

User Documentation for IDA v7.0.0

SUNDIALS v7.0.0

Alan C. Hindmarsh¹, Radu Serban¹, Cody J. Balos¹,
David J. Gardner¹, Daniel R. Reynolds², and Carol S. Woodward¹

¹*Center for Applied Scientific Computing, Lawrence Livermore National Laboratory*

²*Department of Mathematics, Southern Methodist University*

February 29, 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

1	Introduction	1
1.1	Changes from previous versions	2
1.2	Reading this User Guide	26
1.3	SUNDIALS License and Notices	26
2	Mathematical Considerations	29
2.1	Initial Condition	29
2.2	IVP solution	29
2.3	Preconditioning	33
2.4	Rootfinding	34
3	Code Organization	37
3.1	IDA organization	37
4	Using SUNDIALS	41
4.1	Data Types	41
4.2	The SUNContext Type	43
4.3	Error Checking	49
4.4	Status and Error Logging	51
4.5	Performance Profiling	55
4.6	Getting Version Information	59
4.7	Fortran Interface	60
4.8	Features for GPU Accelerated Computing	68
5	Using IDA for IVP Solution	71
5.1	Access to library and header files	71
5.2	A skeleton of the user's main program	72
5.3	User-callable functions	74
5.4	User-supplied functions	117
5.5	Preconditioner modules	124
6	Vector Data Structures	131
6.1	Description of the NVECTOR Modules	131
6.2	Description of the NVECTOR operations	138
6.3	NVECTOR functions used by IDA	150
6.4	The NVECTOR_SERIAL Module	151
6.5	The NVECTOR_PARALLEL Module	154
6.6	The NVECTOR_OPENMP Module	158
6.7	The NVECTOR_PTHREADS Module	162
6.8	The NVECTOR_PARHYP Module	165
6.9	The NVECTOR_PETSC Module	167
6.10	The NVECTOR_CUDA Module	169
6.11	The NVECTOR_HIP Module	175

6.12	The NVECTOR_SYCL Module	180
6.13	The NVECTOR_RAJA Module	185
6.14	The NVECTOR_KOKKOS Module	188
6.15	The NVECTOR_OPENMPDEV Module	190
6.16	The NVECTOR_TRILINOS Module	194
6.17	The NVECTOR_MANYVECTOR Module	195
6.18	The NVECTOR_MPIMANYVECTOR Module	198
6.19	The NVECTOR_MPIPLUSX Module	201
6.20	NVECTOR Examples	203
7	Matrix Data Structures	207
7.1	Description of the SUNMATRIX Modules	207
7.2	Description of the SUNMATRIX operations	209
7.3	The SUNMATRIX_DENSE Module	211
7.4	The SUNMATRIX_MAGMADENSE Module	214
7.5	The SUNMATRIX_ONEMKLDENSE Module	218
7.6	The SUNMATRIX_BAND Module	223
7.7	The SUNMATRIX_CUSPARSE Module	229
7.8	The SUNMATRIX_SPARSE Module	232
7.9	The SUNMATRIX_SLUNRLOC Module	238
7.10	The SUNMATRIX_GINKGO Module	239
7.11	The SUNMATRIX_KOKKOSDENSE Module	241
7.12	SUNMATRIX Examples	245
7.13	SUNMatrix functions used by IDA	245
8	Linear Algebraic Solvers	247
8.1	The SUNLinearSolver API	248
8.2	IDA SUNLinearSolver interface	260
8.3	The SUNLinSol_Band Module	262
8.4	The SUNLinSol_Dense Module	264
8.5	The SUNLinSol_KLU Module	265
8.6	The SUNLinSol_LapackBand Module	268
8.7	The SUNLinSol_LapackDense Module	270
8.8	The SUNLinSol_MagmaDense Module	272
8.9	The SUNLinSol_OneMklDense Module	273
8.10	The SUNLinSol_PCG Module	274
8.11	The SUNLinSol_SPBCGS Module	278
8.12	The SUNLinSol_SPGMR Module	281
8.13	The SUNLinSol_SPGMR Module	285
8.14	The SUNLinSol_SPTFQMR Module	289
8.15	The SUNLinSol_SuperLUDIST Module	292
8.16	The SUNLinSol_SuperLUMT Module	296
8.17	The SUNLinSol_cuSolverSp_batchQR Module	298
8.18	The SUNLINEARSOLVER_GINKGO Module	300
8.19	The SUNLINEARSOLVER_KOKKOSDENSE Module	303
8.20	SUNLinearSolver Examples	304
9	Nonlinear Algebraic Solvers	307
9.1	The SUNNonlinearSolver API	307
9.2	IDA SUNNonlinearSolver interface	316
9.3	The SUNNonlinSol_Newton implementation	318
9.4	The SUNNonlinSol_FixedPoint implementation	321
9.5	The SUNNonlinSol_PetscSNES implementation	324
10	Tools for Memory Management	327

10.1	The SUNMemoryHelper API	327
10.2	The SUNMemoryHelper_Cuda Implementation	332
10.3	The SUNMemoryHelper_Hip Implementation	334
10.4	The SUNMemoryHelper_Sycl Implementation	336
11	Acquiring SUNDIALS	339
12	Building and Installing with CMake	341
12.1	Configuring, building, and installing on Unix-like systems	341
12.2	Configuration options (Unix/Linux)	345
12.3	Configuration examples	357
12.4	Working with external Libraries	358
12.5	Testing the build and installation	362
12.6	Building and Running Examples	362
12.7	Configuring, building, and installing on Windows	363
12.8	Installed libraries and exported header files	363
12.9	Using SUNDIALS in your prpjct	364
12.10	Using SUNDIALS as a Third Party Library in other CMake Projects	364
12.11	Table of SUNDIALS libraries and header files	365
13	IDA Constants	369
13.1	IDA input constants	369
13.2	IDA output constants	369
14	Release History	371
	Bibliography	373
	Index	377

Chapter 1

Introduction

IDA is part of a software family called SUNDIALS: SUite of Nonlinear and Differential/ALgebraic equation Solvers [36]. This suite consists of CVODE, ARKODE, KINSOL, and IDA, and variants of these with sensitivity analysis capabilities, CVODES and IDAS.

IDA is a general purpose solver for the initial value problem (IVP) for systems of differential-algebraic equations (DAEs). The name IDA stands for Implicit Differential-Algebraic solver. IDA is based on DASPK [17, 18], 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 IDA shares significant modules previously written within CASC at LLNL to support the ordinary differential equation (ODE) solvers CVODE [23, 39] and PVODE [21, 22], and also the nonlinear system solver KINSOL [40].

At present, IDA 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) [49], FGMRES (Flexible Generalized Minimum RESidual) [48], Bi-CGStab (Bi-Conjugate Gradient Stabilized) [53], TFQMR (Transpose-Free Quasi-Minimal Residual) [31], and PCG (Preconditioned Conjugate Gradient) [33] 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-CGFStab 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.

1.1 Changes from previous versions

1.1.1 Changes in v7.0.0

Major Features

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](#).

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 functions `IDASetErrFile` and `IDASetHandlerErrFn` have been removed. Users of these functions can use the functions `SUNContext_PushErrorHandler()`, 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
<code>SUNLS_SUCCESS</code>	<code>SUN_SUCCESS</code>
<code>SUNLS_UNRECOV_FAILURE</code>	no replacement (value was unused)
<code>SUNLS_MEM_NULL</code>	<code>SUN_ERR_ARG_CORRUPT</code>
<code>SUNLS_ILL_INPUT</code>	<code>SUN_ERR_ARG_*</code>
<code>SUNLS_MEM_FAIL</code>	<code>SUN_ERR_MEM_FAIL</code>
<code>SUNLS_PACKAGE_FAIL_UNREC</code>	<code>SUN_ERR_EXT_FAIL</code>
<code>SUNLS_VECTOROP_ERR</code>	<code>SUN_ERR_OP_FAIL</code>
<code>SUN-NLS_SUCCESS</code>	<code>SUN_SUCCESS</code>
<code>SUN-NLS_MEM_NULL</code>	<code>SUN_ERR_ARG_CORRUPT</code>
<code>SUN-NLS_MEM_FAIL</code>	<code>SUN_ERR_MEM_FAIL</code>
<code>SUN-NLS_ILL_INPUT</code>	<code>SUN_ERR_ARG_*</code>
<code>SUN-NLS_VECTOROP_ERR</code>	<code>SUN_ERR_OP_FAIL</code>
<code>SUN-NLS_EXT_FAIL</code>	<code>SUN_ERR_EXT_FAIL</code>
<code>SUNMAT_SUCCESS</code>	<code>SUN_SUCCESS</code>
<code>SUNMAT_ILL_INPUT</code>	<code>SUN_ERR_ARG_*</code>
<code>SUNMAT_MEM_FAIL</code>	<code>SUN_ERR_MEM_FAIL</code>
<code>SUNMAT_OPERATION_FAIL</code>	<code>SUN_ERR_OP_FAIL</code>
<code>SUNMAT_MATVEC_SETUP_REQUIRED</code>	<code>SUN_ERR_OP_FAIL</code>

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`
 - `SUNMemorNewEmpty()`
 - `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:

- 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
```

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);
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
```

Minor changes

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](#).

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.

Converted most previous Fortran 77 and 90 examples to use SUNDIALS' Fortran 2003 interface.

1.1.2 Changes in v6.7.0

Improved computational complexity of `SUNMatScaleAddI_Sparse` from $O(M*N)$ to $O(NNZ)$.

Added Fortran support for the LAPACK dense `SUNLinearSolver` implementation.

Fixed a regression introduced by the stop time bug fix in v6.6.1 where CVODE 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 scaling bug in `SUNMatScaleAddI_Sparse` for non-square matrices.

Changed the `SUNProfiler` so that it does not rely on `MPI_WTime` in any case. This fixes [GitHub Issue #312](#).

Fixed missing soversions in some `SUNLinearSolver` and `SUNNonlinearSolver` CMake targets.

1.1.3 Changes in v6.6.2

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.

1.1.4 Changes in v6.6.1

Updated the Tpetra NVector interface to support Trilinos 14.

Fixed a memory leak when destroying a CUDA, HIP, SYCL, or system SUNMemoryHelper object.

Fixed a bug where the stop time may not be cleared when using normal mode if the requested output time is the same as the stop time.

1.1.5 Changes in v6.6.0

Updated the F2003 utility routines *SUNDIALSFileOpen()* and *SUNDIALSFileClose()* to support user specification of stdout and stderr strings for the output file names.

1.1.6 Changes in v6.5.1

Added the function *IDAClearStopTime()* to disable a previously set stop time.

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.

1.1.7 Changes in v6.5.0

Added the functions *IDAGetJac()*, *IDAGetJacCj()*, *IDAGetJacTime()*, *IDAGetJacNumSteps()* to assist in debugging simulations utilizing a matrix-based linear solver.

Added support for the SYCL backend with RAJA 2022.x.y.

Fixed an underflow bug during root finding.

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 the documentation for the SUNMemoryHelper classes for more details.

Added support for CUDA v12.

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 FN_VGetArrayPointer functions as well as the FSUNDenseMatrix_Data, FSUNBandMatrix_Data, FSUNSparseMatrix_Data, FSUNSparseMatrix_IndexValues, and FSUNSparseMatrix_IndexPointers functions. Compiling and running code that uses the SUNDIALS Fortran interfaces with bounds checking will now work.

1.1.8 Changes in v6.4.1

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 SUNLINSOL_LAPACKBAND and SUNLINSOL_LAPACKDENSE modules which would cause the tests to fail on some platforms.

1.1.9 Changes in v6.4.0

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.

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 support for the [Ginkgo](#) linear algebra library. This support includes new `SUNMatrix` and `SUNLinearSolver` implementations, see the sections §7.10 and §8.18.

Added new `NVector`, dense `SUNMatrix`, and dense `SUNLinearSolver` implementations utilizing the [Kokkos Ecosystem](#) for performance portability, see sections §6.14, §7.11, and §8.19 for more information.

Fixed a bug in the CUDA and HIP vectors where `N_VMaxNorm()` would return the minimum positive floating-point value for the zero vector.

1.1.10 Changes in v6.3.0

Added the function `IDAGetUserData()` to retrieve the user data pointer provided to `IDASetUserData()`.

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'.

1.1.11 Changes in v6.2.0

Added the `SUNLogger` API which provides a SUNDIALS-wide mechanism for logging of errors, warnings, informational output, and debugging output.

Deprecated `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()` it is recommended to use the `SUNLogger` API instead. The `SUNLinSolSetInfoFile_*` and `SUNNonlinSolSetInfoFile_*` family of functions are now enabled by setting the CMake option `SUNDIALS_LOGGING_LEVEL` to a value ≥ 3 .

Added the function `SUNProfiler_Reset()` to reset the region timings and counters to zero.

Added the function `IDAPrintAllStats()` to output all of the integrator, nonlinear solver, linear solver, and other statistics in one call. The file `scripts/sundials_csv.py` contains functions for parsing the comma-separated value output files.

Added the function `IDASetDeltaCjLSetup()` to adjust the parameter that determines when a change in c_j requires calling the linear solver setup function.

Added the functions `IDASetEtaFixedStepBounds()`, `IDASetEtaMax()`, `IDASetEtaMin()`, `IDASetEtaLow()`, `IDASetEtaMinErrFail()`, and `IDASetEtaConvFail()` to adjust various parameters controlling changes in step size.

Added the function `IDASetMinStep()` to set a minimum step size.

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.

Fixed the `SUNContext` convenience class for C++ users to disallow copy construction and allow move construction.

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_mpmianyvector.h` has been corrected to enable using both the `ManyVector` and `MPI-ManyVector` `NVector` implementations in the same simulation.

Changed exported SUNDIALS PETSc CMake targets to be `INTERFACE IMPORTED` instead of `UNKNOWN IMPORTED`.

A bug was fixed in the functions `IDAGetNumNonlinSolvConvFails()` and `IDAGetNonlinSolvStats()` where the number of nonlinear solver failures 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. These functions have been updated to return the total number of nonlinear solver failures. As such users may see an increase in the number of failures reported.

The function `IDAGetNumStepSolveFails()` has been added to retrieve the number of failed steps due to a nonlinear solver failure. The count returned by this function will match those previously returned by `IDAGetNumNonlinSolvConvFails()` and `IDAGetNonlinSolvStats()`.

1.1.12 Changes in v6.1.1

Fixed exported `SUNDIALSConfig.cmake`.

1.1.13 Changes in v6.1.0

Added new reduction implementations for the CUDA and HIP `NVECTORs` 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.

Fixed exported `SUNDIALSConfig.cmake` when profiling is enabled without Caliper.

Fixed `sundials_export.h` include in `sundials_config.h`.

Fixed memory leaks in the `SUNLINSOL_SUPERLUMT` linear solver.

1.1.14 Changes in v6.0.0

SUNContext

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 documentation section on the *SUNContext* for more details.

