

User Documentation for CVODES v2.4.0

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

November 6, 2006



UCRL-SM-208111

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 the University of California nor any of their employees, makes any warranty, express 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 the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This research was supported under the auspices of the U.S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.

Contents

List of Tables	vii
List of Figures	ix
1 Introduction	1
1.1 Historical background	1
1.2 Changes from previous versions	2
1.3 Reading this user guide	3
2 CVODES Installation Procedure	5
2.1 Configuration options	6
2.2 Configuration examples	10
2.3 Installed libraries and exported header files	11
2.4 Building SUNDIALS without the configure script	11
3 Mathematical Considerations	15
3.1 IVP solution	15
3.2 Forward sensitivity analysis	18
3.2.1 Forward sensitivity methods	19
3.2.2 Selection of the absolute tolerances for sensitivity variables	20
3.2.3 Evaluation of the sensitivity right-hand side	21
3.3 Adjoint sensitivity analysis	21
3.3.1 Checkpointing scheme	23
3.4 Preconditioning	24
3.5 BDF stability limit detection	24
3.6 Rootfinding	25
4 Code Organization	27
4.1 SUNDIALS organization	27
4.2 CVODES organization	27
5 Using CVODES for IVP Solution	31
5.1 Access to libraries and header files	31
5.2 Data Types	32
5.3 Header files	32
5.4 A skeleton of the user's main program	33
5.5 User-callable functions for IVP solution	35
5.5.1 CVODES initialization and deallocation functions	35
5.5.2 Advice on choice and use of tolerances	36
5.5.3 Linear solver specification functions	37
5.5.4 CVODE solver function	40
5.5.5 Optional input functions	41
5.5.5.1 Main solver optional input functions	43

5.5.5.2	Dense linear solver	48
5.5.5.3	Band linear solver	49
5.5.5.4	SPILS linear solver	49
5.5.6	Interpolated output function	52
5.5.7	Optional output functions	52
5.5.7.1	Main solver optional output functions	52
5.5.7.2	Dense linear solver	59
5.5.7.3	Band linear solver	61
5.5.7.4	Diagonal linear solver	62
5.5.7.5	SPILS linear solvers	63
5.5.8	CVODES reinitialization function	66
5.6	User-supplied functions	67
5.6.1	ODE right-hand side	67
5.6.2	Error message handler function	68
5.6.3	Error weight function	69
5.6.4	Jacobian information (direct method with dense Jacobian)	69
5.6.5	Jacobian information (direct method with banded Jacobian)	70
5.6.6	Jacobian information (matrix-vector product)	71
5.6.7	Preconditioning (linear system solution)	72
5.6.8	Preconditioning (Jacobian data)	72
5.7	Integration of pure quadrature equations	73
5.7.1	Quadrature initialization functions	75
5.7.2	CVODE solver function	76
5.7.3	Quadrature extraction functions	76
5.7.4	Optional inputs for quadrature integration	77
5.7.5	Optional outputs for quadrature integration	78
5.7.6	User-supplied function for quadrature integration	79
5.8	Rootfinding	80
5.8.1	User-callable functions for rootfinding	80
5.8.2	User-supplied function for rootfinding	81
5.9	Preconditioner modules	81
5.9.1	A serial banded preconditioner module	81
5.9.2	A parallel band-block-diagonal preconditioner module	85
6	Using CVODES for Forward Sensitivity Analysis	93
6.1	A skeleton of the user's main program	93
6.2	User-callable routines for forward sensitivity analysis	95
6.2.1	Forward sensitivity initialization and deallocation functions	95
6.2.2	CVODE solver function	98
6.2.3	Forward sensitivity extraction functions	98
6.2.4	Optional inputs for forward sensitivity analysis	100
6.2.5	Optional outputs for forward sensitivity analysis	103
6.3	User-supplied routines for forward sensitivity analysis	107
6.3.1	Sensitivity equations right-hand side (all at once)	107
6.3.2	Sensitivity equations right-hand side (one at a time)	108
6.4	Note on using partial error control	109
7	Using CVODES for Adjoint Sensitivity Analysis	111
7.1	A skeleton of the user's main program	111
7.2	User-callable functions for adjoint sensitivity analysis	113
7.2.1	Adjoint sensitivity allocation and deallocation functions	113
7.2.2	Forward integration function	114
7.2.3	Backward problem initialization functions	116
7.2.4	Linear solver initialization functions for backward problem	117

7.2.5	Backward integration function	118
7.2.6	Optional input functions for the backward problem	119
7.2.6.1	Main solver optional input functions	119
7.2.6.2	Dense linear solver	119
7.2.6.3	Band linear solver	120
7.2.6.4	SPILS linear solvers	120
7.2.7	Optional output functions for the backward problem	123
7.2.8	Backward integration of pure quadrature equations	123
7.2.8.1	Backward quadrature initialization functions	123
7.2.8.2	Backward quadrature extraction function	124
7.2.8.3	Optional input/output functions for backward quadrature integration	124
7.2.9	Optional output from the adjoint module	124
7.2.9.1	Checkpoint information function	124
7.2.9.2	Interpolation data	125
7.2.9.3	Return flag name	126
7.3	User-supplied functions for adjoint sensitivity analysis	126
7.3.1	ODE right-hand side for the backward problem	126
7.3.2	Quadrature right-hand side for the backward problem	127
7.3.3	Jacobian information for the backward problem (direct method with dense Jacobian)	128
7.3.4	Jacobian information for the backward problem (direct method with banded Jacobian)	129
7.3.5	Jacobian information for the backward problem (matrix-vector product)	130
7.3.6	Preconditioning for the backward problem (linear system solution)	130
7.3.7	Preconditioning for the backward problem (Jacobian data)	131
7.4	Using CVODES preconditioner modules for the backward problem	132
7.4.1	Using the banded preconditioner CVBANDPRE	132
7.4.2	Using the band-block-diagonal preconditioner CVBBDPRE	134
7.4.2.1	Usage of CVBBDPRE for the backward problem	134
7.4.2.2	User-supplied functions for CVBBDPRE	137
8	Description of the NVECTOR module	139
8.1	The NVECTOR_SERIAL implementation	143
8.2	The NVECTOR_PARALLEL implementation	145
8.3	NVECTOR functions used by CVODES	147
9	Providing Alternate Linear Solver Modules	149
9.1	Initialization function	150
9.2	Setup function	150
9.3	Solve function	151
9.4	Memory deallocation function	151
10	Generic Linear Solvers in SUNDIALS	153
10.1	The DENSE module	153
10.1.1	Type DenseMat	154
10.1.2	Accessor Macros	154
10.1.3	Functions	154
10.1.4	Small Dense Matrix Functions	155
10.2	The BAND module	156
10.2.1	Type BandMat	157
10.2.2	Accessor Macros	157
10.2.3	Functions	159
10.3	The SPGMR module	159
10.3.1	Functions	160

10.4	The SPBCG module	160
10.4.1	Functions	160
10.5	The SPTFQMR module	161
10.5.1	Functions	161
11	CVODES Constants	163
11.1	CVODES input constants	163
11.2	CVODES output constants	164
	Bibliography	169
	Index	171

List of Tables

2.1	SUNDIALS libraries and header files (names are relative to <i>libdir</i> for libraries and to <i>includedir</i> for header files)	12
5.1	Optional inputs for CVODES, CVDENSE, CVBAND, and CVSPILS	42
5.2	Optional outputs from CVODES, CVDENSE, CVBAND, CVDIAG, and CVSPILS	53
6.1	Forward sensitivity optional inputs	100
6.2	Forward sensitivity optional outputs	103
8.1	Description of the NVECTOR operations	141
8.2	List of vector functions usage by CVODES code modules	148

List of Figures

3.1	Illustration of the checkpointing algorithm for generation of the forward solution during the integration of the adjoint system.	23
4.1	Organization of the SUNDIALS suite	28
4.2	Overall structure of the CVODES package	29
10.1	Diagram of the storage for a matrix of type BandMat	158

Chapter 1

Introduction

CVODES [26] is part of a software family called SUNDIALS: SUite of Nonlinear and Differential/ALgebraic equation Solvers [16]. This suite consists of CVODE, KINSOL and IDA, and variants of these with sensitivity analysis capabilities. CVODES is a solver for stiff and nonstiff initial value problems (IVPs) for systems of ordinary differential equation (ODEs). In addition to solving stiff and nonstiff ODE systems, CVODES has sensitivity analysis capabilities, using either the forward or the adjoint methods.

1.1 Historical background

FORTRAN solvers for ODE initial value problems are widespread and heavily used. Two solvers that were previously written at LLNL are VODE [1] and VODPK [3]. VODE is a general-purpose solver that includes methods for both stiff and nonstiff systems, and in the stiff case uses direct methods (full or banded) for the solution of the linear systems that arise at each implicit step. Externally, VODE is very similar to the well known solver LSODE [24]. VODPK is a variant of VODE that uses a preconditioned Krylov (iterative) method, namely GMRES, for the solution of the linear systems. VODPK is a powerful tool for large stiff systems 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 [2]. The capabilities of both VODE and VODPK were combined in the C-language package CVODE [8, 9].

At present, CVODES contains three Krylov methods that can be used in conjunction with Newton iteration: the GMRES (Generalized Minimal RESidual) [25], Bi-CGStab (Bi-Conjugate Gradient Stabilized) [27], and TFQMR (Transpose-Free Quasi-Minimal Residual) linear iterative methods [11]. 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 stiff ODE systems, the Krylov methods are preferable over direct linear solver methods, and are often the only feasible choice. Among the three Krylov methods in CVODES, 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.

In the process of translating the VODE and VODPK algorithms into C, the overall CVODE organization has changed considerably. One key feature of the CVODE organization is that the linear system solvers comprise a layer of code modules that is separated from the integration algorithm, thus allowing for easy modification and expansion of the linear solver array. A second key feature is a separate module devoted to vector operations; this facilitated the extension to multiprocessor environments with only a minimal impact on the rest of the solver, resulting in PVODE [5], the parallel variant of CVODE.

CVODES is written with a functionality that is a superset of that of the pair CVODE/PVODE. Sensitivity analysis capabilities, both forward and adjoint, have been added to the main integrator.

Enabling forward sensitivity computations in CVODES 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. CVODES provides the infrastructure needed to integrate any final-condition ODE dependent on the solution of the original IVP (in particular the adjoint system).

Development of CVODES was concurrent with a redesign of the vector operations module across the SUNDIALS suite. The key feature of the new NVECTOR module is that it is written in terms of abstract vector operations with the actual vector functions attached by a particular implementation (such as serial or parallel) of NVECTOR. This allows writing the SUNDIALS solvers in a manner independent of the actual NVECTOR implementation (which can be user-supplied), as well as allowing more than one NVECTOR module to be linked into an executable file.

There were several motivations for choosing the C language for CVODE, and later for CVODES. First, a general movement away from FORTRAN and toward C in scientific computing was and still is apparent. Second, the pointer, structure, and dynamic memory allocation features in C are extremely useful in software of this complexity. Finally, we prefer C over C++ for CVODES 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.

1.2 Changes from previous versions

Changes in v2.5.0

The main changes in this release involve a rearrangement of the entire SUNDIALS source tree (see §4.1). At the user interface level, the main impact is in the mechanism of including SUNDIALS header files which must now include the relative path (e.g. `#include <cvode/cvode.h>`). Additional changes were made to the build system: all exported header files are now installed in separate subdirectories of the installation *include* directory.

In the adjoint solver module, the following two bugs were fixed: in `CVodeF` the solver was sometimes incorrectly taking an additional step before returning control to the user (in `CV_NORMAL` mode) thus leading to a failure in the interpolated output function; in `CVodeB`, while searching for the current check point, the solver was sometimes reaching outside the integration interval resulting in a segmentation fault.

The functions in the generic dense linear solver (`sundials_dense` and `sundials_smalldense`) were modified to work for rectangular $m \times n$ matrices ($m \leq n$), while the factorization and solution functions were renamed to `DenseGETRF/denGETRF` and `DenseGETRS/denGETRS`, respectively. The factorization and solution functions in the generic band linear solver were renamed `BandGBTRF` and `BandGBTRS`, respectively.

Changes in v2.4.0

CVSPBCG and CVSPTFQMR modules have been added to interface with the Scaled Preconditioned Bi-CGstab (SPBCG) and Scaled Preconditioned Transpose-Free Quasi-Minimal Residual (SPTFQMR) linear solver modules, respectively (for details see Chapter 5). At the same time, function type names for Scaled Preconditioned Iterative Linear Solvers were added for the user-supplied Jacobian-times-vector and preconditioner setup and solve functions.

A new interpolation method was added to the CVODEA adjoint module. The function `CVadjMalloc` has an additional argument which can be used to select the desired interpolation scheme.

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

To reduce the possibility of conflicts, the names of all header files have been changed by adding unique prefixes (`cvodes_` and `sundials_`). When using the default installation procedure, the header files are exported under various subdirectories of the target `include` directory. For more details see §2.

Changes in v2.3.0

A minor bug was fixed in the interpolation functions of the adjoint CVODEA module.

Changes in v2.2.0

The user interface has been further refined. Several functions used for setting optional inputs were combined into a single one. An optional user-supplied routine for setting the error weight vector was added. Additionally, to resolve potential variable scope issues, all SUNDIALS solvers release user data right after its use. The build systems has been further improved to make it more robust.

Changes in v2.1.2

A bug was fixed in the `CVode` function that was potentially leading to erroneous behaviour of the root finding procedure on the integration first step.

Changes in v2.1.1

This CVODES release includes bug fixes related to forward sensitivity computations (possible loss of accuracy on a BDF order increase and incorrect logic in testing user-supplied absolute tolerances). In addition, we have added the option of activating and deactivating forward sensitivity calculations on successive CVODES runs without memory allocation/deallocation.

Other changes in this minor SUNDIALS release affect the build system.

Changes in v2.1.0

The major changes from the previous version involve a redesign of the user interface across the entire SUNDIALS suite. We have eliminated the mechanism of providing optional inputs and extracting optional statistics from the solver through the `iopt` and `ropt` arrays. Instead, CVODES now provides a set of routines (with prefix `CVodeSet`) to change the default values for various quantities controlling the solver and a set of extraction routines (with prefix `CVodeGet`) to extract statistics after return from the main solver routine. Similarly, each linear solver module provides its own set of `Set`- and `Get`-type routines. For more details see §5.5.5 and §5.5.7.

Additionally, the interfaces to several user-supplied routines (such as those providing Jacobians, preconditioner information, and sensitivity right hand sides) were simplified by reducing the number of arguments. The same information that was previously accessible through such arguments can now be obtained through `Get`-type functions.

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

Installation of CVODES (and all of SUNDIALS) has been completely redesigned and is now based on a configure script.

1.3 Reading this user guide

This user guide is a combination of general usage instructions. Specific example programs are provided as a separate document. We expect that some readers will want to concentrate on the general instructions, while others will refer mostly to the examples.

There are different possible levels of usage of CVODES. The most casual user, with an IVP problem only, can get by with reading §3.1, then §5 through §5.5.4 only, and looking at examples in [18]. In addition, to solve a forward sensitivity problem the user should read §3.2, followed by §6 through §6.2.3 only, and look at examples in [18].

