep 2020	5.4.0	4.4.0	5.4.0	5.4.0	5.4.0	4.4.0	5.4.0
May 2020	5.3.0	4.3.0	5.3.0	5.3.0	5.3.0	4.3.0	5.3.0
Mar 2020	5.2.0	4.2.0	5.2.0	5.2.0	5.2.0	4.2.0	5.2.0
Jan 2020	5.1.0	4.1.0	5.1.0	5.1.0	5.1.0	4.1.0	5.1.0
Oct 2019	5.0.0	4.0.0	5.0.0	5.0.0	5.0.0	4.0.0	5.0.0
Feb 2019	4.1.0	3.1.0	4.1.0	4.1.0	4.1.0	3.1.0	4.1.0
Jan 2019	4.0.2	3.0.2	4.0.2	4.0.2	4.0.2	3.0.2	4.0.2
Dec 2018	4.0.1	3.0.1	4.0.1	4.0.1	4.0.1	3.0.1	4.0.1
Dec 2018	4.0.0	3.0.0	4.0.0	4.0.0	4.0.0	3.0.0	4.0.0
Oct 2018	3.2.1	2.2.1	3.2.1	3.2.1	3.2.1	2.2.1	3.2.1
Sep 2018	3.2.0	2.2.0	3.2.0	3.2.0	3.2.0	2.2.0	3.2.0
Jul 2018	3.1.2	2.1.2	3.1.2	3.1.2	3.1.2	2.1.2	3.1.2
May 2018	3.1.1	2.1.1	3.1.1	3.1.1	3.1.1	2.1.1	3.1.1
Nov 2017	3.1.0	2.1.0	3.1.0	3.1.0	3.1.0	2.1.0	3.1.0

continues on next page

Table 14.1 – continued from previous page

Date	SUNDIALS	ARKODE	CVODE	CVODES	IDA	IDAS	KINSOL
Sep 2017	3.0.0	2.0.0	3.0.0	3.0.0	3.0.0	2.0.0	3.0.0
Sep 2016	2.7.0	1.1.0	2.9.0	2.9.0	2.9.0	1.3.0	2.9.0
Aug 2015	2.6.2	1.0.2	2.8.2	2.8.2	2.8.2	1.2.2	2.8.2
Mar 2015	2.6.1	1.0.1	2.8.1	2.8.1	2.8.1	1.2.1	2.8.1
Mar 2015	2.6.0	1.0.0	2.8.0	2.8.0	2.8.0	1.2.0	2.8.0
Mar 2012	2.5.0	–	2.7.0	2.7.0	2.7.0	1.1.0	2.7.0
May 2009	2.4.0	–	2.6.0	2.6.0	2.6.0	1.0.0	2.6.0
Nov 2006	2.3.0	–	2.5.0	2.5.0	2.5.0	–	2.5.0
Mar 2006	2.2.0	–	2.4.0	2.4.0	2.4.0	–	2.4.0
May 2005	2.1.1	–	2.3.0	2.3.0	2.3.0	–	2.3.0
Apr 2005	2.1.0	–	2.3.0	2.2.0	2.3.0	–	2.3.0
Mar 2005	2.0.2	–	2.2.2	2.1.2	2.2.2	–	2.2.2
Jan 2005	2.0.1	–	2.2.1	2.1.1	2.2.1	–	2.2.1
Dec 2004	2.0.0	–	2.2.0	2.1.0	2.2.0	–	2.2.0
Jul 2002	1.0.0	–	2.0.0	1.0.0	2.0.0	–	2.0.0
Mar 2002	–	–	1.0.0 <sup>3</sup>	–	–	–	–
Feb 1999	–	–	–	–	1.0.0 <sup>4</sup>	–	–
Aug 1998	–	–	–	–	–	–	1.0.0 <sup>5</sup>
Jul 1997	–	–	1.0.0 <sup>2</sup>	–	–	–	–
Sep 1994	–	–	1.0.0 <sup>1</sup>	–	–	–	–

1. CVODE written
2. PVODE written
3. CVODE and PVODE combined
4. IDA written
5. KINSOL written

# Chapter 15

## Changelog

### 15.1 Changes to SUNDIALS in release 7.1.0

#### Major Features

Created shared user interface functions for ARKODE to allow more uniform control over time-stepping algorithms, improved extensibility, and simplified code maintenance. The corresponding stepper-specific user-callable functions are now deprecated and will be removed in a future major release.

Added CMake infrastructure that enables externally maintained addons/plugins to be *optionally* built with SUNDIALS. See [Contributing](#) for details.

#### New Features and Enhancements

Added support for Kokkos Kernels v4.

Added the following Runge-Kutta Butcher tables

- ARKODE\_FORWARD\_EULER\_1\_1
- ARKODE\_RALSTON\_EULER\_2\_1\_2
- ARKODE\_EXPLICIT\_MIDPOINT\_EULER\_2\_1\_2
- ARKODE\_BACKWARD\_EULER\_1\_1
- ARKODE\_IMPLICIT\_MIDPOINT\_1\_2
- ARKODE\_IMPLICIT\_TRAPEZOIDAL\_2\_2

Added the following MRI coupling tables

- ARKODE\_MRI\_GARK\_FORWARD\_EULER
- ARKODE\_MRI\_GARK\_RALSTON2
- ARKODE\_MRI\_GARK\_RALSTON3
- ARKODE\_MRI\_GARK\_BACKWARD\_EULER
- ARKODE\_MRI\_GARK\_IMPLICIT\_MIDPOINT
- ARKODE\_IMEX\_MRI\_GARK\_EULER
- ARKODE\_IMEX\_MRI\_GARK\_TRAPEZOIDAL
- ARKODE\_IMEX\_MRI\_GARK\_MIDPOINT

Added `ARKodeButcherTable_ERKIDToName()` and `ARKodeButcherTable_DIRKIDToName()` to convert a Butcher table ID to a string representation.

Added the function `ARKodeSetAutonomous()` in ARKODE to indicate that the implicit right-hand side function does not explicitly depend on time. When using the trivial predictor, an autonomous problem may reuse implicit function evaluations across stage solves to reduce the total number of function evaluations.

Users may now disable interpolated output in ARKODE by passing `ARK_INTERP_NONE` to `ARKodeSetInterpolantType()`. When interpolation is disabled, rootfinding is not supported, implicit methods must use the trivial predictor (the default option), and interpolation at stop times cannot be used (interpolating at stop times is disabled by default). With interpolation disabled, calling `ARKodeEvolve()` in `ARK_NORMAL` mode will return at or past the requested output time (setting a stop time may still be used to halt the integrator at a specific time). Disabling interpolation will reduce the memory footprint of an integrator by two or more state vectors (depending on the interpolant type and degree) which can be beneficial when interpolation is not needed e.g., when integrating to a final time without output in between or using an explicit fast time scale integrator with an MRI method.

Added “Resize” capability to ARKODE’s `SPRKStep` time-stepping module.

Enabled the Fortran interfaces to build with 32-bit `sunindextype`.

### Bug Fixes

Updated the CMake variable `HIP_PLATFORM` default to `amd` as the previous default, `hcc`, is no longer recognized in ROCm 5.7.0 or newer. The new default is also valid in older version of ROCm (at least back to version 4.3.1).

Renamed the DPCPP value for the `SUNDIALS_GINKGO_BACKENDS` CMake option to `SYCL` to match Ginkgo’s updated naming convention.

Changed the CMake version compatibility mode for SUNDIALS to `AnyNewerVersion` instead of `SameMajorVersion`. This fixes the issue seen [here](#).

Fixed a CMake bug that caused an MPI linking error for our C++ examples in some instances. Fixes [GitHub Issue #464](#).

Fixed the runtime library installation path for windows systems. This fix changes the default library installation path from `CMAKE_INSTALL_PREFIX/CMAKE_INSTALL_LIBDIR` to `CMAKE_INSTALL_PREFIX/CMAKE_INSTALL_BINDIR`.

Fixed conflicting `.lib` files between shared and static libs when using `MSVC` on Windows

Fixed invalid `SUNDIALS_EXPORT` generated macro when building both shared and static libs.

Fixed a bug in some Fortran examples where `c_null_ptr` was passed as an argument to a function pointer instead of `c_null_funptr`. This caused compilation issues with the Cray Fortran compiler.

Fixed a bug in the HIP execution policies where `WARP_SIZE` would not be set with ROCm 6.0.0 or newer.

Fixed a bug that caused error messages to be cut off in some cases. Fixes [GitHub Issue #461](#).

Fixed a memory leak when an error handler was added to a [SUNContext](#). Fixes [GitHub Issue #466](#).

Fixed a bug where `MRIStepEvolve()` would not handle a recoverable error produced from evolving the inner stepper.

Added missing `SetRootDirection` and `SetNoInactiveRootWarn` functions to ARKODE’s `SPRKStep` time-stepping module.

Fixed a bug in `ARKodeSPRKTable_Create()` where the coefficient arrays were not allocated.

Fix bug on LLP64 platforms (like Windows 64-bit) where `KLU_INDEXTYPE` could be 32 bits wide even if `SUNDIALS_INT64_T` is defined.

Check if size of `SuiteSparse_long` is 8 if the size of `sunindextype` is 8 when using `KLU`.

Fixed several build errors with the Fortran interfaces on Windows systems.

### Deprecation Notices

Numerous ARKODE stepper-specific functions are now deprecated in favor of ARKODE-wide functions.

Deprecated the *ARKStepSetOptimalParams* function. Since this function does not have an ARKODE-wide equivalent, instructions have been added to the user guide for how to retain the current functionality using other user-callable functions.

The unsupported implementations of *N\_VGetArrayPointer* and *N\_VSetArrayPointer* for the *hypr*e and PETSc vectors are now deprecated. Users should access the underlying wrapped external library vector objects instead with *N\_VGetVector\_ParHyp* and *N\_VGetVector\_Petsc*, respectively.

## 15.2 Changes to SUNDIALS in release 7.0.0

### Major Feature

SUNDIALS now has more robust and uniform error handling. Non-release builds will be built with additional error checking by default. See §4.3 for details.

### Breaking Changes

#### *Minimum C Standard*

SUNDIALS now requires using a compiler that supports a subset of the C99 standard. Note with the Microsoft C/C++ compiler the subset of C99 features utilized by SUNDIALS are available starting with [Visual Studio 2015](#).

#### *Minimum CMake Version*

CMake 3.18 or newer is now required when building SUNDIALS.

#### *Deprecated Types and Functions Removed*

The previously deprecated types *realtype* and *booleantype* were removed from *sundials\_types.h* and replaced with *sunrealtype* and *sunbooleantype*. The deprecated names for these types can be used by including the header file *sundials\_types\_deprecated.h* but will be fully removed in the next major release. Functions, types and header files that were previously deprecated have also been removed.

#### *Error Handling Changes*

With the addition of the new error handling capability, the *\*SetErrHandlerFn* and *\*SetErrFile* functions in CVODE(S), IDA(S), ARKODE, and KINSOL have been removed. Users of these functions can use the functions *SUNContext\_PushErrHandler()*, and *SUNLogger\_SetErrorFilename()* instead. For further details see Sections §4.3 and §4.4.

In addition the following names/symbols were replaced by *SUN\_ERR\_\** codes:

Removed	Replaced with SUNErrCode
SUNLS_SUCCESS	SUN_SUCCESS
SUNLS_UNRECOV_FAILURE	no replacement (value was unused)
SUNLS_MEM_NULL	SUN_ERR_ARG_CORRUPT
SUNLS_ILL_INPUT	SUN_ERR_ARG_*
SUNLS_MEM_FAIL	SUN_ERR_MEM_FAIL
SUNLS_PACKAGE_FAIL_UNREC	SUN_ERR_EXT_FAIL
SUNLS_VECTOROP_ERR	SUN_ERR_OP_FAIL
SUN-NLS_SUCCESS	SUN_SUCCESS
SUN-NLS_MEM_NULL	SUN_ERR_ARG_CORRUPT
SUN-NLS_MEM_FAIL	SUN_ERR_MEM_FAIL
SUN-NLS_ILL_INPUT	SUN_ERR_ARG_*
SUN-NLS_VECTOROP_ERR	SUN_ERR_OP_FAIL
SUN-NLS_EXT_FAIL	SUN_ERR_EXT_FAIL
SUNMAT_SUCCESS	SUN_SUCCESS
SUNMAT_ILL_INPUT	SUN_ERR_ARG_*
SUNMAT_MEM_FAIL	SUN_ERR_MEM_FAIL
SUNMAT_OPERATION_FAIL	SUN_ERR_OP_FAIL
SUNMAT_MATVEC_SETUP_REQUIRED	SUN_ERR_OP_FAIL

The following functions have had their signature updated to ensure they can leverage the new SUNDIALS error handling capabilities.

- From `sundials_futils.h`
  - `SUNDIALSFileOpen()`
  - `SUNDIALSFileClose()`
- From `sundials_memory.h`
  - `SUNMemoryNewEmpty()`
  - `SUNMemoryHelper_Alias()`
  - `SUNMemoryHelper_Wrap()`
- From `sundials_nvector.h`
  - `N_VNewVectorArray()`

#### *SUNComm Type Added*

We have replaced the use of a type-erased (i.e., `void*`) pointer to a communicator in place of `MPI_Comm` throughout the SUNDIALS API with a `SUNComm`, which is just a typedef to an `int` in builds without MPI and a typedef to a `MPI_Comm` in builds with MPI. As a result:

- When MPI is enabled, all SUNDIALS libraries will include MPI symbols and applications will need to include the path for MPI headers and link against the corresponding MPI library.
- All users will need to update their codes because the call to `SUNContext_Create()` now takes a `SUNComm` instead of type-erased pointer to a communicator. For non-MPI codes, pass `SUN_COMM_NULL` to the `comm` argument instead of `NULL`. For MPI codes, pass the `MPI_Comm` directly.
- The same change must be made for calls to `SUNLogger_Create()` or `SUNProfiler_Create()`.
- Some users will need to update their calls to `N_VGetCommunicator()`, and update any custom `N_Vector` implementations that provide `N_VGetCommunicator()`, since it now returns a `SUNComm`.



The change away from type-erased pointers for [SUNComm](#) fixes problems like the one described in [GitHub Issue #275](#).

The SUNLogger is now always MPI-aware if MPI is enabled in SUNDIALS and the SUNDIALS\_LOGGING\_ENABLE\_MPI CMake option and macro definition were removed accordingly.

### *SUNDIALS Core Library*

Users now need to link to `sundials_core` in addition to the libraries already linked to. This will be picked up automatically in projects that use the SUNDIALS CMake target. The library `sundials_generic` has been superseded by `sundials_core` and is no longer available. This fixes some duplicate symbol errors on Windows when linking to multiple SUNDIALS libraries.

### *Fortran Interface Modules Streamlined*

We have streamlined the Fortran modules that need to be included by users by combining the SUNDIALS core into one Fortran module, `fsundials_core_mod`. Modules for implementations of the core APIs still exist (e.g., for the Dense linear solver there is `fsunlinsol_dense_mod`) as do the modules for the SUNDIALS packages (e.g., `fcvode_mod`). The following modules are the ones that have been consolidated into `fsundials_core_mod`:

```
fsundials_adaptcontroller_mod
fsundials_context_mod
fsundials_futils_mod
fsundials_linearsolver_mod
fsundials_logger_mod
fsundials_matrix_mod
fsundials_nonlinearsolver_mod
fsundials_nvector_mod
fsundials_profiler_mod
fsundials_types_mod
```

### **Minor Changes**

The `CMAKE_BUILD_TYPE` defaults to `RelWithDebInfo` mode now i.e., SUNDIALS will be built with optimizations and debugging symbols enabled by default. Previously the build type was unset by default so no optimization or debugging flags were set.

The advanced CMake options to override the inferred LAPACK name-mangling scheme have been updated from `SUNDIALS_F77_FUNC_CASE` and `SUNDIALS_F77_FUNC_UNDERSCORES` to [SUNDIALS\\_LAPACK\\_CASE](#) and [SUNDIALS\\_LAPACK\\_UNDERSCORES](#), respectively.

As a subset of C99 is now required the CMake option `USE_GENERIC_MATH` as been removed.

The C++ convenience classes (e.g., `sundials::Context`) have been moved to from SUNDIALS `.h` headers to corresponding `.hpp` headers (e.g., `sundials/sundials_context.hpp`) so C++ codes do not need to compile with C++14 support when using the C API.

Converted most previous Fortran 77 and 90 examples to use SUNDIALS' Fortran 2003 interface.

### **Bug Fixes**

Fixed [GitHub Issue #329](#) so that C++20 aggregate initialization can be used.

Fixed integer overflow in the internal SUNDIALS hashmap. This resolves [GitHub Issues #409](#) and [#249](#).

### **Deprecation Notice**

The functions in `sundials_math.h` will be deprecated in the next release.

```
sunrealtype SUNRpowerI(sunrealtype base, int exponent);
sunrealtype SUNRpowerR(sunrealtype base, sunrealtype exponent);
sunbooleantype SUNRCompare(sunrealtype a, sunrealtype b);
```

(continues on next page)

(continued from previous page)

```
sunbooleantype SUNRCompareTol(sunrealtype a, sunrealtype b, sunrealtype tol);  
sunrealtype SUNStrToReal(const char* str);
```

Additionally, the following header files (and everything in them) will be deprecated – users who rely on these are recommended to transition to the corresponding *SUNMatrix* and *SUNLinearSolver* modules:

```
sundials_direct.h  
sundials_dense.h  
sundials_band.h
```

## 15.3 Changes to SUNDIALS in release 6.7.0

### Major Feature

Added the *SUNAdaptController* base class, ported ARKODE’s internal implementations of time step controllers to implementations of this class, and updated ARKODE to use these objects instead of its own implementations. Added *ARKStepSetAdaptController()* and *ERKStepSetAdaptController()* routines so that users can modify controller parameters, or even provide custom implementations.

### New Features

Improved the computational complexity of the sparse matrix *ScaleAddI* function from  $\mathcal{O}(M * N)$  to  $\mathcal{O}(\text{NNZ})$ .

Added Fortran support for the LAPACK dense linear solver implementation.

Added the routines *ARKStepSetAdaptivityAdjustment()* and *ERKStepSetAdaptivityAdjustment()*, that allow users to adjust the value for the method order supplied to the temporal adaptivity controllers. The ARKODE default for this adjustment has been  $-1$  since its initial release, but for some applications a value of  $0$  is more appropriate. Users who notice that their simulations encounter a large number of temporal error test failures may want to experiment with adjusting this value.

Added the third order ERK method *ARKODE\_SHU\_OSHER\_3\_2\_3*, the fourth order ERK method *ARKODE\_SOFRONIOU-SPALETTA\_5\_3\_4*, the sixth order ERK method *ARKODE\_VERNER\_9\_5\_6*, the seventh order ERK method *ARKODE-VERNER\_10\_6\_7*, the eighth order ERK method *ARKODE\_VERNER\_13\_7\_8*, and the ninth order ERK method *ARKODE-VERNER\_16\_8\_9*.

*ARKStep*, *ERKStep*, *MRISStep*, and *SPRKStep* were updated to remove a potentially unnecessary right-hand side evaluation at the end of an integration. *ARKStep* was additionally updated to remove extra right-hand side evaluations when using an explicit method or an implicit method with an explicit first stage.

The *MRISStepInnerStepper* class in *MRISStep* was updated to make supplying an *MRISStepInnerFullRhsFn* optional.

### Bug Fixes

Changed the *SUNProfiler* so that it does not rely on *MPI\_WTime* in any case. This fixes [GitHub Issue #312](#).

Fixed scaling bug in *SUNMatScaleAddI\_Sparse* for non-square matrices.

Fixed a regression introduced by the stop time bug fix in v6.6.1 where ARKODE, CVODE, CVODES, IDA, and IDAS would return at the stop time rather than the requested output time if the stop time was reached in the same step in which the output time was passed.

Fixed a bug in *ERKStep* where methods with  $c_s = 1$  but  $a_{s,j} \neq b_j$  were incorrectly treated as having the first same as last (FSAL) property.

Fixed a bug in ARKODE where *ARKStepSetInterpolateStopTime()* would return an interpolated solution at the stop time in some cases when interpolation was disabled.

Fixed a bug in `ARKStepSetTableNum()` wherein it did not recognize `ARKODE_ARK2_ERK_3_1_2` and `ARKODE_ARK2_DIRK_3_1_2` as a valid additive Runge–Kutta Butcher table pair.

Fixed a bug in `MRISStepCoupling_Write()` where explicit coupling tables were not written to the output file pointer.

Fixed missing soversions in some `SUNLinearSolver` and `SUNNonlinearSolver` CMake targets.

Renamed some internal types in CVODES and IDAS to allow both packages to be built together in the same binary.

## 15.4 Changes to SUNDIALS in release 6.6.2

Fixed the build system support for MAGMA when using a NVIDIA HPC SDK installation of CUDA and fixed the targets used for rocBLAS and rocSPARSE.

## 15.5 Changes to SUNDIALS in release 6.6.1

### New Features

Updated the Trilinos Tpetra `N_Vector` interface to support Trilinos 14.

### Bug Fixes

Fixed a memory leak when destroying a CUDA, HIP, SYCL, or system `SUNMemoryHelper` object.

Fixed a bug in ARKODE, CVODE, CVODES, IDA, and IDAS where the stop time may not be cleared when using normal mode if the requested output time is the same as the stop time. Additionally, with ARKODE, CVODE, and CVODES this fix removes an unnecessary interpolation of the solution at the stop time that could occur in this case.

## 15.6 Changes to SUNDIALS in release 6.6.0

### Major Features

A new time-stepping module, `SPRKStep`, was added to ARKODE. This time-stepper provides explicit symplectic partitioned Runge-Kutta methods up to order 10 for separable Hamiltonian systems.

Added support for relaxation Runge-Kutta methods in ERKStep and ARKStep, see [Relaxation Methods](#), [Relaxation Methods](#), and [Relaxation Methods](#) for more information.

### New Features

Updated the default ARKODE, CVODE, and CVODES behavior when returning the solution when the internal time has reached a user-specified stop time. Previously, the output solution was interpolated to the value of `tstop`; the default is now to copy the internal solution vector. Users who wish to revert to interpolation may call a new routine `CVodeSetInterpolateStopTime()`, `ARKStepSetInterpolateStopTime()`, `ERKStepSetInterpolateStopTime()`, or `MRISStepSetInterpolateStopTime()`.

Added the second order IMEX method from [36] as the default second order IMEX method in ARKStep. The explicit table is given by `ARKODE_ARK2_ERK_3_1_2` (see [ARK2-ERK-3-1-2](#)) and the implicit table by `ARKODE_ARK2_DIRK_3_1_2` (see [ARK2-DIRK-3-1-2](#)).

Updated the F2003 utility routines `SUNDIALSFileOpen()` and `SUNDIALSFileClose()` to support user specification of `stdout` and `stderr` strings for the output file names.

### Bug Fixes

A potential bug was fixed when using inequality constraint handling and calling `ARKStepGetEstLocalErrors()` or `ERKStepGetEstLocalErrors()` after a failed step in which an inequality constraint violation occurred. In this case, the values returned by `ARKStepGetEstLocalErrors()` or `ERKStepGetEstLocalErrors()` may have been invalid.

## 15.7 Changes to SUNDIALS in release 6.5.1

### New Features

Added the following functions to disable a previously set stop time:

- `ARKStepClearStopTime()`
- `ERKStepClearStopTime()`
- `MRISStepClearStopTime()`
- `CVodeClearStopTime()`
- `IDAClearStopTime()`

The default interpolant in ARKODE when using a first order method has been updated to a linear interpolant to ensure values obtained by the integrator are returned at the ends of the time interval. To restore the previous behavior of using a constant interpolant call `ARKStepSetInterpolantDegree()`, `ERKStepSetInterpolantDegree()`, or `MRISStepSetInterpolantDegree()` and set the interpolant degree to zero before evolving the problem.

### Bug Fixes

Fixed build errors when using SuperLU\_DIST with ROCM enabled to target AMD GPUs.

Fixed compilation errors in some SYCL examples when using the `icx` compiler.

## 15.8 Changes to SUNDIALS in release 6.5.0

### New Features

A new capability to keep track of memory allocations made through the `SUNMemoryHelper` classes has been added. Memory allocation stats can be accessed through the `SUNMemoryHelper_GetAllocStats()` function. See §10.1 for more details.

Added the following functions to assist in debugging simulations utilizing matrix-based linear solvers:

- `ARKStepGetJac()`
- `ARKStepGetJacTime()`
- `ARKStepGetJacNumSteps()`
- `MRISStepGetJac()`
- `MRISStepGetJacTime()`
- `MRISStepGetJacNumSteps()`
- `CVodeGetJac()`
- `CVodeGetJacTime()`
- `CVodeGetJacNumSteps()`
- `IDAGetJac()`
- `IDAGetJacCj()`

- `IDAGetJacTime()`
- `IDAGetJacNumSteps()`
- `KINGetJac()`
- `KINGetJacNumIters()`

Added support for CUDA 12.

Added support for the SYCL backend with RAJA 2022.x.y.

### Bug Fixes

Fixed an underflow bug during root finding in ARKODE, CVODE, CVODES, IDA and IDAS. This fixes [GitHub Issue #57](#).

Fixed an issue with finding oneMKL when using the `icpx` compiler with the `-fsycl` flag as the C++ compiler instead of `dpcpp`.

Fixed the shape of the arrays returned by the Fortran interfaces to `N_VGetArrayPointer()`, `SUNDenseMatrix_Data()`, `SUNBandMatrix_Data()`, `SUNSparseMatrix_Data()`, `SUNSparseMatrix_IndexValues()`, and `SUNSparseMatrix_IndexPointers()`. Compiling and running code that uses the SUNDIALS Fortran interfaces with bounds checking will now work.

Fixed an implicit conversion error in the Butcher table for ESDIRK5(4)7L[2]SA2.

## 15.9 Changes to SUNDIALS in release 6.4.1

Fixed a bug with the Kokkos interfaces that would arise when using clang.

Fixed a compilation error with the Intel oneAPI 2022.2 Fortran compiler in the Fortran 2003 interface test for the serial `N_Vector`.

Fixed a bug in the LAPACK band and dense linear solvers which would cause the tests to fail on some platforms.

## 15.10 Changes to SUNDIALS in release 6.4.0

### New Requirements

CMake 3.18.0 or newer is now required for CUDA support.

A C++14 compliant compiler is now required for C++ based features and examples e.g., CUDA, HIP, RAJA, Trilinos, SuperLU\_DIST, MAGMA, Ginkgo, and Kokkos.

### Major Features

Added support for the [Ginkgo](#) linear algebra library. This support includes new SUNDIALS matrix and linear solver implementations, see the sections [§7.10](#) and [§8.18](#).

Added new SUNDIALS vector, dense matrix, and dense linear solver implementations utilizing the [Kokkos Ecosystem](#) for performance portability, see sections [§6.14](#), [§7.11](#), and [§8.19](#) for more information.

### New Features

Added support for GPU enabled SuperLU\_DIST and SuperLU\_DIST v8.x.x. Removed support for SuperLU\_DIST v6.x.x or older. Fix mismatched definition and declaration bug in SuperLU\_DIST matrix constructor.

Added the functions following functions to load a Butcher table from a string:

- `ARKStepSetTableName()`

- `ERKStepSetTableName()`
- `MRISetCoupling_LoadTableByName()`
- `ARKodeButcherTable_LoadDIRKByName()`
- `ARKodeButcherTable_LoadERKByName()`

### Bug Fixes

Fixed a bug in the CUDA and HIP vectors where `N_VMaxNorm()` would return the minimum positive floating-point value for the zero vector.

Fixed memory leaks/out of bounds memory accesses in the ARKODE MRISet module that could occur when attaching a coupling table after reinitialization with a different number of stages than originally selected.

Fixed a memory leak where the projection memory would not be deallocated when calling `CVodeFree()`.

## 15.11 Changes to SUNDIALS in release 6.3.0

### New Features

Added the following functions to retrieve the user data pointer provided with `SetUserData` functions:

- `ARKStepGetUserData()`
- `ERKStepGetUserData()`
- `MRISetGetUserData()`
- `CVodeGetUserData()`
- `IDAGetUserData()`
- `KINGetUserData()`

Added a variety of embedded DIRK methods from [49] and [50].

Updated `MRISetReset()` to call the corresponding `MRISetInnerResetFn` with the same `tR` and `yR` arguments for the `MRISetInnerStepper` object that is used to evolve the MRI “fast” time scale subproblems.

Added a new example (`examples/cvode/serial/cvRocket_dns.c`) which demonstrates using CVODE with a discontinuous right-hand-side function and rootfinding.

### Bug Fixes

Fixed a bug in `ERKStepReset()`, `ERKStepReInit()`, `ARKStepReset()`, `ARKStepReInit()`, `MRISetReset()`, and `MRISetReInit()` where a previously-set value of `tstop` (from a call to `ERKStepSetStopTime()`, `ARKStepSetStopTime()`, or `MRISetSetStopTime()`, respectively) would not be cleared.