A script `upgrade-to-sundials-6-from-5.sh` has been provided with the release (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>
```

SUNProfiler

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.

SUNMemoryHelper

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.

NVector

Two new optional vector operations, *N_VDotProdMultiLocal()* and *N_VDotProdMultiAllReduce()*, have been added to support low-synchronization methods for Anderson acceleration.

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 NVector has been replaced with the `sundials::trilinos::nvector_tpetra` namespace.

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.

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.

SUNLinearSolver

The following previously deprecated functions have been removed:

Removed	Replacement
SUNBandLinearSolver	<i>SUNLinSol_Band()</i>
SUNDenseLinearSolver	<i>SUNLinSol_Dense()</i>
SUNKLU	<i>SUNLinSol_KLU()</i>
SUNKLUReInit	<i>SUNLinSol_KLUReInit()</i>
SUNKLUSetOrdering	<i>SUNLinSol_KLUSetOrdering()</i>
SUNLapackBand	<i>SUNLinSol_LapackBand()</i>
SUNLapackDense	<i>SUNLinSol_LapackDense()</i>
SUNPCG	<i>SUNLinSol_PCG()</i>
SUNPCGSetPrecType	<i>SUNLinSol_PCGSetPrecType()</i>
SUNPCGSetMaxl	<i>SUNLinSol_PCGSetMaxl()</i>
SUNSPBCGS	<i>SUNLinSol_SPBCGS()</i>
SUNSPBCGSSetPrecType	<i>SUNLinSol_SPBCGSSetPrecType()</i>
SUNSPBCGSSetMaxl	<i>SUNLinSol_SPBCGSSetMaxl()</i>
SUNSPFGMR	<i>SUNLinSol_SPFGMR()</i>
SUNSPFGMRSetPrecType	<i>SUNLinSol_SPFGMRSetPrecType()</i>
SUNSPFGMRSetGType	<i>SUNLinSol_SPFGMRSetGType()</i>
SUNSPFGMRSetMaxRestarts	<i>SUNLinSol_SPFGMRSetMaxRestarts()</i>
SUNSPGMR	<i>SUNLinSol_SPGMR()</i>
SUNSPGMRSetPrecType	<i>SUNLinSol_SPGMRSetPrecType()</i>
SUNSPGMRSetGType	<i>SUNLinSol_SPGMRSetGType()</i>
SUNSPGMRSetMaxRestarts	<i>SUNLinSol_SPGMRSetMaxRestarts()</i>
SUNSPTFQMR	<i>SUNLinSol_SPTFQMR()</i>
SUNSPTFQMRSetPrecType	<i>SUNLinSol_SPTFQMRSetPrecType()</i>
SUNSPTFQMRSetMaxl	<i>SUNLinSol_SPTFQMRSetMaxl()</i>
SUNSuperLUMT	<i>SUNLinSol_SuperLUMT()</i>
SUNSuperLUMTSetOrdering	<i>SUNLinSol_SuperLUMTSetOrdering()</i>

IDA

The IDA Fortran 77 interface has been removed. See §4.7 and the F2003 example programs for more details using the SUNDIALS Fortran 2003 module interfaces.

Deprecations

In addition to the deprecations noted elsewhere, 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
realttype	sunrealttype
booleanttype	sunbooleanttype
RCONST	SUN_RCONST
BIG_REAL	SUN_BIG_REAL
SMALL_REAL	SUN_SMALL_REAL
UNIT_ROUNDOFF	SUN_UNIT_ROUNDOFF
PREC_NONE	SUN_PREC_NONE
PREC_LEFT	SUN_PREC_LEFT
PREC_RIGHT	SUN_PREC_RIGHT
PREC_BOTH	SUN_PREC_BOTH
MODIFIED_GS	SUN_MODIFIED_GS
CLASSICAL_GS	SUN_CLASSICAL_GS
ATimesFn	SUNATimesFn
PSetupFn	SUNPSetupFn
PSolveFn	SUNPSolveFn
DlsMat	SUNDlsMat
DENSE_COL	SUNDLS_DENSE_COL
DENSE_ELEM	SUNDLS_DENSE_ELEM
BAND_COL	SUNDLS_BAND_COL
BAND_COL_ELEM	SUNDLS_BAND_COL_ELEM
BAND_ELEM	SUNDLS_BAND_ELEM

In addition, the following functions are now deprecated (compile-time warnings will be thrown if supported by the compiler):

Deprecated Name	New Name
IDASpilsSetLinearSolver	IDASetLinearSolver
IDASpilsSetPreconditioner	IDASetPreconditioner
IDASpilsSetJacTimes	IDASetJacTimes
IDASpilsSetEpsLin	IDASetEpsLin
IDASpilsSetIncrementFactor	IDASetIncrementFactor
IDASpilsGetWorkSpace	IDAGetLinWorkSpace
IDASpilsGetNumPrecEvals	IDAGetNumPrecEvals
IDASpilsGetNumPrecSolves	IDAGetNumPrecSolves
IDASpilsGetNumLinIters	IDAGetNumLinIters
IDASpilsGetNumConvFails	IDAGetNumLinConvFails
IDASpilsGetNumJTSetupEvals	IDAGetNumJTSetupEvals
IDASpilsGetNumJtimesEvals	IDAGetNumJtimesEvals
IDASpilsGetNumResEvals	IDAGetNumLinResEvals
IDASpilsGetLastFlag	IDAGetLastLinFlag
IDASpilsGetReturnFlagName	IDAGetLinReturnFlagName
IDADlsSetLinearSolver	IDASetLinearSolver
IDADlsSetJacFn	IDASetJacFn
IDADlsGetWorkSpace	IDAGetLinWorkSpace
IDADlsGetNumJacEvals	IDAGetNumJacEvals
IDADlsGetNumResEvals	IDAGetNumLinResEvals
IDADlsGetLastFlag	IDAGetLastLinFlag
IDADlsGetReturnFlagName	IDAGetLinReturnFlagName
DenseGETRF	SUNDlsMat_DenseGETRF

continues on next page

Table 1.1 – continued from previous page

Deprecated Name	New Name
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
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.

1.1.15 Changes in v5.8.0

The *RAJA N_Vector* implementation has been updated to support the SYCL backend in addition to the CUDA and HIP backends. Users can choose the backend when configuring SUNDIALS by using the `SUNDIALS_RAJA_BACKENDS` CMake variable. This module remains experimental and is subject to change from version to version.

A 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 module 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.

IDA now supports a new “matrix-embedded” `SUNLinearSolver` 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 the function `IDASetNlsResFn()` to supply an alternative residual side function for use within nonlinear system function evaluations.

The installed `SUNDIALSConfig.cmake` file now supports the `COMPONENTS` option to `find_package`.

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.

1.1.16 Changes in v5.7.0

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 module 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.

1.1.17 Changes in v5.6.1

Fixed a bug in the SUNDIALS CMake 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.

1.1.18 Changes in v5.6.0

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 module is considered experimental and is subject to change from version to version.

The `NVECTOR_RAJA` 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 module 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 `N_Vectors` that utilize dual memory spaces, e.g. the native SUNDIALS CUDA `N_Vector`.

The `SUNMATRIX_CUSPARSE` and `SUNLINEARSOLVER_CUSOLVERS_BATCHQR` implementations no longer require the SUNDIALS 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.

1.1.19 Changes in v5.5.0

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.

1.1.20 Changes in v5.4.0

Added the function `IDASetLSNormFactor()` to specify the factor for converting between integrator tolerances (WRMS norm) and linear solver tolerances (L2 norm) i.e., $\text{tol}_{L2} = \text{nrmfac} * \text{tol}_{\text{WRMS}}$.

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 may be retrieved by calling `IDAGetNumNonlinSolvIters()`, the cumulative number of failures with `IDAGetNumNonlinSolvConvFails()`, or both with `IDAGetNonlinSolvStats()`.

A new API, `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 `NVECTOR_CUDA` and `NVECTOR_RAJA` modules that accept a `SUNMemoryHelper` object. Refer to §4.8 and §10 for more information.

The `NVECTOR_RAJA` module has been updated to mirror the `NVECTOR_CUDA` module. Notably, the update adds managed memory support to the `NVECTOR_RAJA` module. Users of the module will need to update any calls to the `N_VMake_Raja()` function because that signature was changed. This module remains experimental and is subject to change from version to version.

The `NVECTOR_TRILINOS` module 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 v11.

1.1.21 Changes in v5.3.0

Fixed a bug in the iterative linear solver modules where an error is not returned if the `ATimes` function is `NULL` or, if preconditioning is enabled, the `PSolve` function is `NULL`.

Added a new function `IDAGetNonlinearSystemData()` which advanced users might find useful if providing a custom `SUNNonlinSolSysFn`.

Added the ability to control the CUDA kernel launch parameters for the `NVECTOR_CUDA` and `SUNMATRIX_CUSPARSE` modules. These modules remain experimental and are subject to change from version to version. In addition, the `NVECTOR_CUDA` 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 `SUNNONLINSOL_NEWTON` and `SUNNONLINSOL_FIXEDPOINT` modules, and the SUNDIALS iterative linear solver modules. SUNDIALS must be built with the CMake option `SUNDIALS_BUILD_WITH_MONITORING` to use these capabilities.

Added the optional function `IDASetJacTimesResFn()` to specify an alternative residual function for computing Jacobian-vector products with the internal difference quotient approximation.

1.1.22 Changes in v5.2.0

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`, `xlf2003`, or `xlf2003_r`.

Fixed a linkage bug affecting Windows users that stemmed from `dllimport/dllexport` attributes missing on some SUNDIALS API functions.

Added a new `SUNMatrix` implementation, `SUNMATRIX_CUSPARSE`, that interfaces to the sparse matrix implementation from the NVIDIA cuSPARSE library. In addition, the `SUNLINSOL_CUSOLVER_BATCHQR` linear solver has been updated to use this matrix, therefore, users of this module will need to update their code. These modules are still considered to be experimental, thus they are subject to breaking changes even in minor releases.

The function `IDASetLinearSolutionScaling()` was added to enable or disable the scaling applied to linear system solutions with matrix-based linear solvers to account for a lagged value of α in the linear system matrix $J = \frac{\partial F}{\partial y} + \alpha \frac{\partial F}{\partial \dot{y}}$. Scaling is enabled by default when using a matrix-based linear solver.

1.1.23 Changes in v5.1.0

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`.

Added a new build system option, `CUDA_ARCH`, that can be used to specify the CUDA architecture to compile for.

Added two utility functions, `FSUNDIALSfileOpen()` and `FSUNDIALSfileClose()` for creating/destroying file pointers that are useful when using the Fortran 2003 interfaces.

1.1.24 Changes in v5.0.0

1.1.24.1 Build system changes

- 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.*, `SUPERLU_DIST_LIBRARIES` when enabling `SuperLU_DIST`.
- Fixed a bug in the build system that prevented the `NVECTOR_PTHREADS` module from being built.

1.1.24.2 NVECTOR module changes

- 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. See §6.1.1 for more details.
- Two new `N_Vector` implementations, `NVECTOR_MANYVECTOR` and `NVECTOR_MPIMANYVECTOR`, have been created 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. This implementation is accompanied by additions to user documentation and SUNDIALS examples. See §6.17 and §6.18 for more details.
- 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 length of an `N_Vector`. The optional operations have been added to support the new `NVECTOR_MPIMANYVECTOR` implementation. The operation `N_VGetCommunicator()` must be implemented by subvectors that are combined to create an `NVECTOR_MPI_MANYVECTOR`, 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_VWsqSumLocal()`, `N_VWsqSumMaskLocal()`, `N_VInvTestLocal()`, `N_VConstrMaskLocal()`, and `N_VMinQuotientLocal()`. If an `N_Vector` implementation defines any of the local operations as `NULL`, then the `NVECTOR_MPIMANYVECTOR` will call standard `N_Vector` operations to complete the computation. See §6.2.4 for more details.
- An additional `N_Vector` implementation, `NVECTOR_MPIPLUSX`, has been created to support the MPI+X paradigm where X is a type of on-node parallelism (*e.g.*, OpenMP, CUDA). The implementation is accompanied by additions to user documentation and SUNDIALS examples. See §6.19 for more details.
- The `*_MPCuda` and `*_MPIRaja` functions have been removed from the `NVECTOR_CUDA` and `NVECTOR_RAJA` 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 `NVECTOR_MPIPLUSX` module coupled in conjunction with the `NVECTOR_CUDA` or `NVECTOR_RAJA` modules to replace the functionality. The necessary changes are minimal and should require few code modifications. See the programs in `examples/ida/mpicuda` and `examples/ida/mpiraja` for examples of how to use the `NVECTOR_MPIPLUSX` module with the `NVECTOR_CUDA` and `NVECTOR_RAJA` modules respectively.
- Fixed a memory leak in the `NVECTOR_PETSC` module clone function.

- Made performance improvements to the *NVECTOR_CUDA* module. Users who utilize a non-default stream should no longer see default stream synchronizations after memory transfers.
- Added a new constructor to the *NVECTOR_CUDA* module that allows a user to provide custom allocate and free functions for the vector data array and internal reduction buffer. See §6.10 for more details.
- Added new Fortran 2003 interfaces for most *N_Vector* modules. See §6 for more details on how to use the interfaces.
- Added three new *N_Vector* utility functions, *FN_VGetVecAtIndexVectorArray()*, *FN_VSetVecAtIndexVectorArray()*, and *FN_VNewVectorArray()*, for working with *N_Vector* arrays when using the Fortran 2003 interfaces. See §6.1.1 for more details.

1.1.24.3 SUNMatrix module changes

- Two new functions were added to aid in creating custom *SUNMatrix* objects. The constructor *SUNMat-NewEmpty()* 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. See §7.1 for more details.
- 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 custom *SUNMatrix* modules will need to at least update their code to set the corresponding ops structure member, *matvecsetup*, to NULL. See §7.1 for more details.
- The generic *SUNMatrix* API now defines error codes to be returned by *SUNMatrix* 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. See §7.9 for more details.
- Added new Fortran 2003 interfaces for most *SUNMatrix* modules. See §7 for more details on how to use the interfaces.

1.1.24.4 SUNLinearSolver module changes

- A new function was added to aid in creating custom *SUNLinearSolver* objects. The constructor *SUNLinSol-NewEmpty()* 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. See §8.1.8 for more details.
- The return type of the *SUNLinearSolver* API function *SUNLinSolLastFlag()* 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. See §8.15 for more details.

- Added a new `SUNLinearSolver` implementation, `SUNLinearSolver_cuSolverSp_batchQR`, which leverages the NVIDIA cuSOLVER sparse batched QR method for efficiently solving block diagonal linear systems on NVIDIA GPUs. See §8.17 for more details.
- Added three new accessor functions to the `SUNLINSOL_KLU` module, `SUNLinSol_KLUGetSymbolic()`, `SUNLinSol_KLUGetNumeric()`, and `SUNLinSol_KLUGetCommon()`, to provide user access to the underlying KLU solver structures. See §8.5 for more details.
- Added new Fortran 2003 interfaces for most `SUNLinearSolver` modules. See §8 for more details on how to use the interfaces.

1.1.24.5 SUNNonlinearSolver module changes

- 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. See §9.1.7 for more details.
- 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 see §9.
- Added a new `SUNNonlinearSolver` implementation, `SUNNONLINSOL_PETSC`, which interfaces to the PETSc SNES nonlinear solver API. See §9.5 for more details.
- Added new Fortran 2003 interfaces for most `SUNNonlinearSolver` modules. See §9 for more details on how to use the interfaces.

1.1.24.6 IDA changes

- A bug was fixed in the IDA linear solver interface where an incorrect Jacobian-vector product increment was used with iterative solvers other than `SUNLINSOL_SPGMR` and `SUNLINSOL_SPGMR`.
- Fixed a memory leak in FIDA when not using the default nonlinear solver.
- 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, IDA will remove at least one global reduction per time step.
- The IDALS interface has been updated to only zero the Jacobian matrix before calling a user-supplied Jacobian evaluation function when the attached linear solver has type `SUNLINEARSOLVER_DIRECT`.
- Added the new functions, `IDAGetCurrentCj()`, `IDAGetCurrentY()`, `IDAGetCurrentYp()`, `IDAComputeY()`, and `IDAComputeYp()` which may be useful to users who choose to provide their own nonlinear solver implementations.
- Added a Fortran 2003 interface to IDA. See §4.7 for more details.

1.1.25 Changes in v4.1.0

An additional `N_Vector` implementation was added for the TPETRA vector from the TRILINOS library to facilitate interoperability between SUNDIALS and TRILINOS. This implementation is accompanied by additions to user documentation and SUNDIALS examples.

A bug was fixed where a nonlinear solver object could be freed twice in some use cases.

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 the RAJA `N_Vector` is enabled).

The implementation header file `ida_impl.h` is no longer installed. This means users who are directly manipulating the `IDAMem` structure will need to update their code to use IDA's public API.

Python is no longer required to run `make test` and `make test_install`.

1.1.26 Changes in v4.0.2

Added information on how to contribute to SUNDIALS and a contributing agreement.

Moved definitions of DLS and SPILS backwards compatibility functions to a source file. The symbols are now included in the IDA library, `libsundials_ida`.

1.1.27 Changes in v4.0.1

No changes were made in this release.

1.1.28 Changes in v4.0.0

IDA's previous direct and iterative linear solver interfaces, `IDADLS` and `IDASPILS`, have been merged into a single unified linear solver interface, `IDALS`, to support any valid `SUNLinearSolver` module. This includes the "DIRECT" and "ITERATIVE" types as well as the new "MATRIX_ITERATIVE" type. Details regarding how `IDALS` utilizes linear solvers of each type as well as discussion regarding intended use cases for user-supplied `SUNLinearSolver` implementations are included in §8. All IDA example programs and the standalone linear solver examples have been updated to use the unified linear solver interface.

The unified interface for the new `IDALS` module is very similar to the previous `IDADLS` and `IDASPILS` interfaces. To minimize challenges in user migration to the new names, the previous C and Fortran routine names may still be used; these will be deprecated in future releases, so we recommend that users migrate to the new names soon. Additionally, we note that Fortran users, however, may need to enlarge their `iout` array of optional integer outputs, and update the indices that they query for certain linear-solver-related statistics.

The names of all constructor routines for SUNDIALS-provided `SUNLinearSolver` implementations have been updated to follow the naming convention `SUNLinSol_*` where `*` is the name of the linear solver. The new names are `SUNLinSol_Band()`, `SUNLinSol_Dense()`, `SUNLinSol_KLU()`, `SUNLinSol_LapackBand()`, `SUNLinSol_LapackDense()`, `SUNLinSol_PCG()`, `SUNLinSol_SPBCGS()`, `SUNLinSol_SPGMR()`, `SUNLinSol_SPTFQMR()`, and `SUNLinSol_SuperLUMT()`. Solver-specific "set" routine names have been similarly standardized. To minimize challenges in user migration to the new names, the previous routine names may still be used; these will be deprecated in future releases, so we recommend that users migrate to the new names soon. All IDA example programs and the standalone linear solver examples have been updated to use the new naming convention.

The `SUNBandMatrix` constructor has been simplified to remove the storage upper bandwidth argument.

SUNDIALS integrators have been updated to utilize generic nonlinear solver modules defined through the `SUNNonlinearSolver` API. This API will ease the addition of new nonlinear solver options and allow for external or user-supplied nonlinear solvers. The `SUNNonlinearSolver` API and SUNDIALS provided modules are described in §9

and follow the same object oriented design and implementation used by the `N_Vector`, `SUNMatrix`, and `SUNLinearSolver` modules. Currently two `SUNNonlinearSolver` implementations are provided, `SUNNONLINSOL_NEWTON` and `SUNNONLINSOL_FIXEDPOINT`. These replicate the previous integrator specific implementations of a Newton iteration and a fixed-point iteration (previously referred to as a functional iteration), respectively. Note the `SUNNONLINSOL_FIXEDPOINT` module can optionally utilize Anderson's method to accelerate convergence. Example programs using each of these nonlinear solver modules in a standalone manner have been added and all IDA example programs have been updated to use generic `SUNNonlinearSolver` modules.

By default IDA uses the `SUNNONLINSOL_NEWTON` module. Since IDA previously only used an internal implementation of a Newton iteration no changes are required to user programs and functions for setting the nonlinear solver options (e.g., `IDASetMaxNonlinIters()`) or getting nonlinear solver statistics (e.g., `IDAGetNumNonlinSolvIters()`) remain unchanged and internally call generic `SUNNonlinearSolver` functions as needed. While SUNDIALS includes a fixed-point nonlinear solver module, it is not currently supported in IDA. For details on attaching a user-supplied nonlinear solver to IDA see `:numref:IDA.Usage.CC`. Additionally, the example program `idaRoberts_dns.c` explicitly creates an attaches a `SUNNONLINSOL_NEWTON` object to demonstrate the process of creating and attaching a nonlinear solver module (note this is not necessary in general as IDA uses the `SUNNONLINSOL_NEWTON` module by default).

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 an `N_Vector` (see §6 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()`, and `N_VDotProdMulti()` and the vector array operations are `N_VLinearCombinationVectorArray()`, `N_VScaleVectorArray()`, `N_VConstVectorArray()`, `N_VWrmsNormVectorArray()`, `N_VWrmsNormMaskVectorArray()`, `N_VScaleAddMultiVectorArray()`, and `N_VLinearCombinationVectorArray()`.

If an `N_Vector` implementation defines any of these operations as `NULL`, then standard `N_Vector` operations will automatically be called as necessary to complete the computation.

Multiple updates to `NVECTOR_CUDA` were made:

- 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 namespace `suncudavec`.
- Changed the `N_VMake_Cuda()` function to take a host data pointer and a device data pointer instead of an `N_VectorContent_Cuda` object.
- Added the ability to set the `cudaStream_t` used for execution of the `NVECTOR_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 `NVECTOR_CUDA`.

Multiple changes to `NVECTOR_RAJA` 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 namespace `suncudavec`.

A new `N_Vector` implementation for leveraging OpenMP 4.5+ device offloading has been added, `NVECTOR_OPENMPDEV`. See §6.15 for more details.

1.1.29 Changes in v3.2.1

The changes in this minor release include the following:

- Fixed a bug in the *CUDA N_Vector* where the *N_VInvTest()* operation could write beyond the allocated vector data.
- Fixed library installation path for multiarch systems. This fix changes the default library installation path to `CMAKE_INSTALL_PREFIX/CMAKE_INSTALL_LIBDIR` from `CMAKE_INSTALL_PREFIX/lib`. Note `CMAKE_INSTALL_LIBDIR` is automatically set, but is available as a CMake option that can be modified.

1.1.30 Changes in v3.2.0

Fixed a problem with setting `sunindextype` which would occur with some compilers (e.g. `armclang`) that did not define `__STDC_VERSION__`.

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.

Changed the name of the RAJA *N_Vector* library to `libsundials_nveccudaraja.lib` from `libsundials_nvecraja.lib` to better reflect that we only support CUDA as a backend for RAJA currently.

Several changes were made to the build system:

- CMake 3.1.3 is now the minimum required CMake version.
- Deprecate 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.

1.1.31 Changes in v3.1.2

The changes in this minor release include the following:

- Updated the minimum required version of CMake to 2.8.12 and enabled using `rpath` by default to locate shared libraries on OSX.
- Fixed Windows specific problem where `sunindextype` was not correctly defined when using 64-bit integers for the SUNDIALS index type. On Windows `sunindextype` is now defined as the MSVC basic type `__int64`.
- Added sparse SUNMatrix “Reallocate” routine to allow specification of the nonzero storage.

- Updated the KLU SUNLinearSolver module to set constants for the two reinitialization types, and fixed a bug in the full reinitialization approach where the sparse SUNMatrix pointer would go out of scope on some architectures.
- Updated the `SUNMatScaleAdd()` and `SUNMatScaleAddI()` implementations in the sparse SUNMatrix module 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 $I + \gamma J$ manually (with zero entries if needed).
- Changed the LICENSE install path to `instdir/include/sundials`.

1.1.32 Changes in v3.1.1

The changes in this minor release include the following:

- Fixed a potential memory leak in the `SUNLINSOL_SPGMR` and `SUNLINSOL_SPCGMR` linear solvers: if “Initialize” was called multiple times then the solver memory was reallocated (without being freed).
- Updated KLU SUNLinearSolver module to use a typedef for the precision-specific solve function to be used (to avoid compiler warnings).
- Added missing typecasts for some (void*) pointers (again, to avoid compiler warnings).
- Bugfix in `sunmatrix_sparse.c` where we had used `int` instead of `sunindextype` in one location.
- Added missing `#include <stdio.h>` in `N_Vector` and `SUNMatrix` header files.
- Added missing prototype for `IDASpilsGetNumJTSetupEvals()`.
- Fixed an indexing bug in the CUDA `N_Vector` implementation of `N_VWrmsNormMask()` and revised the RAJA `N_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` module (e.g., iterative linear solvers).

In addition to the changes above, minor corrections were also made to the example programs, build system, and user documentation.

1.1.33 Changes in v3.1.0

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.

1.1.34 Changes in v3.0.0

All interfaces to matrix structures and linear solvers have been reworked, and all example programs have been updated. The goal of the redesign of these interfaces was to provide more encapsulation and to ease interfacing of custom linear solvers and interoperability with linear solver libraries. Specific changes include:

- Added generic `SUNMatrix` module with three provided implementations: dense, banded, and sparse. These replicate previous SUNDIALS DIs and SIs matrix structures in a single object-oriented API.
- Added example problems demonstrating use of generic `SUNMatrix` modules.

- Added generic `SUNLinearSolver` module with eleven provided implementations: SUNDIALS native dense, SUNDIALS native banded, LAPACK dense, LAPACK band, KLU, SuperLU_MT, SPGMR, SPBCGS, SPT-FQMR, SPFGMR, and PCG. These replicate previous SUNDIALS generic linear solvers in a single object-oriented API.
- Added example problems demonstrating use of generic `SUNLinearSolver` modules.
- Expanded package-provided direct linear solver (Dls) interfaces and scaled, preconditioned, iterative linear solver (Spils) interfaces to utilize generic `SUNMatrix` and `SUNLinearSolver` objects.
- Removed package-specific, linear solver-specific, solver modules (e.g. `CVDENSE`, `KINBAND`, `IDAKLU`, `ARKSPGMR`) since their functionality is entirely replicated by the generic Dls/Spils interfaces and `SUNLinearSolver` and `SUNMatrix` modules. The exception is `CVDIAG`, a diagonal approximate Jacobian solver available to `CVODE` and `CVODES`.
- Converted all SUNDIALS example problems and files to utilize the new generic `SUNMatrix` and `SUNLinearSolver` objects, along with updated Dls 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.

Two additional `N_Vector` implementations were added – one for CUDA and one for RAJA vectors. 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 or residual 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 web site, <https://software.llnl.gov/RAJA/>. These additions are accompanied by updates to various interface functions and to user documentation.

All indices for data structures were updated to a new `sunindextype` that can be configured to be a 32- or 64-bit integer data index type. `sunindextype` is defined to be `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 their sparse matrix interface that now uses the new `sunindextype`. This new flexible capability for index types includes interfaces to PETSc, hypre, SuperLU_MT, and KLU with either 32-bit or 64-bit capabilities depending how the user configures SUNDIALS.

To avoid potential namespace conflicts, the macros defining `boolean` type 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.

The build system was expanded to support many of the xSDK-compliant keys. 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. More information can be found at, <https://xsdk.info>.

Added functions `SUNDIALSGetVersion()` and `SUNDIALSGetVersionNumber()` to get SUNDIALS release version information at runtime.

In addition, numerous changes were made to the build system. These include the addition of separate `BLAS_ENABLE` and `BLAS_LIBRARIES` CMake variables, additional error checking during CMake configuration, minor bug fixes, and renaming CMake options to enable/disable examples for greater clarity and an added option to enable/disable Fortran 77 examples. These changes included changing `EXAMPLES_ENABLE` to `EXAMPLES_ENABLE_C`, changing `CXX_ENABLE` to `EXAMPLES_ENABLE_CXX`, changing `F90_ENABLE` to `EXAMPLES_ENABLE_F90`, and adding an `EXAMPLES_ENABLE_F77` option.

A bug fix was done to add a missing prototype for `IDASetMaxBacksICC()` in `ida.h`.

Corrections and additions were made to the examples, to installation-related files, and to the user documentation.

1.1.35 Changes in v2.9.0

Two additional `N_Vector` implementations were added – one for Hypre (parallel) `ParVector` vectors, and one for PETSc vectors. These additions are accompanied by additions to various interface functions and to user documentation.

Each `N_Vector` module now includes a function, `N_VGetVectorID()`, that returns the `N_Vector` module name.

An optional input function was added to set a maximum number of line-search backtracks in the initial condition calculation. Also, corrections were made to three Fortran interface functions.

For each linear solver, the various solver performance counters are now initialized to 0 in both the solver specification function and in 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.

A memory leak was fixed in the banded preconditioner interface. In addition, updates were done to return integers from linear solver and preconditioner “free” functions.

The Krylov linear solver Bi-CGstab was enhanced by removing a redundant dot product. Various additions and corrections were made to the interfaces to the sparse solvers KLU and SuperLU_MT, including support for CSR format when using KLU.

New examples were added for use of the OpenMP vector.

Minor corrections and additions were made to the IDA solver, to the Fortran interfaces, to the examples, to installation-related files, and to the user documentation.

1.1.36 Changes in v2.8.0

Two major additions were made to the linear system solvers that are available for use with the IDA solver. 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 IDA.

Otherwise, only relatively minor modifications were made to IDA:

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`.

In `IDALapackBand`, the line `smu = MIN(N-1, mu+m1)` was changed to `smu = mu + m1` to correct an illegal input error for `DGBTRF/DGBTRS`.

A minor bug was fixed regarding the testing of the input `tstop` on the first call to `IDASolve()`.

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 in various example programs.

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.

In all FIDA 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.

Two new `N_Vector` modules have been added for thread-parallel computing environments — one for OpenMP, denoted `NVECTOR_OPENMP`, and one for Pthreads, denoted `NVECTOR_PTHREADS`.

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.

1.1.37 Changes in v2.7.0

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` and `NewLintArray`, for `int` and `long int` arrays, respectively.

A large number of minor errors have been fixed. Among these are the following: After the solver memory is created, it is set to zero before being filled. To be consistent with IDAS, IDA uses the function `IDAGetDky` for optional output retrieval. In each linear solver interface function, the linear solver memory is freed on an error return, and the `**Free` function now includes a line setting to `NULL` the main memory pointer to the linear solver memory. A memory leak was fixed in two of the `IDASp***Free` functions. In the rootfinding functions `IDARcheck1` and `IDARcheck2`, 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. 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.

1.1.38 Changes in v2.6.0

Two new features were added in this release: (a) a new linear solver module, based on BLAS and LAPACK for both dense and banded matrices, and (b) option to specify which direction of zero-crossing is to be monitored while performing rootfinding.

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 already present family of scaled preconditioned iterative linear solvers, the direct solvers, including the 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; (c) a general streamlining of the band-block-diagonal preconditioner module distributed with the solver.

1.1.39 Changes in v2.5.0

The main changes in this release involve a rearrangement of the entire SUNDIALS source tree (see §3). 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>`). Additional changes were made to the build system: all exported header files are now installed in separate subdirectories of the installation *include* directory.

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).

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.

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.

1.1.40 Changes in v2.4.0

FIDA, a Fortran-C interface module, was added.

IDASPBCG and IDASPTFQMR modules have been added to interface with the Scaled Preconditioned Bi-CGstab (SPBCG) and Scaled Preconditioned Transpose-Free Quasi-Minimal Residual (SPTFQMR) linear solver modules, respectively (for details see :numref:IDA.Usage.CC). 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.

The rootfinding feature was added, whereby the roots of a set of given functions may be computed during the integration of the DAE system.

A user-callable routine was added to access the estimated local error vector.

The deallocation functions now take as arguments the address of the respective memory block pointer.

To reduce the possibility of conflicts, the names of all header files have been changed by adding unique prefixes (`ida_` 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.

1.1.41 Changes in v2.3.0

The user interface has been further refined. Several functions used for setting optional inputs were combined into a single one. 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.

1.1.42 Changes in v2.2.2

Minor corrections and improvements were made to the build system. A new chapter in the User Guide was added — with constants that appear in the user interface.

1.1.43 Changes in v2.2.1

The changes in this minor SUNDIALS release affect only the build system.

1.1.44 Changes in v2.2.0

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, IDA now provides a set of routines (with prefix `IDASet`) to change the default values for various quantities controlling the solver and a set of extraction routines (with prefix `IDAGet`) to extract statistics after return from the main solver routine. Similarly, each linear solver module provides its own set of `Set`- and `Get`-type routines. For more details see §5.3.12.

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.

Installation of IDA (and all of SUNDIALS) has been completely redesigned and is now based on configure scripts.

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 IDA 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 structure of the SUNDIALS suite of solvers (§3) and the software organization of the IDA solver (§3.1).
- Chapter §5 is the main usage document for IDA for C and C++ applications. It includes a complete description of the user interface for the integration of DAE initial value problems. This is followed by documentation for using IDA with Fortran applications and on GPU accelerated systems.
- 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 IDA, within the structure of SUNDIALS (Appendix §11), as well as a list of all the constants used for input to and output from IDA 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

IDA 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.

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, IDA provides a routine that computes consistent initial conditions from a user's initial guess [18]. 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 IDA 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, IDA 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 IDA is the variable-order, variable-coefficient BDF (Backward Differentiation Formula), in fixed-leading-coefficient form [13]. The method order ranges from 1 to 5, with the BDF of order q given by the multistep formula

$$\sum_{i=0}^q \alpha_{n,i} y_{n-i} = h_n \dot{y}_n, \quad (2.2)$$

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 (2.2) to the DAE system (2.1) results in a nonlinear algebraic system to be solved at each step:

$$G(y_n) \equiv F \left(t_n, y_n, h_n^{-1} \sum_{i=0}^q \alpha_{n,i} y_{n-i} \right) = 0. \quad (2.3)$$

In the process of controlling errors at various levels, IDA 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

$$W_i = \frac{1}{\text{rtol} \cdot |y_i| + \text{atol}_i}. \quad (2.4)$$

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 IDA solves (2.3) with a Newton iteration but IDA also allows for user-defined nonlinear solvers (see Chapter §9). Each Newton iteration requires the solution of a linear system of the form

$$J[y_{n(m+1)} - y_{n(m)}] = -G(y_{n(m)}), \quad (2.5)$$

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

$$J = \frac{\partial G}{\partial y} = \frac{\partial F}{\partial y} + \alpha \frac{\partial F}{\partial \dot{y}}, \quad (2.6)$$

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

For the solution of the linear systems within the Newton iteration, IDA 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 [50] and an interface to the oneMKL library [2],
- band direct solvers, including an internal implementation or an interface to BLAS/LAPACK,
- sparse direct solver interfaces to various libraries, including KLU [3, 24], SuperLU_MT [8, 26, 44], SuperLU_-Dist [7, 32, 45, 46], and cuSPARSE [6],
- 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 [16]. For the *spils* linear solvers with IDA, 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 α 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 IDA 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

$$S \|\delta_m\| < 0.33, \quad (2.7)$$

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 (2.7) 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 δ_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 (2.6) can be either supplied by the user or IDA 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 (2.4). 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 σ .

2.2.2 Local Error Test

During the course of integrating the system, IDA 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, IDA 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 IDA is

$$\max\{|C|, \bar{C}\}\|\Delta_n\| \leq 1. \quad (2.8)$$

A user option is available by which the algebraic components of the error vector are omitted from the test (2.8), if these have been so identified.

2.2.3 Step Size and Order Selection

In IDA, 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 (2.8) 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, IDA 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 IDA (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 IDA 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 (2.8) 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 η is adjusted so that $\eta_{\text{min_ef}} \leq \eta \leq \eta_{\text{low}}$ (by default $\eta_{\text{min_ef}} = 0.25$ and $\eta_{\text{low}} = 0.9$) before setting $h \leftarrow h' = \eta h$. If the local error test fails a second time, IDA uses $\eta = \eta_{\text{min_ef}}$, and on the third and subsequent failures it uses $q = 1$ and $\eta = \eta_{\text{min_ef}}$. 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, IDA 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 η is adjusted such that

- a. If $\eta_{\min_fx} < \eta < \eta_{\max_fx}$, set $\eta = 1$. The defaults are $\eta_{\min_fx} = 1$ and $\eta_{\max_fx} = 2$.
- b. If $\eta \geq \eta_{\max_fx}$, the step size growth is restricted to $\eta_{\max_fx} \leq \eta \leq \eta_{\max}$ with $\eta_{\max} = 2$ by default.
- c. If $\eta \leq \eta_{\min_fx}$, 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 [13] 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, IDA 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 IDA not to integrate past a given stopping point $t = t_{\text{stop}}$.

2.2.4 Inequality Constraints

IDA 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, IDA 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 §5.3.2 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), IDA 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 IDA, 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 [16] for an extensive study of preconditioners for reaction-transport systems).

Typical preconditioners used with IDA 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 α 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 IDA solver has been augmented to include a rootfinding feature. This means that, while integrating the Initial Value Problem (2.1), IDA 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 IDA. 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 [34]. In addition, each time g is computed, IDA 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 , IDA computes g at $t + \delta$ for a small increment δ , slightly further in the direction of integration, and if any $g_i(t + \delta) = 0$ also, IDA stops and reports an error. This way, each time IDA 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, IDA 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 α a weight parameter. On the first two passes through the loop, α is set to 1, making t_{mid} the secant method value. Thereafter, α 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, α is set to 1. If the two sides were the same, α 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 $\tau/2$.

Chapter 3

Code Organization

SUNDIALS consists of the solvers CVODE and ARKODE for ordinary differential equation (ODE) systems, IDA for differential-algebraic (DAE) systems, and KINSOL for nonlinear algebraic systems. In addition, SUNDIALS also includes variants of CVODE and IDA with sensitivity analysis capabilities (using either forward or adjoint methods), called CVODES and IDAS, respectively. The following is a list summarizes the basic functionality of each SUNDIALS package:

- CVODE, a solver for stiff and nonstiff ODE systems $\dot{y} = f(t, y)$ based on Adams and BDF methods;
- CVODES, a solver for stiff and nonstiff ODE systems with sensitivity analysis capabilities;
- ARKODE, a solver for stiff, nonstiff, mixed stiff-nonstiff, and multirate ODE systems $M(t) \dot{y} = f_1(t, y) + f_2(t, y)$ based on Runge-Kutta methods;
- IDA, a solver for differential-algebraic systems $F(t, y, \dot{y}) = 0$ based on BDF methods;
- IDAS, a solver for differential-algebraic systems with sensitivity analysis capabilities;
- KINSOL, a solver for nonlinear algebraic systems $F(u) = 0$.

The various packages in the suite share many common components and are organized as a family. Fig. 3.1 gives a high-level overview of solver packages, the shared vector, matrix, linear solver, and nonlinear solver interfaces (abstract base classes), and the corresponding class implementations provided with SUNDIALS. For classes that provide interfaces to third-party libraries (i.e., LAPACK, KLU, SuperLU_MT, SuperLU_DIST, *hypre*, PETSc, Trilinos, and Raja) users will need to download and compile those packages independently of SUNDIALS. The directory structure is shown in Fig. 3.2.

3.1 IDA organization

The IDA 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.

The overall organization of the IDA package is shown in Fig. 3.3. IDA utilizes generic linear and nonlinear solvers defined by the `SUNLinearSolver` (see §8) and `SUNNonlinearSolver` interfaces (see §9) respectively. As such, IDA 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.

IDA 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




Fig. 3.1: High-level diagram of the SUNDIALS suite.




Fig. 3.2: Directory structure of the SUNDIALS source tree.




Fig. 3.3: Overall structure diagram of the IDA package. Components specific to IDA begin with “IDA” (IDALS, IDANLS, and IDABBDPRE), all other items correspond to generic SUNDIALS vector, matrix, and solver interfaces.

Jacobian information, or they may be matrix-free. Since IDA can operate on any valid `SUNLinearSolver`, the set of linear solver modules available to IDA will expand as new `SUNLinearSolver` implementations are developed.

For users employing `SUNMATRIX_DENSE` or `SUNMATRIX_BAND` Jacobian matrices, IDA 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, IDA 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 [16, 20], together with the example and demonstration programs included with IDA, 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 IDA module to each of the four associated functions is fixed, thus allowing the central module to be completely independent of the linear system method.

IDA 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 IDA to solve a given problem is stored in `N_Vector` instances. There is no global data in the IDA 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 IDA 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

Using SUNDIALS

As discussed in §3, the six solvers packages (CVODE(S), IDA(S), ARKODE, KINSOL) that make up SUNDIALS are built upon common classes/modules for vectors, matrices, and algebraic solvers. In addition, the six packages all leverage some other common infrastructure, which we discuss in this section.

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_type`

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

The advantage of using the name `sunboolean_type` (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_type` are intended to have only the two values `SUNFALSE` (0) and `SUNTRUE` (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.

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.

The *SUNContext* class/type is defined in the header file `sundials/sundials_context.h` as

