

User Documentation for IDAS v1.2.0 (SUNDIALS v2.6.0)

Radu Serban, Cosmin Petra, and Alan C. Hindmarsh
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory

March 9, 2015



UCRL-SM-234051

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.

Contents

List of Tables	vii
List of Figures	ix
1 Introduction	1
1.1 Changes from previous versions	2
1.2 Reading this User Guide	3
2 Mathematical Considerations	5
2.1 IVP solution	5
2.2 Preconditioning	9
2.3 Rootfinding	9
2.4 Pure quadrature integration	10
2.5 Forward sensitivity analysis	11
2.5.1 Forward sensitivity methods	11
2.5.2 Selection of the absolute tolerances for sensitivity variables	12
2.5.3 Evaluation of the sensitivity right-hand side	13
2.5.4 Quadratures depending on forward sensitivities	14
2.6 Adjoint sensitivity analysis	14
2.6.1 Sensitivity of $G(p)$	14
2.6.2 Sensitivity of $g(T, p)$	15
2.6.3 Checkpointing scheme	16
2.7 Second-order sensitivity analysis	17
3 Code Organization	19
3.1 SUNDIALS organization	19
3.2 IDAS organization	19
4 Using IDAS for IVP Solution	25
4.1 Access to library and header files	25
4.2 Data types	26
4.3 Header files	26
4.4 A skeleton of the user's main program	27
4.5 User-callable functions	30
4.5.1 IDAS initialization and deallocation functions	30
4.5.2 IDAS tolerance specification functions	31
4.5.3 Linear solver specification functions	33
4.5.4 Initial condition calculation function	37
4.5.5 Rootfinding initialization function	38
4.5.6 IDAS solver function	38
4.5.7 Optional input functions	40
4.5.7.1 Main solver optional input functions	40
4.5.7.2 Dense direct linear solvers optional input functions	46

4.5.7.3	Sparse direct linear solvers optional input functions	47
4.5.7.4	Iterative linear solvers optional input functions	49
4.5.7.5	Initial condition calculation optional input functions	51
4.5.7.6	Rootfinding optional input functions	53
4.5.8	Interpolated output function	54
4.5.9	Optional output functions	54
4.5.9.1	Main solver optional output functions	55
4.5.9.2	Initial condition calculation optional output functions	61
4.5.9.3	Rootfinding optional output functions	62
4.5.9.4	Dense direct linear solvers optional output functions	63
4.5.9.5	Sparse direct linear solvers optional output functions	64
4.5.9.6	Iterative linear solvers optional output functions	65
4.5.10	IDAS reinitialization function	68
4.6	User-supplied functions	69
4.6.1	Residual function	69
4.6.2	Error message handler function	70
4.6.3	Error weight function	70
4.6.4	Rootfinding function	71
4.6.5	Jacobian information (direct method with dense Jacobian)	71
4.6.6	Jacobian information (direct method with banded Jacobian)	72
4.6.7	Jacobian information (direct method with sparse Jacobian)	73
4.6.8	Jacobian information (matrix-vector product)	74
4.6.9	Preconditioning (linear system solution)	75
4.6.10	Preconditioning (Jacobian data)	76
4.7	Integration of pure quadrature equations	77
4.7.1	Quadrature initialization and deallocation functions	78
4.7.2	IDAS solver function	79
4.7.3	Quadrature extraction functions	79
4.7.4	Optional inputs for quadrature integration	80
4.7.5	Optional outputs for quadrature integration	81
4.7.6	User-supplied function for quadrature integration	82
4.8	A parallel band-block-diagonal preconditioner module	83
5	Using IDAS for Forward Sensitivity Analysis	89
5.1	A skeleton of the user's main program	89
5.2	User-callable routines for forward sensitivity analysis	92
5.2.1	Forward sensitivity initialization and deallocation functions	92
5.2.2	Forward sensitivity tolerance specification functions	94
5.2.3	Forward sensitivity initial condition calculation function	95
5.2.4	IDAS solver function	95
5.2.5	Forward sensitivity extraction functions	96
5.2.6	Optional inputs for forward sensitivity analysis	97
5.2.7	Optional outputs for forward sensitivity analysis	99
5.2.7.1	Main solver optional output functions	99
5.2.7.2	Initial condition calculation optional output functions	102
5.3	User-supplied routines for forward sensitivity analysis	103
5.4	Integration of quadrature equations depending on forward sensitivities	103
5.4.1	Sensitivity-dependent quadrature initialization and deallocation	105
5.4.2	IDAS solver function	106
5.4.3	Sensitivity-dependent quadrature extraction functions	106
5.4.4	Optional inputs for sensitivity-dependent quadrature integration	108
5.4.5	Optional outputs for sensitivity-dependent quadrature integration	110
5.4.6	User-supplied function for sensitivity-dependent quadrature integration	111
5.5	Note on using partial error control	112

6	Using IDAS for Adjoint Sensitivity Analysis	115
6.1	A skeleton of the user's main program	115
6.2	User-callable functions for adjoint sensitivity analysis	118
6.2.1	Adjoint sensitivity allocation and deallocation functions	118
6.2.2	Adjoint sensitivity optional input	119
6.2.3	Forward integration function	119
6.2.4	Backward problem initialization functions	120
6.2.5	Tolerance specification functions for backward problem	122
6.2.6	Linear solver initialization functions for backward problem	123
6.2.7	Initial condition calculation functions for backward problem	124
6.2.8	Backward integration function	125
6.2.9	Optional input functions for the backward problem	127
6.2.9.1	Main solver optional input functions	127
6.2.9.2	Dense linear solver	127
6.2.9.3	Band linear solver	128
6.2.9.4	Sparse linear solvers	129
6.2.9.5	SPILS linear solvers	129
6.2.10	Optional output functions for the backward problem	132
6.2.10.1	Main solver optional output functions	132
6.2.10.2	Initial condition calculation optional output function	133
6.2.11	Backward integration of quadrature equations	134
6.2.11.1	Backward quadrature initialization functions	134
6.2.11.2	Backward quadrature extraction function	135
6.2.11.3	Optional input/output functions for backward quadrature integration	136
6.3	User-supplied functions for adjoint sensitivity analysis	136
6.3.1	DAE residual for the backward problem	136
6.3.2	DAE residual for the backward problem depending on the forward sensitivities	137
6.3.3	Quadrature right-hand side for the backward problem	138
6.3.4	Sensitivity-dependent quadrature right-hand side for the backward problem . .	138
6.3.5	Jacobian information for the backward problem (direct method with dense Ja- cobian)	139
6.3.6	Jacobian information for the backward problem (direct method with banded Jacobian)	141
6.3.7	Jacobian information for the backward problem (direct method with sparse Jacobian)	143
6.3.8	Jacobian information for the backward problem (matrix-vector product)	145
6.3.9	Preconditioning for the backward problem (linear system solution)	147
6.3.10	Preconditioning for the backward problem (Jacobian data)	148
6.4	Using the band-block-diagonal preconditioner for backward problems	150
6.4.1	Usage of IDABBDPRE for the backward problem	150
6.4.2	User-supplied functions for IDABBDPRE	151
7	Description of the NVECTOR module	155
7.1	The NVECTOR_SERIAL implementation	159
7.2	The NVECTOR_PARALLEL implementation	161
7.3	The NVECTOR_OPENMP implementation	163
7.4	The NVECTOR_PTHREADS implementation	165
7.5	NVECTOR Examples	167
7.6	NVECTOR functions used by IDAS	168

8	Providing Alternate Linear Solver Modules	169
8.1	Initialization function	170
8.2	Setup function	170
8.3	Solve function	171
8.4	Performance monitoring function	171
8.5	Memory deallocation function	172
9	General Use Linear Solver Components in SUNDIALS	173
9.1	The DLS modules: DENSE and BAND	174
9.1.1	Type DlsMat	174
9.1.2	Accessor macros for the DLS modules	177
9.1.3	Functions in the DENSE module	177
9.1.4	Functions in the BAND module	180
9.2	The SLS module	181
9.2.1	Type SlsMat	182
9.2.2	Functions in the SLS module	183
9.2.3	The KLU solver	185
9.2.4	The SUPERLUMT solver	185
9.3	The SPILS modules: SPGMR, SPFGMR, SPBCG, and SPTFQMR	185
9.3.1	The SPGMR module	186
9.3.2	The SPFGMR module	187
9.3.3	The SPBCG module	187
9.3.4	The SPTFQMR module	187
A	SUNDIALS Package Installation Procedure	189
A.1	CMake-based installation	190
A.1.1	Configuring, building, and installing on Unix-like systems	190
A.1.2	Configuration options (Unix/Linux)	192
A.1.3	Configuration examples	195
A.1.4	Working with external Libraries	195
A.2	Building and Running Examples	196
A.3	Configuring, building, and installing on Windows	196
A.4	Installed libraries and exported header files	197
B	IDAS Constants	201
B.1	IDAS input constants	201
B.2	IDAS output constants	201
	Bibliography	207
	Index	209

List of Tables

4.1	Optional inputs for IDAS, IDADLS, IDASLS, and IDASPILS	41
4.2	Optional outputs from IDAS, IDADLS, IDASLS, and IDASPILS	56
5.1	Forward sensitivity optional inputs	98
5.2	Forward sensitivity optional outputs	100
7.1	Description of the NVECTOR operations	157
7.2	List of vector functions usage by IDAS code modules	168
A.1	SUNDIALS libraries and header files	198
A.2	SUNDIALS libraries and header files (cont.)	199

List of Figures

2.1	Illustration of the checkpointing algorithm for generation of the forward solution during the integration of the adjoint system.	17
3.1	Organization of the SUNDIALS suite	20
3.2	Overall structure diagram of the IDAS package	21
9.1	Diagram of the storage for a banded matrix of type DlsMat	176
9.2	Diagram of the storage for a compressed-sparse-column matrix of type SlsMat	184
A.1	Initial <i>ccmake</i> configuration screen	191
A.2	Changing the <i>installdir</i>	192

Chapter 1

Introduction

IDAS is part of a software family called SUNDIALS: SUite of Nonlinear and Differential/ALgebraic equation Solvers [19]. This suite consists of CVODE, ARKODE, KINSOL, and IDA, and variants of these with sensitivity analysis capabilities, CVODES and IDAS.

IDAS is a general purpose solver for the initial value problem (IVP) for systems of differential-algebraic equations (DAEs). The name IDAS stands for Implicit Differential-Algebraic solver with Sensitivity capabilities. IDAS is an extension of the IDA solver within SUNDIALS, itself based on DASPK [5, 6]; however, like all SUNDIALS solvers, IDAS 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; (2) it is written in a *data-independent* manner in that it acts on generic vectors without any assumptions on the underlying organization of the data; and (3) it provides a flexible, extensible framework for sensitivity analysis, using either *forward* or *adjoint* methods. Thus IDAS shares significant modules previously written within CASC at LLNL to support the ordinary differential equation (ODE) solvers CVODE [20, 12] and PVODE [8, 9], the DAE solver IDA [22] on which IDAS is based, the sensitivity-enabled ODE solver CVODES [21, 29], and also the nonlinear system solver KINSOL [13].

The Newton/Krylov methods in IDAS are: the GMRES (Generalized Minimal RESidual) [28], Bi-CGStab (Bi-Conjugate Gradient Stabilized) [31], and TFQMR (Transpose-Free Quasi-Minimal Residual) linear iterative methods [17]. As Krylov methods, these require almost no matrix storage for solving the Newton equations as compared to direct methods. However, the algorithms allow for a user-supplied preconditioner matrix, 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 three Krylov methods in IDAS, we recommend GMRES as the best overall choice. However, users are encouraged to compare all three, especially if encountering convergence failures with GMRES. Bi-CGStab and TFQMR have an advantage in storage requirements, in that the number of workspace vectors they require is fixed, while that number for GMRES depends on the desired Krylov subspace size.

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

There are several motivations for choosing the C language for IDAS. First, a general movement away from FORTRAN and toward C in scientific computing was apparent. Second, the pointer, structure,

and dynamic memory allocation features in C are extremely useful in software of this complexity, with the great variety of method options offered. Finally, we prefer C over C++ for IDAS because of the wider availability of C compilers, the potentially greater efficiency of C, and the greater ease of interfacing the solver to applications written in extended FORTRAN.

1.1 Changes from previous versions

Changes in v1.2.0

Two major additions were made to the linear system solvers that are available for use with the IDAS 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 NVECTOR module. As part of these additions, a sparse matrix (CSC format) structure was added to IDAS.

Otherwise, only relatively minor modifications were made to IDAS:

In `IDARootfind`, a minor bug was corrected, where the input array `rootdir` was ignored, and a line was added to break out of root-search loop if the initial interval size is below the tolerance `ttol`.

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

An option was added in the case of Adjoint Sensitivity Analysis with dense or banded Jacobian: With a call to `IDADlsSetDenseJacFnBS` or `IDADlsSetBandJacFnBS`, the user can specify a user-supplied Jacobian function of type `IDADls***JacFnBS`, for the case where the backward problem depends on the forward sensitivities.

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

For the Adjoint Sensitivity Analysis case in which the backward problem depends on the forward sensitivities, options have been added to allow for user-supplied `pset`, `psolve`, and `jtimes` functions.

In 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`, `SUNRpowerI`, and `SUNRpowerR`, respectively. These names occur in both the solver and in various example programs.

In the User Guide, a paragraph was added in Section 6.2.1 on `IDAAAdjReInit`, and a paragraph was added in Section 6.2.9 on `IDAGetAdjY`.

Two new NVECTOR 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.

Changes in v1.1.0

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/NewLintArray`, for `int` and `long int` arrays, respectively. In a minor change to the user interface, the type of the index `which` in IDAS was changed from `long int` to `int`.

Errors in the logic for the integration of backward problems were identified and fixed.

A large number of minor errors have been fixed. Among these are the following: A missing vector pointer setting was added in `IDASensLineSrch`. In `IDACompleteStep`, conditionals around lines loading a new column of three auxiliary divided difference arrays, for a possible order increase, were fixed. After the solver memory is created, it is set to zero before being filled. In each linear solver interface function, the linear solver memory is freed on an error return, and the `**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`/`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.

1.2 Reading this User Guide

The structure of this document is as follows:

- In Chapter 2, we give short descriptions of the numerical methods implemented by IDAS for the solution of initial value problems for systems of DAEs, continue with short descriptions of preconditioning (§2.2) and rootfinding (§2.3), and then give an overview of the mathematical aspects of sensitivity analysis, both forward (§2.5) and adjoint (§2.6).
- The following chapter describes the structure of the SUNDIALS suite of solvers (§3.1) and the software organization of the IDAS solver (§3.2).
- Chapter 4 is the main usage document for IDAS for simulation applications. It includes a complete description of the user interface for the integration of DAE initial value problems. Readers that are not interested in using IDAS for sensitivity analysis can then skip the next two chapters.
- Chapter 5 describes the usage of IDAS for forward sensitivity analysis as an extension of its IVP integration capabilities. We begin with a skeleton of the user main program, with emphasis on the steps that are required in addition to those already described in Chapter 4. Following that we provide detailed descriptions of the user-callable interface routines specific to forward sensitivity analysis and of the additional optional user-defined routines.
- Chapter 6 describes the usage of IDAS for adjoint sensitivity analysis. We begin by describing the IDAS checkpointing implementation for interpolation of the original IVP solution during integration of the adjoint system backward in time, and with an overview of a user's main program. Following that we provide complete descriptions of the user-callable interface routines for adjoint sensitivity analysis as well as descriptions of the required additional user-defined routines.
- Chapter 7 gives a brief overview of the generic NVECTOR module shared amongst the various components of SUNDIALS, as well as details on the NVECTOR implementations provided with SUNDIALS: a serial implementation (§7.1), a distributed memory parallel implementation based on MPI (§7.2), and two thread-parallel implementations based on openMP (§7.3) and Pthreads (§7.4), respectively.
- Chapter 8 describes the specifications of linear solver modules as supplied by the user.
- Chapter 9 describes in detail the generic linear solvers shared by all SUNDIALS solvers.
- Finally, in the appendices, we provide detailed instructions for the installation of IDAS, within the structure of SUNDIALS (Appendix A), as well as a list of all the constants used for input to and output from IDAS functions (Appendix B).

The reader should be aware of the following notational conventions in this user guide: program listings and identifiers (such as `IDAInit`) within textual explanations appear in typewriter type style; fields in C structures (such as *content*) appear in italics; and packages or modules, such as `IDADENSE`, are written in all capitals. Usage and installation instructions that constitute important warnings are marked with a triangular symbol in the margin.



Chapter 2

Mathematical Considerations

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

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

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

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

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

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

2.1 IVP solution

Prior to integrating a DAE initial-value problem, an important requirement is that the pair of vectors y_0 and \dot{y}_0 are both initialized to satisfy the DAE residual $F(t_0, y_0, \dot{y}_0) = 0$. For a class of problems that includes so-called semi-explicit index-one systems, IDAS provides a routine that computes consistent initial conditions from a user's initial guess [6]. For this, the user must identify sub-vectors of y (not necessarily contiguous), denoted y_d and y_a , which are its differential and algebraic parts, respectively, such that F depends on \dot{y}_d but not on any components of \dot{y}_a . The assumption that the system is “index one” means that for a given t and y_d , the system $F(t, y, \dot{y}) = 0$ defines y_a uniquely. In this case, a solver within IDAS computes y_a and \dot{y}_d at $t = t_0$, given y_d and an initial guess for y_a . A second available option with this solver also computes all of $y(t_0)$ given $\dot{y}(t_0)$; this is intended mainly for quasi-steady-state problems, where $\dot{y}(t_0) = 0$ is given. In both cases, IDAS solves the system $F(t_0, y_0, \dot{y}_0) = 0$ for the unknown components of y_0 and \dot{y}_0 , using Newton iteration augmented with a line search global 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.

The integration method used in IDAS is the variable-order, variable-coefficient BDF (Backward Differentiation Formula), in fixed-leading-coefficient form [3]. The method order ranges from 1 to 5, with the BDF of order q given by the multistep formula

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

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

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

Regardless of the method options, the solution of the nonlinear system (2.4) is accomplished with some form of Newton iteration. This leads to a linear system for each Newton correction, 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 corrections, IDAS provides several choices, including the option of an user-supplied linear solver module. The linear solver modules distributed with SUNDIALS are organized in three families, a *direct* family comprising direct linear solvers for dense or banded matrices, a *sparse* family comprising direct linear solvers for matrices stored in compressed-sparse-column format, and a *spils* family comprising scaled preconditioned iterative (Krylov) linear solvers. The methods offered through these modules are as follows:

- dense direct solvers, using either an internal implementation or a Blas/Lapack implementation (serial or threaded vector modules only),
- band direct solvers, using either an internal implementation or a Blas/Lapack implementation (serial or threaded vector modules only),
- sparse direct solver interfaces, using either the KLU sparse solver library [14, 1], or the thread-enabled SuperLU_MT sparse solver library [25, 15, 2] (serial or threaded vector modules only) [Note that users will need to download and install the KLU or SuperLU_MT packages independent of IDAS],
- SPGMR, a scaled preconditioned GMRES (Generalized Minimal Residual method) solver without restarts,
- SPBCG, a scaled preconditioned Bi-CGStab (Bi-Conjugate Gradient Stable method) solver, or
- SPTFQMR, a scaled preconditioned TFQMR (Transpose-Free Quasi-Minimal Residual method) solver.

For large stiff systems, where direct methods are not feasible, the combination of a BDF integrator and any of the preconditioned Krylov methods (SPGMR, SPBCG, or SPTFQMR) 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 [4]. For the *spils* linear solvers, preconditioning is allowed only on the left (see §2.2). Note that the direct linear solvers (dense, band, and sparse) can only be used with serial or threaded vector representations.

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

$$W_i = 1/[\text{RTOL} \cdot |y_i| + \text{ATOL}_i]. \quad (2.7)$$

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.

In the case of a direct linear solver (dense, band, or sparse), the nonlinear iteration (2.5) is a Modified Newton iteration, in that the Jacobian J is fixed (and usually out of date), with a coefficient $\bar{\alpha}$ in place of α in J . When using one of the Krylov methods SPGMR, SPBCG, or SPTFQMR as the linear solver, the iteration is an Inexact Newton iteration, using the current Jacobian (through matrix-free products Jv), in which the linear residual $J\Delta y + G$ is nonzero but controlled. The Jacobian matrix J (direct cases) or preconditioner matrix P (SPGMR/SPBCG/SPTFQMR case) is updated 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 stopping test for the Newton iteration in IDAS ensures that the iteration error $y_n - y_{n(m)}$ is small relative to y itself. For this, we estimate the linear convergence rate at all iterations $m > 1$ as

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

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

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

where $S = R/(R-1)$ whenever $m > 1$ and $R \leq 0.9$. The user has the option of changing the constant in the convergence test from its default value of 0.33. The quantity S is set to $S = 20$ initially and whenever J or P is updated, and it is reset to $S = 100$ on a step with $\alpha \neq \bar{\alpha}$. Note that at $m = 1$, the convergence test (2.8) uses an old value for S . Therefore, at the first Newton 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 Newton 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 $h_n/4$. The integration is halted after a preset number (default value 10) of convergence failures. Both the maximum allowable Newton iterations and the maximum nonlinear convergence failures can be changed by the user from their default values.

When SPGMR, SPBCG, or SPTFQMR 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 Newton iteration, i.e., $\|P^{-1}(Jx + G)\| < 0.05 \cdot 0.33$. The safety factor 0.05 can be changed by the user.

In the direct linear solver cases, the Jacobian J defined in (2.6) can be either supplied by the user or have IDAS compute one 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 (2.7). In the SPGMR/SPBCG/SPTFQMR case, 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 σ is $1/\|v\|$. As an option, the user can specify a constant factor that is inserted into this expression for σ .

We note that with the sparse direct solvers, the Jacobian *must* be supplied by a user routine in compressed-sparse-column format.

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

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

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

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

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

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

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

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

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

where $\phi(k)$ is a modified divided difference of order k that is retained by IDAS (and behaves asymptotically as h^k). Thus the local truncation errors are estimated as $\text{ELTE}(q') = C(q')\|\phi(q'+1)\|$ to select step sizes. But the choice of order in IDAS is based on the requirement that the scaled derivative norms, $\|h^k y^{(k)}\|$, are monotonically decreasing with k , for k near q . These norms are again estimated using the $\phi(k)$, and in fact

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

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

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

The value of η is adjusted so that $0.25 \leq \eta \leq 0.9$ before setting $h \leftarrow h' = \eta h$. If the local error test fails a second time, IDAS uses $\eta = 0.25$, and on the third and subsequent failures it uses $q = 1$ and $\eta = 0.25$. After 10 failures, IDAS 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 such change is made if $q' = q-1$ from the prior test, if $q = 5$, or if q was increased on the previous step. Otherwise, if the last $q+1$ steps were taken at a constant order $q < 5$ and a constant step size, IDAS considers raising the order to $q+1$. The logic is as follows: (a) If $q = 1$, then reset $q = 2$ if $T(2) < T(1)/2$. (b) If $q > 1$ then

- reset $q \leftarrow q-1$ if $T(q-1) \leq \min\{T(q), T(q+1)\}$;

- else reset $q \leftarrow q + 1$ if $T(q + 1) < T(q)$;
- leave q unchanged otherwise [then $T(q - 1) > T(q) \leq T(q + 1)$].

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 $\eta > 2$, η is reset to 2; (b) if $\eta \leq 1$, η is restricted to $0.5 \leq \eta \leq 0.9$; and (c) if $1 < \eta < 2$ we use $\eta = 1$. Finally h is reset to $h' = \eta h$. Thus we do not increase the step size unless it can be doubled. See [3] for details.

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

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

2.2 Preconditioning

When using a Newton method to solve the nonlinear system (2.5), IDAS makes repeated use of a linear solver to solve linear systems of the form $J\Delta y = -G$. If this linear system solve is done with one of the scaled preconditioned iterative linear solvers, these solvers are rarely successful if used without preconditioning; it is generally necessary to precondition the system in order to obtain acceptable efficiency. A system $Ax = b$ can be preconditioned on the left, on the right, or on both sides. The Krylov method is then applied to a system with the matrix $P^{-1}A$, or AP^{-1} , or $P_L^{-1}AP_R^{-1}$, instead of A . However, within IDAS, preconditioning is allowed *only* on the left, so that the iterative method is applied to systems $(P^{-1}J)\Delta y = -P^{-1}G$. Left preconditioning is required to make the norm of the linear residual in the Newton 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 [4] for an extensive study of preconditioners for reaction-transport systems).

Typical preconditioners used with IDAS are based on approximations to the Newton iteration matrix of the systems involved; in other words, $P \approx \frac{\partial F}{\partial y} + \alpha \frac{\partial F}{\partial y}$, where α is a scalar inversely proportional to the integration step size h . Because the Krylov iteration occurs within a Newton 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.

2.3 Rootfinding

The IDAS solver has been augmented to include a rootfinding feature. This means that, while integrating the Initial Value Problem (2.1), IDAS can also find the roots of a set of user-defined functions

$g_i(t, y, \dot{y})$ that depend on t , the solution vector $y = y(t)$, and its t -derivative $\dot{y}(t)$. The number of these root functions is arbitrary, and if more than one g_i is found to have a root in any given interval, the various root locations are found and reported in the order that they occur on the t axis, in the direction of integration.

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

The basic scheme used is to check for sign changes of any $g_i(t)$ over each time step taken, and then (when a sign change is found) to home in on the root (or roots) with a modified secant method [18]. In addition, each time g is computed, IDAS checks to see if $g_i(t) = 0$ exactly, and if so it reports this as a root. However, if an exact zero of any g_i is found at a point t , IDAS computes g at $t + \delta$ for a small increment δ , slightly further in the direction of integration, and if any $g_i(t + \delta) = 0$ also, IDAS stops and reports an error. This way, each time IDAS takes a time step, it is guaranteed that the values of all g_i are nonzero at some past value of t , beyond which a search for roots is to be done.

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

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

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

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

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

where 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 .1 and .5 (.5 being the midpoint), and the actual distance from the endpoint is at least $\tau/2$.

2.4 Pure quadrature integration

In many applications, and most notably during the backward integration phase of an adjoint sensitivity analysis run (see §2.6) it is of interest to compute integral quantities of the form

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

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

$$\dot{z} = q(t, y, \dot{y}, p), \quad z(t_0) = 0. \quad (2.11)$$

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

This can be done at the “user level” by simply exposing to IDAS the extended DAE system (2.2)+(2.10). However, in the context of an implicit integration solver, this approach is not desirable since the nonlinear solver module will require the Jacobian (or Jacobian-vector product) of this extended DAE. Moreover, since the additional states z do not enter the right-hand side of the ODE (2.10) and therefore the residual of the extended DAE system does not depend on z , it is much more efficient to treat the ODE system (2.10) separately from the original DAE system (2.2) by “taking out” the additional states z from the nonlinear system (2.4) that must be solved in the correction step of the LMM. Instead, “corrected” values z_n are computed explicitly as

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

once the new approximation y_n is available.

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

2.5 Forward sensitivity analysis

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

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

$$\begin{aligned} \frac{\partial F}{\partial y} s_i + \frac{\partial F}{\partial \dot{y}} \dot{s}_i + \frac{\partial F}{\partial p_i} &= 0 \\ s_i(t_0) &= \frac{\partial y_0(p)}{\partial p_i}, \quad \dot{s}_i(t_0) = \frac{\partial \dot{y}_0(p)}{\partial p_i}, \end{aligned} \quad (2.12)$$

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

When performing forward sensitivity analysis, IDAS carries out the time integration of the combined system, (2.2) and (2.12), by viewing it as a DAE system of size $N(N_s + 1)$, where N_s is the number of model parameters p_i , with respect to which sensitivities are desired ($N_s \leq N_p$). However, major improvements in efficiency can be made by taking advantage of the special form of the sensitivity equations as linearizations of the original DAEs. In particular, the original DAE system and all sensitivity systems share the same Jacobian matrix J in (2.6).

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

2.5.1 Forward sensitivity methods

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

- *Staggered Direct* In this approach [11], the nonlinear system (2.4) is first solved and, once an acceptable numerical solution is obtained, the sensitivity variables at the new step are found by directly solving (2.12) after the BDF discretization is used to eliminate \dot{s}_i . Although the system matrix of the above linear system is based on exactly the same information as the matrix J in (2.6), it must be updated and factored at every step of the integration, in contrast to an evaluation of J which is updated only occasionally. For problems with many parameters (relative to the problem size), the staggered direct method can outperform the methods described below [24]. However, the computational cost associated with matrix updates and factorizations makes this method unattractive for problems with many more states than parameters (such as those arising from semidiscretization of PDEs) and is therefore not implemented in IDAS.
- *Simultaneous Corrector* In this method [26], the discretization is applied simultaneously to both the original equations (2.2) and the sensitivity systems (2.12) resulting in an “extended” nonlinear system $\hat{G}(\hat{y}_n) = 0$ where $\hat{y}_n = [y_n, \dots, s_i, \dots]$. This combined nonlinear system can be solved using a modified Newton method as in (2.5) by solving the corrector equation

$$\hat{J}[\hat{y}_{n(m+1)} - \hat{y}_{n(m)}] = -\hat{G}(\hat{y}_{n(m)}) \quad (2.13)$$

at each iteration, where

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

J is defined as in (2.6), and $J_i = (\partial/\partial y)[F_y s_i + F_{\dot{y}} \dot{s}_i + F_{p_i}]$. It can be shown that 2-step quadratic convergence can be retained by using only the block-diagonal portion of \hat{J} in the corrector equation (2.13). This results in a decoupling that allows the reuse of J without additional matrix factorizations. However, the sum $F_y s_i + F_{\dot{y}} \dot{s}_i + F_{p_i}$ must still be reevaluated at each step of the iterative process (2.13) to update the sensitivity portions of the residual \hat{G} .

- *Staggered corrector* In this approach [16], as in the staggered direct method, the nonlinear system (2.4) is solved first using the Newton iteration (2.5). Then, for each sensitivity vector $\xi \equiv s_i$, a separate Newton iteration is used to solve the sensitivity system (2.12):

$$\begin{aligned} J[\xi_{n(m+1)} - \xi_{n(m)}] = \\ - \left[F_y(t_n, y_n, \dot{y}_n) \xi_{n(m)} + F_{\dot{y}}(t_n, y_n, \dot{y}_n) \cdot h_n^{-1} \left(\alpha_{n,0} \xi_{n(m)} + \sum_{i=1}^q \alpha_{n,i} \xi_{n-i} \right) + F_{p_i}(t_n, y_n, \dot{y}_n) \right]. \end{aligned} \quad (2.14)$$

In other words, a modified Newton iteration is used to solve a linear system. In this approach, the matrices $\partial F/\partial y$, $\partial F/\partial \dot{y}$ and vectors $\partial F/\partial p_i$ need be updated only once per integration step, after the state correction phase (2.5) has converged.

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

An important observation is that the staggered corrector method, combined with a Krylov linear solver, effectively results in a staggered direct method. Indeed, the Krylov solver requires only the action of the matrix J on a vector and this can be provided with the current Jacobian information. Therefore, the modified Newton procedure (2.14) will theoretically converge after one iteration.

2.5.2 Selection of the absolute tolerances for sensitivity variables

If the sensitivities are included in the error test, IDAS provides an automated estimation of absolute tolerances for the sensitivity variables based on the absolute tolerance for the corresponding state variable. The relative tolerance for sensitivity variables is set to be the same as for the state variables.

The selection of absolute tolerances for the sensitivity variables is based on the observation that the sensitivity vector s_i will have units of $[y]/[p_i]$. With this, the absolute tolerance for the j -th component of the sensitivity vector s_i is set to $\text{ATOL}_j/|\bar{p}_i|$, where ATOL_j are the absolute tolerances for the state variables and \bar{p} is a vector of scaling factors that are dimensionally consistent with the model parameters p and give an indication of their order of magnitude. This choice of relative and absolute tolerances is equivalent to requiring that the weighted root-mean-square norm of the sensitivity vector s_i with weights based on s_i be the same as the weighted root-mean-square norm of the vector of scaled sensitivities $\bar{s}_i = |\bar{p}_i|s_i$ with weights based on the state variables (the scaled sensitivities \bar{s}_i being dimensionally consistent with the state variables). However, this choice of tolerances for the s_i may be a poor one, and the user of IDAS can provide different values as an option.

2.5.3 Evaluation of the sensitivity right-hand side

There are several methods for evaluating the residual functions in the sensitivity systems (2.12): analytic evaluation, automatic differentiation, complex-step approximation, and finite differences (or directional derivatives). IDAS provides all the software hooks for implementing interfaces to automatic differentiation (AD) or complex-step approximation; future versions will include a generic interface to AD-generated functions. At the present time, besides the option for analytical sensitivity right-hand sides (user-provided), IDAS can evaluate these quantities using various finite difference-based approximations to evaluate the terms $(\partial F/\partial y)s_i + (\partial F/\partial \dot{y})\dot{s}_i$ and $(\partial F/\partial p_i)$, or using directional derivatives to evaluate $[(\partial F/\partial y)s_i + (\partial F/\partial \dot{y})\dot{s}_i + (\partial F/\partial p_i)]$. As is typical for finite differences, the proper choice of perturbations is a delicate matter. IDAS takes into account several problem-related features: the relative DAE error tolerance RTOL , the machine unit roundoff U , the scale factor \bar{p}_i , and the weighted root-mean-square norm of the sensitivity vector s_i .

Using central finite differences as an example, the two terms $(\partial F/\partial y)s_i + (\partial F/\partial \dot{y})\dot{s}_i$ and $\partial F/\partial p_i$ in (2.12) can be evaluated either separately:

$$\frac{\partial F}{\partial y}s_i + \frac{\partial F}{\partial \dot{y}}\dot{s}_i \approx \frac{F(t, y + \sigma_y s_i, \dot{y} + \sigma_y \dot{s}_i, p) - F(t, y - \sigma_y s_i, \dot{y} - \sigma_y \dot{s}_i, p)}{2\sigma_y}, \quad (2.15)$$

$$\frac{\partial F}{\partial p_i} \approx \frac{F(t, y, \dot{y}, p + \sigma_i e_i) - F(t, y, \dot{y}, p - \sigma_i e_i)}{2\sigma_i}, \quad (2.15')$$

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

or simultaneously:

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

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

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

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

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

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

2.5.4 Quadratures depending on forward sensitivities

If pure quadrature variables are also included in the problem definition (see §2.4), IDAS does *not* carry their sensitivities automatically. Instead, we provide a more general feature through which integrals depending on both the states y of (2.2) and the state sensitivities s_i of (2.12) can be evaluated. In other words, IDAS provides support for computing integrals of the form:

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

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

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

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

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

2.6 Adjoint sensitivity analysis

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

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

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

In what follows, we only sketch the analysis for the sensitivity problem for both G and g . For details on the derivation see [10].

2.6.1 Sensitivity of $G(p)$

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

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

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

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

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

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

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

$$\frac{dG}{dp} = \int_{t_0}^T (g_p - \lambda^* F_p) dt - \int_{t_0}^T [-g_y + \lambda^* F_y - (\lambda^* F_{\dot{y}})'] y_p dt - (\lambda^* F_{\dot{y}} y_p)|_{t_0}^T. \quad (2.19)$$

Now by requiring λ to satisfy

$$(\lambda^* F_{\dot{y}})' - \lambda^* F_y = -g_y, \quad (2.20)$$

we obtain

$$\frac{dG}{dp} = \int_{t_0}^T (g_p - \lambda^* F_p) dt - (\lambda^* F_{\dot{y}} y_p)|_{t_0}^T. \quad (2.21)$$

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

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

$$\lambda^* F_{\dot{y}}|_{t=T} = 0, \quad (2.22)$$

yielding the sensitivity equation for dG/dp

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

This choice will not suffice for a Hessenberg index-2 DAE system. For a derivation of proper final conditions in such cases, see [10].

The first thing to notice about the adjoint system (2.20) is that there is no explicit specification of the parameters p ; this implies that, once the solution λ is found, the formula (2.21) can then be used to find the gradient of G with respect to any of the parameters p . The second important remark is that the adjoint system (2.20) is a terminal value problem which depends on the solution $y(t)$ of the original IVP (2.2). Therefore, a procedure is needed for providing the states y obtained during a forward integration phase of (2.2) to IDAS during the backward integration phase of (2.20). The approach adopted in IDAS, based on *checkpointing*, is described in §2.6.3 below.

2.6.2 Sensitivity of $g(T, p)$

Now let us consider the computation of $dg/dp(T)$. From $dg/dp(T) = (d/dT)(dG/dp)$ and equation (2.21), we have

$$\frac{dg}{dp} = (g_p - \lambda^* F_p)(T) - \int_{t_0}^T \lambda_T^* F_p dt + (\lambda_T^* F_{\dot{y}} y_p)|_{t=t_0} - \frac{d(\lambda^* F_{\dot{y}} y_p)}{dT} \quad (2.24)$$

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

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

while for a Hessenberg index-2 DAE system we have

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

The corresponding adjoint equations are

$$(\lambda_T^* F_{\dot{y}})' - \lambda_T^* F_y = 0. \quad (2.25)$$

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

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

Taking the total derivative, we obtain

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

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

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

For the index-one DAE case, the above relation and (2.20) yield

$$(\lambda_T^*F_{\dot{y}})|_{t=T} = [g_y - \lambda^*F_y]|_{t=T}. \quad (2.26)$$

For the regular implicit ODE case, $F_{\dot{y}}$ is invertible; thus we have $\lambda(T, T) = 0$, which leads to $\lambda_T(T) = -\dot{\lambda}(T)$. As with the final conditions for $\lambda(T)$ in (2.20), the above selection for $\lambda_T(T)$ is not sufficient for index-two Hessenberg DAEs (see [10] for details).

2.6.3 Checkpointing scheme

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

However, especially for large-scale problems and long integration intervals, the number and size of the vectors y and \dot{y} that would need to be stored make this approach computationally intractable. Thus, IDAS settles for a compromise between storage space and execution time by implementing a so-called *checkpointing scheme*. At the cost of at most one additional forward integration, this approach offers the best possible estimate of memory requirements for adjoint sensitivity analysis. To begin with, based on the problem size N and the available memory, the user decides on the number N_d of data pairs (y, \dot{y}) if cubic Hermite interpolation is selected, or on the number N_d of y vectors in the case of variable-degree polynomial interpolation, that can be kept in memory for the purpose of interpolation. Then, during the first forward integration stage, after every N_d integration steps a checkpoint is formed by saving enough information (either in memory or on disk) to allow for a hot restart, that is a restart which will exactly reproduce the forward integration. In order to avoid storing Jacobian-related data at each checkpoint, a reevaluation of the iteration matrix is forced before each checkpoint. At the end of this stage, we are left with N_c checkpoints, including one at t_0 . During the backward integration stage, the adjoint variables are integrated backwards from T to t_0 , going from one checkpoint to the previous one. The backward integration from checkpoint $i + 1$ to checkpoint i is preceded by a forward integration from i to $i + 1$ during which the N_d vectors y (and, if necessary \dot{y}) are generated and stored in memory for interpolation¹

This approach transfers the uncertainty in the number of integration steps in the forward integration phase to uncertainty in the final number of checkpoints. However, N_c is much smaller than the number of steps taken during the forward integration, and there is no major penalty for writing/reading the checkpoint data to/from a temporary file. Note that, at the end of the first forward

¹The degree of the interpolation polynomial is always that of the current BDF order for the forward interpolation at the first point to the right of the time at which the interpolated value is sought (unless too close to the i -th checkpoint, in which case it uses the BDF order at the right-most relevant point). However, because of the FLC BDF implementation (see §2.1), the resulting interpolation polynomial is only an approximation to the underlying BDF interpolant.

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




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

integration stage, interpolation data are available from the last checkpoint to the end of the interval of integration. If no checkpoints are necessary (N_d is larger than the number of integration steps taken in the solution of (2.2)), the total cost of an adjoint sensitivity computation can be as low as one forward plus one backward integration. In addition, IDAS provides the capability of reusing a set of checkpoints for multiple backward integrations, thus allowing for efficient computation of gradients of several functionals (2.17).

Finally, we note that the adjoint sensitivity module in IDAS provides the necessary infrastructure to integrate backwards in time any DAE terminal value problem dependent on the solution of the IVP (2.2), including adjoint systems (2.20) or (2.25), as well as any other quadrature ODEs that may be needed in evaluating the integrals in (2.21). In particular, for DAE systems arising from semi-discretization of time-dependent PDEs, this feature allows for integration of either the discretized adjoint PDE system or the adjoint of the discretized PDE.

2.7 Second-order sensitivity analysis

In some applications (e.g., dynamically-constrained optimization) it may be desirable to compute second-order derivative information. Considering the DAE problem (2.2) and some model output functional² $g(y)$, the Hessian d^2g/dp^2 can be obtained in a forward sensitivity analysis setting as

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

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

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

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

where y_p denotes the first-order sensitivity matrix, the solution of N_p systems (2.12), and y_{pp} is a third-order tensor. It is easy to see that, except for situations in which the number of parameters N_p is very small, the computational cost of this so-called *forward-over-forward* approach is exorbitant as it requires the solution of $N_p + N_p^2$ additional DAE systems of the same dimension as (2.2).

A much more efficient alternative is to compute Hessian-vector products using a so-called *forward-over-adjoint* approach. This method is based on using the same “trick” as the one used in computing gradients of pointwise functionals with the adjoint method, namely applying a formal directional forward derivation to the gradient of (2.21) (or the equivalent one for a pointwise functional $g(T, y(T))$).

²For the sake of simplicity in presentation, we do not include explicit dependencies of g on time t or parameters p . Moreover, we only consider the case in which the dependency of the original DAE (2.2) on the parameters p is through its initial conditions only. For details on the derivation in the general case, see [27].

With that, the cost of computing a full Hessian is roughly equivalent to the cost of computing the gradient with forward sensitivity analysis. However, Hessian-vector products can be cheaply computed with one additional adjoint solve.

As an illustration³, consider the ODE problem

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

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

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

where λ and μ are solutions of

$$\begin{aligned} -\dot{\mu} &= f_y^T \mu + (\lambda^T \otimes I_n) f_{yy} s; & \mu(t_f) &= 0 \\ -\dot{\lambda} &= f_y^T \lambda + g_y^T; & \lambda(t_f) &= 0 \\ \dot{s} &= f_y s; & s(t_0) &= y_{0p} u. \end{aligned} \tag{2.27}$$

In the above equation, $s = y_p u$ is a linear combination of the columns of the sensitivity matrix y_p . The *forward-over-adjoint* approach hinges crucially on the fact that s can be computed at the cost of a forward sensitivity analysis with respect to a single parameter (the last ODE problem above) which is possible due to the linearity of the forward sensitivity equations (2.12).

Therefore (and this is also valid for the DAE case), the cost of computing the Hessian-vector product is roughly that of two forward and two backward integrations of a system of DAEs of size N . For more details, including the corresponding formulas for a pointwise model functional output, see the work by Ozyurt and Barton [27] who discuss this problem for ODE initial value problems. As far as we know, there is no published equivalent work on DAE problems. However, the derivations given in [27] for ODE problems can be extended to DAEs with some careful consideration given to the derivation of proper final conditions on the adjoint systems, following the ideas presented in [10].

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

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

³The derivation for the general DAE case is too involved for the purposes of this discussion.

Chapter 3

Code Organization

3.1 SUNDIALS organization

The family of solvers referred to as SUNDIALS consists of the solvers CVODE and ARKODE (for ODE systems), KINSOL (for nonlinear algebraic systems), and IDA (for differential-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 various solvers of this family share many subordinate modules. For this reason, it is organized as a family, with a directory structure that exploits that sharing (see Fig. 3.1). The following is a list of the solver packages presently available, and the basic functionality of each:


- CVODE, a solver for stiff and nonstiff ODE systems $dy/dt = f(t, y)$ based on Adams and BDF methods;
- CVODES, a solver for stiff and nonstiff ODE systems with sensitivity analysis capabilities;
- ARKODE, a solver for ODE systems $Mdy/dt = f(t, y)$ based on additive 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 $F(u) = 0$.

3.2 IDAS organization

The IDAS package is written in the ANSI C language. 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 IDAS package is shown in Figure 3.2. The central integration module, implemented in the files `idas.h`, `idas_impl.h`, and `idas.c`, deals with the evaluation of integration coefficients, the Newton iteration process, estimation of local error, selection of stepsize and order, and interpolation to user output points, among other issues. Although this module contains logic for the basic Newton iteration algorithm, it has no knowledge of the method being used to solve the linear systems that arise. For any given user problem, one of the linear system modules is specified, and is then invoked as needed during the integration.


In addition, if forward sensitivity analysis is turned on, the main module will integrate the forward sensitivity equations simultaneously with the original IVP. The sensitivity variables may be included in the local error control mechanism of the main integrator. IDAS provides two different strategies for dealing with the correction stage for the sensitivity variables: `IDA_SIMULTANEOUS` and `IDA_STAGGERED` (see §2.5). The IDAS package includes an algorithm for the approximation of the sensitivity equations residuals by difference quotients, but the user has the option of supplying these residual functions directly.



(a) High-level diagram (note that none of the Lapack-based linear solver modules are represented.)

* only applies to ARKODE

** only applies to ARKODE and KINSOL



(b) Directory structure of the source tree

Figure 3.1: Organization of the SUNDIALS suite

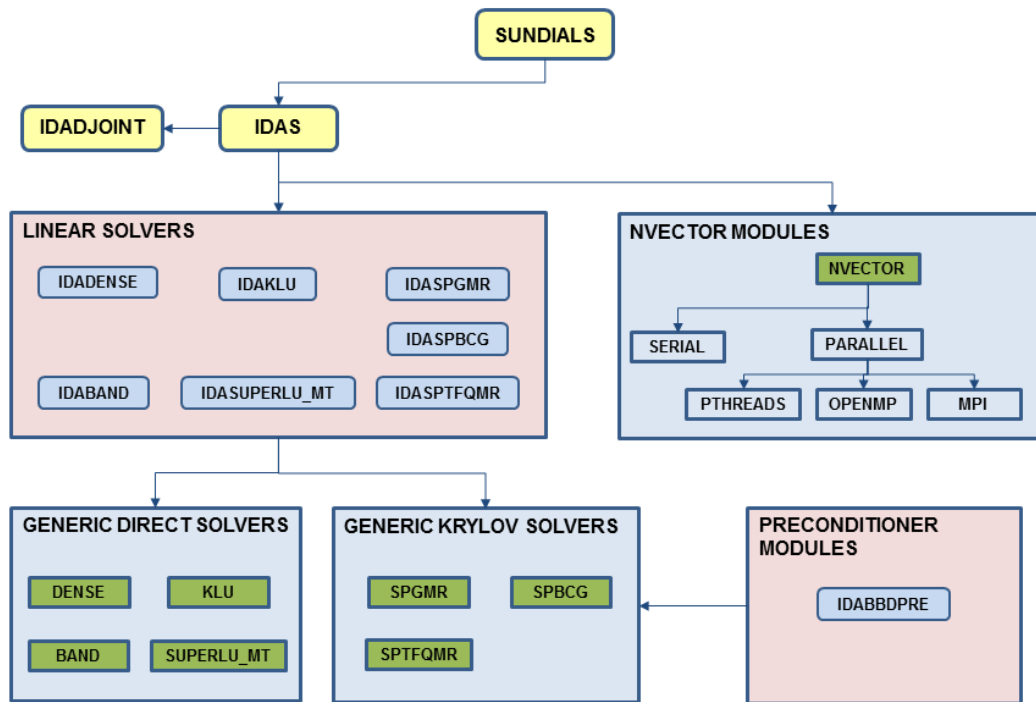


Figure 3.2: Overall structure diagram of the IDAS package. Modules specific to IDAS are distinguished by rounded boxes, while generic solver and auxiliary modules are in square boxes. Note that the direct linear solvers using Lapack implementations are not explicitly represented. Note also that the KLU and SuperLU_MT support is through interfaces to packages. Users will need to download and compile those packages independently.

The adjoint sensitivity module (file `idaa.c`) provides the infrastructure needed for the backward integration of any system of DAEs which depends on the solution of the original IVP, in particular the adjoint system and any quadratures required in evaluating the gradient of the objective functional. This module deals with the setup of the checkpoints, the interpolation of the forward solution during the backward integration, and the backward integration of the adjoint equations.

At present, the package includes the following seven IDAS linear algebra modules, organized into two families. The *direct* family of linear solvers provides solvers for the direct solution of linear systems with dense or banded matrices and includes:

- IDADENSE: LU factorization and backsolving with dense matrices (using either an internal implementation or Blas/Lapack);
- IDABAND: LU factorization and backsolving with banded matrices (using either an internal implementation or Blas/Lapack);
- IDAKLU: LU factorization and backsolving with compressed-sparse-column (CSC) matrices using the KLU linear solver library [14, 1] (KLU to be downloaded and compiled by user independent of IDA);
- IDASUPERLUMT: LU factorization and backsolving with compressed-sparse-column (CSC) matrices using the threaded SuperLU_MT linear solver library [25, 15, 2] (SuperLU_MT to be downloaded and compiled by user independent of IDA).

The *spils* family of linear solvers provides scaled preconditioned iterative linear solvers and includes:

- IDASPGMR: scaled preconditioned GMRES method;
- IDASPCBG: scaled preconditioned Bi-CGStab method;
- IDASPTFQMR: scaled preconditioned TFQMR method.

The set of linear solver modules distributed with IDAS is intended to be expanded in the future as new algorithms are developed. Note that users wishing to employ KLU or SuperLU_MT will need to download and install these libraries independent of SUNDIALS. SUNDIALS provides only the interfaces between itself and these libraries.

In the case of the direct methods IDADENSE and IDABAND the package includes an algorithm for the approximation of the Jacobian by difference quotients, but the user also has the option of supplying the Jacobian (or an approximation to it) directly. When using the sparse direct linear solvers IDAKLU and IDASUPERLUMT the user must supply a routine for the Jacobian (or an approximation to it) in CSC format, since standard difference quotient approximations do not leverage the inherent sparsity of the problem. In the case of the Krylov iterative methods IDASPGMR, IDASPCBG, and IDASPTFQMR, the package includes an algorithm for the approximation by difference quotients of the product between the Jacobian matrix and a vector of appropriate length. Again, the user has the option of providing a routine for this operation. When using any of the Krylov methods, the user must supply the preconditioning in two phases: a setup phase (preprocessing of Jacobian data) and a solve phase. While there is no default choice of preconditioner analogous to the difference quotient approximation in the direct case, the references [4, 7], together with the example and demonstration programs included with IDAS, offer considerable assistance in building preconditioners.

Each IDAS linear solver module consists of five routines, devoted to (1) memory allocation and initialization, (2) setup of the matrix data involved, (3) solution of the system, (4) monitoring performance, and (5) freeing of memory. The setup and solution phases are separate because the evaluation of Jacobians and preconditioners is done only periodically during the integration, as required to achieve convergence. The call list within the central IDAS module to each of the five associated functions is fixed, thus allowing the central module to be completely independent of the linear system method.

These modules are also decomposed in another way. Each of the linear solver modules (IDADENSE etc.) consists of an interface built on top of a generic linear system solver (DENSE etc.). The interface deals with the use of the particular method in the IDAS context, whereas the generic solver is

independent of the context. While some of the generic linear system solvers (DENSE, BAND, SPGMR, SPBCG, and SPTFQMR) were written with SUNDIALS in mind, they are intended to be usable anywhere as general-purpose solvers. This separation also allows for any generic solver to be replaced by an improved version, with no necessity to revise the IDAS package elsewhere.

IDAS also provides a preconditioner module, IDABBDPRE, that works in conjunction with NVECTOR_PARALLEL and generates a preconditioner that is a block-diagonal matrix with each block being a band matrix.

All state information used by IDAS to solve a given problem is saved in a structure, and a pointer to that structure is returned to the user. There is no global data in the IDAS package, and so, in this respect, it is reentrant. State information specific to the linear solver is saved in a separate structure, a pointer to which resides in the IDAS memory structure. The reentrancy of IDAS was motivated by the situation where two or more problems are solved by intermixed calls to the package from one user program.

Chapter 4

Using IDAS for IVP Solution

This chapter is concerned with the use of IDAS for the integration of DAEs. The following sections treat the header files, the layout of the user's main program, description of the IDAS user-callable functions, and description of user-supplied functions. This usage is essentially equivalent to using IDA [22].

The sample programs described in the companion document [30] may also be helpful. Those codes may be used as templates (with the removal of some lines involved in testing), and are included in the IDAS package.

The user should be aware that not all linear solver modules are compatible with all NVECTOR implementations. For example, NVECTOR_PARALLEL is not compatible with the direct dense, direct band or direct sparse linear solvers, since these linear solver modules need to form the complete system Jacobian. The IDADENSE and IDABAND modules (using either the internal implementation or Lapack), as well as the IDAKLU and IDASUPERLUMT modules can only be used with NVECTOR_SERIAL, NVECTOR_OPENMP and NVECTOR_PTHREADS. It is not recommended to use a threaded vector module with SuperLU_MT unless it is the NVECTOR_OPENMP module and SuperLU_MT is also compiled with openMP. The preconditioner module IDABBDPRE can only be used with NVECTOR_PARALLEL.

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

4.1 Access to library and header files

At this point, it is assumed that the installation of IDAS, following the procedure described in Appendix A, has been completed successfully.

Regardless of where the user's application program resides, its associated compilation and load commands must make reference to the appropriate locations for the library and header files required by IDAS. The relevant library files are

- *libdir/libsundials_idas.lib*,
- *libdir/libsundials_nvec*.lib* (one to four files),

where the file extension *.lib* is typically *.so* for shared libraries and *.a* for static libraries. The relevant header files are located in the subdirectories

- *incdir/include/idas*
- *incdir/include/sundials*
- *incdir/include/nvector*

The directories *libdir* and *incdir* are the install 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 (see Appendix A).

Note that an application cannot link to both the IDA and IDAS libraries because both contain user-callable functions with the same names (to ensure that IDAS is backward compatible with IDA). Therefore, applications that contain both DAE problems and DAEs with sensitivity analysis, should use IDAS.

4.2 Data types

The `sundials_types.h` file contains the definition of the type `realtype`, which is used by the SUNDIALS solvers for all floating-point data. The type `realtype` can be `float`, `double`, or `long double`, with the default being `double`. The user can change the precision of the SUNDIALS solvers arithmetic at the configuration stage (see §A.1.2 or §??).

Additionally, based on the current precision, `sundials_types.h` defines `BIG_REAL` to be the largest value representable as a `realtype`, `SMALL_REAL` to be the smallest value representable as a `realtype`, and `UNIT_ROUNDOFF` to be the difference between 1.0 and the minimum `realtype` greater than 1.0.

Within SUNDIALS, real constants are set by way of a macro called `RCONST`. It is this macro that needs the ability to branch on the definition `realtype`. 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 `RCONST(1.0)` automatically expands to 1.0 if `realtype` is `double`, to 1.0F if `realtype` is `float`, or to 1.0L if `realtype` is `long double`. SUNDIALS uses the `RCONST` macro internally to declare all of its floating-point constants.

A user program which uses the type `realtype` and the `RCONST` macro to handle floating-point constants is precision-independent except for any calls to precision-specific standard math library functions. (Our example programs use both `realtype` and `RCONST`.) Users can, however, use the type `double`, `float`, or `long double` in their code (assuming that this usage is consistent with the typedef for `realtype`). Thus, a previously existing piece of ANSI C code can use SUNDIALS without modifying the code to use `realtype`, so long as the SUNDIALS libraries use the correct precision (for details see §??).

4.3 Header files

The calling program must include several header files so that various macros and data types can be used. The header file that is always required is:

- `idas.h`, the header file for IDAS, which defines the several types and various constants, and includes function prototypes.

Note that `idas.h` includes `sundials_types.h`, which defines the types `realtype` and `booleantype` and the constants `FALSE` and `TRUE`.

The calling program must also include an `NVECTOR` implementation header file (see Chapter 7 for details). For the `NVECTOR` implementations that are included in the IDAS package, the corresponding header files are:

- `nvector_serial.h`, which defines the serial implementation `NVECTOR_SERIAL`;
- `nvector_parallel.h`, which defines the parallel MPI implementation, `NVECTOR_PARALLEL`.
- `nvector_openmp.h`, which defines the shared memory parallel openMP implementation,
- `nvector_pthreads.h`, which defines the shared memory parallel Pthreads implementation.

Note that both these files include in turn the header file `sundials_nvector.h` which defines the abstract `N_Vector` type.

Finally, a linear solver module header file is required. The header files corresponding to the various linear solver options in IDAS are as follows:

- `idas_dense.h`, which is used with the dense direct linear solver;
- `idas_band.h`, which is used with the band direct linear solver;
- `idas_lapack.h`, which is used with Lapack implementations of dense or band direct linear solvers;
- `idas_klu.h`, which is used with the KLU sparse direct linear solver;
- `idas_superlumi.h`, which is used with the SuperLU-MT threaded sparse direct linear solver;
- `idas_spgmr.h`, which is used with the scaled, preconditioned GMRES Krylov linear solver SPGMR;
- `idas_spgcrs.h`, which is used with the scaled, preconditioned Bi-CGSTab Krylov linear solver SPBCG;
- `idas_sptfqmr.h`, which is used with the scaled, preconditioned TFQMR Krylov solver SPTFQMR.

The header files for the dense and banded linear solvers (both internal and Lapack) include the file `idas_direct.h`, which defines common functions. This in turn includes a file (`sundials_direct.h`) which defines the matrix type for these direct linear solvers (`DlsMat`), as well as various functions and macros acting on such matrices.

The header files for the KLU and SuperLU-MT sparse linear solvers include the file `idas_sparse.h`, which defines common functions. This in turn includes a file (`sundials_sparse.h`) which defines the matrix type for these sparse direct linear solvers (`SlsMat`), as well as various functions and macros acting on such matrices.

The header files for the Krylov iterative solvers include `idas_spils.h` which defines common functions and which in turn includes a header file (`sundials_iterative.h`) which enumerates the kind of preconditioning and (for the SPGMR solver only) the choices for the Gram-Schmidt process.

Other headers may be needed, according to the choice of preconditioner, etc. For example, in the `idasFoodWeb_kry_p` example (see [30]), preconditioning is done with a block-diagonal matrix. For this, even though the IDASPGMR linear solver is used, the header `sundials_dense.h` is included for access to the underlying generic dense linear solver.

4.4 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. Some steps are independent of the NVECTOR implementation used; where this is not the case, usage specifications are given for the two implementations provided with IDAS: steps marked [P] correspond to NVECTOR_PARALLEL, steps marked [O] correspond to NVECTOR_OPENMP, steps marked [T] correspond to NVECTOR_PTHREADS, while steps marked [S] correspond to NVECTOR_SERIAL.

1. [P] Initialize MPI

Call `MPI_Init(&argc, &argv)` to initialize MPI if used by the user's program, aside from the internal use in NVECTOR_PARALLEL. Here `argc` and `argv` are the command line argument counter and array received by `main`.

2. Set problem dimensions

[S], [O], [T] Set `N`, the problem size N .

[O], [T] Set `num_threads`, the number of threads to use within the threaded vector functions.

[P] Set `Nlocal`, the local vector length (the sub-vector length for this processor); `N`, the global vector length (the problem size N , and the sum of all the values of `Nlocal`); and the active set of processors.

Note: The variables `N` and `Nlocal` should be of type `long int`. The variable `num_threads` should be of type `int`.

3. Set vectors of initial values

To set the vectors `y0` and `yp0` to initial values for y and \dot{y} , use functions defined by the particular `NVECTOR` implementation. For the two `NVECTOR` implementations provided, if a `realtype` array `ydata` already exists, containing the initial values of y , make the calls:

```
[S] y0 = N_VMake_Serial(N, ydata);
[O] y0 = N_VMake_OpenMP(N, num_threads, ydata);
[T] y0 = N_VMake_Pthreads(N, num_threads, ydata);
[P] y0 = N_VMake_Parallel(comm, Nlocal, N, ydata);
```

Otherwise, make the calls:

```
[S] y0 = N_VNew_Serial(N);
[O] y0 = N_VNew_OpenMP(N, num_threads);
[T] y0 = N_VNew_Pthreads(N, num_threads);
[P] y0 = N_VNew_Parallel(comm, Nlocal, N);
```

and load initial values into the structure defined by:

```
[S] NV_DATA_S(y0)
[O] NV_DATA_OMP(y0)
[T] NV_DATA_PT(y0)
[P] NV_DATA_P(y0)
```

Here `comm` is the MPI communicator, set in one of two ways: If a proper subset of active processors is to be used, `comm` must be set by suitable MPI calls. Otherwise, to specify that all processors are to be used, `comm` must be `MPI_COMM_WORLD`.

The initial conditions for \dot{y} are set similarly.

4. Create IDAS object

Call `ida_mem = IDACreate()` to create the IDAS memory block. `IDACreate` returns a pointer to the IDAS memory structure. See §4.5.1 for details. This `void *` pointer must then be passed as the first argument to all subsequent IDAS function calls.

5. Initialize IDAS solver

Call `IDAInit(...)` to provide required problem specifications (residual function, initial time, and initial conditions), allocate internal memory for IDAS, and initialize IDAS. `IDAInit` returns an error flag to indicate success or an illegal argument value. See §4.5.1 for details.

6. Specify integration tolerances

Call `IDASStolerances(...)` or `IDASVtolerances(...)` to specify, respectively, a scalar relative tolerance and scalar absolute tolerance, or a scalar relative tolerance and a vector of absolute tolerances. Alternatively, call `IDAWFtolerances` to specify a function which sets directly the weights used in evaluating WRMS vector norms. See §4.5.2 for details.

7. Set optional inputs

Optionally, call `IDASet*` functions to change from their default values any optional inputs that control the behavior of IDAS. See §4.5.7.1 for details.

8. Attach linear solver module

Initialize the linear solver module with one of the following calls (for details see §4.5.3):

```
[S], [O], [T] flag = IDADense(...);
[S], [O], [T] flag = IDABand(...);
[S], [O], [T] flag = IDALapackDense(...);
[S], [O], [T] flag = IDALapackBand(...);
[S], [O], [T] flag = IDAKLU(...);
[S], [O], [T] flag = IDASuperLUMT(...);
flag = IDASpgmr(...);
flag = IDASpbcg(...);
flag = IDASptfqmr(...);
```

9. Set linear solver optional inputs

Optionally, call `IDA*Set*` functions from the selected linear solver module to change optional inputs specific to that linear solver. See §4.5.7.2 and §4.5.7.4 for details.

10. Correct initial values

Optionally, call `IDACalcIC` to correct the initial values `y0` and `yp0` passed to `IDAInit`. See §4.5.4. Also see §4.5.7.5 for relevant optional input calls.

11. Specify rootfinding problem

Optionally, call `IDARootInit` to initialize a rootfinding problem to be solved during the integration of the DAE system. See §4.5.5 for details, and see §4.5.7.6 for relevant optional input calls.

12. Advance solution in time

For each point at which output is desired, call `flag = IDASolve(ida_mem, tout, &tret, yret, ypret, itask)`. Here `itask` specifies the return mode. The vector `yret` (which can be the same as the vector `y0` above) will contain $y(t)$, while the vector `ypret` will contain $\dot{y}(t)$. See §4.5.6 for details.

13. Get optional outputs

Call `IDA*Get*` functions to obtain optional output. See §4.5.9 for details.

14. Deallocate memory for solution vectors

Upon completion of the integration, deallocate memory for the vectors `yret` and `ypret` by calling the destructor function defined by the `NVECTOR` implementation:

```
[S] N_VDestroy_Serial(yret);
[O] N_VDestroy_OpenMP(y);
[T] N_VDestroy_Pthreads(y);
[P] N_VDestroy_Parallel(yret);
```

and similarly for `ypret`.