Fixed the unintuitive behavior of the `USE_GENERIC_MATH` CMake option which caused the double precision math functions to be used regardless of the value of `SUNDIALS_PRECISION`. Now, SUNDIALS will use precision appropriate math functions when they are available and the user may provide the math library to link to via the advanced CMake option `SUNDIALS_MATH_LIBRARY`.

Changed `SUNDIALS_LOGGING_ENABLE_MPI` CMake option default to be OFF. This fixes [GitHub Issue #177](#).

## 15.12 Changes to SUNDIALS in release 6.2.0

### Major Features

Added the *SUNLogger* API which provides a SUNDIALS-wide mechanism for logging of errors, warnings, informational output, and debugging output.

Added support to CVODES for integrating IVPs with constraints using BDF methods and projecting the solution onto the constraint manifold with a user defined projection function. This implementation is accompanied by additions to the CVODES user documentation and examples.

### New Features

Added the function *SUNProfiler\_Reset()* to reset the region timings and counters to zero.

Added the following functions to output all of the integrator, nonlinear solver, linear solver, and other statistics in one call:

- *ARKStepPrintAllStats()*
- *ERKStepPrintAllStats()*
- *MRIStepPrintAllStats()*
- *CNodePrintAllStats()*
- *IDAPrintAllStats()*
- *KINPrintAllStats()*

The file `scripts/sundials_csv.py` contains functions for parsing the comma-separated value (CSV) output files when using the CSV output format.

Added functions to CVODE, CVODES, IDA, and IDAS to change the default step size adaptivity parameters. For more information see the documentation for:

- *CNodeSetEtaFixedStepBounds()*
- *CNodeSetEtaMaxFirstStep()*
- *CNodeSetEtaMaxEarlyStep()*
- *CNodeSetNumStepsEtaMaxEarlyStep()*
- *CNodeSetEtaMax()*
- *CNodeSetEtaMin()*
- *CNodeSetEtaMinErrFail()*
- *CNodeSetEtaMaxErrFail()*
- *CNodeSetNumFailsEtaMaxErrFail()*
- *CNodeSetEtaConvFail()*
- *IDASetEtaFixedStepBounds()*
- *IDASetEtaMax()*
- *IDASetEtaMin()*
- *IDASetEtaLow()*
- *IDASetEtaMinErrFail()*
- *IDASetEtaConvFail()*

Added the functions `ARKStepSetDeduceImplicitRhs()` and `MRISetSetDeduceImplicitRhs()` to optionally remove an evaluation of the implicit right-hand side function after nonlinear solves. See [Nonlinear solver methods](#), for considerations on using this optimization.

Added the function `MRISetSetOrder()` to select the default MRI method of a given order.

Added the functions `CVodeSetDeltaGammaMaxLSetup()` and `CVodeSetDeltaGammaMaxBadJac()` in CVODE and CVODES to adjust the  $\gamma$  change thresholds to require a linear solver setup or Jacobian/precondition update, respectively.

Added the function `IDASetDeltaCjLSetup()` in IDA and IDAS to adjust the parameter that determines when a change in  $c_j$  requires calling the linear solver setup function.

Added the function `IDASetMinStep()` to set a minimum step size.

### Bug Fixes

Fixed the `SUNContext` convenience class for C++ users to disallow copy construction and allow move construction.

The behavior of `N_VSetKernelExecPolicy_Sycl()` has been updated to be consistent with the CUDA and HIP vectors. The input execution policies are now cloned and may be freed after calling `N_VSetKernelExecPolicy_Sycl()`. Additionally, NULL inputs are now allowed and, if provided, will reset the vector execution policies to the defaults.

A memory leak in the SYCL vector was fixed where the execution policies were not freed when the vector was destroyed.

The include guard in `nvector_mpmmanyvector.h` has been corrected to enable using both the ManyVector and MPI-ManyVector vector implementations in the same simulation.

A bug was fixed in the ARKODE, CVODE(S), and IDA(S) functions to retrieve the number of nonlinear solver failures. The failure count returned was the number of failed *steps* due to a nonlinear solver failure i.e., if a nonlinear solve failed with a stale Jacobian or preconditioner but succeeded after updating the Jacobian or preconditioner, the initial failure was not included in the nonlinear solver failure count. The following functions have been updated to return the total number of nonlinear solver failures:

- `ARKStepGetNumNonlinSolvConvFails()`
- `ARKStepGetNonlinSolvStats()`
- `MRISetGetNumNonlinSolvConvFails()`
- `MRISetGetNonlinSolvStats()`
- `CVodeGetNumNonlinSolvConvFails()`
- `CVodeGetNonlinSolvStats()`
- `CVodeGetSensNumNonlinSolvConvFails()`
- `CVodeGetSensNonlinSolvStats()`
- `CVodeGetStgrSensNumNonlinSolvConvFails()`
- `CVodeGetStgrSensNonlinSolvStats()`
- `IDAGetNumNonlinSolvConvFails()`
- `IDAGetNonlinSolvStats()`
- `IDAGetSensNumNonlinSolvConvFails()`
- `IDAGetSensNonlinSolvStats()`

As a result of this change users may see an increase in the number of failures reported from the above functions. The following functions have been added to retrieve the number of failed steps due to a nonlinear solver failure i.e., the counts previously returned by the above functions:

- `ARKStepGetNumStepSolveFails()`



- `MRIStepGetNumStepSolveFails()`
- `CVodeGetNumStepSolveFails()`
- `CVodeGetNumStepSensSolveFails()`
- `CVodeGetNumStepStgrSensSolveFails()`
- *`IDAGetNumStepSolveFails()`*
- *`IDAGetNumStepSensSolveFails()`*

Changed exported SUNDIALS PETSc CMake targets to be INTERFACE IMPORTED instead of UNKNOWN IMPORTED.

### Deprecation Notice

Deprecated the following functions, it is recommended to use the *SUNLogger* API instead.

- `ARKStepSetDiagnostics`
- `ERKStepSetDiagnostics`
- `MRIStepSetDiagnostics`
- `KINSetInfoFile`
- `SUNNonlinSolSetPrintLevel_Newton`
- `SUNNonlinSolSetInfoFile_Newton`
- `SUNNonlinSolSetPrintLevel_FixedPoint`
- `SUNNonlinSolSetInfoFile_FixedPoint`
- `SUNLinSolSetInfoFile_PCG`
- `SUNLinSolSetPrintLevel_PCG`
- `SUNLinSolSetInfoFile_SPGMR`
- `SUNLinSolSetPrintLevel_SPGMR`
- `SUNLinSolSetInfoFile_SPGMR`
- `SUNLinSolSetPrintLevel_SPGMR`
- `SUNLinSolSetInfoFile_SPTFQM`
- `SUNLinSolSetPrintLevel_SPTFQM`
- `SUNLinSolSetInfoFile_SPBCGS`
- `SUNLinSolSetPrintLevel_SPBCGS`

The `SUNLinSolSetInfoFile_*` and `SUNNonlinSolSetInfoFile_*` family of functions are now enabled by setting the CMake option *SUNDIALS\_LOGGING\_LEVEL* to a value  $\geq 3$ .

## 15.13 Changes to SUNDIALS in release 6.1.1

### New Feature

Added new Fortran example program, `examples/arkode/F2003_serial/ark_kpr_mri_f2003.f90` demonstrating MRI capabilities.

### Bug Fixes

Fixed exported `SUNDIALSConfig.cmake`.

Fixed Fortran interface to `MRISetInnerStepper` and `MRISetCoupling` structures and functions.

## 15.14 Changes to SUNDIALS in release 6.1.0

### New Features

Added new reduction implementations for the CUDA and HIP vectors that use shared memory (local data storage) instead of atomics. These new implementations are recommended when the target hardware does not provide atomic support for the floating point precision that SUNDIALS is being built with. The HIP vector uses these by default, but the `N_VSetKernelExecPolicy_Cuda()` and `N_VSetKernelExecPolicy_Hip()` functions can be used to choose between different reduction implementations.

`SUNDIALS::<lib>` targets with no static/shared suffix have been added for use within the build directory (this mirrors the targets exported on installation).

`CMAKE_C_STANDARD` is now set to 99 by default.

### Bug Fixes

Fixed exported `SUNDIALSConfig.cmake` when profiling is enabled without Caliper.

Fixed `sundials_export.h` include in `sundials_config.h`.

Fixed memory leaks in the SuperLU\_MT linear solver interface.

## 15.15 Changes to SUNDIALS in release 6.0.0

### Breaking Changes

#### *SUNContext Object Added*

SUNDIALS v6.0.0 introduces a new `SUNContext` object on which all other SUNDIALS objects depend. As such, the constructors for all SUNDIALS packages, vectors, matrices, linear solvers, nonlinear solvers, and memory helpers have been updated to accept a context as the last input. Users upgrading to SUNDIALS v6.0.0 will need to call `SUNContext_Create()` to create a context object with before calling any other SUNDIALS library function, and then provide this object to other SUNDIALS constructors. The context object has been introduced to allow SUNDIALS to provide new features, such as the profiling/instrumentation also introduced in this release, while maintaining thread-safety. See the §4.2 for more details.

The script `scripts/upgrade-to-sundials-6-from-5.sh` has been provided with this release (and obtainable from the GitHub release page) to help ease the transition to SUNDIALS v6.0.0. The script will add a `SUNCTX_PLACEHOLDER` argument to all of the calls to SUNDIALS constructors that now require a `SUNContext` object. It can also update deprecated SUNDIALS constants/types to the new names. It can be run like this:

```
./upgrade-to-sundials-6-from-5.sh <files to update>
```

### Updated *SUNMemoryHelper* Function Signatures

The *SUNMemoryHelper* functions *SUNMemoryHelper\_Alloc()*, *SUNMemoryHelper\_Dealloc()*, and *SUNMemoryHelper\_Copy()* have been updated to accept an opaque handle as the last input. At a minimum, user-defined *SUNMemoryHelper* implementations will need to update these functions to accept the additional argument. Typically, this handle is the execution stream (e.g., a CUDA/HIP stream or SYCL queue) for the operation. The CUDA, HIP, and SYCL implementations have been updated accordingly. Additionally, the constructor *SUNMemoryHelper\_Sycl()* has been updated to remove the SYCL queue as an input.

### Deprecated Functions Removed

The previously deprecated constructor *N\_VMakeWithManagedAllocator\_Cuda* and the function *N\_VSetCudaStream\_Cuda* have been removed and replaced with *N\_VNewWithMemHelp\_Cuda()* and *N\_VSetKernelExecPolicy\_Cuda()* respectively.

The previously deprecated macros *PVEC\_REAL\_MPI\_TYPE* and *PVEC\_INTEGER\_MPI\_TYPE* have been removed and replaced with *MPI\_SUNREALTYPE* and *MPI\_SUNINDEXTYPE* respectively.

The following previously deprecated *SUNLinearSolver* functions have been removed:

Removed	Replacement
<i>SUNBandLinearSolver</i>	<i>SUNLinSol_Band()</i>
<i>SUNDenseLinearSolver</i>	<i>SUNLinSol_Dense()</i>
<i>SUNKLU</i>	<i>SUNLinSol_KLU()</i>
<i>SUNKLUReInit</i>	<i>SUNLinSol_KLUReInit()</i>
<i>SUNKLUSetOrdering</i>	<i>SUNLinSol_KLUSetOrdering()</i>
<i>SUNLapackBand</i>	<i>SUNLinSol_LapackBand()</i>
<i>SUNLapackDense</i>	<i>SUNLinSol_LapackDense()</i>
<i>SUNPCG</i>	<i>SUNLinSol_PCG()</i>
<i>SUNPCGSetPrecType</i>	<i>SUNLinSol_PCGSetPrecType()</i>
<i>SUNPCGSetMaxl</i>	<i>SUNLinSol_PCGSetMaxl()</i>
<i>SUNSPBCGS</i>	<i>SUNLinSol_SPBCGS()</i>
<i>SUNSPBCGSSetPrecType</i>	<i>SUNLinSol_SPBCGSSetPrecType()</i>
<i>SUNSPBCGSSetMaxl</i>	<i>SUNLinSol_SPBCGSSetMaxl()</i>
<i>SUNSPFGMR</i>	<i>SUNLinSol_SPFGMR()</i>
<i>SUNSPFGMRSetPrecType</i>	<i>SUNLinSol_SPFGMRSetPrecType()</i>
<i>SUNSPFGMRSetGStype</i>	<i>SUNLinSol_SPFGMRSetGStype()</i>
<i>SUNSPFGMRSetMaxRestarts</i>	<i>SUNLinSol_SPFGMRSetMaxRestarts()</i>
<i>SUNSPGMR</i>	<i>SUNLinSol_SPGMR()</i>
<i>SUNSPGMRSetPrecType</i>	<i>SUNLinSol_SPGMRSetPrecType()</i>
<i>SUNSPGMRSetGStype</i>	<i>SUNLinSol_SPGMRSetGStype()</i>
<i>SUNSPGMRSetMaxRestarts</i>	<i>SUNLinSol_SPGMRSetMaxRestarts()</i>
<i>SUNSPTFQMR</i>	<i>SUNLinSol_SPTFQMR()</i>
<i>SUNSPTFQMRSetPrecType</i>	<i>SUNLinSol_SPTFQMRSetPrecType()</i>
<i>SUNSPTFQMRSetMaxl</i>	<i>SUNLinSol_SPTFQMRSetMaxl()</i>
<i>SUNSuperLUMT</i>	<i>SUNLinSol_SuperLUMT()</i>
<i>SUNSuperLUMTSetOrdering</i>	<i>SUNLinSol_SuperLUMTSetOrdering()</i>

The deprecated functions *MRISetGetCurrentButcherTables* and *MRISetWriteButcher* and the utility functions *MRISetSetTable* and *MRISetSetTableNum* have been removed. Users wishing to create an MRI-GARK method from a Butcher table should use *MRISetCoupling\_MISetMRI()* to create the corresponding MRI coupling table and attach it with *MRISetSetCoupling()*.

The previously deprecated functions *ARKSetSetMaxStepsBetweenLSet* and *ARKSetSetMaxStepsBetweenJac* have been removed and replaced with *ARKSetSetLSetupFrequency()* and *ARKSetSetJacEvalFrequency()*

respectively.

The previously deprecated function `CVodeSetMaxStepsBetweenJac` has been removed and replaced with `CVodeSetJacEvalFrequency()`.

The ARKODE, CVODE, IDA, and KINSOL Fortran 77 interfaces has been removed. See §4.7 and the F2003 example programs for more details using the SUNDIALS Fortran 2003 module interfaces.

#### *Namespace Changes*

The CUDA, HIP, and SYCL execution policies have been moved from the `sundials` namespace to the `sundials::cuda`, `sundials::hip`, and `sundials::sycl` namespaces respectively. Accordingly, the prefixes “Cuda”, “Hip”, and “Sycl” have been removed from the execution policy classes and methods.

The Sundials namespace used by the Trilinos Tpetra `N_Vector` implementation has been replaced with the `sundials::trilinos::nvector_tpetra` namespace.

### Major Features

#### *Profiling Capability*

A capability to profile/instrument SUNDIALS library code has been added. This can be enabled with the CMake option `SUNDIALS_BUILD_WITH_PROFILING`. A built-in profiler will be used by default, but the `Caliper` library can also be used instead with the CMake option `ENABLE_CALIPER`. See the documentation section on profiling for more details.

**Warning:** Profiling will impact performance, and should be enabled judiciously.

#### *IMEX MRI Methods and MRISetInnerStepper Object*

The MRISet module has been extended to support implicit-explicit (ImEx) multirate infinitesimal generalized additive Runge–Kutta (MRI-GARK) methods. As such, `MRISetCreate()` has been updated to include arguments for the slow explicit and slow implicit ODE right-hand side functions. `MRISetCreate()` has also been updated to require attaching an MRISetInnerStepper for evolving the fast time scale. `MRISetReInit()` has been similarly updated to take explicit and implicit right-hand side functions as input. Codes using explicit or implicit MRI methods will need to update `MRISetCreate()` and `MRISetReInit()` calls to pass NULL for either the explicit or implicit right-hand side function as appropriate. If ARKStep is used as the fast time scale integrator, codes will need to call `ARKStepCreateMRISetInnerStepper()` to wrap the ARKStep memory as an MRISetInnerStepper object. Additionally, `MRISetGetNumRhsEvals()` has been updated to return the number of slow implicit and explicit function evaluations. The coupling table, `MRISetCoupling`, and the functions `MRISetCoupling_Alloc()` and `MRISetCoupling_Create()` have also been updated to support IMEX-MRI-GARK methods.

### New Features

Two new optional vector operations, `N_VDotProdMultiLocal()` and `N_VDotProdMultiAllReduce()`, have been added to support low-synchronization methods for Anderson acceleration.

The implementation of solve-decoupled implicit MRI-GARK methods has been updated to remove extraneous slow implicit function calls and reduce the memory requirements.

Added a new function `CVodeGetLinSolveStats()` to get the CVODES linear solver statistics as a group.

Added a new function, `CVodeSetMonitorFn()`, that takes a user-function to be called by CVODES after every `nst` successfully completed time-steps. This is intended to provide a way of monitoring the CVODES statistics throughout the simulation.

New orthogonalization methods were added for use within the KINSOL Anderson acceleration routine. See [Anderson Acceleration QR Factorization](#) and `KINSetOrthAA()` for more details.

### Deprecation Notice

The serial, PThreads, PETSc, *hypre*, Parallel, OpenMP\_DEV, and OpenMP vector functions `N_VCloneVectorArray_*` and `N_VDestroyVectorArray_*` have been deprecated. The generic `N_VCloneVectorArray()` and `N_VDestroyVectorArray()` functions should be used instead.

Many constants, types, and functions have been renamed so that they are properly namespaced. The old names have been deprecated and will be removed in SUNDIALS v7.0.0.

The following constants, macros, and typedefs are now deprecated:

Deprecated Name	New Name
<code>realtype</code>	<code>sunrealtype</code>
<code>booleantype</code>	<code>sunbooleantype</code>
<code>RCONST</code>	<code>SUN_RCONST</code>
<code>BIG_REAL</code>	<code>SUN_BIG_REAL</code>
<code>SMALL_REAL</code>	<code>SUN_SMALL_REAL</code>
<code>UNIT_ROUNDOFF</code>	<code>SUN_UNIT_ROUNDOFF</code>
<code>PREC_NONE</code>	<code>SUN_PREC_NONE</code>
<code>PREC_LEFT</code>	<code>SUN_PREC_LEFT</code>
<code>PREC_RIGHT</code>	<code>SUN_PREC_RIGHT</code>
<code>PREC_BOTH</code>	<code>SUN_PREC_BOTH</code>
<code>MODIFIED_GS</code>	<code>SUN_MODIFIED_GS</code>
<code>CLASSICAL_GS</code>	<code>SUN_CLASSICAL_GS</code>
<code>ATimesFn</code>	<code>SUNATimesFn</code>
<code>PSetupFn</code>	<code>SUNPSetupFn</code>
<code>PSolveFn</code>	<code>SUNPSolveFn</code>
<code>DlsMat</code>	<code>SUNDlsMat</code>
<code>DENSE_COL</code>	<code>SUNDLS_DENSE_COL</code>
<code>DENSE_ELEM</code>	<code>SUNDLS_DENSE_ELEM</code>
<code>BAND_COL</code>	<code>SUNDLS_BAND_COL</code>
<code>BAND_COL_ELEM</code>	<code>SUNDLS_BAND_COL_ELEM</code>
<code>BAND_ELEM</code>	<code>SUNDLS_BAND_ELEM</code>
<code>SDIRK_2_1_2</code>	<code>ARKODE_SDIRK_2_1_2</code>
<code>BILLINGTON_3_3_2</code>	<code>ARKODE_BILLINGTON_3_3_2</code>
<code>TRBDF2_3_3_2</code>	<code>ARKODE_TRBDF2_3_3_2</code>
<code>KVAERNO_4_2_3</code>	<code>ARKODE_KVAERNO_4_2_3</code>
<code>ARK324L2SA_DIRK_4_2_3</code>	<code>ARKODE_ARK324L2SA_DIRK_4_2_3</code>
<code>CASH_5_2_4</code>	<code>ARKODE_CASH_5_2_4</code>
<code>CASH_5_3_4</code>	<code>ARKODE_CASH_5_3_4</code>
<code>SDIRK_5_3_4</code>	<code>ARKODE_SDIRK_5_3_4</code>
<code>KVAERNO_5_3_4</code>	<code>ARKODE_KVAERNO_5_3_4</code>
<code>ARK436L2SA_DIRK_6_3_4</code>	<code>ARKODE_ARK436L2SA_DIRK_6_3_4</code>
<code>KVAERNO_7_4_5</code>	<code>ARKODE_KVAERNO_7_4_5</code>
<code>ARK548L2SA_DIRK_8_4_5</code>	<code>ARKODE_ARK548L2SA_DIRK_8_4_5</code>
<code>ARK437L2SA_DIRK_7_3_4</code>	<code>ARKODE_ARK437L2SA_DIRK_7_3_4</code>
<code>ARK548L2SAb_DIRK_8_4_5</code>	<code>ARKODE_ARK548L2SAb_DIRK_8_4_5</code>
<code>MIN_DIRK_NUM</code>	<code>ARKODE_MIN_DIRK_NUM</code>
<code>MAX_DIRK_NUM</code>	<code>ARKODE_MAX_DIRK_NUM</code>
<code>MIS_KW3</code>	<code>ARKODE_MIS_KW3</code>
<code>MRI_GARK_ERK33a</code>	<code>ARKODE_MRI_GARK_ERK33a</code>
<code>MRI_GARK_ERK45a</code>	<code>ARKODE_MRI_GARK_ERK45a</code>
<code>MRI_GARK_IRK21a</code>	<code>ARKODE_MRI_GARK_IRK21a</code>
<code>MRI_GARK_ESDIRK34a</code>	<code>ARKODE_MRI_GARK_ESDIRK34a</code>
<code>MRI_GARK_ESDIRK46a</code>	<code>ARKODE_MRI_GARK_ESDIRK46a</code>

continues on next page

Table 15.1 – continued from previous page

Deprecated Name	New Name
IMEX_MRI_GARK3a	ARKODE_IMEX_MRI_GARK3a
IMEX_MRI_GARK3b	ARKODE_IMEX_MRI_GARK3b
IMEX_MRI_GARK4	ARKODE_IMEX_MRI_GARK4
MIN_MRI_NUM	ARKODE_MIN_MRI_NUM
MAX_MRI_NUM	ARKODE_MAX_MRI_NUM
DEFAULT_MRI_TABLE_3	MRISTEP_DEFAULT_TABLE_3
DEFAULT_EXPL_MRI_TABLE_3	MRISTEP_DEFAULT_EXPL_TABLE_3
DEFAULT_EXPL_MRI_TABLE_4	MRISTEP_DEFAULT_EXPL_TABLE_4
DEFAULT_IMPL_SD_TABLE_2	MRISTEP_DEFAULT_IMPL_SD_TABLE_2
DEFAULT_IMPL_SD_TABLE_3	MRISTEP_DEFAULT_IMPL_SD_TABLE_3
DEFAULT_IMPL_SD_TABLE_4	MRISTEP_DEFAULT_IMPL_SD_TABLE_4
DEFAULT_IMEX_SD_TABLE_3	MRISTEP_DEFAULT_IMEX_SD_TABLE_3
DEFAULT_IMEX_SD_TABLE_4	MRISTEP_DEFAULT_IMEX_SD_TABLE_4
HEUN_EULER_2_1_2	ARKODE_HEUN_EULER_2_1_2
BOGACKI_SHAMPINE_4_2_3	ARKODE_BOGACKI_SHAMPINE_4_2_3
ARK324L2SA_ERK_4_2_3	ARKODE_ARK324L2SA_ERK_4_2_3
ZONNEVELD_5_3_4	ARKODE_ZONNEVELD_5_3_4
ARK436L2SA_ERK_6_3_4	ARKODE_ARK436L2SA_ERK_6_3_4
SAYFY_ABURUB_6_3_4	ARKODE_SAYFY_ABURUB_6_3_4
CASH_KARP_6_4_5	ARKODE_CASH_KARP_6_4_5
FEHLBERG_6_4_5	ARKODE_FEHLBERG_6_4_5
DORMAND_PRINCE_7_4_5	ARKODE_DORMAND_PRINCE_7_4_5
ARK548L2SA_ERK_8_4_5	ARKODE_ARK548L2SA_ERK_8_4_5
VERNER_8_5_6	ARKODE_VERNER_8_5_6
FEHLBERG_13_7_8	ARKODE_FEHLBERG_13_7_8
KNOTH_WOLKE_3_3	ARKODE_KNOTH_WOLKE_3_3
ARK437L2SA_ERK_7_3_4	ARKODE_ARK437L2SA_ERK_7_3_4
ARK548L2SAb_ERK_8_4_5	ARKODE_ARK548L2SAb_ERK_8_4_5
MIN_ERK_NUM	ARKODE_MIN_ERK_NUM
MAX_ERK_NUM	ARKODE_MAX_ERK_NUM
DEFAULT_ERK_2	ARKSTEP_DEFAULT_ERK_2
DEFAULT_ERK_3	ARKSTEP_DEFAULT_ERK_3
DEFAULT_ERK_4	ARKSTEP_DEFAULT_ERK_4
DEFAULT_ERK_5	ARKSTEP_DEFAULT_ERK_5
DEFAULT_ERK_6	ARKSTEP_DEFAULT_ERK_6
DEFAULT_ERK_8	ARKSTEP_DEFAULT_ERK_8
DEFAULT_DIRK_2	ARKSTEP_DEFAULT_DIRK_2
DEFAULT_DIRK_3	ARKSTEP_DEFAULT_DIRK_3
DEFAULT_DIRK_4	ARKSTEP_DEFAULT_DIRK_4
DEFAULT_DIRK_5	ARKSTEP_DEFAULT_DIRK_5
DEFAULT_ARK_ETABLE_3	ARKSTEP_DEFAULT_ARK_ETABLE_3
DEFAULT_ARK_ETABLE_4	ARKSTEP_DEFAULT_ARK_ETABLE_4
DEFAULT_ARK_ETABLE_5	ARKSTEP_DEFAULT_ARK_ETABLE_4
DEFAULT_ARK_ITABLE_3	ARKSTEP_DEFAULT_ARK_ITABLE_3
DEFAULT_ARK_ITABLE_4	ARKSTEP_DEFAULT_ARK_ITABLE_4
DEFAULT_ARK_ITABLE_5	ARKSTEP_DEFAULT_ARK_ITABLE_5
DEFAULT_ERK_2	ERKSTEP_DEFAULT_2
DEFAULT_ERK_3	ERKSTEP_DEFAULT_3
DEFAULT_ERK_4	ERKSTEP_DEFAULT_4
DEFAULT_ERK_5	ERKSTEP_DEFAULT_5

continues on next page

Table 15.1 – continued from previous page

Deprecated Name	New Name
DEFAULT_ERK_6	ERKSTEP_DEFAULT_6
DEFAULT_ERK_8	ERKSTEP_DEFAULT_8

In addition, the following functions are now deprecated (compile-time warnings will be printed if supported by the compiler):

Deprecated Name	New Name
DenseGETRF	SUNDlsMat_DenseGETRF
DenseGETRS	SUNDlsMat_DenseGETRS
denseGETRF	SUNDlsMat_denseGETRF
denseGETRS	SUNDlsMat_denseGETRS
DensePOTRF	SUNDlsMat_DensePOTRF
DensePOTRS	SUNDlsMat_DensePOTRS
densePOTRF	SUNDlsMat_densePOTRF
densePOTRS	SUNDlsMat_densePOTRS
DenseGEQRF	SUNDlsMat_DenseGEQRF
DenseORMQR	SUNDlsMat_DenseORMQR
denseGEQRF	SUNDlsMat_denseGEQRF
denseORMQR	SUNDlsMat_denseORMQR
DenseCopy	SUNDlsMat_DenseCopy
denseCopy	SUNDlsMat_denseCopy
DenseScale	SUNDlsMat_DenseScale
denseScale	SUNDlsMat_denseScale
denseAddIdentity	SUNDlsMat_denseAddIdentity
DenseMatvec	SUNDlsMat_DenseMatvec
denseMatvec	SUNDlsMat_denseMatvec
BandGBTRF	SUNDlsMat_BandGBTRF
bandGBTRF	SUNDlsMat_bandGBTRF
BandGBTRS	SUNDlsMat_BandGBTRS
bandGBTRS	SUNDlsMat_bandGBTRS
BandCopy	SUNDlsMat_BandCopy
bandCopy	SUNDlsMat_bandCopy
BandScale	SUNDlsMat_BandScale
bandScale	SUNDlsMat_bandScale
bandAddIdentity	SUNDlsMat_bandAddIdentity
BandMatvec	SUNDlsMat_BandMatvec
bandMatvec	SUNDlsMat_bandMatvec
ModifiedGS	SUNModifiedGS
ClassicalGS	SUNClassicalGS
QRfact	SUNQRFact
QRsol	SUNQRsol
DlsMat_NewDenseMat	SUNDlsMat_NewDenseMat
DlsMat_NewBandMat	SUNDlsMat_NewBandMat
DestroyMat	SUNDlsMat_DestroyMat
NewIntArray	SUNDlsMat_NewIntArray
NewIndexArray	SUNDlsMat_NewIndexArray
NewRealArray	SUNDlsMat_NewRealArray
DestroyArray	SUNDlsMat_DestroyArray
AddIdentity	SUNDlsMat_AddIdentity

continues on next page

Table 15.2 – continued from previous page

Deprecated Name	New Name
SetToZero	SUNDlsMat_SetToZero
PrintMat	SUNDlsMat_PrintMat
newDenseMat	SUNDlsMat_newDenseMat
newBandMat	SUNDlsMat_newBandMat
destroyMat	SUNDlsMat_destroyMat
newIntArray	SUNDlsMat_newIntArray
newIndexArray	SUNDlsMat_newIndexArray
newRealArray	SUNDlsMat_newRealArray
destroyArray	SUNDlsMat_destroyArray

In addition, the entire `sundials_lapack.h` header file is now deprecated for removal in SUNDIALS v7.0.0. Note, this header file is not needed to use the SUNDIALS LAPACK linear solvers.