```
typedef struct _SUNContext *SUNContext
```

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);
```

(continues on next page)

(continued from previous page)

```

package_mem = CNodeCreate(..., 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* suncctx)

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

Parameters

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

Returns

SUNErrCode indicating success or failure.

SUNErrCode **SUNContext_ClearErrHandlers**(*SUNContext* suncctx)

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

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

Returns

SUNErrCode indicating success or failure.

SUNErrCode **SUNContext_GetProfiler**(*SUNContext* suncctx, *SUNProfiler* *profiler)

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

Parameters

- **suncctx** – 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* suncctx, *SUNProfiler* profiler)

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

Parameters

- **suncctx** – 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* suncctx, *SUNLogger* logger)

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

Parameters

- **suncctx** – 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 cvode_initialized[num_threads];
void* cvode_mem[num_threads];

// Create
for (int i = 0; i < num_threads; i++) {
    sunctxs[i] = SUNContext_Create(...);
    cvode_mem[i] = CVodeCreate(..., sunctxs[i]);
    cvode_initialized[i] = 0; // not yet initialized
    // set optional cvode inputs...
}

// Solve
#pragma omp parallel for
for (int i = 0; i < num_problems; i++) {
    int retval = 0;
    int tid = omp_get_thread_num();
    if (!cvode_initialized[tid]) {
        retval = CVodeInit(cvode_mem[tid], ...);
        cvode_initialized[tid] = 1;
    }
}
```

(continues on next page)

(continued from previous page)

```

} else {
    retval = CVodeReInit(cvode_mem[tid], ...);
}
CVode(cvode_mem[i], ...);
}

// Destroy
for (int i = 0; i < num_threads; i++) {
    // get optional cvode outputs...
    CVodeFree(&cvode_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* cvode_mem;
    SUNContext sunctx;

    sunctx = SUNContext_Create(...);
    cvode_mem = CVodeCreate(..., sunctx);
    retval = CVodeInit(cvode_mem, ...);

    // set optional cvode inputs...

    CVode(cvode_mem, ...);

    // get optional cvode outputs...

    CVodeFree(&cvode_mem);
    SUNContext_Free(&sunctx);
}

```

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)
{
    sunctx_ = std::make_unique<SUNContext>();
    SUNContext_Create(comm, sunctx_.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 *sunctx_.get();
}
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.
// 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 be disabled by calling *SUNContext_PopErrorHandler()* or *SUNContext_ClearErrHandlers()*. A SUNDIALS error handler function has the type

```
typedef int (*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, *SUNErrorCode* 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:

```
typedef struct SUNLogger_ *SUNLogger
```

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 [12] 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. an instance of the struct

```
typedef struct SUNProfiler_ SUNProfiler
```

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_mrimestep_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_fnvecopenmp_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 ([§6](#), [§7](#), [§8](#), and [§9](#) respectively). 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 [\[42\]](#), 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	<code>fsundials_core_mod</code>

continues on next page

Table 4.1 – continued from previous page

Class/Module	Fortran 2003 Module Name
ARKODE	farkode_mod
ARKODE::ARKSTEP	farkode_arkstep_mod
ARKODE::ERKSTEP	farkode_erkstep_mod
ARKODE::MRISTEP	farkode_mrimestep_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
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_spbcgs_mod
SUNLINSOL_SPTFQMR	fsunlinsol_sptfqmr_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 the `sunindextype` size is 64-bits.

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>
<code>double</code>	in, inout, out, return	<code>real(c_double)</code>
<code>int</code>	in, inout, out, return	<code>integer(c_int)</code>
<code>long</code>	in, inout, out, return	<code>integer(c_long)</code>
<code>sunbooleantype</code>	in, inout, out, return	<code>integer(c_int)</code>
<code>sunrealtype</code>	in, inout, out, return	<code>real(c_double)</code>
<code>sunindextype</code>	in, inout, out, return	<code>integer(c_long)</code>
<code>double*</code>	in, inout, out	<code>real(c_double), dimension(*)</code>
<code>double*</code>	return	<code>real(c_double), pointer, dimension(:)</code>
<code>int*</code>	in, inout, out	<code>real(c_int), dimension(*)</code>
<code>int*</code>	return	<code>real(c_int), pointer, dimension(:)</code>
<code>long*</code>	in, inout, out	<code>real(c_long), dimension(*)</code>
<code>long*</code>	return	<code>real(c_long), pointer, dimension(:)</code>
<code>sunrealtype*</code>	in, inout, out	<code>real(c_double), dimension(*)</code>
<code>sunrealtype*</code>	return	<code>real(c_double), pointer, dimension(:)</code>
<code>sunindextype*</code>	in, inout, out	<code>real(c_long), dimension(*)</code>
<code>sunindextype*</code>	return	<code>real(c_long), pointer, dimension(:)</code>
<code>sunrealtype[]</code>	in, inout, out	<code>real(c_double), dimension(*)</code>
<code>sunindextype[]</code>	in, inout, out	<code>integer(c_long), dimension(*)</code>
<code>N_Vector</code>	in, inout, out	<code>type(N_Vector)</code>
<code>N_Vector</code>	return	<code>type(N_Vector), pointer</code>
<code>SUNMatrix</code>	in, inout, out	<code>type(SUNMatrix)</code>
<code>SUNMatrix</code>	return	<code>type(SUNMatrix), pointer</code>
<code>SUNLinearSolver</code>	in, inout, out	<code>type(SUNLinearSolver)</code>
<code>SUNLinearSolver</code>	return	<code>type(SUNLinearSolver), pointer</code>
<code>SUNNonlinearSolver</code>	in, inout, out	<code>type(SUNNonlinearSolver)</code>
<code>SUNNonlinearSolver</code>	return	<code>type(SUNNonlinearSolver), pointer</code>
<code>FILE*</code>	in, inout, out, return	<code>type(c_ptr)</code>
<code>void*</code>	in, inout, out, return	<code>type(c_ptr)</code>
<code>T**</code>	in, inout, out, return	<code>type(c_ptr)</code>
<code>T***</code>	in, inout, out, return	<code>type(c_ptr)</code>
<code>T****</code>	in, inout, out, return	<code>type(c_ptr)</code>

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.
! Therefore, we cannot pass a c_null_ptr, rather we pass a dissociated 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()` that can be called for accessing a vector in a vector array. The example below demonstrates this:

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)
```

(continues on next page)

(continued from previous page)

```
! 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
```

(continues on next page)

(continued from previous page)

```

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. However, it has only been tested and confirmed to be working with GNU Fortran 4.9+ and Intel Fortran 18.0.1+.

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

```

(continues on next page)

(continued from previous page)

```

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)
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 (see Chapters §6, §7, §8, and §9) 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 [11], 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 [4], AMD ROCm/HIP [1], and Intel oneAPI [2]. 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
<code>NVECTOR_CUDA</code>	X			X	X
<code>NVECTOR_HIP</code>	X	X		X	X
<code>NVECTOR_SYCL</code>	X ³	X ³	X	X	X
<code>NVECTOR_RAJA</code>	X	X	X	X	X
<code>NVECTOR_KOKKOS</code>	X	X	X	X	X
<code>NVECTOR_OPENMPDEV</code>	X	X ²	X ²	X	

Table 4.4: List of SUNDIALS GPU-enabled `SUNMatrix` Modules

Module	CUDA	ROCm/HIP	oneAPI	Unmanaged Memory	UVM
<code>SUNMATRIX_CUSPARSE</code>	X			X	X
<code>SUNMATRIX_ONEMKLDENSE</code>	X ³	X ³	X	X	X
<code>SUNMATRIX_MAGMADENSE</code>	X	X		X	X
<code>SUNMATRIX_GINKGO</code>	X	X		X	X
<code>SUNMATRIX_KOKKOSDENSE</code>	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 ³	X ³	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 ¹	X ¹	X ¹	X ¹	X ¹
<i>SUNLINSOL_SPGFMR</i>	X ¹	X ¹	X ¹	X ¹	X ¹
<i>SUNLINSOL_SPTFQMR</i>	X ¹	X ¹	X ¹	X ¹	X ¹
<i>SUNLINSOL_SPBCGS</i>	X ¹	X ¹	X ¹	X ¹	X ¹
<i>SUNLINSOL_PCG</i>	X ¹	X ¹	X ¹	X ¹	X ¹

Table 4.6: List of SUNDIALS GPU-enabled SUNNonlinearSolver Modules

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

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 IDA for IVP Solution

This chapter is concerned with the use of IDA 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 IDA user-callable functions and user-supplied functions. The sample programs described in the companion document [41] 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 IDA package.

IDA 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 Access to library and header files

At this point, it is assumed that the installation of IDA, 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 IDA. IDA symbols are found in `libdir/libsundials_ida.lib`. Thus, in addition to linking to `libdir/libsundials_core.lib`, IDA users need to link to the IDA 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  
<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 IDA are located in the subdirectories `incdir/include/ida`. To use IDA the application needs to include the header file for IDA in addition to the SUNDIALS core header file:

```
#include <sundials/sundials_core.h> // Provides core SUNDIALS types
#include <ida/ida.h>                // IDA provides methods for DAEs
```

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 (2.3) (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 `idaFood-Web_kry_p` (see [41]), 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.

5.2 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 IDA object

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

8. Initialize IDA solver

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

9. Specify integration tolerances

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

- `IDASStolerances()` 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 §5.3.3 for general advice on selecting tolerances and §5.3.4 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 IDA 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 IDA from their default values. See [§5.3.10](#) for details.

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

Call `IDACalcIC()` to correct the initial values y_0 and yp_0 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 y_0 above) will contain $y(t)$, while the vector `ypret` (which can be the same as the vector yp_0 above) will contain $\dot{y}(t)$.

See `IDASolve()` for details.

18. **Get optional outputs**

Call `IDAGet***` functions to obtain optional output. See [§5.3.12](#) 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 IDA.
- Call `SUNContext_Free()` to free the SUNDIALS context.

20. **Finalize MPI, if used**

Call `MPI_Finalize` to terminate MPI.

5.3 User-callable functions

This section describes the IDA functions that are called by the user to setup and then solve an IVP. Some of these are required. However, starting with [§5.3.10](#), the functions listed involve optional inputs/outputs or restarting, and those paragraphs may be skipped for a casual use of IDA. In any case, refer to [§5.2](#) 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 [§5.3.10.1](#)).

5.3.1 IDA initialization and deallocation functions

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

The function `IDACreate` instantiates an IDA solver object.

Arguments:

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

Return value:

- void* pointer the IDA 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 IDA.

Arguments:

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

Return value:

- void

5.3.2 IDA 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 IDA 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 IDA 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, *IDAwtFn* 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 (2.4).

Arguments:

- *ida_mem* – pointer to the IDA solver object. *IDACreate()*
- *efun* – is the function which defines the *ewt* vector. For full details see *IDAwtFn*.

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.

5.3.3 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 IDA package, and the discussion of it in the IDA Examples document [41]. 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.

5.3.4 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 IDA, 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. IDA 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.

5.3.5 Linear solver interface functions

As previously explained, if the nonlinear solver requires the solution of linear systems of the form (2.5), 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 IDA only left preconditioning is supported. For the specification of a preconditioner, see the iterative linear solver sections in §5.3.10 and §5.4. A preconditioner matrix P must approximate the Jacobian J , at least crudely.

To attach a generic linear solver to IDA, 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 IDA via a call to `IDASetLinearSolver()`. The first argument passed to this function is the IDA 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 IDA 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 IDA, initializing the IDALS linear solver interface.

Arguments:

- `ida_mem` – pointer to the IDA solver object.
- `LS` – `SUNLinearSolver` object to use for solving linear systems of the form (2.5).
- `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).

Warning: The previous routines `IDADlsSetLinearSolver` and `IDASpilsSetLinearSolver` are now wrappers for this routine, and may still be used for backward-compatibility. However, these will be deprecated in future releases, so we recommend that users transition to the new routine name soon.

5.3.6 Nonlinear solver interface function

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

When changing the nonlinear solver in IDA, `IDASetNonlinearSolver()` must be called after `IDAInit()`. If any calls to `IDASolve()` have been made, then IDA 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 IDA memory pointer returned by `IDACreate()` and the second argument is the `SUNNonlinearSolver` object to use for solving the nonlinear system (2.3). A call to this function attaches the nonlinear solver to the main IDA 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 IDA.

Arguments:

- `ida_mem` – pointer to the IDA 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.

5.3.7 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 [18]). 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 \dot{y}_0 at time t_0 .

Arguments:

- `ida_mem` – pointer to the IDA 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()**.

5.3.8 Rootfinding initialization function

While solving the IVP, IDA 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 IDA 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 `IDA-RootFn` 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`.

5.3.9 IDA 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 IDA 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 IDA 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, *trret* 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.

5.3.10 Optional input functions

There are numerous optional input parameters that control the behavior of the IDA solver. IDA 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 IDA 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 IDA 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 IDA, the reader can skip to §5.4.

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.

5.3.10.1 Main solver optional input functions

Table 5.1: Optional inputs for IDA

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>	∞
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 IDA solver object.

Arguments:

- `ida_mem` – pointer to the IDA 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 IDA 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 ≤ 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 IDA 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 IDA 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 IDA 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 IDA solver object.
- `hin` – value of the initial step size to be attempted. Pass 0.0 to have IDA 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, IDA estimates the initial step as the solution of $\|h\dot{y}\|_{WRMS} = 1/2$, with an added restriction that $|h| \leq .001|t_{out} - t_0|$.

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 IDA 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 6.2.0.

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 IDA 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 ∞ .

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 IDA 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 IDA 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 IDA 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 [13] 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 IDA 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 IDA 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.

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.

5.3.10.2 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>IDALSetJacFn()</i>	DQ
Set parameter determining if a c_j change requires a linear solver setup call	<i>IDALSetDeltaCJLSetup()</i>	0.25
Enable or disable linear solution scaling	<i>IDALSetLinearSolutionScaling()</i>	on
Jacobian-times-vector function	<i>IDALSetJacTimes()</i>	NULL, DQ
Preconditioner functions	<i>IDALSetPreconditioner()</i>	NULL, NULL
Ratio between linear and nonlinear tolerances	<i>IDALSetEpsLin()</i>	0.05
Increment factor used in DQ Jv approx.	<i>IDALSetIncrementFactor()</i>	1.0
Jacobian-times-vector DQ Res function	<i>IDALSetJacTimesResFn()</i>	NULL
Newton linear solve tolerance conversion factor	<i>IDALSetLSNormFactor()</i>	vector length

The mathematical explanation of the linear solver methods available to IDA 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 a *SUNMATRIX_DENSE* or *SUNMATRIX_BAND* 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 *IDALSetJacFn()*. 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 *IDALSetUserData()*.