15. Free solver memory

`IDAFree(&ida_mem)` to free the memory allocated for IDAS.

16. [P] Finalize MPI

Call `MPI_Finalize()` to terminate MPI.

4.5 User-callable functions

This section describes the IDAS functions that are called by the user to set up and solve a DAE. Some of these are required. However, starting with §4.5.7, the functions listed involve optional inputs/outputs or restarting, and those paragraphs can be skipped for a casual use of IDAS. In any case, refer to §4.4 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 §4.5.7.1).

4.5.1 IDAS initialization and deallocation functions

The following three functions must be called in the order listed. The last one is to be called only after the DAE solution is complete, as it frees the IDAS memory block created and allocated by the first two calls.

IDACreate

Call `ida_mem = IDACreate();`

Description The function `IDACreate` instantiates an IDAS solver object.

Arguments `IDACreate` has no arguments.

Return value If successful, `IDACreate` returns a pointer to the newly created IDAS memory block (of type `void *`). Otherwise it returns `NULL`.

IDAINit

Call `flag = IDAINit(ida_mem, res, t0, y0, yp0);`

Description The function `IDAINit` provides required problem and solution specifications, allocates internal memory, and initializes IDAS.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block returned by `IDACreate`.
`res` (`IDAResFn`) is the C function which computes the residual function F in the DAE. This function has the form `res(t, yy, yp, resval, user_data)`. For full details see §4.6.1.
`t0` (`realtype`) is the initial value of t .
`y0` (`N_Vector`) is the initial value of y .
`yp0` (`N_Vector`) is the initial value of \dot{y} .

Return value The return value `flag` (of type `int`) will be one of the following:

`IDA_SUCCESS` The call to `IDAINit` was successful.
`IDA_MEM_NULL` The IDAS memory block was not initialized through a previous call to `IDACreate`.
`IDA_MEM_FAIL` A memory allocation request has failed.
`IDA_ILL_INPUT` An input argument to `IDAINit` has an illegal value.

Notes If an error occurred, `IDAINit` also sends an error message to the error handler function.

IDAFree

Call `IDAFree(&ida_mem);`

Description The function `IDAFree` frees the pointer allocated by a previous call to `IDACreate`.

Arguments The argument is the pointer to the IDAS memory block (of type `void *`).

Return value The function `IDAFree` has no return value.

4.5.2 IDAS tolerance specification functions

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

IDASStolerances

Call `flag = IDASStolerances(ida_mem, reltol, abstol);`

Description The function `IDASStolerances` specifies scalar relative and absolute tolerances.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block returned by `IDACreate`.

`reltol` (`realtype`) is the scalar relative error tolerance.

`abstol` (`realtype`) is the scalar absolute error tolerance.

Return value The return value `flag` (of type `int`) will be one of the following:

`IDA_SUCCESS` The call to `IDASStolerances` was successful.

`IDA_MEM_NULL` The IDAS memory block was not initialized through a previous call to `IDACreate`.

`IDA_NO_MALLOC` The allocation function `IDAInit` has not been called.

`IDA_ILL_INPUT` One of the input tolerances was negative.

IDASVtolerances

Call `flag = IDASVtolerances(ida_mem, reltol, abstol);`

Description The function `IDASVtolerances` specifies scalar relative tolerance and vector absolute tolerances.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block returned by `IDACreate`.

`reltol` (`realtype`) is the scalar relative error tolerance.

`abstol` (`N_Vector`) is the vector of absolute error tolerances.

Return value The return value `flag` (of type `int`) will be one of the following:

`IDA_SUCCESS` The call to `IDASVtolerances` was successful.

`IDA_MEM_NULL` The IDAS memory block was not initialized through a previous call to `IDACreate`.

`IDA_NO_MALLOC` The allocation function `IDAInit` has not been called.

`IDA_ILL_INPUT` The relative error tolerance was negative or the absolute tolerance had a negative component.

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

IDAWFtolerances

Call	<code>flag = IDAWFtolerances(ida_mem, efun);</code>
Description	The function <code>IDAWFtolerances</code> specifies a user-supplied function <code>efun</code> that sets the multiplicative error weights W_i for use in the weighted RMS norm, which are normally defined by Eq. (2.7).
Arguments	<p><code>ida_mem</code> (void *) pointer to the IDAS memory block returned by <code>IDACreate</code>.</p> <p><code>efun</code> (<code>IDAEvtFn</code>) is the C function which defines the <code>ewt</code> vector (see §4.6.3).</p>
Return value	<p>The return value <code>flag</code> (of type <code>int</code>) will be one of the following:</p> <p><code>IDA_SUCCESS</code> The call to <code>IDAWFtolerances</code> was successful.</p> <p><code>IDA_MEM_NULL</code> The IDAS memory block was not initialized through a previous call to <code>IDACreate</code>.</p> <p><code>IDA_NO_MALLOC</code> The allocation function <code>IDAINit</code> has not been called.</p>

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`= 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 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 `idasRoberts_dns` in the IDAS package, and the discussion of it in the IDAS Examples document [30]. 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 a 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 `reltol`= 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.

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

(1) The way to control the size of unwanted negative computed values is with tighter absolute tolerances. Again this requires some knowledge of the noise level of these components, which may or may not be different for different components. Some experimentation may be needed.

(2) If output plots or tables are being generated, and it is important to avoid having negative numbers appear there (for the sake of avoiding a long explanation of them, if nothing else), then eliminate them, but only in the context of the output medium. Then the internal values carried by the solver are unaffected. Remember that a small negative value in `yret` returned by IDAS, with magnitude comparable to `abstol` or less, is equivalent to zero as far as the computation is concerned.

(3) The user's residual routine `res` should never change a negative value in the solution vector `yy` to a non-negative value, as a "solution" to this problem. This can cause instability. If the `res` routine cannot tolerate a zero or negative value (e.g. because there is a square root or log of it), then the

offending value should be changed to zero or a tiny positive number in a temporary variable (not in the input `yy` vector) for the purposes of computing $F(t, y, \dot{y})$.

(4) IDAS provides the option of enforcing positivity or non-negativity on components. Also, such constraints can be enforced by use of the recoverable error return feature in the user-supplied residual function. However, because these options involve some extra overhead cost, they should only be exercised if the use of absolute tolerances to control the computed values is unsuccessful.

4.5.3 Linear solver specification functions

As previously explained, Newton iteration requires the solution of linear systems of the form (2.5). There are seven IDAS linear solvers currently available for this task: IDADENSE, IDABAND, IDAKLU, IDASUPERLUMT, IDASPGMR, IDASPBCG, and IDASPTFQMR.

The first two linear solvers are direct and derive their names from the type of approximation used for the Jacobian $J = \partial F / \partial y + \alpha \partial F / \partial \dot{y}$. IDADENSE and IDABAND work with dense and banded approximations to J , respectively. The SUNDIALS suite includes both internal implementations of these two linear solvers and interfaces to Lapack implementations. Together, these linear solvers are referred to as IDADLS (from Direct Linear Solvers).

The second two linear solvers are sparse direct solvers based on Gaussian elimination, and require user-supplied routines to construct the Jacobian $J = \partial F / \partial y + \alpha \partial F / \partial \dot{y}$ in compressed-sparse-column format. The SUNDIALS suite does not include internal implementations of these solver libraries, instead requiring compilation of SUNDIALS to link with existing installations of these libraries (if either is missing, SUNDIALS will install without the corresponding interface routines). Together, these linear solvers are referred to as CVSLs (from Sparse Linear Solvers).

The remaining three IDAS linear solvers, IDASPGMR, IDASPBCG, and IDASPTFQMR, are Krylov iterative solvers. The SPGMR, SPBCG, and SPTFQMR in the names indicate the scaled preconditioned GMRES, scaled preconditioned Bi-CGStab, and scaled preconditioned TFQMR methods, respectively. Together, they are referred to as IDASPILS (from Scaled Preconditioned Iterative Linear Solvers).

When using any of the Krylov linear solvers, preconditioning (on the left) is permitted, and in fact encouraged, for the sake of efficiency. A preconditioner matrix P must approximate the Jacobian J , at least crudely. For the specification of a preconditioner, see §4.5.7.4 and §4.6.

To specify an IDAS linear solver, after the call to `IDACreate` but before any calls to `IDASolve`, the user's program must call one of the functions `IDADense`/`IDALapackDense`, `IDABand`/`IDALapackBand`, `IDAKLU`, `IDASuperLUMT`, `IDASpgmr`, `IDASpbcg`, or `IDASptfqmr`, as documented below. The first argument passed to these functions is the IDAS memory pointer returned by `IDACreate`. A call to one of these functions links the main IDAS integrator to a linear solver and allows the user to specify parameters which are specific to a particular solver, such as the bandwidths in the IDABAND case. The use of each of the linear solvers involves certain constants 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 linear solver, as specified below.

In each case the linear solver module used by IDAS is actually built on top of a generic linear system solver, which may be of interest in itself. These generic solvers, denoted DENSE, BAND, KLU, SUPERLUMT, SPGMR, SPBCG, and SPTFQMR, are described separately in Chapter 9.

IDADense

Call `flag = IDADense(ida_mem, N);`

Description The function `IDADense` selects the IDADENSE linear solver and indicates the use of the internal direct dense linear algebra functions.

The user's main program must include the `idas_dense.h` header file.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`N` (`long int`) problem dimension.

Return value The return value `flag` (of type `int`) is one of

`IDADLS_SUCCESS` The IDADENSE initialization was successful.

IDADLS_MEM_NULL The `ida_mem` pointer is NULL.
 IDADLS_ILL_INPUT The IDADENSE solver is not compatible with the current NVECTOR module.
 IDADLS_MEM_FAIL A memory allocation request failed.

Notes The IDADENSE linear solver is not compatible with all implementations of the NVECTOR module. Of the NVECTOR modules provided with SUNDIALS, only NVECTOR_SERIAL, NVECTOR_OPENMP and NVECTOR_PTHREADS are compatible, while NVECTOR_PARALLEL is not.

IDALapackDense

Call `flag = IDALapackDense(ida_mem, N);`
 Description The function `IDALapackDense` selects the IDADENSE linear solver and indicates the use of Lapack functions.
 The user's main program must include the `idas_lapack.h` header file.
 Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`N` (`int`) problem dimension.
 Return value The values of the returned `flag` (of type `int`) are identical to those of `IDADense`.
 Notes Note that `N` is restricted to be of type `int` here, because of the corresponding type restriction in the Lapack solvers.

IDABand

Call `flag = IDABand(ida_mem, N, mupper, mlower);`
 Description The function `IDABand` selects the IDABAND linear solver and indicates the use of the internal direct band linear algebra functions.
 The user's main program must include the `idas_band.h` header file.
 Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`N` (`long int`) problem dimension.
`mupper` (`long int`) upper half-bandwidth of the problem Jacobian (or of the approximation of it).
`mlower` (`long int`) lower half-bandwidth of the problem Jacobian (or of the approximation of it).
 Return value The return value `flag` (of type `int`) is one of
 IDABAND_SUCCESS The IDABAND initialization was successful.
 IDABAND_MEM_NULL The `ida_mem` pointer is NULL.
 IDABAND_ILL_INPUT The IDABAND solver is not compatible with the current NVECTOR module, or one of the Jacobian half-bandwidths is outside its valid range ($0 \dots N-1$).
 IDABAND_MEM_FAIL A memory allocation request failed.
 Notes The IDABAND linear solver is not compatible with all implementations of the NVECTOR module. Of the NVECTOR modules provided with SUNDIALS, only NVECTOR_SERIAL, NVECTOR_OPENMP and NVECTOR_PTHREADS are compatible, while NVECTOR_PARALLEL is not. The half-bandwidths are to be set so that the nonzero locations (i, j) in the banded (approximate) Jacobian satisfy $-mlower \leq j - i \leq mupper$.

IDALapackBand

- Call** `flag = IDALapackBand(ida_mem, N, mupper, mlower);`
- Description** The function `IDALapackBand` selects the `IDABAND` linear solver and indicates the use of Lapack functions.
- The user's main program must include the `idas_lapack.h` header file.
- Arguments** The input arguments are identical to those of `IDABand`, except that `N`, `mupper`, and `mlower` are of type `int` here.
- Return value** The values of the returned `flag` (of type `int`) are identical to those of `IDABand`.
- Notes** Note that `N`, `mupper`, and `mlower` are restricted to be of type `int` here, because of the corresponding type restriction in the Lapack solvers.

IDAKLU

- Call** `flag = IDAKLU(ida_mem, N, NNZ);`
- Description** The function `IDAKLU` selects the `IDAKLU` linear solver and indicates the use of sparse direct linear algebra functions.
- The user's main program must include the `idas_sparse.h` header file.
- Arguments** `ida_mem` (`void *`) pointer to the IDAS memory block.
`N` (`int`) problem dimension.
`NNZ` (`int`) maximum number of nonzero entries in the system Jacobian.
- Return value** The return value `flag` (of type `int`) is one of
- `IDASLS_SUCCESS` The `IDAKLU` initialization was successful.
 - `IDASLS_MEM_NULL` The `ida_mem` pointer is `NULL`.
 - `IDASLS_ILL_INPUT` The `IDAKLU` solver is not compatible with the current `NVECTOR` module.
 - `IDASLS_MEM_FAIL` A memory allocation request failed.
 - `IDASLS_PACKAGE_FAIL` A call to the `KLU` library returned a failure flag.
- Notes** The `IDAKLU` linear solver is not compatible with all implementations of the `NVECTOR` module. Of the `NVECTOR` modules provided with `SUNDIALS`, only `NVECTOR_SERIAL`, `NVECTOR_OPENMP` and `NVECTOR_PTHREADS` are compatible, while `NVECTOR_PARALLEL` is not.

IDASuperLUMT

- Call** `flag = IDASuperLUMT(ida_mem, num_threads, N, NNZ);`
- Description** The function `IDASuperLUMT` selects the `IDASUPERLUMT` linear solver and indicates the use of sparse direct linear algebra functions.
- The user's main program must include the `idas_superluml.h` header file.
- Arguments** `ida_mem` (`void *`) pointer to the IDAS memory block.
`num_threads` (`int`) the number of threads to use when factoring/solving the linear systems. Note that `SuperLU-MT` is thread-parallel only in the factorization routine.
`N` (`int`) problem dimension.
`NNZ` (`int`) maximum number of nonzero entries in the system Jacobian.
- Return value** The return value `flag` (of type `int`) is one of
- `IDASLS_SUCCESS` The `IDASUPERLUMT` initialization was successful.
 - `IDASLS_MEM_NULL` The `ida_mem` pointer is `NULL`.

IDASLS_ILL_INPUT The IDASUPERLUMT solver is not compatible with the current NVECTOR module.

IDASLS_MEM_FAIL A memory allocation request failed.

IDASLS_PACKAGE_FAIL A call to the SuperLU_MT library returned a failure flag.

Notes The IDASUPERLUMT linear solver is not compatible with all implementations of the NVECTOR module. Of the NVECTOR modules provided with SUNDIALS, only NVECTOR_SERIAL, NVECTOR_OPENMP and NVECTOR_PTHREADS are compatible, while NVECTOR_PARALLEL is not.



Performance will significantly degrade if the user applies the SuperLU_MT package compiled with PThreads while using the NVECTOR_OPENMP module. If a user wants to use a threaded vector kernel with this thread-parallel solver, then SuperLU_MT should be compiled with openMP and the NVECTOR_OPENMP module should be used. Also, note that the expected benefit of using the threaded vector kernel is minimal compared to the potential benefit of the threaded solver, unless very long (greater than 100,000 entries) vectors are used.

IDASpgmr

Call `flag = IDASpgmr(ida_mem, maxl);`

Description The function IDASpgmr selects the IDASPGMR linear solver.

The user's main program must include the `idas_spgmr.h` header file.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.

`maxl` (int) maximum dimension of the Krylov subspace to be used. Pass 0 to use the default value IDA_SPILS_MAXL= 5.

Return value The return value `flag` (of type int) is one of

IDASPILS_SUCCESS The IDASPGMR initialization was successful.

IDASPILS_MEM_NULL The `ida_mem` pointer is NULL.

IDASPILS_MEM_FAIL A memory allocation request failed.

IDASpbcg

Call `flag = IDASpbcg(ida_mem, maxl);`

Description The function IDASpbcg selects the IDASPCBG linear solver.

The user's main program must include the `idas_spgcs.h` header file.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.

`maxl` (int) maximum dimension of the Krylov subspace to be used. Pass 0 to use the default value IDA_SPILS_MAXL= 5.

Return value The return value `flag` (of type int) is one of

IDASPILS_SUCCESS The IDASPCBG initialization was successful.

IDASPILS_MEM_NULL The `ida_mem` pointer is NULL.

IDASPILS_MEM_FAIL A memory allocation request failed.

IDASptfqmr

Call `flag = IDASptfqmr(ida_mem, maxl);`

Description The function IDASptfqmr selects the IDASPTFQMR linear solver.

The user's main program must include the `idas_sptfqmr.h` header file.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.

max1 (int) maximum dimension of the Krylov subspace to be used. Pass 0 to use the default value `IDA_SPILS_MAXL=5`.

Return value The return value **flag** (of type **int**) is one of

`IDASPILS_SUCCESS` The IDASPTFQMR initialization was successful.

`IDASPILS_MEM_NULL` The `ida_mem` pointer is NULL.

`IDASPILS_MEM_FAIL` A memory allocation request failed.

4.5.4 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 §2.1 and Ref. [6].) It uses 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 **y0** and **yp0** 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.

IDACalcIC

Call `flag = IDACalcIC(ida_mem,icopt,tout1);`

Description The function IDACalcIC corrects the initial values **y0** and **yp0** at time **t0**.

Arguments **ida_mem** (void *) pointer to the IDAS memory block.

icopt (int) is one of the following two options for the initial condition calculation.
`icopt=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.
`icopt=IDA_Y_INIT` directs IDACalcIC to compute all components of *y*, given \dot{y} . In this case, **id** is not required.

tout1 (realtype) 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 The return value **flag** (of type **int**) will be one of the following:

<code>IDA_SUCCESS</code>	IDASolve succeeded.
<code>IDA_MEM_NULL</code>	The argument <code>ida_mem</code> was NULL.
<code>IDA_NO_MALLOC</code>	The allocation function IDAInit has not been called.
<code>IDA_ILL_INPUT</code>	One of the input arguments was illegal.
<code>IDA_LSETUP_FAIL</code>	The linear solver's setup function failed in an unrecoverable manner.
<code>IDA_LINIT_FAIL</code>	The linear solver's initialization function failed.
<code>IDA_LSOLVE_FAIL</code>	The linear solver's solve function failed in an unrecoverable manner.
<code>IDA_BAD_EWT</code>	Some component of the error weight vector is zero (illegal), either for the input value of y0 or a corrected value.
<code>IDA_FIRST_RES_FAIL</code>	The user's residual function returned a recoverable error flag on the first call, but IDACalcIC was unable to recover.
<code>IDA_RES_FAIL</code>	The user's residual function returned a nonrecoverable error flag.
<code>IDA_NO_RECOVERY</code>	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 <code>steptol</code> in weighted RMS norm.
	IDA_CONV_FAIL	IDACalcIC failed to get convergence of the Newton iterations.
Notes	All failure return values are negative and therefore a test <code>flag < 0</code> will trap all IDACalcIC failures.	
	Note that 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 (see §4.5.9.2).	

4.5.5 Rootfinding initialization function

While integrating the IVP, IDAS has the capability of finding the roots of a set of user-defined functions. To activate the rootfinding 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.

IDARootInit

Call	<code>flag = IDARootInit(ida_mem, nrtfn, g);</code>	
Description	The function IDARootInit specifies that the roots of a set of functions $g_i(t, y, \dot{y})$ are to be found while the IVP is being solved.	
Arguments	ida_mem (void *) pointer to the IDAS memory block returned by IDACreate. nrtfn (int) is the number of root functions g_i . g (IDARootFn) is the C function which defines the nrtfn functions $g_i(t, y, \dot{y})$ whose roots are sought. See §4.6.4 for details.	
Return value	The return value flag (of type int) is one of <ul style="list-style-type: none"> IDA_SUCCESS The call to IDARootInit 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.	

4.5.6 IDAS solver function

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

IDASolve

Call	<code>flag = IDASolve(ida_mem, tout, &tret, yret, ypret, itask);</code>	
Description	The function IDASolve integrates the DAE over an interval in t .	
Arguments	ida_mem (void *) pointer to the IDAS memory block. tout (realtype) the next time at which a computed solution is desired. tret (realtype) the time reached by the solver (output). yret (N.Vector) the computed solution vector y . ypret (N.Vector) the computed solution vector \dot{y} .	

itask (int) a flag indicating the job of the solver for the next user step. The `IDA_NORMAL` task is to have the solver take internal steps until it has reached or just passed the user specified `tout` parameter. The solver then interpolates in order to return approximate values of $y(tout)$ and $\dot{y}(tout)$. The `IDA_ONE_STEP` option tells the solver to just take one internal step and return the solution at the point reached by that step.

Return value `IDASolve` returns vectors `yret` and `ypret` and a corresponding independent variable value $t = tret$, such that $(yret, ypret)$ are the computed values of $(y(t), \dot{y}(t))$.

In `IDA_NORMAL` mode with no errors, `tret` will be equal to `tout` and $yret = y(tout)$, $ypret = \dot{y}(tout)$.

The return value `flag` (of type `int`) will be one of the following:

<code>IDA_SUCCESS</code>	<code>IDASolve</code> succeeded.
<code>IDA_TSTOP_RETURN</code>	<code>IDASolve</code> succeeded by reaching the stop point specified through the optional input function <code>IDASetStopTime</code> .
<code>IDA_ROOT_RETURN</code>	<code>IDASolve</code> succeeded and found one or more roots. In this case, <code>tret</code> is the location of the root. If <code>nrtfn</code> > 1, call <code>IDAGetRootInfo</code> to see which g_i were found to have a root. See §4.5.9.3 for more information.
<code>IDA_MEM_NULL</code>	The <code>ida_mem</code> argument was NULL.
<code>IDA_ILL_INPUT</code>	One of the inputs to <code>IDASolve</code> was illegal, or some other input to the solver was either illegal or missing. The latter category includes the following situations: (a) The tolerances have not been set. (b) A component of the error weight vector became zero during internal time-stepping. (c) The linear solver initialization function (called by the user after calling <code>IDACreate</code>) failed to set the linear solver-specific <code>lsolve</code> field in <code>ida_mem</code> . (d) 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.
<code>IDA_TOO_MUCH_WORK</code>	The solver took <code>mxstep</code> internal steps but could not reach <code>tout</code> . The default value for <code>mxstep</code> is <code>MXSTEP_DEFAULT</code> = 500.
<code>IDA_TOO_MUCH_ACC</code>	The solver could not satisfy the accuracy demanded by the user for some internal step.
<code>IDA_ERR_FAIL</code>	Error test failures occurred too many times (<code>MXNEF</code> = 10) during one internal time step or occurred with $ h = h_{min}$.
<code>IDA_CONV_FAIL</code>	Convergence test failures occurred too many times (<code>MXNCF</code> = 10) during one internal time step or occurred with $ h = h_{min}$.
<code>IDA_LINIT_FAIL</code>	The linear solver's initialization function failed.
<code>IDA_LSETUP_FAIL</code>	The linear solver's setup function failed in an unrecoverable manner.
<code>IDA_LSOLVE_FAIL</code>	The linear solver's solve function failed in an unrecoverable manner.
<code>IDA_CONSTR_FAIL</code>	The inequality constraints were violated and the solver was unable to recover.
<code>IDA_REP_RES_ERR</code>	The user's residual function repeatedly returned a recoverable error flag, but the solver was unable to recover.
<code>IDA_RES_FAIL</code>	The user's residual function returned a nonrecoverable error flag.
<code>IDA_RTFUNC_FAIL</code>	The rootfinding function failed.

Notes The vector `yret` can occupy the same space as the vector `y0` of initial conditions that was passed to `IDAInit`, and the vector `ypret` can occupy the same space as `yp0`.

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.

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.

4.5.7 Optional input functions

There are numerous optional input parameters that control the behavior of the IDAS solver. IDAS provides functions that can be used to change these optional input parameters from their default values. Table 4.1 lists all optional input functions in IDAS which are then described in detail in the remainder of this section. For the most casual use of IDAS, the reader can skip to §4.6.

We note that, on an error return, all these functions also send an error message to the error handler function. We also note that all error return values are negative, so a test `flag < 0` will catch any error.

4.5.7.1 Main solver optional input functions

The calls listed here can be executed in any order. However, if the user's program calls either `IDASetErrFile` or `IDASetErrHandlerFn`, then that call should appear first, in order to take effect for any later error message.

IDASetErrFile

Call	<code>flag = IDASetErrFile(ida_mem, errfp);</code>
Description	The function <code>IDASetErrFile</code> specifies the pointer to the file where all IDAS messages should be directed when the default IDAS error handler function is used.
Arguments	<code>ida_mem</code> (void *) pointer to the IDAS memory block. <code>errfp</code> (FILE *) pointer to output file.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDA_SUCCESS</code> The optional value has been successfully set. <code>IDA_MEM_NULL</code> The <code>ida_mem</code> pointer is NULL.
Notes	The default value for <code>errfp</code> is <code>stderr</code> . Passing a value NULL disables all future error message output (except for the case in which the IDAS memory pointer is NULL). This use of <code>IDASetErrFile</code> is strongly discouraged. If <code>IDASetErrFile</code> is to be called, it should be called before any other optional input functions, in order to take effect for any later error message.



IDASetErrHandlerFn

Call	<code>flag = IDASetErrHandlerFn(ida_mem, ehfun, eh_data);</code>
Description	The function <code>IDASetErrHandlerFn</code> specifies the optional user-defined function to be used in handling error messages.
Arguments	<code>ida_mem</code> (void *) pointer to the IDAS memory block. <code>ehfun</code> (IDAErrHandlerFn) is the user's C error handler function (see §4.6.2). <code>eh_data</code> (void *) pointer to user data passed to <code>ehfun</code> every time it is called.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDA_SUCCESS</code> The function <code>ehfun</code> and data pointer <code>eh_data</code> have been successfully set. <code>IDA_MEM_NULL</code> The <code>ida_mem</code> pointer is NULL.

Table 4.1: Optional inputs for IDAS, IDADLS, IDASLS, and IDASPILS

Optional input	Function name	Default
IDAS main solver		
Pointer to an error file	IDASetErrFile	stderr
Error handler function	IDASetErrHandlerFn	internal fn.
User data	IDASetUserData	NULL
Maximum order for BDF method	IDASetMaxOrd	5
Maximum no. of internal steps before t_{out}	IDASetMaxNumSteps	500
Initial step size	IDASetInitStep	estimated
Maximum absolute step size	IDASetMaxStep	∞
Value of t_{stop}	IDASetStopTime	∞
Maximum no. of error test failures	IDASetMaxErrTestFails	10
Maximum no. of nonlinear iterations	IDASetMaxNonlinIters	4
Maximum no. of convergence failures	IDASetMaxConvFails	10
Maximum no. of error test failures	IDASetMaxErrTestFails	7
Coeff. in the nonlinear convergence test	IDASetNonlinConvCoef	0.33
Suppress alg. vars. from error test	IDASetSuppressAlg	FALSE
Variable types (differential/algebraic)	IDASetId	NULL
Inequality constraints on solution	IDASetConstraints	NULL
Direction of zero-crossing	IDASetRootDirection	both
Disable rootfinding warnings	IDASetNoInactiveRootWarn	none
IDAS initial conditions calculation		
Coeff. in the nonlinear convergence test	IDASetNonlinConvCoefIC	0.0033
Maximum no. of steps	IDASetMaxNumStepsIC	5
Maximum no. of Jacobian/precond. evals.	IDASetMaxNumJacsIC	4
Maximum no. of Newton iterations	IDASetMaxNumItersIC	10
Turn off linesearch	IDASetLineSearchOffIC	FALSE
Lower bound on Newton step	IDASetStepToleranceIC	around ^{2/3}
IDADLS linear solvers		
Dense Jacobian function	IDADlsSetDenseJacFn	DQ
Band Jacobian function	IDADlsSetBandJacFn	DQ
IDASLS linear solvers		
Sparse Jacobian function	IDASlsSetSparseJacFn	none
Sparse matrix ordering algorithm	IDAKLUSetOrdering	1 for COLAMD
Sparse matrix ordering algorithm	IDASuperLUMTSetOrdering	3 for COLAMD
IDASPILS linear solvers		
Preconditioner functions	IDASpilsSetPreconditioner	NULL, NULL
Jacobian-times-vector function	IDASpilsSetJacTimesVecFn	DQ
Factor in linear convergence test	IDASpilsSetEpsLin	0.05
Factor in DQ increment calculation	IDASpilsSetIncrementFactor	1.0
Maximum no. of restarts (IDASPGMR)	IDASpilsSetMaxRestarts	5
Type of Gram-Schmidt orthogonalization ^(a)	IDASpilsSetGSType	classical GS
Maximum Krylov subspace size ^(b)	IDASpilsSetMaxl	5

^(a) Only for IDASPGMR^(b) Only for IDASPCBG and IDASPTFQMR

Notes Error messages indicating that the IDAS solver memory is NULL will always be directed to `stderr`.

IDASetUserData

Call `flag = IDASetUserData(ida_mem, user_data);`

Description The function `IDASetUserData` specifies the user data block `user_data` and attaches it to the main IDAS memory block.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`user_data` (`void *`) pointer to the user data.

Return value The return value `flag` (of type `int`) is one of
`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.



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

IDASetMaxOrd

Call `flag = IDASetMaxOrd(ida_mem, maxord);`

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

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`maxord` (`int`) value of the maximum method order. This must be positive.

Return value The return value `flag` (of type `int`) is one of
`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 its previous value.

Notes The default value is 5. If the input value exceeds 5, the value 5 will be used. Since `maxord` affects the memory requirements for the internal IDAS memory block, its value cannot be increased past its previous value.

IDASetMaxNumSteps

Call `flag = IDASetMaxNumSteps(ida_mem, mxsteps);`

Description 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` (`void *`) pointer to the IDAS memory block.
`mxsteps` (`long int`) maximum allowed number of steps.

Return value The return value `flag` (of type `int`) is one of
`IDA_SUCCESS` The optional value has been successfully set.
`IDA_MEM_NULL` The `ida_mem` pointer is NULL.

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

IDASetInitStep

Call `flag = IDASetInitStep(ida_mem, hin);`

Description The function `IDASetInitStep` specifies the initial step size.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`hin` (`realtype`) value of the initial step size to be attempted. Pass 0.0 to have IDAS use the default value.

Return value The return value `flag` (of type `int`) is one of
`IDA_SUCCESS` The optional value has been successfully set.
`IDA_MEM_NULL` The `ida_mem` pointer is NULL.

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

IDASetMaxStep

Call `flag = IDASetMaxStep(ida_mem, hmax);`

Description The function `IDASetMaxStep` specifies the maximum absolute value of the step size.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`hmax` (`realtype`) maximum absolute value of the step size.

Return value The return value `flag` (of type `int`) is one of
`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 ∞ .

IDASetStopTime

Call `flag = IDASetStopTime(ida_mem, tstop);`

Description The function `IDASetStopTime` specifies the value of the independent variable t past which the solution is not to proceed.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`tstop` (`realtype`) value of the independent variable past which the solution should not proceed.

Return value The return value `flag` (of type `int`) is one of
`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.

IDASetMaxErrTestFails

Call `flag = IDASetMaxErrTestFails(ida_mem, maxnef);`

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

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`maxnef` (`int`) maximum number of error test failures allowed on one step (> 0).

Return value The return value `flag` (of type `int`) is one of

IDA_SUCCESS The optional value has been successfully set.

IDA_MEM_NULL The `ida_mem` pointer is NULL.

Notes The default value is 7.

IDASsetMaxNonlinIters

Call `flag = IDASsetMaxNonlinIters(ida_mem, maxcor);`

Description The function `IDASsetMaxNonlinIters` specifies the maximum number of nonlinear solver iterations at one step.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`maxcor` (`int`) maximum number of nonlinear solver iterations allowed on one step (> 0).

Return value The return value `flag` (of type `int`) is one of

IDA_SUCCESS The optional value has been successfully set.

IDA_MEM_NULL The `ida_mem` pointer is NULL.

Notes The default value is 3.

IDASsetMaxConvFails

Call `flag = IDASsetMaxConvFails(ida_mem, maxncf);`

Description The function `IDASsetMaxConvFails` specifies the maximum number of nonlinear solver convergence failures at one step.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`maxncf` (`int`) maximum number of allowable nonlinear solver convergence failures on one step (> 0).

Return value The return value `flag` (of type `int`) is one of

IDA_SUCCESS The optional value has been successfully set.

IDA_MEM_NULL The `ida_mem` pointer is NULL.

Notes The default value is 10.

IDASsetNonlinConvCoef

Call `flag = IDASsetNonlinConvCoef(ida_mem, nlscoef);`

Description The function `IDASsetNonlinConvCoef` specifies the safety factor in the nonlinear convergence test; see Chapter 2, Eq. (2.8).

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`nlscoef` (`realtype`) coefficient in nonlinear convergence test (> 0.0).

Return value The return value `flag` (of type `int`) is one of

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.

IDASetSuppressAlg

Call	<code>flag = IDASetSuppressAlg(ida_mem, suppressalg);</code>
Description	The function <code>IDASetSuppressAlg</code> indicates whether or not to suppress algebraic variables in the local error test.
Arguments	<code>ida_mem</code> (void *) pointer to the IDAS memory block. <code>suppressalg</code> (boolean type) indicates whether to suppress (TRUE) or not (FALSE) the algebraic variables in the local error test.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDA_SUCCESS</code> The optional value has been successfully set. <code>IDA_MEM_NULL</code> The <code>ida_mem</code> pointer is NULL.
Notes	The default value is FALSE. If <code>suppressalg=TRUE</code> is selected, then the <code>id</code> vector must be set (through <code>IDASetId</code>) to specify the algebraic components. In general, the use of this option (with <code>suppressalg = TRUE</code>) is <i>discouraged</i> when solving DAE systems of index 1, whereas it is generally <i>encouraged</i> for systems of index 2 or more. See pp. 146-147 of Ref. [3] for more on this issue.

IDASetId

Call	<code>flag = IDASetId(ida_mem, id);</code>
Description	The function <code>IDASetId</code> specifies algebraic/differential components in the y vector.
Arguments	<code>ida_mem</code> (void *) pointer to the IDAS memory block. <code>id</code> (N_Vector) state vector. A value of 1.0 indicates a differential variable, while 0.0 indicates an algebraic variable.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDA_SUCCESS</code> The optional value has been successfully set. <code>IDA_MEM_NULL</code> The <code>ida_mem</code> pointer is NULL.
Notes	The vector <code>id</code> is required if the algebraic variables are to be suppressed from the local error test (see <code>IDASetSuppressAlg</code>) or if <code>IDACalcIC</code> is to be called with <code>icopt = IDA_YA_YDP_INIT</code> (see §4.5.4).

IDASetConstraints

Call	<code>flag = IDASetConstraints(ida_mem, constraints);</code>
Description	The function <code>IDASetConstraints</code> specifies a vector defining inequality constraints for each component of the solution vector y .
Arguments	<code>ida_mem</code> (void *) pointer to the IDAS memory block. <code>constraints</code> (N_Vector) vector of constraint flags. If <code>constraints[i]</code> is 0.0 then no constraint is imposed on y_i . 1.0 then y_i will be constrained to be $y_i \geq 0.0$. -1.0 then y_i will be constrained to be $y_i \leq 0.0$. 2.0 then y_i will be constrained to be $y_i > 0.0$. -2.0 then y_i will be constrained to be $y_i < 0.0$.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDA_SUCCESS</code> The optional value has been successfully set. <code>IDA_MEM_NULL</code> The <code>ida_mem</code> pointer is NULL. <code>IDA_ILL_INPUT</code> 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` will result in an illegal input return.

4.5.7.2 Dense direct linear solvers optional input functions

The IDADENSE solver needs a function to compute a dense approximation to the Jacobian matrix $J(t, y, \dot{y})$. This function must be of type `IDADlsDenseJacFn`. The user can supply his/her own dense Jacobian function, or use the default internal difference quotient approximation that comes with the IDADENSE solver. To specify a user-supplied Jacobian function `djac`, IDADENSE provides the function `IDADlsSetDenseJacFn`. The IDADENSE solver passes the pointer `user_data` to the dense Jacobian function. This allows the user to create an arbitrary structure with relevant problem data and access it during the execution of the user-supplied Jacobian function, without using global data in the program. The pointer `user_data` may be specified through `IDASetUserData`.

`IDADlsSetDenseJacFn`

Call `flag = IDADlsSetDenseJacFn(ida_mem, djac);`

Description The function `IDADlsSetDenseJacFn` specifies the dense Jacobian approximation function to be used.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.
 `djac` (`IDADlsDenseJacFn`) user-defined dense Jacobian approximation function.

Return value The return value `flag` (of type `int`) is one of

`IDADLS_SUCCESS` The optional value has been successfully set.
`IDADLS_MEM_NULL` The `ida_mem` pointer is NULL.
`IDADLS_LMEM_NULL` The IDADENSE linear solver has not been initialized.

Notes By default, IDADENSE uses an internal difference quotient function. If NULL is passed to `djac`, this default function is used.

The function type `IDADlsDenseJacFn` is described in §4.6.5.

The IDABAND solver needs a function to compute a banded approximation to the Jacobian matrix $J(t, y, \dot{y})$. This function must be of type `IDADlsBandJacFn`. The user can supply his/her own banded Jacobian approximation function, or use the default difference quotient function that comes with the IDABAND solver. To specify a user-supplied Jacobian function `bjac`, IDABAND provides the function `IDADlsSetBandJacFn`. The IDABAND solver passes the pointer `user_data` to the banded Jacobian approximation function. This allows the user to create an arbitrary structure with relevant problem data and access it during the execution of the user-supplied Jacobian function, without using global data in the program. The pointer `user_data` may be specified through `IDASetUserData`.

`IDADlsSetBandJacFn`

Call `flag = IDADlsSetBandJacFn(ida_mem, bjac);`

Description The function `IDADlsSetBandJacFn` specifies the banded Jacobian approximation function to be used.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.
 `bjac` (`IDADlsBandJacFn`) user-defined banded Jacobian approximation function.

Return value The return value `flag` (of type `int`) is one of

`IDADLS_SUCCESS` The optional value has been successfully set.
`IDADLS_MEM_NULL` The `ida_mem` pointer is NULL.
`IDADLS_LMEM_NULL` The IDABAND linear solver has not been initialized.

Notes By default, IDABAND uses an internal difference quotient function. If NULL is passed to `bjac`, this default function is used.
The function type `IDADlsBandJacFn` is described in §4.6.6.

4.5.7.3 Sparse direct linear solvers optional input functions

The IDAKLU and IDASUPERLUMT solvers require a function to compute a compressed-sparse-column approximation of the Jacobian matrix $J(t, y, \dot{y})$. This function must be of type `IDASlsSparseJacFn`. The user must supply a custom sparse Jacobian function since a difference quotient approximation would not leverage the underlying sparse matrix structure of the problem. To specify a user-supplied Jacobian function `sjac`, IDAKLU and IDASUPERLUMT provide the function `IDASlsSetSparseJacFn`. The IDAKLU and IDASUPERLUMT solvers pass the pointer `user_data` to the sparse Jacobian function. This mechanism allows the user to create an arbitrary structure with relevant problem data and access it during the execution of the user-supplied Jacobian function, without using global data in the program. The pointer `user_data` may be specified through `IDASetUserData`.

`IDASlsSetSparseJacFn`

Call `flag = IDASlsSetSparseJacFn(ida_mem, sjac);`
 Description The function `IDASlsSetSparseJacFn` specifies the sparse Jacobian approximation function to be used.
 Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
 `sjac` (`IDASlsSparseJacFn`) user-defined sparse Jacobian approximation function.
 Return value The return value `flag` (of type `int`) is one of
 `IDASLS_SUCCESS` The optional value has been successfully set.
 `IDASLS_MEM_NULL` The `ida_mem` pointer is NULL.
 `IDASLS_LMEM_NULL` The IDAKLU or IDASUPERLUMT linear solver has not been initialized.
 Notes The function type `IDASlsSparseJacFn` is described in §4.6.7.

When using a sparse direct solver, there may be instances when the number of state variables does not change, but the number of nonzeros in the Jacobian does change. For the IDAKLU solver, we provide a reinitialization. This reinitialization routine reinitializes the Jacobian matrix memory for the new number of nonzeros and sets 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).

`IDAKLUReInit`

Call `flag = IDAKLUReInit(ida_mem, n, nnz, reinit_type);`
 Description The function `IDAKLUReInit` reinitializes Jacobian matrix memory and flags for new symbolic and numeric KLU factorizations.
 Arguments `ida_mem` (`void *`) pointer to the IDA memory block.
 `n` (`int`) number of state variables in the system.
 `nnz` (`int`) number of nonzeros in the Jacobian matrix.
 `reinit_type` (`int`) type of reinitialization:
 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 prior call to IDAKLU.

Return value The return value `flag` (of type `int`) is one of

`IDASLS_SUCCESS` The reinitialization succeeded.
`IDASLS_MEM_NULL` The `ida_mem` pointer is `NULL`.
`IDASLS_LMEM_NULL` The IDAKLU linear solver has not been initialized.
`IDASLS_ILL_INPUT` The given `reinit_type` has an illegal value.
`IDASLS_MEM_FAIL` A memory allocation failed.

Notes The default value for `reinit_type` is 2.

Both the IDAKLU and IDASUPERLUMT solvers can apply reordering algorithms to minimize fill-in for the resulting sparse LU decomposition internal to the solver. The approximate minimal degree ordering for nonsymmetric matrices given by the COLAMD algorithm is the default algorithm used within both solvers, but alternate orderings may be chosen through one of the following two functions. The input values to these functions are the numeric values used in the respective packages, and the user-supplied value will be passed directly to the package.

`IDAKLUSetOrdering`

Call `flag = IDAKLUSetOrdering(ida_mem, ordering_choice);`

Description The function `IDAKLUSetOrdering` specifies the ordering algorithm used by IDAKLU for reducing fill.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`ordering_choice` (`int`) flag denoting algorithm choice:

0 AMD
1 COLAMD
2 natural ordering

Return value The return value `flag` (of type `int`) is one of

`IDASLS_SUCCESS` The optional value has been successfully set.
`IDASLS_MEM_NULL` The `ida_mem` pointer is `NULL`.
`IDASLS_ILL_INPUT` The supplied value of `ordering_choice` is illegal.

Notes The default ordering choice is 1 for COLAMD.

`IDASuperLUMTSetOrdering`

Call `flag = IDASuperLUMTSetOrdering(ida_mem, ordering_choice);`

Description The function `IDASuperLUMTSetOrdering` specifies the ordering algorithm used by IDASUPERLUMT for reducing fill.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`ordering_choice` (`int`) flag denoting algorithm choice:

0 natural ordering
1 minimal degree ordering on $J^T J$
2 minimal degree ordering on $J^T + J$
3 COLAMD

Return value The return value `flag` (of type `int`) is one of

`IDASLS_SUCCESS` The optional value has been successfully set.
`IDASLS_MEM_NULL` The `ida_mem` pointer is `NULL`.
`IDASLS_ILL_INPUT` The supplied value of `ordering_choice` is illegal.

Notes The default ordering choice is 3 for COLAMD.

4.5.7.4 Iterative linear solvers optional input functions

If preconditioning is to be done with one of the IDASPILS linear solvers, then the user must supply a preconditioner solve function `psolve` and specify its name through a call to `IDASpilsSetPreconditioner`.

The evaluation and preprocessing of any Jacobian-related data needed by the user's preconditioner solve function is done in the optional user-supplied function `psetup`. Both of these functions are fully specified in §4.6. If used, the name of the `psetup` function should be specified in the call to `IDASpilsSetPreconditioner`.

The pointer `user_data` received through `IDASetUserData` (or a pointer to `NULL` if `user_data` was not specified) is passed to the preconditioner `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.

The IDASPILS solvers require a function to compute an approximation to the product between the Jacobian matrix $J(t, y)$ and a vector v . The user can supply his/her own Jacobian-times-vector approximation function, or use the default internal difference quotient function that comes with the IDASPILS solvers. A user-defined Jacobian-vector function must be of type `IDASpilsJacTimesVecFn` and can be specified through a call to `IDASpilsSetJacTimesVecFn` (see §4.6.8 for specification details). As with the preconditioner user-supplied functions, a pointer to the user-defined data structure, `user_data`, specified through `IDASetUserData` (or a `NULL` pointer otherwise) is passed to the Jacobian-times-vector function `jtimes` each time it is called.


`IDASpilsSetPreconditioner`

Call	<code>flag = IDASpilsSetPreconditioner(ida_mem, psetup, psolve);</code>
Description	The function <code>IDASpilsSetPreconditioner</code> specifies the preconditioner setup and solve functions.
Arguments	<code>ida_mem</code> (void *) pointer to the IDAS memory block. <code>psetup</code> (<code>IDASpilsPrecSetupFn</code>) user-defined preconditioner setup function. Pass <code>NULL</code> if no setup is to be done. <code>psolve</code> (<code>IDASpilsPrecSolveFn</code>) user-defined preconditioner solve function.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDASPILS_SUCCESS</code> The optional values have been successfully set. <code>IDASPILS_MEM_NULL</code> The <code>ida_mem</code> pointer is <code>NULL</code> . <code>IDASPILS_LMEM_NULL</code> The IDASPILS linear solver has not been initialized.
Notes	The function type <code>IDASpilsPrecSolveFn</code> is described in §4.6.9. The function type <code>IDASpilsPrecSetupFn</code> is described in §4.6.10.


`IDASpilsSetJacTimesVecFn`

Call	<code>flag = IDASpilsSetJacTimesVecFn(ida_mem, jtimes);</code>
Description	The function <code>IDASpilsSetJacTimesFn</code> specifies the Jacobian-vector function to be used.
Arguments	<code>ida_mem</code> (void *) pointer to the IDAS memory block. <code>jtimes</code> (<code>IDASpilsJacTimesVecFn</code>) user-defined Jacobian-vector product function.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDASPILS_SUCCESS</code> The optional value has been successfully set. <code>IDASPILS_MEM_NULL</code> The <code>ida_mem</code> pointer is <code>NULL</code> . <code>IDASPILS_LMEM_NULL</code> The IDASPILS linear solver has not been initialized.
Notes	By default, the IDASPILS solvers use the difference quotient function. If <code>NULL</code> is passed to <code>jtimes</code> , this default function is used. The function type <code>IDASpilsJacTimesVecFn</code> is described in §4.6.8.

IDASpilsSetGSType

Call	<code>flag = IDASpilsSetGSType(ida_mem, gstype);</code>
Description	The function <code>IDASpilsSetGSType</code> specifies the Gram-Schmidt orthogonalization to be used. This must be one of the enumeration constants <code>MODIFIED_GS</code> or <code>CLASSICAL_GS</code> . These correspond to using modified Gram-Schmidt and classical Gram-Schmidt, respectively.
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block. <code>gstype</code> (<code>int</code>) type of Gram-Schmidt orthogonalization.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDASPILS_SUCCESS</code> The optional value has been successfully set. <code>IDASPILS_MEM_NULL</code> The <code>ida_mem</code> pointer is <code>NULL</code> . <code>IDASPILS_LMEM_NULL</code> The IDASPILS linear solver has not been initialized. <code>IDASPILS_ILL_INPUT</code> The value of <code>gstype</code> is not valid.
Notes	The default value is <code>MODIFIED_GS</code> .  This option is available only for the IDASPGMR linear solver.

IDASpilsSetMaxRestarts

Call	<code>flag = IDASpilsSetMaxRestarts(ida_mem, maxrs);</code>
Description	The function <code>IDASpilsSetMaxRestarts</code> specifies the maximum number of restarts to be used in the GMRES algorithm.
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block. <code>maxrs</code> (<code>int</code>) maximum number of restarts.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDASPILS_SUCCESS</code> The optional value has been successfully set. <code>IDASPILS_MEM_NULL</code> The <code>ida_mem</code> pointer is <code>NULL</code> . <code>IDASPILS_LMEM_NULL</code> The IDASPILS linear solver has not been initialized. <code>IDASPILS_ILL_INPUT</code> The <code>maxrs</code> argument is negative.
Notes	The default value is 5. Pass <code>maxrs = 0</code> to specify no restarts.  This option is available only for the IDASPGMR linear solver.

IDASpilsSetEpsLin

Call	<code>flag = IDASpilsSetEpsLin(ida_mem, eplifac);</code>
Description	The function <code>IDASpilsSetEpsLin</code> specifies the factor by which the Krylov linear solver's convergence test constant is reduced from the Newton iteration test constant. (See §2.1).
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block. <code>eplifac</code> (<code>realtype</code>) linear convergence safety factor (≥ 0.0).
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDASPILS_SUCCESS</code> The optional value has been successfully set. <code>IDASPILS_MEM_NULL</code> The <code>ida_mem</code> pointer is <code>NULL</code> . <code>IDASPILS_LMEM_NULL</code> The IDASPILS linear solver has not been initialized. <code>IDASPILS_ILL_INPUT</code> The value of <code>eplifac</code> is negative.
Notes	The default value is 0.05. Passing a value <code>eplifac = 0.0</code> also indicates using the default value.

IDASpilsSetIncrementFactor

Call	<code>flag = IDASpilsSetIncrementFactor(ida_mem, dqincfac);</code>
Description	The function <code>IDASpilsSetIncrementFactor</code> specifies a factor in the increments to y used in the difference quotient approximations to the Jacobian-vector products. (See §2.1). The increment used to approximate Jv will be $\sigma = \text{dqincfac}/\ v\ $.
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block. <code>dqincfac</code> (<code>realtype</code>) difference quotient increment factor.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDASPILS_SUCCESS</code> The optional value has been successfully set. <code>IDASPILS_MEM_NULL</code> The <code>ida_mem</code> pointer is <code>NULL</code> . <code>IDASPILS_LMEM_NULL</code> The IDASPILS linear solver has not been initialized. <code>IDASPILS_ILL_INPUT</code> The increment factor was non-positive.
Notes	The default value is <code>dqincfac = 1.0</code> .

IDASpilsSetMaxl

Call	<code>flag = IDASpilsSetMaxl(ida_mem, maxl);</code>
Description	The function <code>IDASpilsSetMaxl</code> resets the maximum Krylov subspace dimension for the Bi-CGStab or TFQMR methods.
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block. <code>maxl</code> (<code>int</code>) maximum dimension of the Krylov subspace.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDASPILS_SUCCESS</code> The optional value has been successfully set. <code>IDASPILS_MEM_NULL</code> The <code>ida_mem</code> pointer is <code>NULL</code> . <code>IDASPILS_LMEM_NULL</code> The IDASPILS linear solver has not been initialized.
Notes	The maximum subspace dimension is initially specified in the call to the linear solver specification function (see §4.5.3). This function call is needed only if <code>maxl</code> is being changed from its previous value. An input value <code>maxl ≤ 0</code> will result in the default value, 5. This option is available only for the IDASPBCG and IDASPTFQMR linear solvers.

**4.5.7.5 Initial condition calculation optional input functions**

The following functions can be called just prior to calling `IDACalcIC` to set optional inputs controlling the initial condition calculation.

IDASetNonlinConvCoefIC

Call	<code>flag = IDASetNonlinConvCoefIC(ida_mem, epiccon);</code>
Description	The function <code>IDASetNonlinConvCoefIC</code> specifies the positive constant in the Newton iteration convergence test within the initial condition calculation.
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block. <code>epiccon</code> (<code>realtype</code>) coefficient in the Newton convergence test (> 0).
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDA_SUCCESS</code> The optional value has been successfully set. <code>IDA_MEM_NULL</code> The <code>ida_mem</code> pointer is <code>NULL</code> . <code>IDA_ILL_INPUT</code> The <code>epiccon</code> 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.

IDASetMaxNumStepsIC

Call `flag = IDASetMaxNumStepsIC(ida_mem, maxnh);`

Description 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, $J = \partial F / \partial y + (1/h) \partial F / \partial \dot{y}$.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.
 `maxnh` (int) maximum allowed number of values for h .

Return value The return value `flag` (of type `int`) is one of

`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.

IDASetMaxNumJacsIC

Call `flag = IDASetMaxNumJacsIC(ida_mem, maxnj);`

Description 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` (void *) pointer to the IDAS memory block.
 `maxnj` (int) maximum allowed number of Jacobian or preconditioner evaluations.

Return value The return value `flag` (of type `int`) is one of

`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.

IDASetMaxNumItersIC

Call `flag = IDASetMaxNumItersIC(ida_mem, maxnit);`

Description 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` (void *) pointer to the IDAS memory block.
 `maxnit` (int) maximum number of Newton iterations.

Return value The return value `flag` (of type `int`) is one of

`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.

IDASetLineSearchOffIC

Call	<code>flag = IDASetLineSearchOffIC(ida_mem, lsoff);</code>
Description	The function <code>IDASetLineSearchOffIC</code> specifies whether to turn on or off the linesearch algorithm.
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block. <code>lsoff</code> (<code>boolean_t</code>) a flag to turn off (<code>TRUE</code>) or keep (<code>FALSE</code>) the linesearch algorithm.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDA_SUCCESS</code> The optional value has been successfully set. <code>IDA_MEM_NULL</code> The <code>ida_mem</code> pointer is <code>NULL</code> .
Notes	The default value is <code>FALSE</code> .

IDASetStepToleranceIC

Call	<code>flag = IDASetStepToleranceIC(ida_mem, steptol);</code>
Description	The function <code>IDASetStepToleranceIC</code> specifies a positive lower bound on the Newton step.
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block. <code>steptol</code> (<code>int</code>) Minimum allowed WRMS-norm of the Newton step (> 0.0).
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDA_SUCCESS</code> The optional value has been successfully set. <code>IDA_MEM_NULL</code> The <code>ida_mem</code> pointer is <code>NULL</code> . <code>IDA_ILL_INPUT</code> The <code>steptol</code> tolerance is ≤ 0.0 .
Notes	The default value is $(\text{unit roundoff})^{2/3}$.

4.5.7.6 Rootfinding optional input functions

The following functions can be called to set optional inputs to control the rootfinding algorithm.

IDASetRootDirection

Call	<code>flag = IDASetRootDirection(ida_mem, rootdir);</code>
Description	The function <code>IDASetRootDirection</code> specifies the direction of zero-crossings to be located and returned to the user.
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block. <code>rootdir</code> (<code>int *</code>) state array of length <code>nrtfn</code> , the number of root functions g_i , as specified in the call to the function <code>IDARootInit</code> . A value of 0 for <code>rootdir[i]</code> indicates that crossing in either direction should be reported for g_i . A value of +1 or -1 indicates that the solver should report only zero-crossings where g_i is increasing or decreasing, respectively.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDA_SUCCESS</code> The optional value has been successfully set. <code>IDA_MEM_NULL</code> The <code>ida_mem</code> pointer is <code>NULL</code> . <code>IDA_ILL_INPUT</code> rootfinding has not been activated through a call to <code>IDARootInit</code> .
Notes	The default behavior is to locate both zero-crossing directions.

IDASetNoInactiveRootWarn

Call	<code>flag = IDASetNoInactiveRootWarn(ida_mem);</code>
Description	The function <code>IDASetNoInactiveRootWarn</code> disables issuing a warning if some root function appears to be identically zero at the beginning of the integration.
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDA_SUCCESS</code> The optional value has been successfully set. <code>IDA_MEM_NULL</code> The <code>ida_mem</code> pointer is <code>NULL</code> .
Notes	IDAS will not report the initial conditions as a possible zero-crossing (assuming that one or more components g_i are zero at the initial time). However, if it appears that some g_i is identically zero at the initial time (i.e., g_i is zero at the initial time and after the first step), IDAS will issue a warning which can be disabled with this optional input function.

4.5.8 Interpolated output function

An optional function `IDAGetDky` is available to obtain additional output values. This function must be called after a successful return from `IDASolve` and provides interpolated values of y or its derivatives of order up to the last internal order used for any value of t in the last internal step taken by IDAS.

The call to the `IDAGetDky` function has the following form:

IDAGetDky

Call	<code>flag = IDAGetDky(ida_mem, t, k, dky);</code>
Description	The function <code>IDAGetDky</code> computes the interpolated values of the k^{th} derivative of y for any value of t in the last internal step taken by IDAS. The value of k must be non-negative and smaller than the last internal order used. A value of 0 for k means that the y is interpolated. The value of t must satisfy $t_n - h_u \leq t \leq t_n$, where t_n denotes the current internal time reached, and h_u is the last internal step size used successfully.
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block. <code>t</code> (<code>realtype</code>) time at which to interpolate. <code>k</code> (<code>int</code>) integer specifying the order of the derivative of y wanted. <code>dky</code> (<code>N_Vector</code>) vector containing the interpolated k^{th} derivative of $y(t)$.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDA_SUCCESS</code> <code>IDAGetDky</code> succeeded. <code>IDA_MEM_NULL</code> The <code>ida_mem</code> argument was <code>NULL</code> . <code>IDA_BAD_T</code> <code>t</code> is not in the interval $[t_n - h_u, t_n]$. <code>IDA_BAD_K</code> <code>k</code> is not one of $\{0, 1, \dots, klast\}$. <code>IDA_BAD_DKY</code> <code>dky</code> is <code>NULL</code> .
Notes	It is only legal to call the function <code>IDAGetDky</code> after a successful return from <code>IDASolve</code> . Functions <code>IDAGetCurrentTime</code> , <code>IDAGetLastStep</code> and <code>IDAGetLastOrder</code> (see §4.5.9.1) can be used to access t_n , h_u and $klast$.

4.5.9 Optional output functions

IDAS provides an extensive list of functions that can be used to obtain solver performance information. Table 4.2 lists all optional output functions in IDAS, which are then described in detail in the remainder of this section.

Some of the optional outputs, especially the various counters, can be very useful in determining how successful the IDAS solver is in doing its job. For example, the counters `nsteps` and `nrevals` provide a rough measure of the overall cost of a given run, and can be compared among runs with

differing input options to suggest which set of options is most efficient. The ratio `nriters/nsteps` measures the performance of the Newton iteration in solving the nonlinear systems at each time step; typical values for this range from 1.1 to 1.8. The ratio `njevals/nriters` (in the case of a direct linear solver), and the ratio `npevals/nriters` (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/nriters` 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 `nlriters/nriters` measures the performance of the Krylov iterative linear solver, and thus (indirectly) the quality of the preconditioner.

4.5.9.1 Main solver optional output functions

IDAS provides several user-callable functions that can be used to obtain different quantities that may be of interest to the user, such as solver workspace requirements, solver performance statistics, as well as additional data from the IDAS memory block (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 IDAS nonlinear solver being used. As a convenience, additional extraction functions provide the optional outputs in groups. These optional output functions are described next.

IDAGetWorkSpace

Call `flag = IDAGetWorkSpace(ida_mem, &lenrw, &leniw);`

Description The function `IDAGetWorkSpace` returns the IDAS real and integer workspace sizes.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.

`lenrw` (`long int`) number of real values in the IDAS workspace.

`leniw` (`long int`) number of integer values in the IDAS workspace.

Return value The return value `flag` (of type `int`) is one of

`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 `nrtfn` of root functions (see §4.5.5), the actual size of the real workspace, in `realtype` 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: $\text{leniw} = \text{leniw} + N_i$;

Table 4.2: Optional outputs from IDAS, IDADLS, IDASLS, and IDASPILS

Optional output	Function name
IDAS main solver	
Size of IDAS real and integer workspace	IDAGetWorkSpace
Cumulative number of internal steps	IDAGetNumSteps
No. of calls to residual function	IDAGetNumResEvals
No. of calls to linear solver setup function	IDAGetNumLinSolvSetups
No. of local error test failures that have occurred	IDAGetNumErrTestFails
Order used during the last step	IDAGetLastOrder
Order to be attempted on the next step	IDAGetCurrentOrder
Order reductions due to stability limit detection	IDAGetNumStabLimOrderReds
Actual initial step size used	IDAGetActualInitStep
Step size used for the last step	IDAGetLastStep
Step size to be attempted on the next step	IDAGetCurrentStep
Current internal time reached by the solver	IDAGetCurrentTime
Suggested factor for tolerance scaling	IDAGetTolScaleFactor
Error weight vector for state variables	IDAGetErrWeights
Estimated local errors	IDAGetEstLocalErrors
No. of nonlinear solver iterations	IDAGetNumNonlinSolvIters
No. of nonlinear convergence failures	IDAGetNumNonlinSolvConvFails
Array showing roots found	IDAGetRootInfo
No. of calls to user root function	IDAGetNumGEvals
Name of constant associated with a return flag	IDAGetReturnFlagName
IDAS initial conditions calculation	
Number of backtrack operations	IDAGetNumBacktrackops
Corrected initial conditions	IDAGetConsistentIC
IDADLS linear solver	
Size of real and integer workspace	IDADlsGetWorkSpace
No. of Jacobian evaluations	IDADlsGetNumJacEvals
No. of residual calls for finite diff. Jacobian evals.	IDADlsGetNumResEvals
Last return from a linear solver function	IDADlsGetLastFlag
Name of constant associated with a return flag	IDADlsGetReturnFlagName
IDASLS linear solver	
No. of Jacobian evaluations	IDASlsGetNumJacEvals
Last return from a linear solver function	IDASlsGetLastFlag
Name of constant associated with a return flag	IDASlsGetReturnFlagName
IDASPILS linear solvers	
Size of real and integer workspace	IDASpilsGetWorkSpace
No. of linear iterations	IDASpilsGetNumLinIters
No. of linear convergence failures	IDASpilsGetNumConvFails
No. of preconditioner evaluations	IDASpilsGetNumPrecEvals
No. of preconditioner solves	IDASpilsGetNumPrecSolves
No. of Jacobian-vector product evaluations	IDASpilsGetNumJtimesEvals
No. of residual calls for finite diff. Jacobian-vector evals.	IDASpilsGetNumResEvals
Last return from a linear solver function	IDASpilsGetLastFlag
Name of constant associated with a return flag	IDASpilsGetReturnFlagName

where N_i is the number of integer words in one `N_Vector` (= 1 for `NVECTOR_SERIAL` and $2 \times \text{npes}$ for `NVECTOR_PARALLEL` on `npes` 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: `lenrw` = $55 + 11N$, `leniw` = 49.

Note that additional memory is allocated if quadratures and/or forward sensitivity integration is enabled. See §4.7.1 and §5.2.1 for more details.

IDAGetNumSteps

Call `flag = IDAGetNumSteps(ida_mem, &nsteps);`

Description The function `IDAGetNumSteps` returns the cumulative number of internal steps taken by the solver (total so far).

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`nsteps` (`long int`) number of steps taken by IDAS.

Return value The return value `flag` (of type `int`) is one of

`IDA_SUCCESS` The optional output value has been successfully set.

`IDA_MEM_NULL` The `ida_mem` pointer is NULL.

IDAGetNumResEvals

Call `flag = IDAGetNumResEvals(ida_mem, &nrevals);`

Description The function `IDAGetNumResEvals` returns the number of calls to the user's residual evaluation function.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`nrevals` (`long int`) number of calls to the user's `res` function.

Return value The return value `flag` (of type `int`) is one of

`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.

IDAGetNumLinSolvSetups

Call `flag = IDAGetNumLinSolvSetups(ida_mem, &nlinsetups);`

Description The function `IDAGetNumLinSolvSetups` returns the cumulative number of calls made to the linear solver's setup function (total so far).