Deprecated “bootstrap” and “minimum correction” predictors in `ARKStep` (options 4 and 5 to `ARKStepSetPredictorMethod()`) and the “bootstrap” predictor in `MRISStep` (option 4 to `MRISStepSetPredictorMethod()`). These functions will output a deprecation warning message and will be removed in a future release.

## 15.16 Changes to SUNDIALS in release 5.8.0

### New Features

The *RAJA vector* implementation has been updated to support the SYCL backend in addition to the CUDA and HIP backend. Users can choose the backend when configuring SUNDIALS by using the `SUNDIALS_RAJA_BACKENDS` CMake variable. This vector remains experimental and is subject to change from version to version.

New *SUNMatrix* and *SUNLinearSolver* implementation were added to interface with the Intel oneAPI Math Kernel Library (oneMKL). Both the matrix and the linear solver support general dense linear systems as well as block diagonal linear systems. See §8.9 for more details. This matrix is experimental and is subject to change from version to version.

Added a new *optional* function to the `SUNLinearSolver` API, `SUNLinSolSetZeroGuess()`, to indicate that the next call to `SUNLinSolSolve()` will be made with a zero initial guess. `SUNLinearSolver` implementations that do not use the `SUNLinSolNewEmpty()` constructor will, at a minimum, need set the `setzeroguess` function pointer in the linear solver `ops` structure to `NULL`. The SUNDIALS iterative linear solver implementations have been updated to leverage this new set function to remove one dot product per solve.

The time integrator packages (`ARKODE`, `CVODE(S)`, and `IDA(S)`) all now support a new “matrix-embedded” *SUN-LinearSolver* type. This type supports user-supplied `SUNLinearSolver` implementations that set up and solve the specified linear system at each linear solve call. Any matrix-related data structures are held internally to the linear solver itself, and are not provided by the SUNDIALS package.

Added functions to `ARKODE` and `CVODE(S)` for supplying an alternative right-hand side function and to `IDA(S)` for supplying an alternative residual for use within nonlinear system function evaluations:

- `ARKStepSetNlsRhsFn()`
- `MRISStepSetNlsRhsFn()`
- `CVodeSetNlsRhsFn()`
- `IDASetNlsResFn()`

Support for user-defined inner (fast) integrators has been to the `MRISStep` module. See *MRISStep Custom Inner Steppers* for more information on providing a user-defined integration method.



Added specialized fused HIP kernels to CVODE which may offer better performance on smaller problems when using CVODE with the HIP vector. See the optional input function `CVodeSetUseIntegratorFusedKernels()` for more information. As with other SUNDIALS HIP features, this capability is considered experimental and may change from version to version.

New KINSOL options have been added to apply a constant damping factor in the fixed point and Picard iterations (see `KINSetDamping()`), to delay the start of Anderson acceleration with the fixed point and Picard iterations (see `KINSetDelayAA()`), and to return the newest solution with the fixed point iteration (see `KINSetReturnNewest()`).

The installed `SUNDIALSConfig.cmake` file now supports the `COMPONENTS` option to `find_package`. The exported targets no longer have `IMPORTED_GLOBAL` set.

### Bug Fixes

A bug was fixed in `SUNMatCopyOps()` where the matrix-vector product setup function pointer was not copied.

A bug was fixed in the `SPBCGS` and `SPTFQMR` solvers for the case where a non-zero initial guess and a solution scaling vector are provided. This fix only impacts codes using SPBCGS or SPTFQMR as standalone solvers as all SUNDIALS packages utilize a zero initial guess.

A bug was fixed in the ARKODE stepper modules where the stop time may be passed after resetting the integrator.

A bug was fixed in `IDASSetJacTimesResFn()` in IDAS where the supplied function was used in the dense finite difference Jacobian computation rather than the finite difference Jacobian-vector product approximation.

A bug was fixed in the KINSOL Picard iteration where the value of `KINSetMaxSetupCalls()` would be ignored.

## 15.17 Changes to SUNDIALS in release 5.7.0

A new `N_Vector` implementation based on the SYCL abstraction layer has been added targeting Intel GPUs. At present the only SYCL compiler supported is the DPC++ (Intel oneAPI) compiler. See §6.12 for more details. This vector is considered experimental and is subject to major changes even in minor releases.

A new `SUNMatrix` and `SUNLinearSolver` implementation were added to interface with the MAGMA linear algebra library. Both the matrix and the linear solver support general dense linear systems as well as block diagonal linear systems, and both are targeted at GPUs (AMD or NVIDIA). See §8.8 for more details.

## 15.18 Changes to SUNDIALS in release 5.6.1

Fixed a CMake bug which caused an error if the `CMAKE_CXX_STANDARD` and `SUNDIALS_RAJA_BACKENDS` options were not provided.

Fixed some compiler warnings when using the IBM XL compilers.

## 15.19 Changes to SUNDIALS in release 5.6.0

A new `N_Vector` implementation based on the AMD ROCm HIP platform has been added. This vector can target NVIDIA or AMD GPUs. See §6.11 for more details. This vector is considered experimental and is subject to change from version to version.

The `RAJA vector` implementation has been updated to support the HIP backend in addition to the CUDA backend. Users can choose the backend when configuring SUNDIALS by using the `SUNDIALS_RAJA_BACKENDS` CMake variable. This vector remains experimental and is subject to change from version to version.

A new optional operation, `N_VGetDeviceArrayPointer()`, was added to the `N_Vector` API. This operation is useful for vectors that utilize dual memory spaces, e.g. the native SUNDIALS CUDA `N_Vector`.

The SUNDIALS matrix and linear solver interfaces to the *cuSparse matrix* and *cuSolver batched QR solver* no longer require using the CUDA `N_Vector`. Instead, they require that the vector utilized provides the `N_VGetDeviceArrayPointer()` operation, and that the pointer returned by `N_VGetDeviceArrayPointer()` is a valid CUDA device pointer.

## 15.20 Changes to SUNDIALS in release 5.5.0

Refactored the SUNDIALS build system. CMake 3.12.0 or newer is now required. Users will likely see deprecation warnings, but otherwise the changes should be fully backwards compatible for almost all users. SUNDIALS now exports CMake targets and installs a `SUNDIALSConfig.cmake` file.

Added support for SuperLU DIST 6.3.0 or newer.

## 15.21 Changes to SUNDIALS in release 5.4.0

### Major Features

A new class, *SUNMemoryHelper*, was added to support **GPU users** who have complex memory management needs such as using memory pools. This is paired with new constructors for the CUDA and RAJA vectors that accept a *SUNMemoryHelper* object. Refer to §4.8, §10, §6.10 and §6.13 for more information.

Added full support for time-dependent mass matrices in ARKStep, and expanded existing non-identity mass matrix infrastructure to support use of the fixed point nonlinear solver.

An interface between ARKStep and the XBraid multigrid reduction in time (MGRIT) library [1] has been added to enable parallel-in-time integration. See the *Multigrid Reduction in Time with XBraid* section for more information and the example codes in `examples/arkode/CXX_xbraid`. This interface required the addition of three new `N_Vector` operations to exchange vector data between computational nodes, see `N_VBufSize()`, `N_VBufPack()`, and `N_VBufUnpack()`. These `N_Vector` operations are only used within the XBraid interface and need not be implemented for any other context.

### New Features

The *RAJA vector* has been updated to mirror the CUDA vector. Notably, the update adds managed memory support to the RAJA vector. Users of the vector will need to update any calls to the `N_VMake_Raja()` function because that signature was changed. This vector remains experimental and is subject to change from version to version.

The expected behavior of `SUNNonlinSolGetNumIters()` and `SUNNonlinSolGetNumConvFails()` in the *SUNNonlinearSolver* API have been updated to specify that they should return the number of nonlinear solver iterations and convergence failures in the most recent solve respectively rather than the cumulative number of iterations and failures across all solves respectively. The API documentation and SUNDIALS provided *SUNNonlinearSolver* implementations have been updated accordingly. As before, the cumulative number of nonlinear iterations and failures may be retrieved with the following functions:

- `ARKStepGetNumNonlinSolvIters()`
- `ARKStepGetNumNonlinSolvConvFails()`
- `ARKStepGetNonlinSolvStats()`
- `MRISStepGetNumNonlinSolvIters()`
- `MRISStepGetNumNonlinSolvConvFails()`
- `MRISStepGetNonlinSolvStats()`

- `CCodeGetNumNonlinSolvIters()`
- `CCodeGetNumNonlinSolvConvFails()`
- `CCodeGetNonlinSolvStats()`
- `IDAGetNumNonlinSolvIters()`
- `IDAGetNumNonlinSolvConvFails()`
- `IDAGetNonlinSolvStats()`

Added the following the following functions that advanced users might find useful when providing a custom `SUNNonlinSolSysFn()`:

- `ARKStepComputeState()`
- `ARKStepGetNonlinearSystemData()`
- `MRISStepComputeState()`
- `MRISStepGetNonlinearSystemData()`
- `CCodeComputeState()`
- `CCodeGetNonlinearSystemData()`
- `IDAGetNonlinearSystemData()`

Added new functions to CVODE(S), ARKODE, and IDA(S) to specify the factor for converting between integrator tolerances (WRMS norm) and linear solver tolerances (L2 norm) i.e.,  $\text{tol\_L2} = \text{nrmfac} * \text{tol\_WRMS}$ :

- `ARKStepSetLSNormFactor()`
- `ARKStepSetMassLSNormFactor()`
- `MRISStepSetLSNormFactor()`
- `CCodeSetLSNormFactor()`
- `IDASetLSNormFactor()`

Added new reset functions `ARKStepReset()`, `ERKStepReset()`, and `MRISStepReset()` to reset the stepper time and state vector to user-provided values for continuing the integration from that point while retaining the integration history. These function complement the reinitialization functions `ARKStepReInit()`, `ERKStepReInit()`, and `MRISStepReInit()` which reinitialize the stepper so that the problem integration should resume as if started from scratch.

Updated the MRISStep time-stepping module in ARKODE to support higher-order MRI-GARK methods [61], including methods that involve solve-decoupled, diagonally-implicit treatment of the slow time scale.

The function `CCodeSetLSSetupFrequency()` has been added to CVODE(S) to set the frequency of calls to the linear solver setup function.

The Trilinos Tpetra `N_Vector` interface has been updated to work with Trilinos 12.18+. This update changes the local ordinal type to always be an `int`.

Added support for CUDA 11.

### Bug Fixes

A minor inconsistency in CVODE(S) and a bug ARKODE when checking the Jacobian evaluation frequency has been fixed. As a result codes using using a non-default Jacobian update frequency through a call to `CCodeSetMaxStepsBetweenJac` or `ARKStepSetMaxStepsBetweenJac` will need to increase the provided value by 1 to achieve the same behavior as before.

In IDAS and CVODES, the functions for forward integration with checkpointing (`IDASolveF()`, `CCodeF()`) are now subject to a restriction on the number of time steps allowed to reach the output time. This is the same restriction

applied to `IDASolve()` and `CVode()`. The default maximum number of steps is 500, but this may be changed using the `CVodeSetMaxNumSteps()` and `IDASetMaxNumSteps()` function. This change fixes a bug that could cause an infinite loop in `IDASolveF()` and `CVodeF()`. **This change may cause a runtime error in existing user code.**

Fixed bug in using ERK method integration with static mass matrices.

### Deprecation Notice

For greater clarity the following functions have been deprecated:

- `CVodeSetMaxStepsBetweenJac`
- `ARKStepSetMaxStepsBetweenJac`
- `ARKStepSetMaxStepsBetweenLSet`

The following functions should be used instead:

- `CVodeSetJacEvalFrequency()`
- `ARKStepSetJacEvalFrequency()`
- `ARKStepSetLSetupFrequency()`

## 15.22 Changes to SUNDIALS in release 5.3.0

### Major Feature

Added support to CVODE for integrating IVPs with constraints using BDF methods and projecting the solution onto the constraint manifold with a user defined projection function. This implementation is accompanied by additions to user documentation and CVODE examples. See `CVodeSetProjFn()` for more information.

### New Features

Added the ability to control the CUDA kernel launch parameters for the CUDA vector and sparse matrix implementations. These implementations remain experimental and are subject to change from version to version. In addition, the CUDA vector kernels were rewritten to be more flexible. Most users should see equivalent performance or some improvement, but a select few may observe minor performance degradation with the default settings. Users are encouraged to contact the SUNDIALS team about any performance changes that they notice.

Added new capabilities for monitoring the solve phase in the Newton and fixed-point `SUNNonlinearSolver`, and the SUNDIALS iterative linear solvers. SUNDIALS must be built with the CMake option `SUNDIALS_BUILD_WITH_MONITORING` to use these capabilities.

Added specialized fused CUDA kernels to CVODE which may offer better performance on smaller problems when using CVODE with the CUDA vector. See the optional input function `CVodeSetUseIntegratorFusedKernels()` for more information. As with other SUNDIALS CUDA features, this feature is experimental and may change from version to version.

Added a new function, `CVodeSetMonitorFn()`, that takes a user-function to be called by CVODE after every `nst` successfully completed time-steps. This is intended to provide a way of monitoring the CVODE statistics throughout the simulation.

Added a new function `CVodeGetLinSolveStats()` to get the CVODE linear solver statistics as a group.

Added the following optional functions to provide an alternative ODE right-hand side function (ARKODE and CVODE(S)), DAE residual function (IDA(S)), or nonlinear system function (KINSOL) for use when computing Jacobian-vector products with the internal difference quotient approximation:

- `ARKStepSetJacTimesRhsFn()`
- `CVodeSetJacTimesRhsFn()`

- `CVodeSetJacTimesRhsFnB()`
- `IDASetJacTimesResFn()`
- `IDASetJacTimesResFnB()`
- `KINSetJacTimesVecSysFn()`

### Bug Fixes

Fixed a bug in the iterative linear solvers where an error is not returned if the `Atimes` function is NULL or, if preconditioning is enabled, the `PSolve` function is NULL.

Fixed a bug in ARKODE where the prototypes for `ERKStepSetMinReduction()` and `ARKStepSetMinReduction()` were not included in `arkode_erkstep.h` and `arkode_arkstep.h` respectively.

Fixed a bug in ARKODE where inequality constraint checking would need to be disabled and then re-enabled to update the inequality constraint values after resizing a problem. Resizing a problem will now disable constraints and a call to `ARKStepSetConstraints()` or `ERKStepSetConstraints()` is required to re-enable constraint checking for the new problem size.

## 15.23 Changes to SUNDIALS in release 5.2.0

### New Features

The following functions were added to each of the time integration packages to enable or disable the scaling applied to linear system solutions with matrix-based linear solvers to account for lagged matrix information:

- `ARKStepSetLinearSolutionScaling()`
- `CVodeSetLinearSolutionScaling()`
- `CVodeSetLinearSolutionScalingB()`
- `IDASetLinearSolutionScaling()`
- `IDASetLinearSolutionScalingB()`

When using a matrix-based linear solver with ARKODE, IDA(S), or BDF methods in CVODE(S) scaling is enabled by default.

Added a new *SUNMatrix* implementation that interfaces to the sparse matrix implementation from the NVIDIA cuSPARSE library, see §7.7 for more details. In addition, the CUDA Sparse linear solver has been updated to use the new matrix, as such, users of this matrix will need to update their code. This implementations are still considered to be experimental, thus they are subject to breaking changes even in minor releases.

Added a new “stiff” interpolation module to ARKODE, based on Lagrange polynomial interpolation, that is accessible to each of the ARKStep, ERKStep and MRISStep time-stepping modules. This module is designed to provide increased interpolation accuracy when integrating stiff problems, as opposed to the ARKODE-standard Hermite interpolation module that can suffer when the IVP right-hand side has large Lipschitz constant. While the Hermite module remains the default, the new Lagrange module may be enabled using one of the routines `ARKStepSetInterpolantType()`, `ERKStepSetInterpolantType()`, or `MRISStepSetInterpolantType()`. The serial example problem `ark_brusselator.c` has been converted to use this Lagrange interpolation module. Created accompanying routines `ARKStepSetInterpolantDegree()`, `ARKStepSetInterpolantDegree()` and `ARKStepSetInterpolantDegree()` to provide user control over these interpolating polynomials.

Added two new functions, `ARKStepSetMinReduction()` and `ERKStepSetMinReduction()`, to change the minimum allowed step size reduction factor after an error test failure.

### Bug Fixes

Fixed a build system bug related to the Fortran 2003 interfaces when using the IBM XL compiler. When building the Fortran 2003 interfaces with an XL compiler it is recommended to set `CMAKE_Fortran_COMPILER` to `f2003`, `xl_f2003`, or `xl_f2003_r`.

Fixed a bug in how ARKODE interfaces with a user-supplied, iterative, unscaled linear solver. In this case, ARKODE adjusts the linear solver tolerance in an attempt to account for the lack of support for left/right scaling matrices. Previously, ARKODE computed this scaling factor using the error weight vector, `ewt`; this fix changes that to the residual weight vector, `rwt`, that can differ from `ewt` when solving problems with non-identity mass matrix.

Fixed a linkage bug affecting Windows users that stemmed from `dllimport/dllexport` attribute missing on some SUNDIALS API functions.

Fixed a memory leak in CVODES and IDAS from not deallocating the `atolSmin0` and `atolQsmin0` arrays.

Fixed a bug where a non-default value for the maximum allowed growth factor after the first step would be ignored.

### Deprecation Notice

The routines `ARKStepSetDenseOrder()`, `ARKStepSetDenseOrder()` and `ARKStepSetDenseOrder()` have been deprecated and will be removed in a future release. The new functions `ARKStepSetInterpolantDegree()`, `ARKStepSetInterpolantDegree()`, and `ARKStepSetInterpolantDegree()` should be used instead.

## 15.24 Changes to SUNDIALS in release 5.1.0

### New Features

Added support for a user-supplied function to update the prediction for each implicit stage solution in `ARKStep`. If supplied, this routine will be called *after* any existing `ARKStep` predictor algorithm completes, so that the predictor may be modified by the user as desired. The new user-supplied routine has type `ARKStagePredictFn`, and may be set by calling `ARKStepSetStagePredictFn()`.

The `MRISStep` module has been updated to support attaching different user data pointers to the inner and outer integrators. If applicable, user codes will need to add a call to `ARKStepSetUserData()` to attach their user data pointer to the inner integrator memory as `MRISStepSetUserData()` will not set the pointer for both the inner and outer integrators. The `MRISStep` examples have been updated to reflect this change.

Added support for damping when using Anderson acceleration in `KINSOL`. See the [Mathematical Considerations](#) and the description of the `KINSetDampingAA()` function for more details.

Added support for constant damping to the fixed-point `SUNNonlinearSolver` when using Anderson acceleration. See [SUNNonlinSol\\_FixedPoint description](#) and the `SUNNonlinSolSetDamping_FixedPoint()` for more details.

Added two utility functions, `SUNDIALSFileOpen()` and `SUNDIALSFileClose()` for creating/destroying file pointers. These are useful when using the Fortran 2003 interfaces.

Added a new build system option, `CUDA_ARCH`, to specify the CUDA architecture to target.

### Bug Fixes

Fixed a build system bug related to finding LAPACK/BLAS.

Fixed a build system bug related to checking if the KLU library works.

Fixed a build system bug related to finding PETSc when using the CMake variables `PETSC_INCLUDES` and `PETSC_LIBRARIES` instead of `PETSC_DIR`.

Fixed a bug in the Fortran 2003 interfaces to the ARKODE Butcher table routines and structure. This includes changing the `ARKodeButcherTable` type to be a `type(c_ptr)` in Fortran.



## 15.25 Changes to SUNDIALS in release 5.0.0

### Build System

Increased the minimum required CMake version to 3.5 for most SUNDIALS configurations, and 3.10 when CUDA or OpenMP with device offloading are enabled.

The CMake option `BLAS_ENABLE` and the variable `BLAS_LIBRARIES` have been removed to simplify builds as SUNDIALS packages do not use BLAS directly. For third party libraries that require linking to BLAS, the path to the BLAS library should be included in the `_LIBRARIES` variable for the third party library e.g., `SUPERLUDIST_LIBRARIES` when enabling SuperLU\_DIST.

### NVector

Two new functions were added to aid in creating custom `N_Vector` objects. The constructor `N_VNewEmpty()` allocates an “empty” generic `N_Vector` with the object’s content pointer and the function pointers in the operations structure initialized to NULL. When used in the constructor for custom objects this function will ease the introduction of any new optional operations to the `N_Vector` API by ensuring only required operations need to be set. Additionally, the function `N_VCopyOps()` has been added to copy the operation function pointers between vector objects. When used in clone routines for custom vector objects these functions also will ease the introduction of any new optional operations to the `N_Vector` API by ensuring all operations are copied when cloning objects.

Added new `N_Vector` implementations, `ManyVector` and `MPIManyVector`, to support flexible partitioning of solution data among different processing elements (e.g., CPU + GPU) or for multi-physics problems that couple distinct MPI-based simulations together (see the §6.17 and §6.18 for more details). This implementation is accompanied by additions to user documentation and SUNDIALS examples.

Additionally, an `MPIPlusX vector` implementation has been created to support the MPI+X paradigm where X is a type of on-node parallelism (e.g., OpenMP, CUDA, etc.). The implementation is accompanied by additions to user documentation and SUNDIALS examples.

One new required vector operation and ten new optional vector operations have been added to the `N_Vector` API. The new required operation, `N_VGetLength()`, returns the global vector length. The optional operations have been added to support the new `MPIManyVector` implementation. The operation `N_VGetCommunicator()` must be implemented by subvectors that are combined to create an `MPIManyVector`, but is not used outside of this context. The remaining nine operations are optional local reduction operations intended to eliminate unnecessary latency when performing vector reduction operations (norms, etc.) on distributed memory systems. The optional local reduction vector operations are `N_VDotProdLocal`, `N_VMaxNormLocal`, `N_VMinLocal`, `N_VL1NormLocal`, `N_VWSqrSumLocal`, `N_VWSqrSumMaskLocal`, `N_VInvTestLocal`, `N_VConstrMaskLocal`, and `N_VMinQuotientLocal`. If an `N_Vector` implementation defines any of the local operations as NULL, then the `MPIManyVector` will call standard `N_Vector` operations to complete the computation.

The `*_MPCuda` and `*_MPIRaja` functions have been removed from the CUDA and RAJA vector implementations respectively. Accordingly, the `nvector_mpicuda.h`, `nvector_mpiraja.h`, `libsundials_nvecmpicuda.lib`, and `libsundials_nvecmpicudaraja.lib` files have been removed. Users should use the MPI+X vector in conjunction with the CUDA and RAJA vectors to replace the functionality. The necessary changes are minimal and should require few code modifications. See the example programs in `examples/ida/mpicuda` and `examples/ida/mpiraja` for examples of how to use the MPI+X vector with the CUDA and RAJA vectors, respectively.

Made performance improvements to the CUDA vector. Users who utilize a non-default stream should no longer see default stream synchronizations after memory transfers.

Added a new constructor to the CUDA vector that allows a user to provide custom allocate and free functions for the vector data array and internal reduction buffer.

Added three new `N_Vector` utility functions, `N_VGetVecAtIndexVectorArray()`, `N_VSetVecAtIndexVectorArray()`, and `N_VNewVectorArray()`, for working with `N_Vector` arrays when using the Fortran 2003 interfaces.

### SUNMatrix

Two new functions were added to aid in creating custom *SUNMatrix* objects. The constructor *SUNMatNewEmpty()* allocates an “empty” generic *SUNMatrix* with the object’s content pointer and the function pointers in the operations structure initialized to NULL. When used in the constructor for custom objects this function will ease the introduction of any new optional operations to the *SUNMatrix* API by ensuring only required operations need to be set. Additionally, the function *SUNMatCopyOps()* has been added to copy the operation function pointers between matrix objects. When used in clone routines for custom matrix objects these functions also will ease the introduction of any new optional operations to the *SUNMatrix* API by ensuring all operations are copied when cloning objects.

A new operation, *SUNMatMatvecSetup()*, was added to the *SUNMatrix* API to perform any setup necessary for computing a matrix-vector product. This operation is useful for *SUNMatrix* implementations which need to prepare the matrix itself, or communication structures before performing the matrix-vector product. Users who have implemented a custom *SUNMatrix* will need to at least update their code to set the corresponding ops structure member, *matvecsetup*, to NULL.

The generic *SUNMatrix* API now defines error codes to be returned by matrix operations. Operations which return an integer flag indicating success/failure may return different values than previously.

A new *SUNMatrix* (and *SUNLinearSolver*) implementation was added to facilitate the use of the SuperLU\_DIST library with SUNDIALS.

### **SUNLinearSolver**

A new function was added to aid in creating custom *SUNLinearSolver* objects. The constructor *SUNLinSolNewEmpty()* allocates an “empty” generic *SUNLinearSolver* with the object’s content pointer and the function pointers in the operations structure initialized to NULL. When used in the constructor for custom objects this function will ease the introduction of any new optional operations to the *SUNLinearSolver* API by ensuring only required operations need to be set.

The return type of the *SUNLinSolLastFlag* in the *SUNLinearSolver* has changed from `long int` to *sunindextype* to be consistent with the type used to store row indices in dense and banded linear solver modules.

Added a new optional operation to the *SUNLinearSolver* API, *SUNLinSolGetID()*, that returns a *SUNLinearSolver\_ID* for identifying the linear solver module.

The *SUNLinearSolver* API has been updated to make the initialize and setup functions optional.

A new *SUNLinearSolver* (and *SUNMatrix*) implementation was added to facilitate the use of the SuperLU\_DIST library with SUNDIALS.

Added a new *SUNLinearSolver* implementation, *cuSolverSp\_batchQR*, which leverages the NVIDIA cuSOLVER sparse batched QR method for efficiently solving block diagonal linear systems on NVIDIA GPUs.

Added three new accessor functions to the KLU linear solver to provide user access to the underlying KLU solver structures: *SUNLinSol\_KLUGetSymbolic()*, *SUNLinSol\_KLUGetNumeric()*, and *SUNLinSol\_KLUGetCommon()*.

### **SUNNonlinearSolver**

A new function was added to aid in creating custom *SUNNonlinearSolver* objects. The constructor *SUNNonlinSolNewEmpty()* allocates an “empty” generic *SUNNonlinearSolver* with the object’s content pointer and the function pointers in the operations structure initialized to NULL. When used in the constructor for custom objects this function will ease the introduction of any new optional operations to the *SUNNonlinearSolver* API by ensuring only required operations need to be set.

To facilitate the use of user supplied nonlinear solver convergence test functions the *SUNNonlinSolSetConvTestFn()* function in the *SUNNonlinearSolver* API has been updated to take a `void*` data pointer as input. The supplied data pointer will be passed to the nonlinear solver convergence test function on each call.

The inputs values passed to the first two inputs of the *SUNNonlinSolSolve()* function in the *SUNNonlinearSolver* have been changed to be the predicted state and the initial guess for the correction to that state. Additionally, the definitions of *SUNNonlinSolSetupFn()* and *SUNNonlinSolSolveFn()* in the *SUNNonlinearSolver* API have



been updated to remove unused input parameters. For more information on the nonlinear system formulation and the API functions see *Nonlinear Algebraic Solvers*.

Added a new *SUNNonlinearSolver* implementation for interfacing with the *PETSc SNES* nonlinear solver.

### New Features

A new linear solver interface functions, *ARKLsLinSysFn* and *CVLsLinSysFn*, as added as an alternative method for evaluating the linear systems  $M - \gamma J$  or  $I - \gamma J$ .