In a different direction, a more advanced user with an IVP problem may want to (a) use a package preconditioner (§5.9), (b) supply his/her own Jacobian or preconditioner routines (§5.6), (c) do multiple runs of problems of the same size (§5.5.8), (d) supply a new NVECTOR module (§8), or even (e)

supply a different linear solver module (§4.2). An advanced user with a forward sensitivity problem may also want to (a) provide his/her own sensitivity equations right-hand side routine (§6.3), (b) perform multiple runs with the same number of sensitivity parameters (§6.2.1), or (c) extract additional diagnostic information (§6.2.3). A user with an adjoint sensitivity problem needs to understand the IVP solution approach at the desired level and also go through §3.3 for a short mathematical description of the adjoint approach, §7 for the usage of the adjoint module in CVODES, and the examples in [18].

The structure of this document is as follows:

- In Chapter 2 we begin with instructions for the installation of CVODES, within the structure of SUNDIALS.
- In Chapter 3, we give short descriptions of the numerical methods implemented by CVODES for the solution of initial value problems for systems of ODEs, continue with an overview of the mathematical aspects of sensitivity analysis, both forward (§3.2) and adjoint (§3.3), and conclude with a description of stability limit detection (§3.5).
- The following chapter describes the structure of the SUNDIALS suite of solvers (§4.1) and the software organization of the CVODES solver (§4.2).
- Chapter 5 is the main usage document for CVODES for simulation applications. It includes a complete description of the user interface for the integration of ODE initial value problems. Readers that are not interested in using CVODES for sensitivity analysis can then skip the next two chapters.
- Chapter 6 describes the usage of CVODES for forward sensitivity analysis as an extension of its IVP integration capabilities. We begin with a skeleton of the user main program, with emphasis on the steps that are required in addition to those already described in Chapter 5. 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 7 describes the usage of CVODES for adjoint sensitivity analysis. We begin by describing the CVODES 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 8 gives a brief overview of the generic NVECTOR module shared amongst the various components of SUNDIALS, as well as details on the two NVECTOR implementations provided with SUNDIALS: a serial implementation (§8.1) and a parallel implementation based on MPI (§8.2).
- Chapter 9 describes the specifications of linear solver modules as supplied by the user.
- Chapter 10 describes in detail the generic linear solvers shared by all SUNDIALS solvers.
- Finally, Chapter 11 lists the constants used for input to and output from CVODES.

Finally, the reader should be aware of the following notational conventions in this user guide: Program listings and identifiers (such as `CVodeMalloc`) within textual explanations appear in typewriter type style; fields in C structures (such as *content*) appear in italics; and packages or modules, such as CVDENSE, are written in all capitals. In the Index, page numbers that appear in bold indicate the main reference for that entry.

Chapter 2

CVODES Installation Procedure

The installation of CVODES 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 solvers other than CVODES.¹

The SUNDIALS suite (or individual solvers) are distributed as compressed archives (`.tgz`). The name of the distribution archive is of the form `solver-x.y.z.tgz`, where *solver* is one of: `sundials`, `cvode`, `cvodes`, `ida`, 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.tgz
```

This will extract source files under a directory `solver-x.y.z`.

The installation procedure outlined below will work on commodity LINUX/UNIX systems without modification. However, users are still encouraged to carefully read the entire chapter before attempting to install the SUNDIALS suite, in case non-default choices are desired for compilers, compilation options, or the like. In lieu of reading the option list below, the user may invoke the configuration script with the help flag to view a complete listing of available options, which may be done by issuing

```
% ./configure --help
```

from within the directory created above.

In the remainder of this chapter, we make the following distinctions:

- *srcdir*
is the directory `solver-x.y.z` created above; i.e., the directory containing the SUNDIALS sources.
- *builddir*
is the directory under which SUNDIALS is built; i.e., the directory from within which the `configure` command is issued. Usually, this is the same as *srcdir*.
- *instdir*
is the directory under which the SUNDIALS exported header files and libraries will be installed. Typically, header files are exported under a directory `instdir/include` while libraries are installed under `instdir/lib`, with *instdir* specified with the `--prefix` flag to `configure`. See §2.1 for more details on the installation directories, including the special cases of the SUNDIALS examples and the SUNDIALS/TB toolbox.

Note: The installation directory *instdir* should *not* be the same as the source directory *srcdir*.

¹Files for both the serial and parallel versions of CVODES are included in the distribution. For users in a serial computing environment, the files specific to parallel environments (which may be deleted) are as follows: all files in `src/nvec-par/`; `cvodes_bbdpre.c`, `cvodes_bbdpre_impl.h` (in `src/cvodes/`); `cvodes_bbdpre.h` (in `include/cvodes/`); all files in `examples/cvodes/parallel/`. (By “serial version” of CVODES we mean the CVODES solver with the serial NVECTOR module attached, and similarly for “parallel version”.)



The installation steps for SUNDIALS can be as simple as

```
% tar xzf solver-x.y.z.tgz
% cd solver-x.y.z
% ./configure
% make
% make install
```

in which case the SUNDIALS header files and libraries are installed under `/usr/local/include` and `/usr/local/lib`, respectively. Note that, by default, neither the example programs nor the SUNDIALS toolbox are built and installed.

If disk space is a priority, then to delete all temporary files created by building SUNDIALS, issue

```
% make clean
```

To prepare the SUNDIALS distribution for a new install (using, for example, different options and/or installation destinations), issue

```
% make distclean
```

2.1 Configuration options

The installation procedure given above will generally work without modification; however, if the system includes multiple MPI implementations, then certain configure script-related options may be used to indicate which MPI implementation should be used. Also, if the user wants to use non-default language compilers, then, again, the necessary shell environment variables must be appropriately redefined. The remainder of this section provides explanations of available configure script options.

General options

--prefix=PREFIX

Location for architecture-independent files.

Default: PREFIX=/usr/local

--exec-prefix=EPREFIX

Location for architecture-dependent files.

Default: EPREFIX=/usr/local

--includedir=DIR

Alternate location for installation of header files.

Default: DIR=PREFIX/include

--libdir=DIR

Alternate location for installation of libraries.

Default: DIR=EPREFIX/lib

--disable-solver

Although each existing solver module is built by default, support for a given solver can be explicitly disabled using this option. The valid values for *solver* are: `cvode`, `cvodes`, `ida`, and `kinsol`.

--enable-examples

Available example programs are *not* built by default. Use this option to enable compilation of all pertinent example programs. Upon completion of the `make` command, the example executables will be created under solver-specific subdirectories of `builddir/examples`:

builddir/examples/solver/serial : serial C examples

builddir/examples/solver/parallel : parallel C examples

builddir/examples/solver/fcmix_serial : serial FORTRAN examples

builddir/examples/solver/fcmix_parallel : parallel FORTRAN examples

Note: Some of these subdirectories may not exist depending upon the solver and/or the configuration options given.

--with-examples-instdir=EXINSTDIR

Alternate location for example executables and sample output files (valid only if examples are enabled). Note that installation of example files can be completely disabled by issuing **EXINSTDIR=no** (in case building the examples is desired only as a test of the SUNDIALS libraries).

Default: **DIR=EPREFIX/examples**

--with-cppflags=ARG

Specify additional C preprocessor flags (e.g., **ARG=-I<include_dir>** if necessary header files are located in nonstandard locations).

--with-cflags=ARG

Specify additional C compilation flags.

--with-ldflags=ARG

Specify additional linker flags (e.g., **ARG=-L<lib_dir>** if required libraries are located in nonstandard locations).

--with-libs=ARG

Specify additional libraries to be used (e.g., **ARG=-l<foo>** to link with the library named **libfoo.a** or **libfoo.so**).

--with-precision=ARG

By default, SUNDIALS will define a real number (internally referred to as **realtype**) to be a double-precision floating-point numeric data type (**double** C-type); however, this option may be used to build SUNDIALS with **realtype** alternatively defined as a single-precision floating-point numeric data type (**float** C-type) if **ARG=single**, or as a long double C-type if **ARG=extended**.

Default: **ARG=double**

Users should *not* build SUNDIALS with support for single-precision floating-point arithmetic on 32- or 64-bit systems. This will almost certainly result in unreliable numerical solutions. The configuration option **--with-precision=single** is intended for systems on which single-precision arithmetic involves at least 14 decimal digits.



Options for Fortran support

--disable-fcmix

Using this option will disable all FORTRAN support. The FCVODE, FKINSOL, FIDA, and FNVECTOR modules will not be built, regardless of availability.

--with-fflags=ARG

Specify additional FORTRAN compilation flags.

The configuration script will attempt to automatically determine the function name mangling scheme required by the specified FORTRAN compiler, but the following two options may be used to override the default behavior.

`--with-f77underscore=ARG`

This option pertains to the FCVODE, FKINSOL, FIDA, and FNVECTOR FORTRAN-C interface modules and is used to specify the number of underscores to append to function names so FORTRAN routines can properly link with the associated SUNDIALS libraries. Valid values for **ARG** are: **none**, **one** and **two**.

Default: **ARG=one**

`--with-f77case=ARG`

Use this option to specify whether the external names of the FCVODE, FKINSOL, FIDA, and FNVECTOR FORTRAN-C interface functions should be lowercase or uppercase so FORTRAN routines can properly link with the associated SUNDIALS libraries. Valid values for **ARG** are: **lower** and **upper**.

Default: **ARG=lower**

Options for MPI support

The following configuration options are only applicable to the parallel SUNDIALS packages:

`--disable-mpi`

Using this option will completely disable MPI support.

`--with-mpicc=ARG`

`--with-mpif77=ARG`

By default, the configuration utility script will use the MPI compiler scripts named **mpicc** and **mpif77** to compile the parallelized SUNDIALS subroutines; however, for reasons of compatibility, different executable names may be specified via the above options. Also, **ARG=no** can be used to disable the use of MPI compiler scripts, thus causing the serial C and FORTRAN compilers to be used to compile the parallelized SUNDIALS functions and examples.

`--with-mpi-root=MPIDIR`

This option may be used to specify which MPI implementation should be used. The SUNDIALS configuration script will automatically check under the subdirectories **MPIDIR/include** and **MPIDIR/lib** for the necessary header files and libraries. The subdirectory **MPIDIR/bin** will also be searched for the C and FORTRAN MPI compiler scripts, unless the user uses `--with-mpicc=no` or `--with-mpif77=no`.

`--with-mpi-incdir=INCDIR`

`--with-mpi-libdir=LIBDIR`

`--with-mpi-libs=LIBS`

These options may be used if the user would prefer not to use a preexisting MPI compiler script, but instead would rather use a serial compiler and provide the flags necessary to compile the MPI-aware subroutines in SUNDIALS.

Often an MPI implementation will have unique library names and so it may be necessary to specify the appropriate libraries to use (e.g., **LIBS=-lmpich**).

Default: **INCDIR=MPIDIR/include** and **LIBDIR=MPIDIR/lib**

`--with-mpi-flags=ARG`

Specify additional MPI-specific flags.

Options for library support

By default, only static libraries are built, but the following option may be used to build shared libraries on supported platforms.

`--enable-shared`

Using this particular option will result in both static and shared versions of the available SUNDIALS libraries being built if the system supports shared libraries. To build only shared libraries also specify `--disable-static`.

Note: The FCVODE, FKINSOL, and FIDA libraries can only be built as static libraries because they contain references to externally defined symbols, namely user-supplied FORTRAN subroutines. Although the FORTRAN interfaces to the serial and parallel implementations of the supplied NVECTOR module do not contain any unresolvable external symbols, the libraries are still built as static libraries for the purpose of consistency.

Options for Matlab support

The following options are relevant only for configuring and building the SUNDIALSTB Matlab toolbox:

`--enable-sundialsTB`

The SUNDIALSTB Matlab toolbox is *not* built by default. Use this option to enable configuration and compilation of the `mex` files. Upon completion of the `make` command, the following `mex` files will be created:

`builddir/sundialsTB/cvodes/cvm/cvm.mexext`

`builddir/sundialsTB/idas/idm/idm.mexext`

`builddir/sundialsTB/kinsol/kim/kim.mexext`

where *mexext* is the platform-specific extension of `mex` files.

`--with-sundialsTB-instdir=STBINSTDIR`

Alternate location for the installed SUNDIALSTB toolbox (valid only if SUNDIALSTB is enabled). As for the example programs, installation of SUNDIALSTB can be completely disabled by issuing `STBINSTDIR=no` (in case building the toolbox is desired but its installation will be done manually afterwards). Otherwise, all required SUNDIALSTB files will be installed under the directory `STBINSTDIR/sundialsTB`.

Default: `DIR=MATLAB/toolbox` (see below for the definition of `MATLAB`).

`--with-matlab=MATLAB`

This option can be used to specify the location of the Matlab executable. The default is to search the path.

`--with-mexopts=ARG`

Specify the `mex` options file to be used.

Default: Standard Matlab `mex` options file.

`--with-mexflags=ARG`

Specify the `mex` compiler flags to be used.

Default: `ARG=-O`

`--with-mexldadd=ARG`

Specify additional `mex` linker flags.

Default: none

Environment variables

The following environment variables can be locally (re)defined for use during the configuration of SUNDIALS. See the next section for illustrations of these.

CC

F77

Since the configuration script uses the first C and FORTRAN compilers found in the current executable search path, then each relevant shell variable (CC and F77) must be locally (re)defined in order to use a different compiler. For example, to use `xcc` (executable name of chosen compiler) as the C language compiler, use `CC=xcc` in the configure step.

CFLAGS

FFLAGS

Use these environment variables to override the default C and FORTRAN compilation flags.

2.2 Configuration examples

The following examples are meant to help demonstrate proper usage of the configure options.

To build SUNDIALS using the default C and Fortran compilers, and default `mpicc` and `mpif77` parallel compilers, enable compilation of examples, build the Matlab `mex` files for SUNDIALSB, and install it under `/home/myname/matlab/sundialsTB`, use

```
% configure --prefix=/home/myname/sundials --enable-examples \
--enable-sundialsTB --with-sundialsTB-instdir=/home/myname/matlab
```

To disable installation of the examples, use:

```
% configure --prefix=/home/myname/sundials \
--enable-examples --with-examples-instdir=no \
--enable-sundialsTB --with-sundialsTB-instdir=/home/myname/matlab
```

The following example builds SUNDIALS using `gcc` as the serial C compiler, `g77` as the serial FORTRAN compiler, `mpicc` as the parallel C compiler, `mpif77` as the parallel FORTRAN compiler, and appends the `-g3` compilation flag to the list of default flags:

```
% configure CC=gcc F77=g77 --with-cflags=-g3 --with-fflags=-g3 \
--with-mpicc=/usr/apps/mpich/1.2.4/bin/mpicc \
--with-mpif77=/usr/apps/mpich/1.2.4/bin/mpif77
```

The next example again builds SUNDIALS using `gcc` as the serial C compiler, but the `--with-mpicc=no` option explicitly disables the use of the corresponding MPI compiler script. In addition, since the `--with-mpi-root` option is given, the compilation flags `-I/usr/apps/mpich/1.2.4/include` and `-L/usr/apps/mpich/1.2.4/lib` are passed to `gcc` when compiling the MPI-enabled functions. The `--disable-examples` option explicitly disables the examples (which means a FORTRAN compiler is not required). The `--with-mpi-libs` option is required so that the configure script can check if `gcc` can link with the appropriate MPI library.