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`nlinsetups` (`long int`) number of calls made to the linear solver setup function.

Return value The return value `flag` (of type `int`) is one of

`IDA_SUCCESS` The optional output value has been successfully set.

`IDA_MEM_NULL` The `ida_mem` pointer is NULL.

IDAGetNumErrTestFails

Call `flag = IDAGetNumErrTestFails(ida_mem, &netfails);`

Description The function `IDAGetNumErrTestFails` returns the cumulative number of local error test failures that have occurred (total so far).

Arguments `ida_mem` (void *) pointer to the IDAS memory block.
`netfails` (long int) number of error test failures.

Return value The return value `flag` (of type `int`) is one of
`IDA_SUCCESS` The optional output value has been successfully set.
`IDA_MEM_NULL` The `ida_mem` pointer is NULL.

IDAGetLastOrder

Call `flag = IDAGetLastOrder(ida_mem, &klast);`

Description The function `IDAGetLastOrder` returns the integration method order used during the last internal step.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.
`klast` (int) method order used on the last internal step.

Return value The return value `flag` (of type `int`) is one of
`IDA_SUCCESS` The optional output value has been successfully set.
`IDA_MEM_NULL` The `ida_mem` pointer is NULL.

IDAGetCurrentOrder

Call `flag = IDAGetCurrentOrder(ida_mem, &kcur);`

Description The function `IDAGetCurrentOrder` returns the integration method order to be used on the next internal step.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.
`kcur` (int) method order to be used on the next internal step.

Return value The return value `flag` (of type `int`) is one of
`IDA_SUCCESS` The optional output value has been successfully set.
`IDA_MEM_NULL` The `ida_mem` pointer is NULL.

IDAGetLastStep

Call `flag = IDAGetLastStep(ida_mem, &hlast);`

Description The function `IDAGetLastStep` returns the integration step size taken on the last internal step.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.
`hlast` (realtype) step size taken on the last internal step by IDA, or last artificial step size used in `IDACalcIC`, whichever was called last.

Return value The return value `flag` (of type `int`) is one of
`IDA_SUCCESS` The optional output value has been successfully set.
`IDA_MEM_NULL` The `ida_mem` pointer is NULL.

IDAGetCurrentStep

Call `flag = IDAGetCurrentStep(ida_mem, &hcur);`

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

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`hcur` (`realtype`) step size to be attempted on the next internal step.

Return value The return value `flag` (of type `int`) is one of
`IDA_SUCCESS` The optional output value has been successfully set.
`IDA_MEM_NULL` The `ida_mem` pointer is NULL.

IDAGetActualInitStep

Call `flag = IDAGetActualInitStep(ida_mem, &hinused);`

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

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`hinused` (`realtype`) actual value of initial step size.

Return value The return value `flag` (of type `int`) is one of
`IDA_SUCCESS` The optional output value has been successfully set.
`IDA_MEM_NULL` The `ida_mem` pointer is NULL.

Notes Even if the value of the initial integration step size was specified by the user through a call to `IDASetInitStep`, this value might have been changed by IDAS to ensure that the step size is within the prescribed bounds ($h_{\min} \leq h_0 \leq h_{\max}$), or to meet the local error test.

IDAGetCurrentTime

Call `flag = IDAGetCurrentTime(ida_mem, &tcur);`

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

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`tcur` (`realtype`) current internal time reached.

Return value The return value `flag` (of type `int`) is one of
`IDA_SUCCESS` The optional output value has been successfully set.
`IDA_MEM_NULL` The `ida_mem` pointer is NULL.

IDAGetTolScaleFactor

Call `flag = IDAGetTolScaleFactor(ida_mem, &tolsfac);`

Description 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` (`void *`) pointer to the IDAS memory block.
`tolsfac` (`realtype`) suggested scaling factor for user tolerances.

Return value The return value `flag` (of type `int`) is one of
`IDA_SUCCESS` The optional output value has been successfully set.
`IDA_MEM_NULL` The `ida_mem` pointer is NULL.


IDAGetErrWeights

Call `flag = IDAGetErrWeights(ida_mem, eweight);`

Description The function `IDAGetErrWeights` returns the solution error weights at the current time. These are the W_i given by Eq. (2.7) (or by the user's `IDAwtFn`).

Arguments `ida_mem` (void *) pointer to the IDAS memory block.
 `eweight` (N_Vector) solution error weights at the current time.

Return value The return value `flag` (of type `int`) is one of
 `IDA_SUCCESS` The optional output value has been successfully set.
 `IDA_MEM_NULL` The `ida_mem` pointer is NULL.

 Notes The user must allocate space for `eweight`.


IDAGetEstLocalErrors

Call `flag = IDAGetEstLocalErrors(ida_mem, ele);`

Description The function `IDAGetEstLocalErrors` returns the estimated local errors.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.
 `ele` (N_Vector) estimated local errors at the current time.

Return value The return value `flag` (of type `int`) is one of
 `IDA_SUCCESS` The optional output value has been successfully set.
 `IDA_MEM_NULL` The `ida_mem` pointer is NULL.

 Notes The user must allocate space for `ele`.
 The values returned in `ele` are only valid if `IDASolve` returned a non-negative value.
 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]`.

IDAGetIntegratorStats

Call `flag = IDAGetIntegratorStats(ida_mem, &nsteps, &nrevals, &nlinsetups, &netfails, &klast, &kcur, &hinused, &hlast, &hcur, &tcur);`

Description The function `IDAGetIntegratorStats` returns the IDAS integrator statistics as a group.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.
 `nsteps` (long int) cumulative number of steps taken by IDAS.
 `nrevals` (long int) cumulative number of calls to the user's `res` function.
 `nlinsetups` (long int) cumulative number of calls made to the linear solver setup function.
 `netfails` (long int) cumulative number of error test failures.
 `klast` (int) method order used on the last internal step.
 `kcur` (int) method order to be used on the next internal step.
 `hinused` (realtype) actual value of initial step size.
 `hlast` (realtype) step size taken on the last internal step.
 `hcur` (realtype) step size to be attempted on the next internal step.
 `tcur` (realtype) current internal time reached.

Return value The return value `flag` (of type `int`) is one of

- `IDA_SUCCESS` the optional output values have been successfully set.
- `IDA_MEM_NULL` the `ida_mem` pointer is `NULL`.

IDAGetNumNonlinSolvIters

Call `flag = IDAGetNumNonlinSolvIters(ida_mem, &nniters);`

Description The function `IDAGetNumNonlinSolvIters` returns the cumulative number of nonlinear (functional or Newton) iterations performed.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`nniters` (`long int`) number of nonlinear iterations performed.

Return value The return value `flag` (of type `int`) is one of

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

IDAGetNumNonlinSolvConvFails

Call `flag = IDAGetNumNonlinSolvConvFails(ida_mem, &nncfails);`

Description The function `IDAGetNumNonlinSolvConvFails` returns the cumulative number of nonlinear convergence failures that have occurred.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`nncfails` (`long int`) number of nonlinear convergence failures.

Return value The return value `flag` (of type `int`) is one of

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

IDAGetNonlinSolvStats

Call `flag = IDAGetNonlinSolvStats(ida_mem, &nniters, &nncfails);`

Description The function `IDAGetNonlinSolvStats` returns the IDAS nonlinear solver statistics as a group.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`nniters` (`long int`) cumulative number of nonlinear iterations performed.
`nncfails` (`long int`) cumulative number of nonlinear convergence failures.

Return value The return value `flag` (of type `int`) is one of

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

IDAGetReturnFlagName

Call `name = IDAGetReturnFlagName(flag);`

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

Arguments The only argument, of type `int`, is a return flag from an IDAS function.

Return value The return value is a string containing the name of the corresponding constant.

4.5.9.2 Initial condition calculation optional output functions

IDAGetNumBcktrackOps

Call `flag = IDAGetNumBacktrackOps(ida_mem, &nbacktr);`

Description The function `IDAGetNumBacktrackOps` returns the number of backtrack operations done in the linesearch algorithm in `IDACalcIC`.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`nbacktr` (`long int`) the cumulative number of backtrack operations.

Return value The return value `flag` (of type `int`) is one of
`IDA_SUCCESS` The optional output value has been successfully set.
`IDA_MEM_NULL` The `ida_mem` pointer is NULL.

IDAGetConsistentIC


Call `flag = IDAGetConsistentIC(ida_mem, yy0_mod, yp0_mod);`

Description The function `IDAGetConsistentIC` returns the corrected initial conditions calculated by `IDACalcIC`.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`yy0_mod` (`N_Vector`) consistent solution vector.
`yp0_mod` (`N_Vector`) consistent derivative vector.

Return value The return value `flag` (of type `int`) is one of
`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.

 The user must allocate space for `yy0_mod` and `yp0_mod` (if not NULL).

4.5.9.3 Rootfinding optional output functions

There are two optional output functions associated with rootfinding.

IDAGetRootInfo


Call `flag = IDAGetRootInfo(ida_mem, rootsfound);`

Description The function `IDAGetRootInfo` returns an array showing which functions were found to have a root.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`rootsfound` (`int *`) array of length `nrtfn` with the indices of the user functions g_i found to have a root. For $i = 0, \dots, \text{nrtfn} - 1$, `rootsfound[i]` $\neq 0$ if g_i has a root, and $= 0$ if not.

Return value The return value `flag` (of type `int`) is one of
`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 .

 The user must allocate memory for the vector `rootsfound`.

IDAGetNumGEvals

Call	<code>flag = IDAGetNumGEvals(ida_mem, &ngevals);</code>
Description	The function <code>IDAGetNumGEvals</code> returns the cumulative number of calls to the user root function <code>g</code> .
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block. <code>ngevals</code> (<code>long int</code>) number of calls to the user's function <code>g</code> so far.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDA_SUCCESS</code> The optional output value has been successfully set. <code>IDA_MEM_NULL</code> The <code>ida_mem</code> pointer is NULL.

4.5.9.4 Dense direct linear solvers optional output functions

The following optional outputs are available from the IDADLS modules: workspace requirements, number of calls to the Jacobian routine, number of calls to the residual routine for finite-difference Jacobian approximation, and last return value from an IDADLS function. Note that, where the name of an output would otherwise conflict with the name of an optional output from the main solver, a suffix LS (for Linear Solver) has been added here (e.g. `lenrwLS`).

IDADlsGetWorkSpace

Call	<code>flag = IDADlsGetWorkSpace(ida_mem, &lenrwLS, &leniwLS);</code>
Description	The function <code>IDADlsGetWorkSpace</code> returns the sizes of the real and integer workspaces used by an IDADLS linear solver (<code>IDADENSE</code> or <code>IDABAND</code>).
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block. <code>lenrwLS</code> (<code>long int</code>) the number of real values in the IDADLS workspace. <code>leniwLS</code> (<code>long int</code>) the number of integer values in the IDADLS workspace.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDADLS_SUCCESS</code> The optional output value has been successfully set. <code>IDADLS_MEM_NULL</code> The <code>ida_mem</code> pointer is NULL. <code>IDADLS_LMEM_NULL</code> The IDADLS linear solver has not been initialized.
Notes	For the <code>IDADENSE</code> linear solver, in terms of the problem size N , the actual size of the real workspace is $2N^2$ <code>realtype</code> words, while the actual size of the integer workspace is N integer words. For the <code>IDABAND</code> linear solver, in terms of N and Jacobian half-bandwidths, the actual size of the real workspace is $N(2 \text{ mupper} + 3 \text{ mlower} + 2)$ <code>realtype</code> words, while the actual size of the integer workspace is N integer words.

IDADlsGetNumJacEvals

Call	<code>flag = IDADlsGetNumJacEvals(ida_mem, &njevals);</code>
Description	The function <code>IDADlsGetNumJacEvals</code> returns the cumulative number of calls to the IDADLS (dense or banded) Jacobian approximation function.
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block. <code>njevals</code> (<code>long int</code>) the cumulative number of calls to the Jacobian function (total so far).
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDADLS_SUCCESS</code> The optional output value has been successfully set. <code>IDADLS_MEM_NULL</code> The <code>ida_mem</code> pointer is NULL. <code>IDADLS_LMEM_NULL</code> The <code>IDADENSE</code> linear solver has not been initialized.

IDADlsGetNumResEvals

Call	<code>flag = IDADlsGetNumResEvals(ida_mem, &nrevalsLS);</code>
Description	The function <code>IDADlsGetNumResEvals</code> returns the cumulative number of calls to the user residual function due to the finite difference (dense or band) Jacobian approximation.
Arguments	<code>ida_mem</code> (void *) pointer to the IDAS memory block. <code>nrevalsLS</code> (long int) the cumulative number of calls to the user residual function.
Return value	The return value <code>flag</code> (of type int) is one of IDADLS_SUCCESS The optional output value has been successfully set. IDADLS_MEM_NULL The <code>ida_mem</code> pointer is NULL. IDADLS_LMEM_NULL The IDADENSE linear solver has not been initialized.
Notes	The value <code>nrevalsLS</code> is incremented only if the default internal difference quotient function is used.

IDADlsGetLastFlag

Call	<code>flag = IDADlsGetLastFlag(ida_mem, &lsflag);</code>
Description	The function <code>IDADlsGetLastFlag</code> returns the last return value from an IDADLS routine.
Arguments	<code>ida_mem</code> (void *) pointer to the IDAS memory block. <code>lsflag</code> (long int) the value of the last return flag from an IDADLS function.
Return value	The return value <code>flag</code> (of type int) is one of IDADLS_SUCCESS The optional output value has been successfully set. IDADLS_MEM_NULL The <code>ida_mem</code> pointer is NULL. IDADLS_LMEM_NULL The IDADENSE linear solver has not been initialized.
Notes	If the IDADENSE setup function failed (i.e., <code>IDASolve</code> returned <code>IDA_LSETUP_FAIL</code>), the value <code>lsflag</code> 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 band) Jacobian matrix. For all other failures, the value of <code>lsflag</code> is negative.

IDADlsGetReturnFlagName

Call	<code>name = IDADlsGetReturnFlagName(lsflag);</code>
Description	The function <code>IDADlsGetReturnFlagName</code> returns the name of the IDADLS constant corresponding to <code>lsflag</code> .
Arguments	The only argument, of type long int, is a return flag from an IDADLS function.
Return value	The return value is a string containing the name of the corresponding constant. If $1 \leq \text{lsflag} \leq N$ (LU factorization failed), this function returns "NONE".

4.5.9.5 Sparse direct linear solvers optional output functions

The following optional outputs are available from the IDASLS modules: number of calls to the Jacobian routine and last return value from an IDASLS function.

IDASlsGetNumJacEvals

Call	<code>flag = IDASlsGetNumJacEvals(ida_mem, &njevals);</code>
Description	The function <code>IDASlsGetNumJacEvals</code> returns the cumulative number of calls to the IDASLS sparse Jacobian approximation function.
Arguments	<code>ida_mem</code> (void *) pointer to the IDAS memory block.

njevals (long int) the cumulative number of calls to the Jacobian function (total so far).

Return value The return value `flag` (of type `int`) is one of

`IDASLS_SUCCESS` The optional output value has been successfully set.
`IDASLS_MEM_NULL` The `ida_mem` pointer is `NULL`.
`IDASLS_LMEM_NULL` The IDASLS linear solver has not been initialized.

IDASlsGetLastFlag

Call `flag = IDASlsGetLastFlag(ida_mem, &lsflag);`

Description The function `IDASlsGetLastFlag` returns the last return value from an IDASLS routine.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.

`lsflag` (`long int`) the value of the last return flag from an IDASLS function.

Return value The return value `flag` (of type `int`) is one of

`IDASLS_SUCCESS` The optional output value has been successfully set.
`IDASLS_MEM_NULL` The `ida_mem` pointer is `NULL`.
`IDASLS_LMEM_NULL` The IDASLS linear solver has not been initialized.

Notes

IDASlsGetReturnFlagName

Call `name = IDASlsGetReturnFlagName(lsflag);`

Description The function `IDASlsGetReturnFlagName` returns the name of the IDASLS constant corresponding to `lsflag`.

Arguments The only argument, of type `long int`, is a return flag from an IDASLS function.

Return value The return value is a string containing the name of the corresponding constant.

4.5.9.6 Iterative linear solvers optional output functions

The following optional outputs are available from the IDASPILS modules: workspace requirements, number of linear iterations, number of linear convergence failures, number of calls to the preconditioner setup and solve routines, number of calls to the Jacobian-vector product routine, number of calls to the residual routine for finite-difference Jacobian-vector product approximation, and last return value from a linear solver function. Note that, where the name of an output would otherwise conflict with the name of an optional output from the main solver, a suffix `LS` (for Linear Solver) has been added here (e.g. `lenrwLS`).

IDASpilsGetWorkSpace

Call `flag = IDASpilsGetWorkSpace(ida_mem, &lenrwLS, &leniwLS);`

Description The function `IDASpilsGetWorkSpace` returns the global sizes of the IDASPILS real and integer workspaces.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.

`lenrwLS` (`long int`) global number of real values in the IDASPILS workspace.

`leniwLS` (`long int`) global number of integer values in the IDASPILS workspace.

Return value The return value `flag` (of type `int`) is one of

`IDASPILS_SUCCESS` The optional output value has been successfully set.
`IDASPILS_MEM_NULL` The `ida_mem` pointer is `NULL`.
`IDASPILS_LMEM_NULL` The IDASPILS linear solver has not been initialized.

Notes In terms of the problem size N and maximum subspace size `maxl`, the actual size of the real workspace is roughly:
 $N * (\text{maxl} + 5) + \text{maxl} * (\text{maxl} + 4) + 1$ `realtype` words for IDASPGMR,
 $10 * N$ `realtype` words for IDASPBCG,
and $13 * N$ `realtype` words for IDASPTFQMR.
In a parallel setting, the above values are global, summed over all processors.

IDASpilsGetNumLinIters

Call `flag = IDASpilsGetNumLinIters(ida_mem, &nliters);`

Description The function `IDASpilsGetNumLinIters` returns the cumulative number of linear iterations.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`nliters` (`long int`) the current number of linear iterations.

Return value The return value `flag` (of type `int`) is one of

`IDASPILS_SUCCESS` The optional output value has been successfully set.
`IDASPILS_MEM_NULL` The `ida_mem` pointer is `NULL`.
`IDASPILS_LMEM_NULL` The IDASPILS linear solver has not been initialized.

IDASpilsGetNumConvFails

Call `flag = IDASpilsGetNumConvFails(ida_mem, &nlcfails);`

Description The function `IDASpilsGetNumConvFails` returns the cumulative number of linear convergence failures.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`nlcfails` (`long int`) the current number of linear convergence failures.

Return value The return value `flag` (of type `int`) is one of

`IDASPILS_SUCCESS` The optional output value has been successfully set.
`IDASPILS_MEM_NULL` The `ida_mem` pointer is `NULL`.
`IDASPILS_LMEM_NULL` The IDASPILS linear solver has not been initialized.

IDASpilsGetNumPrecEvals

Call `flag = IDASpilsGetNumPrecEvals(ida_mem, &npevals);`

Description The function `IDASpilsGetNumPrecEvals` returns the cumulative number of preconditioner evaluations, i.e., the number of calls made to `psetup`.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`npevals` (`long int`) the cumulative number of calls to `psetup`.

Return value The return value `flag` (of type `int`) is one of

`IDASPILS_SUCCESS` The optional output value has been successfully set.
`IDASPILS_MEM_NULL` The `ida_mem` pointer is `NULL`.
`IDASPILS_LMEM_NULL` The IDASPILS linear solver has not been initialized.

IDASpilsGetNumPrecSolves

Call `flag = IDASpilsGetNumPrecSolves(ida_mem, &npsolves);`

Description The function `IDASpilsGetNumPrecSolves` returns the cumulative number of calls made to the preconditioner solve function, `psolve`.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.
`npsolves` (long int) the cumulative number of calls to `psolve`.

Return value The return value `flag` (of type `int`) is one of

- `IDASPILS_SUCCESS` The optional output value has been successfully set.
- `IDASPILS_MEM_NULL` The `ida_mem` pointer is `NULL`.
- `IDASPILS_LMEM_NULL` The IDASPILS linear solver has not been initialized.

IDASpilsGetNumJtimesEvals

Call `flag = IDASpilsGetNumJtimesEvals(ida_mem, &njvevals);`

Description The function `IDASpilsGetNumJtimesEvals` returns the cumulative number of calls made to the Jacobian-vector function, `jtimes`.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.
`njvevals` (long int) the cumulative number of calls to `jtimes`.

Return value The return value `flag` (of type `int`) is one of

- `IDASPILS_SUCCESS` The optional output value has been successfully set.
- `IDASPILS_MEM_NULL` The `ida_mem` pointer is `NULL`.
- `IDASPILS_LMEM_NULL` The IDASPILS linear solver has not been initialized.

IDASpilsGetNumResEvals

Call `flag = IDASpilsGetNumResEvals(ida_mem, &nrevalsLS);`

Description The function `IDASpilsGetNumResEvals` returns the cumulative number of calls to the user residual function for finite difference Jacobian-vector product approximation.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.
`nrevalsLS` (long int) the cumulative number of calls to the user residual function.

Return value The return value `flag` (of type `int`) is one of

- `IDASPILS_SUCCESS` The optional output value has been successfully set.
- `IDASPILS_MEM_NULL` The `ida_mem` pointer is `NULL`.
- `IDASPILS_LMEM_NULL` The IDASPILS linear solver has not been initialized.

Notes The value `nrevalsLS` is incremented only if the default `IDASpilsDQJtimes` difference quotient function is used.

IDASpilsGetLastFlag

Call `flag = IDASpilsGetLastFlag(ida_mem, &lsflag);`

Description The function `IDASpilsGetLastFlag` returns the last return value from an IDASPILS routine.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.
`lsflag` (long int) the value of the last return flag from an IDASPILS function.

Return value The return value `flag` (of type `int`) is one of

- `IDASPILS_SUCCESS` The optional output value has been successfully set.
- `IDASPILS_MEM_NULL` The `ida_mem` pointer is `NULL`.

	IDASPILS_LMEM_NULL The IDASPILS linear solver has not been initialized.
Notes	<p>If the IDASPILS setup function failed (IDASolve returned IDA_LSETUP_FAIL), <code>lsflag</code> will be <code>SPGMR_PSET_FAIL_UNREC</code>, <code>SPBCG_PSET_FAIL_UNREC</code>, or <code>SPTFQMR_PSET_FAIL_UNREC</code>.</p> <p>If the IDASPGMR solve function failed (IDASolve returned IDA_LSOLVE_FAIL), <code>lsflag</code> contains the error return flag from <code>SpgmrSolve</code> and will be one of: <code>SPGMR_MEM_NULL</code>, indicating that the SPGMR memory is NULL; <code>SPGMR_ATIMES_FAIL_UNREC</code>, indicating an unrecoverable failure in the $J * v$ function; <code>SPGMR_PSOLVE_FAIL_UNREC</code>, indicating that the preconditioner solve function <code>psolve</code> failed unrecoverably; <code>SPGMR_GS_FAIL</code>, indicating a failure in the Gram-Schmidt procedure; or <code>SPGMR_QRSOL_FAIL</code>, indicating that the matrix R was found to be singular during the QR solve phase.</p> <p>If the IDASPCG solve function failed (IDASolve returned IDA_LSOLVE_FAIL), <code>lsflag</code> contains the error return flag from <code>SpcgSolve</code> and will be one of: <code>SPBCG_MEM_NULL</code>, indicating that the SPBCG memory is NULL; <code>SPBCG_ATIMES_FAIL_UNREC</code>, indicating an unrecoverable failure in the $J * v$ function; or <code>SPBCG_PSOLVE_FAIL_UNREC</code>, indicating that the preconditioner solve function <code>psolve</code> failed unrecoverably.</p> <p>If the IDASPTFQMR solve function failed (IDASolve returned IDA_LSOLVE_FAIL), <code>lsflag</code> contains the error flag from <code>SptfqmrSolve</code> and will be one of: <code>SPTFQMR_MEM_NULL</code>, indicating that the SPTFQMR memory is NULL; <code>SPTFQMR_ATIMES_FAIL_UNREC</code>, indicating an unrecoverable failure in the $J * v$ function; or <code>SPTFQMR_PSOLVE_FAIL_UNREC</code>, indicating that the preconditioner solve function <code>psolve</code> failed unrecoverably.</p>

IDASpilsGetReturnFlagName

Call	<code>name = IDASpilsGetReturnFlagName(lsflag);</code>
Description	The function <code>IDASpilsGetReturnFlagName</code> returns the name of the IDASPILS constant corresponding to <code>lsflag</code> .
Arguments	The only argument, of type <code>long int</code> , is a return flag from an IDASPILS function.
Return value	The return value is a string containing the name of the corresponding constant.

4.5.10 IDAS reinitialization function

The function `IDAREInit` reinitializes the main IDAS solver for the solution of a 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, assuming 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.

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 `NVECTOR` 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 `IDA***` calls, as described in §4.5.3. If there are changes to any optional inputs, make the appropriate `IDASet***` calls, as described in §4.5.7.

IDAREInit

Call	<code>flag = IDAREInit(ida_mem, t0, y0, yp0);</code>
Description	The function <code>IDAREInit</code> provides required problem specifications and reinitializes IDAS.
Arguments	<p><code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block.</p> <p><code>t0</code> (<code>realtype</code>) is the initial value of t.</p> <p><code>y0</code> (<code>N_Vector</code>) is the initial value of y.</p> <p><code>yp0</code> (<code>N_Vector</code>) is the initial value of \dot{y}.</p>

Return value	The return value <code>flag</code> (of type <code>int</code>) will be one of the following:
	<code>IDA_SUCCESS</code> The call to <code>IDAReInit</code> was successful.
	<code>IDA_MEM_NULL</code> The IDAS memory block was not initialized through a previous call to <code>IDACreate</code> .
	<code>IDA_NO_MALLOC</code> Memory space for the IDAS memory block was not allocated through a previous call to <code>IDAINit</code> .
	<code>IDA_ILL_INPUT</code> An input argument to <code>IDAReInit</code> has an illegal value.
Notes	If an error occurred, <code>IDAReInit</code> also sends an error message to the error handler function.

4.6 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) a function that provides Jacobian-related information for the linear solver (if Newton iteration is chosen), and (optionally) one or two functions that define the preconditioner for use in any of the Krylov iteration algorithms.

4.6.1 Residual function

The user must provide a function of type `IDAResFn` defined as follows:

	<div style="border: 1px solid black; padding: 2px; display: inline-block;">IDAResFn</div>
Definition	<pre>typedef int (*IDAResFn)(realtype tt, N_Vector yy, N_Vector yp, N_Vector rr, void *user_data);</pre>
Purpose	This function computes the problem residual for given values of the independent variable t , state vector y , and derivative \dot{y} .
Arguments	<p><code>tt</code> is the current value of the independent variable.</p> <p><code>yy</code> is the current value of the dependent variable vector, $y(t)$.</p> <p><code>yp</code> is the current value of $\dot{y}(t)$.</p> <p><code>rr</code> is the output residual vector $F(t, y, \dot{y})$.</p> <p><code>user_data</code> is a pointer to user data, the same as the <code>user_data</code> parameter passed to <code>IDASetUserData</code>.</p>
Return value	An <code>IDAResFn</code> function type should return a value of 0 if successful, a positive value if a recoverable error occurred (e.g. <code>yy</code> 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	<p>A recoverable failure error return from the <code>IDAResFn</code> is typically used to flag a value of the dependent variable y that is “illegal” in some way (e.g., negative where only a non-negative value is physically meaningful). If such a return is made, IDAS will attempt to recover (possibly repeating the Newton iteration, or reducing the step size) in order to avoid this recoverable error return.</p> <p>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 right-hand side function is called the first time during the following integration step, but a successful step cannot be undone.) However, if the user program also includes quadrature integration, the state variables can be checked for legality in the call to <code>IDAQuadRhsFn</code>, which is called at the converged solution of the nonlinear system, and therefore IDAS can be flagged to attempt to recover from such a situation. Also, if sensitivity analysis is performed with</p>

the staggered method, the DAE residual function is called at the converged solution of the nonlinear system, and a recoverable error at that point can be flagged, and IDAS will then try to correct it.

Allocation of memory for `yp` is handled within IDAS.

4.6.2 Error message handler function

As an alternative to the default behavior of directing error and warning messages to the file pointed to by `errfp` (see `IDASetErrFile`), the user may provide a function of type `IDAErrorHandlerFn` to process any such messages. The function type `IDAErrorHandlerFn` is defined as follows:

IDAErrorHandlerFn

Definition `typedef void (*IDAErrorHandlerFn)(int error_code, const char *module, const char *function, char *msg, void *eh_data);`

Purpose This function processes error and warning messages from IDAS and its sub-modules.

Arguments `error_code` is the error code.
 `module` is the name of the IDAS module reporting the error.
 `function` is the name of the function in which the error occurred.
 `msg` is the error message.
 `eh_data` is a pointer to user data, the same as the `eh_data` parameter passed to `IDASetErrorHandlerFn`.

Return value A `IDAErrorHandlerFn` function has no return value.

Notes `error_code` is negative for errors and positive (`IDA_WARNING`) for warnings. If a function that returns a pointer to memory encounters an error, it sets `error_code` to 0.

4.6.3 Error weight function

As an alternative to providing the relative and absolute tolerances, the user may provide a function of type `IDAWEtFn` to compute a vector `ewt` containing the multiplicative weights W_i used in the WRMS norm $\|v\|_{\text{WRMS}} = \sqrt{(1/N) \sum_1^N (W_i \cdot v_i)^2}$. These weights will be used in place of those defined by Eq. (2.7). The function type `IDAWEtFn` is defined as follows:

IDAWEtFn

Definition `typedef int (*IDAWEtFn)(N_Vector y, N_Vector ewt, void *user_data);`

Purpose This function computes the WRMS error weights for the vector `y`.

Arguments `y` is the value of the dependent variable vector at which the weight vector is to be computed.
 `ewt` is the output vector containing the error weights.
 `user_data` is a pointer to user data, the same as the `user_data` parameter passed to `IDASetUserData`.

Return value An `IDAWEtFn` function type must return 0 if it successfully set the error weights and -1 otherwise.

Notes Allocation of memory for `ewt` is handled within IDAS.

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.



4.6.4 Rootfinding function

If a rootfinding problem is to be solved during the integration of the DAE system, the user must supply a C function of type `IDARootFn`, defined as follows:

`IDARootFn`

Definition	<code>typedef int (*IDARootFn)(realtype t, N_Vector y, N_Vector yp, realtype *gout, void *user_data);</code>
Purpose	This function computes a vector-valued function $g(t, y, \dot{y})$ such that the roots of the <code>nrtnfn</code> components $g_i(t, y, \dot{y})$ are to be found during the integration.
Arguments	<p><code>t</code> is the current value of the independent variable.</p> <p><code>y</code> is the current value of the dependent variable vector, $y(t)$.</p> <p><code>yp</code> is the current value of $\dot{y}(t)$, the t-derivative of y.</p> <p><code>gout</code> is the output array, of length <code>nrtnfn</code>, with components $g_i(t, y, \dot{y})$.</p> <p><code>user_data</code> is a pointer to user data, the same as the <code>user_data</code> parameter passed to <code>IDASetUserData</code>.</p>
Return value	An <code>IDARootFn</code> should return 0 if successful or a non-zero value if an error occurred (in which case the integration is halted and <code>IDASolve</code> returns <code>IDA_RTFUNC_FAIL</code>).
Notes	Allocation of memory for <code>gout</code> is handled within IDAS.

4.6.5 Jacobian information (direct method with dense Jacobian)

If the direct linear solver with dense treatment of the Jacobian is used (i.e. either `IDADense` or `IDALapackDense` is called in Step 8 of §4.4), the user may provide a function of type `IDADlsDenseJacFn` defined by

`IDADlsDenseJacFn`

Definition	<code>typedef int (*IDADlsDenseJacFn)(long int Neq, realtype tt, realtype cj, N_Vector yy, N_Vector yp, N_Vector rr, DlsMat Jac, void *user_data, N_Vector tmp1, N_Vector tmp2, N_Vector tmp3);</code>
Purpose	This function computes the dense Jacobian J of the DAE system (or an approximation to it), defined by Eq. (2.6).
Arguments	<p><code>Neq</code> is the problem size (number of equations).</p> <p><code>tt</code> is the current value of the independent variable t.</p> <p><code>cj</code> is the scalar in the system Jacobian, proportional to the inverse of the step size (α in Eq. (2.6)).</p> <p><code>yy</code> is the current value of the dependent variable vector, $y(t)$.</p> <p><code>yp</code> is the current value of $\dot{y}(t)$.</p> <p><code>rr</code> is the current value of the residual vector $F(t, y, \dot{y})$.</p> <p><code>Jac</code> is the output (approximate) Jacobian matrix, $J = \partial F / \partial y + cj \partial F / \partial \dot{y}$.</p> <p><code>user_data</code> is a pointer to user data, the same as the <code>user_data</code> parameter passed to <code>IDASetUserData</code>.</p> <p><code>tmp1</code> <code>tmp2</code> <code>tmp3</code> are pointers to memory allocated for variables of type <code>N_Vector</code> which can be used by <code>IDADlsDenseJacFn</code> as temporary storage or work space.</p>

Return value An `IDADlsDenseJacFn` function type 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 A user-supplied dense Jacobian function must load the $\text{Neq} \times \text{Neq}$ dense matrix `Jac` with an approximation to the Jacobian matrix $J(t, y, \dot{y})$ at the point `(tt, yy, yp)`. Only nonzero elements need to be loaded into `Jac` because `Jac` is set to the zero matrix before the call to the Jacobian function. The type of `Jac` is `DlsMat` (described below and in §9.1).

The accessor macros `DENSE_ELEM` and `DENSE_COL` allow the user to read and write dense matrix elements without making explicit references to the underlying representation of the `DlsMat` type. `DENSE_ELEM(Jac, i, j)` references the (i, j) -th element of the dense matrix `Jac` ($i, j = 0 \dots \text{Neq}-1$). This macro is for use in small problems in which efficiency of access is not a major concern. Thus, in terms of indices m and n running from 1 to `Neq`, the Jacobian element $J_{m,n}$ can be loaded with the statement `DENSE_ELEM(Jac, m-1, n-1) = J_{m,n}`. Alternatively, `DENSE_COL(Jac, j)` returns a pointer to the storage for the j th column of `Jac` ($j = 0 \dots \text{Neq}-1$), and the elements of the j -th column are then accessed via ordinary array indexing. Thus $J_{m,n}$ can be loaded with the statements `col_n = DENSE_COL(Jac, n-1); col_n[m-1] = J_{m,n}`. For large problems, it is more efficient to use `DENSE_COL` than to use `DENSE_ELEM`. Note that both of these macros number rows and columns starting from 0, not 1.

The `DlsMat` type and the accessor macros `DENSE_ELEM` and `DENSE_COL` are documented in §9.1.

If the user's `IDADlsDenseJacFn` 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 the `ida_mem` structure to their `user_data` and use the `IDAGet*` functions described in §4.5.9.1. The unit roundoff can be accessed as `UNIT_ROUNDOFF` defined in `sundials.types.h`.

For the sake of uniformity, the argument `Neq` is of type `long int`, even in the case that the Lapack dense solver is to be used.

4.6.6 Jacobian information (direct method with banded Jacobian)

If the direct linear solver with banded treatment of the Jacobian is used (i.e. either `IDABand` or `IDALapackBand` is called in Step 8 of §4.4), the user may provide a function of type `IDADlsBandJacFn` defined as follows:

`IDADlsBandJacFn`

Definition

```
typedef int (*IDADlsBandJacFn)(long int Neq, long int mupper,
                                long int mlower, realtype tt, realtype cj,
                                N_Vector yy, N_Vector yp, N_Vector rr,
                                DlsMat Jac, void *user_data,
                                N_Vector tmp1, N_Vector tmp2, N_Vector tmp3);
```

Purpose This function computes the banded Jacobian J of the DAE system (or a banded approximation to it), defined by Eq. (2.6).

Arguments

<code>Neq</code>	is the problem size.
<code>mupper</code>	
<code>mlower</code>	are the upper and lower half bandwidth of the Jacobian.
<code>tt</code>	is the current value of the independent variable.
<code>yy</code>	is the current value of the dependent variable vector, $y(t)$.

	<p>yp is the current value of $\dot{y}(t)$.</p> <p>rr is the current value of the residual vector $F(t, y, \dot{y})$.</p> <p>cj is the scalar in the system Jacobian, proportional to the inverse of the step size (α in Eq. (2.6)).</p> <p>Jac is the output (approximate) Jacobian matrix, $J = \partial F / \partial y + cj \partial F / \partial \dot{y}$.</p> <p>user_data is a pointer to user data, the same as the user_data parameter passed to <code>IDASetUserData</code>.</p> <p>tmp1 tmp2 tmp3 are pointers to memory allocated for variables of type <code>N_Vector</code> which can be used by <code>IDADlsBandJacFn</code> as temporary storage or work space.</p>
Return value	<p>A <code>IDADlsBandJacFn</code> function type should return 0 if successful, a positive value if a recoverable error occurred, or a negative value if a nonrecoverable error occurred.</p> <p>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).</p>
Notes	<p>A user-supplied band Jacobian function must load the band matrix Jac of type <code>DlsMat</code> with the elements of the Jacobian $J(t, y, \dot{y})$ at the point (tt, yy, yp). Only nonzero elements need to be loaded into Jac because Jac is preset to zero before the call to the Jacobian function.</p> <p>The accessor macros <code>BAND_ELEM</code>, <code>BAND_COL</code>, and <code>BAND_COL_ELEM</code> allow the user to read and write band matrix elements without making specific references to the underlying representation of the <code>DlsMat</code> type. <code>BAND_ELEM(Jac, i, j)</code> references the (i, j)th element of the band matrix Jac, counting from 0. This macro is for use in small problems in which efficiency of access is not a major concern. Thus, in terms of indices m and n running from 1 to Neq with (m, n) within the band defined by mupper and mlower, the Jacobian element $J_{m,n}$ can be loaded with the statement <code>BAND_ELEM(Jac, m-1, n-1) = J_{m,n}</code>. The elements within the band are those with $-\text{mupper} \leq m-n \leq \text{mlower}$. Alternatively, <code>BAND_COL(Jac, j)</code> returns a pointer to the diagonal element of the jth column of Jac, and if we assign this address to <code>realtype *col_j</code>, then the ith element of the jth column is given by <code>BAND_COL_ELEM(col_j, i, j)</code>, counting from 0. Thus for (m, n) within the band, $J_{m,n}$ can be loaded by setting <code>col_n = BAND_COL(Jac, n-1); BAND_COL_ELEM(col_n, m-1, n-1) = J_{m,n}</code>. The elements of the jth column can also be accessed via ordinary array indexing, but this approach requires knowledge of the underlying storage for a band matrix of type <code>DlsMat</code>. The array <code>col_n</code> can be indexed from $-\text{mupper}$ to mlower. For large problems, it is more efficient to use the combination of <code>BAND_COL</code> and <code>BAND_COL_ELEM</code> than to use the <code>BAND_ELEM</code>. As in the dense case, these macros all number rows and columns starting from 0, not 1.</p> <p>The <code>DlsMat</code> type and the accessor macros <code>BAND_ELEM</code>, <code>BAND_COL</code>, and <code>BAND_COL_ELEM</code> are documented in §9.1.</p> <p>If the user's <code>IDADlsBandJacFn</code> 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 the <code>ida_mem</code> structure to their user_data and use the <code>IDAGet*</code> functions described in §4.5.9.1. The unit roundoff can be accessed as <code>UNIT_ROUNDOFF</code> defined in <code>sundials.types.h</code>.</p> <p>For the sake of uniformity, the arguments Neq, mlower, and mupper are of type <code>long int</code>, even in the case that the Lapack band solver is to be used.</p>

4.6.7 Jacobian information (direct method with sparse Jacobian)

If the direct linear solver with sparse treatment of the Jacobian is used (i.e. either `IDAKLU` or `IDASuperLUMT` is called in Step 8 of §4.4), the user must provide a function of type `IDASlsSparseJacFn`

Definition	<pre>typedef int (*IDASpilsJacTimesVecFn)(realtype tt, N_Vector yy, N_Vector yp, N_Vector rr, N_Vector v, N_Vector Jv, realtype cj, void *user_data, N_Vector tmp1, N_Vector tmp2);</pre>	
Purpose	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 Eq. (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 Eq. (2.6)).
	user_data	is a pointer to user data, the same as the user_data parameter passed to <code>IDASetUserData</code> .
	tmp1	
	tmp2	are pointers to memory allocated for variables of type <code>N_Vector</code> which can be used by <code>IDASpilsJacTimesVecFn</code> as temporary storage or work space.
Return value	The value to be returned by the Jacobian-times-vector function should be 0 if successful. A nonzero value indicates that a nonrecoverable error occurred.	
	If the user's <code>IDASpilsJacTimesVecFn</code> 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 the <code>ida_mem</code> structure to their user_data and use the <code>IDAGet*</code> functions described in §4.5.9.1. The unit roundoff can be accessed as <code>UNIT_ROUNDOFF</code> defined in <code>sundials.types.h</code> .	

4.6.9 Preconditioning (linear system solution)

If preconditioning is used, then the user must provide a C function to solve the linear system $Pz = r$ where P is a left preconditioner matrix which approximates (at least crudely) the Jacobian matrix $J = \partial F / \partial y + cj \partial F / \partial \dot{y}$. This function must be of type `IDASpilsPrecSolveFn`, defined as follows:

IDASpilsPrecSolveFn

Definition	<pre>typedef int (*IDASpilsPrecSolveFn)(realtype tt, N_Vector yy, N_Vector yp, N_Vector rr, N_Vector rvec, N_Vector zvec, realtype cj, realtype delta, void *user_data, N_Vector tmp);</pre>	
Purpose	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 Eq. (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., $\sqrt{\sum_i (Res_i \cdot ewt_i)^2} < \text{delta}$. To obtain the `N_Vector ewt`, call `IDAGetErrWeights` (see §4.5.9.1).

user_data is a pointer to user data, the same as the `user_data` parameter passed to the function `IDASetUserData`.

tmp is a pointer to memory allocated for a variable of type `N_Vector` which can be used for work space.

Return value The value to be returned by the preconditioner solve function is a flag indicating whether it was successful. This value should be 0 if successful, positive for a recoverable error (in which case the step will be retried), negative for an unrecoverable error (in which case the integration is halted).

4.6.10 Preconditioning (Jacobian data)

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 C function of type `IDASpilsPrecSetupFn`, defined as follows:

<code>IDASpilsPrecSetupFn</code>

Definition `typedef int (*IDASpilsPrecSetupFn)(realtype tt, N_Vector yy,
N_Vector yp, N_Vector rr,
realtype cj, void *user_data,
N_Vector tmp1, N_Vector tmp2,
N_Vector tmp3);`

Purpose This function evaluates and/or preprocesses Jacobian-related data needed by the preconditioner.

Arguments The arguments of an `IDASpilsPrecSetupFn` are as follows:

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 Eq. (2.6)).

user_data is a pointer to user data, the same as the `user_data` parameter passed to the function `IDASetUserData`.

tmp1
tmp2
tmp3 are pointers to memory allocated for variables of type `N_Vector` which can be used by `IDASpilsPrecSetupFn` as temporary storage or work space.

Return value The value to be returned by the preconditioner setup function is a flag indicating whether it was successful. This value should be 0 if successful, positive for a recoverable error (in which case the step will be retried), negative for an unrecoverable error (in which case the integration is halted).

Notes The operations performed by this function might include forming a crude approximate Jacobian, and performing an LU factorization on the resulting approximation.

Each call to the preconditioner setup function is preceded by a call to the `IDAResFn` user function with the same (`tt`, `yy`, `yp`) arguments. Thus the preconditioner setup function can use any auxiliary data that is computed and saved during the evaluation of the DAE residual.

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 Newton iteration.

If the user's `IDASpilsPrecSetupFn` 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 the `ida_mem` structure to their `user_data` and use the `IDAGet*` functions described in §4.5.9.1. The unit roundoff can be accessed as `UNIT_ROUNDOFF` defined in `sundials_types.h`.

4.7 Integration of pure quadrature equations

IDAS allows the DAE system to include *pure quadratures*. In this case, it is more efficient to treat the quadratures separately by excluding them from the nonlinear solution stage. To do this, begin by excluding the quadrature variables from the vectors `yy` and `yp` and the quadrature equations from within `res`. Thus a separate vector `yQ` of quadrature variables is to satisfy $(d/dt)yQ = f_Q(t, y, \dot{y})$. The following is an overview of the sequence of calls in a user's main program in this situation. Steps that are unchanged from the skeleton program presented in §4.4 are grayed out.

1. **[P] Initialize MPI**
2. **Set problem dimensions**
 - [S] Set `N` to the problem size N (excluding quadrature variables), and `Nq` to the number of quadrature variables.
 - [P] Set `Nlocal` to the local vector length (excluding quadrature variables), and `Nqlocal` to the local number of quadrature variables.
3. **Set vectors of initial values**
4. **Create IDAS object**
5. **Allocate internal memory**
6. **Set optional inputs**
7. **Attach linear solver module**
8. **Set linear solver optional inputs**
9. **Set vector of initial values for quadrature variables**
 - Typically, the quadrature variables should be initialized to 0.
10. **Initialize quadrature integration**
 - Call `IDAQuadInit` to specify the quadrature equation right-hand side function and to allocate internal memory related to quadrature integration. See §4.7.1 for details.
11. **Set optional inputs for quadrature integration**
 - Call `IDASetQuadErrCon` to indicate whether or not quadrature variables should be used in the step size control mechanism. If so, one of the `IDAQuad*tolerances` functions must be called to specify the integration tolerances for quadrature variables. See §4.7.4 for details.
12. **Advance solution in time**
13. **Extract quadrature variables**
 - Call `IDAGetQuad` or `IDAGetQuadDky` to obtain the values of the quadrature variables or their derivatives at the current time. See §4.7.3 for details.

14. Get optional outputs

15. Get quadrature optional outputs

Call `IDAGetQuad*` functions to obtain optional output related to the integration of quadratures. See §4.7.5 for details.

16. Deallocate memory for solution vectors and for the vector of quadrature variables

17. Free solver memory

18. [P] Finalize MPI

`IDAQuadInit` can be called and quadrature-related optional inputs (step 11 above) can be set, anywhere between steps 4 and 12.

4.7.1 Quadrature initialization and deallocation functions

The function `IDAQuadInit` activates integration of quadrature equations and allocates internal memory related to these calculations. The form of the call to this function is as follows:

IDAQuadInit	
Call	<code>flag = IDAQuadInit(ida_mem, rhsQ, yQ0);</code>
Description	The function <code>IDAQuadInit</code> provides required problem specifications, allocates internal memory, and initializes quadrature integration.
Arguments	<p><code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block returned by <code>IDACreate</code>.</p> <p><code>rhsQ</code> (<code>IDAQuadRhsFn</code>) is the C function which computes f_Q, the right-hand side of the quadrature equations. This function has the form <code>fQ(t, yy, yp, rhsQ, user_data)</code> (for full details see §4.7.6).</p> <p><code>yQ0</code> (<code>N.Vector</code>) is the initial value of y_Q.</p>
Return value	<p>The return value <code>flag</code> (of type <code>int</code>) will be one of the following:</p> <p><code>IDA_SUCCESS</code> The call to <code>IDAQuadInit</code> was successful.</p> <p><code>IDA_MEM_NULL</code> The IDAS memory was not initialized by a prior call to <code>IDACreate</code>.</p> <p><code>IDA_MEM_FAIL</code> A memory allocation request failed.</p>
Notes	If an error occurred, <code>IDAQuadInit</code> also sends an error message to the error handler function.

In terms of the number of quadrature variables N_q and maximum method order `maxord`, the size of the real workspace is increased as follows:

- Base value: $\text{lenrw} = \text{lenrw} + (\text{maxord}+5)N_q$
- If `IDAQuadSVtolerances` is called: $\text{lenrw} = \text{lenrw} + N_q$

and the size of the integer workspace is increased as follows:

- Base value: $\text{leniw} = \text{leniw} + (\text{maxord}+5)N_q$
- If `IDAQuadSVtolerances` is called: $\text{leniw} = \text{leniw} + N_q$

The function `IDAQuadReInit`, useful during the solution of a sequence of problems of same size, reinitializes the quadrature-related internal memory and must follow a call to `IDAQuadInit` (and maybe a call to `IDAReInit`). The number N_q of quadratures is assumed to be unchanged from the prior call to `IDAQuadInit`. The call to the `IDAQuadReInit` function has the following form:

IDAQuadReInit

Call	<code>flag = IDAQuadReInit(ida_mem, yQ0);</code>
Description	The function <code>IDAQuadReInit</code> provides required problem specifications and reinitializes the quadrature integration.
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block. <code>yQ0</code> (<code>N.Vector</code>) is the initial value of y_Q .
Return value	The return value <code>flag</code> (of type <code>int</code>) will be one of the following: <code>IDA_SUCCESS</code> The call to <code>IDAReInit</code> was successful. <code>IDA_MEM_NULL</code> The IDAS memory was not initialized by a prior call to <code>IDACreate</code> . <code>IDA_NO_QUAD</code> Memory space for the quadrature integration was not allocated by a prior call to <code>IDAQuadInit</code> .
Notes	If an error occurred, <code>IDAQuadReInit</code> also sends an error message to the error handler function.

IDAQuadFree

Call	<code>IDAQuadFree(ida_mem);</code>
Description	The function <code>IDAQuadFree</code> frees the memory allocated for quadrature integration.
Arguments	The argument is the pointer to the IDAS memory block (of type <code>void *</code>).
Return value	The function <code>IDAQuadFree</code> has no return value.
Notes	In general, <code>IDAQuadFree</code> need not be called by the user as it is invoked automatically by <code>IDAFree</code> .

4.7.2 IDAS solver function

Even if quadrature integration was enabled, the call to the main solver function `IDASolve` is exactly the same as in §4.5.6. However, in this case the return value `flag` can also be one of the following:

<code>IDA_QRHS_FAIL</code>	The quadrature right-hand side function failed in an unrecoverable manner.
<code>IDA_FIRST_QRHS_ERR</code>	The quadrature right-hand side function failed at the first call.
<code>IDA_REP_QRHS_ERR</code>	Convergence test failures occurred too many times due to repeated recoverable errors in the quadrature right-hand side function. This value will also be returned if the quadrature right-hand side function had repeated recoverable errors during the estimation of an initial step size (assuming the quadrature variables are included in the error tests).

4.7.3 Quadrature extraction functions

If quadrature integration has been initialized by a call to `IDAQuadInit`, or reinitialized by a call to `IDAQuadReInit`, then IDAS computes both a solution and quadratures at time t . However, `IDASolve` will still return only the solution y in y . Solution quadratures can be obtained using the following function:

IDAGetQuad

Call	<code>flag = IDAGetQuad(ida_mem, &tret, yQ);</code>
Description	The function <code>IDAGetQuad</code> returns the quadrature solution vector after a successful return from <code>IDASolve</code> .
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the memory previously allocated by <code>IDAInit</code> . <code>tret</code> (<code>realtype</code>) the time reached by the solver (output).

`yQ` (N_Vector) the computed quadrature vector.

Return value The return value `flag` of `IDAGetQuad` is one of:

`IDA_SUCCESS` `IDAGetQuad` was successful.

`IDA_MEM_NULL` `ida_mem` was NULL.

`IDA_NO_QUAD` Quadrature integration was not initialized.

`IDA_BAD_DKY` `yQ` is NULL.

The function `IDAGetQuadDky` computes the k -th derivatives of the interpolating polynomials for the quadrature variables at time t . This function is called by `IDAGetQuad` with $k = 0$ and with the current time at which `IDASolve` has returned, but may also be called directly by the user.

`IDAGetQuadDky`

Call `flag = IDAGetQuadDky(ida_mem, t, k, dkyQ);`

Description The function `IDAGetQuadDky` returns derivatives of the quadrature solution vector after a successful return from `IDASolve`.

Arguments `ida_mem` (void *) pointer to the memory previously allocated by `IDAInit`.
`t` (realtype) the time at which quadrature information is requested. The time `t` must fall within the interval defined by the last successful step taken by IDAS.
`k` (int) order of the requested derivative. This must be $\leq klast$.
`dkyQ` (N_Vector) the vector containing the derivative. This vector must be allocated by the user.

Return value The return value `flag` of `IDAGetQuadDky` is one of:

`IDA_SUCCESS` `IDAGetQuadDky` succeeded.

`IDA_MEM_NULL` The pointer to `ida_mem` was NULL.

`IDA_NO_QUAD` Quadrature integration was not initialized.

`IDA_BAD_DKY` The vector `dkyQ` is NULL.

`IDA_BAD_K` `k` is not in the range $0, 1, \dots, klast$.

`IDA_BAD_T` The time `t` is not in the allowed range.

4.7.4 Optional inputs for quadrature integration

IDAS provides the following optional input functions to control the integration of quadrature equations.

`IDASetQuadErrCon`

Call `flag = IDASetQuadErrCon(ida_mem, errconQ);`

Description The function `IDASetQuadErrCon` specifies whether or not the quadrature variables are to be used in the step size control mechanism within IDAS. If they are, the user must call either `IDAQuadSStolerances` or `IDAQuadSVtolerances` to specify the integration tolerances for the quadrature variables.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.
`errconQ` (boolean type) specifies whether quadrature variables are included (TRUE) or not (FALSE) in the error control mechanism.

Return value The return value `flag` (of type int) is one of:

`IDA_SUCCESS` The optional value has been successfully set.

`IDA_MEM_NULL` The `ida_mem` pointer is NULL.

`IDA_NO_QUAD` Quadrature integration has not been initialized.



Notes By default, `errconQ` is set to `FALSE`.
It is illegal to call `IDASetQuadErrCon` before a call to `IDAQuadInit`.

If the quadrature variables are part of the step size control mechanism, one of the following functions must be called to specify the integration tolerances for quadrature variables.

IDAQuadSStolerances

Call `flag = IDAQuadSStolerances(ida_mem, reltolQ, abstolQ);`
 Description The function `IDAQuadSStolerances` specifies scalar relative and absolute tolerances.
 Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`reltolQ` (`realtype`) is the scalar relative error tolerance.
`abstolQ` (`realtype`) is the scalar absolute error tolerance.
 Return value The return value `flag` (of type `int`) is one of:
`IDA_SUCCESS` The optional value has been successfully set.
`IDA_NO_QUAD` Quadrature integration was not initialized.
`IDA_MEM_NULL` The `ida_mem` pointer is `NULL`.
`IDA_ILL_INPUT` One of the input tolerances was negative.

IDAQuadSVtolerances

Call `flag = IDAQuadSVtolerances(ida_mem, reltolQ, abstolQ);`
 Description The function `IDAQuadSVtolerances` specifies scalar relative and vector absolute tolerances.
 Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`reltolQ` (`realtype`) is the scalar relative error tolerance.
`abstolQ` (`N_Vector`) is the vector absolute error tolerance.
 Return value The return value `flag` (of type `int`) is one of:
`IDA_SUCCESS` The optional value has been successfully set.
`IDA_NO_QUAD` Quadrature integration was not initialized.
`IDA_MEM_NULL` The `ida_mem` pointer is `NULL`.
`IDA_ILL_INPUT` One of the input tolerances was negative.

4.7.5 Optional outputs for quadrature integration

IDAS provides the following functions that can be used to obtain solver performance information related to quadrature integration.

IDAGetQuadNumRhsEvals

Call `flag = IDAGetQuadNumRhsEvals(ida_mem, &nrhsQevals);`
 Description The function `IDAGetQuadNumRhsEvals` returns the number of calls made to the user's quadrature right-hand side function.
 Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`nrhsQevals` (`long int`) number of calls made to the user's `rhsQ` function.
 Return value The return value `flag` (of type `int`) is one of:
`IDA_SUCCESS` The optional output value has been successfully set.
`IDA_MEM_NULL` The `ida_mem` pointer is `NULL`.
`IDA_NO_QUAD` Quadrature integration has not been initialized.

IDAGetQuadNumErrTestFails

Call `flag = IDAGetQuadNumErrTestFails(ida_mem, &nQetfails);`

Description The function `IDAGetQuadNumErrTestFails` returns the number of local error test failures due to quadrature variables.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`nQetfails` (`long int`) number of error test failures due to quadrature variables.

Return value The return value `flag` (of type `int`) is one of:

- `IDA_SUCCESS` The optional output value has been successfully set.
- `IDA_MEM_NULL` The `ida_mem` pointer is NULL.
- `IDA_NO_QUAD` Quadrature integration has not been initialized.

IDAGetQuadErrWeights

Call `flag = IDAGetQuadErrWeights(ida_mem, eQweight);`

Description The function `IDAGetQuadErrWeights` returns the quadrature error weights at the current time.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`eQweight` (`N_Vector`) quadrature error weights at the current time.

Return value The return value `flag` (of type `int`) is one of:

- `IDA_SUCCESS` The optional output value has been successfully set.
- `IDA_MEM_NULL` The `ida_mem` pointer is NULL.
- `IDA_NO_QUAD` Quadrature integration has not been initialized.



Notes The user must allocate memory for `eQweight`.

If quadratures were not included in the error control mechanism (through a call to `IDASetQuadErrCon` with `errconQ = TRUE`), `IDAGetQuadErrWeights` does not set the `eQweight` vector.

IDAGetQuadStats

Call `flag = IDAGetQuadStats(ida_mem, &nrhsQevals, &nQetfails);`

Description The function `IDAGetQuadStats` returns the IDAS integrator statistics as a group.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`nrhsQevals` (`long int`) number of calls to the user's `rhsQ` function.
`nQetfails` (`long int`) number of error test failures due to quadrature variables.

Return value The return value `flag` (of type `int`) is one of

- `IDA_SUCCESS` the optional output values have been successfully set.
- `IDA_MEM_NULL` the `ida_mem` pointer is NULL.
- `IDA_NO_QUAD` Quadrature integration has not been initialized.

4.7.6 User-supplied function for quadrature integration

For integration of quadrature equations, the user must provide a function that defines the right-hand side of the quadrature equations (in other words, the integrand function of the integral that must be evaluated). This function must be of type `IDAQuadRhsFn` defined as follows:

IDAQuadRhsFn

Definition	<code>typedef int (*IDAQuadRhsFn)(realtype t, N_Vector yy, N_Vector yp, N_Vector rhsQ, void *user_data);</code>
Purpose	This function computes the quadrature equation right-hand side for a given value of the independent variable t and state vectors y and \dot{y} .
Arguments	<p>t is the current value of the independent variable.</p> <p>yy is the current value of the dependent variable vector, $y(t)$.</p> <p>yp is the current value of the dependent variable derivative vector, $\dot{y}(t)$.</p> <p>$rhsQ$ is the output vector $f_Q(t, y, \dot{y})$.</p> <p>$user_data$ is the <code>user_data</code> pointer passed to <code>IDASetUserData</code>.</p>
Return value	A <code>IDAQuadRhsFn</code> should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and <code>IDA_QRHS_FAIL</code> is returned).
Notes	<p>Allocation of memory for <code>rhsQ</code> is automatically handled within IDAS.</p> <p>Both y and <code>rhsQ</code> are of type <code>N_Vector</code>, but they typically have different internal representations. It is the user's responsibility to access the vector data consistently (including the use of the correct accessor macros from each <code>NVECTOR</code> implementation). For the sake of computational efficiency, the vector functions in the two <code>NVECTOR</code> implementations provided with IDAS do not perform any consistency checks with respect to their <code>N_Vector</code> arguments (see §7.1 and §7.2).</p> <p>There is one situation in which recovery is not possible even if <code>IDAQuadRhsFn</code> function returns a recoverable error flag. This is when this occurs at the very first call to the <code>IDAQuadRhsFn</code> (in which case IDAS returns <code>IDA_FIRST_QRHS_ERR</code>).</p>

4.8 A parallel band-block-diagonal preconditioner module

A principal reason for using a parallel DAE solver such as IDAS lies in the solution of partial differential equations (PDEs). Moreover, the use of a Krylov iterative method for the solution of many such problems is motivated by the nature of the underlying linear system of equations (2.5) 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 [23] and is included in a software module within the IDAS package. This module works with the parallel vector module `NVECTOR_PARALLEL` and generates a preconditioner that is a block-diagonal matrix with each block being a band matrix. The blocks need not have the same number of super- and sub-diagonals and these numbers may vary from block to block. This Band-Block-Diagonal Preconditioner module is called `IDABBDPRE`.

One way to envision these preconditioners is to think of the domain of the computational PDE problem as being subdivided into M non-overlapping sub-domains. Each of these sub-domains is then assigned to one of the M processors to be used to solve the DAE system. The basic idea is to isolate the preconditioning so that it is local to each processor, and also to use a (possibly cheaper) approximate residual function. This requires the definition of a new function $G(t, y, \dot{y})$ which approximates the function $F(t, y, \dot{y})$ in the definition of the DAE system (2.1). However, the user may set $G = F$. Corresponding to the domain decomposition, there is a decomposition of the solution vectors y and \dot{y} into M disjoint blocks y_m and \dot{y}_m , and a decomposition of G into blocks G_m . The block G_m depends on y_m and \dot{y}_m , and also on components of $y_{m'}$ and $\dot{y}_{m'}$ associated with neighboring sub-domains

- `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.

IDABBDCommFn

- Definition `typedef int (*IDABBDCommFn)(long int Nlocal, realtype tt, N_Vector yy, N_Vector yp, void *user_data);`
- Purpose 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.
`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 (`tt`, `yy`, `yp`) 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` (see below).

Besides the header files required for the integration of the DAE problem (see §4.3), to use the `IDABBDPRE` module, the main program must include the header file `idas.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 §4.4 are grayed-out.

1. Initialize MPI
2. Set problem dimensions
3. Set vector of initial values
4. Create IDAS object
5. Allocate internal memory
6. Set optional inputs
7. Attach iterative linear solver, one of:

- (a) `flag = IDASpgmr(ida_mem, maxl);`
- (b) `flag = IDASpbcg(ida_mem, maxl);`

```
(c) flag = IDASptfqmr(ida_mem, maxl);
```

8. Initialize the IDABBDPRE preconditioner module

Specify the upper and lower bandwidths `mudq`, `mldq` and `mukeep`, `mlkeep` and call

```
flag = IDABBDPrecInit(ida_mem, Nlocal, mudq, mldq,
                      mukeep, mlkeep, dq_rel_yy, Gres, Gcomm);
```

to allocate memory and initialize the internal preconditioner data. The last two arguments of `IDABBDPrecInit` are the two user-supplied functions described above.

9. Set linear solver optional inputs

Note that the user should not overwrite the preconditioner setup function or solve function through calls to IDASPILS optional input functions.

10. Correct initial values

11. Specify rootfinding problem

12. Advance solution in time

13. Get optional outputs

Additional optional outputs associated with IDABBDPRE are available by way of two routines described below, `IDABBDPrecGetWorkSpace` and `IDABBDPrecGetNumGfnEvals`.

14. Deallocate memory for solution vector

15. Free solver memory

16. Finalize MPI

The user-callable functions that initialize (step 8 above) or re-initialize the IDABBDPRE preconditioner module are described next.

IDABBDPrecInit

Call `flag = IDABBDPrecInit(ida_mem, Nlocal, mudq, mldq, mukeep, mlkeep, dq_rel_yy, Gres, Gcomm);`

Description The function `IDABBDPrecInit` initializes and allocates (internal) memory for the IDABBDPRE preconditioner.