Added the following functions to get the current state and gamma value to ARKStep, CVODE and CVODES that may be useful to users who choose to provide their own nonlinear solver implementation:

- *ARKStepGetCurrentState()*
- *ARKStepGetCurrentGamma()*
- *CVodeGetCurrentGamma()*
- *CVodeGetCurrentState()*
- *CVodeGetCurrentGamma()*
- *CVodeGetCurrentStateSens()*
- *CVodeGetCurrentSensSolveIndex()*
- *IDAGetCurrentCj()*
- *IDAGetCurrentY()*
- *IDAGetCurrentYp()*
- *IDAComputeY()*
- *IDAComputeYp()*

Removed extraneous calls to *N\_VMin()* for simulations where the scalar valued absolute tolerance, or all entries of the vector-valued absolute tolerance array, are strictly positive. In this scenario ARKODE, CVODE(S), and IDA(S) steppers will remove at least one global reduction per time step.

The ARKODE, CVODE(S), IDA(S), and KINSOL linear solver interfaces have been updated to only zero the Jacobian matrix before calling a user-supplied Jacobian evaluation function when the attached linear solver has type *SUNLIN-EARSOLVER\_DIRECT*.

Added new Fortran 2003 interfaces to all of the SUNDIALS packages (ARKODE, CVODE(S), IDA(S), and KINSOL as well as most of the *N\_Vector*, *SUNMatrix*, *SUNLinearSolver*, and *SUNNonlinearSolver* implementations. See §4.7 section for more details. These new interfaces were generated with SWIG-Fortran and provide a user an idiomatic Fortran 2003 interface to most of the SUNDIALS C API.

The MRISStep module has been updated to support explicit, implicit, or IMEX methods as the fast integrator using the ARKStep module. As a result some function signatures have been changed including *MRISStepCreate()* which now takes an ARKStep memory structure for the fast integration as an input.

The reinitialization functions *ERKStepReInit()*, *ARKStepReInit()*, and *MRISStepReInit()* have been updated to retain the minimum and maximum step size values from before reinitialization rather than resetting them to the default values.

Added two new embedded ARK methods of orders 4 and 5 to ARKODE (from [51]).

Support for optional inequality constraints on individual components of the solution vector has been added the ARKODE ERKStep and ARKStep modules. See the descriptions of *ERKStepSetConstraints()* and *ARKStepSetConstraints()* for more details. Note that enabling constraint handling requires the *N\_Vector* operations *N\_VMinQuotient()*, *N\_VConstrMask()*, and *N\_VCompare()* that were not previously required by ARKODE.

Add two new ‘Set’ functions to MRISStep, `MRISStepSetPreInnerFn()` and `MRISStepSetPostInnerFn()`, for performing communication or memory transfers needed before or after the inner integration.

### Bug Fixes

Fixed a bug in the build system that prevented the PThreads NVECTOR module from being built.

Fixed a memory leak in the PETSc `N_Vector` clone function.

Fixed a memeory leak in the ARKODE, CVODE, and IDA F77 interfaces when not using the default nonlinear solver.

Fixed a bug in the ARKStep time-stepping module in ARKODE that would result in an infinite loop if the nonlinear solver failed to converge more than the maximum allowed times during a single step.

Fixed a bug in ARKODE that would result in a “too much accuracy requested” error when using fixed time step sizes with explicit methods in some cases.

Fixed a bug in ARKStep where the mass matrix linear solver setup function was not called in the Matrix-free case.

Fixed a minor bug in ARKStep where an incorrect flag is reported when an error occurs in the mass matrix setup or Jacobian-vector product setup functions.

Fixed a bug in the CVODE and CVODES constraint handling where the step size could be set below the minimum step size.

Fixed a bug in the CVODE and CVODES nonlinear solver interfaces where the norm of the accumulated correction was not updated when using a non-default convergence test function.

Fixed a bug in the CVODES `cvRescale` function where the loops to compute the array of scalars for the fused vector scale operation stopped one iteration early.

Fixed a bug in CVODES and IDAS where `CVodeF()` and `IDASolveF()` would return the wrong flag under certain circumstances.

Fixed a bug in CVODES and IDAS where `CVodeF()` and `IDASolveF()` would not return a root in NORMAL\_STEP mode if the root occurred after the desired output time.

Fixed a bug in the IDA and IDAS linear solver interfaces where an incorrect Jacobian-vector product increment was used with iterative solvers other than SPGMR and SPFGMR.

Fixed a bug the IDAS `IDAQuadReInitB()` function where an incorrect memory structure was passed to `IDAQuadReInit()`.

Fixed a bug in the KINSOL linear solver interface where the auxiliary scalar `sJpnorm` was not computed when necessary with the Picard iteration and the auxiliary scalar `sFdotJp` was unnecessarily computed in some cases.

## 15.26 Changes to SUNDIALS in release 4.1.0

### Removed Implementation Headers

The implementation header files (`*_impl.h`) are no longer installed. This means users who are directly accessing or manipulating package memory structures will need to update their code to use the package’s public API.

### New Features

An additional `N_Vector` implementation was added for interfacing with the Tpetra vector from Trilinos library to facilitate interoperability between SUNDIALS and Trilinos. This implementation is accompanied by additions to user documentation and SUNDIALS examples.

### Bug Fixes

The `EXAMPLES_ENABLE_RAJA` CMake option has been removed. The option `EXAMPLES_ENABLE_CUDA` enables all examples that use CUDA including the RAJA examples with a CUDA back end (if RAJA is enabled).

Python is no longer required to run `make test` and `make test_install`.

A bug was fixed where a nonlinear solver object could be freed twice in some use cases.

Fixed a bug in `ARKodeButcherTable_Write()` when printing a Butcher table without an embedding.

## 15.27 Changes to SUNDIALS in release 4.0.2

Added information on how to contribute to SUNDIALS and a contributing agreement.

Moved the definitions of backwards compatibility functions for the prior direct linear solver (DLS) and scaled preconditioned iterative linear solvers (SPILS) to a source file. The symbols are now included in the appropriate package library, e.g. `libsundials_cvode.lib`.

## 15.28 Changes to SUNDIALS in release 4.0.1

A bug in ARKODE where single precision builds would fail to compile has been fixed.

## 15.29 Changes to SUNDIALS in release 4.0.0

The direct and iterative linear solver interfaces in all SUNDIALS packages have been merged into a single unified linear solver interface to support any valid [SUNLinearSolver](#). This includes the `DIRECT` and `ITERATIVE` types as well as the new `MATRIX_ITERATIVE` type. Details regarding how SUNDIALS packages utilize linear solvers of each type as well as a discussion regarding the intended use cases for user-supplied linear solver implementations are included in §8. All example programs have been updated to use the new unified linear solver interfaces.

The unified linear solver interface is very similar to the previous DLS (direct linear solver) and SPILS (scaled preconditioned iterative linear solver) interface in each package. To minimize challenges in user migration to the unified linear solver interfaces, the previous DLS and SPILS functions may still be used however, these are now deprecated and will be removed in a future release. Additionally, that Fortran users will need to enlarge their array of optional integer outputs, and update the indices that they query for certain linear solver related statistics.

The names of all SUNDIALS-provided [SUNLinearSolver](#) constructors have been updated to follow the naming convention `SUNLinSol_*` where `*` is the name of the linear solver. The new constructor names are:

- [SUNLinSol\\_Band\(\)](#)
- [SUNLinSol\\_Dense\(\)](#)
- [SUNLinSol\\_KLU\(\)](#)
- [SUNLinSol\\_LapackBand\(\)](#)
- [SUNLinSol\\_LapackDense\(\)](#)
- [SUNLinSol\\_PCG\(\)](#)
- [SUNLinSol\\_SPBCGS\(\)](#)
- [SUNLinSol\\_SPGMR\(\)](#)
- [SUNLinSol\\_SPTFQMR\(\)](#)
- [SUNLinSol\\_SuperLUMT\(\)](#)

Linear solver-specific “set” routine names have been similarly standardized. To minimize challenges in user migration to the new names, the previous function names may still be used however, these are now deprecated and will be removed in a future release. All example programs and the standalone linear solver examples have been updated to use the new naming convention.

The `SUNLinSol_Band()` constructor has been simplified to remove the storage upper bandwidth argument.

SUNDIALS integrators (ARKODE, CVODE(S), and IDA(S)) have been updated to utilize generic nonlinear solvers defined by the `SUNNonlinearSolver` API. This enables the addition of new nonlinear solver options and allows for external or user-supplied nonlinear solvers. The nonlinear solver API and SUNDIALS provided implementations are described in *Nonlinear Algebraic Solvers* and follow the same object oriented design used by the `N_Vector`, `SUNMatrix`, and `SUNLinearSolver` classes. Currently two nonlinear solver implementations are provided, *Newton* and *fixed-point*. These replicate the previous integrator-specific implementations of Newton’s method and a fixed-point iteration (previously referred to as a functional iteration), respectively. Note the new *fixed-point* implementation can optionally utilize Anderson’s method to accelerate convergence. Example programs using each of these nonlinear solvers in a standalone manner have been added and all example programs have been updated accordingly.

The SUNDIALS integrators (ARKODE, CVODE(S), and IDA(S)) all now use the *Newton* `SUNNonlinearSolver` by default. Users that wish to use the *fixed-point* `SUNNonlinearSolver` will need to create the corresponding nonlinear solver object and attach it to the integrator with the appropriate set function:

- `ARKStepSetNonlinearSolver()`
- `CVodeSetNonlinearSolver()`
- `IDASetNonlinearSolver()`

Functions for setting the nonlinear solver options or getting nonlinear solver statistics remain unchanged and internally call generic `SUNNonlinearSolver` functions as needed.

With the introduction of the `SUNNonlinearSolver` class, the input parameter `iter` to `CVodeCreate()` has been removed along with the function `CVodeSetIterType` and the constants `CV_NEWTON` and `CV_FUNCTIONAL`. While SUNDIALS includes a fixed-point nonlinear solver, it is not currently supported in IDA.

Three fused vector operations and seven vector array operations have been added to the `N_Vector` API. These *optional* operations are disabled by default and may be activated by calling vector specific routines after creating a vector (see §6.1 for more details). The new operations are intended to increase data reuse in vector operations, reduce parallel communication on distributed memory systems, and lower the number of kernel launches on systems with accelerators. The fused operations are:

- `N_VLinearCombination()`
- `N_VScaleAddMulti()`
- `N_VDotProdMulti()`

and the vector array operations are:

- `N_VLinearCombinationVectorArray()`
- `N_VScaleVectorArray()`
- `N_VConstVectorArray()`
- `N_VWrmsNormVectorArray()`
- `N_VWrmsNormMaskVectorArray()`
- `N_VScaleAddMultiVectorArray()`
- `N_VLinearCombinationVectorArray()`

If an `N_Vector` implementation defines the implementation any of these operations as `NULL`, then standard vector operations will automatically be called as necessary to complete the computation.

A new *N\_Vector* implementation, *OpenMPDEV*, leveraging OpenMP device offloading has been added.

Multiple updates to the *CUDA* vector were made:

- Changed the *N\_VMake\_Cuda()* function to take a host data pointer and a device data pointer instead of an *N\_VectorContent\_Cuda* object.
- Changed *N\_VGetLength\_Cuda* to return the global vector length instead of the local vector length.
- Added *N\_VGetLocalLength\_Cuda* to return the local vector length.
- Added *N\_VGetMPIComm\_Cuda* to return the MPI communicator used.
- Removed the accessor functions in the *suncudavec* namespace.
- Added the ability to set the *cudaStream\_t* used for execution of the CUDA kernels. See the function *N\_VSetCudaStreams\_Cuda*.
- Added *N\_VNewManaged\_Cuda()*, *N\_VMakeManaged\_Cuda()*, and *N\_VIsManagedMemory\_Cuda()* functions to accommodate using managed memory with the CUDA vector.

Multiple updates to the *RAJA* vector were made:

- Changed *N\_VGetLength\_Raja* to return the global vector length instead of the local vector length.
- Added *N\_VGetLocalLength\_Raja* to return the local vector length.
- Added *N\_VGetMPIComm\_Raja* to return the MPI communicator used.
- Removed the accessor functions in the *sunrajavec* namespace.

Two changes were made in the ARKODE and CVODE(S) initial step size algorithm:

- Fixed an efficiency bug where an extra call to the RHS function was made.
- Changed the behavior of the algorithm if the max-iterations case is hit. Before the algorithm would exit with the step size calculated on the penultimate iteration. Now it will exit with the step size calculated on the final iteration.

Fortran 2003 interfaces to CVODE, the fixed-point and Newton nonlinear solvers, the dense, band, KLU, PCG, SP-BCGS, SPFGMR, SPGMR, and SPTFQMR linear solvers, and the serial, PThreads, and OpenMP vectors have been added.

The ARKODE library has been entirely rewritten to support a modular approach to one-step methods, which should allow rapid research and development of novel integration methods without affecting existing solver functionality. To support this, the existing ARK-based methods have been encapsulated inside the new *ARKStep* time-stepping module. Two new time-stepping modules have been added:

- The *ERKStep* module provides an optimized implementation for explicit Runge–Kutta methods with reduced storage and number of calls to the ODE right-hand side function.
- The *MRISStep* module implements two-rate explicit-explicit multirate infinitesimal step methods utilizing different step sizes for slow and fast processes in an additive splitting.

This restructure has resulted in numerous small changes to the user interface, particularly the suite of “Set” routines for user-provided solver parameters and “Get” routines to access solver statistics, that are now prefixed with the name of time-stepping module (e.g., *ARKStep* or *ERKStep*) instead of ARKODE. Aside from affecting the names of these routines, user-level changes have been kept to a minimum. However, we recommend that users consult both this documentation and the ARKODE example programs for further details on the updated infrastructure.

As part of the ARKODE restructuring an *ARKodeButcherTable* structure has been added for storing Butcher tables. Functions for creating new Butcher tables and checking their analytic order are provided along with other utility routines. For more details see the *Butcher Table Data Structure* section.

ARKODE's dense output infrastructure has been improved to support higher-degree Hermite polynomial interpolants (up to degree 5) over the last successful time step.

## 15.30 Changes to SUNDIALS in release 3.2.1

Fixed a bug in the [CUDA](#) vector where the [N\\_VInvTest\(\)](#) operation could write beyond the allocated vector data.

Fixed the library installation path for multiarch systems. This fix changes the default library installation path from `CMAKE_INSTALL_PREFIX/lib` to `CMAKE_INSTALL_PREFIX/CMAKE_INSTALL_LIBDIR`. The default value library directory name is automatically set to `lib`, `lib64`, or `lib/<multiarch-tuple>` depending on the system, but maybe be overridden by setting [CMAKE\\_INSTALL\\_LIBDIR](#).

## 15.31 Changes to SUNDIALS in release 3.2.0

### Library Name Change

Changed the name of the RAJA nvector library from `libsundials_nvecraja.lib` to `libsundials_nveccudaraja.lib` to better reflect that the RAJA vector only support the CUDA backend currently.

### New Features

Added hybrid MPI+CUDA and MPI+RAJA vectors to allow use of more than one MPI rank when using a GPU system. The vectors assume one GPU device per MPI rank.

Support for optional inequality constraints on individual components of the solution vector has been added to CVODE and CVODES. For more details see the [Mathematical Considerations](#) and [Optional input functions](#) sections. Use of [CNodeSetConstraints\(\)](#) requires the [N\\_Vector](#) operations [N\\_VMinQuotient\(\)](#), [N\\_VConstrMask\(\)](#), and [N\\_VCompare\(\)](#) that were not previously required by CVODE and CVODES.

### CMake Updates

CMake 3.1.3 is now the minimum required CMake version.

Deprecated the behavior of the [SUNDIALS\\_INDEX\\_TYPE](#) CMake option and added the [SUNDIALS\\_INDEX\\_SIZE](#) CMake option to select the [sunindextype](#) integer size.

The native CMake FindMPI module is now used to locate an MPI installation.

If MPI is enabled and MPI compiler wrappers are not set, the build system will check if `CMAKE_<language>_COMPILER` can compile MPI programs before trying to locate and use an MPI installation.

The previous options for setting MPI compiler wrappers and the executable for running MPI programs have been deprecated. The new options that align with those used in native CMake FindMPI module are [MPI\\_C\\_COMPILER](#), [MPI\\_CXX\\_COMPILER](#), [MPI\\_Fortran\\_COMPILER](#), and [MPIEXEC\\_EXECUTABLE](#).

When a Fortran name-mangling scheme is needed (e.g., [ENABLE\\_LAPACK](#) is ON) the build system will infer the scheme from the Fortran compiler. If a Fortran compiler is not available or the inferred or default scheme needs to be overridden, the advanced options [SUNDIALS\\_F77\\_FUNC\\_CASE](#) and [SUNDIALS\\_F77\\_FUNC\\_UNDERSCORES](#) can be used to manually set the name-mangling scheme and bypass trying to infer the scheme.

Parts of the main `CMakeLists.txt` file were moved to new files in the `src` and `example` directories to make the CMake configuration file structure more modular.

### Bug Fixes

Fixed a problem with setting [sunindextype](#) which would occur with some compilers (e.g. `armclang`) that do not define `__STDC_VERSION__`.

Fixed a thread-safety issue in CVODES and IDAS when using adjoint sensitivity analysis.

Fixed a bug in IDAS where the saved residual value used in the nonlinear solve for consistent initial conditions was passed as temporary workspace and could be overwritten.

## 15.32 Changes to SUNDIALS in release 3.1.2

### CMake Updates

Updated the minimum required version of CMake to 2.8.12 and enabled using rpath by default to locate shared libraries on OSX.

### New Features

Added the function `SUNSparseMatrix_Reallocate()` to allow specification of the matrix nonzero storage.

Added named constants for the two reinitialization types for the KLU SUNLinearSolver.

Updated the `SUNMatScaleAdd()` and `SUNMatScaleAddI()` implementations in the sparse SUNMatrix to more optimally handle the case where the target matrix contained sufficient storage for the sum, but had the wrong sparsity pattern. The sum now occurs in-place, by performing the sum backwards in the existing storage. However, it is still more efficient if the user-supplied Jacobian routine allocates storage for the sum  $M + \gamma J$  or  $M + \gamma J$  manually (with zero entries if needed).

The following examples from the usage notes page of the SUNDIALS website, and updated them to work with SUNDIALS 3.x:

- `cvDisc_dns.c` demonstrates using CVODE with discontinuous solutions or RHS.
- `cvRoberts_dns_negsol.c` illustrates the use of the RHS function return value to control unphysical negative concentrations.
- `cvRoberts_FSA_dns_Switch.c` demonstrates switching on/off forward sensitivity computations. This example came from the usage notes page of the SUNDIALS website.

### Bug Fixes

Fixed a Windows specific problem where `sunindextype` was not correctly defined when using 64-bit integers. On Windows `sunindextype` is now defined as the MSVC basic type `__int64`.

Fixed a bug in the full KLU SUNLinearSolver reinitialization approach where the sparse SUNMatrix pointer would go out of scope on some architectures.

The misnamed function `CVSpilsSetJacTimesSetupFnBS` has been deprecated and replaced by `CVSpilsSetJacTimesBS`. The deprecated function `CVSpilsSetJacTimesSetupFnBS` will be removed in the next major release.

Changed LICENSE install path to `instdir/include/sundials`.

## 15.33 Changes to SUNDIALS in release 3.1.1

### Bug Fixes

Fixed a minor bug in the CVODE and CVODES `cvSLdet`, where a return was missing in the error check for three inconsistent roots.

Fixed a potential memory leak in the `SPGMR` and `SPFGMR` linear solvers. If “Initialize” was called multiple times then the solver memory was reallocated (without being freed).

Fixed a minor bug in `ARKReInit`, where a flag was incorrectly set to indicate that the problem had been resized (instead of just re-initialized).

Fixed C++11 compiler errors/warnings about incompatible use of string literals.



Updated the KLU SUNLinearSolver to use a typedef for the precision-specific solve functions to avoid compiler warnings.

Added missing typecasts for some (void\*) pointers to avoid compiler warnings.

Fixed bug in the sparse SUNMatrix where `int` was used instead of `sunindextype` in one location.

Fixed a minor bug in `KINPrintInfo` where a case was missing for `KIN_REPTD_SYSFUNC_ERR` leading to an undefined info message.

Added missing `#include <stdio.h>` in `N_Vector` and `SUNMatrix` header files.

Added missing prototypes for `ARKSpilsGetNumMTSetups` in `ARKODE` and `IDASpilsGetNumJTSetupEvals` in `IDA` and `IDAS`.

Fixed an indexing bug in the CUDA vector implementation of `N_VWrmsNormMask()` and revised the RAJA vector implementation of `N_VWrmsNormMask()` to work with mask arrays using values other than zero or one. Replaced `double` with `realtype` in the RAJA vector test functions.

Fixed compilation issue with GCC 7.3.0 and Fortran programs that do not require a `SUNMatrix` or `SUNLinearSolver` e.g., iterative linear solvers, explicit methods in `ARKODE`, functional iteration in `CVODE`, etc.

## 15.34 Changes to SUNDIALS in release 3.1.0

Added `N_Vector` print functions that write vector data to a specified file (e.g., `N_VPrintFile_Serial()`).

Added `make test` and `make test_install` options to the build system for testing SUNDIALS after building with `make` and installing with `make install` respectively.

## 15.35 Changes to SUNDIALS in release 3.0.0

### Major Feature

Added new linear solver and matrix interfaces for all SUNDIALS packages and updated the existing linear solver and matrix implementations. The goal of the redesign is to provide greater encapsulation and ease interfacing custom linear solvers with linear solver libraries. Specific changes include:

- Added a `SUNMatrix` interface with three provided implementations: dense, banded, and sparse. These replicate previous SUNDIALS direct (DIs) and sparse (SIs) matrix structures.
- Added example problems demonstrating use of the matrices.
- Added a `SUNLinearSolver` interface with eleven provided implementations: dense, banded, LAPACK dense, LAPACK band, KLU, SuperLU\_MT, SPGMR, SPBCGS, SPTFQMR, SPFGMR, PCG. These replicate previous SUNDIALS generic linear solvers.
- Added example problems demonstrating use of the linear solvers.
- Expanded package-provided direct linear solver (DIs) interfaces and scaled, preconditioned, iterative linear solver (Spils) interfaces to utilize `SUNMatrix` and `SUNLinearSolver` objects.
- Removed package-specific, linear solver-specific, solver modules (e.g., CVDENSE, KINBAND, IDAKLU, ARK-SPGMR) since their functionality is entirely replicated by the generic DIs/Spils interfaces and `SUNLinearSolver` / `SUNMatrix` classes. The exception is CVDIAG, a diagonal approximate Jacobian solver available to `CVODE` and `CVODES`.
- Converted all SUNDIALS example problems to utilize new the new matrix and linear solver objects, along with updated DIs and Spils linear solver interfaces.



- Added Spils interface routines to ARKODE, CVODE, CVODES, IDA and IDAS to allow specification of a user-provided JTSetup routine. This change supports users who wish to set up data structures for the user-provided Jacobian-times-vector (JTimes) routine, and where the cost of one JTSetup setup per Newton iteration can be amortized between multiple JTimes calls.

Corresponding updates were made to all the example programs.

### New Features

[CUDA](#) and [RAJA N\\_Vector](#) implementations to support GPU systems. These vectors are supplied to provide very basic support for running on GPU architectures. Users are advised that these vectors both move all data to the GPU device upon construction, and speedup will only be realized if the user also conducts the right-hand-side function evaluation on the device. In addition, these vectors assume the problem fits on one GPU. For further information about RAJA, users are referred to the [RAJA web site](#).

Added the type [sunindextype](#) to support using 32-bit or 64-bit integer types for indexing arrays within all SUNDIALS structures. [sunindextype](#) is defined to `int32_t` or `int64_t` when portable types are supported, otherwise it is defined as `int` or `long int`. The Fortran interfaces continue to use `long int` for indices, except for the sparse matrix interface that now uses [sunindextype](#). Interfaces to PETSc, hypre, SuperLU\_MT, and KLU have been updated with 32-bit or 64-bit capabilities depending how the user configures SUNDIALS.

To avoid potential namespace conflicts, the macros defining `boolean_t` values `TRUE` and `FALSE` have been changed to [SUNTRUE](#) and [SUNFALSE](#) respectively.

Temporary vectors were removed from preconditioner setup and solve routines for all packages. It is assumed that all necessary data for user-provided preconditioner operations will be allocated and stored in user-provided data structures.

The file `include/sundials_fconfig.h` was added. This file contains SUNDIALS type information for use in Fortran programs.

Added support for many xSDK-compliant build system keys. For more information on xSDK compliance the [xSDK policies](#). The xSDK is a movement in scientific software to provide a foundation for the rapid and efficient production of high-quality, sustainable extreme-scale scientific applications. For more information visit the [xSDK web site](#).

Added functions [SUNDIALSGetVersion\(\)](#) and [SUNDIALSGetVersionNumber\(\)](#) to get SUNDIALS release version information at runtime.

Added comments to `arkode_butcher.c` regarding which methods should have coefficients accurate enough for use in quad precision.

### Build System

Renamed CMake options to enable/disable examples for greater clarity and added option to enable/disable Fortran 77 examples:

- Changed `EXAMPLES_ENABLE` to [EXAMPLES\\_ENABLE\\_C](#)
- Changed `CXX_ENABLE` to [EXAMPLES\\_ENABLE\\_CXX](#)
- Changed `F90_ENABLE` to `EXAMPLES_ENABLE_F90`
- Added `EXAMPLES_ENABLE_F77` option

Added separate `BLAS_ENABLE` and `BLAS_LIBRARIES` CMake variables.

Fixed minor CMake bugs and included additional error checking during CMake configuration.

### Bug Fixes

#### ARKODE

Fixed `RCONST` usage in `arkode_butcher.c`.

Fixed bug in `arkInitialSetup` to ensure the mass matrix vector product is set up before the “msetup” routine is called.

Fixed ARKODE `printf`-related compiler warnings when building SUNDIALS with extended precision.

*CVODE and CVODES*

`CVodeFree()` now calls `lfree` unconditionally (if non-NULL).

*IDA and IDAS*

Added missing prototype for `IDASetMaxBacksIC()` in `ida.h` and `idas.h`.

*KINSOL*

Corrected KINSOL Fortran name translation for `FKIN_SPFGMR`.

Renamed `KINLocalFn` and `KINCommFn` to `KINBBDDLocalFn` and `KINBBDDCommFn` respectively in the BBD preconditioner module for consistency with other SUNDIALS solvers.

## 15.36 Changes to SUNDIALS in release 2.7.0

### New Features and Enhancements

Two additional *N\_Vector* implementations were added – one for *hypre parallel vectors* and one for *PETSc vectors*. These additions are accompanied by additions to various interface functions and to user documentation.

Added a new *N\_Vector* function, `N_VGetVectorID()`, that returns an identifier for the vector.

The sparse matrix structure was enhanced to support both CSR and CSC matrix storage formats.

Various additions were made to the KLU and SuperLU\_MT sparse linear solver interfaces, including support for the CSR matrix format when using KLU.

In all packages, the linear solver and preconditioner `free` routines were updated to return an integer.

In all packages, example codes were updated to use `N_VGetArrayPointer_*` rather than the `NV_DATA` macro when using the native vectors shipped with SUNDIALS.

Additional example programs were added throughout including new examples utilizing the OpenMP vector.

*ARKODE*

The ARKODE implicit predictor algorithms were updated: methods 2 and 3 were improved slightly, a new predictor approach was added, and the default choice was modified.

The handling of integer codes for specifying built-in ARKODE Butcher tables was enhanced. While a global numbering system is still used, methods now have `#defined` names to simplify the user interface and to streamline incorporation of new Butcher tables into ARKODE.

The maximum number of Butcher table stages was increased from 8 to 15 to accommodate very high order methods, and an 8th-order adaptive ERK method was added.

Support was added for the explicit and implicit methods in an additive Runge–Kutta method with different stage times to support new SSP-ARK methods.

The FARKODE interface was extended to include a routine to set scalar/array-valued residual tolerances, to support Fortran applications with non-identity mass-matrices.

*IDA and IDAS*

The optional input function `IDASetMaxBacksIC()` was added to set the maximum number of linesearch backtracks in the initial condition calculation.

### Bug Fixes

Various minor fixes to installation-related files.

Fixed some examples with respect to the change to use new macro/function names e.g., SUNRexp, etc.

In all packages, a memory leak was fixed in the banded preconditioner and banded-block-diagonal preconditioner interfaces.

Corrected name `N_VCloneEmptyVectorArray` to `N_VCloneVectorArrayEmpty` in all documentation files.

Various corrections were made to the interfaces to the sparse solvers KLU and SuperLU\_MT.

For each linear solver, the various solver performance counters are now initialized to 0 in both the solver specification function and in the solver `init` function. This ensures that these solver counters are initialized upon linear solver instantiation as well as at the beginning of the problem solution.