```
int IDALSetJacFn(void *ida_mem, IDALSJacFn jac)
```

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

Arguments:

- *ida_mem* – pointer to the IDA 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 *IDALSetLinearSolver()*. 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.

Warning: The previous routine `IDADlsSetJacFn` is now a wrapper for this routine, and may still be used for backward-compatibility. However, this will be deprecated in future releases, so we recommend that users transition to the new routine name soon.

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 α may not be current and a scaling factor is applied to the solution of the linear system to account for the lagged value of α . 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 α 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 ≥ 1 then the default value (0.25) is used.

Arguments:

- `ida_mem` – pointer to the IDA 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 6.2.0.

int `IDASetLinearSolutionScaling`(void *ida_mem, *sunbooleantype* onoff)

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

Arguments:

- `ida_mem` – pointer to the IDA 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 §5.4.6 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 `jtimes`, 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 IDA solver object.
- `jtsetup` – 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 `jtsetup`. If `NULL` is passed to `jtimes`, these defaults are used. A user may specify non-`NULL` `jtimes` and `NULL` `jtsetup` inputs. This function must be called after the IDALS linear solver interface has been initialized through a call to `IDASetLinearSolver()`.

Warning: The previous routine `IDASpilsSetJacTimes` is now a wrapper for this routine, and may still be used for backward-compatibility. However, this will be deprecated in future releases, so we recommend that users transition to the new routine name soon.

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 IDA 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 ≤ 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()`.

Warning: The previous routine `IDASpilsSetIncrementFactor()` is now a wrapper for this routine, and may still be used for backward-compatibility. However, this will be deprecated in future releases, so we recommend that users transition to the new routine name soon.

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 [28] 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 IDA 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 IDA 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 §5.4.7 and §5.4.8). 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 IDA 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()*.

Warning: The previous routine *IDASpilsSetPreconditioner* is now a wrapper for this routine, and may still be used for backward-compatibility. However, this will be deprecated in future releases, so we recommend that users transition to the new routine name soon.

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 ϵ 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 IDA solver object.
- `eplifac` – linear convergence safety factor ≥ 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.

Warning: The previous routine *IDASpilsSetEpsLin* is now a wrapper for this routine, and may still be used for backward-compatibility. However, this will be deprecated in future releases, so we recommend that users transition to the new routine name soon.

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., $\text{tol}_{L2} = \text{fac} * \text{tol}_{WRMS}$.

Arguments:

- `ida_mem` – pointer to the IDA 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 (IDA v5.0.0) the value of `nrmfac` was computed using `N_VDotProd()` i.e., the `nrmfac < 0` case.

5.3.10.3 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	<code>IDASetMaxNonlinIters()</code>	4
Maximum no. of convergence failures	<code>IDASetMaxConvFails()</code>	10
Coeff. in the nonlinear convergence test	<code>IDASetNonlinConvCoef()</code>	0.33
Residual function for nonlinear system evaluations	<code>IDASetNlsResFn()</code>	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 IDA 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 IDA 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 (2.7).

Arguments:

- `ida_mem` – pointer to the IDA 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 ≤ 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 IDA 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()*.

5.3.10.4 Initial condition calculation optional input functions

Table 5.4: Optional inputs for IDA 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 ^{2/3}

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 IDA 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 ≤ 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 IDA 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 IDA 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 IDA 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 IDA 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.

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 IDA 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 IDA 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 ≤ 0.0 .

Notes:

The default value is $(\text{unit roundoff})^{2/3}$.

5.3.10.5 Time step adaptivity optional input functions

Table 5.5: Optional inputs for IDA time step adaptivity

Optional input	Function name	Default
Fixed step size bounds η_{\min_fx} and η_{\max_fx}	<code>IDASetEtaFixedStepBounds()</code>	1.0 and 2.0
Maximum step size growth factor η_{\max}	<code>IDASetEtaMax()</code>	2.0
Minimum step size reduction factor η_{\min}	<code>IDASetEtaMin()</code>	0.5
Maximum step size reduction factor η_{low}	<code>IDASetEtaLow()</code>	0.9
Minimum step size reduction factor after an error test failure η_{\min_ef}	<code>IDASetEtaMinErrFail()</code>	0.25
Step size reduction factor after a nonlinear solver convergence failure η_{cf}	<code>IDASetEtaConvFail()</code>	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 η_{\min_fx} and η_{\max_fx} . If step size change factor η 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 IDA 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 6.2.0.

int **IDASetEtaMax**(void *ida_mem, *sunrealtype* eta_max)

The function `IDASetEtaMax` specifies the maximum step size growth factor η_{max} .

The default value is $\eta_{max} = 2$.

Arguments:

- `ida_mem` – pointer to the IDA 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 6.2.0.

int **IDASetEtaMin**(void *ida_mem, *sunrealtype* eta_min)

The function `IDASetEtaMin` specifies the minimum step size reduction factor η_{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 IDA 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 6.2.0.

int **IDASetEtaLow**(void *ida_mem, *sunrealtype* eta_low)

The function `IDASetEtaLow` specifies the maximum step size reduction factor η_{low} .

The default value is $\eta_{low} = 0.9$.

`eta_low` should satisfy $0 < \eta_{low} \leq 1$, otherwise the default value is used.

Arguments:

- `ida_mem` – pointer to the IDA 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 6.2.0.

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 ≤ 0 or ≥ 1 , the default value is used.

Arguments:

- `ida_mem` – pointer to the IDA 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 6.2.0.

int **IDASetEtaConvFail**(void *ida_mem, *sunrealtype* eta_cf)

The function `IDASetEtaConvFail` specifies the step size reduction factor η_{cf} after a nonlinear solver convergence failure.

The default value is $\eta_{\text{cf}} = 0.25$.

If `eta_cf` is ≤ 0 or ≥ 1 , the default value is used.

Arguments:

- `ida_mem` – pointer to the IDA 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 6.2.0.

5.3.10.6 Rootfinding optional input functions

Table 5.6: Optional inputs for IDA rootfinding

Optional input	Function name	Default
Direction of zero-crossing	<i>IDASetRootDirection()</i>	both
Disable rootfinding warnings	<i>IDASetNoInactiveRootWarn()</i>	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 IDA 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 IDA solver object.

Return value:

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

Notes:

IDA 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), IDA will issue a warning which can be disabled with this optional input function.

5.3.11 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 IDA.

```
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 IDA. 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 IDA 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 $0, 1, \dots, k_{\text{last}}$.
- `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_{last} .

5.3.12 Optional output functions

IDA provides an extensive list of functions that can be used to obtain solver performance information. Table 5.7 lists all optional output functions in IDA, 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 IDA 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 IDA, IDALS, and IDANLS

Optional output	Function name
Size of IDA real and integer workspace	<i>IDAGetWorkSpace()</i>
Cumulative number of internal steps	<i>IDAGetNumSteps()</i>
No. of calls to residual function	<i>IDAGetNumResEvals()</i>
No. of calls to linear solver setup function	<i>IDAGetNumLinSolvSetups()</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>IDAGetLastOrder()</i>
Order to be attempted on the next step	<i>IDAGetCurrentOrder()</i>
Actual initial step size used	<i>IDAGetActualInitStep()</i>
Step size used for the last step	<i>IDAGetLastStep()</i>
Step size to be attempted on the next step	<i>IDAGetCurrentStep()</i>
Current internal time reached by the solver	<i>IDAGetCurrentTime()</i>
Suggested factor for tolerance scaling	<i>IDAGetTolScaleFactor()</i>
Error weight vector for state variables	<i>IDAGetErrWeights()</i>
Estimated local errors	<i>IDAGetEstLocalErrors()</i>
All IDA integrator statistics	<i>IDAGetIntegratorStats()</i>
No. of nonlinear solver iterations	<i>IDAGetNumNonlinSolvIters()</i>
No. of nonlinear convergence failures	<i>IDAGetNumNonlinSolvConvFails()</i>
IDA nonlinear solver statistics	<i>IDAGetNonlinSolvStats()</i>
User data pointer	<i>IDAGetUserData()</i>
Array showing roots found	<i>IDAGetRootInfo()</i>
No. of calls to user root function	<i>IDAGetNumGEvals()</i>
Print all statistics	<i>IDAPrintAllStats()</i>
Name of constant associated with a return flag	<i>IDAGetReturnFlagName()</i>
Number of backtrack operations	<i>IDAGetNumBacktrackOps()</i>
Corrected initial conditions	<i>IDAGetConsistentIC()</i>
Stored Jacobian of the DAE residual function	<i>IDAGetJac()</i>
c_j value used in the Jacobian evaluation	<i>IDAGetJacCj()</i>
Time at which the Jacobian was evaluated	<i>IDAGetJacTime()</i>
Step number at which the Jacobian was evaluated	<i>IDAGetJacNumSteps()</i>
Size of real and integer workspace	<i>IDAGetLinWorkSpace()</i>
No. of Jacobian evaluations	<i>IDAGetNumJacEvals()</i>
No. of residual calls for finite diff. Jacobian-vector evals.	<i>IDAGetNumLinResEvals()</i>
No. of linear iterations	<i>IDAGetNumLinIters()</i>
No. of linear convergence failures	<i>IDAGetNumLinConvFails()</i>
No. of preconditioner evaluations	<i>IDAGetNumPrecEvals()</i>
No. of preconditioner solves	<i>IDAGetNumPrecSolves()</i>
No. of Jacobian-vector setup evaluations	<i>IDAGetNumJTSetupEvals()</i>
No. of Jacobian-vector product evaluations	<i>IDAGetNumJtimesEvals()</i>
Last return from a linear solver function	<i>IDAGetLastLinFlag()</i>
Name of constant associated with a return flag	<i>IDAGetLinReturnFlagName()</i>

5.3.12.1 Main solver optional output functions

IDA 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 IDA 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 IDA real and integer workspace sizes.

Arguments:

- `ida_mem` – pointer to the IDA solver object.
- `lenrw` – number of real values in the IDA workspace.
- `leniw` – number of integer values in the IDA 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 §5.3.8), 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: $\text{leniw} = \text{leniw} + N_i$;
- with `id` specified (see `IDASetId()`): $\text{leniw} = \text{leniw} + N_i$;

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 $\text{leniw} = 49$.

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 IDA solver object.
- `nsteps` – number of steps taken by IDA.

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 IDA 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 IDA 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 IDA 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 IDA 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 IDA 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 IDA 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 IDA solver object.
- `hlast` – step size taken on the last internal step by IDA, 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 **IDAGetCurrentStep**(void *ida_mem, *sunrealtype* *hcur)

The function `IDAGetCurrentStep` returns the integration step size to be attempted on the next internal step.

Arguments:

- `ida_mem` – pointer to the IDA 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 **IDAGetActualInitStep**(void *ida_mem, *sunrealtype* *hinused)

The function `IDAGetActualInitStep` returns the value of the integration step size used on the first step.

Arguments:

- `ida_mem` – pointer to the IDA 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 IDA to ensure that the step size is within the prescribed bounds ($h_0 \leq h_{max}$), or to meet the local error test.

int **IDAGetCurrentTime**(void *ida_mem, *sunrealtype* *tcur)

The function `IDAGetCurrentTime` returns the current internal time reached by the solver.

Arguments:

- `ida_mem` – pointer to the IDA 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 **IDAGetTolScaleFactor**(void *ida_mem, *sunrealtype* *tolsfac)

The function `IDAGetTolScaleFactor` 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 IDA 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 (2.4) (or by the user's `IDAErrFn`).

Arguments:

- `ida_mem` – pointer to the IDA 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 IDA 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 *qlast, int *qcur, *sunrealtype* *hinused, *sunrealtype* *hlast, *sunrealtype* *hcur, *sunrealtype* *tcur)

The function `IDAGetIntegratorStats` returns the IDA integrator stats in one function call.

Arguments:

- `ida_mem` – pointer to the IDA solver object.
- `nsteps` – cumulative number of steps taken by IDA.
- `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.
- `qlast` – method order used on the last internal step.
- `qcur` – 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 IDA 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 IDA 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 IDA nonlinear solver statistics as a group.

Arguments:

- `ida_mem` – pointer to the IDA 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 IDA 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 6.3.0.

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 IDA 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 6.2.0.

char ***IDAGetReturnFlagName**(long int flag)

The function `IDAGetReturnFlagName` returns the name of the IDA constant corresponding to `flag`.

Arguments:

- `flag` – the flag returned by a call to an IDA function

Return value:

- `char*` – the flag name string

5.3.12.2 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 IDA 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 IDA 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).

5.3.12.3 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 IDA solver object.
- `rootsfound` – array of length `nrtfn` with the indices of the user functions g_i found to have a root. For $i = 0, \dots, nrtfn - 1$, `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 IDA 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`.

5.3.12.4 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 IDA 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 IDA 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 IDA 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 IDA 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 IDA 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.

<p>Warning: The previous routines <code>IDADlsGetWorkspace</code> and <code>IDASpilsGetWorkspace</code> are now wrappers for this routine, and may still be used for backward-compatibility. However, these will be deprecated in future releases, so we recommend that users transition to the new routine name soon.</p>

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 IDA 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.

Warning: The previous routine `IDADlsGetNumJacEvals` is now a wrapper for this routine, and may still be used for backward-compatibility. However, this will be deprecated in future releases, so we recommend that users transition to the new routine name soon.

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 IDA 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.

Warning: The previous routines `IDADlsGetNumRhsEvals` and `IDASpilsGetNumRhsEvals` are now wrappers for this routine, and may still be used for backward-compatibility. However, these will be deprecated in future releases, so we recommend that users transition to the new routine name soon.

int **IDAGetNumLinIters**(void *ida_mem, long int *nliters)

The function `IDAGetNumLinIters` returns the cumulative number of linear iterations.

Arguments:

- `ida_mem` – pointer to the IDA 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.

Warning: The previous routine `IDASpilsGetNumLinIters` is now a wrapper for this routine, and may still be used for backward-compatibility. However, this will be deprecated in future releases, so we recommend that users transition to the new routine name soon.

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 IDA 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.

Warning: The previous routine `IDASpilsGetNumConvFails` is now a wrapper for this routine, and may still be used for backward-compatibility. However, this will be deprecated in future releases, so we recommend that users transition to the new routine name soon.

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 IDA 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.

Warning: The previous routine `IDASpilsGetNumPrecEvals` is now a wrapper for this routine, and may still be used for backward-compatibility. However, this will be deprecated in future releases, so we recommend that users transition to the new routine name soon.

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 IDA 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.

Warning: The previous routine `IDASpilsGetNumPrecSolves` is now a wrapper for this routine, and may still be used for backward-compatibility. However, this will be deprecated in future releases, so we recommend that users transition to the new routine name soon.

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 IDA solver object.
- `njtsetup` – the current number of calls to `jtsetup`.

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.

Warning: The previous routine `IDASpilsGetNumJTSetupEvals` is now a wrapper for this routine, and may still be used for backward-compatibility. However, this will be deprecated in future releases, so we recommend that users transition to the new routine name soon.

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 IDA 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.

Warning: The previous routine `IDASpilsGetNumJtimesEvals` is now a wrapper for this routine, and may still be used for backward-compatibility. However, this will be deprecated in future releases, so we recommend that users transition to the new routine name soon.

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 IDA 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.

Warning: The previous routines `IDADlsGetLastFlag` and `IDASpilsGetLastFlag` are now wrappers for this routine, and may still be used for backward-compatibility. However, these will be deprecated in future releases, so we recommend that users transition to the new routine name soon.

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 IDA function

Return value:

- `char*` – the flag name string or if $1 \leq \text{lsflag} \leq N$ (LU factorization failed), this function returns “NONE”.

Warning: The previous routines `IDADlsGetReturnFlagName` and `IDASpilsGetReturnFlagName` are now wrappers for this routine, and may still be used for backward-compatibility. However, these will be deprecated in future releases, so we recommend that users transition to the new routine name soon.

5.3.13 IDA reinitialization function

The function `IDAReInit()` reinitializes the main IDA 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 §5.3.5.

If there are changes to any optional inputs, make the appropriate `IDASet***` calls, as described in §5.3.10.1. 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_{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 IDA.

Arguments:

- `ida_mem` – pointer to the IDA 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 IDA solver object was not initialized through a previous call to `IDACreate()`.
- `IDA_NO_MALLOC` – Memory space for the IDA 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, `IDAREInit()` also sends an error message to the error handler function.

5.4 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.

5.4.1 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, IDA 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.)

5.4.2 Error weight function

```
typedef int (*IDAEwtFn)(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 *IDASectUserData()*.

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 IDA.

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.

5.4.3 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 IDA.

5.4.4 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 *IDASetLinearSolver()*), 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 (2.6).

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 (α in (2.6)).
- `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 α in (2.6).

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 (`tt`, `yy`, `yp`) 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 §5.3.12.1. 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 $N_{eq} \times N_{eq}$ 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*)-th element of the dense matrix `Jac` (with *i*, *j* = 0 . . . *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 . . . *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 $N_{eq} \times N_{eq}$ 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 ≤ m-n ≤ 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) = J_{m,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 $\text{Neq} \times \text{Neq}$ 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.

Warning: The previous function type `IDADlsJacFn` is identical to `IDALsJacFn`, and may still be used for backward-compatibility. However, this will be deprecated in future releases, so we recommend that users transition to the new function type name soon.

5.4.5 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 `IDASetLinearSolver()`), the user may provide a function of type `IDALsJacTimesVecFn` in the following, 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 (2.6).

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 (α in (2.6)).
- `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 §5.3.12.1. The unit roundoff can be accessed as `SUN_UNIT_ROUNDOFF` defined in `sundials_types.h`.

Warning: The previous function type `IDASpilsJacTimesVecFn` is identical to `IDALsJacTimesVecFn`, and may still be used for backward-compatibility. However, this will be deprecated in future releases, so we recommend that users transition to the new function type name soon.

5.4.6 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, ealtype 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 (α in (2.6)).
- `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 §5.3.12.1. The unit roundoff can be accessed as `SUN_UNIT_ROUNDOFF` defined in `sundials_types.h`.

Warning: The previous function type `IDASpilsJacTimesSetupFn` is identical to `IDALsJacTimesSetupFn`, and may still be used for backward-compatibility. However, this will be deprecated in future releases, so we recommend that users transition to the new function type name soon.

5.4.7 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 $J = \frac{\partial F}{\partial y} + cj \frac{\partial F}{\partial \dot{y}}$. 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 (α in (2.6)).
- `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).

5.4.8 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 (α in (2.6)).

- `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 §5.3.12.1. The unit roundoff can be accessed as `SUN_UNIT_ROUNDOFF` defined in `sundials_types.h`.

5.5 Preconditioner modules

A principal reason for using a parallel DAE solver such as IDA 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 (2.5) 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 [37] and is included in a software module within the IDA 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.

5.5.1 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 `IDASetUserData()` 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), 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 §5.2 are not bold.

1. Initialize parallel or multi-threaded environment (*if appropriate*)
2. Create the vector of initial values
3. Create matrix object (*if appropriate*)
4. **Create linear solver object** (*if appropriate*)

When creating the iterative linear solver object, specify the use of left preconditioning (`SUN_PREC_LEFT`) as IDA only supports left preconditioning.

5. Create nonlinear solver object (*if appropriate*)
6. Create IDA object
7. Initialize IDA solver
8. Specify integration tolerances
9. Attach the linear solver (*if appropriate*)
10. **Set linear solver optional inputs** (*if appropriate*)

Note that 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 (*if appropriate*)
13. Set nonlinear solver optional inputs (*if appropriate*)
14. Specify rootfinding problem (*optional*)
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. Deallocate memory
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, sunindex_t Nlocal, sunindex_t mudq, sunindex_t mldq, sunindex_t
    mukeep, sunindex_t mlkeep, sunreal_t 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 IDA 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 `mukkeep` 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`, `mukkeep`, or `mlkeep`. After solving one problem, and after calling `IDAReInit()` to re-initialize IDA 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 IDA 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 IDABBDPrecGetWorkSpace(void *ida_mem, long int *lenrwBBDP, long int *leniwBBDP)`

The function `IDABBDPrecGetWorkSpace` returns the local sizes of the IDABBDPRE real and integer workspaces.

Arguments:

- `ida_mem` – pointer to the IDA 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 `IDABBDPrecGetNumGfnEvals`(void *ida_mem, long int *ngevalsBBDP)

The function `IDABBDPrecGetNumGfnEvals` 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 IDA 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 IDA output (see §5.3.12.1), and `npsolves` and `nrevalsLS` are linear solver optional outputs (see §5.3.12.4).

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

The 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 NVECTOR implementation. Users can provide a custom implementation of the NVECTOR module or use one provided within SUNDIALS. The generic operations are described below. In the sections following, the implementations provided with SUNDIALS are described.

The generic `N_Vector` type is a pointer to a structure that has an implementation-dependent *content* field containing the description and actual data of the vector, and an *ops* field pointing to a structure with generic vector operations. The type `N_Vector` is defined as

```
typedef struct _generic_N_Vector *N_Vector
```

and the generic structure is defined as

```
struct _generic_N_Vector {
    void *content;
    struct _generic_N_Vector_Ops *ops;
};
```

Here, the `_generic_N_Vector_Op` structure is essentially a list of function pointers to the various actual vector operations, and is defined as

```
struct _generic_N_Vector_Ops {
    N_Vector_ID (*nvgetvectorid)(N_Vector);
```

(continues on next page)

(continued from previous page)

```

N_Vector (*nvclone)(N_Vector);
N_Vector (*nvcloneempty)(N_Vector);
void (*nvdestroy)(N_Vector);
void (*nvspace)(N_Vector, sunindextype*, sunindextype*);
sunrealtype* (*nvgetarraypointer)(N_Vector);
sunrealtype* (*nvgetdevicearraypointer)(N_Vector);
void (*nvsetarraypointer)(sunrealtype*, N_Vector);
SUNComm (*nvgetcommunicator)(N_Vector);
sunindextype (*nvgetlength)(N_Vector);
sunindextype (*nvgetlocallength)(N_Vector);
void (*nvlinearsum)(sunrealtype, N_Vector, sunrealtype, N_Vector, N_Vector);
void (*nvconst)(sunrealtype, N_Vector);
void (*nvprod)(N_Vector, N_Vector, N_Vector);
void (*nvdiv)(N_Vector, N_Vector, N_Vector);
void (*nvscale)(sunrealtype, N_Vector, N_Vector);
void (*nvabs)(N_Vector, N_Vector);
void (*nvinv)(N_Vector, N_Vector);
void (*nvaddconst)(N_Vector, sunrealtype, N_Vector);
sunrealtype (*nvdotprod)(N_Vector, N_Vector);
sunrealtype (*nvmaxnorm)(N_Vector);
sunrealtype (*nvwrmsnorm)(N_Vector, N_Vector);
sunrealtype (*nvwrmsnormmask)(N_Vector, N_Vector, N_Vector);
sunrealtype (*nvmin)(N_Vector);
sunrealtype (*nvwl2norm)(N_Vector, N_Vector);
sunrealtype (*nvlnorm)(N_Vector);
void (*nvcompare)(sunrealtype, N_Vector, N_Vector);
sunboolean (*nvinvtest)(N_Vector, N_Vector);
sunboolean (*nvconstrmask)(N_Vector, N_Vector, N_Vector);
sunrealtype (*nvminquotient)(N_Vector, N_Vector);
SUNErrCode (*nvlinearcombination)(int, sunrealtype*, N_Vector*, N_Vector);
SUNErrCode (*nvscaleaddmulti)(int, sunrealtype*, N_Vector, N_Vector*,
                               N_Vector*);
SUNErrCode (*nvdotprodmulti)(int, N_Vector, N_Vector*, sunrealtype*);
SUNErrCode (*nvlinearsumvectorarray)(int, sunrealtype, N_Vector*, sunrealtype,
                                     N_Vector*, N_Vector*);
SUNErrCode (*nvscalevectorarray)(int, sunrealtype*, N_Vector*, N_Vector*);
SUNErrCode (*nvconstvectorarray)(int, sunrealtype, N_Vector*);
SUNErrCode (*nvwrmsnormvectorarray)(int, N_Vector*, N_Vector*, sunrealtype*);
SUNErrCode (*nvwrmsnormmaskvectorarray)(int, N_Vector*, N_Vector*, N_Vector,
                                         sunrealtype*);
SUNErrCode (*nvscaleaddmultivectorarray)(int, int, sunrealtype*, N_Vector*,
                                         N_Vector**, N_Vector**);
SUNErrCode (*nvlinearcombinationvectorarray)(int, int, sunrealtype*,
                                             N_Vector**, N_Vector*);
sunrealtype (*nvdotprodlocal)(N_Vector, N_Vector);
sunrealtype (*nvmaxnormlocal)(N_Vector);
sunrealtype (*nvminlocal)(N_Vector);
sunrealtype (*nvlnormlocal)(N_Vector);
sunboolean (*nvinvtestlocal)(N_Vector, N_Vector);
sunboolean (*nvconstrmasklocal)(N_Vector, N_Vector, N_Vector);
sunrealtype (*nvminquotientlocal)(N_Vector, N_Vector);
sunrealtype (*nvwsqrsumlocal)(N_Vector, N_Vector);

```

(continues on next page)

(continued from previous page)

```

sunrealtype (*nvwqrsummasklocal)(N_Vector, N_Vector, N_Vector);
SUNErrCode (*nvdotprodmultilocal)(int, N_Vector, N_Vector*, sunrealtype*);
SUNErrCode (*nvdotprodmultiallreduce)(int, N_Vector, sunrealtype*);
SUNErrCode (*nvbufsize)(N_Vector, sunindextype*);
SUNErrCode (*nvbufpack)(N_Vector, void*);
SUNErrCode (*nvbufunpack)(N_Vector, void*);
void (*nvprint)(N_Vector);
void (*nvprintfile)(N_Vector, FILE*);
};

```

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()`

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`.

Each NVECTOR 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 ID **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 accessible 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* y)

Returns the value of the dot-product of the *N_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_∞ 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:

$$z_{j,i} = cx_{j,i}, \quad i = 0, \dots, n-1 \quad j = 0, \dots, nv-1,$$

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)
```

6.3 NVECTOR functions used by IDA

In Table 6.2 below, we list the vector functions used in the *N_Vector* module used by the IDA package. The table also shows, for each function, which of the code modules uses the function. The IDA 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 IDA user does not need to know anything about the usage of vector functions by the IDA code modules in order to use IDA. The information is presented as an implementation detail for the interested reader.