Arguments	<code>ida_mem</code>	(void *) pointer to the IDAS memory block.
	<code>Nlocal</code>	(long int) local vector dimension.
	<code>mudq</code>	(long int) upper half-bandwidth to be used in the difference-quotient Jacobian approximation.
	<code>mldq</code>	(long int) lower half-bandwidth to be used in the difference-quotient Jacobian approximation.
	<code>mukeep</code>	(long int) upper half-bandwidth of the retained banded approximate Jacobian block.
	<code>mlkeep</code>	(long int) lower half-bandwidth of the retained banded approximate Jacobian block.
	<code>dq_rel_yy</code>	(realtype) the relative increment in components of <code>y</code> used in the difference quotient approximations. The default is <code>dq_rel_yy = $\sqrt{\text{unit roundoff}}$</code> , which can be specified by passing <code>dq_rel_yy = 0.0</code> .
	<code>Gres</code>	(IDABBDLocalFn) the C function which computes the local residual approximation $G(t, y, \dot{y})$.

	Gcomm	(IDABBDCommFn) the optional C function which performs all inter-process communication required for the computation of $G(t, y, \dot{y})$.
Return value	The return value flag (of type int) is one of	
	IDASPILS_SUCCESS	The call to IDABBDPrecInit was successful.
	IDASPILS_MEM_NULL	The ida_mem pointer was NULL.
	IDASPILS_MEM_FAIL	A memory allocation request has failed.
	IDASPILS_LMEM_NULL	An IDASPILS linear solver memory was not attached.
	IDASPILS_ILL_INPUT	The supplied vector implementation was not compatible with block band preconditioner.
Notes	<p>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.</p> <p>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.</p> <p>Also, the half-bandwidths mkeep and mlkeep of the retained banded approximate Jacobian block may be even smaller, to reduce storage and computation costs further.</p> <p>For all four half-bandwidths, the values need not be the same on every processor.</p>	

The IDABBDPRE module also provides a reinitialization function to allow for a sequence of problems of the same size with IDASPGMR/IDABBDPRE, IDASPCG/IDABBDPRE, or IDASPTFQMR/IDABBDPRE, provided there is no change in **local_N**, **mkeep**, or **mlkeep**. After solving one problem, and after calling **IDReInit** to re-initialize IDAS for a subsequent problem, a call to **IDABBDPrecReInit** can be made to change any of the following: the half-bandwidths **mudq** and **mldq** used in the difference-quotient Jacobian approximations, the relative increment **dq_rel_yy**, or one of the user-supplied functions **Gres** and **Gcomm**.

IDABBDPrecReInit

Call	flag = IDABBDPrecReInit (ida_mem , mudq , mldq , dq_rel_yy);	
Description	The function IDABBDPrecReInit reinitializes the IDABBDPRE preconditioner.	
Arguments	ida_mem	(void *) pointer to the IDAS memory block.
	mudq	(long int) upper half-bandwidth to be used in the difference-quotient Jacobian approximation.
	mldq	(long int) lower half-bandwidth to be used in the difference-quotient Jacobian approximation.
	dq_rel_yy	(realtype) the relative increment in components of y used in the difference quotient approximations. The default is $\text{dq_rel_yy} = \sqrt{\text{unit roundoff}}$, which can be specified by passing dq_rel_yy = 0.0.
Return value	The return value flag (of type int) is one of	
	IDASPILS_SUCCESS	The call to IDABBDPrecReInit was successful.
	IDASPILS_MEM_NULL	The ida_mem pointer was NULL.
	IDASPILS_LMEM_NULL	An IDASPILS linear solver memory was not attached.
	IDASPILS_PMEM_NULL	The function IDABBDPrecInit was not previously called.
Notes	<p>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.</p>	

The following two optional output functions are available for use with the IDABBDPRE module:

IDABBDPrecGetWorkSpace

Call	<code>flag = IDABBDPrecGetWorkSpace(ida_mem, &lenrwBBDP, &leniwBBDP);</code>
Description	The function <code>IDABBDPrecGetWorkSpace</code> returns the local sizes of the IDABBDPRE real and integer workspaces.
Arguments	<code>ida_mem</code> (void *) pointer to the IDAS memory block. <code>lenrwBBDP</code> (long int) local number of real values in the IDABBDPRE workspace. <code>leniwBBDP</code> (long int) local number of integer values in the IDABBDPRE workspace.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDASPILS_SUCCESS</code> The optional output value has been successfully set. <code>IDASPILS_MEM_NULL</code> The <code>ida_mem</code> pointer was NULL. <code>IDASPILS_PMEM_NULL</code> The IDABBDPRE preconditioner has not been initialized.
Notes	In terms of the local vector dimension N_l , and $\text{smu} = \min(N_l - 1, \text{mukeep} + \text{mlkeep})$, the actual size of the real workspace is $N_l(2 \text{mlkeep} + \text{mukeep} + \text{smu} + 2)$ <code>realtype</code> words. The actual size of the integer workspace is N_l integer words.

IDABBDPrecGetNumGfnEvals

Call	<code>flag = IDABBDPrecGetNumGfnEvals(ida_mem, &ngevalsBBDP);</code>
Description	The function <code>IDABBDPrecGetNumGfnEvals</code> returns the cumulative number of calls to the user <code>Gres</code> function due to the finite difference approximation of the Jacobian blocks used within IDABBDPRE's preconditioner setup function.
Arguments	<code>ida_mem</code> (void *) pointer to the IDAS memory block. <code>ngevalsBBDP</code> (long int) the cumulative number of calls to the user <code>Gres</code> function.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDASPILS_SUCCESS</code> The optional output value has been successfully set. <code>IDASPILS_MEM_NULL</code> The <code>ida_mem</code> pointer was NULL. <code>IDASPILS_PMEM_NULL</code> The IDABBDPRE preconditioner has not been initialized.

In addition to the `ngevalsBBDP` `Gres` evaluations, the costs associated with IDABBDPRE also include `nlinsetups` LU factorizations, `nlinsetups` calls to `Gcomm`, `npsolves` banded backsolve calls, and `nrevalsLS` residual function evaluations, where `nlinsetups` is an optional IDAS output (see §4.5.9.1), and `npsolves` and `nrevalsLS` are linear solver optional outputs (see §4.5.9.6).

Chapter 5

Using IDAS for Forward Sensitivity Analysis

This chapter describes the use of IDAS to compute solution sensitivities using forward sensitivity analysis. One of our main guiding principles was to design the IDAS user interface for forward sensitivity analysis as an extension of that for IVP integration. Assuming a user main program and user-defined support routines for IVP integration have already been defined, in order to perform forward sensitivity analysis the user only has to insert a few more calls into the main program and (optionally) define an additional routine which computes the residuals for sensitivity systems (2.12). The only departure from this philosophy is due to the `IDResFn` type definition (§4.6.1). Without changing the definition of this type, the only way to pass values of the problem parameters to the DAE residual function is to require the user data structure `user_data` to contain a pointer to the array of real parameters p .

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

We begin with a brief overview, in the form of a skeleton user program. Following that are detailed descriptions of the interface to the various user-callable routines and of the user-supplied routines that were not already described in Chapter 4.

5.1 A skeleton of the user's main program

The following is a skeleton of the user's main program (or calling program) as an application of IDAS. The user program is to have these steps in the order indicated, unless otherwise noted. For the sake of brevity, we defer many of the details to the later sections. As in §4.4, most steps are independent of the `NVECTOR` implementation used; where this is not the case, usage specifications are given for the four implementations provided with IDAS: steps marked **[P]** correspond to `NVECTOR_PARALLEL`, steps marked **[O]** correspond to `NVECTOR_OPENMP`, steps marked **[T]** correspond to `NVECTOR_PTHREADS`, while steps marked **[S]** correspond to `NVECTOR_SERIAL`. Differences between the user main program in §4.4 and the one below start only at step (11). Steps that are unchanged from the skeleton program presented in §4.4 are grayed out.

First, note that no additional header files need be included for forward sensitivity analysis beyond those for IVP solution (§4.4).

1. **[P]** Initialize MPI
2. Set problem dimensions
3. Set initial values
4. Create IDAS object
5. Allocate internal memory

6. Specify integration tolerances
7. Set optional inputs
8. Attach linear solver module
9. Set linear solver optional inputs
10. Initialize quadrature problem, if not sensitivity-dependent
11. Define the sensitivity problem

- Number of sensitivities (required)

Set $N_s = N_s$, the number of parameters with respect to which sensitivities are to be computed.

- Problem parameters (optional)

If IDAS is to evaluate the residuals of the sensitivity systems, set \mathbf{p} , an array of N_p real parameters upon which the IVP depends. Only parameters with respect to which sensitivities are (potentially) desired need to be included. Attach \mathbf{p} to the user data structure `user_data`. For example, `user_data->p = p`;

If the user provides a function to evaluate the sensitivity residuals, \mathbf{p} need not be specified.

- Parameter list (optional)

If IDAS is to evaluate the sensitivity residuals, set `plist`, an array of N_s integers to specify the parameters \mathbf{p} with respect to which solution sensitivities are to be computed. If sensitivities with respect to the j -th parameter $\mathbf{p}[j]$ ($0 \leq j < N_p$) are desired, set $\text{plist}_i = j$, for some $i = 0, \dots, N_s - 1$.

If `plist` is not specified, IDAS will compute sensitivities with respect to the first N_s parameters; i.e., $\text{plist}_i = i$ ($i = 0, \dots, N_s - 1$).

If the user provides a function to evaluate the sensitivity residuals, `plist` need not be specified.

- Parameter scaling factors (optional)

If IDAS is to estimate tolerances for the sensitivity solution vectors (based on tolerances for the state solution vector) or if IDAS is to evaluate the residuals of the sensitivity systems using the internal difference-quotient function, the results will be more accurate if order of magnitude information is provided.

Set `pbar`, an array of N_s positive scaling factors. Typically, if $p_i \neq 0$, the value $\bar{p}_i = |p_{\text{plist}_i}|$ can be used.

If `pbar` is not specified, IDAS will use $\bar{p}_i = 1.0$.

If the user provides a function to evaluate the sensitivity residual and specifies tolerances for the sensitivity variables, `pbar` need not be specified.

Note that the names for `p`, `pbar`, `plist`, as well as the field `p` of `user_data` are arbitrary, but they must agree with the arguments passed to `IDASetsSensParams` below.

12. Set sensitivity initial conditions

To set the sensitivities vectors `yS0` and `ypS0` to initial values, use functions defined by a particular `NVECTOR` implementation.

For sensitivity vectors `yS0`, set the N_s N -vectors `yS0[i]` of initial values for sensitivities (for $i = 0, \dots, N_s - 1$).

First, create an array of N_s vectors by making the call

```
[S] yS0 = N_VCloneVectorArray_Serial(Ns, y0);
```

```
[O] yS0 = N_VCloneVectorArray_OpenMP(Ns, y0);
```

```
[T] yS0 = N_VCloneVectorArray_Pthreads(Ns, y0);
```

```
[P] yS0 = N_VCloneVectorArray_Parallel(Ns, y0);
```

and, for each $i = 0, \dots, Ns-1$, load initial values for the i -th sensitivity vector into the structure defined by:

```
[S] NV_DATA_S(yS0[i])
```

```
[O] NV_DATA_OMP(yS0[i])
```

```
[T] NV_DATA_PT(yS0[i])
```

```
[P] NV_DATA_P(yS0[i])
```

Here the argument `y0` serves only to provide the `N_Vector` type for cloning.

Alternatively, if the initial values for the sensitivity variables are already available in `realttype` arrays, create an array of `Ns` “empty” vectors by making the call

```
[S] yS0 = N_VCloneEmptyVectorArray_Serial(Ns, y0);
```

```
[O] yS0 = N_VCloneEmptyVectorArray_OpenMP(Ns, y0);
```

```
[T] yS0 = N_VCloneEmptyVectorArray_Pthreads(Ns, y0);
```

```
[P] yS0 = N_VCloneEmptyVectorArray_Parallel(Ns, y0);
```

and then attach the `realttype` array `yS0_i` containing the initial values of the i -th sensitivity vector using

```
[S] N_VSetArrayPointer_Serial(yS0_i, yS0[i]);
```

```
[O] N_VSetArrayPointer_OpenMP(yS0_i, yS0[i]);
```

```
[T] N_VSetArrayPointer_Pthreads(yS0_i, yS0[i]);
```

```
[P] N_VSetArrayPointer_Parallel(yS0_i, yS0[i]);
```

for $i = 0, \dots, Ns - 1$.

The initial conditions for the sensitivity derivatives \dot{y} are set similarly.

13. Activate sensitivity calculations

Call `flag = IDASensInit(...)`; to activate forward sensitivity computations and allocate internal memory for IDAS related to sensitivity calculations (see §5.2.1).

14. Set sensitivity tolerances

Call `IDASensSStolerances`, `IDASensSVtolerances`, or `IDASenseEetolerances`. See §5.2.2.

15. Set sensitivity analysis optional inputs

Call `IDASetSens*` routines to change from their default values any optional inputs that control the behavior of IDAS in computing forward sensitivities. See §5.2.6.

16. Correct initial values

17. Specify rootfinding problem

18. Advance solution in time

19. Extract sensitivity solution

After each successful return from `IDASolve`, the solution of the original IVP is available in the `y` argument of `IDASolve`, while the sensitivity solution can be extracted into `yS` and `ypS` (which can be the same as `yS0` and `ypS0`, respectively) by calling one of the following routines: `IDAGetSens`, `IDAGetSens1`, `IDAGetSensDky` or `IDAGetSensDky1` (see §5.2.5).

20. Deallocate memory for solutions vector

21. Deallocate memory for sensitivity vectors

Upon completion of the integration, deallocate memory for the vectors contained in `yS0` and `ypS0`:

```
[S] N_VDestroyVectorArray_Serial(yS0, Ns);
[O] N_VDestroyVectorArray_OpenMP(yS0, Ns);
[T] N_VDestroyVectorArray_Pthreads(yS0, Ns);
[P] N_VDestroyVectorArray_Parallel(yS0, Ns);
```

and similarly for `ypS0`.

If `yS` was created from `realttype` arrays `yS_i`, it is the user's responsibility to also free the space for the arrays `yS_i`, and likewise for `ypS`.

22. Free user data structure

23. Free solver memory

24. Free vector specification memory

5.2 User-callable routines for forward sensitivity analysis

This section describes the IDAS functions, in addition to those presented in §4.5, that are called by the user to set up and solve a forward sensitivity problem.

5.2.1 Forward sensitivity initialization and deallocation functions

Activation of forward sensitivity computation is done by calling `IDASensInit`. The form of the call to this routine is as follows:

<code>IDASensInit</code>

Call	<code>flag = IDASensInit(ida_mem, Ns, ism, resS, yS0, ypS0);</code>
Description	The routine <code>IDASensInit</code> activates forward sensitivity computations and allocates internal memory related to sensitivity calculations.
Arguments	<div> <div><code>ida_mem</code></div> <div>(<code>void *</code>) pointer to the IDAS memory block returned by <code>IDACreate</code>.</div> </div> <div> <div><code>Ns</code></div> <div>(<code>int</code>) the number of sensitivities to be computed.</div> </div> <div> <div><code>ism</code></div> <div>(<code>int</code>) a flag used to select the sensitivity solution method. Its value can be either <code>IDA_SIMULTANEOUS</code> or <code>IDA_STAGGERED</code>:</div> <div> <ul style="list-style-type: none"> • In the <code>IDA_SIMULTANEOUS</code> approach, the state and sensitivity variables are corrected at the same time. If <code>IDA_NEWTON</code> was selected as the nonlinear system solution method, this amounts to performing a modified Newton iteration on the combined nonlinear system; • In the <code>IDA_STAGGERED</code> approach, the correction step for the sensitivity variables takes place at the same time for all sensitivity equations, but only after the correction of the state variables has converged and the state variables have passed the local error test; </div> </div> <div> <div><code>resS</code></div> <div>(<code>IDASensResFn</code>) is the C function which computes the residual of the sensitivity DAE. For full details see §5.3.</div> </div> <div> <div><code>yS0</code></div> <div>(<code>N_Vector *</code>) a pointer to an array of <code>Ns</code> vectors containing the initial values of the sensitivities of y.</div> </div> <div> <div><code>ypS0</code></div> <div>(<code>N_Vector *</code>) a pointer to an array of <code>Ns</code> vectors containing the initial values of the sensitivities of \dot{y}.</div> </div>

Return value The return value `flag` (of type `int`) will be one of the following:

- `IDA_SUCCESS` The call to `IDASensInit` was successful.
- `IDA_MEM_NULL` The IDAS memory block was not initialized through a previous call to `IDACreate`.
- `IDA_MEM_FAIL` A memory allocation request has failed.
- `IDA_ILL_INPUT` An input argument to `IDASensInit` has an illegal value.

Notes Passing `resS=NULL` indicates using the default internal difference quotient sensitivity residual routine.

If an error occurred, `IDASensInit` also prints an error message to the file specified by the optional input `errfp`.

In terms of the problem size N , number of sensitivity vectors N_s , and maximum method order `maxord`, the size of the real workspace is increased as follows:

- Base value: $\text{lenrw} = \text{lenrw} + (\text{maxord}+5)N_sN$
- With `IDASensSVtolerances`: $\text{lenrw} = \text{lenrw} + N_sN$

the size of the integer workspace is increased as follows:

- Base value: $\text{leniw} = \text{leniw} + (\text{maxord}+5)N_sN_i$
- With `IDASensSVtolerances`: $\text{leniw} = \text{leniw} + N_sN_i$,

where N_i is the number of integer words in one `N_Vector`.

The routine `IDASensReInit`, useful during the solution of a sequence of problems of same size, reinitializes the sensitivity-related internal memory and must follow a call to `IDASensInit` (and maybe a call to `IDAReInit`). The number N_s of sensitivities is assumed to be unchanged since the call to `IDASensInit`. The call to the `IDASensReInit` function has the form:

`IDASensReInit`

Call `flag = IDASensReInit(ida_mem, ism, yS0, ypS0);`

Description The routine `IDASensReInit` reinitializes forward sensitivity computations.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block returned by `IDACreate`.
`ism` (`int`) a flag used to select the sensitivity solution method. Its value can be either `IDA_SIMULTANEOUS` or `IDA_STAGGERED`.
`yS0` (`N_Vector *`) a pointer to an array of N_s variables of type `N_Vector` containing the initial values of the sensitivities of y .
`ypS0` (`N_Vector *`) a pointer to an array of N_s variables of type `N_Vector` containing the initial values of the sensitivities of \dot{y} .

Return value The return value `flag` (of type `int`) will be one of the following:

- `IDA_SUCCESS` The call to `IDAReInit` was successful.
- `IDA_MEM_NULL` The IDAS memory block was not initialized through a previous call to `IDACreate`.
- `IDA_NO_SENS` Memory space for sensitivity integration was not allocated through a previous call to `IDASensInit`.
- `IDA_ILL_INPUT` An input argument to `IDASensReInit` has an illegal value.
- `IDA_MEM_FAIL` A memory allocation request has failed.

Notes All arguments of `IDASensReInit` are the same as those of `IDASensInit`.

If an error occurred, `IDASensReInit` also prints an error message to the file specified by the optional input `errfp`.

To deallocate all forward sensitivity-related memory (allocated in a prior call to `IDASensInit`), the user must call

IDASensFree

Call	<code>IDASensFree(ida_mem);</code>
Description	The function <code>IDASensFree</code> frees the memory allocated for forward sensitivity computations by a previous call to <code>IDASensInit</code> .
Arguments	The argument is the pointer to the IDAS memory block (of type <code>void *</code>).
Return value	The function <code>IDASensFree</code> has no return value.
Notes	In general, <code>IDASensFree</code> need not be called by the user as it is invoked automatically by <code>IDAFree</code> . After a call to <code>IDASensFree</code> , forward sensitivity computations can be reactivated only by calling <code>IDASensInit</code> again.

To activate and deactivate forward sensitivity calculations for successive IDAS runs, without having to allocate and deallocate memory, the following function is provided:

IDASensToggleOff

Call	<code>IDASensToggleOff(ida_mem);</code>
Description	The function <code>IDASensToggleOff</code> deactivates forward sensitivity calculations. It does <i>not</i> deallocate sensitivity-related memory.
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the memory previously allocated by <code>IDAInit</code> .
Return value	The return value <code>flag</code> of <code>IDASensToggle</code> is one of: <code>IDA_SUCCESS</code> <code>IDASensToggleOff</code> was successful. <code>IDA_MEM_NULL</code> <code>ida_mem</code> was <code>NULL</code> .
Notes	Since sensitivity-related memory is not deallocated, sensitivities can be reactivated at a later time (using <code>IDASensReInit</code>).

5.2.2 Forward sensitivity tolerance specification functions

One of the following three functions must be called to specify the integration tolerances for sensitivities. Note that this call must be made after the call to `IDASensInit`.

IDASensSStolerances

Call	<code>flag = IDASensSStolerances(ida_mem, reltolS, abstolS);</code>
Description	The function <code>IDASensSStolerances</code> specifies scalar relative and absolute tolerances.
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block returned by <code>IDACreate</code> . <code>reltolS</code> (<code>realtype</code>) is the scalar relative error tolerance. <code>abstolS</code> (<code>realtype*</code>) is a pointer to an array of length <code>Ns</code> containing the scalar absolute error tolerances.
Return value	The return flag <code>flag</code> (of type <code>int</code>) will be one of the following: <code>IDA_SUCCESS</code> The call to <code>IDASSStolerances</code> was successful. <code>IDA_MEM_NULL</code> The IDAS memory block was not initialized through a previous call to <code>IDACreate</code> . <code>IDA_NO_SENS</code> The sensitivity allocation function <code>IDASensInit</code> has not been called. <code>IDA_ILL_INPUT</code> One of the input tolerances was negative.

IDASensSVtolerances

Call	<code>flag = IDASensSVtolerances(ida_mem, reltolS, abstolS);</code>
Description	The function <code>IDASensSVtolerances</code> specifies scalar relative tolerance and vector absolute tolerances.
Arguments	<p><code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block returned by <code>IDACreate</code>.</p> <p><code>reltolS</code> (<code>realtype</code>) is the scalar relative error tolerance.</p> <p><code>abstolS</code> (<code>N_Vector*</code>) is an array of <code>Ns</code> variables of type <code>N_Vector</code>. The <code>N_Vector</code> from <code>abstolS[is]</code> specifies the vector tolerances for <code>is</code>-th sensitivity.</p>
Return value	<p>The return flag <code>flag</code> (of type <code>int</code>) will be one of the following:</p> <p><code>IDA_SUCCESS</code> The call to <code>IDASVtolerances</code> was successful.</p> <p><code>IDA_MEM_NULL</code> The IDAS memory block was not initialized through a previous call to <code>IDACreate</code>.</p> <p><code>IDA_NO_SENS</code> The sensitivity allocation function <code>IDASensInit</code> has not been called.</p> <p><code>IDA_ILL_INPUT</code> The relative error tolerance was negative or one of the absolute tolerance vectors had a negative component.</p>
Notes	This choice of tolerances is important when the absolute error tolerance needs to be different for each component of any vector <code>yS[i]</code> .

IDASenseEtolerances

Call	<code>flag = IDASenseEtolerances(ida_mem);</code>
Description	When <code>IDASenseEtolerances</code> is called, IDAS will estimate tolerances for sensitivity variables based on the tolerances supplied for states variables and the scaling factors \bar{p} .
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block returned by <code>IDACreate</code> .
Return value	<p>The return flag <code>flag</code> (of type <code>int</code>) will be one of the following:</p> <p><code>IDA_SUCCESS</code> The call to <code>IDASenseEtolerances</code> was successful.</p> <p><code>IDA_MEM_NULL</code> The IDAS memory block was not initialized through a previous call to <code>IDACreate</code>.</p> <p><code>IDA_NO_SENS</code> The sensitivity allocation function <code>IDASensInit</code> has not been called.</p>

5.2.3 Forward sensitivity initial condition calculation function

`IDACalcIC` also calculates corrected initial conditions for sensitivity variables of a DAE system. When used for initial conditions calculation of the forward sensitivities, `IDACalcIC` must be preceded by successful calls to `IDASensInit` (or `IDASensReInit`) and should precede the call(s) to `IDASolve`. For restrictions that apply for initial conditions calculation of the state variables, see §4.5.4.

Calling `IDACalcIC` is optional. It is only necessary when the initial conditions do not satisfy the sensitivity systems. Even if forward sensitivity analysis was enabled, the call to the initial conditions calculation function `IDACalcIC` is exactly the same as for state variables.

```
flag = IDACalcIC(ida_mem, icopt, tout1);
```

See §4.5.4 for a list of possible return values.

5.2.4 IDAS solver function

Even if forward sensitivity analysis was enabled, the call to the main solver function `IDASolve` is exactly the same as in §4.5.6. However, in this case the return value `flag` can also be one of the following:

<code>IDA_SRES_FAIL</code>	The sensitivity residual function failed in an unrecoverable manner.
<code>IDA_REP_SRES_ERR</code>	The user's residual function repeatedly returned a recoverable error flag, but the solver was unable to recover.

5.2.5 Forward sensitivity extraction functions

If forward sensitivity computations have been initialized by a call to `IDASensInit`, or reinitialized by a call to `IDASensReInit`, then IDAS computes both a solution and sensitivities at time `t`. However, `IDASolve` will still return only the solutions y and \dot{y} in `yret` and `ypret`, respectively. Solution sensitivities can be obtained through one of the following functions:

IDAGetSens

Call `flag = IDAGetSens(ida_mem, &tret, yS);`

Description The function `IDAGetSens` returns the sensitivity solution vectors after a successful return from `IDASolve`.

Arguments

- `ida_mem` (`void *`) pointer to the memory previously allocated by `IDAINit`.
- `tret` (`realtype`) the time reached by the solver (output).
- `yS` (`N_Vector *`) the array of `Ns` computed forward sensitivity vectors.

Return value The return value `flag` of `IDAGetSens` is one of:

- `IDA_SUCCESS` `IDAGetSens` was successful.
- `IDA_MEM_NULL` `ida_mem` was `NULL`.
- `IDA_NO_SENS` Forward sensitivity analysis was not initialized.
- `IDA_BAD_DKY` `yS` is `NULL`.

Notes Note that the argument `tret` is an output for this function. Its value will be the same as that returned at the last `IDASolve` call.

The function `IDAGetSensDky` computes the k -th derivatives of the interpolating polynomials for the sensitivity variables at time `t`. This function is called by `IDAGetSens` with $k = 0$, but may also be called directly by the user.

IDAGetSensDky

Call `flag = IDAGetSensDky(ida_mem, t, k, dkyS);`

Description The function `IDAGetSensDky` returns derivatives of the sensitivity solution vectors after a successful return from `IDASolve`.

Arguments

- `ida_mem` (`void *`) pointer to the memory previously allocated by `IDAINit`.
- `t` (`realtype`) specifies the time at which sensitivity information is requested. The time `t` must fall within the interval defined by the last successful step taken by IDAS.
- `k` (`int`) order of derivatives.
- `dkyS` (`N_Vector *`) array of `Ns` vectors containing the derivatives on output. The space for `dkyS` must be allocated by the user.

Return value The return value `flag` of `IDAGetSensDky` is one of:

- `IDA_SUCCESS` `IDAGetSensDky` succeeded.
- `IDA_MEM_NULL` `ida_mem` was `NULL`.
- `IDA_NO_SENS` Forward sensitivity analysis was not initialized.
- `IDA_BAD_DKY` `dkyS` or one of the vectors `dkyS[i]` is `NULL`.
- `IDA_BAD_K` `k` is not in the range $0, 1, \dots, klast$.
- `IDA_BAD_T` The time `t` is not in the allowed range.

Forward sensitivity solution vectors can also be extracted separately for each parameter in turn through the functions `IDAGetSens1` and `IDAGetSensDky1`, defined as follows:

IDAGetSens1

Call	<code>flag = IDAGetSens1(ida_mem, &tret, is, yS);</code>
Description	The function <code>IDAGetSens1</code> returns the <code>is</code> -th sensitivity solution vector after a successful return from <code>IDASolve</code> .
Arguments	<p><code>ida_mem</code> (<code>void *</code>) pointer to the memory previously allocated by <code>IDAInit</code>.</p> <p><code>tret</code> (<code>realtype *</code>) the time reached by the solver (output).</p> <p><code>is</code> (<code>int</code>) specifies which sensitivity vector is to be returned ($0 \leq is < N_s$).</p> <p><code>yS</code> (<code>N_Vector</code>) the computed forward sensitivity vector.</p>
Return value	<p>The return value <code>flag</code> of <code>IDAGetSens1</code> is one of:</p> <p><code>IDA_SUCCESS</code> <code>IDAGetSens1</code> was successful.</p> <p><code>IDA_MEM_NULL</code> <code>ida_mem</code> was <code>NULL</code>.</p> <p><code>IDA_NO_SENS</code> Forward sensitivity analysis was not initialized.</p> <p><code>IDA_BAD_IS</code> The index <code>is</code> is not in the allowed range.</p> <p><code>IDA_BAD_DKY</code> <code>yS</code> is <code>NULL</code>.</p> <p><code>IDA_BAD_T</code> The time <code>t</code> is not in the allowed range.</p>
Notes	Note that the argument <code>tret</code> is an output for this function. Its value will be the same as that returned at the last <code>IDASolve</code> call.


IDAGetSensDky1

Call	<code>flag = IDAGetSensDky1(ida_mem, t, k, is, dkyS);</code>
Description	The function <code>IDAGetSensDky1</code> returns the <code>k</code> -th derivative of the <code>is</code> -th sensitivity solution vector after a successful return from <code>IDASolve</code> .
Arguments	<p><code>ida_mem</code> (<code>void *</code>) pointer to the memory previously allocated by <code>IDAInit</code>.</p> <p><code>t</code> (<code>realtype</code>) specifies the time at which sensitivity information is requested. The time <code>t</code> must fall within the interval defined by the last successful step taken by IDAS.</p> <p><code>k</code> (<code>int</code>) order of derivative.</p> <p><code>is</code> (<code>int</code>) specifies the sensitivity derivative vector to be returned ($0 \leq is < N_s$).</p> <p><code>dkyS</code> (<code>N_Vector</code>) the vector containing the derivative on output. The space for <code>dkyS</code> must be allocated by the user.</p>
Return value	<p>The return value <code>flag</code> of <code>IDAGetSensDky1</code> is one of:</p> <p><code>IDA_SUCCESS</code> <code>IDAGetQuadDky1</code> succeeded.</p> <p><code>IDA_MEM_NULL</code> <code>ida_mem</code> was <code>NULL</code>.</p> <p><code>IDA_NO_SENS</code> Forward sensitivity analysis was not initialized.</p> <p><code>IDA_BAD_DKY</code> <code>dkyS</code> is <code>NULL</code>.</p> <p><code>IDA_BAD_IS</code> The index <code>is</code> is not in the allowed range.</p> <p><code>IDA_BAD_K</code> <code>k</code> is not in the range $0, 1, \dots, klast$.</p> <p><code>IDA_BAD_T</code> The time <code>t</code> is not in the allowed range.</p>

5.2.6 Optional inputs for forward sensitivity analysis

Optional input variables that control the computation of sensitivities can be changed from their default values through calls to `IDASetSens*` functions. Table 5.1 lists all forward sensitivity optional input functions in IDAS which are described in detail in the remainder of this section.

IDASetsensParams

- Call** `flag = IDASetsensParams(ida_mem, p, pbar, plist);`
- Description** The function `IDASetsensParams` specifies problem parameter information for sensitivity calculations.
- Arguments**
- `ida_mem` (`void *`) pointer to the IDAS memory block.
 - `p` (`realtype *`) a pointer to the array of real problem parameters used to evaluate $F(t, y, \dot{y}, p)$. If non-NULL, `p` must point to a field in the user's data structure `user_data` passed to the user's residual function. (See §5.1).
 - `pbar` (`realtype *`) an array of `Ns` positive scaling factors. If non-NULL, `pbar` must have all its components > 0.0 . (See §5.1).
 - `plist` (`int *`) an array of `Ns` non-negative indices to specify which components of `p` to use in estimating the sensitivity equations. If non-NULL, `plist` must have all components ≥ 0 . (See §5.1).
- Return value** The return value `flag` (of type `int`) is one of:
- `IDA_SUCCESS` The optional value has been successfully set.
 - `IDA_MEM_NULL` The `ida_mem` pointer is NULL.
 - `IDA_NO_SENS` Forward sensitivity analysis was not initialized.
 - `IDA_ILL_INPUT` An argument has an illegal value.
-  **Notes** This function must be preceded by a call to `IDASensInit`.

IDASetsensDQMethod

- Call** `flag = IDASetsensDQMethod(ida_mem, DQtype, DQrhomax);`
- Description** The function `IDASetsensDQMethod` specifies the difference quotient strategy in the case in which the residual of the sensitivity equations are to be computed by IDAS.
- Arguments**
- `ida_mem` (`void *`) pointer to the IDAS memory block.
 - `DQtype` (`int`) specifies the difference quotient type and can be either `IDA_CENTERED` or `IDA_FORWARD`.
 - `DQrhomax` (`realtype`) positive value of the selection parameter used in deciding switching between a simultaneous or separate approximation of the two terms in the sensitivity residual.
- Return value** The return value `flag` (of type `int`) is one of:
- `IDA_SUCCESS` The optional value has been successfully set.
 - `IDA_MEM_NULL` The `ida_mem` pointer is NULL.
 - `IDA_ILL_INPUT` An argument has an illegal value.
- Notes** If `DQrhomax = 0.0`, then no switching is performed. The approximation is done simultaneously using either centered or forward finite differences, depending on the value of `DQtype`. For values of `DQrhomax` ≥ 1.0 , the simultaneous approximation is used whenever the estimated finite difference perturbations for states and parameters are within

Table 5.1: Forward sensitivity optional inputs

Optional input	Routine name	Default
Sensitivity scaling factors	<code>IDASetsensParams</code>	NULL
DQ approximation method	<code>IDASetsensDQMethod</code>	centered, 0.0
Error control strategy	<code>IDASetsensErrCon</code>	FALSE
Maximum no. of nonlinear iterations	<code>IDASetsensMaxNonlinIters</code>	3

a factor of `DQrhomax`, and the separate approximation is used otherwise. Note that a value `DQrhomax < 1.0` will effectively disable switching. See §2.5 for more details.

The default value are `DQtype=IDA_CENTERED` and `DQrhomax= 0.0`.

IDASetsensErrCon

Call	<code>flag = IDASetsensErrCon(ida_mem, errconS);</code>
Description	The function <code>IDASetsensErrCon</code> specifies the error control strategy for sensitivity variables.
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block. <code>errconS</code> (<code>boolean_t</code>) specifies whether sensitivity variables are included (<code>TRUE</code>) or not (<code>FALSE</code>) in the error control mechanism.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of: <code>IDA_SUCCESS</code> The optional value has been successfully set. <code>IDA_MEM_NULL</code> The <code>ida_mem</code> pointer is <code>NULL</code> .
Notes	By default, <code>errconS</code> is set to <code>FALSE</code> . If <code>errconS=TRUE</code> then both state variables and sensitivity variables are included in the error tests. If <code>errconS=FALSE</code> then the sensitivity variables are excluded from the error tests. Note that, in any event, all variables are considered in the convergence tests.

IDASetsensMaxNonlinIters

Call	<code>flag = IDASetsensMaxNonlinIters(ida_mem, maxcorS);</code>
Description	The function <code>IDASetsensMaxNonlinIters</code> specifies the maximum number of nonlinear solver iterations for sensitivity variables per step.
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block. <code>maxcorS</code> (<code>int</code>) maximum number of nonlinear solver iterations allowed per step (> 0).
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of: <code>IDA_SUCCESS</code> The optional value has been successfully set. <code>IDA_MEM_NULL</code> The <code>ida_mem</code> pointer is <code>NULL</code> .
Notes	The default value is 3.

5.2.7 Optional outputs for forward sensitivity analysis

5.2.7.1 Main solver optional output functions

Optional output functions that return statistics and solver performance information related to forward sensitivity computations are listed in Table 5.2 and described in detail in the remainder of this section.

IDAGetSensNumResEvals

Call	<code>flag = IDAGetSensNumResEvals(ida_mem, &nfSevals);</code>
Description	The function <code>IDAGetSensNumResEvals</code> returns the number of calls to the sensitivity residual function.
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block. <code>nfSevals</code> (<code>long int</code>) number of calls to the sensitivity residual function.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of: <code>IDA_SUCCESS</code> The optional output value has been successfully set. <code>IDA_MEM_NULL</code> The <code>ida_mem</code> pointer is <code>NULL</code> . <code>IDA_NO_SENS</code> Forward sensitivity analysis was not initialized.

IDAGetNumResEvalsSens

- Call** `flag = IDAGetNumResEvalsSens(ida_mem, &nfevalsS);`
- Description** The function `IDAGetNumResEvalsSens` returns the number of calls to the user's residual function due to the internal finite difference approximation of the sensitivity residuals.
- Arguments** `ida_mem` (`void *`) pointer to the IDAS memory block.
`nfevalsS` (`long int`) number of calls to the user residual function for sensitivity residuals.
- Return value** The return value `flag` (of type `int`) is one of:
`IDA_SUCCESS` The optional output value has been successfully set.
`IDA_MEM_NULL` The `ida_mem` pointer is NULL.
`IDA_NO_SENS` Forward sensitivity analysis was not initialized.
- Notes** This counter is incremented only if the internal finite difference approximation routines are used for the evaluation of the sensitivity residuals.

IDAGetSensNumErrTestFails

- Call** `flag = IDAGetSensNumErrTestFails(ida_mem, &nSetfails);`
- Description** The function `IDAGetSensNumErrTestFails` returns the number of local error test failures for the sensitivity variables that have occurred.
- Arguments** `ida_mem` (`void *`) pointer to the IDAS memory block.
`nSetfails` (`long int`) number of error test failures.
- Return value** The return value `flag` (of type `int`) is one of:
`IDA_SUCCESS` The optional output value has been successfully set.
`IDA_MEM_NULL` The `ida_mem` pointer is NULL.
`IDA_NO_SENS` Forward sensitivity analysis was not initialized.
- Notes** This counter is incremented only if the sensitivity variables have been included in the error test (see `IDASetSensErrCon` in §5.2.6). Even in that case, this counter is not incremented if the `ism=IDA_SIMULTANEOUS` sensitivity solution method has been used.

IDAGetSensNumLinSolvSetups

- Call** `flag = IDAGetSensNumLinSolvSetups(ida_mem, &nlinsetupsS);`
- Description** The function `IDAGetSensNumLinSolvSetups` returns the number of calls to the linear solver setup function due to forward sensitivity calculations.
- Arguments** `ida_mem` (`void *`) pointer to the IDAS memory block.

Table 5.2: Forward sensitivity optional outputs

Optional output	Routine name
No. of calls to sensitivity residual function	<code>IDAGetSensNumResEvals</code>
No. of calls to residual function for sensitivity	<code>IDAGetNumResEvalsSens</code>
No. of sensitivity local error test failures	<code>IDAGetSensNumErrTestFails</code>
No. of calls to lin. solv. setup routine for sens.	<code>IDAGetSensNumLinSolvSetups</code>
Sensitivity-related statistics as a group	<code>IDAGetSensStats</code>
Error weight vector for sensitivity variables	<code>IDAGetSensErrWeights</code>
No. of sens. nonlinear solver iterations	<code>IDAGetSensNumNonlinSolvIters</code>
No. of sens. convergence failures	<code>IDAGetSensNumNonlinSolvConvFails</code>
Sens. nonlinear solver statistics as a group	<code>IDAGetSensNonlinSolvStats</code>

`nlinsetupsS` (long int) number of calls to the linear solver setup function.

Return value The return value `flag` (of type `int`) is one of:

`IDA_SUCCESS` The optional output value has been successfully set.

`IDA_MEM_NULL` The `ida_mem` pointer is `NULL`.

`IDA_NO_SENS` Forward sensitivity analysis was not initialized.

Notes This counter is incremented only if Newton iteration has been used and staggered sensitivity solution method (`ism=IDA_STAGGERED`) was specified in the call to `IDASensInit` (see §5.2.1).

`IDAGetSensStats`

Call `flag = IDAGetSensStats(ida_mem, &nfSevals, &nfevalsS, &nSetfails, &nlinsetupsS);`

Description The function `IDAGetSensStats` returns all of the above sensitivity-related solver statistics as a group.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.
`nfSevals` (long int) number of calls to the sensitivity residual function.
`nfevalsS` (long int) number of calls to the user-supplied residual function.
`nSetfails` (long int) number of error test failures.
`nlinsetupsS` (long int) number of calls to the linear solver setup function.

Return value The return value `flag` (of type `int`) is one of:

`IDA_SUCCESS` The optional output values have been successfully set.

`IDA_MEM_NULL` The `ida_mem` pointer is `NULL`.

`IDA_NO_SENS` Forward sensitivity analysis was not initialized.

`IDAGetSensErrWeights`

Call `flag = IDAGetSensErrWeights(ida_mem, eSweight);`

Description The function `IDAGetSensErrWeights` returns the sensitivity error weight vectors at the current time. These are the reciprocals of the W_i of (2.7) for the sensitivity variables.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.
`eSweight` (N_Vector_S) pointer to the array of error weight vectors.

Return value The return value `flag` (of type `int`) is one of:

`IDA_SUCCESS` The optional output value has been successfully set.

`IDA_MEM_NULL` The `ida_mem` pointer is `NULL`.

`IDA_NO_SENS` Forward sensitivity analysis was not initialized.

Notes The user must allocate memory for `eweightS`.

`IDAGetSensNumNonlinSolvIters`

Call `flag = IDAGetSensNumNonlinSolvIters(ida_mem, &nSniters);`

Description The function `IDAGetSensNumNonlinSolvIters` returns the number of nonlinear iterations performed for sensitivity calculations.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.
`nSniters` (long int) number of nonlinear iterations performed.

Return value The return value `flag` (of type `int`) is one of:

`IDA_SUCCESS` The optional output value has been successfully set.

IDA_MEM_NULL The `ida_mem` pointer is NULL.
IDA_NO_SENS Forward sensitivity analysis was not initialized.
 Notes This counter is incremented only if `ism` was `IDA_STAGGERED` in the call to `IDASensInit` (see §5.2.1).

IDAGetSensNumNonlinSolvConvFails

Call `flag = IDAGetSensNumNonlinSolvConvFails(ida_mem, &nSncfails);`
 Description The function `IDAGetSensNumNonlinSolvConvFails` returns the number of nonlinear convergence failures that have occurred for sensitivity calculations.
 Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
 `nSncfails` (`long int`) number of nonlinear convergence failures.
 Return value The return value `flag` (of type `int`) is one of:
 IDA_SUCCESS The optional output value has been successfully set.
 IDA_MEM_NULL The `ida_mem` pointer is NULL.
 IDA_NO_SENS Forward sensitivity analysis was not initialized.
 Notes This counter is incremented only if `ism` was `IDA_STAGGERED` in the call to `IDASensInit` (see §5.2.1).

IDAGetSensNonlinSolvStats

Call `flag = IDAGetSensNonlinSolvStats(ida_mem, &nSniters, &nSncfails);`
 Description The function `IDAGetSensNonlinSolvStats` returns the sensitivity-related nonlinear solver statistics as a group.
 Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
 `nSniters` (`long int`) number of nonlinear iterations performed.
 `nSncfails` (`long int`) number of nonlinear convergence failures.
 Return value The return value `flag` (of type `int`) is one of:
 IDA_SUCCESS The optional output values have been successfully set.
 IDA_MEM_NULL The `ida_mem` pointer is NULL.
 IDA_NO_SENS Forward sensitivity analysis was not initialized.

5.2.7.2 Initial condition calculation optional output functions

The sensitivity consistent initial conditions found by IDAS (after a successful call to `IDACalcIC`) can be obtained by calling the following function:

IDAGetSensConsistentIC

Call `flag = IDAGetSensConsistentIC(ida_mem, yyS0_mod, ypS0_mod);`
 Description The function `IDAGetSensConsistentIC` returns the corrected initial conditions calculated by `IDACalcIC` for sensitivities variables.
 Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
 `yyS0_mod` (`N_Vector *`) a pointer to an array of `Ns` vectors containing consistent sensitivity vectors.
 `ypS0_mod` (`N_Vector *`) a pointer to an array of `Ns` vectors containing consistent sensitivity derivative vectors.
 Return value The return value `flag` (of type `int`) is one of
 IDA_SUCCESS `IDAGetSensConsistentIC` succeeded.

	IDA_MEM_NULL	The <code>ida_mem</code> pointer is NULL.
	IDA_NO_SENS	The function <code>IDASensInit</code> has not been previously called.
	IDA_ILL_INPUT	<code>IDASolve</code> has been already called.
Notes		If the consistent sensitivity vectors or consistent derivative vectors are not desired, pass NULL for the corresponding argument.
		The user must allocate space for <code>yyS0_mod</code> and <code>ypS0_mod</code> (if not NULL).



5.3 User-supplied routines for forward sensitivity analysis

In addition to the required and optional user-supplied routines described in §4.6, when using IDAS for forward sensitivity analysis, the user has the option of providing a routine that calculates the residual of the sensitivity equations (2.12).

By default, IDAS uses difference quotient approximation routines for the residual of the sensitivity equations. However, IDAS allows the option for user-defined sensitivity residual routines (which also provides a mechanism for interfacing IDAS to routines generated by automatic differentiation).

The user may provide the residuals of the sensitivity equations (2.12), for all sensitivity parameters at once, through a function of type `IDASensResFn` defined by:

IDASensResFn

Definition	<pre>typedef int (*IDASensResFn)(int Ns, realtype t, N_Vector yy, N_Vector yp, N_Vector resval, N_Vector *yS, N_Vector *ypS, N_Vector *resvalS, void *user_data, N_Vector tmp1, N_Vector tmp2, N_Vector tmp3);</pre>	
Purpose	This function computes the sensitivity residual for all sensitivity equations. It must compute the vectors $(\partial F/\partial y)s_i(t) + (\partial F/\partial \dot{y})\dot{s}_i(t) + (\partial F/\partial p_i)$ and store them in <code>resvalS[i]</code> .	
Arguments	t	is the current value of the independent variable.
	yy	is the current value of the state vector, $y(t)$.
	yp	is the current value of $\dot{y}(t)$.
	resval	contains the current value F of the original DAE residual.
	yS	contains the current values of the sensitivities s_i .
	ypS	contains the current values of the sensitivity derivatives \dot{s}_i .
	resvalS	contains the output sensitivity residual vectors.
	user_data	is a pointer to user data.
	tmp1	
	tmp2	
	tmp3	are <code>N_Vectors</code> of length N which can be used as temporary storage.
Return value	An <code>IDASensResFn</code> should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and <code>IDA_SRES_FAIL</code> is returned).	
Notes	There is one situation in which recovery is not possible even if <code>IDASensResFn</code> function returns a recoverable error flag. That is when this occurs at the very first call to the <code>IDASensResFn</code> , in which case IDAS returns <code>IDA_FIRST_RES_FAIL</code> .	

5.4 Integration of quadrature equations depending on forward sensitivities

IDAS provides support for integration of quadrature equations that depends not only on the state variables but also on forward sensitivities.

The following is an overview of the sequence of calls in a user's main program in this situation. Steps that are unchanged from the skeleton program presented in §5.1 are grayed out.

1. [P] Initialize MPI
2. Set problem dimensions
3. Set vectors of initial values
4. Create IDAS object
5. Allocate internal memory
6. Set optional inputs
7. Attach linear solver module
8. Set linear solver optional inputs
9. Initialize sensitivity-independent quadrature problem
10. Define the sensitivity problem
11. Set sensitivity initial conditions
12. Activate sensitivity calculations
13. Set sensitivity analysis optional inputs

14. Set vector of initial values for quadrature variables

Typically, the quadrature variables should be initialized to 0.

15. Initialize sensitivity-dependent quadrature integration

Call `IDAQuadSensInit` to specify the quadrature equation right-hand side function and to allocate internal memory related to quadrature integration. See §5.4.1 for details.

16. Set optional inputs for sensitivity-dependent quadrature integration

Call `IDASetQuadSensErrCon` to indicate whether or not quadrature variables should be used in the step size control mechanism. If so, one of the `IDAQuadSens*tolerances` functions must be called to specify the integration tolerances for quadrature variables. See §5.4.4 for details.

17. Advance solution in time

18. Extract sensitivity-dependent quadrature variables

Call `IDAGetQuadSens`, `IDAGetQuadSens1`, `IDAGetQuadSensDky` or `IDAGetQuadSensDky1` to obtain the values of the quadrature variables or their derivatives at the current time. See §5.4.3 for details.

19. Get optional outputs

20. Extract sensitivity solution

21. Get sensitivity-dependent quadrature optional outputs

Call `IDAGetQuadSens*` functions to obtain optional output related to the integration of sensitivity-dependent quadratures. See §5.4.5 for details.

22. Deallocate memory for solutions vector

23. Deallocate memory for sensitivity vectors

24. Deallocate memory for sensitivity-dependent quadrature variables

25. Free solver memory

26. [P] Finalize MPI

Note: IDAQuadSensInit (step 15 above) can be called and quadrature-related optional inputs (step 16 above) can be set, anywhere between steps 10 and 17.

5.4.1 Sensitivity-dependent quadrature initialization and deallocation

The function IDAQuadSensInit activates integration of quadrature equations depending on sensitivities and allocates internal memory related to these calculations. If rhsQS is input as NULL, then IDAS uses an internal function that computes difference quotient approximations to the functions $\bar{q}_i = (\partial q / \partial y)s_i + (\partial q / \partial \dot{y})\dot{s}_i + \partial q / \partial p_i$, in the notation of (2.10). The form of the call to this function is as follows:

IDAQuadSensInit

Call `flag = IDAQuadSensInit(ida_mem, rhsQS, yQS0);`

Description The function IDAQuadSensInit provides required problem specifications, allocates internal memory, and initializes quadrature integration.

Arguments `ida_mem` (void *) pointer to the IDAS memory block returned by IDACreate.
`rhsQS` (IDAQuadSensRhsFn) is the C function which computes f_{QS} , the right-hand side of the sensitivity-dependent quadrature equations (for full details see §5.4.6).
`yQS0` (N_Vector *) contains the initial values of sensitivity-dependent quadratures.

Return value The return value `flag` (of type int) will be one of the following:

`IDA_SUCCESS` The call to IDAQuadSensInit was successful.
`IDA_MEM_NULL` The IDAS memory was not initialized by a prior call to IDACreate.
`IDA_MEM_FAIL` A memory allocation request failed.
`IDA_NO_SENS` The sensitivities were not initialized by a prior call to IDASensInit.
`IDA_ILL_INPUT` The parameter `yQS0` is NULL.

Notes Before calling IDAQuadSensInit, the user must enable the sensitivities by calling IDASensInit.

If an error occurred, IDAQuadSensInit also sends an error message to the error handler function.



In terms of the number of quadrature variables N_q and maximum method order `maxord`, the size of the real workspace is increased as follows:

- Base value: `lenrw = lenrw + (maxord+5) N_q`
- If IDAQuadSensSVtolerances is called: `lenrw = lenrw + N_qN_s`

and the size of the integer workspace is increased as follows:

- Base value: `leniw = leniw + (maxord+5) N_q`
- If IDAQuadSensSVtolerances is called: `leniw = leniw + N_qN_s`

The function IDAQuadSensReInit, useful during the solution of a sequence of problems of same size, reinitializes the quadrature related internal memory and must follow a call to IDAQuadSensInit. The number `Nq` of quadratures as well as the number `Ns` of sensitivities are assumed to be unchanged from the prior call to IDAQuadSensInit. The call to the IDAQuadSensReInit function has the form:

IDAQuadSensReInit

Call	<code>flag = IDAQuadSensReInit(ida_mem, yQS0);</code>
Description	The function <code>IDAQuadSensReInit</code> provides required problem specifications and reinitializes the sensitivity-dependent quadrature integration.
Arguments	<p><code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block.</p> <p><code>yQS0</code> (<code>N.Vector *</code>) contains the initial values of sensitivity-dependent quadratures.</p>
Return value	The return value <code>flag</code> (of type <code>int</code>) will be one of the following:
	<p><code>IDA_SUCCESS</code> The call to <code>IDAQuadSensReInit</code> was successful.</p> <p><code>IDA_MEM_NULL</code> The IDAS memory was not initialized by a prior call to <code>IDACreate</code>.</p> <p><code>IDA_NO_SENS</code> Memory space for the sensitivity calculation was not allocated by a prior call to <code>IDASensInit</code>.</p> <p><code>IDA_NO_QUADSENS</code> Memory space for the sensitivity quadratures integration was not allocated by a prior call to <code>IDAQuadSensInit</code>.</p> <p><code>IDA_ILL_INPUT</code> The parameter <code>yQS0</code> is <code>NULL</code>.</p>
Notes	If an error occurred, <code>IDAQuadSensReInit</code> also sends an error message to the error handler function.

IDAQuadSensFree

Call	<code>IDAQuadSensFree(ida_mem);</code>
Description	The function <code>IDAQuadSensFree</code> frees the memory allocated for sensitivity quadrature integration.
Arguments	The argument is the pointer to the IDAS memory block (of type <code>void *</code>).
Return value	The function <code>IDAQuadSensFree</code> has no return value.
Notes	In general, <code>IDAQuadSensFree</code> need not be called by the user as it is called automatically by <code>IDAFree</code> .

5.4.2 IDAS solver function

Even if quadrature integration was enabled, the call to the main solver function `IDASolve` is exactly the same as in §4.5.6. However, in this case the return value `flag` can also be one of the following:

<code>IDA_QSRHS_FAIL</code>	The sensitivity quadrature right-hand side function failed in an unrecoverable manner.
<code>IDA_FIRST_QSRHS_ERR</code>	The sensitivity quadrature right-hand side function failed at the first call.
<code>IDA_REP_QSRHS_ERR</code>	Convergence test failures occurred too many times due to repeated recoverable errors in the quadrature right-hand side function. The <code>IDA_REP_RES_ERR</code> will also be returned if the quadrature right-hand side function had repeated recoverable errors during the estimation of an initial step size (assuming the sensitivity quadrature variables are included in the error tests).

5.4.3 Sensitivity-dependent quadrature extraction functions

If sensitivity-dependent quadratures have been initialized by a call to `IDAQuadSensInit`, or reinitialized by a call to `IDAQuadSensReInit`, then IDAS computes a solution, sensitivities, and quadratures depending on sensitivities at time `t`. However, `IDASolve` will still return only the solutions `y` and `ȳ`. Sensitivity-dependent quadratures can be obtained using one of the following functions:

IDAGetQuadSens

Call `flag = IDAGetQuadSens(ida_mem, &tret, yQS);`

Description The function `IDAGetQuadSens` returns the quadrature sensitivity solution vectors after a successful return from `IDASolve`.

Arguments

- `ida_mem` (`void *`) pointer to the memory previously allocated by `IDAInit`.
- `tret` (`realtype`) the time reached by the solver (output).
- `yQS` (`N_Vector *`) array of N_s computed sensitivity-dependent quadrature vectors.

Return value The return value `flag` of `IDAGetQuadSens` is one of:

- `IDA_SUCCESS` `IDAGetQuadSens` was successful.
- `IDA_MEM_NULL` `ida_mem` was NULL.
- `IDA_NO_SENS` Sensitivities were not activated.
- `IDA_NO_QUADSENS` Quadratures depending on the sensitivities were not activated.
- `IDA_BAD_DKY` `yQS` or one of the `yQS[i]` is NULL.

The function `IDAGetQuadSensDky` computes the k -th derivatives of the interpolating polynomials for the sensitivity-dependent quadrature variables at time t . This function is called by `IDAGetQuadSens` with $k = 0$, but may also be called directly by the user.

IDAGetQuadSensDky

Call `flag = IDAGetQuadSensDky(ida_mem, t, k, dkyQS);`

Description The function `IDAGetQuadSensDky` returns derivatives of the quadrature sensitivities solution vectors after a successful return from `IDASolve`.

Arguments

- `ida_mem` (`void *`) pointer to the memory previously allocated by `IDAInit`.
- `t` (`realtype`) the time at which information is requested. The time t must fall within the interval defined by the last successful step taken by IDAS.
- `k` (`int`) order of the requested derivative.
- `dkyQS` (`N_Vector *`) array of N_s vectors containing the derivatives. This vector array must be allocated by the user.

Return value The return value `flag` of `IDAGetQuadSensDky` is one of:

- `IDA_SUCCESS` `IDAGetQuadSensDky` succeeded.
- `IDA_MEM_NULL` `ida_mem` was NULL.
- `IDA_NO_SENS` Sensitivities were not activated.
- `IDA_NO_QUADSENS` Quadratures depending on the sensitivities were not activated.
- `IDA_BAD_DKY` `dkyQS` or one of the vectors `dkyQS[i]` is NULL.
- `IDA_BAD_K` k is not in the range $0, 1, \dots, klast$.
- `IDA_BAD_T` The time t is not in the allowed range.

Quadrature sensitivity solution vectors can also be extracted separately for each parameter in turn through the functions `IDAGetQuadSens1` and `IDAGetQuadSensDky1`, defined as follows:

IDAGetQuadSens1

Call `flag = IDAGetQuadSens1(ida_mem, &tret, is, yQS);`

Description The function `IDAGetQuadSens1` returns the is -th sensitivity of quadratures after a successful return from `IDASolve`.

Arguments

- `ida_mem` (`void *`) pointer to the memory previously allocated by `IDAInit`.
- `tret` (`realtype`) the time reached by the solver (output).
- `is` (`int`) specifies which sensitivity vector is to be returned ($0 \leq is < N_s$).
- `yQS` (`N_Vector`) the computed sensitivity-dependent quadrature vector.

Return value The return value `flag` of `IDAGetQuadSens1` is one of:

<code>IDA_SUCCESS</code>	<code>IDAGetQuadSens1</code> was successful.
<code>IDA_MEM_NULL</code>	<code>ida_mem</code> was NULL.
<code>IDA_NO_SENS</code>	Forward sensitivity analysis was not initialized.
<code>IDA_NO_QUADSENS</code>	Quadratures depending on the sensitivities were not activated.
<code>IDA_BAD_IS</code>	The index <code>is</code> is not in the allowed range.
<code>IDA_BAD_DKY</code>	<code>yQS</code> is NULL.

`IDAGetQuadSensDky1`

Call `flag = IDAGetQuadSensDky1(ida_mem, t, k, is, dkyQS);`

Description The function `IDAGetQuadSensDky1` returns the `k`-th derivative of the `is`-th sensitivity solution vector after a successful return from `IDASolve`.

Arguments `ida_mem` (`void *`) pointer to the memory previously allocated by `IDAInit`.
`t` (`realtype`) specifies the time at which sensitivity information is requested. The time `t` must fall within the interval defined by the last successful step taken by IDAS.
`k` (`int`) order of derivative.
`is` (`int`) specifies the sensitivity derivative vector to be returned ($0 \leq is < N_s$).
`dkyQS` (`N_Vector`) the vector containing the derivative. The space for `dkyQS` must be allocated by the user.

Return value The return value `flag` of `IDAGetQuadSensDky1` is one of:

<code>IDA_SUCCESS</code>	<code>IDAGetQuadDky1</code> succeeded.
<code>IDA_MEM_NULL</code>	<code>ida_mem</code> was NULL.
<code>IDA_NO_SENS</code>	Forward sensitivity analysis was not initialized.
<code>IDA_NO_QUADSENS</code>	Quadratures depending on the sensitivities were not activated.
<code>IDA_BAD_DKY</code>	<code>dkyQS</code> is NULL.
<code>IDA_BAD_IS</code>	The index <code>is</code> is not in the allowed range.
<code>IDA_BAD_K</code>	<code>k</code> is not in the range $0, 1, \dots, klast$.
<code>IDA_BAD_T</code>	The time <code>t</code> is not in the allowed range.

5.4.4 Optional inputs for sensitivity-dependent quadrature integration

IDAS provides the following optional input functions to control the integration of sensitivity-dependent quadrature equations.

`IDASetQuadSensErrCon`

Call `flag = IDASetQuadSensErrCon(ida_mem, errconQS)`

Description The function `IDASetQuadSensErrCon` specifies whether or not the quadrature variables are to be used in the local error control mechanism. If they are, the user must specify the error tolerances for the quadrature variables by calling `IDAQuadSensSStolerances`, `IDAQuadSensSVtolerances`, or `IDAQuadSensEETolerances`.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`errconQS` (`booleantype`) specifies whether sensitivity quadrature variables are included (`TRUE`) or not (`FALSE`) in the error control mechanism.

Return value The return value `flag` (of type `int`) is one of:

<code>IDA_SUCCESS</code>	The optional value has been successfully set.
<code>IDA_MEM_NULL</code>	The <code>ida_mem</code> pointer is NULL.

IDA_NO_SENS Sensitivities were not activated.
IDA_NO_QUADSENS Quadratures depending on the sensitivities were not activated.
 Notes By default, **errconQS** is set to **FALSE**.
 It is illegal to call **IDASetQuadSensErrCon** before a call to **IDAQuadSensInit**.



If the quadrature variables are part of the step size control mechanism, one of the following functions must be called to specify the integration tolerances for quadrature variables.

IDAQuadSensSStolerances

Call **flag = IDAQuadSensSVtolerances(ida_mem, reltolQS, abstolQS);**
 Description The function **IDAQuadSensSStolerances** specifies scalar relative and absolute tolerances.
 Arguments **ida_mem** (void *) pointer to the IDAS memory block.
 reltolQS (realtype) is the scalar relative error tolerance.
 abstolQS (realtype*) is a pointer to an array containing the **Ns** scalar absolute error tolerances.
 Return value The return value **flag** (of type **int**) is one of:
 IDA_SUCCESS The optional value has been successfully set.
 IDA_MEM_NULL The **ida_mem** pointer is **NULL**.
 IDA_NO_SENS Sensitivities were not activated.
 IDA_NO_QUADSENS Quadratures depending on the sensitivities were not activated.
 IDA_ILL_INPUT One of the input tolerances was negative.

IDAQuadSensSVtolerances

Call **flag = IDAQuadSensSVtolerances(ida_mem, reltolQS, abstolQS);**
 Description The function **IDAQuadSensSVtolerances** specifies scalar relative and vector absolute tolerances.
 Arguments **ida_mem** (void *) pointer to the IDAS memory block.
 reltolQS (realtype) is the scalar relative error tolerance.
 abstolQS (**N_Vector***) is an array of **Ns** variables of type **N_Vector**. The **N_Vector** from **abstolS[is]** specifies the vector tolerances for **is**-th quadrature sensitivity.
 Return value The return value **flag** (of type **int**) is one of:
 IDA_SUCCESS The optional value has been successfully set.
 IDA_NO_QUAD Quadrature integration was not initialized.
 IDA_MEM_NULL The **ida_mem** pointer is **NULL**.
 IDA_NO_SENS Sensitivities were not activated.
 IDA_NO_QUADSENS Quadratures depending on the sensitivities were not activated.
 IDA_ILL_INPUT One of the input tolerances was negative.

IDAQuadSensEETolerances

Call **flag = IDAQuadSensEETolerances(ida_mem);**
 Description The function **IDAQuadSensEETolerances** specifies that the tolerances for the sensitivity-dependent quadratures should be estimated from those provided for the pure quadrature variables.
 Arguments **ida_mem** (void *) pointer to the IDAS memory block.
 Return value The return value **flag** (of type **int**) is one of:

IDA_SUCCESS The optional value has been successfully set.
 IDA_MEM_NULL The `ida_mem` pointer is NULL.
 IDA_NO_SENS Sensitivities were not activated.
 IDA_NO_QUADSENS Quadratures depending on the sensitivities were not activated.