#### *ARKODE*

The missing `ARKSpilsGetNumMtimesEvals` function was added – this had been included in the previous documentation but had not been implemented.

The choice of the method vs embedding the Billington and TRBDF2 explicit Runge–Kutta methods were swapped, since in those the lower-order coefficients result in an A-stable method, while the higher-order coefficients do not. This change results in significantly improved robustness when using those methods.

A bug was fixed for the situation where a user supplies a vector of absolute tolerances, and also uses the vector `Resize` functionality.

A bug was fixed wherein a user-supplied Butcher table without an embedding is supplied, and the user is running with either fixed time steps (or they do adaptivity manually); previously this had resulted in an error since the embedding order was below 1.

#### *CVODE*

Corrections were made to three Fortran interface functions.

In FCVODE, fixed argument order bugs in the FCVKLU and FCVSUPERLUMT linear solver interfaces.

Added missing Fortran interface routines for supplying a sparse Jacobian routine with sparse direct solvers.

#### *CVODES*

A bug was fixed in the interpolation functions used in solving backward problems for adjoint sensitivity analysis.

In the interpolation routines for backward problems, added logic to bypass sensitivity interpolation if input sensitivity argument is `NULL`.

Changed each the return type of `*FreeB` functions to `int` and added `return(0)` to each.

#### *IDA*

Corrections were made to three Fortran interface functions.

Corrected the output from the `idaFoodWeb_bnd.c` example, the wrong component was printed in `PrintOutput`.

#### *IDAS*

In the interpolation routines for backward problems, added logic to bypass sensitivity interpolation if input sensitivity argument is `NULL`.

Changed each the return type of `*FreeB` functions to `int` and added `return(0)` to each.

Corrections were made to three Fortran interface functions.

Added missing Fortran interface routines for supplying a sparse Jacobian routine with sparse direct solvers.

#### *KINSOL*

The Picard iteration return was changed to always return the newest iterate upon success.

A minor bug in the line search was fixed to prevent an infinite loop when the beta condition fails and lambda is below the minimum size.

Corrections were made to three Fortran interface functions.

The functions FKINCREATE and FKININIT were added to split the FKINMALLOC routine into two pieces. FKINMALLOC remains for backward compatibility, but documentation for it has been removed.

Added missing Fortran interface routines for supplying a sparse Jacobian routine with sparse direct solvers.

#### **Matlab Interfaces Removed**

Removed the Matlab interface from distribution as it has not been updated since 2009.

## **15.37 Changes to SUNDIALS in release 2.6.2**

### **New Features and Enhancements**

Various minor fixes to installation-related files

In KINSOL and ARKODE, updated the Anderson acceleration implementation with QR updating.

In CVODES and IDAS, added ReInit and SetOrdering wrappers for backward problems.

In IDAS, fixed for-loop bugs in IDAAckpntAllocVectors that could lead to a memory leak.

### **Bug Fixes**

Updated the BiCGStab linear solver to remove a redundant dot product call.

Fixed potential memory leak in KLU ReInit functions in all solvers.

In ARKODE, fixed a bug in the Cash-Karp Butcher table where the method and embedding coefficient were swapped.

In ARKODE, fixed error in arkDoErrorTest in recovery after failure.

In CVODES, added CVKLUB prototype and corrected CVSuperLUMTB prototype.

In the CVODES and IDAS header files, corrected documentation of backward integration functions, especially the which argument.

In IDAS, added missing backward problem support functions IDALapackDenseB, IDALapackDenseFreeB, IDALapackBandB, and IDALapackBandFreeB.

In IDAS, made SuperLUMT call for backward problem consistent with CVODES.

In CVODE, IDA, and ARKODE, fixed Fortran interfaces to enable calls to GetErrWeights, GetEstLocalErrors, and GetDky within a time step.

## **15.38 Changes to SUNDIALS in release 2.6.1**

Fixed loop limit bug in SlsAddMat function.

In all six solver interfaces to KLU and SuperLUMT, added #include lines, and removed redundant KLU structure allocations.

Minor bug fixes in ARKODE.

## 15.39 Changes to SUNDIALS in release 2.6.0

### Autotools Build Option Removed

With this version of SUNDIALS, support and documentation of the Autotools mode of installation is being dropped, in favor of the CMake mode, which is considered more widely portable.

### New Package: ARKODE

Addition of ARKODE package of explicit, implicit, and additive Runge-Kutta methods for ODEs. This package API is close to CVODE so switching between the two should be straightforward. Thanks go to Daniel Reynolds for the addition of this package.

### New Features and Enhancements

Added *OpenMP* and *Pthreads N\_Vector* implementations for thread-parallel computing environments.

Two major additions were made to the linear system solvers available in all packages. First, in the serial case, an interface to the sparse direct solver KLU was added. Second, an interface to SuperLU\_MT, the multi-threaded version of SuperLU, was added as a thread-parallel sparse direct solver option, to be used with the serial version of the N\_Vector module. As part of these additions, a sparse matrix (CSC format) structure was added to CVODE.

### KINSOL

Two major additions were made to the globalization strategy options (KINSOL argument *strategy*). One is fixed-point iteration, and the other is Picard iteration. Both can be accelerated by use of the Anderson acceleration method. See the relevant paragraphs in Chapter [Mathematical Considerations](#).

An interface to the Flexible GMRES iterative linear solver was added.

### Bug Fixes

In order to avoid possible name conflicts, the mathematical macro and function names MIN, MAX, SQR, RAbs, RSqrt, RExp, RPowerI, and RPowerR were changed to SUNMIN, SUNMAX, SUNSQR, SUNRAbs, SUNRSqrt, SUNRexp, SRpowerI, and SUNRpowerR, respectively. These names occur in both the solver and example programs.

In the LAPACK banded linear solver interfaces, the line `smu = MIN(N-1, mu+ml)` was changed to `smu = mu + ml` to correct an illegal input error for DGBTRF and DGBTRS.

In all Fortran examples, integer declarations were revised so that those which must match a C type `long int` are declared `INTEGER*8`, and a comment was added about the type match. All other integer declarations are just `INTEGER`. Corresponding minor corrections were made to the user guide.

### CVODE and CVODES

In `cvRootFind`, a minor bug was corrected, where the input array was ignored, and a line was added to break out of root-search loop if the initial interval size is below the tolerance `ttol`.

Two minor bugs were fixed regarding the testing of input on the first call to `CVode` – one involving `tstop` and one involving the initialization of `*tret`.

The example program `cvAdvDiff_diag_p` was added to illustrate the use of in parallel.

In the FCVODE optional input routines `FCVSETIIN` and `FCVSETRIN`, the optional fourth argument `key_length` was removed, with hardcoded key string lengths passed to all tests.

In order to eliminate or minimize the differences between the sources for private functions in CVODE and CVODES, the names of many private functions were changed from `CV*` to `cv*` and a few other names were also changed.

An option was added in the case of Adjoint Sensitivity Analysis with dense or banded Jacobian. With a call to `CVDlsSetDenseJacFnBS` or `CVDlsSetBandJacFnBS`, the user can specify a user-supplied Jacobian function of type `CVDls***JacFnBS`, for the case where the backward problem depends on the forward sensitivities.

In `CVodeQuadSensInit`, the line `cv_mem->cv_fQS_data = ...` was corrected (missing `Q`).

In the CVODES User Guide, a paragraph was added in Section 6.2.1 on `CVodeAdjReInit`, and a paragraph was added in Section 6.2.9 on `CVodeGetAdjY`. In the example `cvsRoberts_ASAi_dns`, the output was revised to include the use of `CVodeGetAdjY`.

For the Adjoint Sensitivity Analysis case in which the backward problem depends on the forward sensitivities, options have been added to allow for user-supplied `pset`, `psolve`, and `jtimes` functions.

In the example `cvsHessian_ASA_FSA`, an error was corrected in the function `fB2`, `y2` in place of `y3` in the third term of `Ith(yBdot,6)`.

#### *IDA and IDAS*

In `IDARootfind`, a minor bug was corrected, where the input array `rootdir` was ignored, and a line was added to break out of root-search loop if the initial interval size is below the tolerance `ttol`.

A minor bug was fixed regarding the testing of the input `tstop` on the first call to `IDASolve()`.

In the FIDA optional input routines `FIDASETIIN`, `FIDASETRIN`, and `FIDASETVIN`, the optional fourth argument `key_length` was removed, with hardcoded key string lengths passed to all `strncmp` tests.

An option was added in the case of Adjoint Sensitivity Analysis with dense or banded Jacobian. With a call to `IDADlsSetDenseJacFnBS` or `IDADlsSetBandJacFnBS`, the user can specify a user-supplied Jacobian function of type `IDADls***JacFnBS`, for the case where the backward problem depends on the forward sensitivities.

#### *KINSOL*

In function `KINStop`, two return values were corrected to make the values of `uu` and `fval` consistent.

A bug involving initialization of `mxnewtstep` was fixed. The error affects the case of repeated user calls to `KINSol` with no intervening call to `KINSetMaxNewtonStep`.

A bug in the increments for difference quotient Jacobian approximations was fixed in function `kinDlsBandDQJac`.

In the FKINSOL module, an incorrect return value `ier` in `FKINfunc` was fixed.

In the FKINSOL optional input routines `FKINSETIIN`, `FKINSETRIN`, and `FKINSETVIN`, the optional fourth argument `key_length` was removed, with hardcoded key string lengths passed to all `strncmp` tests.

## 15.40 Changes to SUNDIALS in release 2.5.0

### Integer Type Change

One significant design change was made with this release, the problem size and its relatives, bandwidth parameters, related internal indices, pivot arrays, and the optional output `lsflag` have all been changed from type `int` to type `long int`, except for the problem size and bandwidths in user calls to routines specifying BLAS/LAPACK routines for the dense/band linear solvers. The function `NewIntArray` is replaced by a pair `NewIntArray / NewLintArray`, for `int` and `long int` arrays, respectively.

### Bug Fixes

In the installation files, we modified the treatment of the macro `SUNDIALS_USE_GENERIC_MATH`, so that the parameter `GENERIC_MATH_LIB` is either defined (with no value) or not defined.

In all packages, after the solver memory is created, it is set to zero before being filled.

In each linear solver interface function, the linear solver memory is freed on an error return, and the function now includes a line setting to NULL the main memory pointer to the linear solver memory.

#### *Rootfinding*

In `CVODE(S)` and `IDA(S)`, in the functions `Rcheck1` and `Rcheck2`, when an exact zero is found, the array `glo` of  $g$  values at the left endpoint is adjusted, instead of shifting the  $t$  location `tlo` slightly.

*CVODE and CVODES*

In CVSetTqBDF, the logic was changed to avoid a divide by zero.

In a minor change to the CVODES user interface, the type of the index `which` was changed from `long int` to `int`.

Errors in the logic for the integration of backward problems in CVODES were identified and fixed.

*IDA and IDAS*

To be consistent with IDAS, IDA uses the function `IDAGetDky` for optional output retrieval.

A memory leak was fixed in two of the `IDASp***Free` functions.

A missing vector pointer setting was added in `IDASensLineSrch`.

In `IDACompleteStep`, conditionals around lines loading a new column of three auxiliary divided difference arrays, for a possible order increase, were fixed.

*KINSOL*

Three major logic bugs were fixed - involving updating the solution vector, updating the linesearch parameter, and a missing error return.

Three minor errors were fixed - involving setting `etachoice` in the Matlab/KINSOL interface, a missing error case in `KINPrintInfo`, and avoiding an exponential overflow in the evaluation of `omega`.

## 15.41 Changes to SUNDIALS in release 2.4.0

Added a CMake-based build option in addition to the one based on autotools.

The user interface has been further refined. Some of the API changes involve:

- (a) a reorganization of all linear solver modules into two families (besides the existing family of scaled preconditioned iterative linear solvers, the direct solvers, including new LAPACK-based ones, were also organized into a *direct* family);
- (b) maintaining a single pointer to user data, optionally specified through a `Set`-type function; and
- (c) a general streamlining of the preconditioner modules distributed with the solvers.

Added interfaces to LAPACK linear solvers for dense and banded matrices to all packages.

An option was added to specify which direction of zero-crossing is to be monitored while performing rootfinding in `CVODE(S)` and `IDA(S)`.

CVODES includes several new features related to sensitivity analysis, among which are:

- (a) support for integration of quadrature equations depending on both the states and forward sensitivity (and thus support for forward sensitivity analysis of quadrature equations),
- (b) support for simultaneous integration of multiple backward problems based on the same underlying ODE (e.g., for use in an *forward-over-adjoint* method for computing second order derivative information),
- (c) support for backward integration of ODEs and quadratures depending on both forward states and sensitivities (e.g., for use in computing second-order derivative information), and
- (d) support for reinitialization of the adjoint module.

Moreover, the prototypes of all functions related to integration of backward problems were modified to support the simultaneous integration of multiple problems.

All backward problems defined by the user are internally managed through a linked list and identified in the user interface through a unique identifier.

## 15.42 Changes to SUNDIALS in release 2.3.0

### New Features and Enhancements

The main changes in this release involve a rearrangement of the entire SUNDIALS source tree. At the user interface level, the main impact is in the mechanism of including SUNDIALS header files which must now include the relative path e.g., `#include <cvode/cvode.h>` as all exported header files are now installed in separate subdirectories of the installation *include* directory.

The functions in the generic dense linear solver (`sundials_dense` and `sundials_smalldense`) were modified to work for rectangular  $m \times n$  matrices ( $m \leq n$ ), while the factorization and solution functions were renamed to `DenseGETRF` / `denGETRF` and `DenseGETRS` / `denGETRS`, respectively. The factorization and solution functions in the generic band linear solver were renamed `BandGBTRF` and `BandGBTRS`, respectively.

In IDA, the user interface to the consistent initial conditions calculations was modified. The `IDACalcIC()` arguments `t0`, `yy0`, and `yp0` were removed and a new function, `IDAGetConsistentIC()` is provided.

### Bug Fixes

In the CVODES adjoint solver module, the following two bugs were fixed:

- In `CVodeF` the solver was sometimes incorrectly taking an additional step before returning control to the user (in `CV_NORMAL` mode) thus leading to a failure in the interpolated output function.
- In `CVodeB`, while searching for the current check point, the solver was sometimes reaching outside the integration interval resulting in a segmentation fault.

In IDA, a bug was fixed in the internal difference-quotient dense and banded Jacobian approximations, related to the estimation of the perturbation (which could have led to a failure of the linear solver when zero components with sufficiently small absolute tolerances were present).

## 15.43 Changes to SUNDIALS in release 2.2.0

### New Header Files Names

To reduce the possibility of conflicts, the names of all header files have been changed by adding unique prefixes (e.g., `cvode_` and `sundials_`). When using the default installation procedure, the header files are exported under various subdirectories of the target *include* directory. For more details see Appendix §11.

### Build System Changes

Updated configure script and Makefiles for Fortran examples to avoid C++ compiler errors (now use `CC` and `MPICC` to link only if necessary).

The shared object files are now linked into each SUNDIALS library rather than into a separate `libsundials_shared` library.

### New Features and Enhancements

Deallocation functions now take the address of the respective memory block pointer as the input argument.

Interfaces to the Scaled Preconditioned Bi-CGstab (SPBCG) and Scaled Preconditioned Transpose-Free Quasi-Minimal Residual (SPTFQMR) linear solver modules have been added to all packages. At the same time, function type names for Scaled Preconditioned Iterative Linear Solvers were added for the user-supplied Jacobian-times-vector and preconditioner setup and solve functions. Additionally, in KINSOL interfaces have been added to the SUNDIALS DENSE, and BAND linear solvers and include support for nonlinear residual monitoring which can be used to control Jacobian updating.

A new interpolation method was added to the CVODES adjoint module. The function `CVadjMalloc` has an additional argument which can be used to select the desired interpolation scheme.



FIDA, a Fortran-C interface module, was added.

The rootfinding feature was added to IDA, whereby the roots of a set of given functions may be computed during the integration of the DAE system.

In IDA a user-callable routine was added to access the estimated local error vector.

In the KINSOL Fortran interface module, FKINSOL, optional inputs are now set using FKINSETIIN (integer inputs), FKINSETRIN (real inputs), and FKINSETVIN (vector inputs). Optional outputs are still obtained from the IOUT and ROUT arrays which are owned by the user and passed as arguments to FKINMALLOC.

## 15.44 Changes to SUNDIALS in release 2.1.1

The function `N_VCloneEmpty` was added to the global vector operations table.

A minor bug was fixed in the interpolation functions of the adjoint CVODES module.

## 15.45 Changes to SUNDIALS in release 2.1.0

The user interface has been further refined. Several functions used for setting optional inputs were combined into a single one.

In CVODE(S) and IDA, an optional user-supplied routine for setting the error weight vector was added.

Additionally, to resolve potential variable scope issues, all SUNDIALS solvers release user data right after its use.

The build systems has been further improved to make it more robust.

## 15.46 Changes to SUNDIALS in release 2.0.2

Fixed autoconf-related bug to allow configuration with the PGI Fortran compiler.

Modified the build system to use customized detection of the Fortran name mangling scheme (autoconf's `AC_F77_WRAPPERS` routine is problematic on some platforms).

A bug was fixed in the `CVode` function that was potentially leading to erroneous behavior of the rootfinding procedure on the integration first step.

A new chapter in the User Guide was added - with constants that appear in the user interface.

## 15.47 Changes to SUNDIALS in release 2.0.1

### Build System

Changed the order of compiler directives in header files to avoid compilation errors when using a C++ compiler.

Changed the method of generating `sundials_config.h` to avoid potential warnings of redefinition of preprocessor symbols.

### New Features

In CVODES the option of activating and deactivating forward sensitivity calculations on successive runs without memory allocation and deallocation.

### Bug Fixes

In CVODES bug fixes related to forward sensitivity computations (possible loss of accuracy on a BDF order increase and incorrect logic in testing user-supplied absolute tolerances) were made.

## 15.48 Changes to SUNDIALS in release 2.0.0

Installation of all of SUNDIALS packages has been completely redesigned and is now based on configure scripts.

The major changes from the previous version involve a redesign of the user interface across the entire SUNDIALS suite. We have eliminated the mechanism of providing optional inputs and extracting optional statistics from the solver through the `iopt` and `ropt` arrays. Instead, packages now provide `Set` functions to change the default values for various quantities controlling the solver and `Get` functions to extract statistics after return from the main solver routine.

Additionally, the interfaces to several user-supplied routines (such as those providing Jacobians and preconditioner information) were simplified by reducing the number of arguments. The same information that was previously accessible through such arguments can now be obtained through `Get`-type functions.

In CVODE and CVODES a rootfinding feature was added, whereby the roots of a set of given functions may be computed during the integration of the ODE system.

Changes to the `NVector`:

- Removed `machEnv`, redefined table of vector operations (now contained in the `N_Vector` structure itself).
- All SUNDIALS functions create new `N_Vector` variables through cloning, using an `N_Vector` passed by the user as a template.
- A particular vector implementation is supposed to provide user-callable constructor and destructor functions.
- Removed the following functions from the structure of vector operations: `N_VNew`, `N_VNew_S`, `N_VFree`, `N_VFree_S`, `N_VMake`, `N_VDispose`, `N_VGetData`, `N_VSetData`, `N_VConstrProdPos`, and `N_VOneMask`.
- Added the following functions to the structure of vector operations: `N_VClone`, `N_VDestroy`, `N_VSpace`, `N_VGetArrayPointer`, `N_VSetArrayPointer`, and `N_VWrmsNormMask`.
- Note that `nvec_ser` and `nvec_par` are now separate modules outside the shared SUNDIALS module.

Changes to the linear solvers:

- In SPGMR, added a dummy `N_Vector` argument to be used as a template for cloning.
- In SPGMR, removed `N` (problem dimension) from the argument list of `SpgmrMalloc`.
- Replaced `iterativ.{c,h}` with `iterative.{c,h}`.
- Modified constant names in `iterative.h` (preconditioner types are prefixed with `PREC_`).
- Changed numerical values for `MODIFIED_GS` (from 0 to 1) and `CLASSICAL_GS` (from 1 to 2).

Changes to `sundialsmath` submodule:

- Replaced the internal routine for estimating unit roundoff with definition of unit roundoff from `float.h`.
- Modified functions to call the appropriate math routines given the precision level specified by the user.

Changes to `sundialstypes` submodule:

- Removed `integertype`.
- Added definitions for `BIG_REAL`, `SMALL_REAL`, and `UNIT_ROUNDOFF` using values from `float.h` based on the precision.
- Changed definition of macro `RCONST` to depend on the precision level specified by the user.

# Bibliography

- [1] Xbraid: parallel multigrid in time. <http://llnl.gov/casc/xbraid>.
- [2] AMD ROCm Documentation. <https://rocm.docs.amd.com/en/latest/index.html>.
- [3] Intel oneAPI Programming Guide. <https://software.intel.com/content/www/us/en/develop/documentation/oneapi-programming-guide/top.html>.
- [4] KLU Sparse Matrix Factorization Library. <http://faculty.cse.tamu.edu/davis/suitesparse.html>.
- [5] NVIDIA CUDA Programming Guide. <https://docs.nvidia.com/cuda/index.html>.
- [6] NVIDIA cuSOLVER Programming Guide. <https://docs.nvidia.com/cuda/cusolver/index.html>.
- [7] NVIDIA cuSPARSE Programming Guide. <https://docs.nvidia.com/cuda/cusparse/index.html>.
- [8] SuperLU\_DIST Parallel Sparse Matrix Factorization Library. [https://portal.nersc.gov/project/sparse/superlu/#superlu\\_dist](https://portal.nersc.gov/project/sparse/superlu/#superlu_dist).
- [9] SuperLU\_MT Threaded Sparse Matrix Factorization Library. [https://portal.nersc.gov/project/sparse/superlu/#superlu\\_mt](https://portal.nersc.gov/project/sparse/superlu/#superlu_mt).
- [10] D. G. Anderson. Iterative procedures for nonlinear integral equations. *J. Assoc. Comput. Machinery*, 12:547–560, 1965. doi:10.1145/321296.321305.
- [11] Hartwig Anzt, Terry Cojean, Goran Flegar, Fritz Göbel, Thomas Grützmacher, Pratik Nayak, Tobias Ribizel, Yuhsiang Mike Tsai, and Enrique S. Quintana-Ortí. Ginkgo: A Modern Linear Operator Algebra Framework for High Performance Computing. *ACM Transactions on Mathematical Software*, 48(1):2:1–2:33, February 2022. URL: <https://doi.org/10.1145/3480935> (visited on 2022-02-17), doi:10.1145/3480935.
- [12] Cody J Balos, David J Gardner, Carol S Woodward, and Daniel R Reynolds. Enabling GPU accelerated computing in the SUNDIALS time integration library. *Parallel Computing*, 108:102836, 2021. doi:10.1016/j.parco.2021.102836.
- [13] David Boehme, Todd Gamblin, David Beckingsale, Peer-Timo Bremer, Alfredo Gimenez, Matthew LeGendre, Olga Pearce, and Martin Schulz. Caliper: performance introspection for hpc software stacks. In *SC'16: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 550–560. IEEE, 2016. doi:10.1109/SC.2016.46.
- [14] K. E. Brenan, S. L. Campbell, and L. R. Petzold. *Numerical Solution of Initial-Value Problems in Differential-Algebraic Equations*. SIAM, Philadelphia, Pa, 1996. doi:10.1137/1.9781611971224.
- [15] P. N. Brown. A local convergence theory for combined inexact-Newton/finite difference projection methods. *SIAM J. Numer. Anal.*, 24(2):407–434, 1987. doi:10.1137/0724031.
- [16] P. N. Brown, G. D. Byrne, and A. C. Hindmarsh. VODE, a Variable-Coefficient ODE Solver. *SIAM J. Sci. Stat. Comput.*, 10:1038–1051, 1989. doi:10.1137/0910062.
- [17] P. N. Brown and A. C. Hindmarsh. Reduced storage matrix methods in stiff ODE systems. *J. Appl. Math. & Comp.*, 31:49–91, 1989. doi:10.1016/0096-3003(89)90110-0.

- [18] P. N. Brown, A. C. Hindmarsh, and L. R. Petzold. Using Krylov Methods in the Solution of Large-Scale Differential-Algebraic Systems. *SIAM J. Sci. Comput.*, 15:1467–1488, 1994. doi:10.1137/0915088.
- [19] P. N. Brown, A. C. Hindmarsh, and L. R. Petzold. Consistent Initial Condition Calculation for Differential-Algebraic Systems. *SIAM J. Sci. Comput.*, 19:1495–1512, 1998. doi:10.1137/S1064827595289996.
- [20] P. N. Brown and Y. Saad. Hybrid Krylov Methods for Nonlinear Systems of Equations. *SIAM J. Sci. Stat. Comput.*, 11:450–481, 1990. doi:10.1137/0911026.
- [21] G. D. Byrne. Pragmatic Experiments with Krylov Methods in the Stiff ODE Setting. In J.R. Cash and I. Gladwell, editors, *Computational Ordinary Differential Equations*, 323–356. Oxford, 1992. Oxford University Press.
- [22] G. D. Byrne and A. C. Hindmarsh. User Documentation for PVODE, An ODE Solver for Parallel Computers. Technical Report UCRL-ID-130884, LLNL, May 1998.
- [23] G. D. Byrne and A. C. Hindmarsh. PVODE, an ODE Solver for Parallel Computers. *Intl. J. High Perf. Comput. Apps.*, 13(4):254–365, 1999. doi:10.1177/109434209901300405.
- [24] Y. Cao, S. Li, L. R. Petzold, and R. Serban. Adjoint Sensitivity Analysis for Differential-Algebraic Equations: The Adjoint DAE System and its Numerical Solution. *SIAM J. Sci. Comput.*, 24(3):1076–1089, 2003. doi:10.1137/S1064827501380630.
- [25] M. Caracotsios and W. E. Stewart. Sensitivity analysis of initial value problems with mixed odes and algebraic equations. *Computers and Chemical Engineering*, 9(4):359–365, 1985. doi:10.1016/0098-1354(85)85014-6.
- [26] S. D. Cohen and A. C. Hindmarsh. CVODE, A Stiff/Nonstiff ODE Solver in C. *Computers in Physics*, 10(2):138–143, 1996. doi:10.1063/1.4822377.
- [27] T. A. Davis and P. N. Ekanathan. Algorithm 907: KLU, a direct sparse solver for circuit simulation problems. *ACM Trans. Math. Softw.*, 37(3):1–17, 2010. doi:10.1145/1824801.1824814.
- [28] R. S. Dembo, S. C. Eisenstat, and T. Steihaug. Inexact Newton Methods. *SIAM J. Numer. Anal.*, 19(2):400–408, 1982. doi:10.1137/0719025.
- [29] J. W. Demmel, J. R. Gilbert, and X. S. Li. An Asynchronous Parallel Supernodal Algorithm for Sparse Gaussian Elimination. *SIAM J. Matrix Analysis and Applications*, 20(4):915–952, 1999. doi:10.1137/S0895479897317685.
- [30] J. E. Dennis and R. B. Schnabel. *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. SIAM, Philadelphia, 1996. doi:10.1137/1.9781611971200.
- [31] M.R. Dorr, J.-L. Fattebert, M.E. Wickett, J.F. Belak, and P.E.A. Turchi. A numerical algorithm for the solution of a phase-field model of polycrystalline materials. *Journal of Computational Physics*, 229(3):626–641, 2010. doi:10.1016/j.jcp.2009.09.041.
- [32] H. Carter Edwards, Christian R. Trott, and Daniel Sunderland. Kokkos: enabling manycore performance portability through polymorphic memory access patterns. *Journal of Parallel and Distributed Computing*, 74(12):3202–3216, 2014. doi:10.1016/j.jpdc.2014.07.003.
- [33] H. Fang and Y. Saad. Two classes of secant methods for nonlinear acceleration. *Numer. Linear Algebra Appl.*, 16(3):197–221, 2009. doi:10.1002/nla.617.
- [34] W. F. Feehery, J. E. Tolsma, and P. I. Barton. Efficient Sensitivity Analysis of Large-Scale Differential-Algebraic Systems. *Applied Numer. Math.*, 25(1):41–54, 1997. doi:10.1016/S0168-9274(97)00050-0.
- [35] R. W. Freund. A Transpose-Free Quasi-Minimal Residual Algorithm for Non-Hermitian Linear Systems. *SIAM J. Sci. Comp.*, 14(2):470–482, 1993. doi:10.1137/0914029.
- [36] F. X. Giraldo, J. F. Kelly, and E. M. Constantinescu. Implicit-explicit formulations of a three-dimensional nonhydrostatic unified model of the atmosphere (numa). *SIAM Journal on Scientific Computing*, 35(5):B1162–B1194, 2013. doi:10.1137/120876034.
- [37] Laura Grigori, James W. Demmel, and Xiaoye S. Li. Parallel symbolic factorization for sparse LU with static pivoting. *SIAM J. Scientific Computing*, 29(3):1289–1314, 2007. doi:10.1137/050638102.