Table 6.2: List of vector functions usage by IDA code modules

	IDA	IDALS	IDABBDPRE
<i>N_VGetVectorID()</i>			
<i>N_VGetLength()</i>		4	
<i>N_VClone()</i>	x	x	x
<i>N_VCloneEmpty()</i>		1	
<i>N_VDestroy()</i>	x	x	x
<i>N_VSpace()</i>	x	2	
<i>N_VGetArrayPointer()</i>		1	x
<i>N_VSetArrayPointer()</i>		1	
<i>N_VLinearSum()</i>	x	x	
<i>N_VConst()</i>	x	x	
<i>N_VProd()</i>	x		
<i>N_VDiv()</i>	x		
<i>N_VScale()</i>	x	x	x
<i>N_VAbs()</i>	x		
<i>N_VInv()</i>	x		
<i>N_VAddConst()</i>	x		
<i>N_VMaxNorm()</i>	x		
<i>N_VWrmsNorm()</i>	x		
<i>N_VMin()</i>	x		
<i>N_VMinQuotient()</i>	x		
<i>N_VConstrMask()</i>	x		
<i>N_VWrmsNormMask()</i>	x		
<i>N_VCompare()</i>	x		
<i>N_VLinearCombination()</i>	x		
<i>N_VScaleAddMulti()</i>	x		
<i>N_VDotProdMulti()</i>		3	
<i>N_VLinearSumVectorArray()</i>	x		
<i>N_VScaleVectorArray()</i>	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 IDA 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 IDA could omit these operations.
4. This routine is only used when an iterative or matrix iterative `SUNLinearSolver` module is supplied to IDA.

Of the functions listed in §6.2, `N_VWL2Norm()`, `N_VL1Norm()`, `N_VInvTest()`, and `N_VGetCommunicator()` are *not* used by IDA. Therefore a user-supplied `N_Vector` module for IDA 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 = NV_GetArrayPointer(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()`, `N_VMake_Serial()` set the field `own_data` to `SUNFALSE`. The function `N_VDestroy_Serial()` 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, and 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 routine `N_VDestroy_Parallel()` 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_nvopenmp.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, *sunrealtyp*e *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, *sunbooleantyp*e 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, *sunbooleantyp*e 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, *sunbooleantyp*e 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, *sunbooleantyp*e 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, *sunbooleantyp*e 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, *sunbooleantyp*e 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, *sunboolean* 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, *sunboolean* 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, *sunboolean* 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, *sunboolean* 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, *sunboolean* 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 function `N_VDestroy_OpenMP()` 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 out file.

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 function `N_VDestroy_Pthreads()` 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 fnvector_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;
    sunbooleantype own_data;
    sunbooleantype 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.

SUNErrorCode **N_VEnableFusedOps_ParHyp**(*N_Vector* v, *sunboolean* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) all fused and vector array operations in the parhyp vector. The return value is a *SUNErrorCode*.

SUNErrorCode **N_VEnableLinearCombination_ParHyp**(*N_Vector* v, *sunboolean* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear combination fused operation in the parhyp vector. The return value is a *SUNErrorCode*.

SUNErrorCode **N_VEnableScaleAddMulti_ParHyp**(*N_Vector* v, *sunboolean* 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 *SUNErrorCode*.

SUNErrorCode **N_VEnableDotProdMulti_ParHyp**(*N_Vector* v, *sunboolean* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the multiple dot products fused operation in the parhyp vector. The return value is a *SUNErrorCode*.

SUNErrorCode **N_VEnableLinearSumVectorArray_ParHyp**(*N_Vector* v, *sunboolean* tf)

This function enables (SUNTRUE) or disables (SUNFALSE) the linear sum operation for vector arrays in the parhyp vector. The return value is a *SUNErrorCode*.

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 function `N_VDestroy_ParHyp()` 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, *sunbooleantype* 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 routine *N_VDestroy_Petsc()* 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;
```

(continues on next page)

(continued from previous page)

```

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 [38].

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, *sunrealttype* *h_vdata, *sunrealttype* *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, *sunrealttype* *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

sunrealttype ***N_VCopyToDevice_Cuda**(*N_Vector* v)

This function copies host vector data to the device.

sunrealttype ***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, *sunbooleantype* 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, *sunbooleantype* 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, *sunbooleantype* 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, *sunbooleantype* 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, *sunbooleantype* 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 ExecPolicy
{
public:
    ExecPolicy(cudaStream_t stream = 0) : stream_(stream) { }
    virtual size_t gridSize(size_t numWorkUnits = 0, size_t blockDim = 0) const = 0;
    virtual size_t blockSize(size_t numWorkUnits = 0, size_t gridDim = 0) const = 0;
    virtual const cudaStream_t* stream() const { return (&stream_); }
    virtual ExecPolicy* clone() const = 0;
    ExecPolicy* clone_new_stream(cudaStream_t stream) const {
        ExecPolicy* ex = clone();
        ex->stream_ = stream;
        return ex;
    }
    virtual bool atomic() const { return false; }
    virtual ~ExecPolicy() {}
protected:
    cudaStream_t stream_;
};
```

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
    {
```

(continues on next page)

(continued from previous page)

```

    /* 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 [1]. 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;
    sunbooleantype   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, *sunrealttype* *h_vdata, *sunrealttype* *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, *sunrealttype* *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*

sunrealttype ***N_VCopyToDevice_Hip**(*N_Vector* v)

This function copies host vector data to the device.

sunrealttype ***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 ExecPolicy
{
public:
    ExecPolicy(hipStream_t stream = 0) : stream_(stream) { }
    virtual size_t gridSize(size_t numWorkUnits = 0, size_t blockDim = 0) const = 0;
    virtual size_t blockSize(size_t numWorkUnits = 0, size_t gridDim = 0) const = 0;
    virtual const hipStream_t* stream() const { return (&stream_); }
    virtual ExecPolicy* clone() const = 0;
    ExecPolicy* clone_new_stream(hipStream_t stream) const {
        ExecPolicy* ex = clone();
        ex->stream_ = stream;
        return ex;
    }
    virtual bool atomic() const { return false; }
    virtual ~ExecPolicy() {}
protected:
    hipStream_t stream_;
};
```

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
    {
```

(continues on next page)

(continued from previous page)

```

    /* 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;
    sunbooleantype   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, sunrealttype *h_vdata, sunrealttype *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, sunrealttype *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, sunbooleanstype 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.

sunrealttype ***`N_VGetHostArrayPointer_Sycl`**(`N_Vector` v)

This function returns a pointer to the vector host data array.

sunrealttype ***`N_VGetDeviceArrayPointer_Sycl`**(`N_Vector` v)

This function returns a pointer to the vector device data array.

void **`N_VSetHostArrayPointer_Sycl`**(sunrealttype *h_vdata, `N_Vector` v)

This function sets the host array pointer in the vector v.

void **`N_VSetDeviceArrayPointer_Sycl`**(sunrealttype *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.

sunbooleanstype **`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, sunbooleantype 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, sunbooleantype 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 ExecPolicy
{
public:
    virtual size_t gridSize(size_t numWorkUnits = 0, size_t blockDim = 0) const = 0;
    virtual size_t blockSize(size_t numWorkUnits = 0, size_t gridDim = 0) const = 0;
    virtual ExecPolicy* clone() const = 0;
    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)
```

(continues on next page)

(continued from previous page)

```

{}

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);

```

(continues on next page)

(continued from previous page)

```
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_raj.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 [38].

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 [29, 52] 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>
```

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;
    sunbooleanype 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, *sunrealttype* *h_vdata, *sunrealttype* *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.

sunrealttype ***N_VGetHostArrayPointer_OpenMPDEV**(*N_Vector* v)

This function returns a pointer to the host data array.

sunrealttype ***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, *sunboolean*type 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, *sunboolean*type 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, *sunboolean*type 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, *sunboolean*type 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, *sunboolean*type 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, *sunboolean*type 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, *sunboolean*type 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, *sunboolean*type 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, *sunboolean*type 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 function `N_VDestroy_OpenMPDEV()` 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 `sunindex_t` 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 {
    sunindex_t num_subvectors; /* number of vectors attached */
    sunindex_t 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_VLinearCombinationVectorArray()*. 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`. `N_VDestroy_ManyVector()` 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_MPI_MANYVECTOR. Three potential use cases for this module include:

- A. *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`.
- B. *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.
- C. *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_MANYVECTOR 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 */
}
```

(continues on next page)

(continued from previous page)

```

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_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 `MPIManyVector` 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 `MPIManyVector` 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 MPI-ManyVector 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 MPI-ManyVector vector. The return value is a *SUNErrCode*.

Notes

- *N_VNew_MPIManyVector()* and *N_VMake_MPIManyVector()* set the field `own_data = SUNFALSE`. *N_VDestroy_MPIManyVector()* 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 `N_VDestroy_MPIPlusX()` 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. The type `SUNMatrix` is defined as:

```
typedef struct _generic_SUNMatrix *SUNMatrix
```

and the generic structure is defined as

```
struct _generic_SUNMatrix {  
    void *content;  
    struct _generic_SUNMatrix_Ops *ops;  
};
```

Here, the `_generic_SUNMatrix_Ops` structure is essentially a list of function pointers to the various actual matrix operations, and is defined as

```
struct _generic_SUNMatrix_Ops {  
    SUNMatrix_ID (*getid)(SUNMatrix);  
    SUNMatrix (*clone)(SUNMatrix);
```

(continues on next page)

```

void (*destroy)(SUNMatrix);
SUNErrCode (*zero)(SUNMatrix);
SUNErrCode (*copy)(SUNMatrix, SUNMatrix);
SUNErrCode (*scaleadd)(sunrealtype, SUNMatrix, SUNMatrix);
SUNErrCode (*scaleaddi)(sunrealtype, SUNMatrix);
SUNErrCode (*matvecsetup)(SUNMatrix);
SUNErrCode (*matvec)(SUNMatrix, N_Vector, N_Vector);
SUNErrCode (*space)(SUNMatrix, long int*, long int*);
};

```

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()`

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

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
<i>SUNMATRIX_BAND</i>	Band $M \times M$ matrix
<i>SUNMATRIX_CUSPARSE</i>	CUDA sparse CSR matrix
<i>SUNMATRIX_CUSTOM</i>	User-provided custom matrix
<i>SUNMATRIX_DENSE</i>	Dense $M \times N$ matrix
<i>SUNMATRIX_GINKGO</i>	<i>SUNMatrix</i> wrapper for Ginkgo matrices
<i>SUNMATRIX_MAGMADENSE</i>	Dense $M \times N$ matrix
<i>SUNMATRIX_ONEMKLDENSE</i>	oneMKL dense $M \times N$ matrix
<i>SUNMATRIX_SLUNRLOC</i>	<i>SUNMatrix</i> wrapper for SuperLU_DIST SuperMatrix
<i>SUNMATRIX_SPARSE</i>	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 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 $A_data = \text{SUNDenseMatrix_Data}(A)$, or equivalently $A_data = \text{SM_DATA_D}(A)$, and then access $A_data[i]$ within the loop.
 - First obtain the array of column pointers via $A_cols = \text{SUNDenseMatrix_Cols}(A)$, or equivalently $A_cols = \text{SM_COLS_D}(A)$, and then access $A_cols[j][i]$ within the loop.
 - Within a loop over the columns, access the column pointer via $A_colj = \text{SUNDenseMatrix_Column}(A, j)$ 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 $\text{SM_ELEMENT_D}(A, i, j)$ within a double loop.

- Within the $\text{SUNMatMatvec_Dense}$ 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 $\text{SUNMATRIX_MAGMADENSE}$ module interfaces to the [MAGMA](#) linear algebra library and can target NVIDIA's CUDA programming model or AMD's HIP programming model [50]. 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

$$A = \begin{bmatrix} A_0 & 0 & \cdots & 0 \\ 0 & A_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_{n-1} \end{bmatrix}$$

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 $\text{SUNMATRIX_MAGMADENSE}$ module is experimental and subject to change.

7.4.1 SUNMATRIX_MAGMADENSE Functions

The $\text{SUNMATRIX_MAGMADENSE}$ module defines GPU-enabled implementations of all matrix operations listed in §7.2.

- $\text{SUNMatGetID_MagmaDense}$ – returns $\text{SUNMATRIX_MAGMADENSE}$
- $\text{SUNMatClone_MagmaDense}$
- $\text{SUNMatDestroy_MagmaDense}$
- $\text{SUNMatZero_MagmaDense}$
- $\text{SUNMatCopy_MagmaDense}$
- $\text{SUNMatScaleAdd_MagmaDense}$
- $\text{SUNMatScaleAddI_MagmaDense}$

- `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 $M_{\text{block}} \times \text{nblocks}$.

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 $N_{\text{block}} \times \text{nblocks}$.

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 $M_{\text{block}} \times \text{nblocks}$.

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 $N_{\text{block}} \times \text{nblocks}$.

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 $nblocks * N_{block}$.

7.5.1.5 Access Matrix Block Data

sunindextype `SUNMatrix_OneMklDense_BlockLData(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_{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;
    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, $0 \leq \mu < N$
- `ml` - lower half-bandwidth, $0 \leq ml < N$
- `smu` - storage upper bandwidth, $\mu \leq smu < N$. 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 $\min(N-1, \mu+ml)$ because of partial pivoting. The `smu` field holds the upper half-bandwidth allocated for the band matrix.
- `ldim` - leading dimension ($ldim \geq smu + ml + 1$)
- `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 (= `ldim 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 $j - \mu \leq i \leq j + \mu$.

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:

```
#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:




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 ml , 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_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_SUBAND_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 $j - \text{mu} \leq i \leq j + \text{ml}$.

Implementation:

```
#define SM_ELEMENT_B(A,i,j)  ( (SM_CONTENT_B(A)->cols)[j][(i)-(j)+SM_SUBAND_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 $j - \text{mu} \leq i \leq j + \text{ml}$.

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* sunctx)

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 $\min(N-1, \text{mu}+\text{ml})$ if the matrix will be used by the `SUNLinSol_Band` module;
- exactly equal to $\text{mu}+\text{ml}$ 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* = *SUNBandMatrix_Data*(A), or equivalently *A_data* = *SM_DATA_B*(A), and then access *A_data*[*i*] within the loop.
 - First obtain the array of column pointers via *A_cols* = *SUNBandMatrix_Cols*(A), or equivalently *A_cols* = *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* = *SUNBandMatrix_Column*(A, *j*) and then to access the entries within that column using *SM_COLUMN_ELEMENT_B*(*A_colj*, *i*, *j*).

All three of these are more efficient than using *SM_ELEMENT_B*(A, *i*, *j*) within a double loop.

- Within the *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 SUNMATRIX_CUSPARSE module is an interface to the NVIDIA cuSPARSE matrix for use on NVIDIA GPUs [6]. 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 `sunmatrix/sunmatrix_cusparse.h`. The installed library to link to is `libsundials_sunmatrixcusparse.lib` where `.lib` is typically `.so` for shared libraries and `.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

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_0 & 0 & \cdots & 0 \\ 0 & \mathbf{A}_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mathbf{A}_{n-1} \end{bmatrix},$$

where all the block matrices \mathbf{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 `SUNLinearSolver_cuSolverSp_batchQR` linear solver (see §8.17).

The 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* 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 5×4 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;
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.




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.

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 )
```

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)->sparsitytype )
```

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 `CSC_MAT` or `CSR_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)

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*.

sunrealttype ***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 [7, 32, 45, 46]. 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 *A_super, 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 `A_super` 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 [10]. 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 N_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 [29, 52] and KokkosKernels [51] 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 ExecSpace::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
ExecSpace::memory_space>
```

```
class DenseMatrix : public sundials::impl::BaseMatrix, public sundials::ConvertibleTo<SUNMatrix>
```

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
- **exec_space** – a *ExecSpace* instance
- **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
- **exec_space** – a *ExecSpace* instance
- **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 ExecSpace::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 IDA

In [Table 7.2](#), we list the matrix functions in the `SUNMatrix` module used within the IDA package. The table also shows, for each function, which of the code modules uses the function. The main IDA 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 IDA user does not need to know anything about the usage of matrix functions by the IDA code modules in order to use IDA. The information is presented as an implementation detail for the interested reader.

Table 7.2: List of matrix functions usage by IDA 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 IDA are: `SUNMatCopy()`, `SUNMatClone()`, `SUNMatScaleAdd()`, `SUNMatScaleAddI()` and `SUNMatMatvec()`. Therefore a user-supplied `SUNMatrix` module for IDA 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_linear_solver.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.,

$$\left\| \tilde{b} - \tilde{A}\tilde{x} \right\|_2 < \text{tol}.$$

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

$$\left\| P_1^{-1}b - P_1^{-1}Ax \right\|_2 < \text{tol}.$$

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()`).

SUNLinearSolver_Type **SUNLinSolGetType**(SUNLinearSolver LS)

Returns the type identifier for the linear solver LS .

Return value:

- `SUNLINEARSOLVER_DIRECT (0)` – the SUNLinSol module requires a matrix, and computes an “exact” solution to the linear system defined by that matrix.

- `SUNLINEARSOLVER_ITERATIVE` (1) – the `SUNLinSol` module 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.
- `SUNLINEARSOLVER_MATRIX_ITERATIVE` (2) – the `SUNLinSol` 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.
- `SUNLINEARSOLVER_MATRIX_EMBEDDED` (3) – the `SUNLinSol` module 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.

Usage:

```
type = SUNLinSolGetType(LS);
```

Note: See §8.1.8.1 for more information on intended use cases corresponding to the linear solver type.

`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 *N_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 = SUNLinLastFlag(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 (*SUNATimesFn)(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 R 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 specific SUNLinSol implementations through the generic SUNLinearSolver abstract base class. The SUNLinearSolver type is a pointer to a structure containing an implementation-dependent *content* field, and an *ops* field, and is defined as

```
typedef struct _generic_SUNLinearSolver *SUNLinearSolver
```

and the generic structure is defined as

```
struct _generic_SUNLinearSolver {
    void *content;
    struct _generic_SUNLinearSolver_Ops *ops;
};
```

where the _generic_SUNLinearSolver_Ops structure is a list of pointers to the various actual linear solver operations provided by a specific implementation. The _generic_SUNLinearSolver_Ops structure is defined as

```
struct _generic_SUNLinearSolver_Ops {
    SUNLinearSolver_Type (*gettype)(SUNLinearSolver);
    SUNLinearSolver_ID (*getid)(SUNLinearSolver);
    SUNErrCode (*setatimes)(SUNLinearSolver, void*, SUNATimesFn);
    SUNErrCode (*setpreconditioner)(SUNLinearSolver, void*,
                                     SUNPSetupFn, SUNPSolveFn);
    SUNErrCode (*setscalingvectors)(SUNLinearSolver,
                                     N_Vector, N_Vector);
    SUNErrCode (*setzeroguess)(SUNLinearSolver, sunbooleantype);
    SUNErrCode (*initialize)(SUNLinearSolver);
    int (*setup)(SUNLinearSolver, SUNMatrix);
    int (*solve)(SUNLinearSolver, SUNMatrix, N_Vector,
                 N_Vector, sunrealttype);
    int (*numiters)(SUNLinearSolver);
    sunrealttype (*resnorm)(SUNLinearSolver);
    sunindextype (*lastflag)(SUNLinearSolver);
    SUNErrCode (*space)(SUNLinearSolver, long int*, long int*);
    N_Vector (*resid)(SUNLinearSolver);
    SUNErrCode (*free)(SUNLinearSolver);
};
```

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()`

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.

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 SUNLINEARSOLVER_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 *hypre* library; this may be used as a template for other customized implementations (see `examples/arkode/CXX_parhyp/ark_heat2D_hypre.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_SPFQMR()`, `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 `SUNLINEAR-SOLVER_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 `MRIStepGetCurrentGamma`) 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 IDA 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
<code>SUNLinSolGetType()</code>	x	x	x
<code>SUNLinSolSetATimes()</code>	†	x	†
<code>SUNLinSolSetPreconditioner()</code>	†	†	†
<code>SUNLinSolSetScalingVectors()</code>	†	†	†
<code>SUNLinSolInitialize()</code>	x	x	x
<code>SUNLinSolSetup()</code>	x	x	x
<code>SUNLinSolSolve()</code>	x	x	x
<code>SUNLinSolNumIters()</code>		x	x
<code>SUNLinSolResid()</code>		x	x
¹ <code>SUNLinSolLastFlag()</code>			
² <code>SUNLinSolFree()</code>			
<code>SUNLinSolSpace()</code>	†	†	†

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 $J = \left(\frac{\partial F}{\partial y} - c_j \frac{\partial F}{\partial y_j} \right)$, and the *actual* linear system $\bar{J}\bar{x} = b$ with

$$\bar{J} = \frac{\partial \bar{F}}{\partial y} - \bar{c}_j \frac{\partial \bar{F}}{\partial y_j},$$

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 [15, 35]. 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 `SUNLINEARSOLVER_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)):

$$\begin{aligned} & \left\| \tilde{b} - \tilde{A}\tilde{x} \right\|_2 < \text{tol} \\ \Leftrightarrow & \left\| SP_1^{-1}b - SP_1^{-1}Ax \right\|_2 < \text{tol} \\ \Leftrightarrow & \sum_{i=0}^{n-1} [W_i (P_1^{-1}(b - Ax))_i]^2 < \text{tol}^2 \\ \Leftrightarrow & W_{mean}^2 \sum_{i=0}^{n-1} [(P_1^{-1}(b - Ax))_i]^2 < \text{tol}^2 \\ \Leftrightarrow & \sum_{i=0}^{n-1} [(P_1^{-1}(b - Ax))_i]^2 < \left(\frac{\text{tol}}{W_{mean}} \right)^2 \\ \Leftrightarrow & \left\| P_1^{-1}(b - Ax) \right\|_2 < \frac{\text{tol}}{W_{mean}} \end{aligned}$$

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 *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_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.

When SUNDIALS is compiled with 32-bit indices (`SUNDIALS_INDEX_SIZE=32`), `sun_klu_symbolic` is mapped to the KLU type `klu_symbolic`; when SUNDIALS compiled with 64-bit indices (`SUNDIALS_INDEX_SIZE=64`) this is mapped to the KLU type `klu_l_symbolic`.

`sun_klu_numeric *SUNLinSol_KLUGetNumeric(SUNLinearSolver S)`

This function returns a pointer to the KLU numeric factorization stored in the `SUNLinSol_KLU` content structure.

When SUNDIALS is compiled with 32-bit indices (`SUNDIALS_INDEX_SIZE=32`), `sun_klu_numeric` is mapped to the KLU type `klu_numeric`; when SUNDIALS is compiled with 64-bit indices (`SUNDIALS_INDEX_SIZE=64`) this is mapped to the KLU type `klu_l_numeric`.

`sun_klu_common *SUNLinSol_KLUGetCommon(SUNLinearSolver S)`

This function returns a pointer to the KLU common structure stored in the `SUNLinSol_KLU` content structure.

When SUNDIALS is compiled with 32-bit indices (`SUNDIALS_INDEX_SIZE=32`), `sun_klu_common` is mapped to the KLU type `klu_common`; when SUNDIALS is compiled with 64-bit indices (`SUNDIALS_INDEX_SIZE=64`) this is mapped to the KLU type `klu_l_common`.

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;
    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 ([3, 24]). 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 ε is the double-precision unit roundoff), then a new factorization is performed.
- The module includes the routine `SUNKLURenit`, 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 `m1`, then the upper triangular factor U can have upper bandwidth as big as $smu = \text{MIN}(N-1, mu+m1)$. The lower triangular factor L has lower bandwidth `m1`.

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 `SUN-Matrix` 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 [50]. 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;
    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 [33]) 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

$$\tilde{A}\tilde{x} = \tilde{b} \quad (8.4)$$

where

$$\begin{aligned} \tilde{A} &= SP^{-1}AP^{-1}S, \\ \tilde{b} &= SP^{-1}b, \\ \tilde{x} &= S^{-1}Px. \end{aligned} \quad (8.5)$$

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 δ .

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 sunctx)`

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.
- *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 `max1` argument that is ≤ 0 will result in the default value (5).

Since the PCG algorithm is designed to only support symmetric preconditioning, then any of the `pre-type` 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_PCGSetMax1`(*SUNLinearSolver* S, int max1)

This function updates the number of linear solver iterations to allow.