Notes When `IDAQuadSenseEetolerances` is used, before calling `IDASolve`, integration of pure quadratures must be initialized (see 4.7.1) and tolerances for pure quadratures must be also specified (see 4.7.4).

5.4.5 Optional outputs for sensitivity-dependent quadrature integration

IDAS provides the following functions that can be used to obtain solver performance information related to quadrature integration.

`IDAGetQuadSensNumRhsEvals`

Call `flag = IDAGetQuadSensNumRhsEvals(ida_mem, &nrhsQSevals);`
 Description The function `IDAGetQuadSensNumRhsEvals` returns the number of calls made to the user's quadrature right-hand side function.
 Arguments `ida_mem` (void *) pointer to the IDAS memory block.
`nrhsQSevals` (long int) number of calls made to the user's `rhsQS` function.
 Return value The return value `flag` (of type `int`) is one of:
 IDA_SUCCESS The optional output value has been successfully set.
 IDA_MEM_NULL The `ida_mem` pointer is NULL.
 IDA_NO_QUADSENS Sensitivity-dependent quadrature integration has not been initialized.

`IDAGetQuadSensNumErrTestFails`

Call `flag = IDAGetQuadSensNumErrTestFails(ida_mem, &nQSetfails);`
 Description The function `IDAGetQuadSensNumErrTestFails` returns the number of local error test failures due to quadrature variables.
 Arguments `ida_mem` (void *) pointer to the IDAS memory block.
`nQSetfails` (long int) number of error test failures due to quadrature variables.
 Return value The return value `flag` (of type `int`) is one of:
 IDA_SUCCESS The optional output value has been successfully set.
 IDA_MEM_NULL The `ida_mem` pointer is NULL.
 IDA_NO_QUADSENS Sensitivity-dependent quadrature integration has not been initialized.

`IDAGetQuadSensErrWeights`

Call `flag = IDAGetQuadSensErrWeights(ida_mem, eQSweight);`
 Description The function `IDAGetQuadSensErrWeights` returns the quadrature error weights at the current time.
 Arguments `ida_mem` (void *) pointer to the IDAS memory block.
`eQSweight` (`N_Vector *`) array of quadrature error weight vectors at the current time.
 Return value The return value `flag` (of type `int`) is one of:
 IDA_SUCCESS The optional output value has been successfully set.
 IDA_MEM_NULL The `ida_mem` pointer is NULL.
 IDA_NO_QUADSENS Sensitivity-dependent quadrature integration has not been initialized.



Notes The user must allocate memory for `eQSweight`.
If quadratures were not included in the error control mechanism (through a call to `IDASetQuadSensErrCon` with `errconQS=TRUE`), `IDAGetQuadSensErrWeights` does not set the `eQSweight` vector.

`IDAGetQuadSensStats`

Call `flag = IDAGetQuadSensStats(ida_mem, &nrhsQSevals, &nQSetfails);`
Description The function `IDAGetQuadSensStats` returns the IDAS integrator statistics as a group.
Arguments `ida_mem` (void *) pointer to the IDAS memory block.
`nrhsQSevals` (long int) number of calls to the user's `rhsQS` function.
`nQSetfails` (long int) number of error test failures due to quadrature variables.
Return value The return value `flag` (of type `int`) is one of
`IDA_SUCCESS` the optional output values have been successfully set.
`IDA_MEM_NULL` the `ida_mem` pointer is NULL.
`IDA_NO_QUADSENS` Sensitivity-dependent quadrature integration has not been initialized.

5.4.6 User-supplied function for sensitivity-dependent quadrature integration

For the integration of sensitivity-dependent quadrature equations, the user must provide a function that defines the right-hand side of the sensitivity quadrature equations. For sensitivities of quadratures (2.10) with integrands q , the appropriate right-hand side functions are given by $\bar{q}_i = (\partial q / \partial y)s_i + (\partial q / \partial \dot{y})\dot{s}_i + \partial q / \partial p_i$. This user function must be of type `IDAQuadSensRhsFn`, defined as follows:

`IDAQuadSensRhsFn`

Definition `typedef int (*IDAQuadSensRhsFn)(int Ns, realtype t, N_Vector yy, N_Vector yp, N_Vector *yyS, N_Vector *ypS, N_Vector rrQ, N_Vector *rhsvalQS, void *user_data, N_Vector tmp1, N_Vector tmp2, N_Vector tmp3)`
Purpose This function computes the sensitivity quadrature equation right-hand side for a given value of the independent variable t and state vector y .
Arguments `Ns` is the number of sensitivity vectors.
`t` is the current value of the independent variable.
`yy` is the current value of the dependent variable vector, $y(t)$.
`yp` is the current value of the dependent variable vector, $\dot{y}(t)$.
`yyS` is an array of `Ns` variables of type `N_Vector` containing the dependent sensitivity vectors s_i .
`ypS` is an array of `Ns` variables of type `N_Vector` containing the dependent sensitivity derivatives \dot{s}_i .
`rrQ` is the current value of the quadrature right-hand side q .
`rhsvalQS` contains the `Ns` output vectors.
`user_data` is the `user_data` pointer passed to `IDASetUserData`.
`tmp1`
`tmp2`
`tmp3` are `N_Vectors` which can be used as temporary storage.

Return value An `IDAQuadSensRhsFn` should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and `IDA_QRHS_FAIL` is returned).

Notes Allocation of memory for `rhsvalQS` is automatically handled within IDAS.

Both `yy` and `yp` are of type `N_Vector` and both `yyS` and `ypS` are pointers to an array containing `Ns` vectors of type `N_Vector`. It is the user's responsibility to access the vector data consistently (including the use of the correct accessor macros from each `NVECTOR` implementation). For the sake of computational efficiency, the vector functions in the two `NVECTOR` implementations provided with IDAS do not perform any consistency checks with respect to their `N_Vector` arguments (see §7.1 and §7.2).

There is one situation in which recovery is not possible even if `IDAQuadSensRhsFn` function returns a recoverable error flag. That is when this occurs at the very first call to the `IDAQuadSensRhsFn`, in which case IDAS returns `IDA_FIRST_QSRHS_ERR`.

5.5 Note on using partial error control

For some problems, when sensitivities are excluded from the error control test, the behavior of IDAS may appear at first glance to be erroneous. One would expect that, in such cases, the sensitivity variables would not influence in any way the step size selection.

The short explanation of this behavior is that the step size selection implemented by the error control mechanism in IDAS is based on the magnitude of the correction calculated by the nonlinear solver. As mentioned in §5.2.1, even with partial error control selected in the call to `IDASensInit`, the sensitivity variables are included in the convergence tests of the nonlinear solver.

When using the simultaneous corrector method (§2.5), the nonlinear system that is solved at each step involves both the state and sensitivity equations. In this case, it is easy to see how the sensitivity variables may affect the convergence rate of the nonlinear solver and therefore the step size selection. The case of the staggered corrector approach is more subtle. The sensitivity variables at a given step are computed only once the solver for the nonlinear state equations has converged. However, if the nonlinear system corresponding to the sensitivity equations has convergence problems, IDAS will attempt to improve the initial guess by reducing the step size in order to provide a better prediction of the sensitivity variables. Moreover, even if there are no convergence failures in the solution of the sensitivity system, IDAS may trigger a call to the linear solver's setup routine which typically involves reevaluation of Jacobian information (Jacobian approximation in the case of `IDADENSE` and `IDABAND`, or preconditioner data in the case of the Krylov solvers). The new Jacobian information will be used by subsequent calls to the nonlinear solver for the state equations and, in this way, potentially affect the step size selection.

When using the simultaneous corrector method it is not possible to decide whether nonlinear solver convergence failures or calls to the linear solver setup routine have been triggered by convergence problems due to the state or the sensitivity equations. When using one of the staggered corrector methods, however, these situations can be identified by carefully monitoring the diagnostic information provided through optional outputs. If there are no convergence failures in the sensitivity nonlinear solver, and none of the calls to the linear solver setup routine were made by the sensitivity nonlinear solver, then the step size selection is not affected by the sensitivity variables.

Finally, the user must be warned that the effect of appending sensitivity equations to a given system of DAEs on the step size selection (through the mechanisms described above) is problem-dependent and can therefore lead to either an increase or decrease of the total number of steps that IDAS takes to complete the simulation. At first glance, one would expect that the impact of the sensitivity variables, if any, would be in the direction of increasing the step size and therefore reducing the total number of steps. The argument for this is that the presence of the sensitivity variables in the convergence test of the nonlinear solver can only lead to additional iterations (and therefore a smaller iteration error), or to additional calls to the linear solver setup routine (and therefore more up-to-date Jacobian information), both of which will lead to larger steps being taken by IDAS. However, this is true only

locally. Overall, a larger integration step taken at a given time may lead to step size reductions at later times, due to either nonlinear solver convergence failures or error test failures.

Chapter 6

Using IDAS for Adjoint Sensitivity Analysis

This chapter describes the use of IDAS to compute sensitivities of derived functions using adjoint sensitivity analysis. As mentioned before, the adjoint sensitivity module of IDAS provides the infrastructure for integrating backward in time any system of DAEs that depends on the solution of the original IVP, by providing various interfaces to the main IDAS integrator, as well as several supporting user-callable functions. For this reason, in the following sections we refer to the *backward problem* and not to the *adjoint problem* when discussing details relevant to the DAEs that are integrated backward in time. The backward problem can be the adjoint problem (2.20) or (2.25), and can be augmented with some quadrature differential equations.

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

We begin with a brief overview, in the form of a skeleton user program. Following that are detailed descriptions of the interface to the various user-callable functions and of the user-supplied functions that were not already described in Chapter 4.

6.1 A skeleton of the user's main program

The following is a skeleton of the user's main program as an application of IDAS. The user program is to have these steps in the order indicated, unless otherwise noted. For the sake of brevity, we defer many of the details to the later sections. As in §4.4, most steps are independent of the NVECTOR implementation used; where this is not the case, usage specifications are given for the two implementations provided with IDAS: steps marked [P] correspond to NVECTOR_PARALLEL, steps marked [O] correspond to NVECTOR_OPENMP, steps +marked [T] correspond to NVECTOR_PTHREADS, while steps marked [S] correspond to NVECTOR_SERIAL. Steps that are unchanged from the skeleton programs presented in §4.4, §5.1, and §5.4, are grayed out.

1. Include necessary header files

The `idas.h` header file also defines additional types, constants, and function prototypes for the adjoint sensitivity module user-callable functions. In addition, the main program should include an NVECTOR implementation header file (`nvector_serial.h`, `nvector_openmp.h`, `nvector_pthreads.h`, or `nvector_parallel.h` for the implementations provided with IDAS) and, if Newton iteration was selected, the main header file of the desired linear solver module.

2. [P] Initialize MPI

Forward problem

3. Set problem dimensions for the forward problem

4. Set initial conditions for the forward problem
5. Create IDAS object for the forward problem
6. Allocate internal memory for the forward problem
7. Specify integration tolerances for forward problem
8. Set optional inputs for the forward problem
9. Attach linear solver module for the forward problem
10. Set linear solver optional inputs for the forward problem
11. Initialize quadrature problem or problems for forward problems, using IDAQuadInit and/or IDAQuadSensInit.
12. Initialize forward sensitivity problem
13. **Allocate space for the adjoint computation**

Call IDAAAdjInit() to allocate memory for the combined forward-backward problem (see §6.2.1 for details). This call requires Nd, the number of steps between two consecutive checkpoints. IDAAAdjInit also specifies the type of interpolation used (see §2.6.3).

14. **Integrate forward problem**

Call IDASolveF, a wrapper for the IDAS main integration function IDASolve, either in IDA.NORMAL mode to the time tout or in IDA.ONE_STEP mode inside a loop (if intermediate solutions of the forward problem are desired (see §6.2.3)). The final value of tret is then the maximum allowable value for the endpoint T of the backward problem.

Backward problem(s)

15. **Set problem dimensions for the backward problem**

[S], [O], [T] set NB, the number of variables in the backward problem
 [P] set NB and NBlock
 [O], [T] set num_threads

16. **Set initial values for the backward problem**

Set the endpoint time tB0 = T and the corresponding vectors yB0 and ypB0 at which the backward problem starts.

17. **Create the backward problem**

Call IDACreateB, a wrapper for IDACreate, to create the IDAS memory block for the new backward problem. Unlike IDACreate, the function IDACreateB does not return a pointer to the newly created memory block (see §6.2.4). Instead, this pointer is attached to the internal adjoint memory block (created by IDAAAdjInit) and returns an identifier called which that the user must later specify in any actions on the newly created backward problem.

18. **Allocate memory for the backward problem**

Call IDAInitB (or IDAInitBS, when the backward problem depends on the forward sensitivities). The two functions are actually wrappers for IDAInit and allocate internal memory, specify problem data, and initialize IDAS at tB0 for the backward problem (see §6.2.4).

19. **Specify integration tolerances for backward problem**

Call `IDASStolerancesB(...)` or `IDASVtolerancesB(...)` to specify a scalar relative tolerance and scalar absolute tolerance, or a scalar relative tolerance and a vector of absolute tolerances, respectively. The functions are wrappers for `IDASStolerances(...)` and `IDASVtolerances(...)` but they require an extra argument `which`, the identifier of the backward problem returned by `IDACreateB`. See §6.2.5 for more information.

20. Set optional inputs for the backward problem

Call `IDASet*B` functions to change from their default values any optional inputs that control the behavior of IDAS. Unlike their counterparts for the forward problem, these functions take an extra argument `which`, the identifier of the backward problem returned by `IDACreateB` (see §6.2.9).

21. Attach linear solver module for the backward problem

Initialize the linear solver module for the backward problem by calling the appropriate wrapper function: `IDADenseB`, `IDABandB`, `IDALapackDenseB`, `IDALapackBandB`, `IDAKLUB`, `IDASuperLUMTB`, `IDASpgmrB`, `IDASpbcgB`, or `IDASptfqmrB` (see §6.2.6). Note that it is not required to use the same linear solver module for both the forward and the backward problems; for example, the forward problem could be solved with the `IDADENSE` linear solver and the backward problem with `IDASPGMR`.

22. Initialize quadrature calculation

If additional quadrature equations must be evaluated, call `IDAQuadInitB` or `IDAQuadInitBS` (if quadrature depends also on the forward sensitivities) as shown in §6.2.11.1. These functions are wrappers around `IDAQuadInit` and can be used to initialize and allocate memory for quadrature integration. Optionally, call `IDASetQuad*B` functions to change from their default values optional inputs that control the integration of quadratures during the backward phase.

23. Integrate backward problem

Call `IDASolveB`, a second wrapper around the IDAS main integration function `IDASolve`, to integrate the backward problem from `tB0` (see §6.2.8). This function can be called either in `IDA.NORMAL` or `IDA.ONE_STEP` mode. Typically, `IDASolveB` will be called in `IDA.NORMAL` mode with an end time equal to the initial time t_0 of the forward problem.

24. Extract quadrature variables

If applicable, call `IDAGetQuadB`, a wrapper around `IDAGetQuad`, to extract the values of the quadrature variables at the time returned by the last call to `IDASolveB`. See §6.2.11.2.

25. Deallocate memory

Upon completion of the backward integration, call all necessary deallocation functions. These include appropriate destructors for the vectors `y` and `yB`, a call to `IDAFree` to free the IDAS memory block for the forward problem. If additional forward integration(s) are to be done for this problem, a call to `IDAAdjFree` (see §6.2.1) may be made to free and deallocate the memory allocated for the backward problems.

26. Finalize MPI

[P] If MPI was initialized by the user main program, call `MPI_Finalize();`.

The above user interface to the adjoint sensitivity module in IDAS was motivated by the desire to keep it as close as possible in look and feel to the one for DAE IVP integration. Note that if steps (15)-(24) are not present, a program with the above structure will have the same functionality as one described in §4.4 for integration of DAEs, albeit with some overhead due to the checkpointing scheme.

If there are multiple backward problems associated with the same forward problem, repeat steps (15)-(24) above for each successive backward problem. In the process, each call to `IDACreateB` creates a new value of the identifier `which`.

6.2 User-callable functions for adjoint sensitivity analysis

6.2.1 Adjoint sensitivity allocation and deallocation functions

After the setup phase for the forward problem, but before the call to `IDASolveF`, memory for the combined forward-backward problem must be allocated by a call to the function `IDAAdjInit`. The form of the call to this function is

IDAAdjInit

Call	<code>flag = IDAAdjInit(ida_mem, Nd, interpType);</code>
Description	The function <code>IDAAdjInit</code> updates IDAS memory block by allocating the internal memory needed for backward integration. Space is allocated for the $N_d = N_d$ interpolation data points, and a linked list of checkpoints is initialized.
Arguments	<p><code>ida_mem</code> (void *) is the pointer to the IDAS memory block returned by a previous call to <code>IDACreate</code>.</p> <p><code>Nd</code> (long int) is the number of integration steps between two consecutive checkpoints.</p> <p><code>interpType</code> (int) specifies the type of interpolation used and can be <code>IDA_POLYNOMIAL</code> or <code>IDA_HERMITE</code>, indicating variable-degree polynomial and cubic Hermite interpolation, respectively (see §2.6.3).</p>
Return value	<p>The return value <code>flag</code> (of type <code>int</code>) is one of:</p> <p><code>IDA_SUCCESS</code> <code>IDAAdjInit</code> was successful.</p> <p><code>IDA_MEM_FAIL</code> A memory allocation request has failed.</p> <p><code>IDA_MEM_NULL</code> <code>ida_mem</code> was NULL.</p> <p><code>IDA_ILL_INPUT</code> One of the parameters was invalid: <code>Nd</code> was not positive or <code>interpType</code> is not one of the <code>IDA_POLYNOMIAL</code> or <code>IDA_HERMITE</code>.</p>
Notes	<p>The user must set <code>Nd</code> so that all data needed for interpolation of the forward problem solution between two checkpoints fits in memory. <code>IDAAdjInit</code> attempts to allocate space for $(2N_d+3)$ variables of type <code>N_Vector</code>.</p> <p>If an error occurred, <code>IDAAdjInit</code> also sends a message to the error handler function.</p>

IDAAdjReInit

Call	<code>flag = IDAAdjReInit(ida_mem);</code>
Description	The function <code>IDAAdjReInit</code> reinitializes the IDAS memory block for ASA, assuming that the number of steps between check points and the type of interpolation remain unchanged.
Arguments	<code>ida_mem</code> (void *) is the pointer to the IDAS memory block returned by a previous call to <code>IDACreate</code> .
Return value	<p>The return value <code>flag</code> (of type <code>int</code>) is one of:</p> <p><code>IDA_SUCCESS</code> <code>IDAAdjReInit</code> was successful.</p> <p><code>IDA_MEM_NULL</code> <code>ida_mem</code> was NULL.</p> <p><code>IDA_NO_ADJ</code> The function <code>IDAAdjInit</code> was not previously called.</p>
Notes	<p>The list of check points (and associated memory) is deleted.</p> <p>The list of backward problems is kept. However, new backward problems can be added to this list by calling <code>IDACreateB</code>. If a new list of backward problems is also needed, then free the adjoint memory (by calling <code>IDAAdjFree</code>) and reinitialize ASA with <code>IDAAdjInit</code>.</p> <p>The IDAS memory for the forward and backward problems can be reinitialized separately by calling <code>IDAReInit</code> and <code>IDAReInitB</code>, respectively.</p>

IDAAdjFree

Call	<code>IDAAdjFree(ida_mem);</code>
Description	The function <code>IDAAdjFree</code> frees the memory related to backward integration allocated by a previous call to <code>IDAAdjInit</code> .
Arguments	The only argument is the IDAS memory block pointer returned by a previous call to <code>IDACreate</code> .
Return value	The function <code>IDAAdjFree</code> has no return value.
Notes	This function frees all memory allocated by <code>IDAAdjInit</code> . This includes workspace memory, the linked list of checkpoints, memory for the interpolation data, as well as the IDAS memory for the backward integration phase. In general, <code>IDAAdjFree</code> need not be called by the user as it is invoked automatically by <code>IDAFree</code> .

6.2.2 Adjoint sensitivity optional input

At any time during the integration of the forward problem, the user can disable the checkpointing of the forward sensitivities by calling the following function:

IDAAdjSetNoSensi

Call	<code>flag = IDAAdjSetNoSensi(ida_mem);</code>
Description	The function <code>IDAAdjSetNoSensi</code> instructs <code>IDASolveF</code> not to save checkpointing data for forward sensitivities any more.
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block.
Return value	The return <code>flag</code> (of type <code>int</code>) is one of: <code>IDA_SUCCESS</code> The call to <code>IDACreateB</code> was successful. <code>IDA_MEM_NULL</code> The <code>ida_mem</code> was <code>NULL</code> . <code>IDA_NO_ADJ</code> The function <code>IDAAdjInit</code> has not been previously called.

6.2.3 Forward integration function

The function `IDASolveF` is very similar to the IDAS function `IDASolve` (see §4.5.6) in that it integrates the solution of the forward problem and returns the solution (y, \dot{y}) . At the same time, however, `IDASolveF` stores checkpoint data every `Nd` integration steps. `IDASolveF` can be called repeatedly by the user. The call to this function has the form

IDASolveF

Call	<code>flag = IDASolveF(ida_mem, tout, &tret, yret, ypret, itask, &ncheck);</code>
Description	The function <code>IDASolveF</code> integrates the forward problem over an interval in t and saves checkpointing data.
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block. <code>tout</code> (<code>realtype</code>) the next time at which a computed solution is desired. <code>tret</code> (<code>realtype</code>) the time reached by the solver (output). <code>yret</code> (<code>N_Vector</code>) the computed solution vector y . <code>ypret</code> (<code>N_Vector</code>) the computed solution vector \dot{y} . <code>itask</code> (<code>int</code>) a flag indicating the job of the solver for the next step. The <code>IDA_NORMAL</code> task is to have the solver take internal steps until it has reached or just passed the user-specified <code>tout</code> parameter. The solver then interpolates in order to return an approximate value of $y(tout)$ and $\dot{y}(tout)$. The <code>IDA_ONE_STEP</code> option

	tells the solver to take just one internal step and return the solution at the point reached by that step.																								
	<code>ncheck</code> (int) the number of (internal) checkpoints stored so far.																								
Return value	On return, <code>IDASolveF</code> returns vectors <code>yret</code> , <code>ypret</code> and a corresponding independent variable value $t = \text{tret}$, such that <code>yret</code> is the computed value of $y(t)$ and <code>ypret</code> the value of $\dot{y}(t)$. Additionally, it returns in <code>ncheck</code> the number of internal checkpoints saved; the total number of checkpoint intervals is <code>ncheck+1</code> . The return value <code>flag</code> (of type int) will be one of the following. For more details see §4.5.6.																								
	<table> <tr> <td><code>IDA_SUCCESS</code></td><td><code>IDASolveF</code> succeeded.</td></tr> <tr> <td><code>IDA_TSTOP_RETURN</code></td><td><code>IDASolveF</code> succeeded by reaching the optional stopping point.</td></tr> <tr> <td><code>IDA_NO_MALLOC</code></td><td>The function <code>IDAInit</code> has not been previously called.</td></tr> <tr> <td><code>IDA_ILL_INPUT</code></td><td>One of the inputs to <code>IDASolveF</code> is illegal.</td></tr> <tr> <td><code>IDA_TOO_MUCH_WORK</code></td><td>The solver took <code>mxstep</code> internal steps but could not reach <code>tout</code>.</td></tr> <tr> <td><code>IDA_TOO_MUCH_ACC</code></td><td>The solver could not satisfy the accuracy demanded by the user for some internal step.</td></tr> <tr> <td><code>IDA_ERR_FAILURE</code></td><td>Error test failures occurred too many times during one internal time step or occurred with $h = h_{min}$.</td></tr> <tr> <td><code>IDA_CONV_FAILURE</code></td><td>Convergence test failures occurred too many times during one internal time step or occurred with $h = h_{min}$.</td></tr> <tr> <td><code>IDA_LSETUP_FAIL</code></td><td>The linear solver's setup function failed in an unrecoverable manner.</td></tr> <tr> <td><code>IDA_LSOLVE_FAIL</code></td><td>The linear solver's solve function failed in an unrecoverable manner.</td></tr> <tr> <td><code>IDA_NO_ADJ</code></td><td>The function <code>IDAAdjInit</code> has not been previously called.</td></tr> <tr> <td><code>IDA_MEM_FAIL</code></td><td>A memory allocation request has failed (in an attempt to allocate space for a new checkpoint).</td></tr> </table>	<code>IDA_SUCCESS</code>	<code>IDASolveF</code> succeeded.	<code>IDA_TSTOP_RETURN</code>	<code>IDASolveF</code> succeeded by reaching the optional stopping point.	<code>IDA_NO_MALLOC</code>	The function <code>IDAInit</code> has not been previously called.	<code>IDA_ILL_INPUT</code>	One of the inputs to <code>IDASolveF</code> is illegal.	<code>IDA_TOO_MUCH_WORK</code>	The solver took <code>mxstep</code> internal steps but could not reach <code>tout</code> .	<code>IDA_TOO_MUCH_ACC</code>	The solver could not satisfy the accuracy demanded by the user for some internal step.	<code>IDA_ERR_FAILURE</code>	Error test failures occurred too many times during one internal time step or occurred with $ h = h_{min}$.	<code>IDA_CONV_FAILURE</code>	Convergence test failures occurred too many times during one internal time step or occurred with $ h = h_{min}$.	<code>IDA_LSETUP_FAIL</code>	The linear solver's setup function failed in an unrecoverable manner.	<code>IDA_LSOLVE_FAIL</code>	The linear solver's solve function failed in an unrecoverable manner.	<code>IDA_NO_ADJ</code>	The function <code>IDAAdjInit</code> has not been previously called.	<code>IDA_MEM_FAIL</code>	A memory allocation request has failed (in an attempt to allocate space for a new checkpoint).
<code>IDA_SUCCESS</code>	<code>IDASolveF</code> succeeded.																								
<code>IDA_TSTOP_RETURN</code>	<code>IDASolveF</code> succeeded by reaching the optional stopping point.																								
<code>IDA_NO_MALLOC</code>	The function <code>IDAInit</code> has not been previously called.																								
<code>IDA_ILL_INPUT</code>	One of the inputs to <code>IDASolveF</code> is illegal.																								
<code>IDA_TOO_MUCH_WORK</code>	The solver took <code>mxstep</code> internal steps but could not reach <code>tout</code> .																								
<code>IDA_TOO_MUCH_ACC</code>	The solver could not satisfy the accuracy demanded by the user for some internal step.																								
<code>IDA_ERR_FAILURE</code>	Error test failures occurred too many times during one internal time step or occurred with $ h = h_{min}$.																								
<code>IDA_CONV_FAILURE</code>	Convergence test failures occurred too many times during one internal time step or occurred with $ h = h_{min}$.																								
<code>IDA_LSETUP_FAIL</code>	The linear solver's setup function failed in an unrecoverable manner.																								
<code>IDA_LSOLVE_FAIL</code>	The linear solver's solve function failed in an unrecoverable manner.																								
<code>IDA_NO_ADJ</code>	The function <code>IDAAdjInit</code> has not been previously called.																								
<code>IDA_MEM_FAIL</code>	A memory allocation request has failed (in an attempt to allocate space for a new checkpoint).																								
Notes	<p>All failure return values are negative and therefore a test <code>flag < 0</code> will trap all <code>IDASolveF</code> failures.</p> <p>At this time, <code>IDASolveF</code> stores checkpoint information in memory only. Future versions will provide for a safeguard option of dumping checkpoint data into a temporary file as needed. The data stored at each checkpoint is basically a snapshot of the IDAS internal memory block and contains enough information to restart the integration from that time and to proceed with the same step size and method order sequence as during the forward integration.</p> <p>In addition, <code>IDASolveF</code> also stores interpolation data between consecutive checkpoints so that, at the end of this first forward integration phase, interpolation information is already available from the last checkpoint forward. In particular, if no checkpoints were necessary, there is no need for the second forward integration phase.</p> <p>It is illegal to change the integration tolerances between consecutive calls to <code>IDASolveF</code>, as this information is not captured in the checkpoint data.</p>																								



6.2.4 Backward problem initialization functions

The functions `IDACreateB` and `IDAInitB` (or `IDAInitBS`) must be called in the order listed. They instantiate an IDAS solver object, provide problem and solution specifications, and allocate internal memory for the backward problem.

`IDACreateB`

Call `flag = IDACreateB(ida_mem, &which);`

Description The function `IDACreateB` instantiates an IDAS solver object for the backward problem.

Arguments `ida_mem` (void *) pointer to the IDAS memory block returned by `IDACreate`.
`which` (int) contains the identifier assigned by IDAS for the newly created backward problem. Any call to `IDA*B` functions requires such an identifier.

Return value The return `flag` (of type `int`) is one of:

`IDA_SUCCESS` The call to `IDACreateB` was successful.
`IDA_MEM_NULL` The `ida_mem` was `NULL`.
`IDA_NO_ADJ` The function `IDAAdjInit` has not been previously called.
`IDA_MEM_FAIL` A memory allocation request has failed.

There are two initialization functions for the backward problem – one for the case when the backward problem does not depend on the forward sensitivities, and one for the case when it does. These two functions are described next.

The function `IDAInitB` initializes the backward problem when it does not depend on the forward sensitivities. It is essentially wrapper for `IDAInit` with some particularization for backward integration, as described below.

IDAInitB

Call `flag = IDAInitB(ida_mem, which, resB, tB0, yB0, ypB0);`

Description The function `IDAInitB` provides problem specification, allocates internal memory, and initializes the backward problem.

Arguments `ida_mem` (void *) pointer to the IDAS memory block returned by `IDACreate`.
`which` (int) represents the identifier of the backward problem.
`resB` (`IDAResFnB`) is the C function which computes fB , the residual of the backward DAE problem. This function has the form `resB(t, y, yp, yB, ypB, resvalB, user_dataB)` (for full details see §6.3.1).
`tB0` (`realtype`) specifies the endpoint T where final conditions are provided for the backward problem, normally equal to the endpoint of the forward integration.
`yB0` (`N_Vector`) is the initial value (at $t = tB0$) of the backward solution.
`ypB0` (`N_Vector`) is the initial derivative value (at $t = tB0$) of the backward solution.

Return value The return `flag` (of type `int`) will be one of the following:

`IDA_SUCCESS` The call to `IDAInitB` was successful.
`IDA_NO_MALLOC` The function `IDAInit` has not been previously called.
`IDA_MEM_NULL` The `ida_mem` was `NULL`.
`IDA_NO_ADJ` The function `IDAAdjInit` has not been previously called.
`IDA_BAD_TB0` The final time `tB0` was outside the interval over which the forward problem was solved.
`IDA_ILL_INPUT` The parameter `which` represented an invalid identifier, or one of `yB0`, `ypB0`, `resB` was `NULL`.

Notes The memory allocated by `IDAInitB` is deallocated by the function `IDAAdjFree`.

For the case when backward problem also depends on the forward sensitivities, user must call `IDAInitBS` instead of `IDAInitB`. Only the third argument of each function differs between these functions.

IDAInitBS

Call `flag = IDAInitBS(ida_mem, which, resBS, tB0, yB0, ypB0);`

Description The function `IDAInitBS` provides problem specification, allocates internal memory, and initializes the backward problem.

Arguments `ida_mem` (void *) pointer to the IDAS memory block returned by `IDACreate`.

which (int) represents the identifier of the backward problem.
resBS (IDAResFnBS) is the C function which computes fB , the residual or the backward DAE problem. This function has the form **resBS**(t, y, yp, yS, ypS, yB, ypB, resvalB, user_dataB) (for full details see §6.3.2).
tB0 (realtype) specifies the endpoint T where final conditions are provided for the backward problem.
yB0 (N.Vector) is the initial value (at $t = tB0$) of the backward solution.
ypB0 (N.Vector) is the initial derivative value (at $t = tB0$) of the backward solution.

Return value The return flag (of type int) will be one of the following:

IDA_SUCCESS The call to **IDAInitB** was successful.
IDA_NO_MALLOC The function **IDAInit** has not been previously called.
IDA_MEM_NULL The **ida_mem** was NULL.
IDA_NO_ADJ The function **IDAAdjInit** has not been previously called.
IDA_BAD_TB0 The final time **tB0** was outside the interval over which the forward problem was solved.
IDA_ILL_INPUT The parameter **which** represented an invalid identifier, or one of **yB0**, **ypB0**, **resB** was NULL, or sensitivities were not active during the forward integration.

Notes The memory allocated by **IDAInitBS** is deallocated by the function **IDAAdjFree**.

The function **IDAReInitB** reinitializes IDAS for the solution of a series of backward problems, each identified by a value of the parameter **which**. **IDAReInitB** is essentially a wrapper for **IDAReInit**, and so all details given for **IDAReInit** in §4.5.10 apply. Also, **IDAReInitB** can be called to reinitialize a backward problem even if it has been initialized with the sensitivity-dependent version **IDAInitBS**. The call to the **IDAReInitB** function has the form

IDAReInitB

Call **flag** = **IDAReInitB**(ida_mem, which, tB0, yB0, ypB0)

Description The function **IDAReInitB** reinitializes IDAS the backward problem.

Arguments **ida_mem** (void *) pointer to IDAS memory block returned by **IDACreate**.

which (int) represents the identifier of the backward problem.
tB0 (realtype) specifies the endpoint T where final conditions are provided for the backward problem.
yB0 (N.Vector) is the initial value (at $t = tB0$) of the backward solution.
ypB0 (N.Vector) is the initial derivative value (at $t = tB0$) of the backward solution.

Return value The return value **flag** (of type int) will be one of the following:

IDA_SUCCESS The call to **IDAReInitB** was successful.
IDA_NO_MALLOC The function **IDAInit** has not been previously called.
IDA_MEM_NULL The **ida_mem** memory block pointer was NULL.
IDA_NO_ADJ The function **IDAAdjInit** has not been previously called.
IDA_BAD_TB0 The final time **tB0** is outside the interval over which the forward problem was solved.
IDA_ILL_INPUT The parameter **which** represented an invalid identifier, or one of **yB0**, **ypB0** was NULL.

6.2.5 Tolerance specification functions for backward problem

One of the following two functions must be called to specify the integration tolerances for the backward problem. Note that this call must be made after the call to **IDAInitB** or **IDAInitBS**.

IDASStolerancesB

Call `flag = IDASStolerancesB(ida_mem, which, reltolB, abstolB);`

Description The function `IDASStolerancesB` specifies scalar relative and absolute tolerances.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block returned by `IDACreate`.
`which` (`int`) represents the identifier of the backward problem.
`reltolB` (`realtype`) is the scalar relative error tolerance.
`abstolB` (`realtype`) is the scalar absolute error tolerance.

Return value The return `flag` (of type `int`) will be one of the following:

- `IDA_SUCCESS` The call to `IDASStolerancesB` was successful.
- `IDA_MEM_NULL` The IDAS memory block was not initialized through a previous call to `IDACreate`.
- `IDA_NO_MALLOC` The allocation function `IDAInit` has not been called.
- `IDA_NO_ADJ` The function `IDAAdjInit` has not been previously called.
- `IDA_ILL_INPUT` One of the input tolerances was negative.

IDASVtolerancesB

Call `flag = IDASVtolerancesB(ida_mem, which, reltolB, abstolB);`

Description The function `IDASVtolerancesB` specifies scalar relative tolerance and vector absolute tolerances.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block returned by `IDACreate`.
`which` (`int`) represents the identifier of the backward problem.
`reltol` (`realtype`) is the scalar relative error tolerance.
`abstol` (`N_Vector`) is the vector of absolute error tolerances.

Return value The return `flag` (of type `int`) will be one of the following:

- `IDA_SUCCESS` The call to `IDASVtolerancesB` was successful.
- `IDA_MEM_NULL` The IDAS memory block was not initialized through a previous call to `IDACreate`.
- `IDA_NO_MALLOC` The allocation function `IDAInit` has not been called.
- `IDA_NO_ADJ` The function `IDAAdjInit` has not been previously called.
- `IDA_ILL_INPUT` The relative error tolerance was negative or the absolute tolerance had a negative component.

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

6.2.6 Linear solver initialization functions for backward problem

All IDAS linear solver modules available for forward problems provide additional specification functions for backward problems. The initialization functions described in §4.5.3 cannot be directly used since the optional user-defined Jacobian-related functions have different prototypes for the backward problem than for the forward problem (see §6.3).

The following wrapper functions can be used to initialize one of the linear solver modules for the backward problem. Their arguments are identical to those of the functions in §4.5.3 with the exception of the additional second argument, `which`, the identifier of the backward problem.

```
flag = IDADenseB(ida_mem, which, nB);
flag = IDABandB(ida_mem, which, nB, mupperB, mlowerB);
flag = IDALapackDenseB(ida_mem, which, nB);
flag = IDALapackBandB(ida_mem, which, nB, mupperB, mlowerB);
```


Description	The function <code>IDACalcICBS</code> corrects the initial values <code>yB0</code> and <code>ypB0</code> at time <code>tB0</code> for the backward problem.
Arguments	<p><code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block.</p> <p><code>which</code> (<code>int</code>) is the identifier of the backward problem.</p> <p><code>tBout1</code> (<code>realtype</code>) is the first value of t at which a solution will be requested (from <code>IDASolveB</code>). This value is needed here only to determine the direction of integration and rough scale in the independent variable t.</p> <p><code>yfin</code> (<code>N.Vector</code>) the forward solution at the final time <code>tB0</code>.</p> <p><code>ypfin</code> (<code>N.Vector</code>) the forward solution derivative at the final time <code>tB0</code>.</p> <p><code>ySfin</code> (<code>N.Vector *</code>) a pointer to an array of <code>Ns</code> vectors containing the sensitivities of the forward solution at the final time <code>tB0</code>.</p> <p><code>ypSfin</code> (<code>N.Vector *</code>) a pointer to an array of <code>Ns</code> vectors containing the derivatives of the forward solution sensitivities at the final time <code>tB0</code>.</p>
Return value	<p>The return value <code>flag</code> (of type <code>int</code>) can be any that is returned by <code>IDACalcIC</code> (see §4.5.4). However <code>IDACalcICBS</code> can also return one of the following:</p> <p><code>IDA_NO_ADJ</code> <code>IDAAdjInit</code> has not been previously called.</p> <p><code>IDA_ILL_INPUT</code> Parameter <code>which</code> represented an invalid identifier, sensitivities were not active during forward integration, or <code>IDAInitBS</code> (or <code>IDAReInitBS</code>) has not been previously called.</p>
Notes	<p>All failure return values are negative and therefore a test <code>flag < 0</code> will trap all <code>IDACalcICBS</code> failures.</p> <p>Note that <code>IDACalcICBS</code> will correct the values of $yB(tB_0)$ and $\dot{y}B(tB_0)$ which were specified in the previous call to <code>IDAInitBS</code> or <code>IDAReInitBS</code>. To obtain the corrected values, call <code>IDAGetConsistentICB</code> (see §6.2.10.2).</p>

6.2.8 Backward integration function

The function `IDASolveB` performs the integration of the backward problem. It is essentially a wrapper for the IDAS main integration function `IDASolve` and, in the case in which checkpoints were needed, it evolves the solution of the backward problem through a sequence of forward-backward integration pairs between consecutive checkpoints. In each pair, the first run integrates the original IVP forward in time and stores interpolation data; the second run integrates the backward problem backward in time and performs the required interpolation to provide the solution of the IVP to the backward problem.

The function `IDASolveB` does not return the solution `yB` itself. To obtain that, call the function `IDAGetB`, which is also described below.

The call to `IDASolveB` has the form

`IDASolveB`

Call	<code>flag = IDASolveB(ida_mem, tBout, itaskB);</code>
Description	The function <code>IDASolveB</code> integrates the backward DAE problem.
Arguments	<p><code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory returned by <code>IDACreate</code>.</p> <p><code>tBout</code> (<code>realtype</code>) the next time at which a computed solution is desired.</p> <p><code>itaskB</code> (<code>int</code>) a flag indicating the job of the solver for the next step. The <code>IDA_NORMAL</code> task is to have the solver take internal steps until it has reached or just passed the user-specified value <code>tBout</code>. The solver then interpolates in order to return an approximate value of $yB(tBout)$. The <code>IDA_ONE_STEP</code> option tells the solver to take just one internal step in the direction of <code>tBout</code> and return.</p>
Return value	The return value <code>flag</code> (of type <code>int</code>) will be one of the following. For more details see §4.5.6.

IDA_SUCCESS	IDASolveB succeeded.
IDA_MEM_NULL	The <code>ida_mem</code> was NULL.
IDA_NO_ADJ	The function IDAAAdjInit has not been previously called.
IDA_NO_BCK	No backward problem has been added to the list of backward problems by a call to IDACreateB
IDA_NO_FWD	The function IDASolveF has not been previously called.
IDA_ILL_INPUT	One of the inputs to IDASolveB is illegal.
IDA_BAD_ITASK	The <code>itaskB</code> argument has an illegal value.
IDA_TOO_MUCH_WORK	The solver took <code>mxstep</code> internal steps but could not reach <code>tBout</code> .
IDA_TOO_MUCH_ACC	The solver could not satisfy the accuracy demanded by the user for some internal step.
IDA_ERR_FAILURE	Error test failures occurred too many times during one internal time step.
IDA_CONV_FAILURE	Convergence test failures occurred too many times during one internal time step.
IDA_LSETUP_FAIL	The linear solver's setup function failed in an unrecoverable manner.
IDA_SOLVE_FAIL	The linear solver's solve function failed in an unrecoverable manner.
IDA_BCKMEM_NULL	The IDAS memory for the backward problem was not created with a call to IDACreateB.
IDA_BAD_TBOUT	The desired output time <code>tBout</code> is outside the interval over which the forward problem was solved.
IDA_REIFWD_FAIL	Reinitialization of the forward problem failed at the first checkpoint (corresponding to the initial time of the forward problem).
IDA_FWD_FAIL	An error occurred during the integration of the forward problem.

Notes All failure return values are negative and therefore a test `flag < 0` will trap all IDASolveB failures.

In the case of multiple checkpoints and multiple backward problems, a given call to IDASolveB in IDA_ONE_STEP mode may not advance every problem one step, depending on the relative locations of the current times reached. But repeated calls will eventually advance all problems to `tBout`.

To obtain the solution `yB` to the backward problem, call the function IDAGetB as follows:

IDAGetB

Call	<code>flag = IDAGetB(ida_mem, which, &tret, yB, ypB);</code>
Description	The function IDAGetB provides the solution <code>yB</code> of the backward DAE problem.
Arguments	<code>ida_mem</code> (void *) pointer to the IDAS memory returned by IDACreate. <code>which</code> (int) the identifier of the backward problem. <code>tret</code> (realtype) the time reached by the solver (output). <code>yB</code> (N_Vector) the backward solution at time <code>tret</code> . <code>ypB</code> (N_Vector) the backward solution derivative at time <code>tret</code> .
Return value	The return value <code>flag</code> (of type int) will be one of the following.
	IDA_SUCCESS IDAGetB was successful.
	IDA_MEM_NULL <code>ida_mem</code> is NULL.
	IDA_NO_ADJ The function IDAAAdjInit has not been previously called.
	IDA_ILL_INPUT The parameter <code>which</code> is an invalid identifier.
Notes	The user must allocate space for <code>yB</code> and <code>ypB</code> .



6.2.9 Optional input functions for the backward problem

6.2.9.1 Main solver optional input functions

The adjoint module in IDAS provides wrappers for most of the optional input functions defined in §4.5.7.1. The only difference is that the user must specify the identifier **which** of the backward problem within the list managed by IDAS.

The optional input functions defined for the backward problem are:

```
flag = IDASetUserDataB(ida_mem, which, user_dataB);
flag = IDASetMaxOrdB(ida_mem, which, maxordB);
flag = IDASetMaxNumStepsB(ida_mem, which, mxstepsB);
flag = IDASetInitStepB(ida_mem, which, hinB);
flag = IDASetMaxStepB(ida_mem, which, hmaxB);
flag = IDASetSuppressAlgB(ida_mem, which, suppressalgB);
flag = IDASetIdB(ida_mem, which, idB);
flag = IDASetConstraintsB(ida_mem, which, constraintsB);
```

Their return value **flag** (of type **int**) can have any of the return values of their counterparts, but it can also be **IDA_NO_ADJ** if **IDAAdjInit** has not been called, or **IDA_ILL_INPUT** if **which** was an invalid identifier.

6.2.9.2 Dense linear solver

Optional inputs for the **IDADENSE** linear solver module can be set for the backward problem through the following two functions:

IDADlsSetDenseJacFnB

Call	<code>flag = IDADlsSetDenseJacFnB(ida_mem, which, jacB);</code>
Description	The function IDADlsSetDenseJacFnB specifies the dense Jacobian approximation function to be used for the backward problem.
Arguments	ida_mem (void *) pointer to the IDAS memory returned by IDACreate . which (int) represents the identifier of the backward problem. jacB (IDADlsDenseJacFnB) user-defined dense Jacobian approximation function.
Return value	The return value flag (of type int) is one of: IDADLS_SUCCESS IDADlsSetDenseJacFnB succeeded. IDADLS_MEM_NULL The ida_mem was NULL . IDADLS_NO_ADJ The function IDAAdjInit has not been previously called. IDADLS_LMEM_NULL The linear solver has not been initialized with a call to IDADenseB or IDALapackDenseB . IDADLS_ILL_INPUT The parameter which represented an invalid identifier.
Notes	The function type IDADlsDenseJacFnB is described in §6.3.5.

IDADlsSetDenseJacFnBS

Call	<code>flag = IDADlsSetDenseJacFnBS(ida_mem, which, jacBS);</code>
Description	The function IDADlsSetDenseJacFnBS specifies the dense Jacobian approximation function to be used for the backward problem, in the case where the backward problem depends on the forward sensitivities.
Arguments	ida_mem (void *) pointer to the IDAS memory returned by IDACreate . which (int) represents the identifier of the backward problem. jacBS (IDADlsDenseJacFnBS) user-defined dense Jacobian approximation function.

Return value The return value `flag` (of type `int`) is one of:

`IDADLS_SUCCESS` `IDADlsSetDenseJacFnBS` succeeded.
`IDADLS_MEM_NULL` The `ida_mem` was `NULL`.
`IDADLS_NO_ADJ` The function `IDAAdjInit` has not been previously called.
`IDADLS_LMEM_NULL` The linear solver has not been initialized with a call to `IDADenseB` or `IDALapackDenseB`.
`IDADLS_ILL_INPUT` The parameter `which` represented an invalid identifier.

Notes The function type `IDADlsDenseJacFnBS` is described in §6.3.5.

6.2.9.3 Band linear solver

Optional inputs for the `IDABAND` linear solver module can be set for the backward problem through the following two functions:

`IDADlsSetBandJacFnB`

Call `flag = IDADlsSetBandJacFnB(ida_mem, which, jacB);`

Description The function `IDADlsSetBandJacFnB` specifies the banded Jacobian approximation function to be used for the backward problem.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory returned by `IDACreate`.
`which` (`int`) represents the identifier of the backward problem.
`jacB` (`IDADlsBandJacFnB`) user-defined banded Jacobian approximation function.

Return value The return value `flag` (of type `int`) is one of:

`IDADLS_SUCCESS` `IDADlsSetBandJacFnB` succeeded.
`IDADLS_MEM_NULL` The `ida_mem` was `NULL`.
`IDADLS_NO_ADJ` The function `IDAAdjInit` has not been previously called.
`IDADLS_LMEM_NULL` The linear solver has not been initialized with a call to `IDABandB` or `IDALapackBandB`.
`IDADLS_ILL_INPUT` The parameter `which` represented an invalid identifier.

Notes The function type `IDADlsBandJacFnB` is described in §6.3.6.

`IDADlsSetBandJacFnBS`

Call `flag = IDADlsSetBandJacFnBS(ida_mem, which, jacBS);`

Description The function `IDADlsSetBandJacFnBS` specifies the banded Jacobian approximation function to be used for the backward problem, in the case where the backward problem depends on the forward sensitivities.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory returned by `IDACreate`.
`which` (`int`) represents the identifier of the backward problem.
`jacBS` (`IDADlsBandJacFnBS`) user-defined banded Jacobian approximation function.

Return value The return value `flag` (of type `int`) is one of:

`IDADLS_SUCCESS` `IDADlsSetBandJacFnBS` succeeded.
`IDADLS_MEM_NULL` The `ida_mem` was `NULL`.
`IDADLS_NO_ADJ` The function `IDAAdjInit` has not been previously called.
`IDADLS_LMEM_NULL` The linear solver has not been initialized with a call to `IDABandB` or `IDALapackBandB`.
`IDADLS_ILL_INPUT` The parameter `which` represented an invalid identifier.

Notes The function type `IDADlsBandJacFnBS` is described in §6.3.6.

6.2.9.4 Sparse linear solvers

Optional inputs for the IDAKLU and IDASUPERLUMT linear solver modules can be set for the backward problem through the following two functions:

IDASlsSetSparseJacFnB

Call	<code>flag = IDASlsSetSparseJacFnB(ida_mem, which, jacB);</code>
Description	The function <code>IDASlsSetSparseJacFnB</code> specifies the sparse Jacobian approximation function to be used for the backward problem.
Arguments	<p><code>ida_mem</code> (void *) pointer to the IDAS memory returned by <code>IDACreate</code>.</p> <p><code>which</code> (int) represents the identifier of the backward problem.</p> <p><code>jacB</code> (<code>IDASlsSparseJacFnB</code>) user-defined sparse Jacobian approximation function.</p>
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of: <ul style="list-style-type: none"> <code>IDASLS_SUCCESS</code> <code>IDASlsSetSparseJacFnB</code> succeeded. <code>IDASLS_MEM_NULL</code> The <code>ida_mem</code> was <code>NULL</code>. <code>IDASLS_NO_ADJ</code> The function <code>IDAAAdjInit</code> has not been previously called. <code>IDASLS_LMEM_NULL</code> The linear solver has not been initialized with a call to <code>IDAKLUB</code> or <code>IDASuperLUMTB</code>. <code>IDASLS_ILL_INPUT</code> The parameter <code>which</code> represented an invalid identifier.
Notes	The function type <code>IDASlsSparseJacFnB</code> is described in §6.3.7.

IDASlsSetSparseJacFnBS

Call	<code>flag = IDASlsSetSparseJacFnBS(ida_mem, which, jacBS);</code>
Description	The function <code>IDASlsSetSparseJacFnBS</code> specifies the sparse Jacobian approximation function to be used for the backward problem, in the case where the backward problem depends on the forward sensitivities.
Arguments	<p><code>ida_mem</code> (void *) pointer to the IDAS memory returned by <code>IDACreate</code>.</p> <p><code>which</code> (int) represents the identifier of the backward problem.</p> <p><code>jacBS</code> (<code>IDASlsSparseJacFnBS</code>) user-defined sparse Jacobian approximation function.</p>
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of: <ul style="list-style-type: none"> <code>IDASLS_SUCCESS</code> <code>IDASlsSetSparseJacFnBS</code> succeeded. <code>IDASLS_MEM_NULL</code> The <code>ida_mem</code> was <code>NULL</code>. <code>IDASLS_NO_ADJ</code> The function <code>IDAAAdjInit</code> has not been previously called. <code>IDASLS_LMEM_NULL</code> The linear solver has not been initialized with a call to <code>IDAKLUB</code> or <code>IDASuperLUMTB</code>. <code>IDASLS_ILL_INPUT</code> The parameter <code>which</code> represented an invalid identifier.
Notes	The function type <code>IDASlsSparseJacFnBS</code> is described in §6.3.7.

6.2.9.5 SPILS linear solvers

Optional inputs for the IDASPILS linear solver module can be set for the backward problem through the following functions:

IDASpilsSetPreconditionerB

Call `flag = IDASpilsSetPreconditionerB(ida_mem, which, psetupB, psolveB);`

Description The function `IDASpilsSetPrecSolveFnB` specifies the preconditioner setup and solve functions for the backward integration.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.
`which` (int) the identifier of the backward problem.
`psetupB` (`IDASpilsPrecSetupFnB`) user-defined preconditioner setup function.
`psolveB` (`IDASpilsPrecSolveFnB`) user-defined preconditioner solve function.

Return value The return value `flag` (of type `int`) is one of:

`IDASPILS_SUCCESS` The optional value has been successfully set.
`IDASPILS_MEM_NULL` The `ida_mem` memory block pointer was `NULL`.
`IDASPILS_LMEM_NULL` The `IDASPILS` linear solver has not been initialized.
`IDASPILS_NO_ADJ` The function `IDAAAdjInit` has not been previously called.
`IDASPILS_ILL_INPUT` The parameter `which` represented an invalid identifier.

Notes The function types `IDASpilsPrecSolveFnB` and `IDASpilsPrecSetupFnB` are described in §6.3.9 and §6.3.10, respectively. The `psetupB` argument may be `NULL` if no setup operation is involved in the preconditioner.

IDASpilsSetPreconditionerBS

Call `flag = IDASpilsSetPreconditionerBS(ida_mem, which, psetupBS, psolveBS);`

Description The function `IDASpilsSetPrecSolveFnBS` specifies the preconditioner setup and solve functions for the backward integration, in the case where the backward problem depends on the forward sensitivities.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.
`which` (int) the identifier of the backward problem.
`psetupBS` (`IDASpilsPrecSetupFnBS`) user-defined preconditioner setup function.
`psolveBS` (`IDASpilsPrecSolveFnBS`) user-defined preconditioner solve function.

Return value The return value `flag` (of type `int`) is one of:

`IDASPILS_SUCCESS` The optional value has been successfully set.
`IDASPILS_MEM_NULL` The `ida_mem` memory block pointer was `NULL`.
`IDASPILS_LMEM_NULL` The `IDASPILS` linear solver has not been initialized.
`IDASPILS_NO_ADJ` The function `IDAAAdjInit` has not been previously called.
`IDASPILS_ILL_INPUT` The parameter `which` represented an invalid identifier.

Notes The function types `IDASpilsPrecSolveFnBS` and `IDASpilsPrecSetupFnBS` are described in §6.3.9 and §6.3.10, respectively. The `psetupBS` argument may be `NULL` if no setup operation is involved in the preconditioner.

IDASpilsSetJacTimesVecFnB

Call `flag = IDASpilsSetJacTimesVecFnB(ida_mem, which, jtvB);`

Description The function `IDASpilsSetJacTimesVecFnB` specifies the Jacobian-vector product function to be used.

Arguments `ida_mem` (void *) pointer to the IDAS memory block.
`which` (int) the identifier of the backward problem.
`jtvB` (`IDASpilsJacTimesVecFnB`) user-defined Jacobian-vector product function.

Return value The return value `flag` (of type `int`) is one of:

IDASPILS_SUCCESS The optional value has been successfully set.
 IDASPILS_MEM_NULL The `ida_mem` memory block pointer was NULL.
 IDASPILS_LMEM_NULL The IDASPILS linear solver has not been initialized.
 IDASPILS_NO_ADJ The function `IDAAAdjInit` has not been previously called.
 IDASPILS_ILL_INPUT The parameter `which` represented an invalid identifier.

Notes The function type `IDASpilsJacTimesVecFnB` is described in §6.3.8.

IDASpilsSetJacTimesVecFnBS


Call `flag = IDASpilsSetJacTimesVecFnBS(ida_mem, which, jtvBS);`
Description The function `IDASpilsSetJacTimesVecFnBS` specifies the Jacobian-vector product function to be used, in the case where the backward problem depends on the forward sensitivities.
Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`which` (`int`) the identifier of the backward problem.
`jtvBS` (`IDASpilsJacTimesVecFnBS`) user-defined Jacobian-vector product function.
Return value The return value `flag` (of type `int`) is one of:
 IDASPILS_SUCCESS The optional value has been successfully set.
 IDASPILS_MEM_NULL The `ida_mem` memory block pointer was NULL.
 IDASPILS_LMEM_NULL The IDASPILS linear solver has not been initialized.
 IDASPILS_NO_ADJ The function `IDAAAdjInit` has not been previously called.
 IDASPILS_ILL_INPUT The parameter `which` represented an invalid identifier.
Notes The function type `IDASpilsJacTimesVecFnBS` is described in §6.3.8.

IDASpilsSetGSTypeB

Call `flag = IDASpilsSetGSTypeB(ida_mem, which, gstypeB);`
Description The function `IDASpilsSetGSTypeB` specifies the type of Gram-Schmidt orthogonalization to be used with IDASPGMR. This must be one of the enumeration constants `MODIFIED_GS` or `CLASSICAL_GS`. These correspond to using modified Gram-Schmidt and classical Gram-Schmidt, respectively.
Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`which` (`int`) the identifier of the backward problem.
`gstypeB` (`int`) type of Gram-Schmidt orthogonalization.
Return value The return value `flag` (of type `int`) is one of:
 IDASPILS_SUCCESS The optional value has been successfully set.
 IDASPILS_MEM_NULL `ida_mem` was NULL.
 IDASPILS_LMEM_NULL The IDASPILS linear solver has not been initialized.
 IDASPILS_NO_ADJ The function `IDAAAdjInit` has not been previously called.
 IDASPILS_ILL_INPUT The parameter `which` represented an invalid identifier or the value of `gstypeB` was not valid.
Notes The default value is `MODIFIED_GS`.
 This option is available only with IDASPGMR.



IDASpilsSetMaxlB

Call	<code>flag = IDASpilsSetMaxlB(ida_mem, which, maxlB);</code>
Description	The function <code>IDASpilsSetMaxlB</code> resets maximum Krylov subspace dimension for the Bi-CGStab or TFQMR methods.
Arguments	<p><code>ida_mem</code> (void *) pointer to the IDAS memory block.</p> <p><code>which</code> (int) the identifier of the backward problem.</p> <p><code>maxlB</code> (realtype) maximum dimension of the Krylov subspace.</p>
Return value	<p>The return value <code>flag</code> (of type <code>int</code>) is one of:</p> <p><code>IDASPILS_SUCCESS</code> The optional value has been successfully set.</p> <p><code>IDASPILS_MEM_NULL</code> <code>ida_mem</code> was NULL.</p> <p><code>IDASPILS_LMEM_NULL</code> The IDASPILS linear solver has not been initialized.</p> <p><code>IDASPILS_NO_ADJ</code> The function <code>IDAAAdjInit</code> has not been previously called.</p> <p><code>IDASPILS_ILL_INPUT</code> The parameter <code>which</code> represented an invalid identifier.</p>
Notes	<p>The maximum subspace dimension is initially specified in the call to <code>IDASpbcgB</code> or <code>IDASptfqmrB</code>. The call to <code>IDASpilsSetMaxlB</code> is needed only if <code>maxlB</code> is being changed from its previous value.</p> <p> This option is available only for the IDASPCBG and IDASPTFQMR linear solvers.</p>

IDASpilsSetEpsLinB

Call	<code>flag = IDASpilsSetEpsLinB(ida_mem, eplifacB);</code>
Description	The function <code>IDASpilsSetEpsLinB</code> specifies the factor by which the Krylov linear solver's convergence test constant is reduced from the Newton iteration test constant. (See §2.1).
Arguments	<p><code>ida_mem</code> (void *) pointer to the IDAS memory block.</p> <p><code>eplifacB</code> (realtype) linear convergence safety factor (≥ 0.0).</p>
Return value	<p>The return value <code>flag</code> (of type <code>int</code>) is one of</p> <p><code>IDASPILS_SUCCESS</code> The optional value has been successfully set.</p> <p><code>IDASPILS_MEM_NULL</code> The <code>ida_mem</code> pointer is NULL.</p> <p><code>IDASPILS_LMEM_NULL</code> The IDASPILS linear solver has not been initialized.</p> <p><code>IDASPILS_NO_ADJ</code> The function <code>IDAAAdjInit</code> has not been previously called.</p> <p><code>IDASPILS_ILL_INPUT</code> The value of <code>eplifacB</code> is negative.</p>
Notes	<p>The default value is 0.05.</p> <p>Passing a value <code>eplifacB= 0.0</code> also indicates using the default value.</p>

6.2.10 Optional output functions for the backward problem

6.2.10.1 Main solver optional output functions

The user of the adjoint module in IDAS has access to any of the optional output functions described in §4.5.9, both for the main solver and for the linear solver modules. The first argument of these `IDAGet*` and `IDA*Get*` functions is the pointer to the IDAS memory block for the backward problem. In order to call any of these functions, the user must first call the following function to obtain this pointer:

IDAGetAdjIDABmem

Call	<code>ida_memB = IDAGetAdjIDABmem(ida_mem, which);</code>
Description	The function <code>IDAGetAdjIDABmem</code> returns a pointer to the IDAS memory block for the backward problem.
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block created by <code>IDACreate</code> . <code>which</code> (<code>int</code>) the identifier of the backward problem.
Return value	The return value, <code>ida_memB</code> (of type <code>void *</code>), is a pointer to the IDAS memory for the backward problem.
Notes	The user should not modify <code>ida_memB</code> in any way. Optional output calls should pass <code>ida_memB</code> as the first argument; thus, for example, to get the number of integration steps: <code>flag = IDAGetNumSteps(ida_memB, &nsteps)</code> .



To get values of the *forward* solution during a backward integration, use the following function. The input value of `t` would typically be equal to that at which the backward solution has just been obtained with `IDAGetB`. In any case, it must be within the last checkpoint interval used by `IDASolveB`.

IDAGetAdjY

Call	<code>flag = IDAGetAdjY(ida_mem, t, y, yp);</code>
Description	The function <code>IDAGetAdjY</code> returns the interpolated value of the forward solution y and its derivative during a backward integration.
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block created by <code>IDACreate</code> . <code>t</code> (<code>realtype</code>) value of the independent variable at which y is desired (input). <code>y</code> (<code>N_Vector</code>) forward solution $y(t)$. <code>yp</code> (<code>N_Vector</code>) forward solution derivative $\dot{y}(t)$.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of: <code>IDA_SUCCESS</code> <code>IDAGetAdjY</code> was successful. <code>IDA_MEM_NULL</code> <code>ida_mem</code> was <code>NULL</code> . <code>IDA_GETY_BADT</code> The value of <code>t</code> was outside the current checkpoint interval.
Notes	The user must allocate space for <code>y</code> and <code>yp</code> .

**6.2.10.2 Initial condition calculation optional output function****IDAGetConsistentICB**

Call	<code>flag = IDAGetConsistentICB(ida_mem, which, yB0_mod, ypB0_mod);</code>
Description	The function <code>IDAGetConsistentICB</code> returns the corrected initial conditions for backward problem calculated by <code>IDACalcICB</code> .
Arguments	<code>ida_mem</code> (<code>void *</code>) pointer to the IDAS memory block. <code>which</code> is the identifier of the backward problem. <code>yB0_mod</code> (<code>N_Vector</code>) consistent initial vector. <code>ypB0_mod</code> (<code>N_Vector</code>) consistent initial derivative vector.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>IDA_SUCCESS</code> The optional output value has been successfully set. <code>IDA_MEM_NULL</code> The <code>ida_mem</code> pointer is <code>NULL</code> . <code>IDA_NO_ADJ</code> <code>IDAAAdjInit</code> has not been previously called. <code>IDA_ILL_INPUT</code> Parameter <code>which</code> did not refer a valid backward problem identifier.

Notes If the consistent solution vector or consistent derivative vector is not desired, pass NULL for the corresponding argument.

The user must allocate space for `yB0_mod` and `ypB0_mod` (if not NULL).



6.2.11 Backward integration of quadrature equations

Not only the backward problem but also the backward quadrature equations may or may not depend on the forward sensitivities. Accordingly, one of the `IDAQuadInitB` or `IDAQuadInitBS` should be used to allocate internal memory and to initialize backward quadratures. For any other operation (extraction, optional input/output, reinitialization, deallocation), the same function is called regardless of whether or not the quadratures are sensitivity-dependent.