```
% configure CC=gcc --disable-examples --with-mpicc=no \
--with-mpi-root=/usr/apps/mpich/1.2.4 \
--with-mpi-libs=-lmpich
```

2.3 Installed libraries and exported header files

Using the standard 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 §2.1). 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 Table 2.1. The file extension *.lib* is typically *.so* for shared libraries and *.a* for static libraries (see *Options for library support* for additional details).

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_smalldense.h* could be used in building a preconditioner).

2.4 Building SUNDIALS without the configure script

If the `configure` script cannot be used (*e.g.*, when building SUNDIALS under Microsoft Windows without using Cygwin), or if the user prefers to own the build process (*e.g.*, when SUNDIALS is incorporated into a larger project with its own build system), then the header and source files for a given module can be copied from the *srcdir* to some other location and compiled separately.

The following files are required to compile a SUNDIALS solver module:

- public header files located under *srcdir/include/solver*
- implementation header files and source files located under *srcdir/src/solver*
- (optional) FORTRAN/C interface files located under *srcdir/src/solver/fcmix*
- shared public header files located under *srcdir/include/sundials*
- shared source files located under *srcdir/src/sundials*
- (optional) NVECTOR_SERIAL header and source files located under *srcdir/include/nvector* and *srcdir/src/nvec_ser*
- (optional) NVECTOR_PARALLEL header and source files located under *srcdir/include/nvector* and *srcdir/src/nvec_par*
- configuration header file *sundials_config.h* (see below)

A sample header file that, appropriately modified, can be used as *sundials_config.h* (otherwise created automatically by the `configure` script) is provided below. The various preprocessor macros defined within *sundials_config.h* have the following uses:

- Precision of the SUNDIALS `realtype` type

Only one of the macros *SUNDIALS_SINGLE_PRECISION*, *SUNDIALS_DOUBLE_PRECISION* and *SUNDIALS_EXTENDED_PRECISION* should be defined to indicate if the SUNDIALS `realtype` type is an alias for `float`, `double`, or `long double`, respectively.

Table 2.1: SUNDIALS libraries and header files (names are relative to *libdir* for libraries and to *includedir* for header files)

SHARED	Libraries	n/a	
	Header files	sundials/sundials_types.h sundials/sundials_config.h sundials/sundials_smalldense.h sundials/sundials_iterative.h sundials/sundials_spgmr.h sundials/sundials_spgmr.h	sundials/sundials_math.h sundials/sundials_nvector.h sundials/sundials_dense.h sundials/sundials_band.h sundials/sundials_sptfqmr.h
NVECTOR_SERIAL	Libraries	libsundials_nvecserial. <i>lib</i>	libsundials_fnvecserial.a
	Header files	nvector/nvector_serial.h	
NVECTOR_PARALLEL	Libraries	libsundials_nvecparallel. <i>lib</i>	libsundials_fnvecparallel.a
	Header files	nvector/nvector_parallel.h	
CVODE	Libraries	libsundials_cvode. <i>lib</i>	libsundials_fcvode.a
	Header files	cvode/cvode.h cvode/cvode_dense.h cvode/cvode_diag.h cvode/cvode_bandpre.h cvode/cvode_spgmr.h cvode/cvode_sptfqmr.h	cvode/cvode_band.h cvode/cvode_spils.h cvode/cvode_bbdpre.h cvode/cvode_spgmr.h cvode/cvode_impl.h
CVODES	Libraries	libsundials_cvodes. <i>lib</i>	
	Header files	cvodes/cvodes.h cvodes/cvodes_dense.h cvodes/cvodes_diag.h cvodes/cvodes_bandpre.h cvodes/cvodes_spgmr.h cvodes/cvodes_sptfqmr.h cvodes/cvodea_impl.h	cvodes/cvodes_band.h cvodes/cvodes_spils.h cvodes/cvodes_bbdpre.h cvodes/cvodes_spgmr.h cvodes/cvodes_impl.h
IDA	Libraries	libsundials_ida. <i>lib</i>	libsundials_fida.a
	Header files	ida/ida.h ida/ida_dense.h ida/ida_spils.h ida/ida_spgmr.h ida/ida_bbdpre.h	ida/ida_band.h ida/ida_spgmr.h ida/ida_sptfqmr.h ida/ida_impl.h
KINSOL	Libraries	libsundials_kinsol. <i>lib</i>	libsundials_fkinsol.a
	Header files	kinsol/kinsol.h kinsol/kinsol_dense.h kinsol/kinsol_spils.h kinsol/kinsol_spgmr.h kinsol/kinsol_bbdpre.h	kinsol/kinsol_band.h kinsol/kinsol_spgmr.h kinsol/kinsol_sptfqmr.h kinsol/kinsol_impl.h

- Use of generic math functions

If `SUNDIALS_USE_GENERIC_MATH` is defined, then the functions in `sundials_math.h` will use the `pow`, `sqrt`, `fabs`, and `exp` functions from the standard math library (see `math.h`), regardless of the definition of `realtype`. Otherwise, if `realtype` is defined to be an alias for the `float` C-type, then SUNDIALS will use `powf`, `sqrtf`, `fabsf`, and `expf`. If `realtype` is instead defined to be a synonym for the `long double` C-type, then `powl`, `sqrtl`, `fabsl`, and `expl` will be used.

Note: Although the `powf/powl`, `sqrtf/sqrtl`, `fabsf/fabsl`, and `expf/expl` routines are not specified in the ANSI C standard, they are ISO C99 requirements. Consequently, these routines will only be used if available.

- FORTRAN name-mangling scheme

The macros given below are used to transform the C-language function names defined in the FORTRAN-C interface modules in a manner consistent with the preferred FORTRAN compiler, thus allowing native C functions to be called from within a FORTRAN subroutine. The name-mangling scheme can be specified either by appropriately defining the parameterized macros (using the stringization operator, `##`, if necessary)

- `F77_FUNC(name,NAME)`
- `F77_FUNC_(name,NAME)`

or by defining *one* macro from each of the following lists:

- `SUNDIALS_CASE_LOWER` or `SUNDIALS_CASE_UPPER`
- `SUNDIALS_UNDERSCORE_NONE`, `SUNDIALS_UNDERSCORE_ONE`, or `SUNDIALS_UNDERSCORE_TWO`

For example, to specify that mangled C-language function names should be lowercase with one underscore appended include either

```
#define F77_FUNC(name,NAME) name ## _
#define F77_FUNC_(name,NAME) name ## _
```

or

```
#define SUNDIALS_CASE_LOWER 1
#define SUNDIALS_UNDERSCORE_ONE 1
```

in the `sundials_config.h` header file.

- Use of an MPI communicator other than `MPI_COMM_WORLD` in FORTRAN

If the macro `SUNDIALS_MPI_COMM_F2C` is defined, then the MPI implementation used to build SUNDIALS defines the type `MPI_Fint` and the function `MPI_Comm_f2c`, and it is possible to use MPI communicators other than `MPI_COMM_WORLD` with the FORTRAN-C interface modules.