Arguments:

- *S* – `SUNLinSol_PCG` object to update.
- *max1* – 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;
    sunbooleantype zeroguess;
    int numiters;
    sunrealtype 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 [53] 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 ≤ 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_SPBCGS*SetPrecType(*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_SPBCGS*SetMaxl(*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;
    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 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 [48] 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_SPGMR Usage

The header file to be included when using this module is `sunlinsol/sunlinsol_spfgmr.h`. The `SUNLinSol_SPGMR` module is accessible from all SUNDIALS solvers *without* linking to the `libsundials_sunlinsolspfgmr` 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 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 ≤ 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_SPGMR` object for these packages, 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*

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;
```

(continues on next page)

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 [49] 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 ≤ 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 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 *V;
    sunreal_t **Hes;
    sunreal_t *givens;
    N_Vector xcor;
    sunreal_t *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),

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 [31] 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 ≤ 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;
    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 r_star;
    N_Vector q;
    N_Vector d;
    N_Vector v;
    N_Vector p;
    N_Vector *r;
    N_Vector u;
    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 unprefix 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;
    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 [7, 32, 45, 46]. 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 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.

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 [8, 26, 44]. 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 *sunrealtype* 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 [5]. 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 */
    sunbooleantype 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 [10]. 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, sunctx};

```

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 LinearSolver : public ConvertibleTo<SUNLinearSolver>

```

LinearSolver() = default;

Default constructor - means the solver must be moved to.

LinearSolver(std::shared_ptr<typename *GkoSolverType*::Factory> gko_solver_factory, SUNContext sunctx)

Constructs a new `LinearSolver` from a Ginkgo solver factory.

Parameters

- **gko_solver_factory** – The Ginkgo solver factory (typically `gko::matrix::<type>::Factory`)
- **sunctx** – 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 [51] 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 $10 * tol$, 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 (`SUNNonlinSolInitialization()`), setup (`SUNNonlinSolSetup()`), and destruction (`SUNNonlinSolFree()`) are optional.

`SUNNonlinearSolver_Type` **SUNNonlinSolGetType**(*SUNNonlinearSolver* NLS)

This *required* function returns the nonlinear solver type.

Arguments:

- *NLS* – a `SUNNonlinSol` object.

Return value:

The `SUNNonlinSol` type identifier (of type `int`) will be one of the following:

- `SUNNONLINEARSOLVER_ROOTFIND` – 0, the `SUNNonlinSol` module solves $F(y) = 0$.
- `SUNNONLINEARSOLVER_FIXEDPOINT` – 1, the `SUNNonlinSol` module solves $G(y) = y$.

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* callSetup, 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 *SUNNonlinSolLSetupFn*.

Return value:

- A *SUNErrCode*

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 set this operation to NULL.

SUNErrCode **SUNNonlinSolSetLSolveFn**(*SUNNonlinearSolver* NLS, *SUNNonlinSolLSolveFn* 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 *SUNNonlinSolLSolveFn*.

Return value:

- A *SUNErrCode*

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 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 *nitors)

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.
- *nitors* – 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_RECVR	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. The type SUNNonlinearSolver is defined as follows:

```
typedef struct _generic_SUNNonlinearSolver *SUNNonlinearSolver
```

and the generic structure is defined as

```
struct _generic_SUNNonlinearSolver {
    void *content;
    struct _generic_SUNNonlinearSolver_Ops *ops;
};
```

where the _generic_SUNNonlinearSolver_Ops structure is a list of pointers to the various actual nonlinear solver operations provided by a specific implementation. The _generic_SUNNonlinearSolver_Ops structure is defined as

```
struct _generic_SUNNonlinearSolver_Ops {
    SUNNonlinearSolver_Type (*gettype)(SUNNonlinearSolver);
    int (*initialize)(SUNNonlinearSolver);
    int (*setup)(SUNNonlinearSolver, N_Vector, void*);
    int (*solve)(SUNNonlinearSolver, N_Vector, N_Vector,
                N_Vector, sunrealtype, sunbooleantype, void*);
    int (*free)(SUNNonlinearSolver);
    int (*setsysfn)(SUNNonlinearSolver, SUNNonlinSolSysFn);
    int (*setlsetupfn)(SUNNonlinearSolver, SUNNonlinSolLSetupFn);
    int (*setlsolvefn)(SUNNonlinearSolver, SUNNonlinSolLSolveFn);
    int (*setctestfn)(SUNNonlinearSolver, SUNNonlinSolConvTestFn,
                      void*);
    int (*setmaxiters)(SUNNonlinearSolver, int);
    int (*getnumiters)(SUNNonlinearSolver, long int*);
    int (*getcuriter)(SUNNonlinearSolver, int*);
    int (*getnumconvfails)(SUNNonlinearSolver, long int*);
};
```

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 im-

plementation, 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 callSetup, void* mem)
{
    return((int) NLS->ops->solve(NLS, y0, y, w, tol, callSetup, 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()`

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 IDA 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 IDA 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 $\alpha = h_n^{-1} \alpha_{n,0}$.

The nonlinear system function provided by IDA to the nonlinear solver module internally updates the current value of the new state and its derivative based on the input correction vector. The updated vectors are used when calling the DAE residual function and when setting up linear solves (e.g., for updating the Jacobian or preconditioner).

IDA 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, *sunrealt* *cj)

The function IDAGetCurrentCj returns the scalar c_j which is proportional to the inverse of the step size (α in (9.1)).

Arguments:

- ida_mem – pointer to the IDA memory block.
- cj – the value of c_j .

Return value:

- IDA_SUCCESS – The optional output value has been successfully set.
- IDA_MEM_NULL – The IDA 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 IDA memory block.
- y – the current y vector.

Return value:

- IDA_SUCCESS – The optional output value has been successfully set.
- IDA_MEM_NULL – The IDA 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 IDA memory block.
- yp – the current \dot{y} vector.

Return value:

- IDA_SUCCESS – The optional output value has been successfully set.
- IDA_MEM_NULL – The IDA memory block is NULL.

```
int IDAGetNonlinearSystemData(void *ida_mem, sunrealtype *tcur, N_Vector *yypred, N_Vector *yppred,
                             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 IDA memory block.
- `tcur` – current value of the independent variable t_n .
- `yypred` – predicted value of y_{pred} at t_n .
- `yppred` – 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 (α 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 IDA 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 `SUNNonlinSolL-SolveFn` functions supplied by IDA (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.,

$$\begin{aligned} yyn &= y_{pred} + y_{cor}, \\ ypn &= \dot{y}_{pred} + \alpha \dot{y}_{cor}, \\ res &= F(t_n, y_n, \dot{y}_n), \end{aligned}$$

and $f_n = f(t_n, y^n)$ 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` and `fn` are only current when `IDAGetNonlinearSystemData()` 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 IDA 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 IDA memory block is NULL.

int **IDAComputeYp**(void *ida_mem, *N_Vector* ycor, *N_Vector* yp)

The function computes $\dot{y}(t)$.

Arguments:

- `ida_mem` – pointer to the IDA 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 IDA memory block 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 [14, 19, 25, 27, 43]. 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 [9, 30, 47, 54]. 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 [54], 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 β 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 `N_Vectors` used in acceleration algorithm (length m),
- `dg` – array of `N_Vectors` used in acceleration algorithm (length m),
- `q` – array of `N_Vectors` used in acceleration algorithm (length m),
- `Xvecs` – `N_Vector` pointer array used in acceleration algorithm (length $m+1$),
- `fold` – `N_Vector` used in acceleration algorithm, and
- `gold` – `N_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 `SUNNonlinSol_GetNumIters()` 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 *SUNNonlinSolGetNumConverFails()* 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()*, *SUNNonlinSolSetConverTestFn()*, 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 **SUNNonlinSolGetPestcError_PetscSNES**(*SUNNonlinearSolver* NLS, PestcErrorCode *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 a structure containing a pointer to actual data (`ptr`), the data memory type, and a flag indicating ownership of that data pointer. This structure is defined as

```
struct SUNMemory_  
{  
    void*      ptr;  
    SUNMemoryType type;  
    sunboolean_t own;  
    size_t     bytes;  
};
```

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:

```

typedef enum
{
    SUNMEMTYPE_HOST,      /* pageable memory accessible on the host */
    SUNMEMTYPE_PINNED,   /* page-locked memory accesible on the host */
    SUNMEMTYPE_DEVICE,   /* memory accessible from the device */
    SUNMEMTYPE_UVM       /* memory accessible from the host or device */
} SUNMemoryType;

```

typedef struct SUNMemoryHelper_ **SUNMemoryHelper**

The `SUNMemoryHelper` type is a pointer to a structure containing a pointer to the implementation-specific member data (content) and a virtual method table of member functions (ops). This structure is defined as

```

struct SUNMemoryHelper_
{
    void*          content;
    SUNMemoryHelper_Ops ops;
    SUNContext     sunctx;
};

```

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_
{
    /* operations that implementations are required to provide */
    SUNErrCode (*alloc)(SUNMemoryHelper, SUNMemory* memptr, size_t mem_size,
                       SUNMemoryType mem_type, void* queue);
    SUNErrCode (*dealloc)(SUNMemoryHelper, SUNMemory mem, void* queue);
    SUNErrCode (*copy)(SUNMemoryHelper, SUNMemory dst, SUNMemory src,
                      size_t mem_size, void* queue);

    /* operations that provide default implementations */
    SUNErrCode (*copyasnc)(SUNMemoryHelper, SUNMemory dst,
                          SUNMemory src, size_t mem_size, void* queue);
    SUNErrCode (*getallocstats)(SUNMemoryHelper, SUNMemoryType mem_type, unsigned long* num_allocations,
                               unsigned long* num_deallocations, size_t* bytes_allocated,
                               size_t* bytes_high_watermark);
    SUNMemoryHelper (*clone)(SUNMemoryHelper);
    SUNErrCode      (*destroy)(SUNMemoryHelper);
};

```

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.
- `mem_ptr` – 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 [4] 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 [1]. 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](#) (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-X.Y.Z.tar.gz`, where `SOLVER` is one of: `sundials`, `cvode`, `cvodes`, `arkode`, `ida`, `idas`, or `kinsol`, and `X.Y.Z` 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 -zxf SOLVER-X.Y.Z.tar.gz
```

This will extract source files under a directory `SOLVER-X.Y.Z`.

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-X.Y.Z` 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 3.12.0 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_CCODE                 *ON
BUILD_CCODES                *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_ARCODE *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 (Unix/Linux)

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_CVOICE

Build the CVOICE 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

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.

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 DPC++.

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_DIRPath to *hypre* header files

Default: none

HYPRE_LIBRARYPath 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

mpif90 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_UNSUPPORTEDThis 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 built 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 v7.0.0. 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 v7.0.0.

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

The Intel [oneAPI Math Kernel Library \(oneMKL\)](#) includes CPU and DPC++ interfaces for LAPACK dense linear algebra routines. The SUNDIALS oneMKL interface targets the 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:
```

(continues on next page)

(continued from previous page)