- [38] M. R. Hestenes and E. Stiefel. Methods of Conjugate Gradients for Solving Linear Systems. *J. Research of the National Bureau of Standards*, 49(6):409–436, 1952. doi:10.6028/jres.049.044.
- [39] K. L. Hiebert and L. F. Shampine. Implicitly Defined Output Points for Solutions of ODEs. Technical Report SAND80-0180, Sandia National Laboratories, February 1980.
- [40] A. C. Hindmarsh. The PVODE and IDA Algorithms. Technical Report UCRL-ID-141558, LLNL, Dec 2000.
- [41] A. C. Hindmarsh, P. N. Brown, K. E. Grant, S. L. Lee, R. Serban, D. E. Shumaker, and C. S. Woodward. SUN-DIALS: Suite of nonlinear and differential/algebraic equation solvers. *ACM Trans. Math. Softw.*, pages 363–396, 2005. doi:10.1145/1089014.1089020.
- [42] A. C. Hindmarsh and A. G. Taylor. PVODE and KINSOL: Parallel Software for Differential and Nonlinear Systems. Technical Report UCRL-ID-129739, LLNL, February 1998.
- [43] Alan C. Hindmarsh and Radu Serban. Example Programs for CVODE v7.1.0. Technical Report, LLNL, 2024. UCRL-SM-208110.
- [44] Alan C. Hindmarsh, Radu Serban, Cody J. Balos, David J. Gardner, Daniel R. Reynolds, and Carol S. Woodward. User Documentation for CVODE v7.1.0. Technical Report UCRL-SM-208108, LLNL, 2024.
- [45] Alan C. Hindmarsh, Radu Serban, Cody J. Balos, David J. Gardner, Daniel R. Reynolds, and Carol S. Woodward. User Documentation for KINSOL v7.1.0. Technical Report UCRL-SM-208116, LLNL, 2024.
- [46] Alan C. Hindmarsh, Radu Serban, and Aaron Collier. Example Programs for IDA v7.1.0. Technical Report UCRL-SM-208113, LLNL, 2024.
- [47] Seth R. Johnson, Andrey Prokopenko, and Katherine J. Evans. Automated fortran-c++ bindings for large-scale scientific applications. *Computing in Science & Engineering*, 22(5):84–94, 2020. doi:10.1109/MCSE.2019.2924204.
- [48] C. T. Kelley. *Iterative Methods for Solving Linear and Nonlinear Equations*. SIAM, Philadelphia, 1995. doi:10.1137/1.9781611970944.
- [49] C.A. Kennedy and M.H. Carpenter. Diagonally implicit Runge–Kutta methods for ordinary differential equations. a review. Technical Report TM-2016-219173, NASA, 2016.
- [50] C.A. Kennedy and M.H. Carpenter. Diagonally implicit Runge–Kutta methods for stiff ODEs. *Applied Numerical Mathematics*, 146():221–244, 2019. doi:10.1016/j.apnum.2019.07.008.
- [51] C.A. Kennedy and M.H. Carpenter. Higher-order additive runge-kutta schemes for ordinary differential equations. *Applied Numerical Mathematics*, 136:183–205, 2019. doi:10.1016/j.apnum.2018.10.007.
- [52] S. Li, L. R. Petzold, and W. Zhu. Sensitivity Analysis of Differential-Algebraic Equations: A Comparison of Methods on a Special Problem. *Applied Num. Math.*, 32(2):161–174, 2000. doi:10.1016/S0168-9274(99)00020-3.
- [53] X. S. Li. An overview of SuperLU: algorithms, implementation, and user interface. *ACM Trans. Math. Softw.*, 31(3):302–325, September 2005. doi:10.1145/1089014.1089017.
- [54] X.S. Li, J.W. Demmel, J.R. Gilbert, L. Grigori, M. Shao, and I. Yamazaki. SuperLU Users' Guide. Technical Report LBNL-44289, Lawrence Berkeley National Laboratory, September 1999. <http://crd.lbl.gov/protect/unhbox\voidb@x\penalty\@M\xiaoye/SuperLU/>. Last update: August 2011.
- [55] Xiaoye S. Li and James W. Demmel. SuperLU\_DIST: A scalable distributed-memory sparse direct solver for unsymmetric linear systems. *ACM Trans. Mathematical Software*, 29(2):110–140, June 2003. doi:10.1145/779359.779361.
- [56] P. A. Lott, H. F. Walker, C. S. Woodward, and U. M. Yang. An accelerated Picard method for nonlinear systems related to variably saturated flow. *Adv. Wat. Resour.*, 38:92–101, 2012. doi:10.1016/j.advwatres.2011.12.013.
- [57] T. Maly and L. R. Petzold. Numerical Methods and Software for Sensitivity Analysis of Differential-Algebraic Systems. *Applied Numerical Mathematics*, 20(1-2):57–79, 1996. doi:10.1016/0168-9274(95)00117-4.

- [58] D.B. Ozyurt and P.I. Barton. Cheap second order directional derivatives of stiff ODE embedded functionals. *SIAM J. of Sci. Comp.*, 26(5):1725–1743, 2005. doi:[10.1137/030601582](https://doi.org/10.1137/030601582).
- [59] Y. Saad. A flexible inner-outer preconditioned GMRES algorithm. *SIAM J. Sci. Comput.*, 14(2):461–469, 1993. doi:[10.1137/0914028](https://doi.org/10.1137/0914028).
- [60] Y. Saad and M. H. Schultz. GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems. *SIAM J. Sci. Stat. Comp.*, 7(3):856–869, 1986. doi:[10.1137/0907058](https://doi.org/10.1137/0907058).
- [61] A. Sandu. A class of multirate infinitesimal gark methods. *SIAM Journal of Numerical Analysis*, 57(5):2300–2327, 2019. doi:[10.1137/18M1205492](https://doi.org/10.1137/18M1205492).
- [62] Stanimire Tomov, Jack Dongarra, and Marc Baboulin. Towards dense linear algebra for hybrid GPU accelerated manycore systems. *Parallel Computing*, 36(5-6):232–240, June 2010. doi:[10.1016/j.parco.2009.12.005](https://doi.org/10.1016/j.parco.2009.12.005).
- [63] Christian Trott, Luc Berger-Vergiat, David Poliakoff, Sivasankaran Rajamanickam, Damien Lebrun-Grandie, Jonathan Madsen, Nader Al Awar, Milos Gligoric, Galen Shipman, and Geoff Womeldorff. The kokkos ecosystem: comprehensive performance portability for high performance computing. *Computing in Science Engineering*, 23(5):10–18, 2021. doi:[10.1109/MCSE.2021.3098509](https://doi.org/10.1109/MCSE.2021.3098509).
- [64] Christian R. Trott, Damien Lebrun-Grandié, Daniel Arndt, Jan Ciesko, Vinh Dang, Nathan Ellingwood, Rahulku-mar Gayatri, Evan Harvey, Daisy S. Hollman, Dan Ibanez, Nevin Liber, Jonathan Madsen, Jeff Miles, David Poliakoff, Amy Powell, Sivasankaran Rajamanickam, Mikael Simberg, Dan Sunderland, Bruno Turcksin, and Jeremiah Wilke. Kokkos 3: programming model extensions for the exascale era. *IEEE Transactions on Parallel and Distributed Systems*, 33(4):805–817, 2022. doi:[10.1109/TPDS.2021.3097283](https://doi.org/10.1109/TPDS.2021.3097283).
- [65] H. A. Van Der Vorst. Bi-CGSTAB: A Fast and Smoothly Converging Variant of Bi-CG for the Solution of Non-symmetric Linear Systems. *SIAM J. Sci. Stat. Comp.*, 13(2):631–644, 1992. doi:[10.1137/0913035](https://doi.org/10.1137/0913035).
- [66] H. F. Walker and P. Ni. Anderson Acceleration for Fixed-Point Iterations. *SIAM Jour. Num. Anal.*, 49(4):1715–1735, 2011. doi:[10.1137/10078356X](https://doi.org/10.1137/10078356X).

# Index

## Symbols

`_generic_N_Vector` (*C struct*), 187  
`_generic_N_Vector.content` (*C member*), 187  
`_generic_N_Vector.ops` (*C member*), 187  
`_generic_N_Vector.sunctx` (*C member*), 187  
`_generic_N_Vector_Ops` (*C struct*), 187  
`_generic_N_Vector_Ops.nvabs` (*C member*), 188  
`_generic_N_Vector_Ops.nvaddconst` (*C member*), 188  
`_generic_N_Vector_Ops.nvbufpack` (*C member*), 190  
`_generic_N_Vector_Ops.nvbufsize` (*C member*), 190  
`_generic_N_Vector_Ops.nvbufunpack` (*C member*), 190  
`_generic_N_Vector_Ops.nvclone` (*C member*), 188  
`_generic_N_Vector_Ops.nvcloneempty` (*C member*), 188  
`_generic_N_Vector_Ops.nvcompare` (*C member*), 189  
`_generic_N_Vector_Ops.nvconst` (*C member*), 188  
`_generic_N_Vector_Ops.nvconstrmask` (*C member*), 189  
`_generic_N_Vector_Ops.nvconstrmasklocal` (*C member*), 190  
`_generic_N_Vector_Ops.nvconstvectorarray` (*C member*), 189  
`_generic_N_Vector_Ops.nvdestroy` (*C member*), 188  
`_generic_N_Vector_Ops.nvdiv` (*C member*), 188  
`_generic_N_Vector_Ops.nvdotprod` (*C member*), 189  
`_generic_N_Vector_Ops.nvdotprodllocal` (*C member*), 190  
`_generic_N_Vector_Ops.nvdotprodmulti` (*C member*), 189  
`_generic_N_Vector_-  
Ops.nvdotprodmultiallreduce` (*C member*), 190  
`_generic_N_Vector_Ops.nvdotprodmultilocal` (*C member*), 190  
`_generic_N_Vector_Ops.nvgetarraypointer` (*C member*), 188  
`_generic_N_Vector_Ops.nvgetcommunicator` (*C member*), 188  
`_generic_N_Vector_-  
Ops.nvgetdevicearraypointer` (*C member*), 188  
`_generic_N_Vector_Ops.nvgetlength` (*C member*), 188  
`_generic_N_Vector_Ops.nvgetlocallength` (*C member*), 188  
`_generic_N_Vector_Ops.nvgetvectorid` (*C member*), 188  
`_generic_N_Vector_Ops.nvinv` (*C member*), 188  
`_generic_N_Vector_Ops.nvinvtest` (*C member*), 189  
`_generic_N_Vector_Ops.nvinvtestlocal` (*C member*), 190  
`_generic_N_Vector_Ops.nvlnorm` (*C member*), 189  
`_generic_N_Vector_Ops.nvlnormlocal` (*C member*), 190  
`_generic_N_Vector_Ops.nvlinearcombination` (*C member*), 189  
`_generic_N_Vector_-  
Ops.nvlinearcombinationvectorarray` (*C member*), 190  
`_generic_N_Vector_Ops.nvlinearsum` (*C member*), 188  
`_generic_N_Vector_-  
Ops.nvlinearsumvectorarray` (*C member*), 189  
`_generic_N_Vector_Ops.nvmaxnorm` (*C member*), 189  
`_generic_N_Vector_Ops.nvmaxnormlocal` (*C member*), 190  
`_generic_N_Vector_Ops.nvmin` (*C member*), 189  
`_generic_N_Vector_Ops.nvminlocal` (*C member*), 190  
`_generic_N_Vector_Ops.nvminquotient` (*C member*), 189  
`_generic_N_Vector_Ops.nvminquotientlocal` (*C member*), 190  
`_generic_N_Vector_Ops.nvprint` (*C member*), 190  
`_generic_N_Vector_Ops.nvprintfile` (*C member*), 190  
`_generic_N_Vector_Ops.nvprod` (*C member*), 188  
`_generic_N_Vector_Ops.nvscale` (*C member*), 188  
`_generic_N_Vector_Ops.nvscaleaddmulti` (*C member*), 189



`_generic_N_Vector_-`  
    `Ops.nvscaleaddmultivectorarray` (C member), 190  
`_generic_N_Vector_Ops.nvscalevectorarray` (C member), 189  
`_generic_N_Vector_Ops.nvsetarraypointer` (C member), 188  
`_generic_N_Vector_Ops.nvspace` (C member), 188  
`_generic_N_Vector_Ops.nvwl2norm` (C member), 189  
`_generic_N_Vector_Ops.nvwrmsnorm` (C member), 189  
`_generic_N_Vector_Ops.nvwrmsnormmask` (C member), 189  
`_generic_N_Vector_-`  
    `Ops.nvwrmsnormmaskvectorarray` (C member), 189  
`_generic_N_Vector_Ops.nvwrmsnormvectorarray` (C member), 189  
`_generic_N_Vector_Ops.nvwsqrsumlocal` (C member), 190  
`_generic_N_Vector_Ops.nvwsqrsummasklocal` (C member), 190  
`_generic_SUNLinearSolver` (C struct), 313  
`_generic_SUNLinearSolver.content` (C member), 313  
`_generic_SUNLinearSolver.ops` (C member), 313  
`_generic_SUNLinearSolver.sunctx` (C member), 313  
`_generic_SUNLinearSolver_Ops` (C struct), 313  
`_generic_SUNLinearSolver_Ops.free` (C member), 314  
`_generic_SUNLinearSolver_Ops.getid` (C member), 313  
`_generic_SUNLinearSolver_Ops.gettype` (C member), 313  
`_generic_SUNLinearSolver_Ops.initialize` (C member), 313  
`_generic_SUNLinearSolver_Ops.lastflag` (C member), 313  
`_generic_SUNLinearSolver_Ops.numiters` (C member), 313  
`_generic_SUNLinearSolver_Ops.resid` (C member), 314  
`_generic_SUNLinearSolver_Ops.resnorm` (C member), 313  
`_generic_SUNLinearSolver_Ops.setatimes` (C member), 313  
`_generic_SUNLinearSolver_-`  
    `Ops.setpreconditioner` (C member), 313  
`_generic_SUNLinearSolver_-`  
    `Ops.setscalingvectors` (C member), 313  
`_generic_SUNLinearSolver_Ops.setup` (C member), 313  
`_generic_SUNLinearSolver_Ops.setzeroguess` (C member), 313  
`_generic_SUNLinearSolver_Ops.solve` (C member), 313  
`_generic_SUNLinearSolver_Ops.space` (C member), 314  
`_generic_SUNMatrix` (C struct), 265  
`_generic_SUNMatrix.content` (C member), 265  
`_generic_SUNMatrix.ops` (C member), 265  
`_generic_SUNMatrix.sunctx` (C member), 265  
`_generic_SUNMatrix_Ops` (C struct), 265  
`_generic_SUNMatrix_Ops.clone` (C member), 266  
`_generic_SUNMatrix_Ops.copy` (C member), 266  
`_generic_SUNMatrix_Ops.destroy` (C member), 266  
`_generic_SUNMatrix_Ops.getid` (C member), 266  
`_generic_SUNMatrix_Ops.matvec` (C member), 266  
`_generic_SUNMatrix_Ops.matvecsetup` (C member), 266  
`_generic_SUNMatrix_Ops.scaleadd` (C member), 266  
`_generic_SUNMatrix_Ops.scaleaddi` (C member), 266  
`_generic_SUNMatrix_Ops.space` (C member), 266  
`_generic_SUNMatrix_Ops.zero` (C member), 266  
`_generic_SUNNonlinearSolver` (C struct), 372  
`_generic_SUNNonlinearSolver.content` (C member), 372  
`_generic_SUNNonlinearSolver.ops` (C member), 372  
`_generic_SUNNonlinearSolver.sunctx` (C member), 372  
`_generic_SUNNonlinearSolver_Ops` (C struct), 372  
`_generic_SUNNonlinearSolver_Ops.free` (C member), 372  
`_generic_SUNNonlinearSolver_Ops.getcuriter` (C member), 373  
`_generic_SUNNonlinearSolver_-`  
    `Ops.getnumconvfails` (C member), 373  
`_generic_SUNNonlinearSolver_Ops.getnumiters` (C member), 373  
`_generic_SUNNonlinearSolver_Ops.gettype` (C member), 372  
`_generic_SUNNonlinearSolver_Ops.initialize` (C member), 372  
`_generic_SUNNonlinearSolver_Ops.setctestfn` (C member), 373  
`_generic_SUNNonlinearSolver_Ops.setlsetupfn` (C member), 373  
`_generic_SUNNonlinearSolver_Ops.setlsolvefn` (C member), 373  
`_generic_SUNNonlinearSolver_Ops.setmaxiters` (C member), 373



`_generic_SUNNonlinearSolver_Ops.setsysfn` (*C member*), 373  
`_generic_SUNNonlinearSolver_Ops.setup` (*C member*), 372  
`_generic_SUNNonlinearSolver_Ops.solve` (*C member*), 372

## C

`ccmake`, 401

`cmake`, 405

CMake options

- `adiak_DIR`, 416
- `AMDGPU_TARGETS`, 410
- `BUILD_ARKODE`, 405
- `BUILD_CVODE`, 405
- `BUILD_CVODES`, 405
- `BUILD_FORTRAN_MODULE_INTERFACE`, 409
- `BUILD_IDA`, 405
- `BUILD_IDAS`, 405
- `BUILD_KINSOL`, 405
- `BUILD_SHARED_LIBS`, 405
- `BUILD_STATIC_LIBS`, 406
- `CALIPER_DIR`, 415
- `CMAKE_BUILD_TYPE`, 406
- `CMAKE_C_COMPILER`, 406
- `CMAKE_C_EXTENSIONS`, 406
- `CMAKE_C_FLAGS`, 406
- `CMAKE_C_FLAGS_DEBUG`, 406
- `CMAKE_C_FLAGS_MINSIZEREL`, 406
- `CMAKE_C_FLAGS_RELEASE`, 406
- `CMAKE_C_STANDARD`, 406
- `CMAKE_CUDA_ARCHITECTURES`, 408
- `CMAKE_CXX_COMPILER`, 406
- `CMAKE_CXX_EXTENSIONS`, 407
- `CMAKE_CXX_FLAGS`, 407
- `CMAKE_CXX_FLAGS_DEBUG`, 407
- `CMAKE_CXX_FLAGS_MINSIZEREL`, 407
- `CMAKE_CXX_FLAGS_RELEASE`, 407
- `CMAKE_CXX_STANDARD`, 407
- `CMAKE_Fortran_COMPILER`, 407
- `CMAKE_Fortran_FLAGS`, 407
- `CMAKE_Fortran_FLAGS_DEBUG`, 407
- `CMAKE_Fortran_FLAGS_MINSIZEREL`, 407
- `CMAKE_Fortran_FLAGS_RELEASE`, 407
- `CMAKE_INSTALL_LIBDIR`, 408
- `CMAKE_INSTALL_PREFIX`, 408
- `ENABLE_ADIK`, 416
- `ENABLE_CALIPER`, 415
- `ENABLE_CUDA`, 408
- `ENABLE_GINKGO`, 410
- `ENABLE_HIP`, 410
- `ENABLE_HYPRE`, 411
- `ENABLE_KLU`, 411
- `ENABLE_KOKKOS`, 410

- `ENABLE_KOKKOS_KERNELS`, 410
- `ENABLE_LAPACK`, 411
- `ENABLE_MAGMA`, 411
- `ENABLE_MPI`, 412
- `ENABLE_ONEMKL`, 412
- `ENABLE_OPENMP`, 413
- `ENABLE_PETSC`, 413
- `ENABLE_PTHREAD`, 413
- `ENABLE_RAJA`, 413
- `ENABLE_SUPERLUDIST`, 414
- `ENABLE_SUPERLUMT`, 415
- `ENABLE_SYCL`, 415
- `ENABLE_XBRAID`, 417
- `EXAMPLES_ENABLE_C`, 408
- `EXAMPLES_ENABLE_CUDA`, 408
- `EXAMPLES_ENABLE_CXX`, 408
- `EXAMPLES_ENABLE_F2003`, 408
- `EXAMPLES_INSTALL`, 408
- `EXAMPLES_INSTALL_PATH`, 409
- `Ginkgo_DIR`, 410
- `HYPRE_INCLUDE_DIR`, 411
- `HYPRE_LIBRARY`, 411
- `KLU_INCLUDE_DIR`, 411
- `KLU_LIBRARY_DIR`, 411
- `Kokkos_DIR`, 410
- `KokkosKernels_DIR`, 410
- `LAPACK_LIBRARIES`, 411
- `MAGMA_DIR`, 412
- `MPI_C_COMPILER`, 412
- `MPI_CXX_COMPILER`, 412
- `MPI_Fortran_COMPILER`, 412
- `MPIEXEC_EXECUTABLE`, 412
- `ONEMKL_DIR`, 413
- `PETSC_DIR`, 413
- `PETSC_INCLUDES`, 413
- `PETSC_LIBRARIES`, 413
- `SUNDIALS_BUILD_WITH_MONITORING`, 409
- `SUNDIALS_BUILD_WITH_PROFILING`, 409
- `SUNDIALS_ENABLE_ERROR_CHECKS`, 410
- `SUNDIALS_ENABLE_EXTERNAL_ADDONS`, 410
- `SUNDIALS_GINKGO_BACKENDS`, 410
- `SUNDIALS_INDEX_SIZE`, 416
- `SUNDIALS_INDEX_TYPE`, 416
- `SUNDIALS_INSTALL_CMAKEDIR`, 417
- `SUNDIALS_LAPACK_CASE`, 416
- `SUNDIALS_LAPACK_UNDESCORES`, 416
- `SUNDIALS_LOGGING_LEVEL`, 409
- `SUNDIALS_MAGMA_BACKENDS`, 412
- `SUNDIALS_MATH_LIBRARY`, 417
- `SUNDIALS_ONEMKL_USE_GETRF_LOOP`, 413
- `SUNDIALS_ONEMKL_USE_GETRS_LOOP`, 413
- `SUNDIALS_PRECISION`, 416
- `SUNDIALS_RAJA_BACKENDS`, 414
- `SUNDIALS_SYCL_2020_UNSUPPORTED`, 415

SUPERLUDIST\_DIR, 414  
 SUPERLUDIST\_INCLUDE\_DIR, 414  
 SUPERLUDIST\_INCLUDE\_DIRS, 414  
 SUPERLUDIST\_LIBRARIES, 414  
 SUPERLUDIST\_LIBRARY\_DIR, 414  
 SUPERLUDIST\_OpenMP, 414  
 SUPERLUMT\_INCLUDE\_DIR, 415  
 SUPERLUMT\_LIBRARY\_DIR, 415  
 SUPERLUMT\_THREAD\_TYPE, 415  
 USE\_XSDK\_DEFAULTS, 417  
 XBRAID\_DIR, 417  
 XBRAID\_INCLUDES, 417  
 XBRAID\_LIBRARIES, 417

cmake-gui, 401

CopyFromDevice (C++ function), 247

CopyToDevice (C++ function), 247

## D

DenseLinearSolver (C++ class), 362

DenseLinearSolver::~DenseLinearSolver (C++ function), 362

DenseLinearSolver::Convert (C++ function), 362

DenseLinearSolver::DenseLinearSolver (C++ function), 362

DenseLinearSolver::operator SUNLinearSolver (C++ function), 362

DenseLinearSolver::operator= (C++ function), 362

DenseMatrix (C++ class), 301

DenseMatrix::~DenseMatrix (C++ function), 302

DenseMatrix::BlockCols (C++ function), 302

DenseMatrix::BlockRows (C++ function), 302

DenseMatrix::Blocks (C++ function), 302

DenseMatrix::Cols (C++ function), 302

DenseMatrix::Convert (C++ function), 303

DenseMatrix::DenseMatrix (C++ function), 301, 302

DenseMatrix::exec\_space (C++ type), 301

DenseMatrix::ExecSpace (C++ function), 302

DenseMatrix::member\_type (C++ type), 301

DenseMatrix::memory\_space (C++ type), 301

DenseMatrix::operator SUNMatrix (C++ function), 302, 303

DenseMatrix::operator= (C++ function), 302

DenseMatrix::range\_policy (C++ type), 301

DenseMatrix::Rows (C++ function), 302

DenseMatrix::size\_type (C++ type), 301

DenseMatrix::team\_policy (C++ type), 301

DenseMatrix::View (C++ function), 302

DenseMatrix::view\_type (C++ type), 301

## G

GetDenseMat (C++ function), 303

GetVec (C++ function), 247

IDAAdjFree (C function), 151

IDAAdjInit (C function), 150

IDAAdjReInit (C function), 150

IDAAdjSetNoSensi (C function), 151

IDABBDCommFn (C type), 117

IDABBDCommFnB (C type), 184

IDABBDLocalFn (C type), 116

IDABBDLocalFnB (C type), 184

IDABBDPrecGetNumGfnEvals (C function), 120

IDABBDPrecGetWorkSpace (C function), 119

IDABBDPrecInit (C function), 118

IDABBDPrecInitB (C function), 183

IDABBDPrecReInit (C function), 119

IDABBDPrecReInitB (C function), 183

IDACalcIC (C function), 65

IDACalcICB (C function), 158

IDACalcICBS (C function), 158

IDAClearStopTime (C function), 72

IDAComputeY (C function), 377

IDAComputeYp (C function), 377

IDAComputeYpSens (C function), 378

IDAComputeYSens (C function), 378

IDACreate (C function), 61

IDACreateB (C function), 153

IDAEwtFn (C type), 103

IDAFree (C function), 61

IDAGetActualInitStep (C function), 91

IDAGetAdjCheckPointsInfo (C function), 168

IDAGetAdjCheckPointsInfo.IDAadjCheckPointRec (C struct), 168

IDAGetAdjCheckPointsInfo.IDAadjCheckPointRec.my\_addr (C member), 168

IDAGetAdjCheckPointsInfo.IDAadjCheckPointRec.next\_addr (C member), 168

IDAGetAdjCheckPointsInfo.IDAadjCheckPointRec.nstep (C member), 168

IDAGetAdjCheckPointsInfo.IDAadjCheckPointRec.order (C member), 168

IDAGetAdjCheckPointsInfo.IDAadjCheckPointRec.step (C member), 168

IDAGetAdjCheckPointsInfo.IDAadjCheckPointRec.t0 (C member), 168

IDAGetAdjCheckPointsInfo.IDAadjCheckPointRec.t1 (C member), 168

IDAGetAdjIDABmem (C function), 167

IDAGetAdjY (C function), 167

IDAGetB (C function), 160

IDAGetConsistentIC (C function), 95

IDAGetConsistentICB (C function), 168

IDAGetCurrentCj (C function), 374

IDAGetCurrentOrder (C function), 90

IDAGetCurrentStep (C function), 91

IDAGetCurrentTime (C function), 91

- 
- IDAGetCurrentY (*C function*), 375
  - IDAGetCurrentYp (*C function*), 375
  - IDAGetCurrentYpSens (*C function*), 375
  - IDAGetCurrentYSens (*C function*), 375
  - IDAGetDky (*C function*), 86
  - IDAGetErrWeights (*C function*), 92
  - IDAGetEstLocalErrors (*C function*), 92
  - IDAGetIntegratorStats (*C function*), 92
  - IDAGetJac (*C function*), 96
  - IDAGetJacCj (*C function*), 96
  - IDAGetJacNumSteps (*C function*), 97
  - IDAGetJacTime (*C function*), 97
  - IDAGetLastLinFlag (*C function*), 100
  - IDAGetLastOrder (*C function*), 90
  - IDAGetLastStep (*C function*), 90
  - IDAGetLinReturnFlagName (*C function*), 101
  - IDAGetLinWorkSpace (*C function*), 97
  - IDAGetNonlinearSystemData (*C function*), 375
  - IDAGetNonlinearSystemDataSens (*C function*), 376
  - IDAGetNonlinSolvStats (*C function*), 93
  - IDAGetNumBacktrackOps (*C function*), 95
  - IDAGetNumErrTestFails (*C function*), 89
  - IDAGetNumGEvals (*C function*), 96
  - IDAGetNumJacEvals (*C function*), 98
  - IDAGetNumJtimesEvals (*C function*), 100
  - IDAGetNumJTSetupEvals (*C function*), 99
  - IDAGetNumLinConvFails (*C function*), 99
  - IDAGetNumLinIters (*C function*), 98
  - IDAGetNumLinResEvals (*C function*), 98
  - IDAGetNumLinSolvSetups (*C function*), 89
  - IDAGetNumNonlinSolvConvFails (*C function*), 93
  - IDAGetNumNonlinSolvIters (*C function*), 93
  - IDAGetNumPrecEvals (*C function*), 99
  - IDAGetNumPrecSolves (*C function*), 99
  - IDAGetNumResEvals (*C function*), 89
  - IDAGetNumResEvalsSens (*C function*), 133
  - IDAGetNumSteps (*C function*), 89
  - IDAGetNumStepSensSolveFails (*C function*), 133
  - IDAGetNumStepSolveFails (*C function*), 90
  - IDAGetQuad (*C function*), 111
  - IDAGetQuadB (*C function*), 171
  - IDAGetQuadDky (*C function*), 111
  - IDAGetQuadErrWeights (*C function*), 114
  - IDAGetQuadNumErrTestFails (*C function*), 113
  - IDAGetQuadNumRhsEvals (*C function*), 113
  - IDAGetQuadSens (*C function*), 140
  - IDAGetQuadSens1 (*C function*), 141
  - IDAGetQuadSensDky (*C function*), 140
  - IDAGetQuadSensDky1 (*C function*), 141
  - IDAGetQuadSensErrWeights (*C function*), 144
  - IDAGetQuadSensNumErrTestFails (*C function*), 144
  - IDAGetQuadSensNumRhsEvals (*C function*), 144
  - IDAGetQuadSensStats (*C function*), 145
  - IDAGetQuadStats (*C function*), 114
  - IDAGetReturnFlagName (*C function*), 94
  - IDAGetRootInfo (*C function*), 95
  - IDAGetSens (*C function*), 128
  - IDAGetSens1 (*C function*), 129
  - IDAGetSensConsistentIC (*C function*), 136
  - IDAGetSensDky (*C function*), 128
  - IDAGetSensDky1 (*C function*), 129
  - IDAGetSensErrWeights (*C function*), 134
  - IDAGetSensNonlinSolvStats (*C function*), 135
  - IDAGetSensNumErrTestFails (*C function*), 133
  - IDAGetSensNumLinSolvSetups (*C function*), 134
  - IDAGetSensNumNonlinSolvConvFails (*C function*), 135
  - IDAGetSensNumNonlinSolvIters (*C function*), 135
  - IDAGetSensNumResEvals (*C function*), 132
  - IDAGetSensStats (*C function*), 134
  - IDAGetTolScaleFactor (*C function*), 91
  - IDAGetUserData (*C function*), 94
  - IDAGetWorkSpace (*C function*), 88
  - IDAInit (*C function*), 61
  - IDAInitB (*C function*), 153
  - IDAInitBS (*C function*), 154
  - IDALsJacFn (*C type*), 104
  - IDALsJacFnB (*C type*), 175
  - IDALsJacFnBS (*C type*), 176
  - IDALsJacTimesSetupFn (*C type*), 106
  - IDALsJacTimesSetupFnB (*C type*), 178
  - IDALsJacTimesSetupFnBS (*C type*), 179
  - IDALsJacTimesVecFn (*C type*), 105
  - IDALsJacTimesVecFnB (*C type*), 177
  - IDALsJacTimesVecFnBS (*C type*), 177
  - IDALsPrecSetupFn (*C type*), 108
  - IDALsPrecSetupFnB (*C type*), 181
  - IDALsPrecSetupFnBS (*C type*), 182
  - IDALsPrecSolveFn (*C type*), 107
  - IDALsPrecSolveFnB (*C type*), 180
  - IDALsPrecSolveFnBS (*C type*), 180
  - IDAPrintAllStats (*C function*), 94
  - IDAQuadFree (*C function*), 110
  - IDAQuadInit (*C function*), 110
  - IDAQuadInitB (*C function*), 169
  - IDAQuadInitBS (*C function*), 169
  - IDAQuadReInit (*C function*), 110
  - IDAQuadReInitB (*C function*), 170
  - IDAQuadRhsFn (*C type*), 114
  - IDAQuadRhsFnB (*C type*), 173
  - IDAQuadRhsFnBS (*C type*), 174
  - IDAQuadSenseEtolerances (*C function*), 143
  - IDAQuadSensFree (*C function*), 139
  - IDAQuadSensInit (*C function*), 138
  - IDAQuadSensReInit (*C function*), 139
  - IDAQuadSensRhsFn (*C type*), 145
  - IDAQuadSensSStolerances (*C function*), 143
  - IDAQuadSensSVtolerances (*C function*), 143