6.2.11.1 Backward quadrature initialization functions

The function `IDAQuadInitB` initializes and allocates memory for the backward integration of quadrature equations that do not depend on forward sensitivities. It has the following form:

IDAQuadInitB

Call `flag = IDAQuadInitB(ida_mem, which, rhsQB, yQB0);`

Description The function `IDAQuadInitB` provides required problem specifications, allocates internal memory, and initializes backward quadrature integration.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`which` (`int`) the identifier of the backward problem.
`rhsQB` (`IDAQuadRhsFnB`) is the C function which computes fQB , the residual of the backward quadrature equations. This function has the form `rhsQB(t, y, yp, yB, ypB, rhsvalBQ, user_dataB)` (see §6.3.3).
`yQB0` (`N_Vector`) is the value of the quadrature variables at `tB0`.

Return value The return value `flag` (of type `int`) will be one of the following:

`IDA_SUCCESS` The call to `IDAQuadInitB` was successful.
`IDA_MEM_NULL` `ida_mem` was NULL.
`IDA_NO_ADJ` The function `IDAAdjInit` has not been previously called.
`IDA_MEM_FAIL` A memory allocation request has failed.
`IDA_ILL_INPUT` The parameter `which` is an invalid identifier.

The function `IDAQuadInitBS` initializes and allocates memory for the backward integration of quadrature equations that depend on the forward sensitivities.

IDAQuadInitBS

Call `flag = IDAQuadInitBS(ida_mem, which, rhsQBS, yQBS0);`

Description The function `IDAQuadInitBS` provides required problem specifications, allocates internal memory, and initializes backward quadrature integration.

Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
`which` (`int`) the identifier of the backward problem.
`rhsQBS` (`IDAQuadRhsFnBS`) is the C function which computes $fQBS$, the residual of the backward quadrature equations. This function has the form `rhsQBS(t, y, yp, yS, ypS, yB, ypB, rhsvalBQS, user_dataB)` (see §6.3.4).
`yQBS0` (`N_Vector`) is the value of the sensitivity-dependent quadrature variables at `tB0`.

Return value The return value `flag` (of type `int`) will be one of the following:

`IDA_SUCCESS` The call to `IDAQuadInitBS` was successful.

IDA_MEM_NULL `ida_mem` was NULL.
 IDA_NO_ADJ The function `IDAAdjInit` has not been previously called.
 IDA_MEM_FAIL A memory allocation request has failed.
 IDA_ILL_INPUT The parameter `which` is an invalid identifier.

The integration of quadrature equations during the backward phase can be re-initialized by calling

IDAQuadReInitB

Call `flag = IDAQuadReInitB(ida_mem, which, yQB0);`
Description The function `IDAQuadReInitB` re-initializes the backward quadrature integration.
Arguments `ida_mem` (`void *`) pointer to the IDAS memory block.
 `which` (`int`) the identifier of the backward problem.
 `yQB0` (`N.Vector`) is the value of the quadrature variables at `tB0`.
Return value The return value `flag` (of type `int`) will be one of the following:
 IDA_SUCCESS The call to `IDAQuadReInitB` was successful.
 IDA_MEM_NULL `ida_mem` was NULL.
 IDA_NO_ADJ The function `IDAAdjInit` has not been previously called.
 IDA_MEM_FAIL A memory allocation request has failed.
 IDA_NO_QUAD Quadrature integration was not activated through a previous call to `IDAQuadInitB`.
 IDA_ILL_INPUT The parameter `which` is an invalid identifier.
Notes `IDAQuadReInitB` can be used after a call to either `IDAQuadInitB` or `IDAQuadInitBS`.

6.2.11.2 Backward quadrature extraction function

To extract the values of the quadrature variables at the last return time of `IDASolveB`, IDAS provides a wrapper for the function `IDAGetQuad` (see §4.7.3). The call to this function has the form

IDAGetQuadB

Call `flag = IDAGetQuadB(ida_mem, which, &tret, yQB);`
Description The function `IDAGetQuadB` returns the quadrature solution vector after a successful return from `IDASolveB`.
Arguments `ida_mem` (`void *`) pointer to the IDAS memory.
 `tret` (`realtype`) the time reached by the solver (output).
 `yQB` (`N.Vector`) the computed quadrature vector.

Return value

Notes T

he user must allocate space for `yQB`. The return value `flag` of `IDAGetQuadB` is one of:

IDA_SUCCESS `IDAGetQuadB` was successful.
 IDA_MEM_NULL `ida_mem` is NULL.
 IDA_NO_ADJ The function `IDAAdjInit` has not been previously called.
 IDA_NO_QUAD Quadrature integration was not initialized.
 IDA_BAD_DKY `yQB` was NULL.
 IDA_ILL_INPUT The parameter `which` is an invalid identifier.



responsibility to access the vector data consistently (including the use of the correct accessor macros from each NVECTOR implementation). For the sake of computational efficiency, the vector functions in the two NVECTOR implementations provided with IDAS do not perform any consistency checks with respect to their `N_Vector` arguments (see §7.1 and §7.2).

The `user_dataB` pointer is passed to the user's `resB` function every time it is called and can be the same as the `user_data` pointer used for the forward problem.

Before calling the user's `resB` function, IDAS needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, IDAS triggers an unrecoverable failure in the residual function which will halt the integration and `IDASolveB` will return `IDA_RESFUNC_FAIL`.



6.3.2 DAE residual for the backward problem depending on the forward sensitivities

The user must provide a `resBS` function of type `IDAResFnBS` defined as follows:

IDAResFnBS

Definition	<pre>typedef int (*IDAResFnBS)(realtype t, N_Vector y, N_Vector yp, N_Vector *yS, N_Vector *ypS, N_Vector yB, N_Vector ypB, N_Vector resvalB, void *user_dataB);</pre>	
Purpose	This function evaluates the residual of the backward problem DAE system. This could be (2.20) or (2.25).	
Arguments	<code>t</code>	is the current value of the independent variable.
	<code>y</code>	is the current value of the forward solution vector.
	<code>yp</code>	is the current value of the forward solution derivative vector.
	<code>yS</code>	a pointer to an array of <code>Ns</code> vectors containing the sensitivities of the forward solution.
	<code>ypS</code>	a pointer to an array of <code>Ns</code> vectors containing the derivatives of the forward sensitivities.
	<code>yB</code>	is the current value of the backward dependent variable vector.
	<code>ypB</code>	is the current value of the backward dependent derivative vector.
	<code>resvalB</code>	is the output vector containing the residual for the backward DAE problem.
	<code>user_dataB</code>	is a pointer to user data, same as passed to <code>IDASetUserDataB</code> .
Return value	An <code>IDAResFnBS</code> should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct), or a negative value if an unrecoverable error occurred (in which case the integration stops and <code>IDASolveB</code> returns <code>IDA_RESFUNC_FAIL</code>).	
Notes	Allocation of memory for <code>resvalB</code> is handled within IDAS.	

The `y`, `yp`, `yB`, `ypB`, and `resvalB` arguments are all of type `N_Vector`, but `yB`, `ypB`, and `resvalB` typically have different internal representations from `y` and `yp`. Likewise for each `yS[i]` and `ypS[i]`. It is the user's responsibility to access the vector data consistently (including the use of the correct accessor macros from each NVECTOR implementation). For the sake of computational efficiency, the vector functions in the two NVECTOR implementations provided with IDAS do not perform any consistency checks with respect to their `N_Vector` arguments (see §7.1 and §7.2).

The `user_dataB` pointer is passed to the user's `resBS` function every time it is called and can be the same as the `user_data` pointer used for the forward problem.

Before calling the user's `resBS` function, IDAS needs to evaluate (through interpolation)



the values of the states from the forward integration. If an error occurs in the interpolation, IDAS triggers an unrecoverable failure in the residual function which will halt the integration and `IDASolveB` will return `IDA_RESFUNC_FAIL`.

6.3.3 Quadrature right-hand side for the backward problem

The user must provide an `fQB` function of type `IDAQuadRhsFnB` defined by

`IDAQuadRhsFnB`

Definition	<pre>typedef int (*IDAQuadRhsFnB)(realtype t, N_Vector y, N_Vector yp, N_Vector yB, N_Vector ypB, N_Vector rhsvalBQ, void *user_dataB);</pre>	
Purpose	This function computes the quadrature equation right-hand side for the backward problem.	
Arguments	t	is the current value of the independent variable.
	y	is the current value of the forward solution vector.
	yp	is the current value of the forward solution derivative vector.
	yB	is the current value of the backward dependent variable vector.
	ypB	is the current value of the backward dependent derivative vector.
	rhsvalBQ	is the output vector containing the residual for the backward quadrature equations.
	user_dataB is a pointer to user data, same as passed to <code>IDASetUserDataB</code> .	
Return value	An <code>IDAQuadRhsFnB</code> should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and <code>IDASolveB</code> returns <code>IDA_QRHSFUNC_FAIL</code>).	
Notes	<p>Allocation of memory for rhsvalBQ is handled within IDAS.</p> <p>The y, yp, yB, ypB, and rhsvalBQ arguments are all of type <code>N_Vector</code>, but they typically all have different internal representations. It is the user's responsibility to access the vector data consistently (including the use of the correct accessor macros from each <code>NVECTOR</code> implementation). For the sake of computational efficiency, the vector functions in the two <code>NVECTOR</code> implementations provided with IDAS do not perform any consistency checks with respect to their <code>N_Vector</code> arguments (see §7.1 and §7.2).</p> <p>The user_dataB pointer is passed to the user's <code>fQB</code> function every time it is called and can be the same as the user_data pointer used for the forward problem.</p> <p>Before calling the user's <code>fQB</code> function, IDAS needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, IDAS triggers an unrecoverable failure in the quadrature right-hand side function which will halt the integration and <code>IDASolveB</code> will return <code>IDA_QRHSFUNC_FAIL</code>.</p>	



6.3.4 Sensitivity-dependent quadrature right-hand side for the backward problem

The user must provide an `fQBS` function of type `IDAQuadRhsFnBS` defined by

`IDAQuadRhsFnBS`

Definition	<pre>typedef int (*IDAQuadRhsFnBS)(realtype t, N_Vector y, N_Vector yp, N_Vector *yS, N_Vector *ypS, N_Vector yB, N_Vector ypB, N_Vector rhsvalBQS, void *user_dataB);</pre>	
------------	--	--

Purpose	This function computes the quadrature equation residual for the backward problem.	
Arguments	<code>t</code>	is the current value of the independent variable.
	<code>y</code>	is the current value of the forward solution vector.
	<code>yp</code>	is the current value of the forward solution derivative vector.
	<code>yS</code>	a pointer to an array of <code>Ns</code> vectors containing the sensitivities of the forward solution.
	<code>ypS</code>	a pointer to an array of <code>Ns</code> vectors containing the derivatives of the forward sensitivities.
	<code>yB</code>	is the current value of the backward dependent variable vector.
	<code>ypB</code>	is the current value of the backward dependent derivative vector.
	<code>rhsvalBQS</code>	is the output vector containing the residual for the backward quadrature equations.
	<code>user_dataB</code>	is a pointer to user data, same as passed to <code>IDASetUserDataB</code> .
Return value	An <code>IDAQuadRhsFnBS</code> should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and <code>IDASolveB</code> returns <code>IDA_QRHSFUNC_FAIL</code>).	
Notes	Allocation of memory for <code>rhsvalBQS</code> is handled within IDAS.	
	The <code>y</code> , <code>yp</code> , <code>yB</code> , <code>ypB</code> , and <code>rhsvalBQS</code> arguments are all of type <code>N_Vector</code> , but they typically do not all have the same internal representations. Likewise for each <code>yS[i]</code> and <code>ypS[i]</code> . It is the user's responsibility to access the vector data consistently (including the use of the correct accessor macros from each <code>NVECTOR</code> implementation). For the sake of computational efficiency, the vector functions in the two <code>NVECTOR</code> implementations provided with IDAS do not perform any consistency checks with respect to their <code>N_Vector</code> arguments (see §7.1 and §7.2).	
	The <code>user_dataB</code> pointer is passed to the user's <code>fQBS</code> function every time it is called and can be the same as the <code>user_data</code> pointer used for the forward problem.	
	Before calling the user's <code>fQBS</code> function, IDAS needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, IDAS triggers an unrecoverable failure in the quadrature right-hand side function which will halt the integration and <code>IDASolveB</code> will return <code>IDA_QRHSFUNC_FAIL</code> .	



6.3.5 Jacobian information for the backward problem (direct method with dense Jacobian)

If the direct linear solver with dense treatment of the Jacobian is selected for the backward problem (i.e. `IDADenseB` or `IDALapackDenseB` is called in step 21 of §6.1), the user may provide, through a call to `IDADlsSetDenseJacFnB` or `IDADlsSetDenseJacFnBS` (see §6.2.9), a function of one of the following two types:

`IDADlsDenseJacFnB`

Definition	<pre>typedef int (*IDADlsDenseJacFnB)(long int NeqB, realtype tt, realtype cjB, N_Vector yy, N_Vector yp, N_Vector yB, N_Vector ypB, N_Vector resvalB, DlsMat JacB, void *user_dataB, N_Vector tmp1B, N_Vector tmp2B, N_Vector tmp3B);</pre>
Purpose	This function computes the dense Jacobian of the backward problem (or an approximation to it).

Arguments	NeqB	is the backward problem size (number of equations).
	tt	is the current value of the independent variable.
	cjB	is the scalar in the system Jacobian, proportional to the inverse of the step size (α in Eq. (2.6)).
	yy	is the current value of the forward solution vector.
	yp	is the current value of the forward solution derivative vector.
	yB	is the current value of the backward dependent variable vector.
	ypB	is the current value of the backward dependent derivative vector.
	resvalB	is the current value of the residual for the backward problem.
	JacB	is the output approximate dense Jacobian matrix.
	user_dataB	is a pointer to user data — the parameter passed to <code>IDASetUserDataB</code> .
	tmp1B	
	tmp2B	
	tmp3B	are pointers to memory allocated for variables of type <code>N_Vector</code> which can be used by <code>IDADlsDenseJacFnB</code> as temporary storage or work space.
	Return value	An <code>IDADlsDenseJacFnB</code> should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct, while <code>IDADENSE</code> sets <code>last_flag</code> to <code>IDADLS_JACFUNC_RECVR</code>), or a negative value if it failed unrecoverably (in which case the integration is halted, <code>IDASolveB</code> returns <code>IDA_LSETUP_FAIL</code> and <code>IDADENSE</code> sets <code>last_flag</code> to <code>IDADLS_JACFUNC_UNRECVR</code>).
Notes	A user-supplied dense Jacobian function must load the <code>NeqB</code> by <code>NeqB</code> dense matrix <code>JacB</code> with an approximation to the Jacobian matrix at the point <code>(tt,yy,yB)</code> , where <code>yy</code> is the solution of the original IVP at time <code>tt</code> and <code>yB</code> is the solution of the backward problem at the same time. Only nonzero elements need to be loaded into <code>JacB</code> as this matrix is set to zero before the call to the Jacobian function. The type of <code>JacB</code> is <code>DlsMat</code> . The user is referred to §4.6.5 for details regarding accessing a <code>DlsMat</code> object.	
	Before calling the user's <code>IDADlsDenseJacFnB</code> , IDAS needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, IDAS triggers an unrecoverable failure in the Jacobian function which will halt the integration (<code>IDASolveB</code> returns <code>IDA_LSETUP_FAIL</code> and <code>IDADENSE</code> sets <code>last_flag</code> to <code>IDADLS_JACFUNC_UNRECVR</code>).	



`IDADlsDenseJacFnBS`

Definition	<pre>typedef int (*IDADlsDenseJacFnBS)(long int NeqB, realtype tt, realtype cjB, N_Vector yy, N_Vector yp, N_Vector *yS, N_Vector *ypS, N_Vector yB, N_Vector ypB, N_Vector resvalB, DlsMat JacB, void *user_dataB, N_Vector tmp1B, N_Vector tmp2B, N_Vector tmp3B);</pre>	
Purpose	This function computes the dense Jacobian of the backward problem (or an approximation to it), in the case where the backward problem depends on the forward sensitivities.	
Arguments	NeqB	is the backward problem size (number of equations).
	tt	is the current value of the independent variable.
	cjB	is the scalar in the system Jacobian, proportional to the inverse of the step size (α in Eq. (2.6)).
	yy	is the current value of the forward solution vector.
	yp	is the current value of the forward solution derivative vector.

<code>yS</code>	a pointer to an array of <code>Ns</code> vectors containing the sensitivities of the forward solution.
<code>ypS</code>	a pointer to an array of <code>Ns</code> vectors containing the derivatives of the forward solution sensitivities.
<code>yB</code>	is the current value of the backward dependent variable vector.
<code>ypB</code>	is the current value of the backward dependent derivative vector.
<code>resvalB</code>	is the current value of the residual for the backward problem.
<code>JacB</code>	is the output approximate dense Jacobian matrix.
<code>user_dataB</code>	is a pointer to user data — the parameter passed to <code>IDASetUserDataB</code> .
<code>tmp1B</code>	
<code>tmp2B</code>	
<code>tmp3B</code>	are pointers to memory allocated for variables of type <code>N_Vector</code> which can be used by <code>IDADlsDenseJacFnBS</code> as temporary storage or work space.
Return value	An <code>IDADlsDenseJacFnBS</code> should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct, while <code>IDADENSE</code> sets <code>last_flag</code> to <code>IDADLS_JACFUNC_RECVR</code>), or a negative value if it failed unrecoverably (in which case the integration is halted, <code>IDASolveB</code> returns <code>IDA_LSETUP_FAIL</code> and <code>IDADENSE</code> sets <code>last_flag</code> to <code>IDADLS_JACFUNC_UNRECVR</code>).
Notes	<p>A user-supplied dense Jacobian function must load the <code>NeqB</code> by <code>NeqB</code> dense matrix <code>JacB</code> with an approximation to the Jacobian matrix at the point <code>(tt,yy,yS,yB)</code>, where <code>yy</code> is the solution of the original IVP at time <code>tt</code>, <code>yS</code> is the array of forward sensitivities at time <code>tt</code>, and <code>yB</code> is the solution of the backward problem at the same time. Only nonzero elements need to be loaded into <code>JacB</code> as this matrix is set to zero before the call to the Jacobian function. The type of <code>JacB</code> is <code>DlsMat</code>. The user is referred to §4.6.5 for details regarding accessing a <code>DlsMat</code> object.</p> <p>Before calling the user's <code>IDADlsDenseJacFnBS</code>, IDAS needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, IDAS triggers an unrecoverable failure in the Jacobian function which will halt the integration (<code>IDASolveB</code> returns <code>IDA_LSETUP_FAIL</code> and <code>IDADENSE</code> sets <code>last_flag</code> to <code>IDADLS_JACFUNC_UNRECVR</code>).</p>



6.3.6 Jacobian information for the backward problem (direct method with banded Jacobian)

If the direct linear solver with banded treatment of the Jacobian is selected for the backward problem (i.e. `IDABandB` or `IDALapackBandB` is called in step 21 of §6.1), the user may provide, through a call to `IDADlsSetBandJacFnB` or `IDADlsSetBandJacFnBS` (see §6.2.9), a function of one of the following two types:

`IDADlsBandJacFnB`

```

Definition    typedef int (*IDADlsBandJacFnB)(long int NeqB,
                                              long int mupperB, long int mlowerB,
                                              realtype tt, realtype cjB,
                                              N_Vector yy, N_Vector yp,
                                              N_Vector yB, N_Vector ypB,
                                              N_Vector resvalB, DlsMat JacB,
                                              void *user_dataB,
                                              N_Vector tmp1B, N_Vector tmp2B,
                                              N_Vector tmp3B);

```

Purpose This function computes the banded Jacobian of the backward problem (or a banded approximation to it).

Arguments	NeqB	is the backward problem size.
	mlowerB	
	mupperB	are the lower and upper half-bandwidth of the Jacobian.
	tt	is the current value of the independent variable.
	cjB	is the scalar in the system Jacobian, proportional to the inverse of the step size (α in Eq. (2.6)).
	yy	is the current value of the forward solution vector.
	yp	is the current value of the forward solution derivative vector.
	yB	is the current value of the backward dependent variable vector.
	ypB	is the current value of the backward dependent derivative vector.
	resvalB	is the current value of the residual for the backward problem.
	JacB	is the output approximate band Jacobian matrix.
	user_dataB	is a pointer to user data — the parameter passed to IDASetUserDataB.
	tmp1B	
	tmp2B	
	tmp3B	are pointers to memory allocated for variables of type <code>N_Vector</code> which can be used by <code>IDADlsBandJacFnB</code> as temporary storage or work space.

Return value An `IDADlsBandJacFnB` should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct, while IDABAND sets `last_flag` to `IDADLS_JACFUNC_RECVR`), or a negative value if it failed unrecoverably (in which case the integration is halted, `IDASolveB` returns `IDA_LSETUP_FAIL` and IDADENSE sets `last_flag` to `IDADLS_JACFUNC_UNRECVR`).

Notes A user-supplied band Jacobian function must load the band matrix `JacB` (of type `DlsMat`) with the elements of the Jacobian at the point `(tt,yy,yB)`, where `yy` is the solution of the original IVP at time `tt` and `yB` is the solution of the backward problem at the same time. Only nonzero elements need to be loaded into `JacB` because `JacB` is preset to zero before the call to the Jacobian function. More details on the accessor macros provided for a `DlsMat` object and on the rest of the arguments passed to a function of type `IDADlsBandJacFnB` are given in §4.6.6.



Before calling the user's `IDADlsBandJacFnB`, IDAS needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, IDAS triggers an unrecoverable failure in the Jacobian function which will halt the integration (`IDASolveB` returns `IDA_LSETUP_FAIL` and IDABAND sets `last_flag` to `IDADLS_JACFUNC_UNRECVR`).

`IDADlsBandJacFnBS`

Definition

```
typedef int (*IDADlsBandJacFnBS)(long int NeqB,
                                long int mupperB, long int mlowerB,
                                realtype tt, realtype cjB,
                                N_Vector yy, N_Vector yp,
                                N_Vector *yS, N_Vector *ypS,
                                N_Vector yB, N_Vector ypB,
                                N_Vector resvalB, DlsMat JacB,
                                void *user_dataB,
                                N_Vector tmp1B, N_Vector tmp2B,
                                N_Vector tmp3B);
```

Purpose This function computes the banded Jacobian of the backward problem (or a banded approximation to it), in the case where the backward problem depends on the forward sensitivities.

Arguments `NeqB` is the backward problem size.

<code>mlowerB</code>	
<code>mupperB</code>	are the lower and upper half-bandwidth of the Jacobian.
<code>tt</code>	is the current value of the independent variable.
<code>cjB</code>	is the scalar in the system Jacobian, proportional to the inverse of the step size (α in Eq. (2.6)).
<code>yy</code>	is the current value of the forward solution vector.
<code>yp</code>	is the current value of the forward solution derivative vector.
<code>yS</code>	a pointer to an array of <code>Ns</code> vectors containing the sensitivities of the forward solution.
<code>ypS</code>	a pointer to an array of <code>Ns</code> vectors containing the derivatives of the forward sensitivities.
<code>yB</code>	is the current value of the backward dependent variable vector.
<code>ypB</code>	is the current value of the backward dependent derivative vector.
<code>resvalB</code>	is the current value of the residual for the backward problem.
<code>JacB</code>	is the output approximate band Jacobian matrix.
<code>user_dataB</code>	is a pointer to user data — the parameter passed to <code>IDASetUserDataB</code> .
<code>tmp1B</code>	
<code>tmp2B</code>	
<code>tmp3B</code>	are pointers to memory allocated for variables of type <code>N_Vector</code> which can be used by <code>IDADlsBandJacFnBS</code> as temporary storage or work space.

Return value An `IDADlsBandJacFnBS` should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct, while `IDABAND` sets `last_flag` to `IDADLS_JACFUNC_RECVR`), or a negative value if it failed unrecoverably (in which case the integration is halted, `IDASolveB` returns `IDA_LSETUP_FAIL` and `IDADENSE` sets `last_flag` to `IDADLS_JACFUNC_UNRECVR`).

Notes A user-supplied band Jacobian function must load the band matrix `JacB` (of type `DlsMat`) with the elements of the Jacobian at the point (tt, yy, yS, yB) , where `yy` is the solution of the original IVP at time `tt`, `yS` is the array of forward sensitivities at time `tt`, and `yB` is the solution of the backward problem at the same time. Only nonzero elements need to be loaded into `JacB` because `JacB` is preset to zero before the call to the Jacobian function. More details on the accessor macros provided for a `DlsMat` object and on the rest of the arguments passed to a function of type `IDADlsBandJacFnBS` are given in §4.6.6.

Before calling the user's `IDADlsBandJacFnBS`, IDAS needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, IDAS triggers an unrecoverable failure in the Jacobian function which will halt the integration (`IDASolveB` returns `IDA_LSETUP_FAIL` and `IDABAND` sets `last_flag` to `IDADLS_JACFUNC_UNRECVR`).



6.3.7 Jacobian information for the backward problem (direct method with sparse Jacobian)

If the direct linear solver with sparse treatment of the Jacobian is selected for the backward problem (i.e. `IDAKLUB` or `IDASuperLUMTB` is called in step 21 of §6.1), the user must provide, through a call to `IDASlsSetSparseJacFnB` or `IDASlsSetSparseJacFnBS` (see §6.2.9), a function of one of the following two types:

`IDASlsSparseJacFnB`

Definition	<pre>typedef int (*IDASlsSparseJacFnB)(realtype tt, realtype cjB, N_Vector yy, N_Vector yp, N_Vector yB, N_Vector ypB, N_Vector rrB, SlsMat JacB, void *user_dataB, N_Vector tmp1B, N_Vector tmp2B, N_Vector tmp3B);</pre>
Purpose	This function computes the sparse Jacobian of the backward problem (or an approximation to it).
Arguments	<p>tt is the current value of the independent variable.</p> <p>cjB is the scalar in the system Jacobian, proportional to the inverse of the step size (α in Eq. (2.6)).</p> <p>yy is the current value of the forward solution vector.</p> <p>yp is the current value of the forward solution derivative vector.</p> <p>yB is the current value of the backward dependent variable vector.</p> <p>ypB is the current value of the backward dependent derivative vector.</p> <p>rrB is the current value of the residual for the backward problem.</p> <p>JacB is the output approximate sparse Jacobian matrix.</p> <p>user_dataB is a pointer to user data — the parameter passed to IDASetUserDataB.</p> <p>tmp1B</p> <p>tmp2B</p> <p>tmp3B are pointers to memory allocated for variables of type <code>N_Vector</code> which can be used by <code>IDASlsSparseJacFnB</code> as temporary storage or work space.</p>
Return value	An <code>IDASlsSparseJacFnB</code> should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct, while IDAKLU or IDASUPERLUMT sets <code>last_flag</code> to <code>IDASLS_JACFUNC_RECVR</code>), or a negative value if it failed unrecoverably (in which case the integration is halted, <code>IDASolveB</code> returns <code>IDA_LSETUP_FAIL</code> and IDAKLU or IDASUPERLUMT sets <code>last_flag</code> to <code>IDASLS_JACFUNC_UNRECVR</code>).
Notes	<p>A user-supplied sparse Jacobian function must load the compressed-sparse-column matrix <code>JacB</code> with an approximation to the Jacobian matrix at the point (tt, yy, yB), where <code>yy</code> is the solution of the original IVP at time <code>tt</code> and <code>yB</code> is the solution of the backward problem at the same time. Storage for <code>JacB</code> already exists on entry to this function, although the user should ensure that sufficient space is allocated in <code>JacB</code> to hold the nonzero values to be set; if the existing space is insufficient the user may reallocate the data and row index arrays as needed. The type of <code>JacB</code> is <code>SlsMat</code>, and the amount of allocated space is available within the <code>SlsMat</code> structure as <code>NNZ</code>. The <code>SlsMat</code> type is further documented in the Section §9.2. The user is referred to §4.6.7 for details regarding accessing a <code>SlsMat</code> object.</p> <p>Before calling the user's <code>IDASlsSparseJacFnB</code>, IDAS needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, IDAS triggers an unrecoverable failure in the Jacobian function which will halt the integration (<code>IDASolveB</code> returns <code>IDA_LSETUP_FAIL</code> and IDAKLU or IDASUPERLUMT sets <code>last_flag</code> to <code>IDASLS_JACFUNC_UNRECVR</code>).</p>



IDASlsSparseJacFnBS

Definition	<pre>typedef int (*IDASlsSparseJacFnBS)(realtype tt, realtype cjB, N_Vector yy, N_Vector yp, N_Vector *yS, N_Vector *ypS, N_Vector yB, N_Vector ypB, N_Vector rrB, SlsMat JacB, void *user_dataB, N_Vector tmp1B, N_Vector tmp2B, N_Vector tmp3B);</pre>
------------	--

Purpose	This function computes the sparse Jacobian of the backward problem (or an approximation to it), in the case where the backward problem depends on the forward sensitivities.	
Arguments	<code>tt</code>	is the current value of the independent variable.
	<code>cjB</code>	is the scalar in the system Jacobian, proportional to the inverse of the step size (α in Eq. (2.6)).
	<code>yy</code>	is the current value of the forward solution vector.
	<code>yp</code>	is the current value of the forward solution derivative vector.
	<code>yS</code>	a pointer to an array of <code>Ns</code> vectors containing the sensitivities of the forward solution.
	<code>ypS</code>	a pointer to an array of <code>Ns</code> vectors containing the derivatives of the forward solution sensitivities.
	<code>yB</code>	is the current value of the backward dependent variable vector.
	<code>ypB</code>	is the current value of the backward dependent derivative vector.
	<code>rrB</code>	is the current value of the residual for the backward problem.
	<code>JacB</code>	is the output approximate sparse Jacobian matrix.
	<code>user_dataB</code>	is a pointer to user data — the parameter passed to <code>IDASetUserDataB</code> .
	<code>tmp1B</code>	
	<code>tmp2B</code>	
	<code>tmp3B</code>	are pointers to memory allocated for variables of type <code>N_Vector</code> which can be used by <code>IDASlsSparseJacFnBS</code> as temporary storage or work space.
Return value	An <code>IDASlsSparseJacFnBS</code> should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct, while IDAKLU or IDASUPERLUMT sets <code>last_flag</code> to <code>IDASLS_JACFUNC_RECVR</code>), or a negative value if it failed unrecoverably (in which case the integration is halted, <code>IDASolveB</code> returns <code>IDA_LSETUP_FAIL</code> and IDAKLU or IDASUPERLUMT sets <code>last_flag</code> to <code>IDASLS_JACFUNC_UNRECVR</code>).	
Notes	<p>A user-supplied sparse Jacobian function must load the compressed-sparse-column matrix <code>JacB</code> with an approximation to the Jacobian matrix at the point (tt, yy, yS, yB), where <code>yy</code> is the solution of the original IVP at time <code>tt</code>, <code>yS</code> is the array of forward sensitivities at time <code>tt</code>, and <code>yB</code> is the solution of the backward problem at the same time. Storage for <code>JacB</code> already exists on entry to this function, although the user should ensure that sufficient space is allocated in <code>JacB</code> to hold the nonzero values to be set; if the existing space is insufficient the user may reallocate the data and row index arrays as needed. The type of <code>JacB</code> is <code>SlsMat</code>, and the amount of allocated space is available within the <code>SlsMat</code> structure as <code>NNZ</code>. The <code>SlsMat</code> type is further documented in the Section §9.2. The user is referred to §4.6.7 for details regarding accessing a <code>SlsMat</code> object.</p> <p>Before calling the user's <code>IDASlsSparseJacFnBS</code>, IDAS needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, IDAS triggers an unrecoverable failure in the Jacobian function which will halt the integration (<code>IDASolveB</code> returns <code>IDA_LSETUP_FAIL</code> and IDAKLU or IDASUPERLUMT sets <code>last_flag</code> to <code>IDASLS_JACFUNC_UNRECVR</code>).</p>	



6.3.8 Jacobian information for the backward problem (matrix-vector product)

If one of the Krylov iterative linear solvers SPGMR, SPBCG, or SPTFQMR is selected (`IDASp*B` is called in step 21 of §6.1), the user may provide a function of one of the following two forms:

`IDASpilsJacTimesVecFnB`

Definition	<pre>typedef int (*IDASpilsJacTimesVecFnB)(realtype t, N_Vector yy, N_Vector yp, N_Vector yB, N_Vector ypB, N_Vector resvalB, N_Vector vB, N_Vector JvB, realtype cjB, void *user_dataB, N_Vector tmp1B, N_Vector tmp2B);</pre>
Purpose	This function computes the action of the backward problem Jacobian JB on a given vector vB.
Arguments	<p>t is the current value of the independent variable.</p> <p>yy is the current value of the forward solution vector.</p> <p>yp is the current value of the forward solution derivative vector.</p> <p>yB is the current value of the backward dependent variable vector.</p> <p>ypB is the current value of the backward dependent derivative vector.</p> <p>resvalB is the current value of the residual for the backward problem.</p> <p>vB is the vector by which the Jacobian must be multiplied.</p> <p>JvB is the computed output vector, $JB \cdot vB$.</p> <p>cjB is the scalar in the system Jacobian, proportional to the inverse of the step size (α in Eq. (2.6)).</p> <p>user_dataB is a pointer to user data — the same as the user_dataB parameter passed to IDASetUserDataB.</p> <p>tmp1B tmp2B are pointers to memory allocated for variables of type N_Vector which can be used by IDASpilsJacTimesVecFnB as temporary storage or work space.</p>
Return value	The return value of a function of type IDASpilsJtimesFnB should be 0 if successful or nonzero if an error was encountered, in which case the integration is halted.
Notes	A user-supplied Jacobian-vector product function must load the vector JvB with the product of the Jacobian of the backward problem at the point (t, y, yB) and the vector vB. Here, y is the solution of the original IVP at time t and yB is the solution of the backward problem at the same time. The rest of the arguments are equivalent to those passed to a function of type IDASpilsJacTimesVecFn (see §4.6.8). If the backward problem is the adjoint of $\dot{y} = f(t, y)$, then this function is to compute $-(\partial f / \partial y)^T v_B$.

IDASpilsJacTimesVecFnBS

Definition	<pre>typedef int (*IDASpilsJacTimesVecFnBS)(realtype t, N_Vector yy, N_Vector yp, N_Vector *yyS, N_Vector *ypS, N_Vector yB, N_Vector ypB, N_Vector resvalB, N_Vector vB, N_Vector JvB, realtype cjB, void *user_dataB, N_Vector tmp1B, N_Vector tmp2B);</pre>
Purpose	This function computes the action of the backward problem Jacobian JB on a given vector vB, in the case where the backward problem depends on the forward sensitivities.
Arguments	<p>t is the current value of the independent variable.</p> <p>yy is the current value of the forward solution vector.</p> <p>yp is the current value of the forward solution derivative vector.</p> <p>yyS a pointer to an array of Ns vectors containing the sensitivities of the forward solution.</p>

<code>ypS</code>	a pointer to an array of <code>Ns</code> vectors containing the derivatives of the forward sensitivities.
<code>yB</code>	is the current value of the backward dependent variable vector.
<code>ypB</code>	is the current value of the backward dependent derivative vector.
<code>resvalB</code>	is the current value of the residual for the backward problem.
<code>vB</code>	is the vector by which the Jacobian must be multiplied.
<code>JvB</code>	is the computed output vector, $JB \cdot vB$.
<code>cjB</code>	is the scalar in the system Jacobian, proportional to the inverse of the step size (α in Eq. (2.6)).
<code>user_dataB</code>	is a pointer to user data — the same as the <code>user_dataB</code> parameter passed to <code>IDASetUserDataB</code> .
<code>tmp1B</code>	
<code>tmp2B</code>	are pointers to memory allocated for variables of type <code>N_Vector</code> which can be used by <code>IDASpilsJacTimesVecFnBS</code> as temporary storage or work space.
Return value	The return value of a function of type <code>IDASpilsJtimesFnBS</code> should be 0 if successful or nonzero if an error was encountered, in which case the integration is halted.
Notes	A user-supplied Jacobian-vector product function must load the vector <code>JvB</code> with the product of the Jacobian of the backward problem at the point (t, y, yB) and the vector <code>vB</code> . Here, y is the solution of the original IVP at time t and yB is the solution of the backward problem at the same time. The rest of the arguments are equivalent to those passed to a function of type <code>IDASpilsJacTimesVecFn</code> (see §4.6.8).

6.3.9 Preconditioning for the backward problem (linear system solution)

If preconditioning is used during integration of the backward problem, then the user must provide a `C` function to solve the linear system $Pz = r$, where P is a left preconditioner matrix. This function must have one of the following two forms:

`IDASpilsPrecSolveFnB`

Definition	<pre>typedef int (*IDASpilsPrecSolveFnB)(realtype t, N_Vector yy, N_Vector yp, N_Vector yB, N_Vector ypB, N_Vector resvalB, N_Vector rvecB, N_Vector zvecB, realtype cjB, realtype deltaB, void *user_dataB, N_Vector tmpB);</pre>	
Purpose	This function solves the preconditioning system $Pz = r$ for the backward problem.	
Arguments	<code>t</code>	is the current value of the independent variable.
	<code>yy</code>	is the current value of the forward solution vector.
	<code>yp</code>	is the current value of the forward solution derivative vector.
	<code>yB</code>	is the current value of the backward dependent variable vector.
	<code>ypB</code>	is the current value of the backward dependent derivative vector.
	<code>resvalB</code>	is the current value of the residual for the backward problem.
	<code>rvecB</code>	is the right-hand side vector r of the linear system to be solved.
	<code>zvecB</code>	is the computed output vector.
	<code>cjB</code>	is the scalar in the system Jacobian, proportional to the inverse of the step size (α in Eq. (2.6)).
	<code>deltaB</code>	is an input tolerance to be used if an iterative method is employed in the solution.

IDASpilsPrecSetupFnB

Definition	<pre>typedef int (*IDASpilsPrecSetupFnB)(realtype t, N_Vector yy, N_Vector yp, N_Vector yB, N_Vector ypB, N_Vector resvalB, realtype cjB, void *user_dataB, N_Vector tmp1B, N_Vector tmp2B, N_Vector tmp3B);</pre>
Purpose	This function preprocesses and/or evaluates Jacobian-related data needed by the preconditioner for the backward problem.
Arguments	<p>The arguments of an <code>IDASpilsPrecSetupFnB</code> are as follows:</p> <p><code>t</code> is the current value of the independent variable.</p> <p><code>yy</code> is the current value of the forward solution vector.</p> <p><code>yp</code> is the current value of the forward solution vector.</p> <p><code>yB</code> is the current value of the backward dependent variable vector.</p> <p><code>ypB</code> is the current value of the backward dependent derivative vector.</p> <p><code>resvalB</code> is the current value of the residual for the backward problem.</p> <p><code>cjB</code> is the scalar in the system Jacobian, proportional to the inverse of the step size (α in Eq. (2.6)).</p> <p><code>user_dataB</code> is a pointer to user data — the same as the <code>user_dataB</code> parameter passed to the function <code>IDASetUserDataB</code>.</p> <p><code>tmp1B</code> <code>tmp2B</code> <code>tmp3B</code> are pointers to memory allocated for vectors which can be used as temporary storage or work space.</p>
Return value	The return value of a preconditioner setup function for the backward problem 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).

IDASpilsPrecSetupFnBS

Definition	<pre>typedef int (*IDASpilsPrecSetupFnBS)(realtype t, N_Vector yy, N_Vector yp, N_Vector *yyS, N_Vector *ypS, N_Vector yB, N_Vector ypB, N_Vector resvalB, realtype cjB, void *user_dataB, N_Vector tmp1B, N_Vector tmp2B, N_Vector tmp3B);</pre>
Purpose	This function preprocesses and/or evaluates Jacobian-related data needed by the preconditioner for the backward problem, in the case where the backward problem depends on the forward sensitivities.
Arguments	<p>The arguments of an <code>IDASpilsPrecSetupFnBS</code> are as follows:</p> <p><code>t</code> is the current value of the independent variable.</p> <p><code>yy</code> is the current value of the forward solution vector.</p> <p><code>yp</code> is the current value of the forward solution vector.</p> <p><code>yyS</code> a pointer to an array of <code>Ns</code> vectors containing the sensitivities of the forward solution.</p>

ypS	a pointer to an array of Ns vectors containing the derivatives of the forward sensitivities.
yB	is the current value of the backward dependent variable vector.
ypB	is the current value of the backward dependent derivative vector.
resvalB	is the current value of the residual for the backward problem.
cjB	is the scalar in the system Jacobian, proportional to the inverse of the step size (α in Eq. (2.6)).
user_dataB	is a pointer to user data — the same as the user_dataB parameter passed to the function IDASetUserDataB .
tmp1B	
tmp2B	
tmp3B	are pointers to memory allocated for vectors which can be used as temporary storage or work space.

Return value The return value of a preconditioner setup function for the backward problem 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).

6.4 Using the band-block-diagonal preconditioner for backward problems

As on the forward integration phase, the efficiency of Krylov iterative methods for the solution of linear systems can be greatly enhanced through preconditioning. The band-block-diagonal preconditioner module **IDABBDPRE**, provides interface functions through which it can be used on the backward integration phase.

The adjoint module in IDAS offers an interface to the band-block-diagonal preconditioner module **IDABBDPRE** described in section §4.8. This generates a preconditioner that is a block-diagonal matrix with each block being a band matrix and can be used with one of the Krylov linear solvers and with the MPI-parallel vector module **NVECTOR_PARALLEL**.

In order to use the **IDABBDPRE** module in the solution of the backward problem, the user must define one or two additional functions, described at the end of this section.

6.4.1 Usage of **IDABBDPRE** for the backward problem

The **IDABBDPRE** module is initialized by calling the following function, *after* one of the **IDASPILS** linear solvers has been specified, by calling the appropriate function (see §6.2.6).

IDABBDPrecInitB

Call	<code>flag = IDABBDPrecInitB(ida_mem, which, NlocalB, mudqB, mldqB, mukeepB, mlkeepB, dqrelyB, GresB, GcommB);</code>
Description	The function IDABBDPrecInitB initializes and allocates memory for the IDABBDPRE preconditioner for the backward problem.
Arguments	<p>ida_mem (void *) pointer to the IDAS memory block.</p> <p>which (int) the identifier of the backward problem.</p> <p>NlocalB (long int) local vector dimension for the backward problem.</p> <p>mudqB (long int) upper half-bandwidth to be used in the difference-quotient Jacobian approximation.</p> <p>mldqB (long int) lower half-bandwidth to be used in the difference-quotient Jacobian approximation.</p> <p>mukeepB (long int) upper half-bandwidth of the retained banded approximate Jacobian block.</p>

- mkeepB** (`long int`) lower half-bandwidth of the retained banded approximate Jacobian block.
- dqrelyB** (`realtype`) the relative increment in components of **yB** used in the difference quotient approximations. The default is $\text{dqrelyB} = \sqrt{\text{unit roundoff}}$, which can be specified by passing $\text{dqrelyB} = 0.0$.
- GresB** (`IDABBDLocalFnB`) the C function which computes $G_B(t, y, \dot{y}, y_B, \dot{y}_B)$, the function approximating the residual of the backward problem.
- GcommB** (`IDABBDCommFnB`) the optional C function which performs all interprocess communication required for the computation of G_B .

Return value If successful, `IDABBDPrecInitB` creates, allocates, and stores (internally in the IDAS solver block) a pointer to the newly created IDABBDPRE memory block. The return value **flag** (of type `int`) is one of:

- `IDASPILS_SUCCESS` The call to `IDABBDPrecInitB` was successful.
- `IDASPILS_MEM_FAIL` A memory allocation request has failed.
- `IDASPILS_MEM_NULL` The `ida_mem` argument was NULL.
- `IDASPILS_LMEM_NULL` No linear solver has been attached.
- `IDASPILS_ILL_INPUT` An invalid parameter has been passed.

To reinitialize the IDABBDPRE preconditioner module for the backward problem, possibly with a change in `mudqB`, `mldqB`, or `dqrelyB`, call the following function:

`IDABBDPrecReInitB`

- Call** `flag = IDABBDPrecReInitB(ida_mem, which, mudqB, mldqB, dqrelyB);`
- Description** The function `IDABBDPrecReInitB` reinitializes the IDABBDPRE preconditioner for the backward problem.
- Arguments**
- ida_mem** (`void *`) pointer to the IDAS memory block returned by `IDACreate`.
 - which** (`int`) the identifier of the backward problem.
 - mudqB** (`long int`) upper half-bandwidth to be used in the difference-quotient Jacobian approximation.
 - mldqB** (`long int`) lower half-bandwidth to be used in the difference-quotient Jacobian approximation.
 - dqrelyB** (`realtype`) the relative increment in components of **yB** used in the difference quotient approximations.

Return value The return value **flag** (of type `int`) is one of:

- `IDASPILS_SUCCESS` The call to `IDABBDPrecReInitB` was successful.
- `IDASPILS_MEM_FAIL` A memory allocation request has failed.
- `IDASPILS_MEM_NULL` The `ida_mem` argument was NULL.
- `IDASPILS_PMEM_NULL` The `IDABBDPrecInitB` has not been previously called.
- `IDASPILS_LMEM_NULL` No linear solver has been attached.
- `IDASPILS_ILL_INPUT` An invalid parameter has been passed.

For more details on IDABBDPRE see §4.8.

6.4.2 User-supplied functions for IDABBDPRE

To use the IDABBDPRE module, the user must supply one or two functions which the module calls to construct the preconditioner: a required function **GresB** (of type `IDABBDLocalFnB`) which approximates the residual of the backward problem and which is computed locally, and an optional function **GcommB** (of type `IDABBDCommFnB`) which performs all interprocess communication necessary to evaluate this approximate residual (see §4.8). The prototypes for these two functions are described below.

IDABBDLocalFnB

Definition	<pre>typedef int (*IDABBDLocalFnB)(long int NlocalB, realtype t, N_Vector y, N_Vector yp, N_Vector yB, N_Vector ypB, N_Vector gB, void *user_dataB);</pre>		
Purpose	This GresB function loads the vector gB , an approximation to the residual of the backward problem, as a function of t , y , yp , and yB and ypB .		
Arguments	NlocalB	is the local vector length for the backward problem.	
	t	is the value of the independent variable.	
	y	is the current value of the forward solution vector.	
	yp	is the current value of the forward solution derivative vector.	
	yB	is the current value of the backward dependent variable vector.	
	ypB	is the current value of the backward dependent derivative vector.	
	gB	is the output vector, $G_B(t, y, \dot{y}, y_B, \dot{y}_B)$.	
	user_dataB	is a pointer to user data — the same as the user_dataB parameter passed to IDASetUserDataB .	
Return value	An IDABBDLocalFnB should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and IDASolveB returns IDA_LSETUP_FAIL).		
Notes	<p>This routine must assume that all interprocess communication of data needed to calculate gB has already been done, and this data is accessible within user_dataB.</p> <p>Before calling the user's IDABBDLocalFnB, IDAS needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, IDAS triggers an unrecoverable failure in the preconditioner setup function which will halt the integration (IDASolveB returns IDA_LSETUP_FAIL).</p>		

**IDABBDCommFnB**

Definition	<pre>typedef int (*IDABBDCommFnB)(long int NlocalB, realtype t, N_Vector y, N_Vector yp, N_Vector yB, N_Vector ypB, void *user_dataB);</pre>		
Purpose	This GcommB function performs all interprocess communications necessary for the execution of the GresB function above, using the input vectors y , yp , yB and ypB .		
Arguments	NlocalB	is the local vector length.	
	t	is the value of the independent variable.	
	y	is the current value of the forward solution vector.	
	yp	is the current value of the forward solution derivative vector.	
	yB	is the current value of the backward dependent variable vector.	
	ypB	is the current value of the backward dependent derivative vector.	
	user_dataB	is a pointer to user data — the same as the user_dataB parameter passed to IDASetUserDataB .	
Return value	An IDABBDCommFnB should return 0 if successful, a positive value if a recoverable error occurred (in which case IDAS will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and IDASolveB returns IDA_LSETUP_FAIL).		

Notes The `GcommB` function is expected to save communicated data in space defined within the structure `user_dataB`.

Each call to the `GcommB` function is preceded by a call to the function that evaluates the residual of the backward problem with the same `t`, `y`, `yp`, `yB` and `ypB` arguments. If there is no additional communication needed, then pass `GcommB = NULL` to `IDABBDPrecInitB`.

Chapter 7

Description of the NVECTOR module

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 the particular NVECTOR implementation. Users can provide their own specific implementation of the NVECTOR module, or use one of four provided within SUNDIALS – a serial implementation and three parallel implementations. 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;
```

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

The `_generic_N_Vector_Ops` structure is essentially a list of pointers to the various actual vector operations, and is defined as

```
struct _generic_N_Vector_Ops {  
    N_Vector      (*nvclone)(N_Vector);  
    N_Vector      (*nvcloneempty)(N_Vector);  
    void          (*nvdestroy)(N_Vector);  
    void          (*nvspace)(N_Vector, long int *, long int *);  
    realtype*     (*nvgetarraypointer)(N_Vector);  
    void          (*nvsetarraypointer)(realtype *, N_Vector);  
    void          (*nvlinearsum)(realtype, N_Vector, realtype, N_Vector, N_Vector);  
    void          (*nvconst)(realtype, N_Vector);  
    void          (*nvprod)(N_Vector, N_Vector, N_Vector);  
    void          (*nvdiv)(N_Vector, N_Vector, N_Vector);  
    void          (*nvscale)(realtype, N_Vector, N_Vector);  
    void          (*nvabs)(N_Vector, N_Vector);  
    void          (*nvinv)(N_Vector, N_Vector);  
    void          (*nvaddconst)(N_Vector, realtype, N_Vector);  
    realtype      (*nvdotprod)(N_Vector, N_Vector);  
    realtype      (*nvmaxnorm)(N_Vector);  
    realtype      (*nvwrmsnorm)(N_Vector, N_Vector);
```

```

realtype    (*nvwrmsnormmask)(N_Vector, N_Vector, N_Vector);
realtype    (*nvmin)(N_Vector);
realtype    (*nvwl2norm)(N_Vector, N_Vector);
realtype    (*nv1lnorm)(N_Vector);
void        (*nvcompare)(realtype, N_Vector, N_Vector);
boolean_t   (*nvintest)(N_Vector, N_Vector);
boolean_t   (*nvconstrmask)(N_Vector, N_Vector, N_Vector);
realtype    (*nvminquotient)(N_Vector, N_Vector);
};

```

The generic NVECTOR module defines and implements the vector operations acting on `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 scaling of a vector `x` by a scalar `c`:

```

void N_VScale(realtype c, N_Vector x, N_Vector z)
{
    z->ops->nvscale(c, x, z);
}

```

Table 7.1 contains a complete list of all vector operations defined by the generic NVECTOR module.

Finally, note that the generic NVECTOR module defines the functions `N_VCloneVectorArray` and `N_VCloneEmptyVectorArray`. Both functions create (by cloning) an array of `count` variables of type `N_Vector`, each of the same type as an existing `N_Vector`. Their prototypes are

```

N_Vector *N_VCloneVectorArray(int count, N_Vector w);
N_Vector *N_VCloneEmptyVectorArray(int count, N_Vector w);

```

and their definitions are based on the implementation-specific `N_VClone` and `N_VCloneEmpty` operations, respectively.

An array of variables of type `N_Vector` can be destroyed by calling `N_VDestroyVectorArray`, whose prototype is

```

void N_VDestroyVectorArray(N_Vector *vs, int count);

```

and whose definition is based on the implementation-specific `N_VDestroy` operation.

A particular implementation of the NVECTOR module must:

- Specify the *content* field of `N_Vector`.
- 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`.

Table 7.1: Description of the NVECTOR operations

Name	Usage and Description
N_VClone	<code>v = N_VClone(w);</code> Creates a new N_Vector of the same type as an existing vector w and sets the <i>ops</i> field. It does not copy the vector, but rather allocates storage for the new vector.
N_VCloneEmpty	<code>v = N_VCloneEmpty(w);</code> Creates a new N_Vector of the same type as an existing vector w and sets the <i>ops</i> field. It does not allocate storage for data.
N_VDestroy	<code>N_VDestroy(v);</code> Destroys the N_Vector v and frees memory allocated for its internal data.
N_VSpace	<code>N_VSpace(nvSpec, &lrw, &liw);</code> Returns storage requirements for one N_Vector . lrw contains the number of realtype words and 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.
N_VGetArrayPointer	<code>vdata = N_VGetArrayPointer(v);</code> Returns a pointer to a realtype array from the N_Vector v . Note that this assumes that the internal data in N_Vector is a contiguous array of realtype . This routine is only used in the solver-specific interfaces to the dense and banded (serial) linear solvers, the sparse linear solvers (serial and threaded), and in the interfaces to the banded (serial) and band-block-diagonal (parallel) preconditioner modules provided with SUNDIALS.
N_VSetArrayPointer	<code>N_VSetArrayPointer(vdata, v);</code> Overwrites the data in an N_Vector with a given array of realtype . Note that this assumes that the internal data in N_Vector is a contiguous array of realtype . This routine is only used in the interfaces to the dense (serial) linear solver, hence need not exist in a user-supplied NVECTOR module for a parallel environment.
N_VLinearSum	<code>N_VLinearSum(a, x, b, y, z);</code> Performs the operation $z = ax + by$, where <i>a</i> and <i>b</i> are realtype scalars and <i>x</i> and <i>y</i> are of type N_Vector : $z_i = ax_i + by_i$, $i = 0, \dots, n - 1$.
N_VConst	<code>N_VConst(c, z);</code> Sets all components of the N_Vector z to realtype <i>c</i> : $z_i = c$, $i = 0, \dots, n - 1$.

continued on next page

<i>continued from last page</i>	
Name	Usage and Description
N_VProd	<code>N_VProd(x, y, z);</code> Sets the <code>N_Vector</code> <code>z</code> to be the component-wise product of the <code>N_Vector</code> inputs <code>x</code> and <code>y</code> : $z_i = x_i y_i$, $i = 0, \dots, n-1$.
N_VDiv	<code>N_VDiv(x, y, z);</code> Sets the <code>N_Vector</code> <code>z</code> to be the component-wise ratio of the <code>N_Vector</code> inputs <code>x</code> and <code>y</code> : $z_i = x_i / y_i$, $i = 0, \dots, n-1$. The y_i may not be tested for 0 values. It should only be called with a <code>y</code> that is guaranteed to have all nonzero components.
N_VScale	<code>N_VScale(c, x, z);</code> Scales the <code>N_Vector</code> <code>x</code> by the <code>realtype</code> scalar <code>c</code> and returns the result in <code>z</code> : $z_i = c x_i$, $i = 0, \dots, n-1$.
N_VAbs	<code>N_VAbs(x, z);</code> Sets the components of the <code>N_Vector</code> <code>z</code> to be the absolute values of the components of the <code>N_Vector</code> <code>x</code> : $y_i = x_i $, $i = 0, \dots, n-1$.
N_VInv	<code>N_VInv(x, z);</code> Sets the components of the <code>N_Vector</code> <code>z</code> to be the inverses of the components of the <code>N_Vector</code> <code>x</code> : $z_i = 1.0 / x_i$, $i = 0, \dots, n-1$. This routine may not check for division by 0. It should be called only with an <code>x</code> which is guaranteed to have all nonzero components.
N_VAddConst	<code>N_VAddConst(x, b, z);</code> Adds the <code>realtype</code> scalar <code>b</code> to all components of <code>x</code> and returns the result in the <code>N_Vector</code> <code>z</code> : $z_i = x_i + b$, $i = 0, \dots, n-1$.
N_VDotProd	<code>d = N_VDotProd(x, y);</code> Returns the value of the ordinary dot product of <code>x</code> and <code>y</code> : $d = \sum_{i=0}^{n-1} x_i y_i$.
N_VMaxNorm	<code>m = N_VMaxNorm(x);</code> Returns the maximum norm of the <code>N_Vector</code> <code>x</code> : $m = \max_i x_i $.
N_VWrmsNorm	<code>m = N_VWrmsNorm(x, w)</code> Returns the weighted root-mean-square norm of the <code>N_Vector</code> <code>x</code> with <code>realtype</code> weight vector <code>w</code> : $m = \sqrt{(\sum_{i=0}^{n-1} (x_i w_i)^2) / n}$.
N_VWrmsNormMask	<code>m = N_VWrmsNormMask(x, w, id);</code> Returns the weighted root mean square norm of the <code>N_Vector</code> <code>x</code> with <code>realtype</code> weight vector <code>w</code> built using only the elements of <code>x</code> corresponding to nonzero elements of the <code>N_Vector</code> <code>id</code> : $m = \sqrt{(\sum_{i=0}^{n-1} (x_i w_i \text{sign}(id_i))^2) / n}.$
N_VMin	<code>m = N_VMin(x);</code> Returns the smallest element of the <code>N_Vector</code> <code>x</code> : $m = \min_i x_i$.
<i>continued on next page</i>	

continued from last page	
Name	Usage and Description
N_VWL2Norm	<code>m = N_VWL2Norm(x, w);</code> Returns the weighted Euclidean ℓ_2 norm of the N_Vector <code>x</code> with <code>realtype</code> weight vector <code>w</code> : $m = \sqrt{\sum_{i=0}^{n-1} (x_i w_i)^2}$.
N_VL1Norm	<code>m = N_VL1Norm(x);</code> Returns the ℓ_1 norm of the N_Vector <code>x</code> : $m = \sum_{i=0}^{n-1} x_i $.
N_VCompare	<code>N_VCompare(c, x, z);</code> Compares the components of the N_Vector <code>x</code> to the <code>realtype</code> scalar <code>c</code> and returns an N_Vector <code>z</code> such that: $z_i = 1.0$ if $ x_i \geq c$ and $z_i = 0.0$ otherwise.
N_VInvTest	<code>t = N_VInvTest(x, z);</code> Sets the components of the N_Vector <code>z</code> to be the inverses of the components of the N_Vector <code>x</code> , with prior testing for zero values: $z_i = 1.0/x_i$, $i = 0, \dots, n-1$. This routine returns a boolean assigned to <code>TRUE</code> if all components of <code>x</code> are nonzero (successful inversion) and returns <code>FALSE</code> otherwise.
N_VConstrMask	<code>t = N_VConstrMask(c, x, m);</code> Performs the following constraint tests: $x_i > 0$ if $c_i = 2$, $x_i \geq 0$ if $c_i = 1$, $x_i \leq 0$ if $c_i = -1$, $x_i < 0$ if $c_i = -2$. There is no constraint on x_i if $c_i = 0$. This routine returns a boolean assigned to <code>FALSE</code> if any element failed the constraint test and assigned to <code>TRUE</code> if all passed. It also sets a mask vector <code>m</code> , 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.
N_VMinQuotient	<code>minq = N_VMinQuotient(num, denom);</code> This routine returns the minimum of the quotients obtained by term-wise dividing <code>num_i</code> by <code>denom_i</code> . A zero element in <code>denom</code> will be skipped. If no such quotients are found, then the large value <code>BIG_REAL</code> (defined in the header file <code>sundials.types.h</code>) is returned.

7.1 The NVECTOR_SERIAL implementation

The serial implementation of the NVECTOR module provided with SUNDIALS, NVECTOR_SERIAL, 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, and a boolean flag *own_data* which specifies the ownership of *data*.

```
struct _N_VectorContent_Serial {
    long int length;
    booleantype own_data;
    realtype *data;
};
```

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`

This routine gives access to the contents of the serial vector N_Vector.

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`, `NV_DATA_S`, `NV_LENGTH_S`

These macros give individual access to the parts of the content of a serial `N_Vector`.

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`. The assignment `NV_DATA_S(v) = v_data` sets the component array of `v` to be `v_data` by storing the pointer `v_data`.

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_OWN_DATA_S(v) ( NV_CONTENT_S(v)->own_data )
```

```
#define NV_DATA_S(v) ( NV_CONTENT_S(v)->data )
```

```
#define NV_LENGTH_S(v) ( NV_CONTENT_S(v)->length )
```

- `NV_Ith_S`

This macro gives access to the individual components of the data array of an `N_Vector`.

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] )
```

The `NVECTOR_SERIAL` module defines serial implementations of all vector operations listed in Table 7.1. Their names are obtained from those in Table 7.1 by appending the suffix `_Serial`. The module `NVECTOR_SERIAL` provides the following additional user-callable routines:

- `N_VNew_Serial`

This function creates and allocates memory for a serial `N_Vector`. Its only argument is the vector length.

```
N_Vector N_VNew_Serial(long int vec_length);
```

- `N_VNewEmpty_Serial`

This function creates a new serial `N_Vector` with an empty (NULL) data array.

```
N_Vector N_VNewEmpty_Serial(long int vec_length);
```

- `N_VMake_Serial`

This function creates and allocates memory for a serial vector with user-provided data array.

```
N_Vector N_VMake_Serial(long int vec_length, realtype *v_data);
```

- `N_VCloneVectorArray_Serial`

This function creates (by cloning) an array of count serial vectors.

```
N_Vector *N_VCloneVectorArray_Serial(int count, N_Vector w);
```

- `N_VCloneEmptyVectorArray_Serial`

This function creates (by cloning) an array of count serial vectors, each with an empty (NULL) data array.

```
N_Vector *N_VCloneEmptyVectorArray_Serial(int count, N_Vector w);
```

- `N_VDestroyVectorArray_Serial`

This function frees memory allocated for the array of `count` variables of type `N_Vector` created with `N_VCloneVectorArray_Serial` or with `N_VCloneEmptyVectorArray_Serial`.

```
void N_VDestroyVectorArray_Serial(N_Vector *vs, int count);
```

- `N_VPrint_Serial`

This function prints the content of a serial vector to `stdout`.

```
void N_VPrint_Serial(N_Vector v);
```

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)` 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`, and `N_VCloneEmptyVectorArray_Serial` set the field `own_data = FALSE`. `N_VDestroy_Serial` and `N_VDestroyVectorArray_Serial` will not attempt to free the pointer `data` for any `N_Vector` with `own_data` set to `FALSE`. 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 internal representations.



7.2 The NVECTOR_PARALLEL implementation

The `NVECTOR_PARALLEL` implementation of the `NVECTOR` module provided with `SUNDIALS` is based on `MPI`. It defines the `content` field of `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 {
    long int local_length;
    long int global_length;
    booleantype own_data;
    realtype *data;
    MPI_Comm comm;
};
```

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`

This macro gives access to the contents of the parallel vector `N_Vector`.

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_VectorParallelContent`.

Implementation:

```
#define NV_CONTENT_P(v) ( (N_VectorContent_Parallel)(v->content) )
```

- `NV_OWN_DATA_P`, `NV_DATA_P`, `NV_LOCLENGTH_P`, `NV_GLOBLENGTH_P`

These macros give individual access to the parts of the content of a parallel `N_Vector`.

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`.