```
# cmake -D SUNDIALS_DIR=/path/to/sundials/installation

find_package(SUNDIALS 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::cvoid SUNDIALS::nvecpetsc)
```

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

continues on next page

Table 12.1 – continued from previous page

HIP	Libraries	libsundials_nvechip.LIB
	Headers	nvector/nvector_hip.h
RAJA	Libraries	libsundials_nveccudaraja.LIB
		libsundials_nvechipraja.LIB
SYCL	Headers	nvector/nvector_rajah.h
	Libraries	libsundials_nvecsycl.LIB
MANYVECTOR	Headers	nvector/nvector_sycl.h
	Libraries	libsundials_nvecmanyvector.LIB
MPIMANYVECTOR	Headers	nvector/nvector_manyvector.h
	Libraries	libsundials_nvecmpimanyvector.LIB
MPIPLUSX	Headers	nvector/nvector_mpimanyvector.h
	Libraries	libsundials_nvecmpiplusx.LIB
	Headers	nvector/nvector_mpiplusx.h
SUNMATRIX Modules		
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
SUNLINSOL Modules		
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

continues on next page

Table 12.1 – continued from previous page

SPFGMR	Libraries	libsundials_sunlinsolspfgmr.LIB
	Headers	sunlinsol/sunlinsol_spfgmr.h
SPGMR	Libraries	libsundials_sunlinsolspgmr.LIB
	Headers	sunlinsol/sunlinsol_spgmr.h
SPTFQMR	Libraries	libsundials_sunlinsolsptfqr.LIB
	Headers	sunlinsol/sunlinsol_sptfqr.h
SUPERLUDIST	Libraries	libsundials_sunlinsolsuperludist.LIB
	Headers	sunlinsol/sunlinsol_superludist.h
SUPERLUMT	Libraries	libsundials_sunlinsolsuperlumt.LIB
	Headers	sunlinsol/sunlinsol_superlumt.h
SUNNONLINSOL Modules		
NEWTON	Libraries	libsundials_sunnonlinsolnewton.LIB
	Headers	sunnonlinsol/sunnonlinsol_newton.h
FIXEDPOINT	Libraries	libsundials_sunnonlinsolfixedpoint.LIB
	Headers	sunnonlinsol/sunnonlinsol_fixedpoint.h
PETSCSNES	Libraries	libsundials_sunnonlinsolpetscsnes.LIB
	Headers	sunnonlinsol/sunnonlinsol_petscsnes.h
SUNMEMORY Modules		
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
SUNDIALS Packages		
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
		cvode/cvode_proj.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
ARKODE	Libraries	libsundials_arkode.LIB
		libsundials_xbraid.LIB
	Headers	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	

continues on next page

Table 12.1 – continued from previous page

IDA	Libraries	arkode/arkode_erkstep.h
		arkode/arkode_impl.h
		arkode/arkode_ls.h
		arkode/arkode_mrstep.h
		arkode/arkode_xbraid.h
IDAS	Headers	libsundials_ida.LIB
		ida/ida.h
		ida/ida_bbdpre.h
		ida/ida_impl.h
KINSOL	Libraries	ida/ida_ls.h
		libsundials_idas.LIB
		idas/idas.h
		idas/idas_bbdpre.h
		idas/idas_impl.h
KINSOL	Headers	libsundials_kinsol.LIB
		kinsol/kinsol.h
		kinsol/kinsol_bbdpre.h
		kinsol/kinsol_impl.h
		kinsol/kinsol_ls.h

Chapter 13

IDA 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 IDA input constants

Table 13.1: IDA Input Constants

Constant	Value	Description
IDA_NORMAL	1	Solver returns at specified output time.
IDA_ONE_STEP	2	Solver returns after each successful step.
IDA_YA_YDP_INIT	1	Compute y_a and \dot{y}_d , given y_d .
IDA_Y_INIT	2	Compute y , given \dot{y} .
SUN_PREC_NONE	0	No preconditioning
SUN_PREC_LEFT	1	Preconditioning on the left only.
SUN_MODIFIED_GS	1	Use modified Gram-Schmidt procedure.
SUN_CLASSICAL_GS	2	Use classical Gram-Schmidt procedure.

13.2 IDA output constants

Table 13.2: IDA Output Constants

Constant	Value	Description
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.

continues on next page

Table 13.2 – continued from previous page

Constant	Value	Description
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 IDA 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.
IDALS_SUCCESS	0	Successful function return.
IDALS_MEM_NULL	-1	The <code>ida_mem</code> 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 <code>N_Vector</code> module.
IDALS_MEM_FAIL	-4	A memory allocation request failed.
IDALS_PMEM_NULL	-5	The preconditioner module has not been initialized.
IDALS_JACFUNC_UNRECVR	-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 <code>SUNMatrix</code> module.
IDALS_SUNLS_FAIL	-9	An error occurred with the current <code>SUNLinearSolver</code> module.

Chapter 14

Release History

Date	SUNDIALS	ARKODE	CVODE	CVODES	IDA	IDAS	KINSOL
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
Sep 2017	3.0.0	2.0.0	3.0.0	3.0.0	3.0.0	2.0.0	3.0.0

continues on next page

Table 14.1 – continued from previous page

Date	SUNDIALS	ARKODE	CVODE	CVODES	IDA	IDAS	KINSOL
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 ³	–	–	–	–
Feb 1999	–	–	–	–	1.0.0 ⁴	–	–
Aug 1998	–	–	–	–	–	–	1.0.0 ⁵
Jul 1997	–	–	1.0.0 ²	–	–	–	–
Sep 1994	–	–	1.0.0 ¹	–	–	–	–

1. CVODE written
2. PVODE written
3. CVODE and PVODE combined
4. IDA written
5. KINSOL written

Bibliography

- [1] AMD ROCm Documentation. <https://rocmdocs.amd.com/en/latest/index.html>.
- [2] Intel oneAPI Programming Guide. <https://software.intel.com/content/www/us/en/develop/documentation/oneapi-programming-guide/top.html>.
- [3] KLU Sparse Matrix Factorization Library. <http://faculty.cse.tamu.edu/davis/suitesparse.html>.
- [4] NVIDIA CUDA Programming Guide. <https://docs.nvidia.com/cuda/index.html>.
- [5] NVIDIA cuSOLVER Programming Guide. <https://docs.nvidia.com/cuda/cusolver/index.html>.
- [6] NVIDIA cuSPARSE Programming Guide. <https://docs.nvidia.com/cuda/cusparse/index.html>.
- [7] SuperLU_DIST Parallel Sparse Matrix Factorization Library. https://portal.nersc.gov/project/sparse/superlu/#superlu_dist.
- [8] SuperLU_MT Threaded Sparse Matrix Factorization Library. https://portal.nersc.gov/project/sparse/superlu/#superlu_mt.
- [9] D. G. Anderson. Iterative procedures for nonlinear integral equations. *J. Assoc. Comput. Machinery*, 12:547–560, 1965. doi:10.1145/321296.321305.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] 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.
- [14] 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.
- [15] 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.
- [16] 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.
- [17] 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.

- [18] 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.
- [19] 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.
- [20] 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.
- [21] 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.
- [22] 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.
- [23] 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.
- [24] 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.
- [25] 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.
- [26] 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.
- [27] J. E. Dennis and R. B. Schnabel. *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. SIAM, Philadelphia, 1996. doi:10.1137/1.9781611971200.
- [28] 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.
- [29] 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.
- [30] 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.
- [31] 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.
- [32] 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.
- [33] 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.
- [34] K. L. Hiebert and L. F. Shampine. Implicitly Defined Output Points for Solutions of ODEs. Technical Report SAND80-0180, Sandia National Laboratories, February 1980.
- [35] A. C. Hindmarsh. The PVODE and IDA Algorithms. Technical Report UCRL-ID-141558, LLNL, Dec 2000.
- [36] 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.
- [37] 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.
- [38] Alan C. Hindmarsh and Radu Serban. Example Programs for CVODE v7.0.0. Technical Report, LLNL, 2024. UCRL-SM-208110.

- [39] Alan C. Hindmarsh, Radu Serban, Cody J. Balos, David J. Gardner, Daniel R. Reynolds, and Carol S. Woodward. User Documentation for CVODE v7.0.0. Technical Report UCRL-SM-208108, LLNL, 2024.
- [40] Alan C. Hindmarsh, Radu Serban, Cody J. Balos, David J. Gardner, Daniel R. Reynolds, and Carol S. Woodward. User Documentation for KINSOL v7.0.0. Technical Report UCRL-SM-208116, LLNL, 2024.
- [41] Alan C. Hindmarsh, Radu Serban, and Aaron Collier. Example Programs for IDA v7.0.0. Technical Report UCRL-SM-208113, LLNL, 2024.
- [42] 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.
- [43] C. T. Kelley. *Iterative Methods for Solving Linear and Nonlinear Equations*. SIAM, Philadelphia, 1995. doi:10.1137/1.9781611970944.
- [44] 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.
- [45] 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.
- [46] 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.
- [47] 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.
- [48] Y. Saad. A flexible inner-outer preconditioned GMRES algorithm. *SIAM J. Sci. Comput.*, 14(2):461–469, 1993. doi:10.1137/0914028.
- [49] 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.
- [50] 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.
- [51] 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.
- [52] Christian R. Trott, Damien Lebrun-Grandié, Daniel Arndt, Jan Ciesko, Vinh Dang, Nathan Ellingwood, Rahul Kumar 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.
- [53] H. A. Van Der Vorst. Bi-CGSTAB: A Fast and Smoothly Converging Variant of Bi-CG for the Solution of Nonsymmetric Linear Systems. *SIAM J. Sci. Stat. Comp.*, 13(2):631–644, 1992. doi:10.1137/0913035.
- [54] 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.

Index

C

ccmake, 341

cmake, 345

CMake options

adiak_DIR, 356

AMDGPU_TARGETS, 350

BUILD_ARKODE, 345

BUILD_CVOICE, 345

BUILD_CVOICES, 345

BUILD_FORTRAN_MODULE_INTERFACE, 349

BUILD_IDA, 345

BUILD_IDAS, 345

BUILD_KINSOL, 345

BUILD_SHARED_LIBS, 345

BUILD_STATIC_LIBS, 346

CALIPER_DIR, 355

CMAKE_BUILD_TYPE, 346

CMAKE_C_COMPILER, 346

CMAKE_C_EXTENSIONS, 346

CMAKE_C_FLAGS, 346

CMAKE_C_FLAGS_DEBUG, 346

CMAKE_C_FLAGS_MINSIZEREL, 346

CMAKE_C_FLAGS_RELEASE, 346

CMAKE_C_STANDARD, 346

CMAKE_CUDA_ARCHITECTURES, 348

CMAKE_CXX_COMPILER, 346

CMAKE_CXX_EXTENSIONS, 347

CMAKE_CXX_FLAGS, 347

CMAKE_CXX_FLAGS_DEBUG, 347

CMAKE_CXX_FLAGS_MINSIZEREL, 347

CMAKE_CXX_FLAGS_RELEASE, 347

CMAKE_CXX_STANDARD, 347

CMAKE_Fortran_COMPILER, 347

CMAKE_Fortran_FLAGS, 347

CMAKE_Fortran_FLAGS_DEBUG, 347

CMAKE_Fortran_FLAGS_MINSIZEREL, 347

CMAKE_Fortran_FLAGS_RELEASE, 347

CMAKE_INSTALL_LIBDIR, 348

CMAKE_INSTALL_PREFIX, 348

ENABLE_ADIK, 355

ENABLE_CALIPER, 355

ENABLE_CUDA, 348

ENABLE_GINKGO, 350

ENABLE_HIP, 350

ENABLE_HYPRE, 350

ENABLE_KLU, 351

ENABLE_KOKKOS, 350

ENABLE_KOKKOS_KERNELS, 350

ENABLE_LAPACK, 351

ENABLE_MAGMA, 351

ENABLE_MPI, 352

ENABLE_ONEMKL, 352

ENABLE_OPENMP, 353

ENABLE_PETSC, 353

ENABLE_PTHREAD, 353

ENABLE_RAJA, 353

ENABLE_SUPERLUDIST, 353

ENABLE_SUPERLUMT, 354

ENABLE_SYCL, 355

ENABLE_XBRAID, 357

EXAMPLES_ENABLE_C, 348

EXAMPLES_ENABLE_CUDA, 348

EXAMPLES_ENABLE_CXX, 348

EXAMPLES_ENABLE_F2003, 348

EXAMPLES_INSTALL, 348

EXAMPLES_INSTALL_PATH, 349

Ginkgo_DIR, 350

HYPRE_INCLUDE_DIR, 350

HYPRE_LIBRARY, 351

KLU_INCLUDE_DIR, 351

KLU_LIBRARY_DIR, 351

Kokkos_DIR, 350

KokkosKernels_DIR, 350

LAPACK_LIBRARIES, 351

MAGMA_DIR, 351

MPI_C_COMPILER, 352

MPI_CXX_COMPILER, 352

MPI_Fortran_COMPILER, 352

MPIEXEC_EXECUTABLE, 352

ONEMKL_DIR, 352

PETSC_DIR, 353

PETSC_INCLUDES, 353

PETSC_LIBRARIES, 353

SUNDIALS_BUILD_WITH_MONITORING, 349

SUNDIALS_BUILD_WITH_PROFILING, 349

SUNDIALS_ENABLE_ERROR_CHECKS, 349

SUNDIALS_GINKGO_BACKENDS, 350

SUNDIALS_INDEX_SIZE, 356

SUNDIALS_INDEX_TYPE, 356

SUNDIALS_INSTALL_CMAKEDIR, 357

SUNDIALS_LAPACK_CASE, 356
 SUNDIALS_LAPACK_UNDESCORES, 356
 SUNDIALS_LOGGING_LEVEL, 349
 SUNDIALS_MAGMA_BACKENDS, 351
 SUNDIALS_MATH_LIBRARY, 357
 SUNDIALS_ONEMKL_USE_GETRF_LOOP, 352
 SUNDIALS_ONEMKL_USE_GETRS_LOOP, 353
 SUNDIALS_PRECISION, 356
 SUNDIALS_RAJA_BACKENDS, 353
 SUNDIALS_SYCL_2020_UNSUPPORTED, 355
 SUPERLUDIST_DIR, 354
 SUPERLUDIST_INCLUDE_DIR, 354
 SUPERLUDIST_INCLUDE_DIRS, 354
 SUPERLUDIST_LIBRARIES, 354
 SUPERLUDIST_LIBRARY_DIR, 354
 SUPERLUDIST_OpenMP, 354
 SUPERLUMT_INCLUDE_DIR, 355
 SUPERLUMT_LIBRARY_DIR, 355
 SUPERLUMT_THREAD_TYPE, 355
 USE_XSDK_DEFAULTS, 357
 XBRAID_DIR, 357
 XBRAID_INCLUDES, 357
 XBRAID_LIBRARIES, 357

cmake-gui, 341

CopyFromDevice (C++ function), 190

CopyToDevice (C++ function), 190

D

DenseLinearSolver (C++ class), 303

DenseLinearSolver::~DenseLinearSolver (C++ function), 304

DenseLinearSolver::Convert (C++ function), 304

DenseLinearSolver::DenseLinearSolver (C++ function), 303, 304

DenseLinearSolver::operator SUNLinearSolver (C++ function), 304

DenseLinearSolver::operator=(C++ function), 304

DenseMatrix (C++ class), 243

DenseMatrix::~DenseMatrix (C++ function), 244

DenseMatrix::BlockCols (C++ function), 244

DenseMatrix::BlockRows (C++ function), 244

DenseMatrix::Blocks (C++ function), 244

DenseMatrix::Cols (C++ function), 244

DenseMatrix::Convert (C++ function), 244

DenseMatrix::DenseMatrix (C++ function), 243, 244

DenseMatrix::ExecSpace (C++ function), 244

DenseMatrix::operator SUNMatrix (C++ function), 244

DenseMatrix::operator=(C++ function), 244

DenseMatrix::Rows (C++ function), 244

DenseMatrix::View (C++ function), 244

G

GetDenseMat (C++ function), 244

GetVec (C++ function), 190

I

IDABBDCommFn (C type), 126

IDABBDLocalFn (C type), 125

IDABBDPrecGetNumGfnEvals (C function), 129

IDABBDPrecGetWorkSpace (C function), 128

IDABBDPrecInit (C function), 127

IDABBDPrecReInit (C function), 128

IDACalcIC (C function), 79

IDAClearStopTime (C function), 86

IDAComputeY (C function), 317

IDAComputeYp (C function), 318

IDACreate (C function), 75

IDAEwtFn (C type), 118

IDAFree (C function), 75

IDAGetActualInitStep (C function), 105

IDAGetConsistentIC (C function), 109

IDAGetCurrentCj (C function), 316

IDAGetCurrentOrder (C function), 105

IDAGetCurrentStep (C function), 105

IDAGetCurrentTime (C function), 106

IDAGetCurrentY (C function), 316

IDAGetCurrentYp (C function), 316

IDAGetDky (C function), 101

IDAGetErrWeights (C function), 106

IDAGetEstLocalErrors (C function), 107

IDAGetIntegratorStats (C function), 107

IDAGetJac (C function), 111

IDAGetJacCj (C function), 111

IDAGetJacNumSteps (C function), 112

IDAGetJacTime (C function), 111

IDAGetLastLinFlag (C function), 115

IDAGetLastOrder (C function), 105

IDAGetLastStep (C function), 105

IDAGetLinReturnFlagName (C function), 116

IDAGetLinWorkSpace (C function), 112

IDAGetNonlinearSystemData (C function), 317

IDAGetNonlinSolvStats (C function), 108

IDAGetNumBacktrackOps (C function), 109

IDAGetNumErrTestFails (C function), 104

IDAGetNumGEvals (C function), 110

IDAGetNumJacEvals (C function), 112

IDAGetNumJtimesEvals (C function), 115

IDAGetNumJTSetupEvals (C function), 115

IDAGetNumLinConvFails (C function), 113

IDAGetNumLinIters (C function), 113

IDAGetNumLinResEvals (C function), 113

IDAGetNumLinSolvSetups (C function), 104

IDAGetNumNonlinSolvConvFails (C function), 108

IDAGetNumNonlinSolvIters (C function), 108

IDAGetNumPrecEvals (C function), 114

IDAGetNumPrecSolves (C function), 114

IDAGetNumResEvals (C function), 104

IDAGetNumSteps (C function), 103
 IDAGetNumStepSolveFails (C function), 104
 IDAGetReturnFlagName (C function), 109
 IDAGetRootInfo (C function), 110
 IDAGetTolScaleFactor (C function), 106
 IDAGetUserData (C function), 108
 IDAGetWorkspace (C function), 103
 IDAInit (C function), 75
 IDALsJacFn (C type), 119
 IDALsJacTimesSetupFn (C type), 122
 IDALsJacTimesVecFn (C type), 121
 IDALsPrecSetupFn (C type), 123
 IDALsPrecSolveFn (C type), 123
 IDAPrintAllStats (C function), 108
 IDAReInit (C function), 117
 IDAResFn (C type), 118
 IDARootFn (C type), 119
 IDARootInit (C function), 81
 IDASetConstraints (C function), 87
 IDASetDeltaCjLSetup (C function), 89
 IDASetEpsLin (C function), 92
 IDASetEtaConvFail (C function), 99
 IDASetEtaFixedStepBounds (C function), 97
 IDASetEtaLow (C function), 98
 IDASetEtaMax (C function), 98
 IDASetEtaMin (C function), 98
 IDASetEtaMinErrFail (C function), 99
 IDASetId (C function), 87
 IDASetIncrementFactor (C function), 90
 IDASetInitStep (C function), 84
 IDASetJacFn (C function), 88
 IDASetJacTimes (C function), 90
 IDASetJacTimesResFn (C function), 91
 IDASetLinearSolutionScaling (C function), 89
 IDASetLinearSolver (C function), 78
 IDASetLineSearchOffIC (C function), 96
 IDASetLSNormFactor (C function), 92
 IDASetMaxBacksIC (C function), 96
 IDASetMaxConvFails (C function), 93
 IDASetMaxErrTestFails (C function), 86
 IDASetMaxNonlinIters (C function), 93
 IDASetMaxNumItersIC (C function), 96
 IDASetMaxNumJacsIC (C function), 95
 IDASetMaxNumSteps (C function), 84
 IDASetMaxNumStepsIC (C function), 95
 IDASetMaxOrd (C function), 84
 IDASetMaxStep (C function), 85
 IDASetMinStep (C function), 85
 IDASetNlsResFn (C function), 94
 IDASetNoInactiveRootWarn (C function), 100
 IDASetNonlinConvCoef (C function), 94
 IDASetNonlinConvCoefIC (C function), 95
 IDASetNonlinearSolver (C function), 79
 IDASetPreconditioner (C function), 91

IDASetRootDirection (C function), 100
 IDASetStepToleranceIC (C function), 97
 IDASetStopTime (C function), 85
 IDASetSuppressAlg (C function), 86
 IDASetUserData (C function), 83
 IDASolve (C function), 81
 IDASStolerances (C function), 76
 IDASVtolerances (C function), 76
 IDAWFtolerances (C function), 76

L

LinearSolver (C++ class), 301
 LinearSolver::~LinearSolver (C++ function), 301
 LinearSolver::Convert (C++ function), 302
 LinearSolver::GkoExec (C++ function), 302
 LinearSolver::GkoFactory (C++ function), 302
 LinearSolver::GkoSolver (C++ function), 302
 LinearSolver::LinearSolver (C++ function), 301
 LinearSolver::NumIters (C++ function), 302
 LinearSolver::operator SUNLinearSolver (C++ function), 302
 LinearSolver::operator= (C++ function), 301
 LinearSolver::ResNorm (C++ function), 302
 LinearSolver::Setup (C++ function), 302
 LinearSolver::Solve (C++ function), 302

M

Matrix (C++ class), 240
 Matrix::~Matrix (C++ function), 241
 Matrix::Convert (C++ function), 241
 Matrix::GkoExec (C++ function), 241
 Matrix::GkoMtx (C++ function), 241
 Matrix::GkoSize (C++ function), 241
 Matrix::Matrix (C++ function), 240, 241
 Matrix::operator SUNMatrix (C++ function), 241
 Matrix::operator= (C++ function), 241

N

N_VAbs (C function), 141
 N_VAddConst (C function), 141
 N_VBufPack (C function), 149
 N_VBufSize (C function), 149
 N_VBufUnpack (C function), 150
 N_VClone (C function), 138
 N_VCloneEmpty (C function), 138
 N_VCloneVectorArray (C function), 134
 N_VCloneVectorArrayEmpty (C function), 134
 N_VCompare (C function), 143
 N_VConst (C function), 140
 N_VConstrMask (C function), 143
 N_VConstrMaskLocal (C function), 148
 N_VConstVectorArray (C function), 145
 N_VCopyFromDevice_Cuda (C function), 171

- N_VCopyFromDevice_Hip (C function), 176
 N_VCopyFromDevice_OpenMPDEV (C function), 192
 N_VCopyFromDevice_Raja (C function), 186
 N_VCopyFromDevice_Sycl (C++ function), 181
 N_VCopyOps (C function), 136
 N_VCopyToDevice_Cuda (C function), 171
 N_VCopyToDevice_Hip (C function), 176
 N_VCopyToDevice_OpenMPDEV (C function), 192
 N_VCopyToDevice_Raja (C function), 186
 N_VCopyToDevice_Sycl (C++ function), 181
 N_VDestroy (C function), 138
 N_VDestroyVectorArray (C function), 134
 N_VDiv (C function), 140
 N_VDotProd (C function), 141
 N_VDotProdLocal (C function), 146
 N_VDotProdMulti (C function), 144
 N_VDotProdMultiAllReduce (C function), 149
 N_VDotProdMultiLocal (C function), 149
 N_Vector (C type), 131
 N_VEnableConstVectorArray_Cuda (C function), 172
 N_VEnableConstVectorArray_Hip (C function), 177
 N_VEnableConstVectorArray_ManyVector (C function), 197
 N_VEnableConstVectorArray_MPIManyVector (C function), 201
 N_VEnableConstVectorArray_OpenMP (C function), 160
 N_VEnableConstVectorArray_OpenMPDEV (C function), 193
 N_VEnableConstVectorArray_Parallel (C function), 157
 N_VEnableConstVectorArray_ParHyp (C function), 167
 N_VEnableConstVectorArray_Petsc (C function), 169
 N_VEnableConstVectorArray_Pthreads (C function), 164
 N_VEnableConstVectorArray_Raja (C function), 187
 N_VEnableConstVectorArray_Serial (C function), 153
 N_VEnableConstVectorArray_Sycl (C++ function), 182
 N_VEnableDotProdMulti_Cuda (C function), 172
 N_VEnableDotProdMulti_Hip (C function), 177
 N_VEnableDotProdMulti_ManyVector (C function), 197
 N_VEnableDotProdMulti_MPIManyVector (C function), 201
 N_VEnableDotProdMulti_OpenMP (C function), 160
 N_VEnableDotProdMulti_OpenMPDEV (C function), 193
 N_VEnableDotProdMulti_Parallel (C function), 157
 N_VEnableDotProdMulti_ParHyp (C function), 166
 N_VEnableDotProdMulti_Petsc (C function), 168
 N_VEnableDotProdMulti_Pthreads (C function), 164
 N_VEnableDotProdMulti_Serial (C function), 153
 N_VEnableDotProdMulti_Sycl (C++ function), 182
 N_VEnableFusedOps_Cuda (C function), 172
 N_VEnableFusedOps_Hip (C function), 177
 N_VEnableFusedOps_ManyVector (C function), 197
 N_VEnableFusedOps_MPIManyVector (C function), 200
 N_VEnableFusedOps_OpenMP (C function), 160
 N_VEnableFusedOps_OpenMPDEV (C function), 192
 N_VEnableFusedOps_Parallel (C function), 156
 N_VEnableFusedOps_ParHyp (C function), 166
 N_VEnableFusedOps_Petsc (C function), 168
 N_VEnableFusedOps_Pthreads (C function), 164
 N_VEnableFusedOps_Raja (C function), 187
 N_VEnableFusedOps_Serial (C function), 153
 N_VEnableFusedOps_Sycl (C++ function), 182
 N_VEnableLinearCombination_Cuda (C function), 172
 N_VEnableLinearCombination_Hip (C function), 177
 N_VEnableLinearCombination_ManyVector (C function), 197
 N_VEnableLinearCombination_MPIManyVector (C function), 200
 N_VEnableLinearCombination_OpenMP (C function), 160
 N_VEnableLinearCombination_OpenMPDEV (C function), 192
 N_VEnableLinearCombination_Parallel (C function), 156
 N_VEnableLinearCombination_ParHyp (C function), 166
 N_VEnableLinearCombination_Petsc (C function), 168
 N_VEnableLinearCombination_Pthreads (C function), 164
 N_VEnableLinearCombination_Raja (C function), 187
 N_VEnableLinearCombination_Serial (C function), 153
 N_VEnableLinearCombination_Sycl (C++ function), 182
 N_VEnableLinearCombinationVectorArray_Cuda (C function), 172
 N_VEnableLinearCombinationVectorArray_Hip (C function), 177
 N_VEnableLinearCombinationVectorArray_OpenMP (C function), 161
 N_VEnableLinearCombinationVectorArray_OpenMPDEV (C function), 193
 N_VEnableLinearCombinationVectorArray_Parallel (C function), 157
 N_VEnableLinearCombinationVectorArray_ParHyp (C function), 167
 N_VEnableLinearCombinationVectorArray_Petsc

- (*C function*), 169
- N_VEnableLinearCombinationVectorArray_Pthreads (*C function*), 164
- N_VEnableLinearCombinationVectorArray_Raja (*C function*), 187
- N_VEnableLinearCombinationVectorArray_Serial (*C function*), 153
- N_VEnableLinearCombinationVectorArray_Sycl (*C++ function*), 182
- N_VEnableLinearSumVectorArray_Cuda (*C function*), 172
- N_VEnableLinearSumVectorArray_Hip (*C function*), 177
- N_VEnableLinearSumVectorArray_ManyVector (*C function*), 197
- N_VEnableLinearSumVectorArray_MPIManyVector (*C function*), 201
- N_VEnableLinearSumVectorArray_OpenMP (*C function*), 160
- N_VEnableLinearSumVectorArray_OpenMPDEV (*C function*), 193
- N_VEnableLinearSumVectorArray_Parallel (*C function*), 157
- N_VEnableLinearSumVectorArray_ParHyp (*C function*), 166
- N_VEnableLinearSumVectorArray_Petsc (*C function*), 168
- N_VEnableLinearSumVectorArray_Pthreads (*C function*), 164
- N_VEnableLinearSumVectorArray_Raja (*C function*), 187
- N_VEnableLinearSumVectorArray_Serial (*C function*), 153
- N_VEnableLinearSumVectorArray_Sycl (*C++ function*), 182
- N_VEnableScaleAddMulti_Cuda (*C function*), 172
- N_VEnableScaleAddMulti_Hip (*C function*), 177
- N_VEnableScaleAddMulti_ManyVector (*C function*), 197
- N_VEnableScaleAddMulti_MPIManyVector (*C function*), 200
- N_VEnableScaleAddMulti_OpenMP (*C function*), 160
- N_VEnableScaleAddMulti_OpenMPDEV (*C function*), 192
- N_VEnableScaleAddMulti_Parallel (*C function*), 157
- N_VEnableScaleAddMulti_ParHyp (*C function*), 166
- N_VEnableScaleAddMulti_Petsc (*C function*), 168
- N_VEnableScaleAddMulti_Pthreads (*C function*), 164