IDAQuadSStolerances (*C function*), 112  
IDAQuadSVtolerances (*C function*), 113  
IDAReInit (*C function*), 101  
IDAReInitB (*C function*), 155  
IDAResFn (*C type*), 102  
IDAResFnB (*C type*), 172  
IDAResFnBS (*C type*), 173  
IDARootFn (*C type*), 103  
IDARootInit (*C function*), 67  
IDASenseEtolerances (*C function*), 126  
IDASensFree (*C function*), 124  
IDASensInit (*C function*), 123  
IDASensReInit (*C function*), 124  
IDASensResFn (*C type*), 136  
IDASensSStolerances (*C function*), 125  
IDASensSVtolerances (*C function*), 126  
IDASensToggleOff (*C function*), 125  
IDASetConstraints (*C function*), 73  
IDASetDeltaCjLSetup (*C function*), 75  
IDASetEpsLin (*C function*), 78  
IDASetEpsLinB (*C function*), 166  
IDASetEtaConvFail (*C function*), 85  
IDASetEtaFixedStepBounds (*C function*), 83  
IDASetEtaLow (*C function*), 84  
IDASetEtaMax (*C function*), 83  
IDASetEtaMin (*C function*), 84  
IDASetEtaMinErrFail (*C function*), 84  
IDASetId (*C function*), 73  
IDASetIncrementFactor (*C function*), 76  
IDASetIncrementFactorB (*C function*), 163  
IDASetInitStep (*C function*), 70  
IDASetJacFn (*C function*), 74  
IDASetJacFnB (*C function*), 161  
IDASetJacFnBS (*C function*), 162  
IDASetJacTimes (*C function*), 75  
IDASetJacTimesB (*C function*), 163  
IDASetJacTimesBS (*C function*), 163  
IDASetJacTimesResFn (*C function*), 77  
IDASetJacTimesResFnB (*C function*), 164  
IDASetLinearSolutionScaling (*C function*), 75  
IDASetLinearSolutionScalingB (*C function*), 162  
IDASetLinearSolver (*C function*), 64  
IDASetLinearSolverB (*C function*), 156  
IDASetLineSearchOffIC (*C function*), 82  
IDASetLSNormFactor (*C function*), 78  
IDASetLSNormFactorB (*C function*), 166  
IDASetMaxBacksIC (*C function*), 82  
IDASetMaxConvFails (*C function*), 79  
IDASetMaxErrTestFails (*C function*), 72  
IDASetMaxNonlinIters (*C function*), 79  
IDASetMaxNumItersIC (*C function*), 81  
IDASetMaxNumJacIC (*C function*), 81  
IDASetMaxNumSteps (*C function*), 70  
IDASetMaxNumStepsIC (*C function*), 81

IDASetMaxOrd (*C function*), 70  
IDASetMaxStep (*C function*), 71  
IDASetMinStep (*C function*), 71  
IDASetNlsResFn (*C function*), 80  
IDASetNoInactiveRootWarn (*C function*), 86  
IDASetNonlinConvCoef (*C function*), 79  
IDASetNonlinConvCoefIC (*C function*), 80  
IDASetNonlinearSolver (*C function*), 65  
IDASetNonlinearSolverB (*C function*), 157  
IDASetNonlinearSolverSensSim (*C function*), 127  
IDASetNonlinearSolverSensStg (*C function*), 127  
IDASetPreconditioner (*C function*), 77  
IDASetPreconditionerB (*C function*), 165  
IDASetPreconditionerBS (*C function*), 165  
IDASetQuadErrCon (*C function*), 112  
IDASetQuadSensErrCon (*C function*), 142  
IDASetRootDirection (*C function*), 85  
IDASetSensDQMethod (*C function*), 131  
IDASetSensErrCon (*C function*), 131  
IDASetSensMaxNonlinIters (*C function*), 132  
IDASetSensParams (*C function*), 130  
IDASetStepToleranceIC (*C function*), 82  
IDASetStopTime (*C function*), 71  
IDASetSuppressAlg (*C function*), 72  
IDASetUserData (*C function*), 69  
IDASolve (*C function*), 67  
IDASolveB (*C function*), 159  
IDASolveF (*C function*), 151  
IDASStolerances (*C function*), 62  
IDASStolerancesB (*C function*), 155  
IDASVtolerances (*C function*), 62  
IDASVtolerancesB (*C function*), 156  
IDAWFtolerances (*C function*), 62

## M

Matrix (*C++ class*), 299  
Matrix::~Matrix (*C++ function*), 299  
Matrix::Convert (*C++ function*), 299  
Matrix::GkoExec (*C++ function*), 299  
Matrix::GkoMtx (*C++ function*), 299  
Matrix::GkoSize (*C++ function*), 299  
Matrix::Matrix (*C++ function*), 299  
Matrix::operator SUNMatrix (*C++ function*), 299  
Matrix::operator= (*C++ function*), 299

## N

N\_VAbs (*C function*), 198  
N\_VAddConst (*C function*), 198  
N\_VBufPack (*C function*), 207  
N\_VBufSize (*C function*), 207  
N\_VBufUnpack (*C function*), 207  
N\_VClone (*C function*), 195  
N\_VCloneEmpty (*C function*), 195  
N\_VCloneVectorArray (*C function*), 191



- N\_VCloneVectorArrayEmpty (C function), 192  
 N\_VCompare (C function), 200  
 N\_VConst (C function), 197  
 N\_VConstrMask (C function), 200  
 N\_VConstrMaskLocal (C function), 205  
 N\_VConstVectorArray (C function), 202  
 N\_VCopyFromDevice\_Cuda (C function), 229  
 N\_VCopyFromDevice\_Hip (C function), 234  
 N\_VCopyFromDevice\_OpenMPDEV (C function), 249  
 N\_VCopyFromDevice\_Raja (C function), 244  
 N\_VCopyFromDevice\_Sycl (C++ function), 239  
 N\_VCopyOps (C function), 194  
 N\_VCopyToDevice\_Cuda (C function), 229  
 N\_VCopyToDevice\_Hip (C function), 234  
 N\_VCopyToDevice\_OpenMPDEV (C function), 249  
 N\_VCopyToDevice\_Raja (C function), 243  
 N\_VCopyToDevice\_Sycl (C++ function), 239  
 N\_VDestroy (C function), 196  
 N\_VDestroyVectorArray (C function), 192  
 N\_VDiv (C function), 198  
 N\_VDotProd (C function), 198  
 N\_VDotProdLocal (C function), 204  
 N\_VDotProdMulti (C function), 201  
 N\_VDotProdMultiAllReduce (C function), 206  
 N\_VDotProdMultiLocal (C function), 206  
 N\_Vector (C type), 187  
 N\_Vector\_ID (C enum), 194  
 N\_Vector\_Ops (C type), 187  
 N\_VEnableConstVectorArray\_Cuda (C function), 230  
 N\_VEnableConstVectorArray\_Hip (C function), 234  
 N\_VEnableConstVectorArray\_ManyVector (C function), 254  
 N\_VEnableConstVectorArray\_MPIManyVector (C function), 258  
 N\_VEnableConstVectorArray\_OpenMP (C function), 218  
 N\_VEnableConstVectorArray\_OpenMPDEV (C function), 250  
 N\_VEnableConstVectorArray\_Parallel (C function), 215  
 N\_VEnableConstVectorArray\_ParHyp (C function), 224  
 N\_VEnableConstVectorArray\_Petsc (C function), 226  
 N\_VEnableConstVectorArray\_Pthreads (C function), 222  
 N\_VEnableConstVectorArray\_Raja (C function), 244  
 N\_VEnableConstVectorArray\_Serial (C function), 211  
 N\_VEnableConstVectorArray\_Sycl (C++ function), 240  
 N\_VEnableDotProdMulti\_Cuda (C function), 229  
 N\_VEnableDotProdMulti\_Hip (C function), 234  
 N\_VEnableDotProdMulti\_ManyVector (C function), 254  
 N\_VEnableDotProdMulti\_MPIManyVector (C function), 258  
 N\_VEnableDotProdMulti\_OpenMP (C function), 218  
 N\_VEnableDotProdMulti\_OpenMPDEV (C function), 250  
 N\_VEnableDotProdMulti\_Parallel (C function), 214  
 N\_VEnableDotProdMulti\_ParHyp (C function), 224  
 N\_VEnableDotProdMulti\_Petsc (C function), 226  
 N\_VEnableDotProdMulti\_Pthreads (C function), 221  
 N\_VEnableDotProdMulti\_Serial (C function), 211  
 N\_VEnableFusedOps\_Cuda (C function), 229  
 N\_VEnableFusedOps\_Hip (C function), 234  
 N\_VEnableFusedOps\_ManyVector (C function), 254  
 N\_VEnableFusedOps\_MPIManyVector (C function), 257  
 N\_VEnableFusedOps\_OpenMP (C function), 218  
 N\_VEnableFusedOps\_OpenMPDEV (C function), 250  
 N\_VEnableFusedOps\_Parallel (C function), 214  
 N\_VEnableFusedOps\_ParHyp (C function), 224  
 N\_VEnableFusedOps\_Petsc (C function), 226  
 N\_VEnableFusedOps\_Pthreads (C function), 221  
 N\_VEnableFusedOps\_Raja (C function), 244  
 N\_VEnableFusedOps\_Serial (C function), 211  
 N\_VEnableFusedOps\_Sycl (C++ function), 239  
 N\_VEnableLinearCombination\_Cuda (C function), 229  
 N\_VEnableLinearCombination\_Hip (C function), 234  
 N\_VEnableLinearCombination\_ManyVector (C function), 254  
 N\_VEnableLinearCombination\_MPIManyVector (C function), 258  
 N\_VEnableLinearCombination\_OpenMP (C function), 218  
 N\_VEnableLinearCombination\_OpenMPDEV (C function), 250  
 N\_VEnableLinearCombination\_Parallel (C function), 214  
 N\_VEnableLinearCombination\_ParHyp (C function), 224  
 N\_VEnableLinearCombination\_Petsc (C function), 226  
 N\_VEnableLinearCombination\_Pthreads (C function), 221  
 N\_VEnableLinearCombination\_Raja (C function), 244  
 N\_VEnableLinearCombination\_Serial (C function), 211  
 N\_VEnableLinearCombination\_Sycl (C++ function), 239  
 N\_VEnableLinearCombinationVectorArray\_Cuda (C function), 230  
 N\_VEnableLinearCombinationVectorArray\_Hip (C

- function*), 235
- N\_VEnableLinearCombinationVectorArray\_OpenMP (*C function*), 218
- N\_VEnableLinearCombinationVectorArray\_OpenMPDEV (*C function*), 250
- N\_VEnableLinearCombinationVectorArray\_Parallel (*C function*), 215
- N\_VEnableLinearCombinationVectorArray\_ParHyp (*C function*), 224
- N\_VEnableLinearCombinationVectorArray\_Petsc (*C function*), 226
- N\_VEnableLinearCombinationVectorArray\_Pthreads (*C function*), 222
- N\_VEnableLinearCombinationVectorArray\_Raja (*C function*), 244
- N\_VEnableLinearCombinationVectorArray\_Serial (*C function*), 211
- N\_VEnableLinearCombinationVectorArray\_Sycl (*C++ function*), 240
- N\_VEnableLinearSumVectorArray\_Cuda (*C function*), 229
- N\_VEnableLinearSumVectorArray\_Hip (*C function*), 234
- N\_VEnableLinearSumVectorArray\_ManyVector (*C function*), 254
- N\_VEnableLinearSumVectorArray\_MPIManyVector (*C function*), 258
- N\_VEnableLinearSumVectorArray\_OpenMP (*C function*), 218
- N\_VEnableLinearSumVectorArray\_OpenMPDEV (*C function*), 250
- N\_VEnableLinearSumVectorArray\_Parallel (*C function*), 214
- N\_VEnableLinearSumVectorArray\_ParHyp (*C function*), 224
- N\_VEnableLinearSumVectorArray\_Petsc (*C function*), 226
- N\_VEnableLinearSumVectorArray\_Pthreads (*C function*), 222
- N\_VEnableLinearSumVectorArray\_Raja (*C function*), 244
- N\_VEnableLinearSumVectorArray\_Serial (*C function*), 211
- N\_VEnableLinearSumVectorArray\_Sycl (*C++ function*), 240
- N\_VEnableScaleAddMulti\_Cuda (*C function*), 229
- N\_VEnableScaleAddMulti\_Hip (*C function*), 234
- N\_VEnableScaleAddMulti\_ManyVector (*C function*), 254
- N\_VEnableScaleAddMulti\_MPIManyVector (*C function*), 258
- N\_VEnableScaleAddMulti\_OpenMP (*C function*), 218
- N\_VEnableScaleAddMulti\_OpenMPDEV (*C function*), 250
- N\_VEnableScaleAddMulti\_Parallel (*C function*), 214
- N\_VEnableScaleAddMulti\_ParHyp (*C function*), 224
- N\_VEnableScaleAddMulti\_Petsc (*C function*), 226
- N\_VEnableScaleAddMulti\_Pthreads (*C function*), 222
- N\_VEnableScaleAddMulti\_Raja (*C function*), 244
- N\_VEnableScaleAddMulti\_Serial (*C function*), 211
- N\_VEnableScaleAddMulti\_Sycl (*C++ function*), 239
- N\_VEnableScaleAddMultiVectorArray\_Cuda (*C function*), 230
- N\_VEnableScaleAddMultiVectorArray\_Hip (*C function*), 235
- N\_VEnableScaleAddMultiVectorArray\_OpenMP (*C function*), 218
- N\_VEnableScaleAddMultiVectorArray\_OpenMPDEV (*C function*), 250
- N\_VEnableScaleAddMultiVectorArray\_Parallel (*C function*), 215
- N\_VEnableScaleAddMultiVectorArray\_ParHyp (*C function*), 224
- N\_VEnableScaleAddMultiVectorArray\_Petsc (*C function*), 226
- N\_VEnableScaleAddMultiVectorArray\_Pthreads (*C function*), 222
- N\_VEnableScaleAddMultiVectorArray\_Raja (*C function*), 244
- N\_VEnableScaleAddMultiVectorArray\_Serial (*C function*), 211
- N\_VEnableScaleAddMultiVectorArray\_Sycl (*C++ function*), 240
- N\_VEnableScaleVectorArray\_Cuda (*C function*), 229
- N\_VEnableScaleVectorArray\_Hip (*C function*), 234
- N\_VEnableScaleVectorArray\_ManyVector (*C function*), 254
- N\_VEnableScaleVectorArray\_MPIManyVector (*C function*), 258
- N\_VEnableScaleVectorArray\_OpenMP (*C function*), 218
- N\_VEnableScaleVectorArray\_OpenMPDEV (*C function*), 250
- N\_VEnableScaleVectorArray\_Parallel (*C function*), 214
- N\_VEnableScaleVectorArray\_ParHyp (*C function*), 224
- N\_VEnableScaleVectorArray\_Petsc (*C function*), 226
- N\_VEnableScaleVectorArray\_Pthreads (*C function*), 222
- N\_VEnableScaleVectorArray\_Raja (*C function*), 244
- N\_VEnableScaleVectorArray\_Serial (*C function*), 211
- N\_VEnableScaleVectorArray\_Sycl (*C++ function*), 240

- N\_VEnableWrmsNormMaskVectorArray\_Cuda (C function), 230  
 N\_VEnableWrmsNormMaskVectorArray\_Hip (C function), 235  
 N\_VEnableWrmsNormMaskVectorArray\_ManyVector (C function), 254  
 N\_VEnableWrmsNormMaskVectorArray\_MPI-ManyVector (C function), 258  
 N\_VEnableWrmsNormMaskVectorArray\_OpenMP (C function), 218  
 N\_VEnableWrmsNormMaskVectorArray\_OpenMPDEV (C function), 250  
 N\_VEnableWrmsNormMaskVectorArray\_Parallel (C function), 215  
 N\_VEnableWrmsNormMaskVectorArray\_ParHyp (C function), 224  
 N\_VEnableWrmsNormMaskVectorArray\_Petsc (C function), 226  
 N\_VEnableWrmsNormMaskVectorArray\_Pthreads (C function), 222  
 N\_VEnableWrmsNormMaskVectorArray\_Serial (C function), 211  
 N\_VEnableWrmsNormVectorArray\_Cuda (C function), 230  
 N\_VEnableWrmsNormVectorArray\_Hip (C function), 234  
 N\_VEnableWrmsNormVectorArray\_ManyVector (C function), 254  
 N\_VEnableWrmsNormVectorArray\_MPIManyVector (C function), 258  
 N\_VEnableWrmsNormVectorArray\_OpenMP (C function), 218  
 N\_VEnableWrmsNormVectorArray\_OpenMPDEV (C function), 250  
 N\_VEnableWrmsNormVectorArray\_Parallel (C function), 215  
 N\_VEnableWrmsNormVectorArray\_ParHyp (C function), 224  
 N\_VEnableWrmsNormVectorArray\_Petsc (C function), 226  
 N\_VEnableWrmsNormVectorArray\_Pthreads (C function), 222  
 N\_VEnableWrmsNormVectorArray\_Serial (C function), 211  
 N\_VFreeEmpty (C function), 193  
 N\_VGetArrayPointer (C function), 196  
 N\_VGetArrayPointer\_MPIPlusX (C function), 260  
 N\_VGetCommunicator (C function), 197  
 N\_VGetDeviceArrayPointer (C function), 196  
 N\_VGetDeviceArrayPointer\_Cuda (C function), 228  
 N\_VGetDeviceArrayPointer\_Hip (C function), 233  
 N\_VGetDeviceArrayPointer\_OpenMPDEV (C function), 249  
 N\_VGetDeviceArrayPointer\_Raja (C function), 243  
 N\_VGetDeviceArrayPointer\_Sycl (C++ function), 238  
 N\_VGetHostArrayPointer\_Cuda (C function), 228  
 N\_VGetHostArrayPointer\_Hip (C function), 233  
 N\_VGetHostArrayPointer\_OpenMPDEV (C function), 249  
 N\_VGetHostArrayPointer\_Raja (C function), 243  
 N\_VGetHostArrayPointer\_Sycl (C++ function), 238  
 N\_VGetLength (C function), 197  
 N\_VGetLocalLength (C function), 197  
 N\_VGetLocalLength\_MPIPlusX (C function), 259  
 N\_VGetLocalLength\_Parallel (C function), 214  
 N\_VGetLocalVector\_MPIPlusX (C function), 259  
 N\_VGetNumSubvectors\_ManyVector (C function), 254  
 N\_VGetNumSubvectors\_MPIManyVector (C function), 257  
 N\_VGetSubvector\_ManyVector (C function), 253  
 N\_VGetSubvector\_MPIManyVector (C function), 257  
 N\_VGetSubvectorArrayPointer\_ManyVector (C function), 253  
 N\_VGetSubvectorArrayPointer\_MPIManyVector (C function), 257  
 N\_VGetSubvectorLocalLength\_ManyVector (C function), 253  
 N\_VGetSubvectorLocalLength\_MPIManyVector (C function), 257  
 N\_VGetVecAtIndexVectorArray (C function), 192  
 N\_VGetVector\_ParHyp (C function), 223  
 N\_VGetVector\_Petsc (C function), 226  
 N\_VGetVector\_Trilinos (C++ function), 251  
 N\_VGetVectorID (C function), 195  
 N\_VInv (C function), 198  
 N\_VInvTest (C function), 200  
 N\_VInvTestLocal (C function), 205  
 N\_VIsManagedMemory\_Cuda (C function), 228  
 N\_VIsManagedMemory\_Hip (C function), 233  
 N\_VIsManagedMemory\_Raja (C function), 243  
 N\_VIsManagedMemory\_Sycl (C++ function), 239  
 N\_VL1Norm (C function), 200  
 N\_VL1NormLocal (C function), 204  
 N\_VLinearCombination (C function), 201  
 N\_VLinearCombinationVectorArray (C function), 203  
 N\_VLinearSum (C function), 197  
 N\_VLinearSumVectorArray (C function), 202  
 N\_VMake\_Cuda (C function), 228  
 N\_VMake\_Hip (C function), 233  
 N\_VMake\_MPIManyVector (C function), 256  
 N\_VMake\_MPIPlusX (C function), 259  
 N\_VMake\_OpenMP (C function), 217  
 N\_VMake\_OpenMPDEV (C function), 249  
 N\_VMake\_Parallel (C function), 214  
 N\_VMake\_ParHyp (C function), 223  
 N\_VMake\_Petsc (C function), 225

`N_VMake_Pthreads` (C function), 221  
`N_VMake_Raja` (C function), 243  
`N_VMake_Serial` (C function), 210  
`N_VMake_Sycl` (C++ function), 238  
`N_VMake_Trilinos` (C++ function), 252  
`N_VMakeManaged_Cuda` (C function), 228  
`N_VMakeManaged_Hip` (C function), 233  
`N_VMakeManaged_Raja` (C function), 243  
`N_VMakeManaged_Sycl` (C++ function), 238  
`N_VMakeWithManagedAllocator_Cuda` (C function), 228  
`N_VMaxNorm` (C function), 199  
`N_VMaxNormLocal` (C function), 204  
`N_VMin` (C function), 199  
`N_VMinLocal` (C function), 204  
`N_VMinQuotient` (C function), 201  
`N_VMinQuotientLocal` (C function), 206  
`N_VNew_Cuda` (C function), 228  
`N_VNew_Hip` (C function), 233  
`N_VNew_ManyVector` (C function), 253  
`N_VNew_MPIManyVector` (C function), 256  
`N_VNew_OpenMP` (C function), 217  
`N_VNew_OpenMPDEV` (C function), 249  
`N_VNew_Parallel` (C function), 214  
`N_VNew_Pthreads` (C function), 221  
`N_VNew_Raja` (C function), 243  
`N_VNew_Serial` (C function), 210  
`N_VNew_Sycl` (C++ function), 238  
`N_VNewEmpty` (C function), 193  
`N_VNewEmpty_Cuda` (C function), 228  
`N_VNewEmpty_Hip` (C function), 233  
`N_VNewEmpty_OpenMP` (C function), 217  
`N_VNewEmpty_OpenMPDEV` (C function), 249  
`N_VNewEmpty_Parallel` (C function), 214  
`N_VNewEmpty_ParHyp` (C function), 223  
`N_VNewEmpty_Petsc` (C function), 225  
`N_VNewEmpty_Pthreads` (C function), 221  
`N_VNewEmpty_Raja` (C function), 243  
`N_VNewEmpty_Serial` (C function), 210  
`N_VNewEmpty_Sycl` (C++ function), 238  
`N_VNewManaged_Cuda` (C function), 228  
`N_VNewManaged_Hip` (C function), 233  
`N_VNewManaged_Raja` (C function), 243  
`N_VNewManaged_Sycl` (C++ function), 238  
`N_VNewVectorArray` (C function), 192  
`N_VNewWithMemHelp_Cuda` (C function), 228  
`N_VNewWithMemHelp_Hip` (C function), 233  
`N_VNewWithMemHelp_Raja` (C function), 243  
`N_VNewWithMemHelp_Sycl` (C++ function), 238  
`N_VPrint` (C function), 207  
`N_VPrint_Cuda` (C function), 229  
`N_VPrint_Hip` (C function), 234  
`N_VPrint_OpenMP` (C function), 217  
`N_VPrint_OpenMPDEV` (C function), 249  
`N_VPrint_Parallel` (C function), 214  
`N_VPrint_ParHyp` (C function), 223  
`N_VPrint_Petsc` (C function), 226  
`N_VPrint_Pthreads` (C function), 221  
`N_VPrint_Raja` (C function), 244  
`N_VPrint_Serial` (C function), 210  
`N_VPrint_Sycl` (C++ function), 239  
`N_VPrintFile` (C function), 207  
`N_VPrintFile_Cuda` (C function), 229  
`N_VPrintFile_Hip` (C function), 234  
`N_VPrintFile_OpenMP` (C function), 218  
`N_VPrintFile_OpenMPDEV` (C function), 249  
`N_VPrintFile_Parallel` (C function), 214  
`N_VPrintFile_ParHyp` (C function), 224  
`N_VPrintFile_Petsc` (C function), 226  
`N_VPrintFile_Pthreads` (C function), 221  
`N_VPrintFile_Raja` (C function), 244  
`N_VPrintFile_Serial` (C function), 210  
`N_VPrintFile_Sycl` (C++ function), 239  
`N_VProd` (C function), 197  
`N_VScale` (C function), 198  
`N_VScaleAddMulti` (C function), 201  
`N_VScaleAddMultiVectorArray` (C function), 203  
`N_VScaleVectorArray` (C function), 202  
`N_VSetArrayPointer` (C function), 196  
`N_VSetArrayPointer_MPIPlusX` (C function), 260  
`N_VSetDeviceArrayPointer_Sycl` (C++ function), 238  
`N_VSetHostArrayPointer_Sycl` (C++ function), 238  
`N_VSetKernelExecPolicy_Cuda` (C function), 228  
`N_VSetKernelExecPolicy_Hip` (C function), 233  
`N_VSetKernelExecPolicy_Sycl` (C++ function), 239  
`N_VSetSubvectorArrayPointer_ManyVector` (C function), 254  
`N_VSetSubvectorArrayPointer_MPIManyVector` (C function), 257  
`N_VSetVecAtIndexVectorArray` (C function), 193  
`N_VSpace` (C function), 196  
`N_VWL2Norm` (C function), 199  
`N_VWrmsNorm` (C function), 199  
`N_VWrmsNormMask` (C function), 199  
`N_VWrmsNormMaskVectorArray` (C function), 203  
`N_VWrmsNormVectorArray` (C function), 202  
`N_VWSqrSumLocal` (C function), 205  
`N_VWSqrSumMaskLocal` (C function), 205  
`NV_COMM_P` (C macro), 213  
`NV_CONTENT_OMP` (C macro), 216  
`NV_CONTENT_OMPDEV` (C macro), 248  
`NV_CONTENT_P` (C macro), 212  
`NV_CONTENT_PT` (C macro), 220  
`NV_CONTENT_S` (C macro), 209  
`NV_DATA_DEV_OMPDEV` (C macro), 248  
`NV_DATA_HOST_OMPDEV` (C macro), 248  
`NV_DATA_OMP` (C macro), 216