The assignment `v_llen = NV_LOCLENGTH_P(v)` sets `v_llen` to be the length of the local part of `v`. The call `NV_LENGTH_P(v) = llen_v` sets the local length of `v` to be `llen_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_OWN_DATA_P(v)    ( NV_CONTENT_P(v)->own_data )
#define NV_DATA_P(v)        ( NV_CONTENT_P(v)->data )
#define NV_LOCLENGTH_P(v)   ( NV_CONTENT_P(v)->local_length )
#define NV_GLOBLENGTH_P(v) ( NV_CONTENT_P(v)->global_length )
```

- `NV_COMM_P`

This macro provides access to the MPI communicator used by the `NVECTOR_PARALLEL` vectors.

Implementation:

```
#define NV_COMM_P(v) ( NV_CONTENT_P(v)->comm )
```

- `NV_Ith_P`

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] )
```

The `NVECTOR_PARALLEL` module defines parallel implementations of all vector operations listed in Table 7.1 Their names are obtained from those in Table 7.1 by appending the suffix `_Parallel`. The module `NVECTOR_PARALLEL` provides the following additional user-callable routines:

- `N_VNew_Parallel`

This function creates and allocates memory for a parallel vector.

```
N_Vector N_VNew_Parallel(MPI_Comm comm,
                        long int local_length,
                        long int global_length);
```

- `N_VNewEmpty_Parallel`

This function creates a new parallel `N_Vector` with an empty (NULL) data array.

```
N_Vector N_VNewEmpty_Parallel(MPI_Comm comm,
                             long int local_length,
                             long int global_length);
```

- `N_VMake_Parallel`

This function creates and allocates memory for a parallel vector with user-provided data array.

```
N_Vector N_VMake_Parallel(MPI_Comm comm,
                          long int local_length,
                          long int global_length,
                          realtype *v_data);
```

- **N_VCloneVectorArray_Parallel**

This function creates (by cloning) an array of `count` parallel vectors.

```
N_Vector *N_VCloneVectorArray_Parallel(int count, N_Vector w);
```

- **N_VCloneEmptyVectorArray_Parallel**

This function creates (by cloning) an array of `count` parallel vectors, each with an empty (NULL) data array.

```
N_Vector *N_VCloneEmptyVectorArray_Parallel(int count, N_Vector w);
```

- **N_VDestroyVectorArray_Parallel**

This function frees memory allocated for the array of `count` variables of type `N_Vector` created with `N_VCloneVectorArray_Parallel` or with `N_VCloneEmptyVectorArray_Parallel`.

```
void N_VDestroyVectorArray_Parallel(N_Vector *vs, int count);
```

- **N_VPrint_Parallel**

This function prints the content of a parallel vector to stdout.

```
void N_VPrint_Parallel(N_Vector v);
```

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 = 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`, `N_VMake_Parallel`, and `N_VCloneEmptyVectorArray_Parallel` set the field `own_data = FALSE`. `N_VDestroy_Parallel` and `N_VDestroyVectorArray_Parallel` will not attempt to free the pointer `data` for any `N_Vector` with `own_data` set to `FALSE`. 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.



7.3 The NVECTOR_OPENMP implementation

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.

```

struct _N_VectorContent_OpenMP {
    long int length;
    boolean_t own_data;
    realtype *data;
    int num_threads;
};

```

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**

This routine gives access to the contents of the OpenMP vector `N_Vector`.

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, NV_DATA_OMP, NV_LENGTH_OMP, NV_NUM_THREADS_OMP**

These macros give individual access to the parts of the content of a OpenMP `N_Vector`.

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`. The assignment `NV_DATA_OMP(v) = v_data` sets the component array of `v` to be `v_data` by storing the pointer `v_data`.

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`.

The assignment `v_num_threads = NV_NUM_THREADS_OMP(v)` sets `v_num_threads` to be the number of threads from `v`. On the other hand, the call `NV_NUM_THREADS_OMP(v) = num_threads_v` sets the number of threads for `v` to be `num_threads_v`.

Implementation:

```
#define NV_OWN_DATA_OMP(v) ( NV_CONTENT_OMP(v)->own_data )
```

```
#define NV_DATA_OMP(v) ( NV_CONTENT_OMP(v)->data )
```

```
#define NV_LENGTH_OMP(v) ( NV_CONTENT_OMP(v)->length )
```

```
#define NV_NUM_THREADS_OMP(v) ( NV_CONTENT_OMP(v)->num_threads )
```

- **NV_Ith_OMP**

This macro gives access to the individual components of the data array of an `N_Vector`.

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] )
```

The NVECTOR_OPENMP module defines OpenMP implementations of all vector operations listed in Table 7.1. Their names are obtained from those in Table 7.1 by appending the suffix `_OpenMP`. The module NVECTOR_OPENMP provides the following additional user-callable routines:

- **N_VNew_OpenMP**

This function creates and allocates memory for a OpenMP `N_Vector`. Arguments are the vector length and number of threads.

```
N_Vector N_VNew_OpenMP(long int vec_length, int num_threads);
```

- **N_VNewEmpty_OpenMP**
This function creates a new OpenMP **N_Vector** with an empty (NULL) data array.
`N_Vector N_VNewEmpty_OpenMP(long int vec_length, int num_threads);`
- **N_VMake_OpenMP**
This function creates and allocates memory for a OpenMP vector with user-provided data array.
`N_Vector N_VMake_OpenMP(long int vec_length, realtype *v_data, int num_threads);`
- **N_VCloneVectorArray_OpenMP**
This function creates (by cloning) an array of `count` OpenMP vectors.
`N_Vector *N_VCloneVectorArray_OpenMP(int count, N_Vector w);`
- **N_VCloneEmptyVectorArray_OpenMP**
This function creates (by cloning) an array of `count` OpenMP vectors, each with an empty (NULL) data array.
`N_Vector *N_VCloneEmptyVectorArray_OpenMP(int count, N_Vector w);`
- **N_VDestroyVectorArray_OpenMP**
This function frees memory allocated for the array of `count` variables of type **N_Vector** created with **N_VCloneVectorArray_OpenMP** or with **N_VCloneEmptyVectorArray_OpenMP**.
`void N_VDestroyVectorArray_OpenMP(N_Vector *vs, int count);`
- **N_VPrint_OpenMP**
This function prints the content of a OpenMP vector to `stdout`.
`void N_VPrint_OpenMP(N_Vector v);`

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_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**, **N_VMake_OpenMP**, and **N_VCloneEmptyVectorArray_OpenMP** set the field `own_data = FALSE`. **N_VDestroy_OpenMP** and **N_VDestroyVectorArray_OpenMP** will not attempt to free the pointer `data` for any **N_Vector** with `own_data` set to `FALSE`. 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.



7.4 The NVECTOR_PTHREADS implementation

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, **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 {
    long int length;
    boolean_t own_data;
    realtype *data;
    int num_threads;
};

```

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**

This routine gives access to the contents of the Pthreads vector `N_Vector`.

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, NV_DATA_PT, NV_LENGTH_PT, NV_NUM_THREADS_PT**

These macros give individual access to the parts of the content of a Pthreads `N_Vector`.

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`. The assignment `NV_DATA_PT(v) = v_data` sets the component array of `v` to be `v_data` by storing the pointer `v_data`.

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`.

The assignment `v_num_threads = NV_NUM_THREADS_PT(v)` sets `v_num_threads` to be the number of threads from `v`. On the other hand, the call `NV_NUM_THREADS_PT(v) = num_threads_v` sets the number of threads for `v` to be `num_threads_v`.

Implementation:

```
#define NV_OWN_DATA_PT(v) ( NV_CONTENT_PT(v)->own_data )
```

```
#define NV_DATA_PT(v) ( NV_CONTENT_PT(v)->data )
```

```
#define NV_LENGTH_PT(v) ( NV_CONTENT_PT(v)->length )
```

```
#define NV_NUM_THREADS_PT(v) ( NV_CONTENT_PT(v)->num_threads )
```

- **NV_Ith_PT**

This macro gives access to the individual components of the data array of an `N_Vector`.

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] )
```

The NVECTOR_PTHREADS module defines Pthreads implementations of all vector operations listed in Table 7.1. Their names are obtained from those in Table 7.1 by appending the suffix `_Pthreads`. The module NVECTOR_PTHREADS provides the following additional user-callable routines:

- **N_VNew_Pthreads**

This function creates and allocates memory for a Pthreads `N_Vector`. Arguments are the vector length and number of threads.

```
N_Vector N_VNew_Pthreads(long int vec_length, int num_threads);
```


- **N_VNewEmpty_Pthreads**

This function creates a new Pthreads **N_Vector** with an empty (NULL) data array.

```
N_Vector N_VNewEmpty_Pthreads(long int vec_length, int num_threads);
```

- **N_VMake_Pthreads**

This function creates and allocates memory for a Pthreads vector with user-provided data array.

```
N_Vector N_VMake_Pthreads(long int vec_length, realtype *v_data, int num_threads);
```

- **N_VCloneVectorArray_Pthreads**

This function creates (by cloning) an array of **count** Pthreads vectors.

```
N_Vector *N_VCloneVectorArray_Pthreads(int count, N_Vector w);
```

- **N_VCloneEmptyVectorArray_Pthreads**

This function creates (by cloning) an array of **count** Pthreads vectors, each with an empty (NULL) data array.

```
N_Vector *N_VCloneEmptyVectorArray_Pthreads(int count, N_Vector w);
```

- **N_VDestroyVectorArray_Pthreads**

This function frees memory allocated for the array of **count** variables of type **N_Vector** created with **N_VCloneVectorArray_Pthreads** or with **N_VCloneEmptyVectorArray_Pthreads**.

```
void N_VDestroyVectorArray_Pthreads(N_Vector *vs, int count);
```

- **N_VPrint_Pthreads**

This function prints the content of a Pthreads vector to **stdout**.

```
void N_VPrint_Pthreads(N_Vector v);
```

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_PT(v)** and then access **v_data[i]** within the loop than it is to use **NV_Ith_PT(v,i)** within the loop.
- **N_VNewEmpty_Pthreads**, **N_VMake_Pthreads**, and **N_VCloneEmptyVectorArray_Pthreads** set the field **own_data = FALSE**. **N_VDestroy_Pthreads** and **N_VDestroyVectorArray_Pthreads** will not attempt to free the pointer **data** for any **N_Vector** with **own_data** set to **FALSE**. 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.



7.5 NVECTOR Examples

Eddy to put text here.

7.6 NVECTOR functions used by IDAS

In Table 7.2 below, we list the vector functions used in the NVECTOR module used by the IDAS package. The table also shows, for each function, which of the code modules uses the function. The IDAS column shows function usage within the main integrator module, while the remaining five columns show function usage within each of the five IDAS linear solvers, the IDABBDPRE preconditioner module, and the IDAS adjoint sensitivity module (denoted here by IDAA). Here IDADLS stands for IDADENSE and IDABAND; IDASPILS stands for IDASPGMR, IDASPCBG, and IDASPTFQMR; and IDASLS stands for IDAKLU and IDASUPERLUMT.

There is one subtlety in the IDASPILS column hidden by the table, explained here for the case of the IDASPGMR module. The `N_VDotProd` function is called both within the interface file `ida_spgmr.c` and within the implementation files `sundials_spgmr.c` and `sundials_iterative.c` for the generic SPGMR solver upon which the IDASPGMR solver is built. Also, although `N_VDiv` and `N_VProd` are not called within the interface file `ida_spgmr.c`, they are called within the implementation file `sundials_spgmr.c`, and so are required by the IDASPGMR solver module. Analogous statements apply to the IDASPCBG and IDASPTFQMR modules, except that they do not use `sundials_iterative.c`. This issue does not arise for the direct IDAS linear solvers because the generic DENSE and BAND solvers (used in the implementation of IDADENSE and IDABAND) do not make calls to any vector functions.

Table 7.2: List of vector functions usage by IDAS code modules

	IDAS	IDADLS	IDASPILS	IDASLS	IDABBDPRE	IDAA
<code>N_VClone</code>	✓		✓		✓	✓
<code>N_VDestroy</code>	✓		✓		✓	✓
<code>N_VCloneVectorArray</code>	✓					✓
<code>N_VDestroyVectorArray</code>	✓					✓
<code>N_VSpace</code>	✓					
<code>N_VGetArrayPointer</code>		✓		✓	✓	
<code>N_VSetArrayPointer</code>		✓				
<code>N_VLinearSum</code>	✓	✓	✓			✓
<code>N_VConst</code>	✓		✓			✓
<code>N_VProd</code>	✓		✓			
<code>N_VDiv</code>	✓		✓			
<code>N_VScale</code>	✓	✓	✓	✓	✓	✓
<code>N_VAbs</code>	✓					
<code>N_VInv</code>	✓					
<code>N_VAddConst</code>	✓					
<code>N_VDotProd</code>			✓			
<code>N_VMaxNorm</code>	✓					
<code>N_VWrmsNorm</code>	✓		✓			
<code>N_VMin</code>	✓					
<code>N_VMinQuotient</code>	✓					
<code>N_VConstrMask</code>	✓					
<code>N_VWrmsNormMask</code>	✓					
<code>N_VCompare</code>	✓					

Of the functions listed in Table 7.1, `N_VWL2Norm`, `N_VL1Norm`, `N_VCloneEmpty`, and `N_VInvTest` are *not* used by IDAS. Therefore a user-supplied NVECTOR module for IDAS could omit these four functions.

Chapter 8

Providing Alternate Linear Solver Modules

The central IDAS module interfaces with a linear solver module by way of calls to five functions. These are denoted here by `linit`, `lsetup`, `lsolve`, `lperf`, and `lfree`. Briefly, their purposes are as follows:

- `linit`: initialize memory specific to the linear solver;
- `lsetup`: evaluate and preprocess the Jacobian or preconditioner;
- `lsolve`: solve the linear system;
- `lperf`: monitor performance and issue warnings;
- `lfree`: free the linear solver memory.

A linear solver module must also provide a user-callable **specification function** (like those described in §4.5.3) which will attach the above five functions to the main IDAS memory block. The IDAS memory block is a structure defined in the header file `idas_impl.h`. A pointer to such a structure is defined as the type `IDAMem`. The five fields in an `IDAMem` structure that must point to the linear solver's functions are `ida_linit`, `ida_lsetup`, `ida_lsolve`, `ida_lperf`, and `ida_lfree`, respectively. Note that of the five interface functions, only `lsolve` is required. The `lfree` function must be provided only if the solver specification function makes any memory allocation. For any of the functions that are *not* provided, the corresponding field should be set to `NULL`. The linear solver specification function must also set the value of the field `ida_setupNonNull` in the IDAS memory block — to `TRUE` if `lsetup` is used, or `FALSE` otherwise.

Typically, the linear solver will require a block of memory specific to the solver, and a principal function of the specification function is to allocate that memory block, and initialize it. Then the field `ida_lmem` in the IDAS memory block is available to attach a pointer to that linear solver memory. This block can then be used to facilitate the exchange of data between the five interface functions.

If the linear solver involves adjustable parameters, the specification function should set the default values of those. User-callable functions may be defined that could, optionally, override the default parameter values.

We encourage the use of performance counters in connection with the various operations involved with the linear solver. Such counters would be members of the linear solver memory block, would be initialized in the `linit` function, and would be incremented by the `lsetup` and `lsolve` functions. Then, user-callable functions would be needed to obtain the values of these counters.

For consistency with the existing IDAS linear solver modules, we recommend that the return value of the specification function be 0 for a successful return, and a negative value if an error occurs. Possible error conditions include: the pointer to the main IDAS memory block is `NULL`, an input is illegal, the `NVECTOR` implementation is not compatible, or a memory allocation fails.

To be used during the backward integration with the IDAS module, a linear solver module must also provide an additional user-callable specification function (like those described in §6.2.6) which will attach the five functions to the IDAS memory block for each backward integration. Note that this block, of type `IDAMem`, is not directly accessible to the specification function, but rather is itself a field in the IDAS memory block. For a given backward problem identifier `which`, the corresponding memory block must be located in the linked list starting at `ida_mem->ida_adj_mem->IDAB_mem`; see for example the function `IDADenseB` for specific details. This specification function must also allocate the linear solver memory for the backward problem, and attach that, as well as a corresponding memory free function, to the above block `IDAB_mem`, of type `struct IDABMemRec`. The specification function for backward integration should return a negative value if it encounters an illegal input, if backward integration has not been initialized, or if its memory allocation failed.

These five functions, which interface between IDAS and the linear solver module, necessarily have fixed call sequences. Thus a user wishing to implement another linear solver within the IDAS package must adhere to this set of interfaces. The following is a complete description of the call list for each of these functions. Note that the call list of each function includes a pointer to the main IDAS memory block, by which the function can access various data related to the IDAS solution. The contents of this memory block are given in the file `idas_impl.h` (but not reproduced here, for the sake of space).

8.1 Initialization function

The type definition of `linit` is

`linit`

Definition `int (*linit)(IDAMem IDA_mem);`

Purpose The purpose of `linit` is to complete initializations for the specific linear solver, such as counters and statistics. It should also set pointers to data blocks that will later be passed to functions associated with the linear solver. The `linit` function is called once only, at the start of the problem, during the first call to `IDASolve`.

Arguments `IDA_mem` is the IDAS memory pointer of type `IDAMem`.

Return value An `linit` function should return 0 if it has successfully initialized the IDAS linear solver and a negative value otherwise.

8.2 Setup function

The type definition of `lsetup` is

`lsetup`

Definition `int (*lsetup)(IDAMem IDA_mem, N_Vector ypp, N_Vector ypp, N_Vector resp, N_Vector vtemp1, N_Vector vtemp2, N_Vector vtemp3);`

Purpose The job of `lsetup` is to prepare the linear solver for subsequent calls to `lsolve`, in the solution of systems $Mx = b$, where M is some approximation to the system Jacobian, $J = \partial F / \partial y + cj \partial F / \partial y$. (See Eqn. (2.6), in which $\alpha = cj$). Here `cj` is available as `IDA_mem->ida_cj`.

The `lsetup` function may call a user-supplied function, or a function within the linear solver module, to compute Jacobian-related data that is required by the linear solver. It may also preprocess that data as needed for `lsolve`, which may involve calling a generic function (such as for LU factorization). This data may be intended either for direct use (in a direct linear solver) or for use in a preconditioner (in a preconditioned iterative linear solver).

The `lsetup` function is not called at every time step, but only as frequently as the solver determines that it is appropriate to perform the setup task. In this way, Jacobian-related data generated by `lsetup` is expected to be used over a number of time steps.

Arguments `IDA_mem` is the IDAS memory pointer of type `IDAMem`.
`yyp` is the predicted y vector for the current IDAS internal step.
`ypp` is the predicted \dot{y} vector for the current IDAS internal step.
`resp` is the value of the residual function at `yyp` and `ypp`, i.e. $F(t_n, y_{pred}, \dot{y}_{pred})$.
`vtemp1`
`vtemp2`
`vtemp3` are temporary variables of type `N_Vector` provided for use by `lsetup`.

Return value The `lsetup` function should return 0 if successful, a positive value for a recoverable error, and a negative value for an unrecoverable error. On a recoverable error return, the solver will attempt to recover by reducing the step size.

8.3 Solve function

The type definition of `lsolve` is

`lsolve`

Definition `int (*lsolve)(IDAMem IDA_mem, N_Vector b, N_Vector weight, N_Vector ycur, N_Vector ypcur, N_Vector rescur);`

Purpose The `lsolve` function must solve the linear system, $Mx = b$, where M is some approximation to the system Jacobian, $J = \partial F / \partial y + cj \partial F / \partial \dot{y}$ (see Eqn. (2.6), in which $\alpha = cj$), and the right-hand side vector, b , is input. Here cj is available as `IDA_mem->ida_cj`.

`lsolve` is called once per Newton iteration, hence possibly several times per time step.

If there is an `lsetup` function, this `lsolve` function should make use of any Jacobian data that was computed and preprocessed by `lsetup`, either for direct use, or for use in a preconditioner.

Arguments `IDA_mem` is the IDAS memory pointer of type `IDAMem`.
`b` is the right-hand side vector b . The solution is to be returned in the vector `b`.
`weight` is a vector that contains the error weights. These are the W_i of (2.7). This weight vector is included here to enable the computation of weighted norms needed to test for the convergence of iterative methods (if any) within the linear solver.
`ycur` is a vector that contains the solver's current approximation to $y(t_n)$.
`ypcur` is a vector that contains the solver's current approximation to $\dot{y}(t_n)$.
`rescur` is a vector that contains the current residual, $F(t_n, ycur, ypcur)$.

Return value The `lsolve` function should return a positive value for a recoverable error and a negative value for an unrecoverable error. Success is indicated by a 0 return value. On a recoverable error return, the solver will attempt to recover, such as by calling the `lsetup` function with current arguments.

8.4 Performance monitoring function

The type definition of `lperf` is

lperf

Definition `int (*lperf)(IDAMem IDA_mem, int perftask);`

Purpose The `lperf` function is to monitor the performance of the linear solver. It can be used to compute performance metrics related to the linear solver and issue warnings if these indicate poor performance of the linear solver. The `lperf` function is called with `perftask = 0` at the start of each call to `IDASolve`, and then is called with `perftask = 1` just before each internal time step.

Arguments `IDA_mem` is the IDAS memory pointer of type `IDAMem`.
`perftask` is a task flag. `perftask = 0` means initialize needed counters. `perftask = 1` means evaluate performance and issue warnings if needed. Counters that are used to compute performance metrics (e.g. counts of iterations within the `lsolve` function) should be initialized here when `perftask = 0`, and used for the calculation of metrics when `perftask = 1`.

Return value The `lperf` return value is ignored.

8.5 Memory deallocation function

The type definition of `lfree` is

lfree

Definition `void (*lfree)(IDAMem IDA_mem);`

Purpose The `lfree` function should free up any memory allocated by the linear solver.

Arguments The argument `IDA_mem` is the IDAS memory pointer of type `IDAMem`.

Return value The `lfree` function has no return value.

Notes This function is called once a problem has been completed and the linear solver is no longer needed.

Chapter 9

General Use Linear Solver Components in SUNDIALS

In this chapter, we describe seven generic linear solver code modules that are included in SUNDIALS, but which are of potential use as generic packages in themselves, either in conjunction with the use of SUNDIALS or separately.

These generic modules in SUNDIALS are organized in three families, the *dls* family, which includes direct linear solvers appropriate for sequential computations; the *sls* family, which includes sparse matrices; and the *spils* family, which includes scaled preconditioned iterative (Krylov) linear solvers. The solvers in each family share common data structures and functions.

The *dls* family contains the following two generic linear solvers:

- The DENSE package, a linear solver for dense matrices either specified through a matrix type (defined below) or as simple arrays.
- The BAND package, a linear solver for banded matrices either specified through a matrix type (defined below) or as simple arrays.

Note that this family also includes the Blas/Lapack linear solvers (dense and band) available to the SUNDIALS solvers, but these are not discussed here.

The *sls* family contains a sparse matrix package and interfaces between it and two sparse direct solver packages:

- The KLU package, a linear solver for compressed-sparse-column matrices, [1, 14].
- The SUPERLUMT package, a threaded linear solver for compressed-sparse-column matrices, [2, 25, 15].

The *spils* family contains the following generic linear solvers:

- The SPGMR package, a solver for the scaled preconditioned GMRES method.
- The SPFGMR package, a solver for the scaled preconditioned Flexible GMRES method.
- The SPBCG package, a solver for the scaled preconditioned Bi-CGStab method.
- The SPTFQMR package, a solver for the scaled preconditioned TFQMR method.

For reasons related to installation, the names of the files involved in these packages begin with the prefix `sundials_`. But despite this, each of the *dls* and *spils* solvers is in fact generic, in that it is usable completely independently of SUNDIALS.

For the sake of space, the functions for the `dense` and `band` modules that work with a matrix type and the functions in the SPGMR, SPFGMR, SPBCG, and SPTFQMR modules are only summarized briefly, since they are less likely to be of direct use in connection with a SUNDIALS solver. However, the

functions for dense matrices treated as simple arrays and sparse matrices are fully described, because we expect that they will be useful in the implementation of preconditioners used with the combination of one of the SUNDIALS solvers and one of the *spils* linear solvers.

9.1 The DLS modules: DENSE and BAND

The files comprising the DENSE generic linear solver, and their locations in the SUNDIALS *srcdir*, are as follows:

- header files (located in *srcdir/include/sundials*)
`sundials_direct.h`, `sundials_dense.h`,
`sundials_types.h`, `sundials_math.h`, `sundials_config.h`
- source files (located in *srcdir/src/sundials*)
`sundials_direct.c`, `sundials_dense.c`, `sundials_math.c`

The files comprising the BAND generic linear solver are as follows:

- header files (located in *srcdir/include/sundials*)
`sundials_direct.h`, `sundials_band.h`,
`sundials_types.h`, `sundials_math.h`, `sundials_config.h`
- source files (located in *srcdir/src/sundials*)
`sundials_direct.c`, `sundials_band.c`, `sundials_math.c`

Only two of the preprocessing directives in the header file `sundials_config.h` are relevant to the DENSE and BAND packages by themselves (see §?? for details):

- (required) definition of the precision of the SUNDIALS type `realtype`. One of the following lines must be present:
`#define SUNDIALS_DOUBLE_PRECISION 1`
`#define SUNDIALS_SINGLE_PRECISION 1`
`#define SUNDIALS_EXTENDED_PRECISION 1`
- (optional) use of generic math functions: `#define SUNDIALS_USE_GENERIC_MATH 1`

The `sundials_types.h` header file defines the SUNDIALS `realtype` and `booleantype` types and the macro `RCONST`, while the `sundials_math.h` header file is needed for the macros `SUNMIN` and `SUNMAX`, and the function `SUNRabs`.

The files listed above for either module can be extracted from the SUNDIALS *srcdir* and compiled by themselves into a separate library or into a larger user code.

9.1.1 Type `DlsMat`

The type `DlsMat`, defined in `sundials_direct.h` is a pointer to a structure defining a generic matrix, and is used with all linear solvers in the *dls* family:

```
typedef struct _DlsMat {
    int type;
    long int M;
    long int N;
    long int ldim;
    long int mu;
    long int ml;
    long int s_mu;
    realtype *data;
    long int ldata;
    realtype **cols;
} *DlsMat;
```


For the DENSE module, the relevant fields of this structure are as follows. Note that a dense matrix of type `DlsMat` need not be square.

type - `SUNDIALS_DENSE` (=1)

M - number of rows

N - number of columns

ldim - leading dimension ($\text{ldim} \geq M$)

data - pointer to a contiguous block of `realtype` variables

ldata - length of the data array ($= \text{ldim} \cdot N$). The (i,j) -th element of a dense matrix **A** of type `DlsMat` (with $0 \leq i < M$ and $0 \leq j < N$) is given by the expression `(A->data)[0][j*M+i]`

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) -th element of a dense matrix **A** of type `DlsMat` (with $0 \leq i < M$ and $0 \leq j < N$) is given by the expression `(A->cols)[j][i]`

For the BAND module, the relevant fields of this structure are as follows (see Figure 9.1 for a diagram of the underlying data representation in a banded matrix of type `DlsMat`). Note that only square band matrices are allowed.

type - `SUNDIALS_BAND` (=2)

M - number of rows

N - number of columns ($N = M$)

mu - upper half-bandwidth, $0 \leq \text{mu} < \min(M,N)$

ml - lower half-bandwidth, $0 \leq \text{ml} < \min(M,N)$

s_mu - storage upper bandwidth, $\text{mu} \leq \text{s_mu} < N$. The LU decomposition routine writes the LU factors into the storage for **A**. The upper triangular factor **U**, however, may have an upper bandwidth as big as $\min(N-1, \text{mu}+\text{ml})$ because of partial pivoting. The **s_mu** field holds the upper half-bandwidth allocated for **A**.

ldim - leading dimension ($\text{ldim} \geq \text{s_mu}$)

data - pointer to a contiguous block of `realtype` variables. The elements of a banded matrix of type `DlsMat` 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 band of **A**.

ldata - length of the data array ($= \text{ldim} \cdot (\text{s_mu} + \text{ml} + 1)$)

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 $\text{s_mu}-\text{mu}$ (to access the uppermost element within the band in the j -th column) to $\text{s_mu}+\text{ml}$ (to access the lowest element within the band in the j -th column). Indices from 0 to $\text{s_mu}-\text{mu}-1$ give access to extra storage elements required by the LU decomposition function. Finally, `cols[j][i-j+s_mu]` is the (i,j) -th element, $j-\text{mu} \leq i \leq j+\text{ml}$.

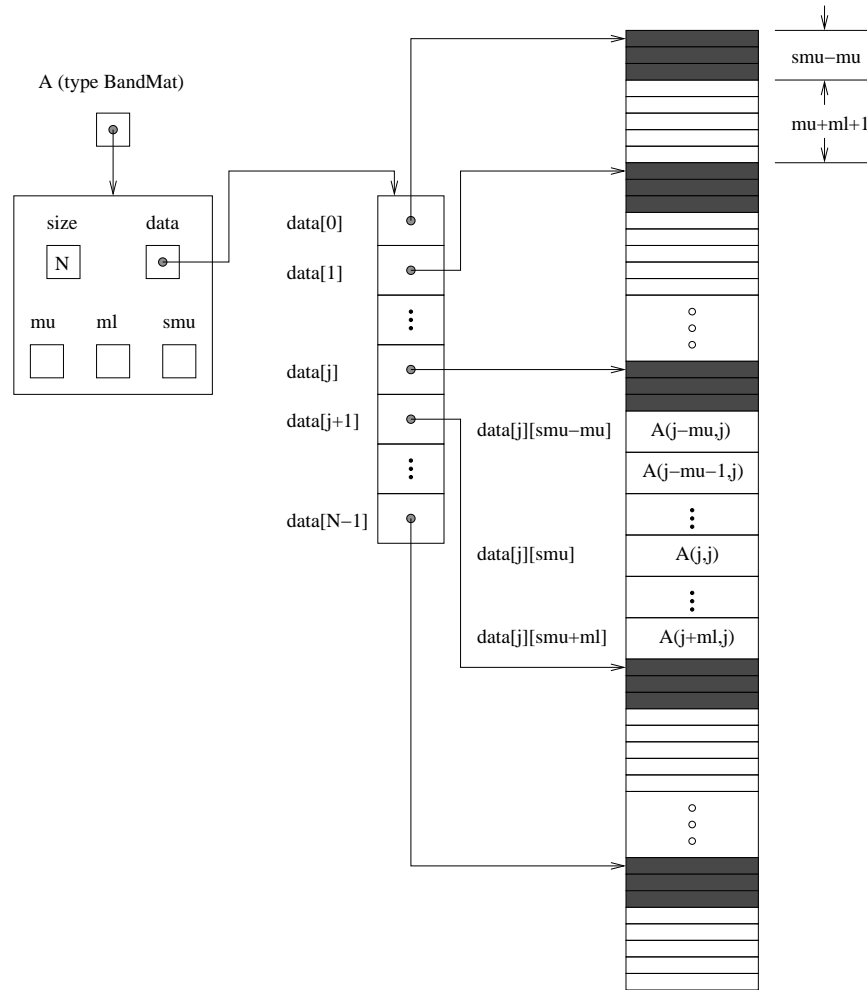


Figure 9.1: Diagram of the storage for a banded matrix of type `DlsMat`. Here A is an $N \times N$ band matrix of type `DlsMat` with upper and lower half-bandwidths μ and m , respectively. The rows and columns of A are numbered from 0 to $N - 1$ and the (i, j) -th element of A is denoted $A(i, j)$. The greyed out areas of the underlying component storage are used by the `BandGBTRF` and `BandGBTRS` routines.

9.1.2 Accessor macros for the DLS modules

The macros below allow a user to efficiently access individual matrix elements without writing out explicit data structure references and without knowing too much about the underlying element storage. The only storage assumption needed is that elements are stored columnwise and that a pointer to the j -th column of elements can be obtained via the `DENSE_COL` or `BAND_COL` macros. Users should use these macros whenever possible.

The following two macros are defined by the `DENSE` module to provide access to data in the `DlsMat` type:

- `DENSE_ELEM`

Usage : `DENSE_ELEM(A,i,j) = a_ij`; or `a_ij = DENSE_ELEM(A,i,j)`;

`DENSE_ELEM` references the (i,j) -th element of the $M \times N$ `DlsMat` `A`, $0 \leq i < M$, $0 \leq j < N$.

- `DENSE_COL`

Usage : `col_j = DENSE_COL(A,j)`;

`DENSE_COL` references the j -th column of the $M \times N$ `DlsMat` `A`, $0 \leq j < N$. The type of the expression `DENSE_COL(A,j)` is `realtype *`. After the assignment in the usage above, `col_j` may be treated as an array indexed from 0 to $M - 1$. The (i, j) -th element of `A` is referenced by `col_j[i]`.

The following three macros are defined by the `BAND` module to provide access to data in the `DlsMat` type:

- `BAND_ELEM`

Usage : `BAND_ELEM(A,i,j) = a_ij`; or `a_ij = BAND_ELEM(A,i,j)`;

`BAND_ELEM` references 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 $j - (A \rightarrow \text{mu}) \leq i \leq j + (A \rightarrow \text{ml})$.

- `BAND_COL`

Usage : `col_j = BAND_COL(A,j)`;

`BAND_COL` references 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 `BAND_COL(A,j)` is `realtype *`. The pointer returned by the call `BAND_COL(A,j)` can be treated as an array which is indexed from $-(A \rightarrow \text{mu})$ to $(A \rightarrow \text{ml})$.

- `BAND_COL_ELEM`

Usage : `BAND_COL_ELEM(col_j,i,j) = a_ij`; or `a_ij = BAND_COL_ELEM(col_j,i,j)`;

This macro references the (i,j) -th entry of the band matrix `A` when used in conjunction with `BAND_COL` to reference the j -th column through `col_j`. The index (i,j) should satisfy $j - (A \rightarrow \text{mu}) \leq i \leq j + (A \rightarrow \text{ml})$.

9.1.3 Functions in the DENSE module

The `DENSE` module defines two sets of functions with corresponding names. The first set contains functions (with names starting with a capital letter) that act on dense matrices of type `DlsMat`. The second set contains functions (with names starting with a lower case letter) that act on matrices represented as simple arrays.

The following functions for `DlsMat` dense matrices are available in the `DENSE` package. For full details, see the header files `sundials_direct.h` and `sundials_dense.h`.

- `NewDenseMat`: allocation of a `DlsMat` dense matrix;
- `DestroyMat`: free memory for a `DlsMat` matrix;

- **PrintMat**: print a **DlsMat** matrix to standard output.
- **NewLintArray**: allocation of an array of **long int** integers for use as pivots with **DenseGETRF** and **DenseGETRS**;
- **NewIntArray**: allocation of an array of **int** integers for use as pivots with the Lapack dense solvers;
- **NewRealArray**: allocation of an array of **realtype** for use as right-hand side with **DenseGETRS**;
- **DestroyArray**: free memory for an array;
- **SetToZero**: load a matrix with zeros;
- **AddIdentity**: increment a square matrix by the identity matrix;
- **DenseCopy**: copy one matrix to another;
- **DenseScale**: scale a matrix by a scalar;
- **DenseGETRF**: LU factorization with partial pivoting;
- **DenseGETRS**: solution of $Ax = b$ using LU factorization (for square matrices A);
- **DensePOTRF**: Cholesky factorization of a real symmetric positive matrix;
- **DensePOTRS**: solution of $Ax = b$ using the Cholesky factorization of A ;
- **DenseGEQRF**: QR factorization of an $m \times n$ matrix, with $m \geq n$;
- **DenseORMQR**: compute the product $w = Qv$, with Q calculated using **DenseGEQRF**;
- **DenseMatvec**: compute the product $y = Ax$, for an M by N matrix A ;

The following functions for small dense matrices are available in the DENSE package:

- **newDenseMat**
newDenseMat(m,n) allocates storage for an m by n dense matrix. It returns a pointer to the newly allocated storage if successful. If the memory request cannot be satisfied, then **newDenseMat** returns NULL. The underlying type of the dense matrix returned is **realtype****. If we allocate a dense matrix **realtype** a** by **a = newDenseMat(m,n)**, then **a[j][i]** references the (i,j) -th element of the matrix **a**, $0 \leq i < m$, $0 \leq j < n$, and **a[j]** is a pointer to the first element in the j -th column of **a**. The location **a[0]** contains a pointer to $m \times n$ contiguous locations which contain the elements of **a**.
- **destroyMat**
destroyMat(a) frees the dense matrix **a** allocated by **newDenseMat**;
- **newLintArray**
newLintArray(n) allocates an array of n integers, all **long int**. It returns a pointer to the first element in the array if successful. It returns NULL if the memory request could not be satisfied.
- **newIntArray**
newIntArray(n) allocates an array of n integers, all **int**. It returns a pointer to the first element in the array if successful. It returns NULL if the memory request could not be satisfied.
- **newRealArray**
newRealArray(n) allocates an array of n **realtype** values. It returns a pointer to the first element in the array if successful. It returns NULL if the memory request could not be satisfied.

- **destroyArray**
`destroyArray(p)` frees the array `p` allocated by `newLintArray`, `newIntArray`, or `newRealArray`;
- **denseCopy**
`denseCopy(a,b,m,n)` copies the `m` by `n` dense matrix `a` into the `m` by `n` dense matrix `b`;
- **denseScale**
`denseScale(c,a,m,n)` scales every element in the `m` by `n` dense matrix `a` by the scalar `c`;
- **denseAddIdentity**
`denseAddIdentity(a,n)` increments the *square* `n` by `n` dense matrix `a` by the identity matrix I_n ;
- **denseGETRF**
`denseGETRF(a,m,n,p)` factors the `m` by `n` dense matrix `a`, using Gaussian elimination with row pivoting. It overwrites the elements of `a` with its LU factors and keeps track of the pivot rows chosen in the pivot array `p`.
A successful LU factorization leaves the matrix `a` and the pivot array `p` with the following information:
 1. `p[k]` contains the row number of the pivot element chosen at the beginning of elimination step `k`, $k = 0, 1, \dots, n-1$.
 2. If the unique LU factorization of `a` is given by $Pa = LU$, where P is a permutation matrix, L is an `m` by `n` lower trapezoidal matrix with all diagonal elements equal to 1, and U is an `n` by `n` upper triangular matrix, then the upper triangular part of `a` (including its diagonal) contains U and the strictly lower trapezoidal part of `a` contains the multipliers, $I - L$. If `a` is square, L is a unit lower triangular matrix.`denseGETRF` returns 0 if successful. Otherwise it encountered a zero diagonal element during the factorization, indicating that the matrix `a` does not have full column rank. In this case it returns the column index (numbered from one) at which it encountered the zero.
- **denseGETRS**
`denseGETRS(a,n,p,b)` solves the `n` by `n` linear system $ax = b$. It assumes that `a` (of size $n \times n$) has been LU-factored and the pivot array `p` has been set by a successful call to `denseGETRF(a,n,n,p)`. The solution x is written into the `b` array.
- **densePOTRF**
`densePOTRF(a,m)` calculates the Cholesky decomposition of the `m` by `m` dense matrix `a`, assumed to be symmetric positive definite. Only the lower triangle of `a` is accessed and overwritten with the Cholesky factor.
- **densePOTRS**
`densePOTRS(a,m,b)` solves the `m` by `m` linear system $ax = b$. It assumes that the Cholesky factorization of `a` has been calculated in the lower triangular part of `a` by a successful call to `densePOTRF(a,m)`.
- **denseGEQRF**
`denseGEQRF(a,m,n,beta,wrk)` calculates the QR decomposition of the `m` by `n` matrix `a` ($m \geq n$) using Householder reflections. On exit, the elements on and above the diagonal of `a` contain the `n` by `n` upper triangular matrix R ; the elements below the diagonal, with the array `beta`, represent the orthogonal matrix Q as a product of elementary reflectors. The real array `wrk`, of length `m`, must be provided as temporary workspace.

- **denseORMQR**

denseORMQR(a,m,n,beta,v,w,wrk) calculates the product $w = Qv$ for a given vector **v** of length **n**, where the orthogonal matrix Q is encoded in the **m** by **n** matrix **a** and the vector **beta** of length **n**, after a successful call to **denseGEQRF(a,m,n,beta,wrk)**. The real array **wrk**, of length **m**, must be provided as temporary workspace.

- **denseMatvec**

denseMatvec(a,x,y,m,n) calculates the product $y = ax$ for a given vector **x** of length **n**, and **m** by **n** matrix **a**.

9.1.4 Functions in the BAND module

The BAND module defines two sets of functions with corresponding names. The first set contains functions (with names starting with a capital letter) that act on band matrices of type **DlsMat**. The second set contains functions (with names starting with a lower case letter) that act on matrices represented as simple arrays.

The following functions for **DlsMat** banded matrices are available in the BAND package. For full details, see the header files **sundials_direct.h** and **sundials_band.h**.

- **NewBandMat**: allocation of a **DlsMat** band matrix;
- **DestroyMat**: free memory for a **DlsMat** matrix;
- **PrintMat**: print a **DlsMat** matrix to standard output.
- **NewLintArray**: allocation of an array of **int** integers for use as pivots with **BandGBRF** and **BandGBRS**;
- **NewIntArray**: allocation of an array of **int** integers for use as pivots with the Lapack band solvers;
- **NewRealArray**: allocation of an array of **realtype** for use as right-hand side with **BandGBRS**;
- **DestroyArray**: free memory for an array;
- **SetToZero**: load a matrix with zeros;
- **AddIdentity**: increment a square matrix by the identity matrix;
- **BandCopy**: copy one matrix to another;
- **BandScale**: scale a matrix by a scalar;
- **BandGBTRF**: LU factorization with partial pivoting;
- **BandGBTRS**: solution of $Ax = b$ using LU factorization;
- **BandMatvec**: compute the product $y = Ax$, for a square band matrix A ;

The following functions for small band matrices are available in the BAND package:

- **newBandMat**
newBandMat(n, smu, ml) allocates storage for an **n** by **n** band matrix with lower half-bandwidth **ml**.
- **destroyMat**
destroyMat(a) frees the band matrix **a** allocated by **newBandMat**;

- **newLintArray**
`newLintArray(n)` allocates an array of `n` integers, all `long int`. It returns a pointer to the first element in the array if successful. It returns `NULL` if the memory request could not be satisfied.
- **newIntArray**
`newIntArray(n)` allocates an array of `n` integers, all `int`. It returns a pointer to the first element in the array if successful. It returns `NULL` if the memory request could not be satisfied.
- **newRealArray**
`newRealArray(n)` allocates an array of `n` `realtype` values. It returns a pointer to the first element in the array if successful. It returns `NULL` if the memory request could not be satisfied.
- **destroyArray**
`destroyArray(p)` frees the array `p` allocated by `newLintArray`, `newIntArray`, or `newRealArray`;
- **bandCopy**
`bandCopy(a,b,n,a_smu, b_smu,copymu, copyml)` copies the `n` by `n` band matrix `a` into the `n` by `n` band matrix `b`;
- **bandScale**
`bandScale(c,a,n,mu,ml,smu)` scales every element in the `n` by `n` band matrix `a` by `c`;
- **bandAddIdentity**
`bandAddIdentity(a,n,smu)` increments the `n` by `n` band matrix `a` by the identity matrix;
- **bandGETRF**
`bandGETRF(a,n,mu,ml,smu,p)` factors the `n` by `n` band matrix `a`, using Gaussian elimination with row pivoting. It overwrites the elements of `a` with its LU factors and keeps track of the pivot rows chosen in the pivot array `p`.
- **bandGETRS**
`bandGETRS(a,n,smu,ml,p,b)` solves the `n` by `n` linear system $ax = b$. It assumes that `a` (of size $n \times n$) has been LU-factored and the pivot array `p` has been set by a successful call to `bandGETRF(a,n,mu,ml,smu,p)`. The solution `x` is written into the `b` array.
- **bandMatvec**
`bandMatvec(a,x,y,n,mu,ml,smu)` calculates the product $y = ax$ for a given vector `x` of length `n`, and `n` by `n` band matrix `a`.

9.2 The SLS module

SUNDIALS provides a compressed-sparse-column matrix type and sparse matrix support functions. In addition, SUNDIALS provides interfaces to the publically available KLU and SuperLU_MT sparse direct solver packages. The files comprising the SLS matrix module, used in the KLU and SUPERLUMT linear solver packages, and their locations in the SUNDIALS *srcdir*, are as follows:

- header files (located in *srcdir/include/sundials*)
`sundials_sparse.h`, `sundials_klu_impl.h`,
`sundials_superlumlmt_impl.h`, `sundials_types.h`,
`sundials_math.h`, `sundials_config.h`
- source files (located in *srcdir/src/sundials*)
`sundials_sparse.c`, `sundials_math.c`

Only two of the preprocessing directives in the header file `sundials_config.h` are relevant to the SLS package by itself (see §?? for details):

- (required) definition of the precision of the SUNDIALS type `realtype`. One of the following lines must be present:


```
#define SUNDIALS_DOUBLE_PRECISION 1
#define SUNDIALS_SINGLE_PRECISION 1
#define SUNDIALS_EXTENDED_PRECISION 1
```
- (optional) use of generic math functions: `#define SUNDIALS_USE_GENERIC_MATH 1`

The `sundials_types.h` header file defines the SUNDIALS `realtype` and `booleantype` types and the macro `RCONST`, while the `sundials_math.h` header file is needed for the macros `SUNMIN` and `SUNMAX`, and the function `SUNRabs`.

9.2.1 Type SlsMat

The type `SlsMat`, defined in `sundials_sparse.h` is a pointer to a structure defining a generic compressed-sparse-column matrix, and is used with all linear solvers in the *sls* family:

```
typedef struct _SlsMat {
    int M;
    int N;
    int NNZ;
    realtype *data;
    int *rowvals;
    int *colptrs;
} *SlsMat;
```

The fields of this structure are as follows (see Figure 9.2 for a diagram of the underlying compressed-sparse-column representation in a sparse matrix of type `SlsMat`). Note that a sparse matrix of type `SlsMat` need not be square.

M - number of rows

N - number of columns

NNZ - maximum number of nonzero entries in the matrix (allocated length of `data` and `rowvals` arrays)

data - pointer to a contiguous block of `realtype` variables (of length `NNZ`), containing the values of the nonzero entries in the matrix

rowvals - pointer to a contiguous block of `int` variables (of length `NNZ`), containing the row indices of each nonzero entry held in `data`

colptrs - pointer to a contiguous block of `int` variables (of length `N+1`). Each entry provides the index of the first column entry into the `data` and `rowvals` arrays, e.g. if `colptr[3]=7`, then the first nonzero entry in the fourth column of the matrix is located in `data[7]`, and is located in row `rowvals[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 `rowvals` arrays.

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 in a `SlsMat` structure as either


```

M = 5;
N = 4;
NNZ = 8;
data = {3.0, 1.0, 3.0, 7.0, 1.0, 2.0, 9.0, 5.0};
rowvals = {1, 3, 0, 2, 0, 1, 3, 4};
colptrs = {0, 2, 4, 5, 8};

```

or

```

M = 5;
N = 4;
NNZ = 10;
data = {3.0, 1.0, 3.0, 7.0, 1.0, 2.0, 9.0, 5.0, *, *};
rowvals = {1, 3, 0, 2, 0, 1, 3, 4, *, *};
colptrs = {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 `colptrs` is 8. The work associated with operations on the sparse matrix is proportional to this value and so one should use the best understanding of the number of nonzeros here.

9.2.2 Functions in the SLS module

The SLS module defines functions that act on sparse matrices of type `SlsMat`. For full details, see the header file `sundials_sparse.h`.

- **NewSparseMat**
NewSparseMat(M, N, NNZ) allocates storage for an M by N sparse matrix, with storage for up to NNZ nonzero entries.
- **SlsConvertDls**
SlsConvertDls(A) converts a dense or band matrix A of type `DlsMat` into a new sparse matrix of type `SlsMat` by retaining only the nonzero values of the matrix A.
- **DestroySparseMat**
DestroySparseMat(A) frees the memory for a sparse matrix A allocated by either **NewSparseMat** or **SlsConvertDls**.
- **SlsSetToZero**(A) zeros out the `SlsMat` matrix A. The storage for A is left unchanged.
- **CopySparseMat**
CopySparseMat(A, B) copies the `SlsMat` A into the `SlsMat` B. It is assumed that the matrices have the same row/column dimensions. If B has insufficient storage to hold all the nonzero entries of A, the data and row index arrays in B are reallocated to match those in A.
- **ScaleSparseMat**
ScaleSparseMat(c, A) scales every element in the `SlsMat` A by the `realtype` scalar c.
- **AddIdentitySparseMat**
AddIdentitySparseMat(A) increments the `SlsMat` A by the identity matrix. If A is not square, only the existing diagonal values are incremented. Resizes the `data` and `rowvals` arrays of A to allow for new nonzero entries on the diagonal.
- **SlsAddMat**
SlsAddMat(A, B) adds two `SlsMat` matrices A and B, placing the result back in A. Resizes the `data` and `rowvals` arrays of A upon completion to exactly match the nonzero storage for the result. Upon successful completion, the return value is zero; otherwise -1 is returned.

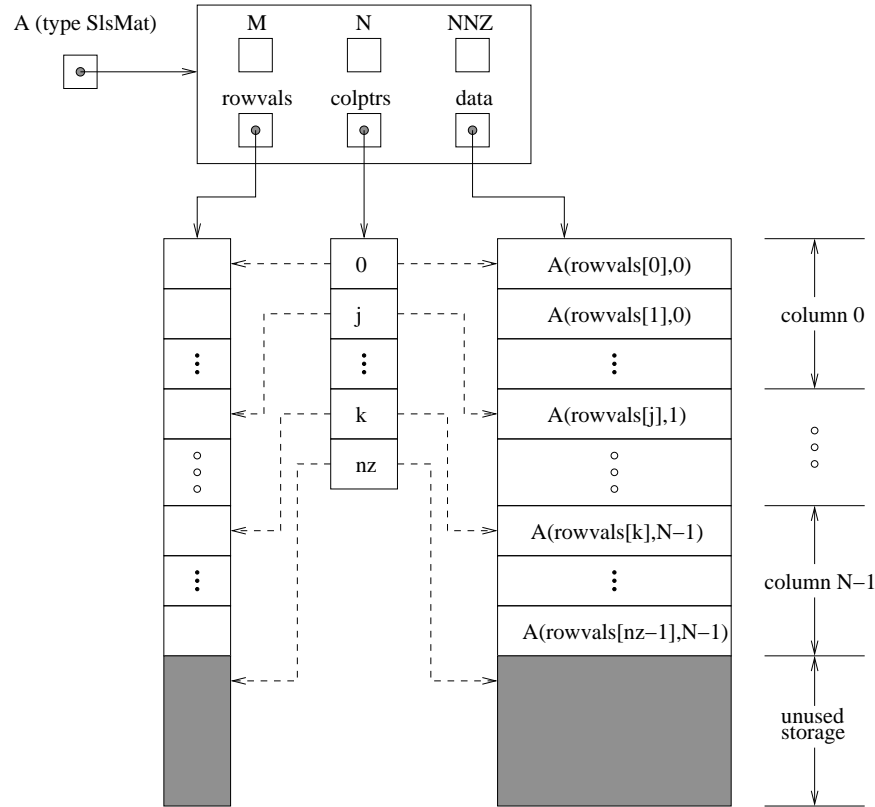


Figure 9.2: Diagram of the storage for a compressed-sparse-column matrix of type `SlsMat`. Here A is an $M \times N$ sparse matrix of type `SlsMat` with storage for up to `NNZ` nonzero entries (the allocated length of both `data` and `rowvals`). The entries in `rowvals` may assume values from 0 to $M - 1$, 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 `colptrs` array contains $N + 1$ entries; the first N denote the starting index of each column within the `rowvals` 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 `rowvals` indicate extra allocated space.

- **ReallocSparseMat**

ReallocSparseMat(A) eliminates unused storage in the **SlsMat** **A** by resizing the internal **data** and **rowvals** arrays to contain exactly **colptrs[N]** values.

- **SlsMatvec**

SlsMatvec(A, x, y) computes the sparse matrix-vector product, $y = Ax$. If the **SlsMat** **A** is a sparse matrix of dimension $M \times N$, then it is assumed that **x** is a **realtype** array of length N , and **y** is a **realtype** array of length M . Upon successful completion, the return value is zero; otherwise -1 is returned.

- **PrintSparseMat**

PrintSparseMat(A) Prints the **SlsMat** matrix **A** to standard output.

9.2.3 The KLU solver

KLU is a sparse matrix factorization and solver library written by Tim Davis [1, 14]. KLU 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). Note that SUNDIALS uses the COLAMD ordering by default with KLU.

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.

The KLU interface in SUNDIALS will perform the symbolic factorization once. It then calls the numerical factorization once and will call the refactor routine until estimates of the numerical conditioning suggest a new factorization should be completed. The KLU interface also has a **ReInit** routine that can be used to force a full refactorization at the next solver setup call.

In order to use the SUNDIALS 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 Appendix A for details).

Designed for serial calculations only, KLU is supported for calculations employing SUNDIALS' serial or shared-memory parallel NVECTOR modules (see Sections 7.1, 7.3 and 7.4 for details).

9.2.4 The SUPERLUMT solver

SUPERLUMT is a threaded sparse matrix factorization and solver library written by X. Sherry Li [2, 25, 15]. 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 SUNDIALS interface to SUPERLUMT, it is assumed that SUPERLUMT has been installed on the system prior to installation of SUNDIALS, and that SUNDIALS has been configured appropriately to link with SUPERLUMT (see Appendix A for details).

Designed for serial and threaded calculations only, SUPERLUMT is supported for calculations employing SUNDIALS' serial or shared-memory parallel NVECTOR modules (see Sections 7.1, 7.3 and 7.4 for details).

9.3 The SPILS modules: SPGMR, SPFGMR, SPBCG, and SPTFQMR

The *spils* modules contain implementations of some of the most commonly use scaled preconditioned Krylov solvers. A linear solver module from the *spils* family can only be used in conjunction with any

NVECTOR implementation library.

9.3.1 The SPGMR module

The SPGMR package, in the files `sundials_spgmr.h` and `sundials_spgmr.c`, includes an implementation of the scaled preconditioned GMRES method. A separate code module, implemented in `sundials_iterative.(h,c)`, contains auxiliary functions that support SPGMR, as well as the other Krylov solvers in SUNDIALS (SPFGMR, SPBCG, and SPTFQMR). For full details, including usage instructions, see the header files `sundials_spgmr.h` and `sundials_iterative.h`.

The files comprising the SPGMR generic linear solver, and their locations in the SUNDIALS *srcdir*, are as follows:

- header files (located in *srcdir/include/sundials*)
`sundials_spgmr.h`, `sundials_iterative.h`, `sundials_nvector.h`,
`sundials_types.h`, `sundials_math.h`, `sundials_config.h`
- source files (located in *srcdir/src/sundials*)
`sundials_spgmr.c`, `sundials_iterative.c`, `sundials_nvector.c`

Only two of the preprocessing directives in the header file `sundials_config.h` are required to use the SPGMR package by itself (see §?? for details):

- (required) definition of the precision of the SUNDIALS type `realtype`. One of the following lines must be present:

```
#define SUNDIALS_DOUBLE_PRECISION 1
#define SUNDIALS_SINGLE_PRECISION 1
#define SUNDIALS_EXTENDED_PRECISION 1
```
- (optional) use of generic math functions:

```
#define SUNDIALS_USE_GENERIC_MATH 1
```

The `sundials_types.h` header file defines the SUNDIALS `realtype` and `booleantype` types and the macro `RCONST`, while the `sundials_math.h` header file is needed for the macros `SUNMIN`, `SUNMAX`, and `SUNSQR`, and the functions `SUNRabs` and `SUNRsqr`.

The generic NVECTOR files, `sundials_nvector.(h,c)` are needed for the definition of the generic `N_Vector` type and functions. The NVECTOR functions used by the SPGMR module are: `N_VDotProd`, `N_VLinearSum`, `N_VScale`, `N_VProd`, `N_VDiv`, `N_VConst`, `N_VClone`, `N_VCloneVectorArray`, `N_VDestroy`, and `N_VDestroyVectorArray`.

The nine files listed above can be extracted from the SUNDIALS *srcdir* and compiled by themselves into an SPGMR library or into a larger user code.

The following functions are available in the SPGMR package:

- `SpgmrMalloc`: allocation of memory for `SpgmrSolve`;
- `SpgmrSolve`: solution of $Ax = b$ by the SPGMR method;
- `SpgmrFree`: free memory allocated by `SpgmrMalloc`.

The following functions are available in the support package `sundials_iterative.(h,c)`:

- `ModifiedGS`: performs modified Gram-Schmidt procedure;
- `ClassicalGS`: performs classical Gram-Schmidt procedure;
- `QRfact`: performs QR factorization of Hessenberg matrix;
- `QRsol`: solves a least squares problem with a Hessenberg matrix factored by `QRfact`.

9.3.2 The SPFGMR module

The SPFGMR package, in the files `sundials_spfgmr.h` and `sundials_spfgmr.c`, includes an implementation of the scaled preconditioned Flexible GMRES method. For full details, including usage instructions, see the file `sundials_spfgmr.h`.

The files needed to use the SPFGMR module by itself are the same as for the SPGMR module, but with `sundials_spfgmr.h(c)` in place of `sundials_spgmr.h(c)`.

The following functions are available in the SPFGMR package:

- `SpfgmrMalloc`: allocation of memory for `SpfgmrSolve`;
- `SpfgmrSolve`: solution of $Ax = b$ by the SPFGMR method;
- `SpfgmrFree`: free memory allocated by `SpfgmrMalloc`.

9.3.3 The SPBCG module

The SPBCG package, in the files `sundials_spgbcs.h` and `sundials_spgbcs.c`, includes an implementation of the scaled preconditioned Bi-CGStab method. For full details, including usage instructions, see the file `sundials_spgbcs.h`.

The files needed to use the SPBCG module by itself are the same as for the SPGMR module, but with `sundials_spgbcs.h(c)` in place of `sundials_spgmr.h(c)`.

The following functions are available in the SPBCG package:

- `SpbcgMalloc`: allocation of memory for `SpbcgSolve`;
- `SpbcgSolve`: solution of $Ax = b$ by the SPBCG method;
- `SpbcgFree`: free memory allocated by `SpbcgMalloc`.

9.3.4 The SPTFQMR module

The SPTFQMR package, in the files `sundials_sptfqmr.h` and `sundials_sptfqmr.c`, includes an implementation of the scaled preconditioned TFQMR method. For full details, including usage instructions, see the file `sundials_sptfqmr.h`.

The files needed to use the SPTFQMR module by itself are the same as for the SPGMR module, but with `sundials_sptfqmr.h(c)` in place of `sundials_spgmr.h(c)`.

The following functions are available in the SPTFQMR package:

- `SptfqmrMalloc`: allocation of memory for `SptfqmrSolve`;
- `SptfqmrSolve`: solution of $Ax = b$ by the SPTFQMR method;
- `SptfqmrFree`: free memory allocated by `SptfqmrMalloc`.

Appendix A

SUNDIALS Package Installation Procedure

The installation of any SUNDIALS package is accomplished by installing the SUNDIALS suite as a whole, according to the instructions that follow. The same procedure applies whether or not the downloaded file contains one or all solvers in SUNDIALS.

The SUNDIALS suite (or individual solvers) are distributed as compressed archives (`.tar.gz`). The name of the distribution archive is of the form `solver-x.y.z.tar.gz`, where *solver* is one of: `sundials`, `cvode`, `cvodes`, `ida`, `idas`, or `kinsol`, and `x.y.z` represents the version number (of the SUNDIALS suite or of the individual solver). To begin the installation, first uncompress and expand the sources, by issuing

```
% tar xzf 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 on the installation procedure begins with a few common observations:

- The remainder of this chapter will follow these conventions:

srcdir 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.

installdir is the directory under which the SUNDIALS exported header files and libraries will be installed. Typically, header files are exported under a directory `installdir/include` while libraries are installed under `installdir/lib`, with *installdir* specified at configuration time.

- For SUNDIALS CMake-based installation, in-source builds are prohibited; in other words, the build directory *builddir* can **not** be the same as *srcdir* 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.
- The installation directory *installdir* can **not** be the same as the source directory *srcdir*.
- By default, only the libraries and header files are exported to the installation directory *installdir*. 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) makefiles. 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.)

- Even if generation of shared libraries is enabled, only static libraries are created for the FCMIX modules. (Because of the use of fixed names for the Fortran user-provided subroutines, FCMIX shared libraries would result in "undefined symbol" errors at link time.)

A.1 CMake-based installation

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 2.8.1 or higher and a working compiler. On Unix-like operating systems, it also requires Make (and **curses**, including its development libraries, for the GUI front end to CMake, **ccmake**), while on Windows it requires Visual Studio. While many Linux distributions offer CMake, the version included is probably out of date. Many new CMake features have been added recently, 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**, 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.

A.1.1 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 *installdir* 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. Examples for using both 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 (...)sundials/installdir
% mkdir (...)sundials/builddir
% cd (...)sundials/builddir
```

Building with the GUI

Using CMake with the GUI follows this general process:

- Select and modify values, run configure (c key)
- New values are denoted with an asterisk
- 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
- Repeat until all values are set as desired and the generate option is available (g key)
- Some variables (advanced variables) are not visible right away
- To see advanced variables, toggle to advanced mode (t key)
- To search for a variable press / key, and to repeat the search, press the n key

To build the default configuration using the GUI, from the *builddir* enter the `ccmake` command and point to the *sourcedir*:

```
% ccmake ../sourcedir
```

The default configuration screen is shown in Figure A.1.

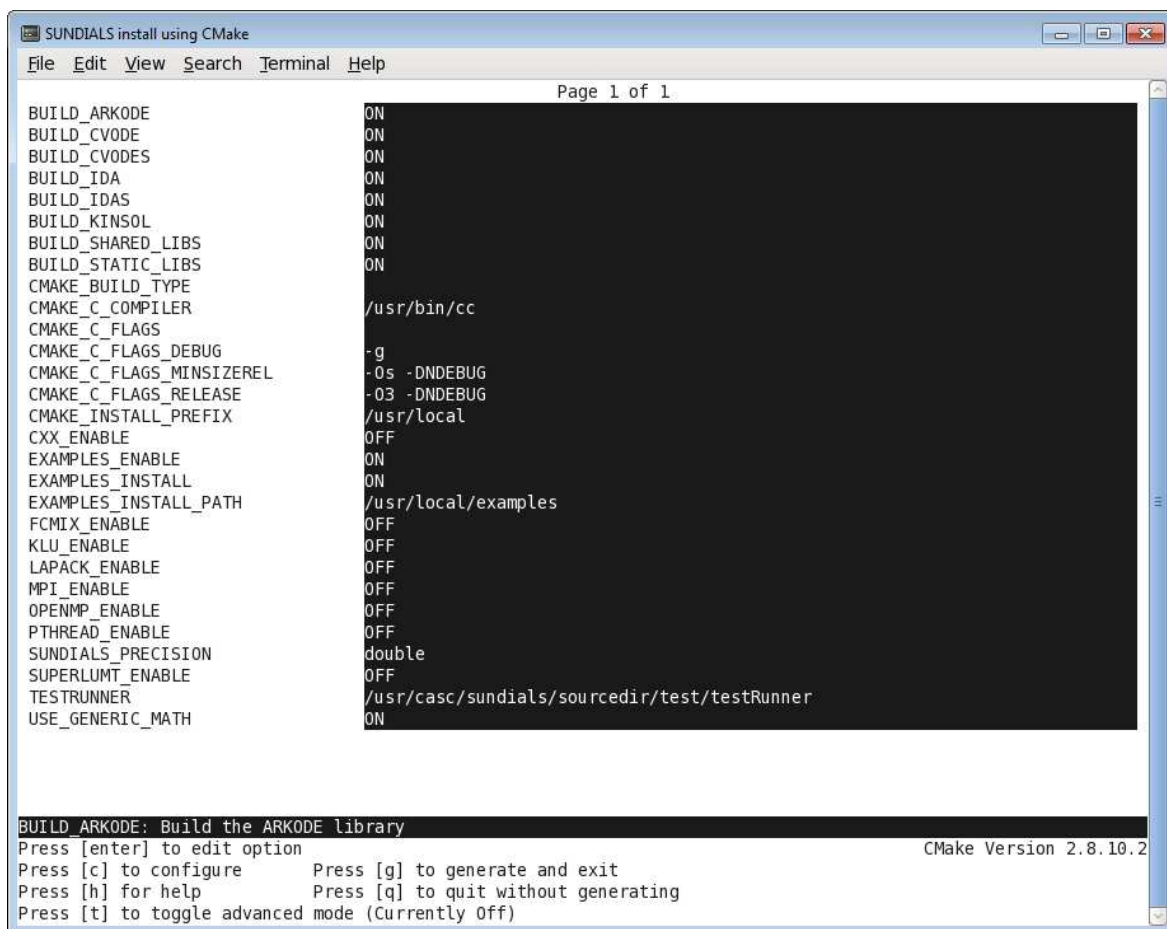


Figure A.1: Default configuration screen. Note: Initial screen is empty. To get this default configuration, press 'c' repeatedly (accepting default values denoted with asterisk) until the 'g' option is available.