```

1  /*
2  * -----
3  * Copyright (c) 2005, The Regents of the University of California.
4  * Produced at the Lawrence Livermore National Laboratory.
5  * All rights reserved.
6  * For details, see sundials/shared/LICENSE.
7  * -----
8  * SUNDIALS configuration header file
9  * -----
10 /*
11
12
13 /* Define SUNDIALS version number
14 * ----- */
15
16 #define SUNDIALS_PACKAGE_VERSION "2.2.1"
17
18 /* Define precision of SUNDIALS data type 'realtype'
19 * ----- */
20
21 /* Define SUNDIALS data type 'realtype' as 'double' */
22 #define SUNDIALS_DOUBLE_PRECISION 1
23
24 /* Define SUNDIALS data type 'realtype' as 'float' */
25 /* #define SUNDIALS_SINGLE_PRECISION 1 */
26
27 /* Define SUNDIALS data type 'realtype' as 'long double' */
28 /* #define SUNDIALS_EXTENDED_PRECISION 1 */
29
30 /* Use generic math functions
31 * ----- */
32
33 #define SUNDIALS_USE_GENERIC_MATH 1
34
35 /* FCMIX: Define Fortran name-mangling macro
36 * ----- */
37
38 #define F77_FUNC(name,NAME) name ## _
39 #define F77_FUNC_(name,NAME) name ## _
40
41 /* FCMIX: Define case of function names
42 * ----- */
43
44 /* FCMIX: Make function names lowercase */
45 /* #define SUNDIALS_CASE_LOWER 1 */
46
47 /* FCMIX: Make function names uppercase */
48 /* #define SUNDIALS_CASE_UPPER 1 */
49
50 /* FCMIX: Define number of underscores to append to function names
51 * ----- */
52
53 /* FCMIX: Do NOT append any underscores to functions names */
54 /* #define SUNDIALS_UNDERSCORE_NONE 1 */
55
56 /* FCMIX: Append ONE underscore to function names */
57 /* #define SUNDIALS_UNDERSCORE_ONE 1 */
58
59 /* FCMIX: Append TWO underscores to function names */
60 /* #define SUNDIALS_UNDERSCORE_TWO 1 */
61
62 /* FNVECTOR: Allow user to specify different MPI communicator
63 * ----- */
64
65 #define SUNDIALS_MPI_COMM_F2C 1

```


Chapter 3

Mathematical Considerations

CVODES solves ODE initial value problems (IVPs) in real N -space, which we write in the abstract form

$$\dot{y} = f(t, y), \quad y(t_0) = y_0, \quad (3.1)$$

where $y \in \mathbf{R}^N$. Here we use \dot{y} to denote dy/dt . While we use t to denote the independent variable, and usually this is time, it certainly need not be. CVODES solves both stiff and non-stiff systems. Roughly speaking, stiffness is characterized by the presence of at least one rapidly damped mode, whose time constant is small compared to the time scale of the solution itself.

Additionally, if (3.1) depends on some parameters $p \in \mathbf{R}^{N_p}$, i.e.

$$\begin{aligned} \dot{y} &= f(t, y, p) \\ y(t_0) &= y_0(p), \end{aligned} \quad (3.2)$$

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

3.1 IVP solution

The methods used in CVODES are variable-order, variable-step multistep methods, based on formulas of the form

$$\sum_{i=0}^{K_1} \alpha_{n,i} y^{n-i} + h_n \sum_{i=0}^{K_2} \beta_{n,i} \dot{y}^{n-i} = 0. \quad (3.3)$$

Here the y^n are computed approximations to $y(t_n)$, and $h_n = t_n - t_{n-1}$ is the step size. The user of CVODES must appropriately choose one of two multistep methods. For non-stiff problems, CVODES includes the Adams-Moulton formulas, characterized by $K_1 = 1$ and $K_2 = q$ above, where the order q varies between 1 and 12. For stiff problems, CVODES includes the Backward Differentiation Formulas (BDF) in so-called fixed-leading coefficient (FLC) form, given by $K_1 = q$ and $K_2 = 0$, with order q varying between 1 and 5. The coefficients are uniquely determined by the method type, its order, the recent history of the step sizes, and the normalization $\alpha_{n,0} = -1$. See [4] and [21].

For either choice of formula, the nonlinear system

$$G(y^n) \equiv y^n - h_n \beta_{n,0} f(t_n, y^n) - a_n = 0, \quad (3.4)$$

where $a_n \equiv \sum_{i>0} (\alpha_{n,i} y^{n-i} + h_n \beta_{n,i} \dot{y}^{n-i})$, must be solved (approximately) at each integration step. For this, CVODES offers the choice of either *functional iteration*, suitable only for non-stiff systems, and various versions of *Newton iteration*. Functional iteration, given by

$$y^{n(m+1)} = h_n \beta_{n,0} f(t_n, y^{n(m)}) + a_n,$$

involves evaluations of f only. In contrast, Newton iteration requires the solution of linear systems

$$M[y^{n(m+1)} - y^{n(m)}] = -G(y^{n(m)}), \quad (3.5)$$

in which

$$M \approx I - \gamma J, \quad J = \partial f / \partial y, \quad \text{and} \quad \gamma = h_n \beta_{n,0}. \quad (3.6)$$

The initial guess for the iteration is a predicted value $y^{n(0)}$ computed explicitly from the available history data. For the Newton corrections, CVODES provides a choice of six methods:

- dense direct solver,
- band direct solver,
- diagonal approximate Jacobian solver,
- scaled preconditioned GMRES (Generalized Minimal Residual method) without restarts,
- scaled preconditioned Bi-CGStab (Bi-Conjugate Gradient Stable method), or
- scaled preconditioned TFQMR (Transpose-Free Quasi-Minimal Residual method).

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 [2]. Note that the direct linear solvers (dense and band) can only be used with serial vector representations.

In the process of controlling errors at various levels, CVODES 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

$$W_i = 1/[\text{RTOL} \cdot |y_i| + \text{ATOL}_i]. \quad (3.7)$$

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 cases of a direct solver (dense, band, diagonal), the iteration is a modified Newton iteration since the iteration matrix M is fixed throughout the nonlinear iterations. However, for any of the Krylov methods, it is an Inexact Newton iteration, in which M is applied in a matrix-free manner, with matrix-vector products Jv obtained by either difference quotients or a user-supplied routine. The matrix M (for the direct solvers) or preconditioner matrix P (Krylov cases) is updated as infrequently as possible to balance the high costs of matrix operations against other costs. Specifically, this matrix update occurs when:

- starting the problem,
- more than 20 steps have been taken since the last update,
- the value $\bar{\gamma}$ of γ at the last update satisfies $|\gamma/\bar{\gamma} - 1| > 0.3$,
- a non-fatal convergence failure just occurred, or
- an error test failure just occurred.

When forced by a convergence failure, an update of M or P may involve a reevaluation of J (in M) or of Jacobian data (in P) if Jacobian error was the likely cause of the failure. More generally, the decision is made to reevaluate J (or instruct the user to reevaluate Jacobian data in P) when:

- starting the problem,
- more than 50 steps have been taken since the last evaluation,

- a convergence failure occurred with an outdated matrix, and the value $\bar{\gamma}$ (γ at the last update) satisfies $|\gamma/\bar{\gamma} - 1| < 0.2$, or
- a convergence failure occurred that forced a reduction of the step size.

The stopping test for the Newton iteration is related to the subsequent local error test, with the goal of keeping the nonlinear iteration errors from interfering with local error control. As described below, the final computed value $y^{n(m)}$ will have to satisfy a local error test $\|y^{n(m)} - y^{n(0)}\| \leq \epsilon$. Letting y^n denote the exact solution of (3.4), we want to ensure that the iteration error $y^n - y^{n(m)}$ is small relative to ϵ , specifically that it is less than 0.1ϵ . (The safety factor 0.1 can be changed by the user.) For this, we also estimate the linear convergence rate constant R as follows. We initialize R to 1, and reset $R = 1$ when M or P is updated. After computing a correction $\delta_m = y^{n(m)} - y^{n(m-1)}$, we update R if $m > 1$ as

$$R \leftarrow \max\{0.3R, \|\delta_m\|/\|\delta_{m-1}\|\}.$$

Now we use the estimate

$$\|y^n - y^{n(m)}\| \approx \|y^{n(m+1)} - y^{n(m)}\| \approx R\|y^{n(m)} - y^{n(m-1)}\| = R\|\delta_m\|.$$

Therefore the convergence (stopping) test is

$$R\|\delta_m\| < 0.1\epsilon.$$

We allow at most 3 iterations, but this limit can be changed by the user. We also declare the iteration diverged if any $\|\delta_m\|/\|\delta_{m-1}\| > 2$ with $m > 1$. If convergence fails with J or P current, we are forced to reduce the step size, and we replace h_n by $h_n/4$. The integration is halted after a preset number of convergence failures; the default value of this limit is 10, but this can be changed by the user.

When a Krylov method is used to solve the linear system, its errors must also be controlled, and this also involves the local error test constant. The linear iteration error in the solution vector δ_m is approximated by the preconditioned residual vector. Thus to ensure (or attempt to ensure) that the linear iteration errors do not interfere with the nonlinear error and local integration error controls, we require that the norm of the preconditioned residual be less than $0.05 \cdot (0.1\epsilon)$.

With the direct dense and band methods, the Jacobian may be supplied by a user routine, or approximated by difference quotients, at the user's option. In the latter case, we use the usual approximation

$$J_{ij} = [f_i(t, y + \sigma_j e_j) - f_i(t, y)]/\sigma_j.$$

The increments σ_j are given by

$$\sigma_j = \max\left\{\sqrt{U} |y_j|, \sigma_0/W_j\right\},$$

where U is the unit roundoff, σ_0 is a dimensionless value, and W_j is the error weight defined in (3.7). In the dense case, this scheme requires N evaluations of f , one for each column of J . In the band case, the columns of J are computed in groups by the Curtis-Powell-Reid algorithm, with the number of f evaluations equal to the bandwidth.

In the case of a Krylov method, preconditioning may be used on the left, on the right, or both, with user-supplied routines for the preconditioning setup and solve operations, and optionally also for the required matrix-vector products Jv . If a routine for Jv is not supplied, these products are computed as

$$Jv = [f(t, y + \sigma v) - f(t, y)]/\sigma. \quad (3.8)$$

The increment σ is $1/\|v\|$, so that σv has norm 1.

A critical part of CVODES, that makes it an ODE “solver” rather than just an ODE method, is its control of local error. At every step, the local error is estimated and required to satisfy tolerance conditions, and the step is redone with reduced step size whenever that error test fails. As with any linear multistep method, the local truncation error LTE, at order q and step size h , satisfies an asymptotic relation

$$\text{LTE} = Ch^{q+1}y^{(q+1)} + O(h^{q+2})$$

for some constant C , under mild assumptions on the step sizes. A similar relation holds for the error in the predictor $y^{n(0)}$. These are combined to get a relation

$$\text{LTE} = C'[y^n - y^{n(0)}] + O(h^{q+2}).$$

The local error test is simply $\|\text{LTE}\| \leq 1$. Using the above, it is performed on the predictor-corrector difference $\Delta_n \equiv y^{n(m)} - y^{n(0)}$ (with $y^{n(m)}$ the final iterate computed), and takes the form

$$\|\Delta_n\| \leq \epsilon \equiv 1/|C'|.$$

If this test passes, the step is considered successful. If it fails, the step is rejected and a new step size h' is computed based on the asymptotic behavior of the local error, namely by the equation

$$(h'/h)^{q+1} \|\Delta_n\| = \epsilon/6.$$

Here $1/6$ is a safety factor. A new attempt at the step is made, and the error test repeated. If it fails three times, the order q is reset to 1 (if $q > 1$), or the step is restarted from scratch (if $q = 1$). The ratio h'/h is limited above to 0.2 after two error test failures, and limited below to 0.1 after three. After seven failures, CVODES returns to the user with a give-up message.

In addition to adjusting the step size to meet the local error test, CVODES periodically adjusts the order, with the goal of maximizing the step size. The integration starts out at order 1, but the order is varied dynamically after that. The basic idea is to pick the order q for which a polynomial of order q best fits the discrete data involved in the multistep method. However, if either a convergence failure or an error test failure occurred on the step just completed, no change is made to the step size or order. At the current order q , selecting a new step size is done exactly as when the error test fails, giving a tentative step size ratio

$$h'/h = (\epsilon/6 \|\Delta_n\|)^{1/(q+1)} \equiv \eta_q.$$

We consider changing order only after taking $q+1$ steps at order q , and then we consider only orders $q' = q-1$ (if $q > 1$) or $q' = q+1$ (if $q < 5$). The local truncation error at order q' is estimated using the history data. Then a tentative step size ratio is computed on the basis that this error, $\text{LTE}(q')$, behaves asymptotically as $h^{q'+1}$. With safety factors of $1/6$ and $1/10$ respectively, these ratios are:

$$h'/h = [1/6 \|\text{LTE}(q-1)\|]^{1/q} \equiv \eta_{q-1}$$

and

$$h'/h = [1/10 \|\text{LTE}(q+1)\|]^{1/(q+2)} \equiv \eta_{q+1}.$$

The new order and step size are then set according to

$$\eta = \max\{\eta_{q-1}, \eta_q, \eta_{q+1}\}, \quad h' = \eta h,$$

with q' set to the index achieving the above maximum. However, if we find that $\eta < 1.5$, we do not bother with the change. Also, h'/h is always limited to 10, except on the first step, when it is limited to 10^4 .

The various algorithmic features of CVODES described above, as inherited from VODE and VODPK, are documented in [1, 3, 15]. They are also summarized in [16].

Normally, CVODES takes steps until a user-defined output value $t = t_{\text{out}}$ is overtaken, and then it 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 CVODES not to integrate past a given stopping point $t = t_{\text{stop}}$.

3.2 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 (3.2). In

addition to numerically solving the ODEs, 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 in short *sensitivity equations*):

$$\dot{s}_i = \frac{\partial f}{\partial y} s_i + \frac{\partial f}{\partial p_i}, \quad s_i(t_0) = \frac{\partial y_0(p)}{\partial p_i}, \quad (3.9)$$

obtained by applying the chain rule of differentiation to the original ODEs (3.2).

When performing forward sensitivity analysis, CVODES carries out the time integration of the combined system, (3.2) and (3.9), by viewing it as an ODE system of size $N(N_s + 1)$, where N_s is the number of model parameters p_i , with respect to the desired sensitivities ($N_s \leq N_p$). However, major improvements in efficiency can be made by taking advantage of the special form of the sensitivity equations as linearizations of the original ODEs. In particular, for stiff systems, for which CVODES employs a Newton iteration, the original ODE system and all sensitivity systems share the same Jacobian matrix, and therefore the same iteration matrix M in (3.6).

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

3.2.1 Forward sensitivity methods

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

- *Staggered Direct*

In this approach [7], the nonlinear system (3.4) is first solved and, once an acceptable numerical solution is obtained, the sensitivity variables at the new step are found by directly solving (3.9) after the (BDF or Adams) 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 M in (3.6), it must be updated and factored at every step of the integration, in contrast to M 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 [22]. 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 CVODES.

- *Simultaneous Corrector*

In this method [23], the discretization is applied simultaneously to both the original equations (3.2) and the sensitivity systems (3.9) resulting in the following nonlinear system

$$\hat{G}(\hat{y}_n) \equiv \hat{y}_n - h_n \beta_{n,0} \hat{f}(t_n, \hat{y}_n) - \hat{a}_n = 0,$$

where $\hat{f} = [f(t, y, p), \dots, (\partial f/\partial y)(t, y, p)s_i + (\partial f/\partial p_i)(t, y, p), \dots]$, and \hat{a}_n is comprised of the terms in the discretization that depend on the solution at previous integration steps. This combined nonlinear system can be solved using a modified Newton method as in (3.5) by solving the corrector equation

$$\hat{M}[\hat{y}_{n(m+1)} - \hat{y}_{n(m)}] = -\hat{G}(\hat{y}_{n(m)}) \quad (3.10)$$

at each iteration, where

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

M is defined as in (3.6), and $J_i = (\partial/\partial y)[(\partial f/\partial y)s_i + (\partial f/\partial p_i)]$. It can be shown that 2-step quadratic convergence can be attained by using only the block-diagonal portion of \hat{M} in the corrector equation (3.10). This results in a decoupling that allows the reuse of M without additional matrix factorizations. However, the products $(\partial f/\partial y)s_i$ and the vectors $\partial f/\partial p_i$ must still be reevaluated at each step of the iterative process (3.10) to update the sensitivity portions of the residual \hat{G} .

- *Staggered corrector*

In this approach [10], as in the staggered direct method, the nonlinear system (3.4) is solved first using the Newton iteration (3.5). Then a separate Newton iteration is used to solve the sensitivity system (3.9):

$$M[s_i^{n(m+1)} - s_i^{n(m)}] = - \left[s_i^{n(m)} - \gamma \left(\frac{\partial f}{\partial y}(t_n, y^n, p) s_i^{n(m)} + \frac{\partial f}{\partial p_i}(t_n, y^n, p) \right) - a_{i,n} \right], \quad (3.11)$$

where $a_{i,n} = \sum_{j>0} (\alpha_{n,j} s_i^{n-j} + h_n \beta_{n,j} \dot{s}_i^{n-j})$. In other words, a modified Newton iteration is used to solve a linear system. In this approach, the vectors $\partial f/\partial p_i$ need be updated only once per integration step, after the state correction phase (3.5) has converged. Note also that Jacobian-related data can be reused at all iterations (3.11) to evaluate the products $(\partial f/\partial y)s_i$.

CVODES implements the simultaneous corrector method and two flavors of the staggered corrector method which differ only if the sensitivity variables are included in the error control test. In the *full error control* case, the first variant of the staggered corrector method requires the convergence of the iterations (3.11) for all N_s sensitivity systems and then performs the error test on the sensitivity variables. The second variant of the method will perform the error test for each sensitivity vector $s_i, (i = 1, 2, \dots, N_s)$ individually, as they pass the convergence test. Differences in performance between the two variants may therefore be noticed whenever one of the sensitivity vectors s_i fails a convergence or error test.

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 M on a vector and this can be provided with the current Jacobian information. Therefore, the modified Newton procedure (3.11) will theoretically converge after one iteration.

3.2.2 Selection of the absolute tolerances for sensitivity variables

If the sensitivities are included in the error test, CVODES 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 CVODES can provide different values as an option.

3.2.3 Evaluation of the sensitivity right-hand side

There are several methods for evaluating the right-hand side of the sensitivity systems (3.9): analytic evaluation, automatic differentiation, complex-step approximation, and finite differences (or directional derivatives). CVODES 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), CVODES can evaluate these quantities using various finite difference-based approximations to evaluate the terms $(\partial f/\partial y)s_i$ and $(\partial f/\partial p_i)$, or using directional derivatives to evaluate $[(\partial f/\partial y)s_i + (\partial f/\partial p_i)]$. As is typical for finite differences, the proper choice of perturbations is a delicate matter. CVODES takes into account several problem-related features: the relative ODE 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 $(\partial f/\partial y)s_i$ and $\partial f/\partial p_i$ in the right-hand side of (3.9) can be evaluated separately:

$$\frac{\partial f}{\partial y}s_i \approx \frac{f(t, y + \sigma_y s_i, p) - f(t, y - \sigma_y s_i, p)}{2\sigma_y}, \quad (3.12)$$

$$\frac{\partial f}{\partial p_i} \approx \frac{f(t, y, p + \sigma_i e_i) - f(t, y, p - \sigma_i e_i)}{2\sigma_i}, \quad (3.12')$$

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

simultaneously:

$$\frac{\partial f}{\partial y}s_i + \frac{\partial f}{\partial p_i} \approx \frac{f(t, y + \sigma s_i, p + \sigma e_i) - f(t, y - \sigma s_i, p - \sigma e_i)}{2\sigma}, \quad (3.13)$$

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

or by adaptively switching between (3.12)+(3.12') and (3.13), depending on the relative size of the estimated finite difference increments σ_i and σ_y .

These procedures for choosing the perturbations $(\delta_i, \delta_y, \delta)$ and switching (ρ_{\max}) between finite difference and directional derivative formulas have also been implemented for first-order formulas. Forward finite differences can be applied to $(\partial f/\partial y)s_i$ and $\partial f/\partial p_i$ separately, or the single directional derivative formula

$$\frac{\partial f}{\partial y}s_i + \frac{\partial f}{\partial p_i} \approx \frac{f(t, y + \delta s_i, p + \delta e_i) - f(t, y, p)}{\delta}$$

can be used. In CVODES, the default value of $\rho_{\max} = 0$ indicates the use of the second-order centered directional derivative formula (3.13) 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.

3.3 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 ODE 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 (3.2), we wish to evaluate the gradient dG/dp of

$$G(p) = \int_{t_0}^{t_f} g(t, y, p) dt, \quad (3.14)$$

or, alternatively, the gradient dg/dp of the function $g(t, x, p)$ at time t_f . 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 [6]. Introducing a Lagrange multiplier λ , we form the augmented objective function

$$I(p) = G(p) - \int_{t_0}^{t_1} \lambda^* (\dot{y} - f(t, y, p)) dt, \quad (3.15)$$

where $*$ denotes the conjugate transpose. The gradient of G with respect to p is

$$\frac{dG}{dp} = \frac{dI}{dp} = \int_{t_0}^{t_1} (g_p + g_y s) dt - \int_{t_0}^{t_1} \lambda^* (\dot{s} - f_y s - f_p) dt, \quad (3.16)$$

where subscripts on functions such as f or g are used to denote partial derivatives and $s = [s_1, \dots, s_{N_s}]$ is the matrix of solution sensitivities. Applying integration by parts to the term $\lambda^* \dot{s}$ and selecting λ such that