NV\_DATA\_P (C macro), 213  
 NV\_DATA\_PT (C macro), 220  
 NV\_DATA\_S (C macro), 209  
 NV\_GLOBLENGTH\_P (C macro), 213  
 NV\_Ith\_OMP (C macro), 217  
 NV\_Ith\_P (C macro), 213  
 NV\_Ith\_PT (C macro), 220  
 NV\_Ith\_S (C macro), 210  
 NV\_LENGTH\_OMP (C macro), 217  
 NV\_LENGTH\_OMPDEV (C macro), 249  
 NV\_LENGTH\_PT (C macro), 220  
 NV\_LENGTH\_S (C macro), 210  
 NV\_LOCLENGTH\_P (C macro), 213  
 NV\_NUM\_THREADS\_OMP (C macro), 217  
 NV\_NUM\_THREADS\_PT (C macro), 220  
 NV\_OWN\_DATA\_OMP (C macro), 216  
 NV\_OWN\_DATA\_OMPDEV (C macro), 248  
 NV\_OWN\_DATA\_P (C macro), 212  
 NV\_OWN\_DATA\_PT (C macro), 220  
 NV\_OWN\_DATA\_S (C macro), 209

## S

SM\_COLS\_B (C macro), 285  
 SM\_COLS\_D (C macro), 271  
 SM\_COLUMN\_B (C macro), 285  
 SM\_COLUMN\_D (C macro), 271  
 SM\_COLUMN\_ELEMENT\_B (C macro), 285  
 SM\_COLUMNS\_B (C macro), 284  
 SM\_COLUMNS\_D (C macro), 270  
 SM\_COLUMNS\_S (C macro), 292  
 SM\_CONTENT\_B (C macro), 282  
 SM\_CONTENT\_D (C macro), 270  
 SM\_CONTENT\_S (C macro), 292  
 SM\_DATA\_B (C macro), 284  
 SM\_DATA\_D (C macro), 271  
 SM\_DATA\_S (C macro), 294  
 SM\_ELEMENT\_B (C macro), 285  
 SM\_ELEMENT\_D (C macro), 271  
 SM\_INDEXPTRS\_S (C macro), 294  
 SM\_INDEXVALS\_S (C macro), 294  
 SM\_LBAND\_B (C macro), 284  
 SM\_LDATA\_B (C macro), 284  
 SM\_LDATA\_D (C macro), 270  
 SM\_LDIM\_B (C macro), 284  
 SM\_NNZ\_S (C macro), 292  
 SM\_NP\_S (C macro), 294  
 SM\_ROWS\_B (C macro), 282  
 SM\_ROWS\_D (C macro), 270  
 SM\_ROWS\_S (C macro), 292  
 SM\_SPARSETYPE\_S (C macro), 294  
 SM\_SUBAND\_B (C macro), 284  
 SM\_UBAND\_B (C macro), 284  
 SUN\_COMM\_NULL (C macro), 28  
 SUN\_LOGLEVEL\_ALL (C enumerator), 37  
 SUN\_LOGLEVEL\_DEBUG (C enumerator), 37  
 SUN\_LOGLEVEL\_ERROR (C enumerator), 37  
 SUN\_LOGLEVEL\_INFO (C enumerator), 37  
 SUN\_LOGLEVEL\_NONE (C enumerator), 37  
 SUN\_LOGLEVEL\_WARNING (C enumerator), 37  
 SUN\_OUTPUTFORMAT\_CSV (C enumerator), 27  
 SUN\_OUTPUTFORMAT\_TABLE (C enumerator), 27  
 SUNAbortErrorHandlerFn (C function), 35  
 SUNATimesFn (C type), 311  
 SUNBandMatrix (C function), 285  
 SUNBandMatrix\_Cols (C function), 286  
 SUNBandMatrix\_Column (C function), 286  
 SUNBandMatrix\_Columns (C function), 286  
 SUNBandMatrix\_Data (C function), 286  
 SUNBandMatrix\_LData (C function), 286  
 SUNBandMatrix\_LDim (C function), 286  
 SUNBandMatrix\_LowerBandwidth (C function), 286  
 SUNBandMatrix\_Print (C function), 286  
 SUNBandMatrix\_Rows (C function), 286  
 SUNBandMatrix\_StoredUpperBandwidth (C function), 286  
 SUNBandMatrix\_UpperBandwidth (C function), 286  
 SUNBandMatrixStorage (C function), 285  
 sunbooleantype (C type), 27  
 SUNComm (C type), 28  
 SUNContext (C type), 28  
 SUNContext\_ClearErrHandlers (C function), 29  
 SUNContext\_Create (C function), 28  
 SUNContext\_Free (C function), 28  
 SUNContext\_GetLastError (C function), 29  
 SUNContext\_GetLogger (C function), 30  
 SUNContext\_GetProfiler (C function), 30  
 SUNContext\_PeekLastError (C function), 29  
 SUNContext\_PopErrorHandler (C function), 29  
 SUNContext\_PushErrorHandler (C function), 29  
 SUNContext\_SetLogger (C function), 30  
 SUNContext\_SetProfiler (C function), 30  
 SUNCudaBlockReduceAtomicExecPolicy (C++ function), 232  
 SUNCudaBlockReduceExecPolicy (C++ function), 231  
 SUNCudaExecPolicy (C++ type), 230  
 SUNCudaGridStrideExecPolicy (C++ function), 231  
 SUNCudaThreadDirectExecPolicy (C++ function), 231  
 SUNDenseMatrix (C function), 271  
 SUNDenseMatrix\_Cols (C function), 272  
 SUNDenseMatrix\_Column (C function), 272  
 SUNDenseMatrix\_Columns (C function), 272  
 SUNDenseMatrix\_Data (C function), 272  
 SUNDenseMatrix\_LData (C function), 272  
 SUNDenseMatrix\_Print (C function), 271  
 SUNDenseMatrix\_Rows (C function), 271  
 sundials::cuda::ExecPolicy (C++ class), 230

`sundials::cuda::ExecPolicy::~~ExecPolicy` (C++ function), 231  
`sundials::cuda::ExecPolicy::atomic` (C++ function), 231  
`sundials::cuda::ExecPolicy::blockSize` (C++ function), 230  
`sundials::cuda::ExecPolicy::clone` (C++ function), 230  
`sundials::cuda::ExecPolicy::clone_new_stream` (C++ function), 231  
`sundials::cuda::ExecPolicy::ExecPolicy` (C++ function), 230  
`sundials::cuda::ExecPolicy::gridSize` (C++ function), 230  
`sundials::cuda::ExecPolicy::stream` (C++ function), 230  
`sundials::ginkgo::LinearSolver` (C++ class), 360  
`sundials::ginkgo::LinearSolver::~~LinearSolver` (C++ function), 360  
`sundials::ginkgo::LinearSolver::Convert` (C++ function), 360  
`sundials::ginkgo::LinearSolver::GkoExec` (C++ function), 360  
`sundials::ginkgo::LinearSolver::GkoFactory` (C++ function), 360  
`sundials::ginkgo::LinearSolver::GkoSolver` (C++ function), 360  
`sundials::ginkgo::LinearSolver::LinearSolver` (C++ function), 360  
`sundials::ginkgo::LinearSolver::NumIters` (C++ function), 360  
`sundials::ginkgo::LinearSolver::operator SUNLinearSolver` (C++ function), 360  
`sundials::ginkgo::LinearSolver::operator=` (C++ function), 360  
`sundials::ginkgo::LinearSolver::ResNorm` (C++ function), 361  
`sundials::ginkgo::LinearSolver::Setup` (C++ function), 361  
`sundials::ginkgo::LinearSolver::Solve` (C++ function), 361  
`sundials::hip::ExecPolicy` (C++ class), 235  
`sundials::hip::ExecPolicy::~~ExecPolicy` (C++ function), 235  
`sundials::hip::ExecPolicy::atomic` (C++ function), 235  
`sundials::hip::ExecPolicy::blockSize` (C++ function), 235  
`sundials::hip::ExecPolicy::clone` (C++ function), 235  
`sundials::hip::ExecPolicy::clone_new_stream` (C++ function), 235  
`sundials::hip::ExecPolicy::ExecPolicy` (C++ function), 235  
`sundials::hip::ExecPolicy::gridSize` (C++ function), 235  
`sundials::hip::ExecPolicy::stream` (C++ function), 235  
`sundials::sycl::ExecPolicy` (C++ class), 240  
`sundials::sycl::ExecPolicy::~~ExecPolicy` (C++ function), 240  
`sundials::sycl::ExecPolicy::blockSize` (C++ function), 240  
`sundials::sycl::ExecPolicy::clone` (C++ function), 240  
`sundials::sycl::ExecPolicy::gridSize` (C++ function), 240  
`SUNDIALSFileClose` (C function), 51  
`SUNDIALSFileOpen` (C function), 51  
`SUNDIALSGetVersion` (C function), 43  
`SUNDIALSGetVersionNumber` (C function), 44  
`SUNErrCode` (C type), 33  
`SUNErrorHandlerFn` (C type), 34  
`SUNFALSE` (C macro), 27  
`SUNGetErrMsg` (C function), 34  
`SUNHipBlockReduceAtomicExecPolicy` (C++ function), 237  
`SUNHipBlockReduceExecPolicy` (C++ function), 236  
`SUNHipExecPolicy` (C++ type), 235  
`SUNHipGridStrideExecPolicy` (C++ function), 236  
`SUNHipThreadDirectExecPolicy` (C++ function), 236  
`sunindextype` (C type), 27  
`SUNLinearSolver` (C type), 313  
`SUNLinearSolver_ID` (C enum), 316  
`SUNLinearSolver_Ops` (C type), 313  
`SUNLinearSolver_Type` (C enum), 306  
`SUNLinearSolver_Type.SUNLINEARSOLVER_DIRECT` (C enumerator), 306  
`SUNLinearSolver_Type.SUNLINEARSOLVER_ITERATIVE` (C enumerator), 307  
`SUNLinearSolver_Type.SUNLINEARSOLVER_MATRIX_EMBEDDED` (C enumerator), 307  
`SUNLinearSolver_Type.SUNLINEARSOLVER_MATRIX_ITERATIVE` (C enumerator), 307  
`SUNLinSol_Band` (C function), 321  
`SUNLinSol_cuSolverSp_batchQR` (C function), 357  
`SUNLinSol_cuSolverSp_batchQR_GetDescription` (C function), 358  
`SUNLinSol_cuSolverSp_batchQR_GetDeviceSpace` (C function), 358  
`SUNLinSol_cuSolverSp_batchQR_SetDescription` (C function), 358  
`SUNLinSol_Dense` (C function), 322  
`SUNLinSol_KLU` (C function), 324  
`SUNLinSol_KLUGetCommon` (C function), 325  
`SUNLinSol_KLUGetCommon.sun_klu_common` (C type), 325  
`SUNLinSol_KLUGetNumeric` (C function), 325

- SUNLinSol\_KLUGetNumeric.sun\_klu\_numeric (C type), 325
- SUNLinSol\_KLUGetSymbolic (C function), 325
- SUNLinSol\_KLUGetSymbolic.sun\_klu\_symbolic (C type), 325
- SUNLinSol\_KLUReInit (C function), 324
- SUNLinSol\_KLUSetOrdering (C function), 324
- SUNLinSol\_LapackBand (C function), 327
- SUNLinSol\_LapackDense (C function), 329
- SUNLinSol\_MagmaDense (C function), 331
- SUNLinSol\_MagmaDense\_SetAsync (C function), 331
- SUNLinSol\_OneMklDense (C function), 332
- SUNLinSol\_PCG (C function), 334
- SUNLinSol\_PCGSetMaxl (C function), 335
- SUNLinSol\_PCGSetPrecType (C function), 334
- SUNLinSol\_SPBCGS (C function), 337
- SUNLinSol\_SPBCGSSetMaxl (C function), 338
- SUNLinSol\_SPBCGSSetPrecType (C function), 337
- SUNLinSol\_SPFQMR (C function), 340
- SUNLinSol\_SPFQMRSetGSType (C function), 341
- SUNLinSol\_SPFQMRSetMaxRestarts (C function), 341
- SUNLinSol\_SPFQMRSetPrecType (C function), 340
- SUNLinSol\_SPGMR (C function), 344
- SUNLinSol\_SPGMRSetGSType (C function), 345
- SUNLinSol\_SPGMRSetMaxRestarts (C function), 345
- SUNLinSol\_SPGMRSetPrecType (C function), 344
- SUNLinSol\_SPTFQMR (C function), 348
- SUNLinSol\_SPTFQMRSetMaxl (C function), 349
- SUNLinSol\_SPTFQMRSetPrecType (C function), 348
- SUNLinSol\_SuperLUDIST (C function), 351
- SUNLinSol\_SuperLUDIST\_GetBerr (C function), 352
- SUNLinSol\_SuperLUDIST\_GetGridinfo (C function), 352
- SUNLinSol\_SuperLUDIST\_GetLUstruct (C function), 352
- SUNLinSol\_SuperLUDIST\_GetScalePermstruct (C function), 352
- SUNLinSol\_SuperLUDIST\_GetSOLVEstruct (C function), 352
- SUNLinSol\_SuperLUDIST\_GetSuperLUOptions (C function), 352
- SUNLinSol\_SuperLUDIST\_GetSuperLUStat (C function), 352
- SUNLinSol\_SuperLUMT (C function), 354
- SUNLinSol\_SuperLUMTSetOrdering (C function), 355
- SUNLinSolFree (C function), 308
- SUNLinSolFreeEmpty (C function), 315
- SUNLinSolGetID (C function), 307
- SUNLinSolGetType (C function), 307
- SUNLinSolInitialize (C function), 307
- SUNLinSolLastFlag (C function), 310
- SUNLinSolNewEmpty (C function), 315
- SUNLinSolNumIters (C function), 310
- SUNLinSolResid (C function), 310
- SUNLinSolResNorm (C function), 310
- SUNLinSolSetATimes (C function), 309
- SUNLinSolSetPreconditioner (C function), 309
- SUNLinSolSetScalingVectors (C function), 309
- SUNLinSolSetup (C function), 307
- SUNLinSolSetZeroGuess (C function), 309
- SUNLinSolSolve (C function), 308
- SUNLinSolSpace (C function), 310
- SUNLogErrorHandlerFn (C function), 35
- SUNLogger (C type), 37
- SUNLogger\_Create (C function), 37
- SUNLogger\_CreateFromEnv (C function), 37
- SUNLogger\_Destroy (C function), 39
- SUNLogger\_Flush (C function), 39
- SUNLogger\_GetOutputRank (C function), 39
- SUNLogger\_QueueMsg (C function), 38
- SUNLogger\_SetDebugFilename (C function), 38
- SUNLogger\_SetErrorFilename (C function), 38
- SUNLogger\_SetInfoFilename (C function), 38
- SUNLogger\_SetWarningFilename (C function), 38
- SUNLogLevel (C enum), 37
- SUNMatClone (C function), 268
- SUNMatCopy (C function), 268
- SUNMatCopyOps (C function), 267
- SUNMatDestroy (C function), 268
- SUNMatFreeEmpty (C function), 267
- SUNMatGetID (C function), 268
- SUNMatMatvec (C function), 269
- SUNMatMatvecSetup (C function), 269
- SUNMatNewEmpty (C function), 267
- SUNMatrix (C type), 265
- SUNMatrix\_cuSparse\_BlockColumns (C function), 289
- SUNMatrix\_cuSparse\_BlockData (C function), 289
- SUNMatrix\_cuSparse\_BlockNNZ (C function), 289
- SUNMatrix\_cuSparse\_BlockRows (C function), 289
- SUNMatrix\_cuSparse\_Columns (C function), 288
- SUNMatrix\_cuSparse\_CopyFromDevice (C function), 289
- SUNMatrix\_cuSparse\_CopyToDevice (C function), 289
- SUNMatrix\_cuSparse\_Data (C function), 289
- SUNMatrix\_cuSparse\_IndexPointers (C function), 289
- SUNMatrix\_cuSparse\_IndexValues (C function), 289
- SUNMatrix\_cuSparse\_MakeCSR (C function), 288
- SUNMatrix\_cuSparse\_MatDescr (C function), 289
- SUNMatrix\_cuSparse\_NewBlockCSR (C function), 288
- SUNMatrix\_cuSparse\_NewCSR (C function), 288
- SUNMatrix\_cuSparse\_NNZ (C function), 288
- SUNMatrix\_cuSparse\_NumBlocks (C function), 289
- SUNMatrix\_cuSparse\_Rows (C function), 288
- SUNMatrix\_cuSparse\_SetFixedPattern (C function), 289

SUNMatrix\_cuSparse\_SetKernelExecPolicy (C function), 290

SUNMatrix\_cuSparse\_SparseType (C function), 288

SUNMatrix\_ID (C type), 267

SUNMatrix\_MagmaDense (C function), 273

SUNMatrix\_MagmaDense\_Block (C function), 275

SUNMatrix\_MagmaDense\_BlockColumn (C function), 275

SUNMatrix\_MagmaDense\_BlockColumns (C function), 274

SUNMatrix\_MagmaDense\_BlockData (C function), 275

SUNMatrix\_MagmaDense\_BlockRows (C function), 274

SUNMatrix\_MagmaDense\_Column (C function), 275

SUNMatrix\_MagmaDense\_Columns (C function), 274

SUNMatrix\_MagmaDense\_CopyFromDevice (C function), 276

SUNMatrix\_MagmaDense\_CopyToDevice (C function), 276

SUNMatrix\_MagmaDense\_Data (C function), 275

SUNMatrix\_MagmaDense\_LData (C function), 274

SUNMatrix\_MagmaDense\_NumBlocks (C function), 275

SUNMatrix\_MagmaDense\_Rows (C function), 274

SUNMatrix\_MagmaDenseBlock (C function), 273

SUNMatrix\_OneMklDense (C++ function), 277

SUNMatrix\_OneMklDense\_Block (C function), 280

SUNMatrix\_OneMklDense\_BlockColumn (C function), 280

SUNMatrix\_OneMklDense\_BlockColumns (C function), 279

SUNMatrix\_OneMklDense\_BlockData (C function), 280

SUNMatrix\_OneMklDense\_BlockLData (C function), 280

SUNMatrix\_OneMklDense\_BlockRows (C function), 279

SUNMatrix\_OneMklDense\_Column (C function), 279

SUNMatrix\_OneMklDense\_Columns (C function), 278

SUNMatrix\_OneMklDense\_CopyFromDevice (C function), 281

SUNMatrix\_OneMklDense\_CopyToDevice (C function), 281

SUNMatrix\_OneMklDense\_Data (C function), 279

SUNMatrix\_OneMklDense\_LData (C function), 279

SUNMatrix\_OneMklDense\_NumBlocks (C function), 279

SUNMatrix\_OneMklDense\_Rows (C function), 278

SUNMatrix\_OneMklDenseBlock (C++ function), 278

SUNMatrix\_SLUNRloc (C function), 297

SUNMatrix\_SLUNRloc\_OwnData (C function), 297

SUNMatrix\_SLUNRloc\_Print (C function), 297

SUNMatrix\_SLUNRloc\_ProcessGrid (C function), 297

SUNMatrix\_SLUNRloc\_SuperMatrix (C function), 297

SUNMatScaleAdd (C function), 269

SUNMatScaleAddI (C function), 269

SUNMatSpace (C function), 268

SUNMatZero (C function), 268

SUNMemory (C type), 387

SUNMemory.SUNMemory\_ (C struct), 387

SUNMemory.SUNMemory\_.bytes (C member), 387

SUNMemory.SUNMemory\_.own (C member), 387

SUNMemory.SUNMemory\_.ptr (C member), 387

SUNMemory.SUNMemory\_.type (C member), 387

SUNMemoryHelper (C type), 388

SUNMemoryHelper.SUNMemoryHelper\_ (C struct), 388

SUNMemoryHelper.SUNMemoryHelper\_.content (C member), 388

SUNMemoryHelper.SUNMemoryHelper\_.ops (C member), 388

SUNMemoryHelper.SUNMemoryHelper\_.sunctx (C member), 388

SUNMemoryHelper\_Alias (C function), 390

SUNMemoryHelper\_Alloc (C function), 389

SUNMemoryHelper\_Alloc\_Cuda (C function), 393

SUNMemoryHelper\_Alloc\_Hip (C function), 395

SUNMemoryHelper\_Alloc\_Sycl (C function), 397

SUNMemoryHelper\_Clone (C function), 391

SUNMemoryHelper\_Copy (C function), 389

SUNMemoryHelper\_Copy\_Cuda (C function), 393

SUNMemoryHelper\_Copy\_Hip (C function), 395

SUNMemoryHelper\_Copy\_Sycl (C function), 397

SUNMemoryHelper\_CopyAsync (C function), 391

SUNMemoryHelper\_CopyAsync\_Cuda (C function), 394

SUNMemoryHelper\_CopyAsync\_Hip (C function), 395

SUNMemoryHelper\_CopyAsync\_Sycl (C function), 398

SUNMemoryHelper\_CopyOps (C function), 390

SUNMemoryHelper\_Cuda (C function), 392

SUNMemoryHelper\_Dealloc (C function), 389

SUNMemoryHelper\_Dealloc\_Cuda (C function), 393

SUNMemoryHelper\_Dealloc\_Hip (C function), 395

SUNMemoryHelper\_Dealloc\_Sycl (C function), 397

SUNMemoryHelper\_Destroy (C function), 392

SUNMemoryHelper\_GetAllocStats (C function), 391

SUNMemoryHelper\_GetAllocStats\_Cuda (C function), 394

SUNMemoryHelper\_GetAllocStats\_Hip (C function), 396

SUNMemoryHelper\_GetAllocStats\_Sycl (C function), 398

SUNMemoryHelper\_Hip (C function), 394

SUNMemoryHelper\_NewEmpty (C function), 390

SUNMemoryHelper\_Ops (C type), 388

SUNMemoryHelper\_Ops.SUNMemoryHelper\_Ops\_ (C struct), 388

SUNMemoryHelper\_Ops.SUNMemoryHelper\_Ops\_.alloc (C member), 388

SUNMemoryHelper\_Ops.SUNMemoryHelper\_Ops\_.clone (C member), 388



- SUNMemoryHelper\_Ops.SUNMemoryHelper\_Ops\_-.copy (C member), 388
- SUNMemoryHelper\_Ops.SUNMemoryHelper\_Ops\_-.copyasync (C member), 388
- SUNMemoryHelper\_Ops.SUNMemoryHelper\_Ops\_-.dealloc (C member), 388
- SUNMemoryHelper\_Ops.SUNMemoryHelper\_Ops\_-.destroy (C member), 388
- SUNMemoryHelper\_Ops.SUNMemoryHelper\_Ops\_-.getallocstats (C member), 388
- SUNMemoryHelper\_Sycl (C function), 396
- SUNMemoryHelper\_Wrap (C function), 390
- SUNMemoryNewEmpty (C function), 387
- SUNMemoryType (C enum), 387
- SUNMemoryType.SUNMEMTYPE\_DEVICE (C enumerator), 388
- SUNMemoryType.SUNMEMTYPE\_HOST (C enumerator), 388
- SUNMemoryType.SUNMEMTYPE\_PINNED (C enumerator), 388
- SUNMemoryType.SUNMEMTYPE\_UVM (C enumerator), 388
- SUNMPIAbortErrHandlerFn (C function), 35
- SUNNonlinearSolver (C type), 372
- SUNNonlinearSolver\_Ops (C type), 372
- SUNNonlinearSolver\_Type (C enum), 366
- SUNNonlinearSolver\_-Type.SUNNONLINEARSOLVER\_FIXEDPOINT (C enumerator), 366
- SUNNonlinearSolver\_-Type.SUNNONLINEARSOLVER\_ROOTFIND (C enumerator), 366
- SUNNonlinSol\_FixedPoint (C function), 382
- SUNNonlinSol\_Newton (C function), 379
- SUNNonlinSol\_PetscSNES (C function), 385
- SUNNonlinSolConvTestFn (C type), 371
- SUNNonlinSolFree (C function), 367
- SUNNonlinSolFreeEmpty (C function), 374
- SUNNonlinSolGetCurIter (C function), 369
- SUNNonlinSolGetNumConvFails (C function), 369
- SUNNonlinSolGetNumIters (C function), 369
- SUNNonlinSolGetPetscError\_PetscSNES (C function), 385
- SUNNonlinSolGetSNES\_PetscSNES (C function), 385
- SUNNonlinSolGetSysFn\_FixedPoint (C function), 382
- SUNNonlinSolGetSysFn\_Newton (C function), 379
- SUNNonlinSolGetSysFn\_PetscSNES (C function), 386
- SUNNonlinSolGetType (C function), 366
- SUNNonlinSolInitialize (C function), 366
- SUNNonlinSolLSetupFn (C type), 370
- SUNNonlinSolLSolveFn (C type), 371
- SUNNonlinSolNewEmpty (C function), 373
- SUNNonlinSolSetConvTestFn (C function), 368
- SUNNonlinSolSetDamping\_FixedPoint (C function), 382
- SUNNonlinSolSetLSetupFn (C function), 367
- SUNNonlinSolSetLSolveFn (C function), 368
- SUNNonlinSolSetMaxIters (C function), 368
- SUNNonlinSolSetSysFn (C function), 367
- SUNNonlinSolSetup (C function), 366
- SUNNonlinSolSolve (C function), 366
- SUNNonlinSolSysFn (C type), 370
- SUNOutputFormat (C enum), 27
- SUNProfiler (C type), 41
- SUNProfiler\_Begin (C function), 41
- SUNProfiler\_Create (C function), 41
- SUNProfiler\_End (C function), 42
- SUNProfiler\_Free (C function), 41
- SUNProfiler\_GetElapsedTime (C function), 42
- SUNProfiler\_GetTimerResolution (C function), 42
- SUNProfiler\_Print (C function), 42
- SUNProfiler\_Reset (C function), 42
- SUNPSolveFn (C type), 311
- SUNPSolveFn (C type), 311
- sunrealtype (C type), 26
- SUNSparseFromBandMatrix (C function), 295
- SUNSparseFromDenseMatrix (C function), 295
- SUNSparseMatrix (C function), 294
- SUNSparseMatrix\_Columns (C function), 295
- SUNSparseMatrix\_Data (C function), 296
- SUNSparseMatrix\_IndexPointers (C function), 296
- SUNSparseMatrix\_IndexValues (C function), 296
- SUNSparseMatrix\_NNZ (C function), 295
- SUNSparseMatrix\_NP (C function), 295
- SUNSparseMatrix\_Print (C function), 295
- SUNSparseMatrix\_Realloc (C function), 295
- SUNSparseMatrix\_Reallocate (C function), 295
- SUNSparseMatrix\_Rows (C function), 295
- SUNSparseMatrix\_SparseType (C function), 295
- SUNSyclBlockReduceExecPolicy (C++ function), 242
- SUNSyclExecPolicy (C++ type), 240
- SUNSyclGridStrideExecPolicy (C++ function), 241
- SUNSyclThreadDirectExecPolicy (C++ function), 241
- SUNTRUE (C macro), 27

## V

- Vector (C++ class), 246
- Vector::~Vector (C++ function), 247
- Vector::Convert (C++ function), 247
- Vector::exec\_space (C++ type), 246
- Vector::host\_view\_type (C++ type), 246
- Vector::HostView (C++ function), 247
- Vector::Length (C++ function), 247
- Vector::memory\_space (C++ type), 246
- Vector::operator N\_Vector (C++ function), 247
- Vector::operator= (C++ function), 247

`Vector::range_policy` (C++ *type*), 246  
`Vector::size_type` (C++ *type*), 246  
`Vector::Vector` (C++ *function*), 246, 247  
`Vector::View` (C++ *function*), 247  
`Vector::view_type` (C++ *type*), 246  
`vector_type` (C++ *type*), 251