The default *installdir* for both SUNDIALS and corresponding examples can be changed by setting the `CMAKE_INSTALL_PREFIX` and the `EXAMPLES_INSTALL_PATH` as shown in figure A.2.

Pressing the (g key) 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
```

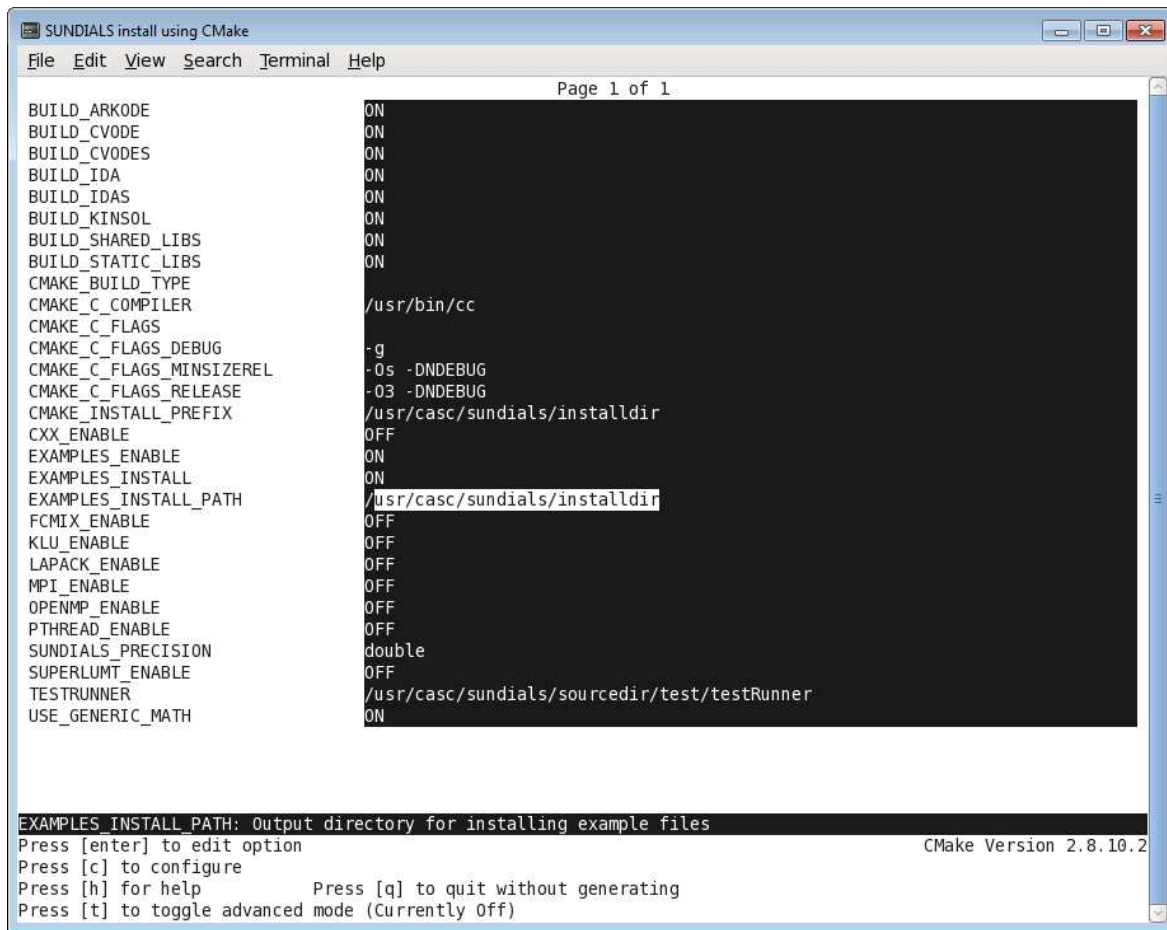


Figure A.2: Changing the *installdir* for SUNDIALS and corresponding examples

To install SUNDIALS in the installation directory specified in the configuration, simply run:

```
% make install
```

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/installdir \
> -DEXAMPLES_INSTALL_PATH=/home/myname/sundials/installdir/examples \
> ../sourcedir
% make
% make test
  
```

A.1.2 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_CVODE - Build the CVODE library
Default: ON

BUILD_CVODES - Build the CVODES library
Default: ON

BUILD_IDA - Build the IDA library
Default: ON

BUILD_IDAS - Build the IDAS library
Default: ON

BUILD_KINSOL - Build the KINSOL library
Default: ON

BUILD_SHARED_LIBS - Build shared libraries
Default: OFF

BUILD_STATIC_LIBS - Build static libraries
Default: ON

CMAKE_BUILD_TYPE - Choose the type of build, options are: None (CMAKE_C_FLAGS used) Debug Release RelWithDebInfo MinSizeRel
Default:

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 compiler during debug builds
Default: -g

CMAKE_C_FLAGS_MINSIZEREL - Flags used by the compiler during release minsize builds
Default: -Os -DNDEBUG

CMAKE_C_FLAGS_RELEASE - Flags used by the compiler during release builds
Default: -O3 -DNDEBUG

CMAKE_Fortran_COMPILER - Fortran compiler
Default: /usr/bin/gfortran
Note: Fortran support (and all related options) are triggered only if either Fortran-C support is enabled (FCMIX_ENABLE is ON) or Blas/Lapack support is enabled (LAPACK_ENABLE is ON).

CMAKE_Fortran_FLAGS - Flags for Fortran compiler
Default:

CMAKE_Fortran_FLAGS_DEBUG - Flags used by the compiler during debug builds
Default:

CMAKE_Fortran_FLAGS_MINSIZEREL - Flags used by the compiler during release minsize builds
Default:

CMAKE_Fortran_FLAGS_RELEASE - Flags used by the compiler during release builds
Default:

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.

EXAMPLES_ENABLE - Build the SUNDIALS examples

Default: ON

EXAMPLES_INSTALL - Install example files

Default: ON

Note: This option is triggered only if building example programs is enabled (**EXAMPLES_ENABLE** 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 have an **examples** subdirectory created under **CMAKE_INSTALL_PREFIX**.

FCMIX_ENABLE - Enable Fortran-C support

Default: OFF

KLU_ENABLE - Enable KLU support

Default: OFF

LAPACK_ENABLE - Enable Lapack support

Default: OFF

Note: Setting this option to ON will trigger the two additional options see below.

LAPACK_LIBRARIES - Lapack (and Blas) libraries

Default: /usr/lib/liblapack.so;/usr/lib/libblas.so

Note: CMake will search for these libraries in your **LD_LIBRARY_PATH** prior to searching default system paths.

MPI_ENABLE - Enable MPI support

Default: OFF

Note: Setting this option to ON will trigger several additional options related to MPI.

MPI_MPICC - mpicc program

Default:

MPI_RUN_COMMAND - Specify run command for MPI

Default: mpirun

Note: This can either be set to **mpirun** for OpenMPI or **srun** if jobs are managed by SLURM - Simple Linux Utility for Resource Management as exists on LLNL's high performance computing clusters.

MPI_MPIF77 - mpif77 program

Default:

Note: This option is triggered only if using MPI compiler scripts (**MPI_USE_MPISCRIPITS** is ON) and Fortran-C support is enabled (**FCMIX_ENABLE** is ON).

OPENMP_ENABLE - Enable OpenMP support

Default: OFF

Turn on support for the OpenMP based nvector.

PTHREAD_ENABLE - Enable Pthreads support

Default: OFF

Turn on support for the Pthreads based nvector.

SUNDIALS_PRECISION - Precision used in SUNDIALS, options are: double, single or extended

Default: double

SUPERLUMT_ENABLE - Enable SUPERLU_MT support

Default: OFF

TESTRUNNER - Location of testRunner script

Default: sourcedir/testRunner

USE_GENERIC_MATH - Use generic (stdc) math libraries

Default: ON

A.1.3 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 mpif77 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/installdir \
> -DEXAMPLES_INSTALL_PATH=/home/myname/sundials/installdir/examples \
> -DMPI_ENABLE=ON \
> -DFCMIX_ENABLE=ON \
> /home/myname/sundials/sourcedir
%
% make install
%
```

To disable installation of the examples, use:

```
% cmake \
> -DCMAKE_INSTALL_PREFIX=/home/myname/sundials/installdir \
> -DEXAMPLES_INSTALL_PATH=/home/myname/sundials/installdir/examples \
> -DMPI_ENABLE=ON \
> -DFCMIX_ENABLE=ON \
> -DEXAMPLES_INSTALL=OFF \
> /home/myname/sundials/sourcedir
%
% make install
%
```

A.1.4 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.

Building with LAPACK and BLAS

To enable LAPACK and BLAS libraries, set the LAPACK_ENABLE option to ON. If the directory containing the LAPACK and BLAS libraries is in the LD_LIBRARY_PATH environment variable, CMake will set the LAPACK_LIBRARIES variable accordingly, otherwise CMake will attempt to find the LAPACK

and BLAS libraries in standard system locations. To explicitly tell CMake what libraries to use, the `LAPACK_LIBRARIES` variable can be set to the desired libraries. Example:

```
% cmake \
> -DCMAKE_INSTALL_PREFIX=/home/myname/sundials/installdir \
> -DEXAMPLES_INSTALL_PATH=/home/myname/sundials/installdir/examples \
> -DLAPACK_LIBRARIES=/mypath/lib/liblapack.so;/mypath/lib/libblas.so \
> /home/myname/sundials/sourcedir
%
% make install
%
```

Building with KLU

The KLU libraries are part of SuiteSparse, a suite of sparse matrix software, available from the Texas A&M University website: <http://faculty.cse.tamu.edu/davis/suitesparse.html>. SUNDIALS has been tested with SuiteSparse version 4.2.1. To enable KLU, set `KLU_ENABLE` 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. The CMake configure will result in populating the variables: `AMD_LIBRARY`, `AMD_LIBRARY_DIR`, `BTF_LIBRARY`, `BTF_LIBRARY_DIR`, `COLAMD_LIBRARY`, `COLAMD_LIBRARY_DIR`, and `KLU_LIBRARY`.

Building with SuperLU_MT

The SuperLU_MT libraries are available for download from the Lawrence Berkeley National Laboratory website: http://crd-legacy.lbl.gov/~xiaoye/SuperLU/#superlu_mt. SUNDIALS has been tested with SuperLU_MT version 2.4. To enable SuperLU_MT, set `SUPERLUMT_ENABLE` to `ON`, set `SUPERLUMT_INCLUDE_DIR` to the `SRC` path of the SuperLU_MT installation and set `SUPERLUMT_LIBRARY_DIR` to the `lib` path of the SuperLU_MT installation. Also, the `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 `OPENMP_ENABLE` or `PTHREAD_ENABLE` set to `ON` then SuperLU_MT should be set to use the same threading type.



A.2 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 both `EXAMPLES_ENABLE` and `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.

From within the installed example directory, run CMake (either with the GUI or command line) to compile the example code. The resulting output from running the example can be compared with example output bundled in the SUNDIALS distribution.

NOTE: There will likely be differences in the output due to machine architecture, compiler versions, use of third party libraries etc.

These examples with the installed configuration files can be used as "templates" for user developed solutions.



A.3 Configuring, building, and installing on Windows

Use `CMakeSetup` from the CMake install location. Make sure to select the appropriate source and the build directory. Also, make sure to pick the appropriate generator (on Visual Studio 6, pick the

Visual Studio 6 generator). Some CMake versions will ask you to select the generator the first time you press Configure instead of having a drop-down menu in the main dialog.

CMake will now create Visual Studio project files. You should now be able to open the SUNDIALS project (or workspace) file. Make sure to select the appropriate build type (Debug, Release, ...). To build SUNDIALS, simply build the `ALL_BUILD` target. To install SUNDIALS, simply run the `INSTALL` target within the build system.

A.4 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 default values for these directories are *instdir/lib* and *instdir/include*, respectively, but can be changed using the configure script options `--prefix`, `--exec-prefix`, `--includedir` and `--libdir` (see the appropriate CMake options). For example, a global installation of SUNDIALS on a *NIX system could be accomplished using

```
% configure --prefix=/opt/sundials-2.1.1
```

Although all installed libraries reside under *libdir*, the public header files are further organized into subdirectories under *includedir*.

The installed libraries and exported header files are listed for reference in Tables [A.1](#) and [A.2](#). The file extension *.lib* is typically *.so* for shared libraries and *.a* for static libraries. Note that, in the Tables, names are relative to *libdir* for libraries and to *includedir* for header files.

A typical user program need not explicitly include any of the shared SUNDIALS header files from under the *includedir/sundials* directory since they are explicitly included by the appropriate solver header files (*e.g.*, `cvode_dense.h` includes `sundials_dense.h`). However, it is both legal and safe to do so (*e.g.*, the functions declared in `sundials_dense.h` could be used in building a preconditioner).

Table A.1: SUNDIALS libraries and header files

SHARED	Libraries	n/a	
	Header files	sundials/sundials_config.h sundials/sundials_math.h sundials/sundials_nvector.h sundials/sundials_direct.h sundials/sundials_dense.h sundials/sundials_sparse.h sundials/sundials_iterative.h sundials/sundials_spgmr.h sundials/sundials_spgmrs.h sundials/sundials_spgmrs.h	sundials/sundials_types.h sundials/sundials_fnvector.h sundials/sundials_lapack.h sundials/sundials_band.h sundials/sundials_spgmr.h sundials/sundials_sptfqmr.h sundials/sundials_spgmrs.h
NVECTOR_SERIAL	Libraries	libsundials_nvecserial. <i>lib</i>	libsundials_nvecserial.a
	Header files	nvector/nvector_serial.h	
NVECTOR_PARALLEL	Libraries	libsundials_nvecparallel. <i>lib</i>	libsundials_nvecparallel.a
	Header files	nvector/nvector_parallel.h	
NVECTOR_OPENMP	Libraries	libsundials_nvecopenmp. <i>lib</i>	libsundials_nvecopenmp.a
	Header files	nvector/nvector_openmp.h	
NVECTOR_PTHREADS	Libraries	libsundials_nvecpthreads. <i>lib</i>	libsundials_nvecpthreads.a
	Header files	nvector/nvector_pthreads.h	
CVMODE	Libraries	libsundials_cvmode. <i>lib</i>	libsundials_cvmode.a
	Header files	cvmode/cvmode.h cvmode/cvmode_direct.h cvmode/cvmode_dense.h cvmode/cvmode_diag.h cvmode/cvmode_sparse.h cvmode/cvmode_superlumt.h cvmode/cvmode_spils.h cvmode/cvmode_sptfqmr.h cvmode/cvmode_bandpre.h	cvmode/cvmode_impl.h cvmode/cvmode_lapack.h cvmode/cvmode_band.h cvmode/cvmode_klu.h cvmode/cvmode_spgmr.h cvmode/cvmode_spgmrs.h cvmode/cvmode_spgmrs.h cvmode/cvmode_spgmrs.h
CVMODES	Libraries	libsundials_cvmodes. <i>lib</i>	
	Header files	cvmodes/cvmodes.h cvmodes/cvmodes_direct.h cvmodes/cvmodes_dense.h cvmodes/cvmodes_diag.h cvmodes/cvmodes_sparse.h cvmodes/cvmodes_superlumt.h cvmodes/cvmodes_spils.h cvmodes/cvmodes_sptfqmr.h cvmodes/cvmodes_bandpre.h	cvmodes/cvmodes_impl.h cvmodes/cvmodes_lapack.h cvmodes/cvmodes_band.h cvmodes/cvmodes_klu.h cvmodes/cvmodes_spgmr.h cvmodes/cvmodes_spgmrs.h cvmodes/cvmodes_spgmrs.h cvmodes/cvmodes_spgmrs.h
ARKODE	Libraries	libsundials_arkode. <i>lib</i>	libsundials_farkode.a
	Header files	arkode/arkode.h arkode/arkode_direct.h arkode/arkode_dense.h arkode/arkode_sparse.h arkode/arkode_superlumt.h arkode/arkode_spils.h arkode/arkode_sptfqmr.h arkode/arkode_pcg.h arkode/arkode_bandpre.h	arkode/arkode_impl.h arkode/arkode_lapack.h arkode/arkode_band.h arkode/arkode_klu.h arkode/arkode_spgmr.h arkode/arkode_spgmrs.h arkode/arkode_spgmrs.h arkode/arkode_spgmrs.h arkode/arkode_spgmrs.h

Table A.2: SUNDIALS libraries and header files (cont.)

IDA	Libraries	libsundials_ida. <i>lib</i>	libsundials_fida.a
	Header files	ida/ida.h ida/ida_direct.h ida/ida_dense.h ida/ida_sparse.h ida/ida_superlumt.h ida/ida_spils.h ida/ida_spgmr.h ida/ida_sptfqmr.h	ida/ida_impl.h ida/ida_lapack.h ida/ida_band.h ida/ida_klu.h ida/ida_spgmr.h ida/ida_sptfqmr.h
IDAS	Libraries	libsundials_idas. <i>lib</i>	
	Header files	idas/idas.h idas/idas_direct.h idas/idas_dense.h idas/idas_sparse.h idas/idas_superlumt.h idas/idas_spils.h idas/idas_spgmr.h idas/idas_sptfqmr.h	idas/idas_impl.h idas/idas_lapack.h idas/idas_band.h idas/idas_klu.h idas/idas_spgmr.h idas/idas_sptfqmr.h
KINSOL	Libraries	libsundials_kinsol. <i>lib</i>	libsundials_fkinsol.a
	Header files	kinsol/kinsol.h kinsol/kinsol_direct.h kinsol/kinsol_dense.h kinsol/kinsol_sparse.h kinsol/kinsol_superlumt.h kinsol/kinsol_spils.h kinsol/kinsol_spgmr.h kinsol/kinsol_sptfqmr.h	kinsol/kinsol_impl.h kinsol/kinsol_lapack.h kinsol/kinsol_band.h kinsol/kinsol_klu.h kinsol/kinsol_spgmr.h kinsol/kinsol_sptfqmr.h kinsol/kinsol_sptfqmr.h

Appendix B

IDAS Constants

Below we list all input and output constants used by the main solver and linear solver modules, together with their numerical values and a short description of their meaning.

B.1 IDAS input constants

IDAS main solver module		
IDA_NORMAL	1	Solver returns at specified output time.
IDA_ONE_STEP	2	Solver returns after each successful step.
IDA_SIMULTANEOUS	1	Simultaneous corrector forward sensitivity method.
IDA_STAGGERED	2	Staggered corrector forward sensitivity method.
IDA_CENTERED	1	Central difference quotient approximation (2^{nd} order) of the sensitivity RHS.
IDA_FORWARD	2	Forward difference quotient approximation (1^{st} order) of the sensitivity RHS.
IDA_YA_YDP_INIT	1	Compute y_a and \dot{y}_d , given y_d .
IDA_Y_INIT	2	Compute y , given \dot{y} .
IDAS adjoint solver module		
IDA_HERMITE	1	Use Hermite interpolation.
IDA_POLYNOMIAL	2	Use variable-degree polynomial interpolation.
Iterative linear solver module		
PREC_NONE	0	No preconditioning
PREC_LEFT	1	Preconditioning on the left.
MODIFIED_GS	1	Use modified Gram-Schmidt procedure.
CLASSICAL_GS	2	Use classical Gram-Schmidt procedure.

B.2 IDAS output constants

IDAS main solver module		
IDA_SUCCESS	0	Successful function return.
IDA_TSTOP_RETURN	1	IDASolve succeeded by reaching the specified stopping point.

IDA_ROOT_RETURN	2	IDASolve succeeded and found one or more roots.
IDA_WARNING	99	IDASolve succeeded but an unusual situation occurred.
IDA_TOO_MUCH_WORK	-1	The solver took <code>mxstep</code> internal steps but could not reach tout.
IDA_TOO_MUCH_ACC	-2	The solver could not satisfy the accuracy demanded by the user for some internal step.
IDA_ERR_FAIL	-3	Error test failures occurred too many times during one internal time step or minimum step size was reached.
IDA_CONV_FAIL	-4	Convergence test failures occurred too many times during one internal time step or minimum step size was reached.
IDA_LINIT_FAIL	-5	The linear solver's initialization function failed.
IDA_LSETUP_FAIL	-6	The linear solver's setup function failed in an unrecoverable manner.
IDA_LSOLVE_FAIL	-7	The linear solver's solve function failed in an unrecoverable manner.
IDA_RES_FAIL	-8	The user-provided residual function failed in an unrecoverable manner.
IDA_REP_RES_FAIL	-9	The user-provided residual function repeatedly returned a recoverable error flag, but the solver was unable to recover.
IDA_RTFUNC_FAIL	-10	The rootfinding function failed in an unrecoverable manner.
IDA_CONSTR_FAIL	-11	The inequality constraints were violated and the solver was unable to recover.
IDA_FIRST_RES_FAIL	-12	The user-provided residual function failed recoverably on the first call.
IDA_LINESEARCH_FAIL	-13	The line search failed.
IDA_NO_RECOVERY	-14	The residual function, linear solver setup function, or linear solver solve function had a recoverable failure, but IDACalcIC could not recover.
IDA_MEM_NULL	-20	The <code>ida_mem</code> argument was NULL.
IDA_MEM_FAIL	-21	A memory allocation failed.
IDA_ILL_INPUT	-22	One of the function inputs is illegal.
IDA_NO_MALLOC	-23	The IDAS memory was not allocated by a call to IDAInit.
IDA_BAD_EWT	-24	Zero value of some error weight component.
IDA_BAD_K	-25	The k -th derivative is not available.
IDA_BAD_T	-26	The time t is outside the last step taken.
IDA_BAD_DKY	-27	The vector argument where derivative should be stored is NULL.
IDA_NO_QUAD	-30	Quadratures were not initialized.
IDA_QRHS_FAIL	-31	The user-provided right-hand side function for quadratures failed in an unrecoverable manner.
IDA_FIRST_QRHS_ERR	-32	The user-provided right-hand side function for quadratures failed in an unrecoverable manner on the first call.
IDA_REP_QRHS_ERR	-33	The user-provided right-hand side repeatedly returned a recoverable error flag, but the solver was unable to recover.
IDA_NO_SENS	-40	Sensitivities were not initialized.
IDA_SRES_FAIL	-41	The user-provided sensitivity residual function failed in an unrecoverable manner.
IDA_REP_SRES_ERR	-42	The user-provided sensitivity residual function repeatedly returned a recoverable error flag, but the solver was unable to recover.

IDA_BAD_IS	-43	The sensitivity identifier is not valid.
IDA_NO_QUADSENS	-50	Sensitivity-dependent quadratures were not initialized.
IDA_QSRHS_FAIL	-51	The user-provided sensitivity-dependent quadrature right-hand side function failed in an unrecoverable manner.
IDA_FIRST_QSRHS_ERR	-52	The user-provided sensitivity-dependent quadrature right-hand side function failed in an unrecoverable manner on the first call.
IDA_REP_QSRHS_ERR	-53	The user-provided sensitivity-dependent quadrature right-hand side repeatedly returned a recoverable error flag, but the solver was unable to recover.
<hr/> IDAS adjoint solver module <hr/>		
IDA_NO_ADJ	-101	The combined forward-backward problem has not been initialized.
IDA_NO_FWD	-102	IDASolveF has not been previously called.
IDA_NO_BCK	-103	No backward problem was specified.
IDA_BAD_TBO	-104	The desired output for backward problem is outside the interval over which the forward problem was solved.
IDA_REIFWD_FAIL	-105	No checkpoint is available for this hot start.
IDA_FWD_FAIL	-106	IDASolveB failed because IDASolve was unable to store data between two consecutive checkpoints.
IDA_GETY_BADT	-107	Wrong time in interpolation function.
<hr/> IDADLS linear solver modules <hr/>		
IDADLS_SUCCESS	0	Successful function return.
IDADLS_MEM_NULL	-1	The <code>ida_mem</code> argument was NULL.
IDADLS_LMEM_NULL	-2	The IDADLS linear solver has not been initialized.
IDADLS_ILL_INPUT	-3	The IDADLS solver is not compatible with the current NVECTOR module.
IDADLS_MEM_FAIL	-4	A memory allocation request failed.
IDADLS_JACFUNC_UNRECVR	-5	The Jacobian function failed in an unrecoverable manner.
IDADLS_JACFUNC_RECVR	-6	The Jacobian function had a recoverable error.
IDADLS_NO_ADJ	-101	The combined forward-backward problem has not been initialized.
IDADLS_LMEMB_NULL	-102	The linear solver was not initialized for the backward phase.
<hr/> IDASLS linear solver module <hr/>		
IDASLS_SUCCESS	0	Successful function return.
IDASLS_MEM_NULL	-1	The <code>ida_mem</code> argument was NULL.
IDASLS_LMEM_NULL	-2	The IDASLS linear solver has not been initialized.
IDASLS_ILL_INPUT	-3	The IDASLS solver is not compatible with the current NVECTOR module or other input is invalid.
IDASLS_MEM_FAIL	-4	A memory allocation request failed.
IDASLS_JAC_NOSET	-5	The Jacobian evaluation routine was not been set before the linear solver setup routine was called.

IDASLS_PACKAGE_FAIL	-6	An external package call return a failure error code.
IDASLS_JACFUNC_UNRECV	-7	The Jacobian function failed in an unrecoverable manner.
IDASLS_JACFUNC_RECVR	-8	The Jacobian function had a recoverable error.
IDASLS_NO_ADJ	-101	The combined forward-backward problem has not been initialized.
IDASLS_LMEMB_NULL	-102	The linear solver was not initialized for the backward phase.
<hr/> IDASPILS linear solver modules <hr/>		
IDASPILS_SUCCESS	0	Successful function return.
IDASPILS_MEM_NULL	-1	The <code>ida_mem</code> argument was NULL.
IDASPILS_LMEM_NULL	-2	The IDASPILS linear solver has not been initialized.
IDASPILS_ILL_INPUT	-3	The IDASPILS solver is not compatible with the current NVECTOR module.
IDASPILS_MEM_FAIL	-4	A memory allocation request failed.
IDASPILS_PMEM_NULL	-5	The preconditioner module has not been initialized.
IDASPILS_NO_ADJ	-101	The combined forward-backward problem has not been initialized.
IDASPILS_LMEMB_NULL	-102	The linear solver was not initialized for the backward phase.
<hr/> SPGMR generic linear solver module <hr/>		
SPGMR_SUCCESS	0	Converged.
SPGMR_RES_REDUCED	1	No convergence, but the residual norm was reduced.
SPGMR_CONV_FAIL	2	Failure to converge.
SPGMR_QRFACT_FAIL	3	A singular matrix was found during the QR factorization.
SPGMR_PSOLVE_FAIL_REC	4	The preconditioner solve function failed recoverably.
SPGMR_ATIMES_FAIL_REC	5	The Jacobian-times-vector function failed recoverably.
SPGMR_PSET_FAIL_REC	6	The preconditioner setup routine failed recoverably.
SPGMR_MEM_NULL	-1	The SPGMR memory is NULL
SPGMR_ATIMES_FAIL_UNREC	-2	The Jacobian-times-vector function failed unrecoverably.
SPGMR_PSOLVE_FAIL_UNREC	-3	The preconditioner solve function failed unrecoverably.
SPGMR_GS_FAIL	-4	Failure in the Gram-Schmidt procedure.
SPGMR_QRSOL_FAIL	-5	The matrix R was found to be singular during the QR solve phase.
SPGMR_PSET_FAIL_UNREC	-6	The preconditioner setup routine failed unrecoverably.
<hr/> SPFGMR generic linear solver module (only available in KINSOL and ARKODE) <hr/>		
SPFGMR_SUCCESS	0	Converged.
SPFGMR_RES_REDUCED	1	No convergence, but the residual norm was reduced.
SPFGMR_CONV_FAIL	2	Failure to converge.
SPFGMR_QRFACT_FAIL	3	A singular matrix was found during the QR factorization.
SPFGMR_PSOLVE_FAIL_REC	4	The preconditioner solve function failed recoverably.
SPFGMR_ATIMES_FAIL_REC	5	The Jacobian-times-vector function failed recoverably.
SPFGMR_PSET_FAIL_REC	6	The preconditioner setup routine failed recoverably.
SPFGMR_MEM_NULL	-1	The SPFGMR memory is NULL
SPFGMR_ATIMES_FAIL_UNREC	-2	The Jacobian-times-vector function failed unrecoverably.

SPFGMR_PSOLVE_FAIL_UNREC	-3	The preconditioner solve function failed unrecoverably.
SPFGMR_GS_FAIL	-4	Failure in the Gram-Schmidt procedure.
SPFGMR_QRSOL_FAIL	-5	The matrix R was found to be singular during the QR solve phase.
SPFGMR_PSET_FAIL_UNREC	-6	The preconditioner setup routine failed unrecoverably.
<hr/>		
SPBCG generic linear solver module		
<hr/>		
SPBCG_SUCCESS	0	Converged.
SPBCG_RES_REDUCED	1	No convergence, but the residual norm was reduced.
SPBCG_CONV_FAIL	2	Failure to converge.
SPBCG_PSOLVE_FAIL_REC	3	The preconditioner solve function failed recoverably.
SPBCG_ATIMES_FAIL_REC	4	The Jacobian-times-vector function failed recoverably.
SPBCG_PSET_FAIL_REC	5	The preconditioner setup routine failed recoverably.
SPBCG_MEM_NULL	-1	The SPBCG memory is NULL
SPBCG_ATIMES_FAIL_UNREC	-2	The Jacobian-times-vector function failed unrecoverably.
SPBCG_PSOLVE_FAIL_UNREC	-3	The preconditioner solve function failed unrecoverably.
SPBCG_PSET_FAIL_UNREC	-4	The preconditioner setup routine failed unrecoverably.
<hr/>		
SPTFQMR generic linear solver module		
<hr/>		
SPTFQMR_SUCCESS	0	Converged.
SPTFQMR_RES_REDUCED	1	No convergence, but the residual norm was reduced.
SPTFQMR_CONV_FAIL	2	Failure to converge.
SPTFQMR_PSOLVE_FAIL_REC	3	The preconditioner solve function failed recoverably.
SPTFQMR_ATIMES_FAIL_REC	4	The Jacobian-times-vector function failed recoverably.
SPTFQMR_PSET_FAIL_REC	5	The preconditioner setup routine failed recoverably.
SPTFQMR_MEM_NULL	-1	The SPTFQMR memory is NULL
SPTFQMR_ATIMES_FAIL_UNREC	-2	The Jacobian-times-vector function failed.
SPTFQMR_PSOLVE_FAIL_UNREC	-3	The preconditioner solve function failed unrecoverably.
SPTFQMR_PSET_FAIL_UNREC	-4	The preconditioner setup routine failed unrecoverably.

Bibliography

- [1] KLU Sparse Matrix Factorization Library. <http://faculty.cse.tamu.edu/davis/suitesparse.html>.
- [2] SuperLU_MT Threaded Sparse Matrix Factorization Library. <http://crd-legacy.lbl.gov/~xiaoye/-SuperLU/>.
- [3] K. E. Brennan, S. L. Campbell, and L. R. Petzold. *Numerical Solution of Initial-Value Problems in Differential-Algebraic Equations*. SIAM, Philadelphia, Pa, 1996.
- [4] P. N. Brown and A. C. Hindmarsh. Reduced Storage Matrix Methods in Stiff ODE Systems. *J. Appl. Math. & Comp.*, 31:49–91, 1989.
- [5] 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.
- [6] 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.
- [7] 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*, pages 323–356, Oxford, 1992. Oxford University Press.
- [8] 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.
- [9] 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.
- [10] Y. Cao, S. Li, L. R. Petzold, and R. Serban. Adjoint Sensitivity Analysis for Differential-Algebraic Equations: The Adjoint DAE System and its Numerical Solution. *SIAM J. Sci. Comput.*, 24(3):1076–1089, 2003.
- [11] M. Caracotsios and W. E. Stewart. Sensitivity Analysis of Initial Value Problems with Mixed ODEs and Algebraic Equations. *Computers and Chemical Engineering*, 9:359–365, 1985.
- [12] S. D. Cohen and A. C. Hindmarsh. CVODE, a Stiff/Nonstiff ODE Solver in C. *Computers in Physics*, 10(2):138–143, 1996.
- [13] A. M. Collier, A. C. Hindmarsh, R. Serban, and C.S. Woodward. User Documentation for KINSOL v2.7.0. Technical Report UCRL-SM-208116, LLNL, 2011.
- [14] T. A. Davis and P. N. Ekanathan. Algorithm 907: KLU, a direct sparse solver for circuit simulation problems. *ACM Trans. Math. Softw.*, 37(3), 2010.
- [15] 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.
- [16] W. F. Feehery, J. E. Tolsma, and P. I. Barton. Efficient Sensitivity Analysis of Large-Scale Differential-Algebraic Systems. *Applied Numer. Math.*, 25(1):41–54, 1997.

- [17] R. W. Freund. A Transpose-Free Quasi-Minimal Residual Algorithm for Non-Hermitian Linear Systems. *SIAM J. Sci. Comp.*, 14:470–482, 1993.
- [18] K. L. Hiebert and L. F. Shampine. Implicitly Defined Output Points for Solutions of ODEs. Technical Report SAND80-0180, Sandia National Laboratories, February 1980.
- [19] A. C. Hindmarsh, P. N. Brown, K. E. Grant, S. L. Lee, R. Serban, D. E. Shumaker, and C. S. Woodward. SUNDIALS, suite of nonlinear and differential/algebraic equation solvers. *ACM Trans. Math. Softw.*, (31):363–396, 2005.
- [20] A. C. Hindmarsh and R. Serban. User Documentation for CVODE v2.7.0. Technical Report UCRL-SM-208108, LLNL, 2011.
- [21] A. C. Hindmarsh and R. Serban. User Documentation for CVODES v2.6.0. Technical report, LLNL, 2011. UCRL-SM-208111.
- [22] A. C. Hindmarsh, R. Serban, and A. Collier. User Documentation for IDA v2.7.0. Technical Report UCRL-SM-208112, LLNL, 2011.
- [23] 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.
- [24] S. Li, L. R. Petzold, and W. Zhu. Sensitivity Analysis of Differential-Algebraic Equations: A Comparison of Methods on a Special Problem. *Applied Num. Math.*, 32:161–174, 2000.
- [25] X. S. Li. An overview of SuperLU: Algorithms, implementation, and user interface. *ACM Trans. Math. Softw.*, 31(3):302–325, September 2005.
- [26] T. Maly and L. R. Petzold. Numerical Methods and Software for Sensitivity Analysis of Differential-Algebraic Systems. *Applied Numerical Mathematics*, 20:57–79, 1997.
- [27] D.B. Ozyurt and P.I. Barton. Cheap second order directional derivatives of stiff ODE embedded functionals. *SIAM J. of Sci. Comp.*, 26(5):1725–1743, 2005.
- [28] Y. Saad and M. H. Schultz. GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems. *SIAM J. Sci. Stat. Comp.*, 7:856–869, 1986.
- [29] R. Serban and A. C. Hindmarsh. CVODES, the sensitivity-enabled ODE solver in SUNDIALS. In *Proceedings of the 5th International Conference on Multibody Systems, Nonlinear Dynamics and Control*, Long Beach, CA, 2005. ASME.
- [30] R. Serban and A. C. Hindmarsh. Example Programs for IDAS v1.1.0. Technical Report LLNL-TR-437091, LLNL, 2011.
- [31] 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.*, 13:631–644, 1992.

Index

- AddIdentitySparseMat, 183
- adjoint sensitivity analysis
 - checkpointing, 16
 - implementation in IDAS, 17, 19–22
 - mathematical background, 14–17
 - quadrature evaluation, 138
 - residual evaluation, 136, 137
 - sensitivity-dependent quadrature evaluation, 138
- BAND generic linear solver
 - functions, 180
 - small matrix, 180–181
 - macros, 177
 - type DlsMat, 174–177
- BAND_COL, 73, 177
- BAND_COL_ELEM, 73, 177
- BAND_ELEM, 73, 177
- bandAddIdentity, 181
- bandCopy, 181
- bandGETRF, 181
- bandGETRS, 181
- bandMatvec, 181
- bandScale, 181
- Bi-CGSTab method, 51, 132, 187
- BIG_REAL, 26, 159
- CLASSICAL_GS, 50, 131
- CopySparseMat, 183
- DENSE generic linear solver
 - functions
 - large matrix, 177–178
 - small matrix, 178–180
 - macros, 177
 - type DlsMat, 174–177
- DENSE_COL, 72, 177
- DENSE_ELEM, 72, 177
- denseAddIdentity, 179
- denseCopy, 179
- denseGEQRF, 179
- denseGETRF, 179
- denseGETRS, 179
- denseMatvec, 180
- denseORMQR, 180
- densePOTRF, 179
- densePOTRS, 179
- denseScale, 179
- destroyArray, 179, 181
- destroyMat, 178, 180
- DestroySparseMat, 183
- DlsMat, 72, 73, 140–143, 174
- eh_data, 70
- error control
 - sensitivity variables, 12
- error messages, 40
 - redirecting, 40
 - user-defined handler, 40, 70
- FGMRES method, 187
- forward sensitivity analysis
 - absolute tolerance selection, 12–13
 - correction strategies, 11–12, 19, 92, 93
 - mathematical background, 11–14
 - residual evaluation, 103
 - right hand side evaluation, 13
 - right-hand side evaluation, 13
- generic linear solvers
 - BAND, 174
 - DENSE, 174
 - KLU, 181
 - SLS, 181
 - SPBCG, 187
 - SPFGMR, 187
 - SPGMR, 186
 - SPTFQMR, 187
 - SUPERLUMT, 181
 - use in IDAS, 22–23
- GMRES method, 186
- Gram-Schmidt procedure, 50, 131
- half-bandwidths, 34, 72–73, 86
- header files, 26, 85
- IDA_BAD_DKY, 54, 80, 96, 97, 107, 108
- IDA_BAD_EWT, 37

- IDA_BAD_IS, 97, 108
- IDA_BAD_ITASK, 126
- IDA_BAD_K, 54, 80, 96, 97, 107, 108
- IDA_BAD_T, 54, 80, 96, 97, 107, 108
- IDA_BAD_TBO, 121, 122
- IDA_BAD_TBOUT, 126
- IDA_BCKMEM_NULL, 126
- IDA_CENTERED, 98
- IDA_CONSTR_FAIL, 38, 39
- IDA_CONV_FAIL, 38, 39
- IDA_CONV_FAILURE, 120, 126
- IDA_ERR_FAIL, 39
- IDA_ERR_FAILURE, 120, 126
- IDA_FIRST_QRHS_ERR, 79, 83
- IDA_FIRST_QSRHS_ERR, 106, 112
- IDA_FIRST_RES_FAIL, 37, 103
- IDA_FORWARD, 98
- IDA_FWD_FAIL, 126
- IDA_GETY_BADT, 133
- IDA_HERMITE, 118
- IDA_ILL_INPUT, 30, 31, 37, 39, 42–45, 51–53, 62, 69, 81, 93–95, 98, 103, 105, 106, 109, 118, 120–123, 126, 127, 133–136
- IDA_LINESEARCH_FAIL, 38
- IDA_LINIT_FAIL, 37, 39
- IDA_LSETUP_FAIL, 37, 39, 120, 126, 140–145, 152
- IDA_LSOLVE_FAIL, 37, 39, 120
- IDA_MEM_FAIL, 30, 78, 93, 105, 118, 120, 121, 134, 135
- IDA_MEM_NULL, 30–32, 37, 39, 40, 42–45, 51–55, 57–63, 69, 78–82, 93–103, 105–111, 119, 121–123, 126, 133–135
- IDA_NO_ADJ, 118–127, 133–136
- IDA_NO_BCK, 126
- IDA_NO_FWD, 126
- IDA_NO_MALLOC, 31, 32, 37, 69, 120–123
- IDA_NO_QUAD, 79–82, 109, 135
- IDA_NO_QUADSENS, 106–111
- IDA_NO_RECOVERY, 37
- IDA_NO_SENS, 93–103, 105, 106, 108
- IDA_NORMAL, 39, 116, 119, 125
- IDA_ONE_STEP, 39, 116, 119, 125
- IDA_POLYNOMIAL, 118
- IDA_QRHS_FAIL, 79, 83, 112
- IDA_QRHSFUNC_FAIL, 138, 139
- IDA_QSRHS_FAIL, 106
- IDA_REIFWD_FAIL, 126
- IDA_REP_QRHS_ERR, 79
- IDA_REP_QSRHS_ERR, 106
- IDA_REP_RES_ERR, 39
- IDA_REP_SRES_ERR, 95
- IDA_RES_FAIL, 37, 39
- IDA_RESFUNC_FAIL, 136, 137
- IDA_ROOT_RETURN, 39
- IDA_RTFUNC_FAIL, 39, 71
- IDA_SIMULTANEOUS, 19, 92
- IDA_SOLVE_FAIL, 126
- IDA_SRES_FAIL, 95, 103
- IDA_STAGGERED, 19, 92
- IDA_SUCCESS, 30–32, 37, 39, 40, 42–45, 51–54, 62, 63, 69, 78–82, 93–102, 105–111, 118–123, 126, 133–135
- IDA_TOO_MUCH_ACC, 39, 120, 126
- IDA_TOO_MUCH_WORK, 39, 120, 126
- IDA_TSTOP_RETURN, 39, 120
- IDA_WARNING, 70
- IDA_Y_INIT, 37
- IDA_YA_YDP_INIT, 37
- IDAAdjFree, 119
- IDAAdjInit, 116, 118
- IDAAdjReInit, 118
- IDAAdjSetNoSensi, 119
- IDABAND linear solver
 - Jacobian approximation used by, 46
 - memory requirements, 63
 - NVECTOR compatibility, 34
 - optional input, 46–47, 128
 - optional output, 63–64
 - selection of, 34
- IDABand, 29, 33, 34, 72
- IDABAND_ILL_INPUT, 34
- IDABAND_MEM_FAIL, 34
- IDABAND_MEM_NULL, 34
- IDABAND_SUCCESS, 34
- IDABandB, 141
- IDABBDPRE preconditioner
 - description, 83–84
 - optional output, 87–88
 - usage, 85–86
 - usage with adjoint module, 150–153
 - user-callable functions, 86–87, 150–151
 - user-supplied functions, 84–85, 151–153
- IDABBDPrecGetNumGfnEvals, 88
- IDABBDPrecGetWorkSpace, 88
- IDABBDPrecInit, 86
- IDABBDPrecInitB, 150
- IDABBDPrecReInit, 87
- IDABBDPrecReInitB, 151
- IDACalcIC, 37
- IDACalcICB, 124
- IDACalcICBS, 124, 125
- IDACreate, 30
- IDACreateB, 116, 120
- IDADENSE linear solver
 - Jacobian approximation used by, 46
 - memory requirements, 63
 - NVECTOR compatibility, 33
 - optional input, 46–47, 127–128

- optional output, 63–64
 - selection of, 33
- IDADense, 29, 33, 71
- IDADenseB, 139
- IDADLS_ILL_INPUT, 34, 127, 128
- IDADLS_JACFUNC_RECVR, 140–143
- IDADLS_JACFUNC_UNRECVR, 140–143
- IDADLS_LMEM_NULL, 46, 63, 64, 127, 128
- IDADLS_MEM_FAIL, 34
- IDADLS_MEM_NULL, 34, 46, 63, 64, 127, 128
- IDADLS_NO_ADJ, 127, 128
- IDADLS_SUCCESS, 33, 46, 64, 127, 128
- IDADlsBandJacFn, 72
- IDADlsDenseJacFn, 71
- IDADlsGetLastFlag, 64
- IDADlsGetNumJacEvals, 63
- IDADlsGetNumResEvals, 64
- IDADlsGetReturnFlagName, 64
- IDADlsGetWorkspace, 63
- IDADlsSetBandJacFn, 46
- IDADlsSetBandJacFnB, 128
- IDADlsSetBandJacFnBS, 128
- IDADlsSetDenseJacFn, 46
- IDADlsSetDenseJacFnB, 127
- IDADlsSetDenseJacFnBS, 127
- IDAErrHandlerFn, 70
- IDA EwtFn, 70
- IDAFree, 29, 31
- IDAGetActualInitStep, 59
- IDAGetAdjIDABmem, 133
- IDAGetAdjY, 133
- IDAGetB, 126
- IDAGetConsistentIC, 62
- IDAGetConsistentICB, 133
- IDAGetCurrentOrder, 58
- IDAGetCurrentStep, 59
- IDAGetCurrentTime, 59
- IDAGetDky, 54
- IDAGetErrWeights, 60
- IDAGetEstLocalErrors, 60
- IDAGetIntegratorStats, 60
- IDAGetLastOrder, 58
- IDAGetLastStep, 58
- IDAGetNonlinSolvStats, 61
- IDAGetNumBacktrackOps, 62
- IDAGetNumErrTestFails, 58
- IDAGetNumGEvals, 63
- IDAGetNumLinSolvSetups, 57
- IDAGetNumNonlinSolvConvFails, 61
- IDAGetNumNonlinSolvIters, 61
- IDAGetNumResEvals, 57
- IDAGetNumResEvalsSEns, 100
- IDAGetNumSteps, 57
- IDAGetQuad, 79, 135
 - IDAGetQuadB, 117, 135
 - IDAGetQuadDky, 80
 - IDAGetQuadErrWeights, 82
 - IDAGetQuadNumErrTestFails, 82
 - IDAGetQuadNumRhsEvals, 81
 - IDAGetQuadSens, 107
 - IDAGetQuadSens1, 107
 - IDAGetQuadSensDky, 107
 - IDAGetQuadSensDky1, 108
 - IDAGetQuadSensErrWeights, 110
 - IDAGetQuadSensNumErrTestFails, 110
 - IDAGetQuadSensNumRhsEvals, 110
 - IDAGetQuadSensStats, 111
 - IDAGetQuadStats, 82
 - IDAGetReturnFlagName, 61
 - IDAGetRootInfo, 62
 - IDAGetSens, 91, 96
 - IDAGetSens1, 91, 97
 - IDAGetSensConsistentIC, 102
 - IDAGetSensDky, 91, 96
 - IDAGetSensDky1, 91, 97
 - IDAGetSensErrWeights, 101
 - IDAGetSensNonlinSolvStats, 102
 - IDAGetSensNumErrTestFails, 100
 - IDAGetSensNumLinSolvSetups, 100
 - IDAGetSensNumNonlinSolvConvFails, 102
 - IDAGetSensNumNonlinSolvIters, 101
 - IDAGetSensNumResEvals, 99
 - IDAGetSensStats, 101
 - IDAGetTolScaleFactor, 59
 - IDAGetWorkspace, 55
- IDAInit, 30, 68
- IDAInitB, 116, 121
- IDAInitBS, 116, 121
- IDAKLU, 29, 33, 35, 73
 - IDAKLU linear solver
 - Jacobian approximation used by, 47
 - matrix reordering algorithm specification, 48
 - NVECTOR compatibility, 35
 - optional input, 47–48, 129
 - optional output, 64–65
 - reinitialization, 47
 - selection of, 35
- IDAKLUB, 143
- IDAKLUReInit, 47
- IDAKLUSetOrdering, 48
- IDALapackBand, 29, 33, 35, 72
- IDALapackBandB, 141
- IDALapackDense, 29, 33, 34, 71
- IDALapackDenseB, 139
- IDAQuadFree, 79
- IDAQuadInit, 78
- IDAQuadInitB, 134
- IDAQuadInitBS, 134

- IDAQuadReInit, 79
- IDAQuadReInitB, 135
- IDAQuadRhsFn, 78, 82
- IDAQuadRhsFnB, 134, 138
- IDAQuadRhsFnBS, 134, 138
- IDAQuadSenseETolerances, 109, 110
- IDAQuadSensFree, 106
- IDAQuadSensInit, 105
- IDAQuadSensReInit, 106
- IDAQuadSensRhsFn, 105, 111
- IDAQuadSensSSTolerances, 109
- IDAQuadSensSVtolerances, 109
- IDAQuadSSTolerances, 81
- IDAQuadSVtolerances, 81
- IDAReInit, 68
- IDAReInitB, 122
- IDAResFn, 30, 69
- IDAResFnB, 121, 136
- IDAResFnBS, 122, 137
- IDARootFn, 71
- IDARootInit, 38
- IDAS
 - motivation for writing in C, 1–2
 - package structure, 19
 - relationship to IDA, 1
- IDAS linear solvers
 - built on generic solvers, 33
 - header files, 27
 - IDABAND, 34
 - IDADENSE, 33
 - IDAKLU, 35
 - IDASPCG, 36
 - IDASPGMR, 36
 - IDASPTFQMR, 36
 - IDASUPERLUMT, 35
 - implementation details, 22
 - list of, 22
 - NVECTOR compatibility, 25
 - selecting one, 33
 - usage with adjoint module, 123
- idas.h, 26
- idas.band.h, 27
- idas.dense.h, 27
- idas.klu.h, 27
- idas.lapack.h, 27
- idas.spbcgs.h, 27
- idas.spgmr.h, 27
- idas.sptfqmr.h, 27
- idas.superlunt.h, 27
- IDASenseETolerances, 95
- IDASensFree, 94
- IDASensInit, 91–93
- IDASensReInit, 93
- IDASensResFn, 92, 103
- IDASensSSTolerances, 94
- IDASensSVtolerances, 95
- IDASensToggleOff, 94
- IDASetConstraints, 45
- IDASetErrFile, 40
- IDASetErrHandlerFn, 40
- IDASetId, 45
- IDASetInitStep, 43
- IDASetLineSearchOffIC, 53
- IDASetMaxConvFails, 44
- IDASetMaxErrTestFails, 43
- IDASetMaxNonlinIters, 44
- IDASetMaxNumItersIC, 52
- IDASetMaxNumJacIC, 52
- IDASetMaxNumSteps, 42
- IDASetMaxNumStepsIC, 52
- IDASetMaxOrd, 42
- IDASetMaxStep, 43
- IDASetNoInactiveRootWarn, 54
- IDASetNonlinConvCoef, 44
- IDASetNonlinConvCoefIC, 51
- IDASetQuadErrCon, 80
- IDASetQuadSensErrCon, 108
- IDASetRootDirection, 53
- IDASetSensDQMethod, 98
- IDASetSensErrCon, 99
- IDASetSensMaxNonlinIters, 99
- IDASetSensParams, 98
- IDASetStepToleranceIC, 53
- IDASetStopTime, 43
- IDASetSuppressAlg, 45
- IDASetUserData, 42
- IDASLS_ILL_INPUT, 35, 36, 48, 129
- IDASLS_JACFUNC_RECVR, 144, 145
- IDASLS_JACFUNC_UNRECVR, 144, 145
- IDASLS_LMEM_NULL, 47, 48, 65, 129
- IDASLS_MEM_FAIL, 35, 36, 48
- IDASLS_MEM_NULL, 35, 47, 48, 65, 129
- IDASLS_NO_ADJ, 129
- IDASLS_PACKAGE_FAIL, 35, 36
- IDASLS_SUCCESS, 35, 47, 48, 65, 129
- IDASlsGetLastFlag, 65
- IDASlsGetNumJacEvals, 64
- IDASlsGetReturnFlagName, 65
- IDASlsSetSparseJacFn, 47
- IDASlsSetSparseJacFnB, 129
- IDASlsSetSparseJacFnBS, 129
- IDASlsSparseJacFn, 73
- IDASolve, 29, 38, 110
- IDASolveB, 117, 125
- IDASolveF, 116, 119
- IDASPCG linear solver
 - Jacobian approximation used by, 49
 - memory requirements, 65

- optional input, [49–51](#), [129–132](#)
- optional output, [65–68](#)
- preconditioner setup function, [49](#), [76](#), [148](#)
- preconditioner solve function, [49](#), [75](#), [147](#)
- selection of, [36](#)
- IDASpbcg, [29](#), [33](#), [36](#)
- IDASPGMR linear solver
 - Jacobian approximation used by, [49](#)
 - memory requirements, [65](#)
 - optional input, [49–51](#), [129–132](#)
 - optional output, [65–68](#)
 - preconditioner setup function, [49](#), [76](#), [148](#)
 - preconditioner solve function, [49](#), [75](#), [147](#)
 - selection of, [36](#)
- IDASpgmr, [29](#), [33](#), [36](#)
- IDASPILS_ILL_INPUT, [50](#), [51](#), [87](#), [130–132](#), [151](#)
- IDASPILS_LMEM_NULL, [49–51](#), [65–68](#), [87](#), [130–132](#), [151](#)
- IDASPILS_MEM_FAIL, [36](#), [37](#), [87](#), [151](#)
- IDASPILS_MEM_NULL, [36](#), [37](#), [49–51](#), [65–67](#), [130–132](#), [151](#)
- IDASPILS_NO_ADJ, [130–132](#)
- IDASPILS_PMEM_NULL, [87](#), [88](#), [151](#)
- IDASPILS_SUCCESS, [36](#), [37](#), [49–51](#), [67](#), [130–132](#), [151](#)
- IDASpilsGetLastFlag, [67](#)
- IDASpilsGetNumConvFails, [66](#)
- IDASpilsGetNumJtimesEvals, [67](#)
- IDASpilsGetNumLinIters, [66](#)
- IDASpilsGetNumPrecEvals, [66](#)
- IDASpilsGetNumPrecSolves, [67](#)
- IDASpilsGetNumResEvals, [67](#)
- IDASpilsGetReturnFlagName, [68](#)
- IDASpilsGetWorkSpace, [65](#)
- IDASpilsJacTimesVecFn, [74](#)
- IDASpilsPrecSetupFn, [76](#)
- IDASpilsPrecSolveFn, [75](#)
- IDASpilsSetEpsLin, [50](#)
- IDASpilsSetEpsLinB, [132](#)
- IDASpilsSetGSType, [50](#)
- IDASpilsSetGSTypeB, [131](#)
- IDASpilsSetIncrementFactor, [51](#)
- IDASpilsSetJacTimesFn, [49](#)
- IDASpilsSetJacTimesVecFnB, [130](#)
- IDASpilsSetJacTimesVecFnBS, [131](#)
- IDASpilsSetMaxl, [51](#)
- IDASpilsSetMaxlB, [132](#)
- IDASpilsSetMaxRestarts, [50](#)
- IDASpilsSetPreconditioner, [49](#)
- IDASpilsSetPrecSolveFnB, [130](#)
- IDASpilsSetPrecSolveFnBS, [130](#)
- IDASPTFQMR linear solver
 - Jacobian approximation used by, [49](#)
 - memory requirements, [65](#)
 - optional input, [49–51](#), [129–132](#)
 - optional output, [65–68](#)
 - preconditioner setup function, [49](#), [76](#), [148](#)
 - preconditioner solve function, [49](#), [75](#), [147](#)
 - selection of, [36](#)
- IDASptfqmr, [29](#), [33](#), [36](#)
- IDASStolerances, [31](#)
- IDASStolerancesB, [123](#)
- IDASUPERLUMT linear solver
 - Jacobian approximation used by, [47](#)
 - matrix reordering algorithm specification, [48](#)
 - NVECTOR compatibility, [35](#)
 - optional input, [47–48](#), [129](#)
 - optional output, [64–65](#)
 - selection of, [35](#)
- IDASuperLUMT, [29](#), [33](#), [35](#), [73](#)
- IDASuperLUMTB, [143](#)
- IDASuperLUMTSetOrdering, [48](#)
- IDASVtolerances, [31](#)
- IDASVtolerancesB, [123](#)
- IDAWFtolerances, [32](#)
- itask, [39](#), [119](#)
- Jacobian approximation function
 - band
 - difference quotient, [46](#)
 - user-supplied, [46](#), [72–73](#)
 - user-supplied (backward), [128](#), [141](#)
 - dense
 - difference quotient, [46](#)
 - user-supplied, [46](#), [71–72](#)
 - user-supplied (backward), [127](#), [139](#), [140](#)
 - Jacobian times vector
 - difference quotient, [49](#)
 - user-supplied, [49](#), [74–75](#)
 - Jacobian-vector product
 - user-supplied (backward), [130](#), [145](#)
 - sparse
 - user-supplied, [47](#), [73–74](#)
 - user-supplied (backward), [129](#), [143](#), [144](#)
- KLU sparse linear solver
 - type SlsMat, [182](#)
- maxl, [36](#), [37](#)
- maxord, [68](#)
- memory requirements
 - IDABAND linear solver, [63](#)
 - IDABBDPRE preconditioner, [87](#)
 - IDADENSE linear solver, [63](#)
 - IDAS solver, [78](#), [93](#), [105](#)
 - IDAS solver, [55](#)
 - IDASPGMR linear solver, [65](#)
- MODIFIED_GS, [50](#), [131](#)
- MPI, [3](#)

- N_VCloneEmptyVectorArray, 156
- N_VCloneEmptyVectorArray_OpenMP, 165
- N_VCloneEmptyVectorArray_Parallel, 163
- N_VCloneEmptyVectorArray_Pthreads, 167
- N_VCloneEmptyVectorArray_Serial, 160
- N_VCloneVectorArray, 156
- N_VCloneVectorArray_OpenMP, 165
- N_VCloneVectorArray_Parallel, 163
- N_VCloneVectorArray_Pthreads, 167
- N_VCloneVectorArray_Serial, 160
- N_VDestroyVectorArray, 156
- N_VDestroyVectorArray_OpenMP, 165
- N_VDestroyVectorArray_Parallel, 163
- N_VDestroyVectorArray_Pthreads, 167
- N_VDestroyVectorArray_Serial, 161
- N_Vector, 27, 155
- N_VMake_OpenMP, 165
- N_VMake_Parallel, 162
- N_VMake_Pthreads, 167
- N_VMake_Serial, 160
- N_VNew_OpenMP, 164
- N_VNew_Parallel, 162
- N_VNew_Pthreads, 166
- N_VNew_Serial, 160
- N_VNewEmpty_OpenMP, 165
- N_VNewEmpty_Parallel, 162
- N_VNewEmpty_Pthreads, 167
- N_VNewEmpty_Serial, 160
- N_VPrint_OpenMP, 165
- N_VPrint_Parallel, 163
- N_VPrint_Pthreads, 167
- N_VPrint_Serial, 161
- newBandMat, 180
- newDenseMat, 178
- newIntArray, 178, 181
- newLintArray, 178, 181
- newRealArray, 178, 181
- NewSparseMat, 183
- NV_COMM_P, 162
- NV_CONTENT_OMP, 164
- NV_CONTENT_P, 161
- NV_CONTENT_PT, 166
- NV_CONTENT_S, 159
- NV_DATA_OMP, 164
- NV_DATA_P, 161
- NV_DATA_PT, 166
- NV_DATA_S, 160
- NV_GLOBLLENGTH_P, 161
- NV_Ith_OMP, 164
- NV_Ith_P, 162
- NV_Ith_PT, 166
- NV_Ith_S, 160
- NV_LENGTH_OMP, 164
- NV_LENGTH_PT, 166
- NV_LENGTH_S, 160
- NV_LOCLENGTH_P, 161
- NV_NUM_THREADS_OMP, 164
- NV_NUM_THREADS_PT, 166
- NV_OWN_DATA_OMP, 164
- NV_OWN_DATA_P, 161
- NV_OWN_DATA_PT, 166
- NV_OWN_DATA_S, 160
- NVECTOR module, 155
- nvector_openmp.h, 26
- nvector_parallel.h, 26
- nvector_pthreads.h, 26
- nvector_serial.h, 26
- openMP, 3
- optional input
 - backward solver, 127
 - band linear solver, 46–47, 128
 - dense linear solver, 46–47, 127–128
 - forward sensitivity, 97–99
 - initial condition calculation, 51–53
 - iterative linear solver, 49–51, 129–132
 - quadrature integration, 80–81, 136
 - rootfinding, 53–54
 - sensitivity-dependent quadrature integration, 108–110
 - solver, 40–46
 - sparse linear solver, 47–48, 129
- optional output
 - backward initial condition calculation, 133–134
 - backward solver, 132–133
 - band linear solver, 63–64
 - band-block-diagonal preconditioner, 87–88
 - dense linear solver, 63–64
 - forward sensitivity, 99–102
 - initial condition calculation, 61–62, 102
 - interpolated quadratures, 80
 - interpolated sensitivities, 96
 - interpolated sensitivity-dep. quadratures, 107
 - interpolated solution, 54
 - iterative linear solver, 65–68
 - quadrature integration, 81–82, 136
 - sensitivity-dependent quadrature integration, 110–111
 - solver, 55–61
 - sparse linear solver, 64–65
- output mode, 119, 125
- partial error control
 - explanation of IDAS behavior, 112
- portability, 26
- preconditioning
 - advice on, 9, 22

- band-block diagonal, 83
- setup and solve phases, 22
- user-supplied, 49, 75, 76, 129–130, 147, 148
- `PrintSparseMat`, 185
- Pthreads, 3
- quadrature integration, 10
 - forward sensitivity analysis, 14
- RCONST, 26
- `ReallocSparseMat`, 185
- `realtype`, 26
- reinitialization, 68, 122
- residual function, 69
 - backward problem, 136, 137
 - forward sensitivity, 103
 - quadrature backward problem, 138
 - sensitivity-dep. quadrature backward problem, 138
- right-hand side function
 - quadrature equations, 82
 - sensitivity-dependent quadrature equations, 111
- Rootfinding, 29, 38
- rootfinding, 9
- `ScaleSparseMat`, 183
- second-order sensitivity analysis, 17
 - support in IDAS, 18
- SLS sparse linear solver
 - functions
 - small matrix, 183–185
- `SlsAddMat`, 183
- `SlsConvertDls`, 183
- `SlsMat`, 182
- `SlsMatvec`, 185
- SMALL_REAL, 26
- SPBCG generic linear solver
 - description of, 187
 - functions, 187
- SPFGMR generic linear solver
 - description of, 187
 - functions, 187
- SPGMR generic linear solver
 - description of, 186
 - functions, 186
 - support functions, 186
- SPTFQMR generic linear solver
 - description of, 187
 - functions, 187
- step size bounds, 43
- `sundials_nvector.h`, 27
- `sundials_types.h`, 26
- SUPERLUMT sparse linear solver
 - type `SlsMat`, 182
- TFQMR method, 51, 132, 187
- tolerances, 6, 32, 70, 81, 109
- UNIT_ROUNDOFF, 26
- User main program
 - Adjoint sensitivity analysis, 115
 - forward sensitivity analysis, 89
 - IDABBDPRE usage, 85
 - IDAS usage, 27
 - integration of quadratures, 77
 - integration of sensitivity-dependent quadratures, 104
- `user_data`, 42, 69–71, 83, 85, 111
- `user_dataB`, 152
- weighted root-mean-square norm, 6–7