$$\begin{aligned} \dot{\lambda} &= - \left(\frac{\partial f}{\partial y} \right)^* \lambda - \left(\frac{\partial g}{\partial y} \right)^* \\ \lambda(t_1) &= 0, \end{aligned} \quad (3.17)$$

the gradient of G with respect to p is nothing but

$$\frac{dG}{dp} = \lambda^*(t_0) s(t_0) + \int_{t_0}^{t_1} (g_p + \lambda^* f_p) dt. \quad (3.18)$$

The gradient of $g(t_1, y, p)$ with respect to p can be then obtained by using the Leibnitz differentiation rule. Indeed, from (3.14),

$$\frac{dg}{dp}(t_1) = \frac{d}{dt_1} \frac{dG}{dp}$$

and therefore, taking into account that dG/dp in (3.18) depends on t_1 both through the upper integration limit and through λ and that $\lambda(t_1) = 0$,

$$\frac{dg}{dp}(t_1) = \mu^*(t_0) s(t_0) + g_p(t_1) + \int_{t_0}^{t_1} \mu^* f_p dt, \quad (3.19)$$

where μ is the sensitivity of λ with respect to the final integration limit and thus satisfies the following equation, obtained by taking the total derivative with respect to t_1 of (3.17):

$$\begin{aligned} \dot{\mu} &= - \left(\frac{\partial f}{\partial y} \right)^* \mu \\ \mu(t_1) &= \left(\frac{\partial g}{\partial y} \right)^*_{t=t_1}. \end{aligned} \quad (3.20)$$

The final condition on $\mu(t_1)$ follows from $(\partial \lambda / \partial t) + (\partial \lambda / \partial t_1) = 0$ at t_1 , and therefore, $\mu(t_1) = -\dot{\lambda}(t_1)$.

The first thing to notice about the adjoint system (3.17) is that there is no explicit specification of the parameters p ; this implies that, once the solution λ is found, the formula (3.18) can then be used to find the gradient of G with respect to any of the parameters p . The same holds true for the system (3.20) and the formula (3.19) for gradients of $g(t_1, y, p)$. The second important remark is that the adjoint systems (3.17) and (3.20) are terminal value problems which depend on the solution $y(t)$ of the original IVP (3.2). Therefore, a procedure is needed for providing the states y obtained during a forward integration phase of (3.2) to CVODES during the backward integration phase of (3.17) or (3.20). The approach adopted in CVODES, based on *checkpointing*, is described below.




Figure 3.1: Illustration of the checkpointing algorithm for generation of the forward solution during the integration of the adjoint system.

3.3.1 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 CVODES 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 CVODES 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. CVODES 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, CVODES 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 from t_1 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 N_d the vectors y (and, if necessary \dot{y}) are generated and stored in memory for interpolation¹ (see Fig. 3.1).

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

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 §3.1), 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. The accuracy differences are minor when using BDF (since the maximum method order cannot exceed 5), but can be significant for the Adams method for which the order can reach 12.

in the solution of (3.2)), the total cost of an adjoint sensitivity computation can be as low as one forward plus one backward integration. In addition, CVODES provides the capability of reusing a set of checkpoints for multiple backward integrations, thus allowing for efficient computation of gradients of several functionals (3.14).

Finally, we note that the adjoint sensitivity module in CVODES provides the necessary infrastructure to integrate backwards in time any ODE terminal value problem dependent on the solution of the IVP (3.2), including adjoint systems (3.17) or (3.20), as well as any other quadrature ODEs that may be needed in evaluating the integrals in (3.18) or (3.19). In particular, for ODE 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.

3.4 Preconditioning

When using a Newton method to solve the nonlinear system (3.4), CVODES makes repeated use of a linear solver to solve linear systems of the form $Mx = -r$, where x is a correction vector and r is a residual vector. 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, as $(P^{-1}A)x = P^{-1}b$; on the right, as $(AP^{-1})Px = b$; or on both sides, as $(P_L^{-1}AP_R^{-1})P_Rx = P_L^{-1}b$. 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 . In order to improve the convergence of the Krylov iteration, the preconditioner matrix P , or the product $P_L P_R$ in the last case, should in some sense approximate the system matrix A . Yet at the same time, in order to be cost-effective, the matrix P , or matrices P_L and P_R , 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 [2] for an extensive study of preconditioners for reaction-transport systems).

The CVODES solver allow for preconditioning either side, or on both sides, although we know of no situation where preconditioning on both sides is clearly superior to preconditioning on one side only (with the product $P_L P_R$). Moreover, for a given preconditioner matrix, the merits of left vs. right preconditioning are unclear in general, and the user should experiment with both choices. Performance will differ because the inverse of the left preconditioner is included in the linear system residual whose norm is being tested in the Krylov algorithm. As a rule, however, if the preconditioner is the product of two matrices, we recommend that preconditioning be done either on the left only or the right only, rather than using one factor on each side.

Typical preconditioners used with CVODES are based on approximations to the Newton iteration matrix of the systems involved; in other words, $P \approx I - \gamma J$, where γ is a scalar proportional to the integration step size h and J is the system Jacobian, i.e. $J = \partial f / \partial y$. 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.

3.5 BDF stability limit detection

CVODES includes an algorithm, STALD (STability Limit Detection), which provides protection against potentially unstable behavior of the BDF multistep integration methods in certain situations, as described below.

When the BDF option is selected, CVODES uses Backward Differentiation Formula methods of orders 1 to 5. At order 1 or 2, the BDF method is A-stable, meaning that for any complex constant

λ in the open left half-plane the method is unconditionally stable (for any step size) for the standard scalar model problem $\dot{y} = \lambda y$. For an ODE system, this means that, roughly speaking, as long as all modes in the system are stable, the method is also stable for any choice of step size, at least in the sense of a local linear stability analysis.

At orders 3 to 5, the BDF methods are not A-stable, although they are *stiffly stable*. In each case, in order for the method to be stable at step size h on the scalar model problem, the product $h\lambda$ must lie within a *region of absolute stability*. That region excludes a portion of the left half-plane that is concentrated near the imaginary axis. The size of that region of instability grows as the order increases from 3 to 5. What this means is that, when running BDF at any of these orders, if an eigenvalue λ of the system lies close enough to the imaginary axis, the step sizes h for which the method is stable are limited (at least according to the linear stability theory) to a set that prevents $h\lambda$ from leaving the stability region. The meaning of *close enough* depends on the order. At order 3, the unstable region is much narrower than at order 5, so the potential for unstable behavior grows with order.

System eigenvalues that are likely to run into this instability are ones that correspond to weakly damped oscillations. A pure undamped oscillation corresponds to an eigenvalue on the imaginary axis. Problems with modes of that kind call for different considerations since the oscillation generally must be followed by the solver, but this requires step sizes ($h \sim 1/\nu$, where ν is the frequency) that are stable for BDF anyway. But for a weakly damped oscillatory mode, the oscillation in the solution is eventually damped to the noise level, and at that time it is important that the solver not be restricted to step sizes on the order of $1/\nu$. It is in this situation that the new option may be of great value.

In terms of partial differential equations, the typical problems for which the stability limit detection option is appropriate are ODE systems resulting from semi-discretized PDEs (i.e., PDEs discretized in space) with advection and diffusion, but with advection dominating over diffusion. Diffusion alone produces pure decay modes, while advection tends to produce undamped oscillatory modes. A mix of the two with advection dominant will have weakly damped oscillatory modes.

The STALD algorithm attempts to detect, in a direct manner, the presence of a stability region boundary that is limiting the step sizes in the presence of a weakly damped oscillation [13]. The algorithm supplements (but differs greatly from) the existing algorithms in CVODES for choosing step size and order based on estimated local truncation errors. The STALD algorithm works directly with history data that is readily available in CVODES. If it concludes that the step size is in fact stability-limited, it dictates a reduction in the method order regardless of the outcome of the error-based algorithm. The STALD algorithm has been tested in combination with the VODE solver on linear advection-dominated advection-diffusion problems [14], where it works well. The implementation in CVODES has been successfully tested on linear and nonlinear advection-diffusion problems, among others.

This stability limit detection option adds some computational overhead to the CVODES solution. (In timing tests, these overhead costs have ranged from 2% to 7% of the total, depending on the size and complexity of the problem, with lower relative costs for larger problems.) Therefore, it should be activated only when there is reasonable expectation of modes in the user's system for which it is appropriate. In particular, if a CVODES solution with this option turned off appears to take an inordinately large number of steps for orders between 3 and 5 for no apparent reason in terms of the solution time scale, then there is a good chance that step sizes are being limited by stability, and that turning on the option will improve efficiency.

3.6 Rootfinding

The CVODES solver has been augmented to include a rootfinding feature. This means that, while integrating the Initial Value Problem (3.1), CVODES can also find the roots of a set of user-defined functions $g_i(t, y)$ that depend both on t and on the solution vector $y = y(t)$. 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 $g_i(t, y(t))$, 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 CVODES. 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 hone in on the root(s) with a modified secant method [12]. In addition, each time g is computed, CVODES 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 , CVODES 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, CVODES stops and reports an error. This way, each time CVODES 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, CVODES 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} , 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_i at t_{hi} for zeros and for sign changes in (t_{lo}, t_{hi}) . If no sign changes were 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 include 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 is 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 $\tau/2$.

Chapter 4

Code Organization

4.1 SUNDIALS organization

The family of solvers referred to as SUNDIALS consists of the solvers CVODE (for ODE systems), KINSOL (for nonlinear algebraic systems), and IDA (for differential-algebraic systems). In addition, variants of these which also do sensitivity analysis calculations are available or in development. CVODES, an extension of CVODE that provides both forward and adjoint sensitivity capabilities is available, while IDAS is currently in development.

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. 4.1). The following is a list of the solver packages presently available:

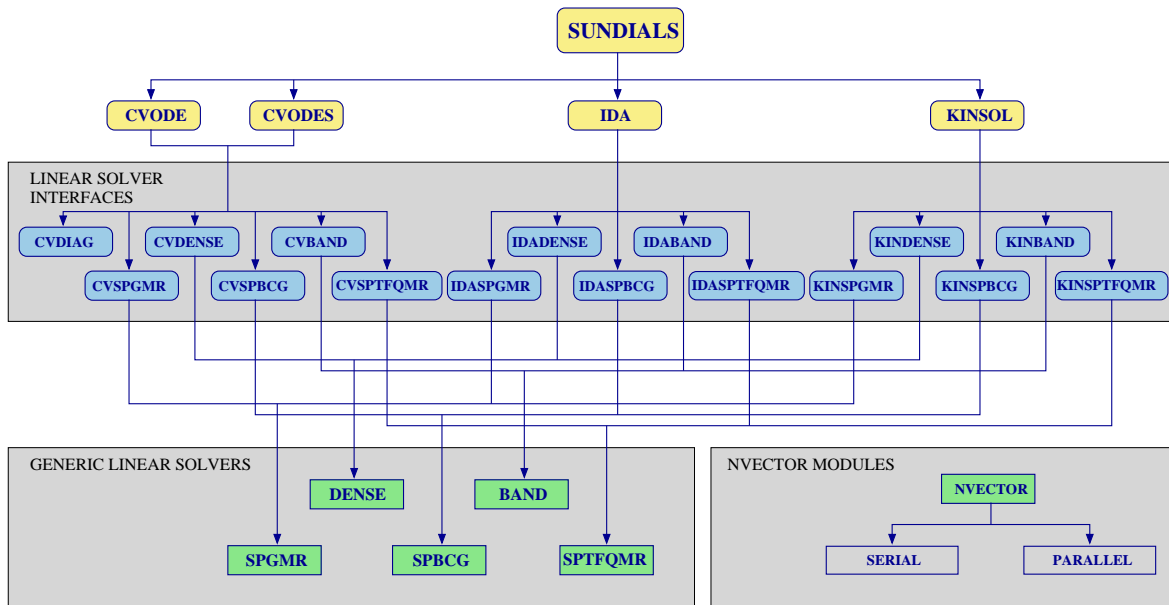
- CVODE, a solver for stiff and nonstiff ODEs $dy/dt = f(t, y)$;
- CVODES, a solver for stiff and nonstiff ODEs $dy/dt = f(t, y, p)$ with sensitivity analysis capabilities;
- KINSOL, a solver for nonlinear algebraic systems $F(u) = 0$;
- IDA, a solver for differential-algebraic systems $F(t, y, y') = 0$.

4.2 CVODES organization


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

The overall organization of the CVODES package is shown in Figure 4.2. The basic elements of the structure are a module for the basic integration algorithm (including forward sensitivity analysis), a module for adjoint sensitivity analysis, and a set of modules for the solution of linear systems that arise in the case of a stiff system. The central integration module, implemented in the files `cvodes.h`, `cvodes_impl.h`, `cvodes.c`, and `cvodea.c`, deals with the evaluation of integration coefficients, the functional or Newton iteration process, estimation of local error, selection of step size 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. CVODES provides three different strategies for dealing with the correction stage for the sensitivity variables: `CV_SIMULTANEOUS`, `CV_STAGGERED`



(a) High-level diagram



(b) Directory structure of the source tree

Figure 4.1: Organization of the SUNDIALS suite




Figure 4.2: Overall structure of the CVODES package. Modules specific to CVODES are distinguished by rounded boxes, while generic solver and auxiliary modules are in rectangular boxes.

and `CV_STAGGERED1` (see §3.2 and §6.2.1). The CVODES package includes an algorithm for the approximation of the sensitivity equations right-hand sides by difference quotients, but the user has the option of supplying these right-hand sides directly.

The adjoint sensitivity module provides the infrastructure needed for the backward integration of any system of ODEs 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 six CVODES linear solver modules:

- `CVDENSE`: LU factorization and backsolving with dense matrices;
- `CVBAND`: LU factorization and backsolving with banded matrices;
- `CVDIAG`: an internally generated diagonal approximation to the Jacobian;
- `CVSPGMR`: scaled preconditioned GMRES method;
- `CVSPBCG`: scaled preconditioned Bi-CGStab method;
- `CVSPTFQMR`: scaled preconditioned TFQMR method.

This set of linear solver modules is intended to be expanded in the future as new algorithms are developed.

In the case of the direct methods `CVDENSE` and `CVBAND`, 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. In the case of the Krylov iterative methods `CVSPGMR`, `CVSPBCG`, and `CVSPTFQMR`, the package includes an algorithm for the approximation of the product between the Jacobian matrix and a vector of appropriate length by difference quotients. Again, the user has the option of providing a routine for this operation. For the Krylov methods, the preconditioner must be supplied by the user in two parts: setup (preprocessing of Jacobian data) and solve. While there is no default choice of preconditioner analogous to the difference-quotient

approximation in the direct case, the references [2, 3], together with the example programs included with CVODES, offer considerable assistance in building preconditioners.

Each CVODES linear solver module consists of four routines devoted to (1) memory allocation and initialization, (2) setup of the matrix data involved, (3) solution of the system, and (4) freeing of memory. The setup and solution phases are separate because the evaluation of Jacobians and preconditioners is done only periodically during the integration, and only as required to achieve convergence. The call list within the central CVODES module for 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 modules CVDENSE, CVBAND, CVSPGMR, CVSPBCG, and CVSPTFQMR is a set of interface routines built on top of a generic solver module, named DENSE, BAND, SPGMR, SPBCG, and SPTFQMR, respectively. The interfaces deal with the use of these methods in the CVODES context, whereas the generic solver is independent of the context. While the generic solvers here were generated with SUNDIALS in mind, our intention is that they be usable in other applications 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 CVODES package elsewhere.

CVODES also provides two preconditioner modules for use with any of the Krylov iterative linear solvers. The first one, CVBANDPRE, is intended to be used with NVECTOR_SERIAL and provides banded difference-quotient Jacobian-based preconditioner and solver routines. The second preconditioner module, CVBBDPRE, 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 CVODES 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 CVODES 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 CVODES memory structure. The reentrancy of CVODES was motivated by the anticipated multicomputer extension, but is also essential in a uniprocessor setting where two or more problems are solved by intermixed calls to the package from within a single user program.

Chapter 5

Using CVODES for IVP Solution

This chapter is concerned with the use of CVODES for the solution of initial value problems (IVPs). The following sections treat the header files and the layout of the user's main program, and provide descriptions of the CVODES user-callable functions and user-supplied functions.

This usage is essentially equivalent to using CVODE [19]. The listings of the sample programs in the companion document [17] may also be helpful. Those codes may be used as templates and are included in the CVODES 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 or direct band linear solvers since these linear solver modules need to form the complete system Jacobian. The following CVODES modules can only be used with NVECTOR_SERIAL: CVDENSE, CVBAND and CVBANDPRE. Also, the preconditioner module CVBBDPRE can only be used with NVECTOR_PARALLEL.

CVODES uses various constants for both input and output. These are defined as needed in this chapter, but for convenience are also listed separately in Chapter 11.

5.1 Access to libraries and header files

At this point, it is assumed that the installation of CVODES, following the procedure described in Chapter 2, 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 CVODES. The relevant library files are

- *libdir/libsundials_cvodes.lib*,
- *libdir/libsundials_nvec*.lib* (one or two 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*
- *incdir/include/cvodes*
- *incdir/include/sundials*

The directories *libdir* and *incdir* are the install library and include directories. For a default installation, these are *builddir/lib* and *builddir/include*, respectively, where *builddir* was defined in Chapter 2.

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

5.2 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 §2.1).

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 the typedef for `realtype` matches this choice). 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 §2.1).

5.3 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:

- `cvodes.h`, the main header file for CVODES, which defines the several types and various constants, and includes function prototypes.

Note that `cvodes.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 §8 for details). For the two `NVECTOR` implementations that are included in the CVODES 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`.

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

Finally, if the user chooses Newton iteration for the solution of the nonlinear systems, then a linear solver module header file will be required. The header files corresponding to the various linear solvers available for use with CVODES are:

- `cvodes_dense.h`, which is used with the dense direct linear solver in the context of CVODES. This in turn includes a header file (`sundials_dense.h`) which defines the `DenseMat` type and corresponding accessor macros;

- `cvodes.band.h`, which is used with the band direct linear solver in the context of CVODES. This in turn includes a header file (`sundials.band.h`) which defines the `BandMat` type and corresponding accessor macros;
- `cvodes.diag.h`, which is used with the diagonal linear solver in the context of CVODES;
- `cvodes.spgmr.h`, which is used with the Krylov solver SPGMR in the context of CVODES;
- `cvodes.spbcgs.h`, which is used with the Krylov solver SPBCG in the context of CVODES;
- `cvodes.sptfqmr.h`; which is used with the Krylov solver SPTFQMR in the context of CVODES;

The header files for the Krylov iterative solvers include `cvodes.spils.h` which defined common functions and which in turn includes a header file (`sundials.iterative.h`) which enumerates the kind of preconditioning and for the choices for the Gram-Schmidt process for SPGMR.

Other headers may be needed, depending upon the choice of preconditioner, etc. In one of the examples in [18], preconditioning is done with a block-diagonal matrix. For this, the header `sundials.smalldense.h` is included.

5.4 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 an ODE 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 CVODES: steps marked with [P] correspond to NVECTOR_PARALLEL, while steps marked with [S] correspond to NVECTOR_SERIAL.

1. [P] Initialize MPI

Call `MPI_Init(&argc, &argv)`; to initialize MPI if used by the user's program. Here `argc` and `argv` are the command line argument counter and array received by `main`, respectively.

2. Set problem dimensions

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

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

3. Set vector of initial values

To set the vector `y0` of initial values, use the appropriate functions defined by a particular NVECTOR implementation. If a `realtype` array `ydata` containing the initial values of y already exists, then make the call:

```
[S] y0 = N_VMake_Serial(N, ydata);
```

```
[P] y0 = N_VMake_Parallel(comm, Nlocal, N, ydata);
```

Otherwise, make the call:

```
[S] y0 = N_VNew_Serial(N);
```

```
[P] y0 = N_VNew_Parallel(comm, Nlocal, N);
```

and load initial values into the structure defined by:

```
[S] NV_DATA_S(y0)
```

```
[P] NV_DATA_P(y0)
```

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

4. Create CVODES object

Call `cvode_mem = CVodeCreate(lmm, iter);` to create the CVODES memory block and to specify the solution method (linear multistep method and nonlinear solver iteration type). `CVodeCreate` returns a pointer to the CVODES memory structure. See §5.5.1 for details.

5. Allocate internal memory

Call `CVodeMalloc(...);` to provide required problem specifications, allocate internal memory for CVODES, and initialize CVODES. `CVodeMalloc` returns a flag, the value of which indicates either success or an illegal argument value. See §5.5.1 for details.

6. Set optional inputs

Call `CVodeSet*` functions to change any optional inputs that control the behavior of CVODES from their default values. See §5.5.5.1 for details.

7. Attach linear solver module

If Newton iteration is chosen, initialize the linear solver module with one of the following calls (for details see §5.5.3):

```
[S] ier = CVDense(...);
[S] ier = CVBand(...);
ier = CVDiag(...);
ier = CVSpqr(...);
ier = CVSpbcg(...);
ier = CVSpqr(...);
```

8. Set linear solver optional inputs

Call `CV*Set*` functions from the selected linear solver module to change optional inputs specific to that linear solver. See §5.5.5 for details.

9. Specify rootfinding problem

Optionally, call `CVodeRootInit` to initialize a rootfinding problem to be solved during the integration of the ODE system. See §5.8.1 for details.

10. Advance solution in time

For each point at which output is desired, call `ier = CVode(cvode_mem, tout, yout, &tret, itask);` Set `itask` to specify the return mode. The vector `y` (which can be the same as the vector `y0` above) will contain $y(t)$. See §5.5.4 for details.

11. Get optional outputs

Call `CV*Get*` functions to obtain optional output. See §5.5.7 and §5.8.1 for details.

12. Deallocate memory for solution vector

Upon completion of the integration, deallocate memory for the vector `y` by calling the destructor function defined by the NVECTOR implementation:

```
[S] N_VDestroy_Serial(y);
[P] N_VDestroy_Parallel(y);
```

13. Free solver memory

Call `CVodeFree(&cvode_mem);` to free the memory allocated for CVODES.

14. [P] Finalize MPI

Call `MPI_Finalize()`; to terminate MPI.

5.5 User-callable functions for IVP solution

This section describes the CVODES functions that are called by the user to setup and then solve an IVP. Some of these are required. However, starting with §5.5.5, the functions listed involve optional inputs/outputs or restarting, and those paragraphs may be skipped for a casual use of CVODES. In any case, refer to §5.4 for the correct order of these calls. Calls related to rootfinding are described in §5.8.

5.5.1 CVODES 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 IVP solution is complete, as it frees the CVODES memory block created and allocated by the first two calls.

CVodeCreate

Call `cvode_mem = CVodeCreate(lmm, iter);`

Description The function `CVodeCreate` instantiates a CVODES solver object and specifies the solution method.

Arguments `lmm` (`int`) specifies the linear multistep method and may be one of two possible values: `CV_ADAMS` or `CV_BDF`.
`iter` (`int`) specifies the type of nonlinear solver iteration and may be either `CV_NEWTON` or `CV_FUNCTIONAL`.

The recommended choices for `(lmm, iter)` are `(CV_ADAMS, CV_FUNCTIONAL)` for nonstiff problems and `(CV_BDF, CV_NEWTON)` for stiff problems.

Return value If successful, `CVodeCreate` returns a pointer to the newly created CVODES memory block (of type `void *`). If an error occurred, `CVodeCreate` prints an error message to `stderr` and returns `NULL`.

CVodeMalloc

Call `flag = CVodeMalloc(cvode_mem, f, t0, y0, itol, reltol, abstol);`

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

Arguments `cvode_mem` (`void *`) pointer to the CVODES memory block returned by `CVodeCreate`.
`f` (`CVRhsFn`) is the C function which computes f in the ODE. This function has the form `f(t, y, ydot, f_data)` (for full details see §5.6.1).
`t0` (`realtype`) is the initial value of t .
`y0` (`N_Vector`) is the initial value of y .
`itol` (`int`) is one of `CV_SS`, `CV_SV`, or `CV_WF`. Here `itol = SS` indicates scalar relative error tolerance and scalar absolute error tolerance, while `itol = CV_SV` indicates scalar relative error tolerance and vector absolute error tolerance. The latter choice is important when the absolute error tolerance needs to be different for each component of the ODE. If `itol = CV_WF`, the arguments `reltol` and `abstol` are ignored and the user is expected to provide a function to evaluate the error weight vector W , replacing (3.7). See `CVodeSetEwtFn` in §5.5.5.1.
`reltol` (`realtype`) is the relative error tolerance.

abstol (void *) is a pointer to the absolute error tolerance. If **itol** = **CV_SS**, **abstol** must be a pointer to a **realtype** variable. If **itol** = **CV_SV**, **abstol** must be an **N_Vector** variable.

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

CV_SUCCESS The call to **CVodeMalloc** was successful.

CV_MEM_NULL The CVODES memory block was not initialized through a previous call to **CVodeCreate**.

CV_MEM_FAIL A memory allocation request has failed.

CV_ILLINPUT An input argument to **CVodeMalloc** has an illegal value.

Notes See also §5.5.2 for advice on tolerances.

The tolerance values in **reltol** and **abstol** may be changed between calls to **CVode** (see **CVodeSetTolerances** in §5.5.5.1).

It is the user's responsibility to provide compatible **itol** and **abstol** arguments.

If an error occurred, **CVodeMalloc** also sends an error message to the error handler function.



CVodeFree

Call **CVodeFree(&cvode_mem);**

Description The function **CVodeFree** frees the memory allocated by a previous call to **CVodeMalloc**.

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

Return value The function **CVodeFree** has no return value.

5.5.2 Advice on choice and use of tolerances

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** = 1.0E-4 means that errors are controlled to .01%. We do not recommend using **reltol** larger than 1.0E-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 1.0E-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 **cvsdex** in the CVODE package, and the discussion of it in the CVODE Examples document [17]. In that problem, the three components vary between 0 and 1, and have different noise levels; hence the **abstol** vector. It is impossible to give any general advice on **abstol** values, because the appropriate noise levels are completely problem-dependent. The user or modeler hopefully has some idea as to what those noise levels are.

(3) Finally, it is important to pick all the tolerance values conservatively, because they control the error committed on each individual time step. The final (global) errors are some sort of accumulation of those per-step errors. A good rule of thumb is to reduce the tolerances by a factor of .01 from the actual desired limits on errors. So if you want .01% accuracy (globally), a good choice is **reltol** = 1.0E-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 y returned by CVODES, with magnitude comparable to `abstol` or less, is equivalent to zero as far as the computation is concerned.

(3) The user's right-hand side routine f should never change a negative value in the solution vector y to a non-negative value, as a "solution" to this problem. This can cause instability. If the f 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 y vector) for the purposes of computing $f(t, y)$.

5.5.3 Linear solver specification functions

As previously explained, Newton iteration requires the solution of linear systems of the form (3.5). There are six CVODES linear solvers currently available for this task: CVDENSE, CVBAND, CVDIAG, CVSPGMR, CVSPBCG, and CVSPTFQMR. The first three are direct solvers and their names indicate the type of approximation used for the Jacobian $J = \partial f / \partial y$; CVDENSE, CVBAND, and CVDIAG work with dense, banded, and diagonal approximations to J , respectively. The last three CVODES linear solvers — CVSPGMR, CVSPBCG, and CVSPTFQMR — are Krylov iterative solvers, which use scaled preconditioned GMRES, scaled preconditioned Bi-CGStab, and scaled preconditioned TFQMR, respectively. Together, they are referred to as CVSPILS (from scaled preconditioned iterative linear solvers).

With any of the Krylov methods, preconditioning can be done on the left only, on the right only, on both the left and the right, or not at all. For the specification of a preconditioner, see the iterative linear solver sections in §5.5.5 and §5.6.

If preconditioning is done, user-supplied functions define left and right preconditioner matrices P_1 and P_2 (either of which could be the identity matrix), such that the product $P_1 P_2$ approximates the Newton matrix $M = I - \gamma J$ of (3.6).

To specify a CVODES linear solver, after the call to `CVodeCreate` but before any calls to `CVode`, the user's program must call one of the functions `CVDense`, `CVBand`, `CVDiag`, `CVSpgmr`, `CVSpbcg`, or `CVSptfqmr`, as documented below. The first argument passed to these functions is the CVODES memory pointer returned by `CVodeCreate`. A call to one of these functions links the main CVODES integrator to a linear solver and allows the user to specify parameters which are specific to a particular solver, such as the half-bandwidths in the CVBAND 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 except the diagonal approximation case CVDIAG, the linear solver module used by CVODES is actually built on top of a generic linear system solver, which may be of interest in itself. These generic solvers, denoted DENSE, BAND, SPGMR, SPBCG, and SPTFQMR, are described separately in Chapter 10.

CVDense

Call `flag = CVDense(cvode_mem, N);`

Description The function `CVDense` selects the CVDENSE linear solver.

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

Arguments `cvode_mem` (void *) pointer to the CVODES memory block.

`N` (long int) problem dimension.

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

`CVDENSE_SUCCESS` The CVDENSE initialization was successful.
`CVDENSE_MEM_NULL` The `cvode_mem` pointer is NULL.
`CVDENSE_ILL_INPUT` The CVDENSE solver is not compatible with the current NVECTOR module.
`CVDENSE_MEM_FAIL` A memory allocation request failed.

Notes The CVDENSE linear solver may not be compatible with a particular implementation of the NVECTOR module. Of the two NVECTOR modules provided with SUNDIALS, only NVECTOR_SERIAL is compatible.

CVBand

Call `flag = CVBand(cvode_mem, N, mupper, mlower);`

Description The function CVBand selects the CVBAND linear solver.
The user's main program must include the `cvodes_band.h` header file.

Arguments `cvode_mem` (`void *`) pointer to the CVODES 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:

`CVBAND_SUCCESS` The CVBAND initialization was successful.
`CVBAND_MEM_NULL` The `cvode_mem` pointer is NULL.
`CVBAND_ILL_INPUT` The CVBAND solver is not compatible with the current NVECTOR module, or one of the Jacobian half-bandwidths is outside of its valid range ($0 \dots N-1$).
`CVBAND_MEM_FAIL` A memory allocation request failed.

Notes The CVBAND linear solver may not be compatible with a particular implementation of the NVECTOR module. Of the two NVECTOR modules provided with SUNDIALS, only NVECTOR_SERIAL is compatible. The half-bandwidths are to be set such that the nonzero locations (i, j) in the banded (approximate) Jacobian satisfy $-\text{mlower} \leq j - i \leq \text{mupper}$.

CVDiag

Call `flag = CVDiag(cvode_mem);`

Description The function CVDiag selects the CVDIAG linear solver.
The user's main function must include the `cvodes_diag.h` header file.

Arguments `cvode_mem` (`void *`) pointer to the CVODES memory block.

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

`CVDIAG_SUCCESS` The CVDIAG initialization was successful.
`CVDIAG_MEM_NULL` The `cvode_mem` pointer is NULL.
`CVDIAG_ILL_INPUT` The CVDIAG solver is not compatible with the current NVECTOR module.
`CVDIAG_MEM_FAIL` A memory allocation request failed.

Notes The CVDIAG solver is the simplest of all of the current CVODES linear solvers. The CVDIAG solver uses an approximate diagonal Jacobian formed by way of a difference quotient. The user does *not* have the option of supplying a function to compute an approximate diagonal Jacobian.

CVSpgmr

Call `flag = CVSpgmr(cvode_mem, pretype, maxl);`

Description The function CVSpgmr selects the CVSPGMR linear solver.
The user's main function must include the `cvodes_spgmr.h` header file.

Arguments

- `cvode_mem` (`void *`) pointer to the CVODES memory block.
- `pretype` (`int`) specifies the preconditioning type and must be one of: `PREC_NONE`, `PREC_LEFT`, `PREC_RIGHT`, or `PREC_BOTH`.
- `maxl` (`int`) maximum dimension of the Krylov subspace to be used. Pass 0 to use the default value `CVSPILS_MAXL = 5`.

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

- `CVSPILS_SUCCESS` The CVSPGMR initialization was successful.
- `CVSPILS_MEM_NULL` The `cvode_mem` pointer is NULL.
- `CVSPILS_ILL_INPUT` The preconditioner type `pretype` is not valid.
- `CVSPILS_MEM_FAIL` A memory allocation request failed.

Notes The CVSPGMR solver uses a scaled preconditioned GMRES iterative method to solve the linear system (3.5).

CVSpgbcg

Call `flag = CVSpgbcg(cvode_mem, pretype, maxl);`

Description The function CVSpgbcg selects the CVSPBCG linear solver.
The user's main function must include the `cvodes_spgbcgs.h` header file.

Arguments

- `cvode_mem` (`void *`) pointer to the CVODES memory block.
- `pretype` (`int`) specifies the preconditioning type and must be one of: `PREC_NONE`, `PREC_LEFT`, `PREC_RIGHT`, or `PREC_BOTH`.
- `maxl` (`int`) maximum dimension of the Krylov subspace to be used. Pass 0 to use the default value `CVSPILS_MAXL = 5`.

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

- `CVSPILS_SUCCESS` The CVSPBCG initialization was successful.
- `CVSPILS_MEM_NULL` The `cvode_mem` pointer is NULL.
- `CVSPILS_ILL_INPUT` The preconditioner type `pretype` is not valid.
- `CVSPILS_MEM_FAIL` A memory allocation request failed.

Notes The CVSPBCG solver uses a scaled preconditioned Bi-CGStab iterative method to solve the linear system (3.5).

CVSptfqmr

Call `flag = CVSptfqmr(cvode_mem, pretype, maxl);`

Description The function CVSptfqmr selects the CVSPTFQMR linear solver.
The user's main function must include the `cvodes_sptfqmr.h` header file.

Arguments `cvode_mem` (`void *`) pointer to the CVODES memory block.

pretype	(int) specifies the preconditioning type and must be one of: <code>PREC_NONE</code> , <code>PREC_LEFT</code> , <code>PREC_RIGHT</code> , or <code>PREC_BOTH</code> .
maxl	(int) maximum dimension of the Krylov subspace to be used. Pass 0 to use the default value <code>CVSPILS_MAXL = 5</code> .
Return value	The return value flag (of type <code>int</code>) is one of <ul style="list-style-type: none"> <code>CVSPILS_SUCCESS</code> The CVSPTFQMR initialization was successful. <code>CVSPILS_MEM_NULL</code> The <code>cnode_mem</code> pointer is NULL. <code>CVSPILS_ILL_INPUT</code> The preconditioner type pretype is not valid. <code>CVSPILS_MEM_FAIL</code> A memory allocation request failed.
Notes	The CVSPTFQMR solver uses a scaled preconditioned TFQMR iterative method to solve the linear system (3.5).

5.5.4 CVODE solver function

This is the central step in the solution process - the call to perform the integration of the IVP.

CVode	
Call	<code>flag = CVode(cnode_mem, tout, yout, tret, itask);</code>
Description	The function <code>CVode</code> integrates the ODE over an interval in t .
Arguments	<p><code>cnode_mem</code> (void *) pointer to the CVODES memory block.</p> <p><code>tout</code> (realtype) the next time at which a computed solution is desired.</p> <p><code>yout</code> (N_Vector) the computed solution vector.</p> <p><code>tret</code> (realtype *) the time reached by the solver.</p> <p><code>itask</code> (int) a flag indicating the job of the solver for the next user step. The <code>CV_NORMAL</code> option causes the solver to 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 $y(tout)$. The <code>CV_ONE_STEP</code> option tells the solver to take just one internal step and then return the solution at the point reached by that step. The <code>CV_NORMAL_TSTOP</code> and <code>CV_ONE_STEP_TSTOP</code> modes are similar to <code>CV_NORMAL</code> and <code>CV_ONE_STEP</code>, respectively, except that the integration never proceeds past the value <code>tstop</code> (specified through the function <code>CVodeSetStopTime</code>).</p>
Return value	<p>On return, <code>CVode</code> returns a vector <code>yout</code> and a corresponding independent variable value $t = *tret$, such that <code>yout</code> is the computed value of $y(t)$.</p> <p>In <code>CV_NORMAL</code> mode (with no errors), <code>*tret</code> will be equal to <code>tout</code> and <code>yout = y(tout)</code>.</p> <p>The return value flag (of type <code>int</code>) will be one of the following:</p> <ul style="list-style-type: none"> <code>CV_SUCCESS</code> <code>CVode</code> succeeded and no roots were found. <code>CV_TSTOP_RETURN</code> <code>CVode</code> succeeded by reaching the stopping point specified through the optional input function <code>CVodeSetStopTime</code> (see §5.5.5.1). <code>CV_ROOT_RETURN</code> <code>CVode</code> succeeded and found one or more roots. If <code>nrtfn > 1</code>, call <code>CVodeGetRootInfo</code> to see which g_i were found to have a root. See §5.8 for more information. <code>CV_MEM_NULL</code> The <code>cnode_mem</code> argument was NULL. <code>CV_NO_MALLOC</code> The CVODES memory was not allocated by a call to <code>CVodeMalloc</code>. <code>CV_ILL_INPUT</code> One of the inputs to <code>CVode</code> is illegal. This includes the situation where a root of one of the root functions was found both at a point t and also very near t. It also includes the situation where a component of the error weight vector becomes negative during internal time-stepping. The <code>CV_ILL_INPUT</code> flag will also be returned if the

	linear solver initialization function (called by the user after calling <code>CVodeCreate</code>) failed to set the linear solver-specific <code>lsolve</code> field in <code>cvode_mem</code> . In any case, the user should see the error message for details.
<code>CV_TOO_CLOSE</code>	The initial time t_0 and the final time t_{out} are too close to each other and the user did not specify an initial step size.
<code>CV_TOO_MUCH_WORK</code>	The solver took <code>mxstep</code> internal steps but still could not reach <code>tout</code> . The default value for <code>mxstep</code> is <code>MXSTEP_DEFAULT</code> = 500.
<code>CV_TOO_MUCH_ACC</code>	The solver could not satisfy the accuracy demanded by the user for some internal step.
<code>CV_ERR_FAILURE</code>	Either error test failures occurred too many times (<code>MXNEF</code> = 7) during one internal time step, or with $ h = h_{min}$.
<code>CV_CONV_FAILURE</code>	Either convergence test failures occurred too many times (<code>MXNCF</code> = 10) during one internal time step, or with $ h = h_{min}$.
<code>CV_LINIT_FAIL</code>	The linear solver's initialization function failed.
<code>CV_LSETUP_FAIL</code>	The linear solver's setup function failed in an unrecoverable manner.
<code>CV_LSOLVE_FAIL</code>	The linear solver's solve function failed in an unrecoverable manner.
<code>CV_RHSFUNC_FAIL</code>	The right-hand side function failed in an unrecoverable manner.
<code>CV_FIRST_RHSFUNC_FAIL</code>	The right-hand side function had a recoverable error at the first call.
<code>CV_REPTD_RHSFUNC_ERR</code>	Convergence tests occurred too many times due to repeated recoverable errors in the right-hand side function. The <code>CV_REPTD_RHSFUNC_ERR</code> will also be returned if the right-hand side function had repeated recoverable errors during the estimation of an initial step size.
<code>CV_UNREC_RHSFUNC_ERR</code>	The right-hand function had a recoverable error, but no recovery was possible. This failure mode is rare, as it can occur only if the right-hand side function fails recoverably after an error test failed while at order one.
<code>CV_RTFUNC_FAIL</code>	The rootfinding function failed.
Notes	<p>The vector <code>yout</code> can occupy the same space as the <code>y0</code> vector of initial conditions that was passed to <code>CVodeMalloc</code>.</p> <p>In the <code>CV_ONE_STEP</code> mode, <code>tout</code> is only used on the first call to get the direction and a rough scale of the independent variable.</p> <p>All failure return values are negative and so the test <code>ier < 0</code> will trap all <code>CVode</code> failures.</p> <p>On any error return in which one or more internal steps were taken by <code>CVode</code>, the returned values of <code>tret</code> and <code>yout</code> correspond to the farthest point reached in the integration. On all other error returns, <code>tret</code> and <code>yout</code> are left unchanged from the previous <code>CVode</code> return.</p>

5.5.5 Optional input functions

CVODES provides an extensive set of functions that can be used to change from their default values various optional input parameters that control the behavior of the CVODES solver. Table 5.1 lists all optional input functions in CVODES which are then described in detail in the remainder of this section, beginning with those for the main CVODES solver and continuing with those for the linear solver modules. Note that the diagonal linear solver module has no optional inputs. For the most casual use of CVODES, the reader can skip to §5.6.

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

Table 5.1: Optional inputs for CVODES, CVDENSE, CVBAND, and CVSPILS

Optional input	Function name	Default
CVODE main solver		
Error handler function	CVodeSetErrHandlerFn	internal fn.
Pointer to an error file	CVodeSetErrFile	stderr
Data for right-hand side function	CVodeSetFdata	NULL
Maximum order for BDF method	CVodeSetMaxOrd	5
Maximum order for Adams method	CVodeSetMaxOrd	12
Maximum no. of internal steps before t_{out}	CVodeSetMaxNumSteps	500
Maximum no. of warnings for $t_n + h = t_n$	CVodeSetMaxHnilWarns	10
Flag to activate stability limit detection	CVodeSetStabLimDet	FALSE
Initial step size	CVodeSetInitStep	estimated
Minimum absolute step size	CVodeSetMinStep	0.0
Maximum absolute step size	CVodeSetMaxStep	∞
Value of t_{stop}	CVodeSetStopTime	undefined
Maximum no. of error test failures	CVodeSetMaxErrTestFails	7
Maximum no. of nonlinear iterations	CVodeSetMaxNonlinIters	3
Maximum no. of convergence failures	CVodeSetMaxConvFails	10
Coefficient in the nonlinear convergence test	CVodeSetNonlinConvCoef	0.1
Nonlinear iteration type	CVodeSetIterType	none
Integration tolerances	CVodeSetTolerances	none
Ewt computation function	CVodeSetEwtFn	internal fn.
CVDENSE linear solver		
Dense Jacobian function and data	CVDenseSetJacFn	internal DQ, NULL
CVBAND linear solver		
Band Jacobian function and data	CVBandSetJacFn	internal DQ, NULL
CVSPILS linear solvers		
Preconditioner functions and data	CVSpilsSetPreconditioner	all NULL
Jacobian-times-vector function and data	CVSpilsSetJacTimesVecFn	internal DQ, NULL
Preconditioning type	CVSpilsSetPrecType	none
Ratio between linear and nonlinear tolerances	CVSpilsSetDelt	0.05
Type of Gram-Schmidt orthogonalization ^(a)	CVSpilsSetGSType	classical GS
Maximum Krylov subspace size ^(b)	CVSpilsSetMaxl	5

^(a) Only for CVSPGMR^(b) Only for CVSPBCG and CVSPTFQMR

5.5.5.1 Main solver optional input functions

The calls listed here can be executed in any order.

However, if `CVodeSetErrHandlerFn` or `CVodeSetErrFile` are to be called, that call should be first, in order to take effect for any later error message.

`CVodeSetErrHandlerFn`

Call	<code>flag = CVodeSetErrHandlerFn(cvode_mem, ehfun, eh_data);</code>
Description	The function <code>CVodeSetErrHandlerFn</code> specifies the optional user-defined function to be used in handling error messages.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODES memory block. <code>ehfun</code> (CErrorHandlerFn) is the C error handler function (see §5.6.2). <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>CV_SUCCESS</code> The function <code>ehfun</code> and data pointer <code>eh_data</code> have been successfully set. <code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL.
Notes	The default internal error handler function directs error messages to the file specified by the file pointer <code>errfp</code> (see <code>CVodeSetErrFile</code> below). Error messages indicating that the CVODES solver memory is NULL will always be directed to <code>stderr</code> .

`CVodeSetErrFile`

Call	<code>flag = CVodeSetErrFile(cvode_mem, errfp);</code>