- N_VEnableScaleAddMulti_Raja (*C function*), 187
- N_VEnableScaleAddMulti_Serial (*C function*), 153
- N_VEnableScaleAddMulti_Sycl (*C++ function*), 182
- N_VEnableScaleAddMultiVectorArray_Cuda (*C function*), 172
- N_VEnableScaleAddMultiVectorArray_Hip (*C function*), 177
- N_VEnableScaleAddMultiVectorArray_OpenMP (*C function*), 161
- N_VEnableScaleAddMultiVectorArray_OpenMPDEV (*C function*), 193
- N_VEnableScaleAddMultiVectorArray_Parallel (*C function*), 157
- N_VEnableScaleAddMultiVectorArray_ParHyp (*C function*), 167
- N_VEnableScaleAddMultiVectorArray_Petsc (*C function*), 169
- N_VEnableScaleAddMultiVectorArray_Pthreads (*C function*), 164
- N_VEnableScaleAddMultiVectorArray_Raja (*C function*), 187
- N_VEnableScaleAddMultiVectorArray_Serial (*C function*), 153
- N_VEnableScaleAddMultiVectorArray_Sycl (*C++ function*), 182
- N_VEnableScaleVectorArray_Cuda (*C function*), 172
- N_VEnableScaleVectorArray_Hip (*C function*), 177
- N_VEnableScaleVectorArray_ManyVector (*C function*), 197
- N_VEnableScaleVectorArray_MPIManyVector (*C function*), 201
- N_VEnableScaleVectorArray_OpenMP (*C function*), 160
- N_VEnableScaleVectorArray_OpenMPDEV (*C function*), 193
- N_VEnableScaleVectorArray_Parallel (*C function*), 157
- N_VEnableScaleVectorArray_ParHyp (*C function*), 166
- N_VEnableScaleVectorArray_Petsc (*C function*), 169
- N_VEnableScaleVectorArray_Pthreads (*C function*), 164
- N_VEnableScaleVectorArray_Raja (*C function*), 187
- N_VEnableScaleVectorArray_Serial (*C function*), 153
- N_VEnableScaleVectorArray_Sycl (*C++ function*), 182
- N_VEnableWrmsNormMaskVectorArray_Cuda (*C function*), 172
- N_VEnableWrmsNormMaskVectorArray_Hip (*C function*), 177
- N_VEnableWrmsNormMaskVectorArray_ManyVector (*C function*), 197
- N_VEnableWrmsNormMaskVectorArray_MPIManyVector (*C function*), 201
- N_VEnableWrmsNormMaskVectorArray_OpenMP (*C function*), 161

- N_VEnableWrmsNormMaskVectorArray_OpenMPDEV (C function), 193
 N_VEnableWrmsNormMaskVectorArray_Parallel (C function), 157
 N_VEnableWrmsNormMaskVectorArray_ParHyp (C function), 167
 N_VEnableWrmsNormMaskVectorArray_Petsc (C function), 169
 N_VEnableWrmsNormMaskVectorArray_Pthreads (C function), 164
 N_VEnableWrmsNormMaskVectorArray_Serial (C function), 153
 N_VEnableWrmsNormVectorArray_Cuda (C function), 172
 N_VEnableWrmsNormVectorArray_Hip (C function), 177
 N_VEnableWrmsNormVectorArray_ManyVector (C function), 197
 N_VEnableWrmsNormVectorArray_MPIManyVector (C function), 201
 N_VEnableWrmsNormVectorArray_OpenMP (C function), 161
 N_VEnableWrmsNormVectorArray_OpenMPDEV (C function), 193
 N_VEnableWrmsNormVectorArray_Parallel (C function), 157
 N_VEnableWrmsNormVectorArray_ParHyp (C function), 167
 N_VEnableWrmsNormVectorArray_Petsc (C function), 169
 N_VEnableWrmsNormVectorArray_Pthreads (C function), 164
 N_VEnableWrmsNormVectorArray_Serial (C function), 153
 N_VFreeEmpty (C function), 136
 N_VGetArrayPointer (C function), 139
 N_VGetArrayPointer_MPIPlusX (C function), 202
 N_VGetCommunicator (C function), 139
 N_VGetDeviceArrayPointer (C function), 139
 N_VGetDeviceArrayPointer_Cuda (C function), 170
 N_VGetDeviceArrayPointer_Hip (C function), 175
 N_VGetDeviceArrayPointer_OpenMPDEV (C function), 192
 N_VGetDeviceArrayPointer_Raja (C function), 186
 N_VGetDeviceArrayPointer_Sycl (C++ function), 181
 N_VGetHostArrayPointer_Cuda (C function), 170
 N_VGetHostArrayPointer_Hip (C function), 175
 N_VGetHostArrayPointer_OpenMPDEV (C function), 192
 N_VGetHostArrayPointer_Raja (C function), 186
 N_VGetHostArrayPointer_Sycl (C++ function), 181
 N_VGetLength (C function), 139
 N_VGetLocalLength (C function), 140
 N_VGetLocalLength_MPIPlusX (C function), 202
 N_VGetLocalLength_Parallel (C function), 156
 N_VGetLocalVector_MPIPlusX (C function), 202
 N_VGetNumSubvectors_ManyVector (C function), 196
 N_VGetNumSubvectors_MPIManyVector (C function), 200
 N_VGetSubvector_ManyVector (C function), 196
 N_VGetSubvector_MPIManyVector (C function), 200
 N_VGetSubvectorArrayPointer_ManyVector (C function), 196
 N_VGetSubvectorArrayPointer_MPIManyVector (C function), 200
 N_VGetSubvectorLocalLength_ManyVector (C function), 196
 N_VGetSubvectorLocalLength_MPIManyVector (C function), 200
 N_VGetVecAtIndexVectorArray (C function), 135
 N_VGetVector_ParHyp (C function), 166
 N_VGetVector_Petsc (C function), 168
 N_VGetVector_Trilinos (C++ function), 194
 N_VGetVectorID (C function), 138
 N_VInv (C function), 141
 N_VInvTest (C function), 143
 N_VInvTestLocal (C function), 148
 N_VIsManagedMemory_Cuda (C function), 170
 N_VIsManagedMemory_Hip (C function), 175
 N_VIsManagedMemory_Raja (C function), 186
 N_VIsManagedMemory_Sycl (C++ function), 181
 N_VL1Norm (C function), 142
 N_VL1NormLocal (C function), 147
 N_VLinearCombination (C function), 144
 N_VLinearCombinationVectorArray (C function), 146
 N_VLinearSum (C function), 140
 N_VLinearSumVectorArray (C function), 144
 N_VMake_Cuda (C function), 171
 N_VMake_Hip (C function), 176
 N_VMake_MPIManyVector (C function), 199
 N_VMake_MPIPlusX (C function), 202
 N_VMake_OpenMP (C function), 160
 N_VMake_OpenMPDEV (C function), 192
 N_VMake_Parallel (C function), 156
 N_VMake_ParHyp (C function), 166
 N_VMake_Petsc (C function), 168
 N_VMake_Pthreads (C function), 163
 N_VMake_Raja (C function), 186
 N_VMake_Serial (C function), 152
 N_VMake_Sycl (C++ function), 181
 N_VMake_Trilinos (C++ function), 194
 N_VMakeManaged_Cuda (C function), 171
 N_VMakeManaged_Hip (C function), 176
 N_VMakeManaged_Raja (C function), 186
 N_VMakeManaged_Sycl (C++ function), 181

N_VMakeWithManagedAllocator_Cuda (C function), 171
 N_VMaxNorm (C function), 141
 N_VMaxNormLocal (C function), 147
 N_VMin (C function), 142
 N_VMinLocal (C function), 147
 N_VMinQuotient (C function), 143
 N_VMinQuotientLocal (C function), 148
 N_VNew_Cuda (C function), 170
 N_VNew_Hip (C function), 175
 N_VNew_ManyVector (C function), 196
 N_VNew_MPIManyVector (C function), 199
 N_VNew_OpenMP (C function), 160
 N_VNew_OpenMPDEV (C function), 192
 N_VNew_Parallel (C function), 156
 N_VNew_Pthreads (C function), 163
 N_VNew_Raja (C function), 186
 N_VNew_Serial (C function), 152
 N_VNew_Sycl (C++ function), 180
 N_VNewEmpty (C function), 136
 N_VNewEmpty_Cuda (C function), 171
 N_VNewEmpty_Hip (C function), 176
 N_VNewEmpty_OpenMP (C function), 160
 N_VNewEmpty_OpenMPDEV (C function), 192
 N_VNewEmpty_Parallel (C function), 156
 N_VNewEmpty_ParHyp (C function), 166
 N_VNewEmpty_Petsc (C function), 168
 N_VNewEmpty_Pthreads (C function), 163
 N_VNewEmpty_Raja (C function), 186
 N_VNewEmpty_Serial (C function), 152
 N_VNewEmpty_Sycl (C++ function), 181
 N_VNewManaged_Cuda (C function), 170
 N_VNewManaged_Hip (C function), 176
 N_VNewManaged_Raja (C function), 186
 N_VNewManaged_Sycl (C++ function), 180
 N_VNewVectorArray (C function), 134
 N_VNewWithMemHelp_Cuda (C function), 171
 N_VNewWithMemHelp_Hip (C function), 176
 N_VNewWithMemHelp_Raja (C function), 186
 N_VNewWithMemHelp_Sycl (C++ function), 181
 N_VPrint_Cuda (C function), 171
 N_VPrint_Hip (C function), 176
 N_VPrint_OpenMP (C function), 160
 N_VPrint_OpenMPDEV (C function), 192
 N_VPrint_Parallel (C function), 156
 N_VPrint_ParHyp (C function), 166
 N_VPrint_Petsc (C function), 168
 N_VPrint_Pthreads (C function), 163
 N_VPrint_Raja (C function), 186
 N_VPrint_Serial (C function), 152
 N_VPrint_Sycl (C++ function), 182
 N_VPrintFile_Cuda (C function), 171
 N_VPrintFile_Hip (C function), 176
 N_VPrintFile_OpenMP (C function), 160
 N_VPrintFile_OpenMPDEV (C function), 192
 N_VPrintFile_Parallel (C function), 156
 N_VPrintFile_ParHyp (C function), 166
 N_VPrintFile_Petsc (C function), 168
 N_VPrintFile_Pthreads (C function), 163
 N_VPrintFile_Raja (C function), 187
 N_VPrintFile_Serial (C function), 152
 N_VPrintFile_Sycl (C++ function), 182
 N_VProd (C function), 140
 N_VScale (C function), 140
 N_VScaleAddMulti (C function), 144
 N_VScaleAddMultiVectorArray (C function), 146
 N_VScaleVectorArray (C function), 145
 N_VSetArrayPointer (C function), 139
 N_VSetArrayPointer_MPIPlusX (C function), 202
 N_VSetDeviceArrayPointer_Sycl (C++ function), 181
 N_VSetHostArrayPointer_Sycl (C++ function), 181
 N_VSetKernelExecPolicy_Cuda (C function), 171
 N_VSetKernelExecPolicy_Hip (C function), 176
 N_VSetKernelExecPolicy_Sycl (C++ function), 181
 N_VSetSubvectorArrayPointer_ManyVector (C function), 196
 N_VSetSubvectorArrayPointer_MPIManyVector (C function), 200
 N_VSetVecAtIndexVectorArray (C function), 135
 N_VSpace (C function), 138
 N_VWl2Norm (C function), 142
 N_VWrmsNorm (C function), 142
 N_VWrmsNormMask (C function), 142
 N_VWrmsNormMaskVectorArray (C function), 145
 N_VWrmsNormVectorArray (C function), 145
 N_VWSqrSumLocal (C function), 147
 N_VWSqrSumMaskLocal (C function), 147
 NV_COMM_P (C macro), 155
 NV_CONTENT_OMP (C macro), 158
 NV_CONTENT_OMPDEV (C macro), 191
 NV_CONTENT_P (C macro), 155
 NV_CONTENT_PT (C macro), 162
 NV_CONTENT_S (C macro), 151
 NV_DATA_DEV_OMPDEV (C macro), 191
 NV_DATA_HOST_OMPDEV (C macro), 191
 NV_DATA_OMP (C macro), 159
 NV_DATA_P (C macro), 155
 NV_DATA_PT (C macro), 162
 NV_DATA_S (C macro), 151
 NV_GLOBLENGTH_P (C macro), 155
 NV_Ith_OMP (C macro), 159
 NV_Ith_P (C macro), 155
 NV_Ith_PT (C macro), 163
 NV_Ith_S (C macro), 152
 NV_LENGTH_OMP (C macro), 159
 NV_LENGTH_OMPDEV (C macro), 191
 NV_LENGTH_PT (C macro), 162

NV_LENGTH_S (C macro), 152
 NV_LOCLENGTH_P (C macro), 155
 NV_NUM_THREADS_OMP (C macro), 159
 NV_NUM_THREADS_PT (C macro), 163
 NV_OWN_DATA_OMP (C macro), 158
 NV_OWN_DATA_OMPDEV (C macro), 191
 NV_OWN_DATA_P (C macro), 155
 NV_OWN_DATA_PT (C macro), 162
 NV_OWN_DATA_S (C macro), 151

S

SM_COLS_B (C macro), 226
 SM_COLS_D (C macro), 212
 SM_COLUMN_B (C macro), 226
 SM_COLUMN_D (C macro), 213
 SM_COLUMN_ELEMENT_B (C macro), 227
 SM_COLUMNS_B (C macro), 224
 SM_COLUMNS_D (C macro), 212
 SM_COLUMNS_S (C macro), 235
 SM_CONTENT_B (C macro), 224
 SM_CONTENT_D (C macro), 211
 SM_CONTENT_S (C macro), 233
 SM_DATA_B (C macro), 226
 SM_DATA_D (C macro), 212
 SM_DATA_S (C macro), 235
 SM_ELEMENT_B (C macro), 227
 SM_ELEMENT_D (C macro), 213
 SM_INDEXPTRS_S (C macro), 236
 SM_INDEXVALS_S (C macro), 235
 SM_LBAND_B (C macro), 224
 SM_LDATA_B (C macro), 226
 SM_LDATA_D (C macro), 212
 SM_LDIM_B (C macro), 226
 SM_NNZ_S (C macro), 235
 SM_NP_S (C macro), 235
 SM_ROWS_B (C macro), 224
 SM_ROWS_D (C macro), 212
 SM_ROWS_S (C macro), 235
 SM_SPARSETYPE_S (C macro), 235
 SM_SUBAND_B (C macro), 226
 SM_UBAND_B (C macro), 224
 SUN_LOGLEVEL_ALL (C enumerator), 52
 SUN_LOGLEVEL_DEBUG (C enumerator), 52
 SUN_LOGLEVEL_ERROR (C enumerator), 52
 SUN_LOGLEVEL_INFO (C enumerator), 52
 SUN_LOGLEVEL_NONE (C enumerator), 52
 SUN_LOGLEVEL_WARNING (C enumerator), 52
 SUN_OUTPUTFORMAT_CSV (C enumerator), 43
 SUN_OUTPUTFORMAT_TABLE (C enumerator), 43
 SUNAbortErrorHandlerFn (C function), 50
 SUNATimesFn (C type), 253
 SUNBandMatrix (C function), 227
 SUNBandMatrix_Cols (C function), 228
 SUNBandMatrix_Column (C function), 228

SUNBandMatrix_Columns (C function), 228
 SUNBandMatrix_Data (C function), 228
 SUNBandMatrix_LData (C function), 228
 SUNBandMatrix_LDim (C function), 228
 SUNBandMatrix_LowerBandwidth (C function), 228
 SUNBandMatrix_Print (C function), 227
 SUNBandMatrix_Rows (C function), 227
 SUNBandMatrix_StoredUpperBandwidth (C function), 228
 SUNBandMatrix_UpperBandwidth (C function), 228
 SUNBandMatrixStorage (C function), 227
 sunbooleantype (C type), 42
 SUNComm (C type), 43
 SUNContext (C type), 43
 SUNContext_ClearErrHandlers (C function), 45
 SUNContext_Create (C function), 43
 SUNContext_Free (C function), 44
 SUNContext_GetLastError (C function), 44
 SUNContext_GetLogger (C function), 46
 SUNContext_GetProfiler (C function), 45
 SUNContext_PeekLastError (C function), 44
 SUNContext_PopErrorHandler (C function), 45
 SUNContext_PushErrorHandler (C function), 44
 SUNContext_SetLogger (C function), 45
 SUNContext_SetProfiler (C function), 45
 SUNCudaBlockReduceAtomicExecPolicy (C++ function), 174
 SUNCudaBlockReduceExecPolicy (C++ function), 174
 SUNCudaExecPolicy (C++ type), 173
 SUNCudaGridStrideExecPolicy (C++ function), 174
 SUNCudaThreadDirectExecPolicy (C++ function), 174
 SUNDenseMatrix (C function), 213
 SUNDenseMatrix_Cols (C function), 213
 SUNDenseMatrix_Column (C function), 213
 SUNDenseMatrix_Columns (C function), 213
 SUNDenseMatrix_Data (C function), 213
 SUNDenseMatrix_LData (C function), 213
 SUNDenseMatrix_Print (C function), 213
 SUNDenseMatrix_Rows (C function), 213
 SUNDIALSFileClose (C function), 67
 SUNDIALSFileOpen (C function), 66
 SUNDIALSGetVersion (C function), 59
 SUNDIALSGetVersionNumber (C function), 59
 SUNErrCode (C type), 49
 SUNErrorHandlerFn (C type), 50
 SUNGetErrMsg (C function), 49
 SUNHipBlockReduceAtomicExecPolicy (C++ function), 179
 SUNHipBlockReduceExecPolicy (C++ function), 179
 SUNHipExecPolicy (C++ type), 178
 SUNHipGridStrideExecPolicy (C++ function), 179
 SUNHipThreadDirectExecPolicy (C++ function), 179
 sunindextype (C type), 42

- SUNLinearSolver (*C type*), 255
- SUNLinSol_Band (*C function*), 262
- SUNLinSol_cuSolverSp_batchQR (*C function*), 299
- SUNLinSol_cuSolverSp_batchQR_GetDescription (*C function*), 299
- SUNLinSol_cuSolverSp_batchQR_GetDeviceSpace (*C function*), 299
- SUNLinSol_cuSolverSp_batchQR_SetDescription (*C function*), 299
- SUNLinSol_Dense (*C function*), 264
- SUNLinSol_KLU (*C function*), 265
- SUNLinSol_KLUGetCommon (*C function*), 267
- SUNLinSol_KLUGetNumeric (*C function*), 267
- SUNLinSol_KLUGetSymbolic (*C function*), 266
- SUNLinSol_KLUReInit (*C function*), 266
- SUNLinSol_KLUSetOrdering (*C function*), 266
- SUNLinSol_LapackBand (*C function*), 269
- SUNLinSol_LapackDense (*C function*), 270
- SUNLinSol_MagmaDense (*C function*), 272
- SUNLinSol_MagmaDense_SetAsync (*C function*), 273
- SUNLinSol_OneMklDense (*C function*), 274
- SUNLinSol_PCG (*C function*), 275
- SUNLinSol_PCGSetMax1 (*C function*), 276
- SUNLinSol_PCGSetPrecType (*C function*), 276
- SUNLinSol_SPBCGS (*C function*), 278
- SUNLinSol_SPBCGSSetMax1 (*C function*), 279
- SUNLinSol_SPBCGSSetPrecType (*C function*), 279
- SUNLinSol_SPFQMR (*C function*), 282
- SUNLinSol_SPFQMRSetGStype (*C function*), 283
- SUNLinSol_SPFQMRSetMaxRestarts (*C function*), 283
- SUNLinSol_SPFQMRSetPrecType (*C function*), 282
- SUNLinSol_SPGMR (*C function*), 286
- SUNLinSol_SPGMRSetGStype (*C function*), 286
- SUNLinSol_SPGMRSetMaxRestarts (*C function*), 287
- SUNLinSol_SPGMRSetPrecType (*C function*), 286
- SUNLinSol_SPTFQMR (*C function*), 289
- SUNLinSol_SPTFQMRSetMax1 (*C function*), 290
- SUNLinSol_SPTFQMRSetPrecType (*C function*), 290
- SUNLinSol_SuperLUDIST (*C function*), 293
- SUNLinSol_SuperLUDIST_GetBerr (*C function*), 293
- SUNLinSol_SuperLUDIST_GetGridinfo (*C function*), 293
- SUNLinSol_SuperLUDIST_GetLUstruct (*C function*), 293
- SUNLinSol_SuperLUDIST_GetScalePermstruct (*C function*), 294
- SUNLinSol_SuperLUDIST_GetSOLVEstruct (*C function*), 294
- SUNLinSol_SuperLUDIST_GetSuperLUOptions (*C function*), 294
- SUNLinSol_SuperLUDIST_GetSuperLUstat (*C function*), 294
- SUNLinSol_SuperLUMT (*C function*), 296
- SUNLinSol_SuperLUMTSetOrdering (*C function*), 296
- SUNLinSolFree (*C function*), 250
- SUNLinSolFreeEmpty (*C function*), 257
- SUNLinSolGetID (*C function*), 249
- SUNLinSolGetType (*C function*), 248
- SUNLinSolInitialize (*C function*), 249
- SUNLinSolLastFlag (*C function*), 252
- SUNLinSolNewEmpty (*C function*), 257
- SUNLinSolNumIters (*C function*), 252
- SUNLinSolResid (*C function*), 252
- SUNLinSolResNorm (*C function*), 252
- SUNLinSolSetATimes (*C function*), 251
- SUNLinSolSetPreconditioner (*C function*), 251
- SUNLinSolSetScalingVectors (*C function*), 251
- SUNLinSolSetup (*C function*), 249
- SUNLinSolSetZeroGuess (*C function*), 251
- SUNLinSolSolve (*C function*), 250
- SUNLinSolSpace (*C function*), 252
- SUNLogErrorHandlerFn (*C function*), 50
- SUNLogger (*C type*), 52
- SUNLogger_Create (*C function*), 52
- SUNLogger_CreateFromEnv (*C function*), 53
- SUNLogger_Destroy (*C function*), 55
- SUNLogger_Flush (*C function*), 54
- SUNLogger_GetOutputRank (*C function*), 55
- SUNLogger_QueueMsg (*C function*), 54
- SUNLogger_SetDebugFilename (*C function*), 54
- SUNLogger_SetErrorFilename (*C function*), 53
- SUNLogger_SetInfoFilename (*C function*), 53
- SUNLogger_SetWarningFilename (*C function*), 53
- SUNLogLevel (*C enum*), 52
- SUNMatClone (*C function*), 209
- SUNMatCopy (*C function*), 210
- SUNMatCopyOps (*C function*), 208
- SUNMatDestroy (*C function*), 210
- SUNMatFreeEmpty (*C function*), 209
- SUNMatGetID (*C function*), 209
- SUNMatMatvec (*C function*), 211
- SUNMatMatvecSetup (*C function*), 211
- SUNMatNewEmpty (*C function*), 208
- SUNMatrix (*C type*), 207
- SUNMatrix_cuSparse_BlockColumns (*C function*), 230
- SUNMatrix_cuSparse_BlockData (*C function*), 231
- SUNMatrix_cuSparse_BlockNNZ (*C function*), 230
- SUNMatrix_cuSparse_BlockRows (*C function*), 230
- SUNMatrix_cuSparse_Columns (*C function*), 230
- SUNMatrix_cuSparse_CopyFromDevice (*C function*), 231
- SUNMatrix_cuSparse_CopyToDevice (*C function*), 231
- SUNMatrix_cuSparse_Data (*C function*), 230
- SUNMatrix_cuSparse_IndexPointers (*C function*), 230
- SUNMatrix_cuSparse_IndexValues (*C function*), 230

- SUNMatrix_cuSparse_MakeCSR (C function), 230
- SUNMatrix_cuSparse_MatDescr (C function), 231
- SUNMatrix_cuSparse_NewBlockCSR (C function), 229
- SUNMatrix_cuSparse_NewCSR (C function), 229
- SUNMatrix_cuSparse_NNZ (C function), 230
- SUNMatrix_cuSparse_NumBlocks (C function), 230
- SUNMatrix_cuSparse_Rows (C function), 230
- SUNMatrix_cuSparse_SetFixedPattern (C function), 231
- SUNMatrix_cuSparse_SetKernelExecPolicy (C function), 231
- SUNMatrix_cuSparse_SparseType (C function), 230
- SUNMatrix_MagmaDense (C function), 215
- SUNMatrix_MagmaDense_Block (C function), 217
- SUNMatrix_MagmaDense_BlockColumn (C function), 217
- SUNMatrix_MagmaDense_BlockColumns (C function), 216
- SUNMatrix_MagmaDense_BlockData (C function), 216
- SUNMatrix_MagmaDense_BlockRows (C function), 216
- SUNMatrix_MagmaDense_Column (C function), 217
- SUNMatrix_MagmaDense_Columns (C function), 216
- SUNMatrix_MagmaDense_CopyFromDevice (C function), 218
- SUNMatrix_MagmaDense_CopyToDevice (C function), 217
- SUNMatrix_MagmaDense_Data (C function), 216
- SUNMatrix_MagmaDense_LData (C function), 216
- SUNMatrix_MagmaDense_NumBlocks (C function), 216
- SUNMatrix_MagmaDense_Rows (C function), 215
- SUNMatrix_MagmaDenseBlock (C function), 215
- SUNMatrix_OneMklDense (C++ function), 219
- SUNMatrix_OneMklDense_Block (C function), 222
- SUNMatrix_OneMklDense_BlockColumn (C function), 222
- SUNMatrix_OneMklDense_BlockColumns (C function), 220
- SUNMatrix_OneMklDense_BlockData (C function), 221
- SUNMatrix_OneMklDense_BlockLData (C function), 221
- SUNMatrix_OneMklDense_BlockRows (C function), 220
- SUNMatrix_OneMklDense_Column (C function), 221
- SUNMatrix_OneMklDense_Columns (C function), 220
- SUNMatrix_OneMklDense_CopyFromDevice (C function), 222
- SUNMatrix_OneMklDense_CopyToDevice (C function), 222
- SUNMatrix_OneMklDense_Data (C function), 221
- SUNMatrix_OneMklDense_LData (C function), 221
- SUNMatrix_OneMklDense_NumBlocks (C function), 220
- SUNMatrix_OneMklDense_Rows (C function), 220
- SUNMatrix_OneMklDenseBlock (C++ function), 219
- SUNMatrix_SLUNRloc (C function), 238
- SUNMatrix_SLUNRloc_OwnData (C function), 239
- SUNMatrix_SLUNRloc_Print (C function), 238
- SUNMatrix_SLUNRloc_ProcessGrid (C function), 238
- SUNMatrix_SLUNRloc_SuperMatrix (C function), 238
- SUNMatScaleAdd (C function), 210
- SUNMatScaleAddI (C function), 210
- SUNMatSpace (C function), 210
- SUNMatZero (C function), 210
- SUNMemory (C type), 327
- SUNMemoryHelper (C type), 328
- SUNMemoryHelper_Alias (C function), 330
- SUNMemoryHelper_Alloc (C function), 328
- SUNMemoryHelper_Alloc_Cuda (C function), 333
- SUNMemoryHelper_Alloc_Hip (C function), 335
- SUNMemoryHelper_Alloc_Sycl (C function), 337
- SUNMemoryHelper_Clone (C function), 331
- SUNMemoryHelper_Copy (C function), 329
- SUNMemoryHelper_Copy_Cuda (C function), 333
- SUNMemoryHelper_Copy_Hip (C function), 335
- SUNMemoryHelper_Copy_Sycl (C function), 337
- SUNMemoryHelper_CopyAsync (C function), 331
- SUNMemoryHelper_CopyAsync_Cuda (C function), 334
- SUNMemoryHelper_CopyAsync_Hip (C function), 335
- SUNMemoryHelper_CopyAsync_Sycl (C function), 338
- SUNMemoryHelper_CopyOps (C function), 330
- SUNMemoryHelper_Cuda (C function), 332
- SUNMemoryHelper_Dealloc (C function), 329
- SUNMemoryHelper_Dealloc_Cuda (C function), 333
- SUNMemoryHelper_Dealloc_Hip (C function), 335
- SUNMemoryHelper_Dealloc_Sycl (C function), 337
- SUNMemoryHelper_Destroy (C function), 332
- SUNMemoryHelper_GetAllocStats (C function), 330
- SUNMemoryHelper_GetAllocStats_Cuda (C function), 334
- SUNMemoryHelper_GetAllocStats_Hip (C function), 336
- SUNMemoryHelper_GetAllocStats_Sycl (C function), 338
- SUNMemoryHelper_Hip (C function), 334
- SUNMemoryHelper_NewEmpty (C function), 330
- SUNMemoryHelper_Ops (C type), 328
- SUNMemoryHelper_Sycl (C function), 336
- SUNMemoryHelper_Wrap (C function), 330
- SUNMemoryNewEmpty (C function), 327
- SUNMemoryType (C enum), 327
- SUNMPIAbortErrorHandlerFn (C function), 51
- SUNNonlinearSolver (C type), 314
- SUNNonlinSol_FixedPoint (C function), 322
- SUNNonlinSol_Newton (C function), 319
- SUNNonlinSol_PetscSNES (C function), 325
- SUNNonlinSolConvTestFn (C type), 313
- SUNNonlinSolFree (C function), 309

- SUNNonlinSolFreeEmpty (C function), 315
 SUNNonlinSolGetCurIter (C function), 311
 SUNNonlinSolGetNumConvFails (C function), 311
 SUNNonlinSolGetNumIters (C function), 311
 SUNNonlinSolGetPetscError_PetscSNES (C function), 325
 SUNNonlinSolGetSNES_PetscSNES (C function), 325
 SUNNonlinSolGetSysFn_FixedPoint (C function), 322
 SUNNonlinSolGetSysFn_Newton (C function), 319
 SUNNonlinSolGetSysFn_PetscSNES (C function), 326
 SUNNonlinSolGetType (C function), 308
 SUNNonlinSolInitialize (C function), 308
 SUNNonlinSolLSetupFn (C type), 312
 SUNNonlinSolLSolveFn (C type), 313
 SUNNonlinSolNewEmpty (C function), 315
 SUNNonlinSolSetConvTestFn (C function), 310
 SUNNonlinSolSetDamping_FixedPoint (C function), 322
 SUNNonlinSolSetLSetupFn (C function), 309
 SUNNonlinSolSetLSolveFn (C function), 310
 SUNNonlinSolSetMaxIters (C function), 310
 SUNNonlinSolSetSysFn (C function), 309
 SUNNonlinSolSetup (C function), 308
 SUNNonlinSolSolve (C function), 308
 SUNNonlinSolSysFn (C type), 312
 SUNOutputFormat (C enum), 43
 SUNProfiler (C type), 56
 SUNProfiler_Begin (C function), 57
 SUNProfiler_Create (C function), 56
 SUNProfiler_End (C function), 57
 SUNProfiler_Free (C function), 57
 SUNProfiler_GetElapsedTime (C function), 57
 SUNProfiler_GetTimerResolution (C function), 57
 SUNProfiler_Print (C function), 58
 SUNProfiler_Reset (C function), 58
 SUNPSetupFn (C type), 253
 SUNPSolveFn (C type), 253
 sunrealtype (C type), 41
 SUNSparseFromBandMatrix (C function), 236
 SUNSparseFromDenseMatrix (C function), 236
 SUNSparseMatrix (C function), 236
 SUNSparseMatrix_Columns (C function), 237
 SUNSparseMatrix_Data (C function), 237
 SUNSparseMatrix_IndexPointers (C function), 237
 SUNSparseMatrix_IndexValues (C function), 237
 SUNSparseMatrix_NNZ (C function), 237
 SUNSparseMatrix_NP (C function), 237
 SUNSparseMatrix_Print (C function), 237
 SUNSparseMatrix_Realloc (C function), 236
 SUNSparseMatrix_Reallocate (C function), 237
 SUNSparseMatrix_Rows (C function), 237
 SUNSparseMatrix_SparseType (C function), 237
 SUNSyclBlockReduceExecPolicy (C++ function), 184
 SUNSyclExecPolicy (C++ type), 183
 SUNSyclGridStrideExecPolicy (C++ function), 184
 SUNSyclThreadDirectExecPolicy (C++ function), 184
- ## V
- Vector (C++ class), 189
 Vector::~~Vector (C++ function), 189
 Vector::Convert (C++ function), 190
 Vector::HostView (C++ function), 190
 Vector::Length (C++ function), 189
 Vector::operator N_Vector (C++ function), 190
 Vector::operator= (C++ function), 189
 Vector::Vector (C++ function), 189
 Vector::View (C++ function), 190
 vector_type (C++ type), 194