Description	The function <code>CVodeSetErrFile</code> specifies a pointer to the file where all CVODES messages should be directed in case the default CVODES error handler function is used.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODES 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>CV_SUCCESS</code> The optional value has been successfully set. <code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL.
Notes	The default value for <code>errfp</code> is <code>stderr</code> . Passing a value of NULL disables all future error message output (except for the case in which the CVODES memory pointer is NULL). If <code>CVodeSetErrFile</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.



`CVodeSetFdata`

Call	<code>flag = CVodeSetFdata(cvode_mem, f_data);</code>
Description	The function <code>CVodeSetFdata</code> specifies the user-defined data block <code>f_data</code> to be passed to the user-supplied right-hand side function <code>f</code> , and attaches it to the main CVODES memory block.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODES memory block. <code>f_data</code> (void *) pointer to the user-defined data block.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CV_SUCCESS</code> The optional value has been successfully set. <code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL.
Notes	If <code>f_data</code> is not specified, a NULL pointer is passed to the <code>f</code> function.

CVodeSetMaxOrd

Call	<code>flag = CVodeSetMaxOrder(cvode_mem, maxord);</code>
Description	The function <code>CVodeSetMaxOrder</code> specifies the maximum order of the linear multistep method.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODES memory block. <code>maxord</code> (int) value of the maximum method order.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CV_SUCCESS</code> The optional value has been successfully set. <code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL. <code>CV_ILL_INPUT</code> The specified value <code>maxord</code> is negative or larger than its previous value.
Notes	The default value is <code>ADAMS_Q_MAX = 12</code> for the Adams-Moulton method and <code>BDF_Q_MAX = 5</code> for the BDF method. Since <code>maxord</code> affects the memory requirements for the internal CVODES memory block, its value cannot be increased past its previous value.

CVodeSetMaxNumSteps

Call	<code>flag = CVodeSetMaxNumSteps(cvode_mem, mxsteps);</code>
Description	The function <code>CVodeSetMaxNumSteps</code> specifies the maximum number of steps to be taken by the solver in its attempt to reach the next output time.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODES memory block. <code>mxsteps</code> (long int) maximum allowed number of steps.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CV_SUCCESS</code> The optional value has been successfully set. <code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL. <code>CV_ILL_INPUT</code> <code>mxsteps</code> is non-positive.
Notes	Passing <code>mxsteps=0</code> results in CVODES using the default value (500).

CVodeSetMaxHnilWarns

Call	<code>flag = CVodeSetMaxHnilWarns(cvode_mem, mxhnil);</code>
Description	The function <code>CVodeSetMaxHnilWarns</code> specifies the maximum number of messages issued by the solver warning that $t + h = t$ on the next internal step.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODES memory block. <code>mxhnil</code> (int) maximum number of warning messages
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CV_SUCCESS</code> The optional value has been successfully set. <code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL.
Notes	The default value is 10. A negative value for <code>mxhnil</code> indicates that no warning messages should be issued.

CVodeSetStabLimDet

Call	<code>flag = CVodeSetstabLimDet(cvode_mem, stldet);</code>
Description	The function <code>CVodeSetStabLimDet</code> indicates if the BDF stability limit detection algorithm should be used. See §3.5 for further details.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODES memory block. <code>stldet</code> (booleantype) flag controlling stability limit detection (TRUE = on; FALSE = off).

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

- `CV_SUCCESS` The optional value has been successfully set.
- `CV_MEM_NULL` The `cvode_mem` pointer is `NULL`.
- `CV_ILL_INPUT` The linear multistep method is not set to `CV_BDF`.

Notes The default value is `FALSE`. If `stldet = TRUE` when BDF is used and the method order is greater than or equal to 3, then an internal function, `CVsldet`, is called to detect a possible stability limit. If such a limit is detected, then the order is reduced.

CVodeSetInitStep

Call `flag = CVodeSetInitStep(cvode_mem, hin);`

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

Arguments `cvode_mem` (`void *`) pointer to the CVODES memory block.
`hin` (`realtype`) value of the initial step size.

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

- `CV_SUCCESS` The optional value has been successfully set.
- `CV_MEM_NULL` The `cvode_mem` pointer is `NULL`.

Notes By default, CVODES estimates the initial step size to be the solution h of the equation $\|0.5h^2\ddot{y}\|_{\text{WRMS}} = 1$, where \ddot{y} is an estimated second derivative of the solution at `t0`.

CVodeSetMinStep

Call `flag = CVodeSetMinStep(cvode_mem, hmin);`

Description The function `CVodeSetMinStep` specifies a lower bound on the magnitude of the step size.

Arguments `cvode_mem` (`void *`) pointer to the CVODES memory block.
`hmin` (`realtype`) minimum absolute value of the step size.

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

- `CV_SUCCESS` The optional value has been successfully set.
- `CV_MEM_NULL` The `cvode_mem` pointer is `NULL`.
- `CV_ILL_INPUT` Either `hmin` is nonpositive or it exceeds the maximum allowable step size.

Notes The default value is 0.0.

CVodeSetMaxStep

Call `flag = CVodeSetMaxStep(cvode_mem, hmax);`

Description The function `CVodeSetMaxStep` specifies an upper bound on the magnitude of the step size.

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

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

- `CV_SUCCESS` The optional value has been successfully set.
- `CV_MEM_NULL` The `cvode_mem` pointer is `NULL`.
- `CV_ILL_INPUT` Either `hmax` is nonpositive or it is smaller than the minimum allowable step size.

Notes Pass `hmax=0` to obtain the default value ∞ .

CVodeSetStopTime

Call	<code>flag = CVodeSetStopTime(cvode_mem, tstop);</code>
Description	The function <code>CVodeSetStopTime</code> specifies the value of the independent variable t past which the solution is not to proceed.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODES memory block. <code>tstop</code> (realtype) value of the independent variable past which the solution should not proceed.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CV_SUCCESS</code> The optional value has been successfully set. <code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL.
Notes	The default value is ∞ .

CVodeSetMaxErrTestFails

Call	<code>flag = CVodeSetMaxErrTestFails(cvode_mem, maxnef);</code>
Description	The function <code>CVodeSetMaxErrTestFails</code> specifies the maximum number of error test failures permitted in attempting one step.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODES memory block. <code>maxnef</code> (int) maximum number of error test failures allowed on one step.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CV_SUCCESS</code> The optional value has been successfully set. <code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL.
Notes	The default value is 7.

CVodeSetMaxNonlinIters

Call	<code>flag = CVodeSetMaxNonlinIters(cvode_mem, maxcor);</code>
Description	The function <code>CVodeSetMaxNonlinIters</code> specifies the maximum number of nonlinear solver iterations permitted per step.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODES memory block. <code>maxcor</code> (int) maximum number of nonlinear solver iterations allowed per step.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CV_SUCCESS</code> The optional value has been successfully set. <code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL.
Notes	The default value is 3.

CVodeSetMaxConvFails

Call	<code>flag = CVodeSetMaxConvFails(cvode_mem, maxncf);</code>
Description	The function <code>CVodeSetMaxConvFails</code> specifies the maximum number of nonlinear solver convergence failures permitted during one step.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODES memory block. <code>maxncf</code> (int) maximum number of allowable nonlinear solver convergence failures per step.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CV_SUCCESS</code> The optional value has been successfully set. <code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL.
Notes	The default value is 10.

CVodeSetNonlinConvCoef

Call	<code>flag = CVodeSetNonlinConvCoef(cvode_mem, nlscoef);</code>
Description	The function <code>CVodeSetNonlinConvCoef</code> specifies the safety factor used in the nonlinear convergence test (see §3.1).
Arguments	<code>cvode_mem</code> (<code>void *</code>) pointer to the CVODES memory block. <code>nlscoef</code> (<code>realtype</code>) coefficient in nonlinear convergence test.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CV_SUCCESS</code> The optional value has been successfully set. <code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is <code>NULL</code> .
Notes	The default value is 0.1.

CVodeSetIterType

Call	<code>flag = CVodeSetIterType(cvode_mem, iter);</code>
Description	The function <code>CVodeSetIterType</code> resets the nonlinear solver iteration type to <code>iter</code> .
Arguments	<code>cvode_mem</code> (<code>void *</code>) pointer to the CVODES memory block. <code>iter</code> (<code>int</code>) specifies the type of nonlinear solver iteration and may be either <code>CV_NEWTON</code> or <code>CV_FUNCTIONAL</code> .
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CV_SUCCESS</code> The optional value has been successfully set. <code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is <code>NULL</code> . <code>CV_ILL_INPUT</code> The <code>iter</code> value passed is neither <code>CV_NEWTON</code> nor <code>CV_FUNCTIONAL</code> .
Notes	The nonlinear solver iteration type is initially specified in the call to <code>CVodeCreate</code> (see §5.5.1). This function call is needed only if <code>iter</code> is being changed from its value in the prior call to <code>CVodeCreate</code> .

CVodeSetTolerances

Call	<code>flag = CVodeSetTolerances(cvode_mem, itol, reltol, abstol);</code>
Description	The function <code>CVodeSetTolerances</code> resets the integration tolerances.
Arguments	<code>cvode_mem</code> (<code>void *</code>) pointer to the CVODES memory block. <code>itol</code> (<code>int</code>) is either <code>CV_SS</code> or <code>CV_SV</code> , where <code>itol = CV_SS</code> indicates scalar relative error tolerance and scalar absolute error tolerance, while <code>itol=CV_SV</code> indicates scalar relative error tolerance and vector absolute error tolerance. The latter choice is important when the absolute error tolerance needs to be different for each component of the ODE. <code>reltol</code> (<code>realtype</code>) the relative error tolerance. <code>abstol</code> (<code>void *</code>) is a pointer to the absolute error tolerance. If <code>itol=CV_SS</code> , <code>abstol</code> must be a pointer to a <code>realtype</code> variable. If <code>itol=CV_SV</code> , <code>abstol</code> must be an <code>N_Vector</code> variable.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CV_SUCCESS</code> The tolerances have been successfully set. <code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is <code>NULL</code> . <code>CV_ILL_INPUT</code> An input argument has an illegal value.
Notes	The integration tolerances are initially specified in the call to <code>CVodeMalloc</code> (see §5.5.1). This function call is needed only if the tolerances are being changed from their values between successive calls to <code>CVode</code> .

It is the user's responsibility to provide compatible `itol` and `abstol` arguments.

It is illegal to call `CVodeSetTolerances` before a call to `CVodeMalloc`.



CVodeSetEwtFn

Call `flag = CVodeSetEwtFn(cvode_mem, efun, e_data);`

Description The function `CVodeSetEwtFn` specifies the user-defined function to be used in computing the error weight vector W , which is normally defined by Eq.(3.7).

Arguments `cvode_mem` (void *) pointer to the CVODES memory block.

`efun` (CWEwtFn) is the C function which defines the `ewt` vector (see §5.6.3).

`e_data` (void *) pointer to user data passed to `efun` every time it is called.

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

`CV_SUCCESS` The function `efun` and data pointer `e_data` have been successfully set.

`CV_MEM_NULL` The `cvode_mem` pointer is NULL.

Notes This function can be called between successive calls to `CVode`.

If not needed, pass NULL for `edata`.

It is illegal to call `CVodeSetEwtFn` before a call to `CVodeMalloc`.



5.5.5.2 Dense linear solver

The CVDENSE solver needs a function to compute a dense approximation to the Jacobian matrix $J(t, y)$. This function must be of type `CVDenseJacFn`. The user can supply his/her own dense Jacobian function, or use the default difference quotient function `CVDenseDQJac` that comes with the CVDENSE solver. To specify a user-supplied Jacobian function `djac` and associated user data `jac_data`, CVDENSE provides the function `CVDenseSetJacFn`. The CVDENSE solver passes the pointer `jac_data` to its 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 `jac_data` may be identical to `f_data`, if the latter was specified through `CVodeSetFdata`.

CVDenseSetJacFn

Call `flag = CVDenseSetJacFn(cvode_mem, djac, jac_data);`

Description The function `CVDenseSetJacFn` specifies the dense Jacobian approximation function to be used and the pointer to user data.

Arguments `cvode_mem` (void *) pointer to the CVODES memory block.

`djac` (CVDenseJacFn) user-defined dense Jacobian approximation function.

`jac_data` (void *) pointer to the user-defined data structure.

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

`CVDENSE_SUCCESS` The optional value has been successfully set.

`CVDENSE_MEM_NULL` The `cvode_mem` pointer is NULL.

`CVDENSE_LMEM_NULL` The CVDENSE linear solver has not been initialized.

Notes By default, CVDENSE uses the difference quotient function `CVDenseDQJac`. If NULL is passed to `djac`, this default function is used.

The function type `CVDenseJacFn` is described in §5.6.4.

5.5.5.3 Band linear solver

The CVBAND solver needs a function to compute a banded approximation to the Jacobian matrix $J(t, y)$. This function must be of type `CVBandJacFn`. The user can supply his/her own banded Jacobian approximation function, or use the default difference quotient function `CVBandDQJac` that comes with the CVBAND solver. To specify a user-supplied Jacobian function `bjac` and associated user data `jac_data`, CVBAND provides the function `CVBandSetJacFn`. The CVBAND solver passes the pointer `jac_data` to its 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 `jac_data` may be identical to `f_data`, if the latter was specified through `CVodeSetFdata`.

CVBandSetJacFn

Call	<code>flag = CVBandSetJacFn(cvode_mem, bjac, jac_data);</code>
Description	The function <code>CVBandSetJacFn</code> specifies the banded Jacobian approximation function to be used and the pointer to user data.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODES memory block. <code>bjac</code> (CVBandJacFn) user-defined banded Jacobian approximation function. <code>jac_data</code> (void *) pointer to the user-defined data structure.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CVBAND_SUCCESS</code> The optional value has been successfully set. <code>CVBAND_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL. <code>CVBAND_LMEM_NULL</code> The CVBAND linear solver has not been initialized.
Notes	By default, CVBAND uses the difference quotient function <code>CVBandDQJac</code> . If NULL is passed to <code>bjac</code> , this default function is used. The function type <code>CVBandJacFn</code> is described in §5.6.5.

5.5.5.4 SPILS linear solver

If any preconditioning is to be done within one of the CVSPILS linear solvers, then the user must supply a preconditioner solve function `psolve` and specify its name in a call to `CVSpilsSetPreconditioner`.

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 §5.6. If used, the `psetup` function should also be specified in the call to `CVSpilsSetPreconditioner`. Optionally, a CVSPILS solver passes the pointer `p_data` received through `CVSpilsSetPreconditioner` 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 pointer `p_data` may be identical to `f_data`, if the latter was specified through `CVodeSetFdata`.

The CVSPILS 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 difference quotient function `CVSpilsDQJtimes` that comes with the CVSPILS solvers. A user-defined Jacobian-vector function must be of type `CVSpilsJacTimesVecFn` and can be specified through a call to `CVSpilsSetJacTimesVecFn` (see §5.6.6 for specification details). As with the preconditioner user data structure `p_data`, the user can also specify, in the call to `CVSpilsSetJacTimesVecFn`, a pointer to a user-defined data structure, `jac_data`, which the CVSPILS solver passes to the Jacobian-times-vector function `jtimes` each time it is called. The pointer `jac_data` may be identical to `p_data` and/or `f_data`.

CVSpilsSetPreconditioner

Call	<code>flag = CVSpilsSetPreconditioner(cvode_mem, psetup, psolve, p_data);</code>
------	--

Description	The function <code>CVSpilsSetPreconditioner</code> specifies the preconditioner setup and solve functions and the pointer to user data.
Arguments	<p><code>cvode_mem</code> (void *) pointer to the CVODES memory block.</p> <p><code>psetup</code> (CVSpilsPrecSetupFn) user-defined preconditioner setup function.</p> <p><code>psolve</code> (CVSpilsPrecSolveFn) user-defined preconditioner solve function.</p> <p><code>p_data</code> (void *) pointer to the user-defined data structure.</p>
Return value	<p>The return value <code>flag</code> (of type <code>int</code>) is one of</p> <p><code>CVSPILS_SUCCESS</code> The optional value has been successfully set.</p> <p><code>CVSPILS_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL.</p> <p><code>CVSPILS_LMEM_NULL</code> The CVSPILS linear solver has not been initialized.</p>
Notes	The function type <code>CVSpilsPrecSolveFn</code> is described in §5.6.7. The function type <code>CVSpilsPrecSetupFn</code> is described in §5.6.8.

CVSpilsSetJacTimesVecFn

Call	<code>flag = CVSpilsSetJacTimesVecFn(cvode_mem, jtimes, jac_data);</code>
Description	The function <code>CVSpilsSetJacTimesFn</code> specifies the Jacobian-vector function to be used and the pointer to user data.
Arguments	<p><code>cvode_mem</code> (void *) pointer to the CVODES memory block.</p> <p><code>jtimes</code> (CVSpilsJacTimesVecFn) user-defined Jacobian-vector product function.</p> <p><code>jac_data</code> (void *) pointer to the user-defined data structure.</p>
Return value	<p>The return value <code>flag</code> (of type <code>int</code>) is one of</p> <p><code>CVSPILS_SUCCESS</code> The optional value has been successfully set.</p> <p><code>CVSPILS_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL.</p> <p><code>CVSPILS_LMEM_NULL</code> The CVSPILS linear solver has not been initialized.</p>
Notes	<p>By default, the CVSPILS linear solvers use an internal difference quotient function <code>CVSpilsDQJtimes</code>. If NULL is passed to <code>jtimes</code>, this default function is used.</p> <p>The function type <code>CVSpilsJacTimesVecFn</code> is described in §5.6.6.</p>

CVSpilsSetPrecType

Call	<code>flag = CVSpilsSetPrecType(cvode_mem, pretype);</code>
Description	The function <code>CVSpilsSetPrecType</code> resets the type of preconditioning to be used.
Arguments	<p><code>cvode_mem</code> (void *) pointer to the CVODES memory block.</p> <p><code>pretype</code> (int) specifies the type of preconditioning and must be one of: <code>PREC_NONE</code>, <code>PREC_LEFT</code>, <code>PREC_RIGHT</code>, or <code>PREC_BOTH</code>.</p>
Return value	<p>The return value <code>flag</code> (of type <code>int</code>) is one of</p> <p><code>CVSPILS_SUCCESS</code> The optional value has been successfully set.</p> <p><code>CVSPILS_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL.</p> <p><code>CVSPILS_LMEM_NULL</code> The CVSPILS linear solver has not been initialized.</p> <p><code>CVSPILS_ILL_INPUT</code> The preconditioner type <code>pretype</code> is not valid.</p>
Notes	The preconditioning type is initially set in the call to the linear solver's specification function (see §5.5.3). This function call is needed only if <code>pretype</code> is being changed from its original value.

CVSpilsSetGSType

Call	<code>flag = CVSpilsSetGSType(cvode_mem, gstype);</code>
Description	The function <code>CVSpilsSetGSType</code> specifies the Gram-Schmidt orthogonalization to be used with the CVSPGMR solver (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>cvode_mem</code> (void *) pointer to the CVODES memory block. <code>gstype</code> (int) type of Gram-Schmidt orthogonalization.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CVSPILS_SUCCESS</code> The optional value has been successfully set. <code>CVSPILS_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL. <code>CVSPILS_LMEM_NULL</code> The CVSPILS linear solver has not been initialized. <code>CVSPILS_ILL_INPUT</code> The Gram-Schmidt orthogonalization type <code>gstype</code> is not valid.
Notes	The default value is <code>MODIFIED_GS</code> . This option is available only for the CVSPGMR linear solver.

**CVSpilsSetDelt**

Call	<code>flag = CVSpilsSetDelt(cvode_mem, delt);</code>
Description	The function <code>CVSpilsSetDelt</code> specifies the factor by which the Krylov linear solver's convergence test constant is reduced from the Newton iteration test constant.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODES memory block. <code>delt</code> (realtype)
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CVSPILS_SUCCESS</code> The optional value has been successfully set. <code>CVSPILS_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL. <code>CVSPILS_LMEM_NULL</code> The CVSPILS linear solver has not been initialized. <code>CVSPILS_ILL_INPUT</code> The factor <code>delt</code> is negative.
Notes	The default value is 0.05. Passing a value <code>delt= 0.0</code> also indicates using the default value.

CVSpilsSetMaxl

Call	<code>flag = CVSpilsSetMaxl(cv_mem, maxl);</code>
Description	The function <code>CVSpilsSetMaxl</code> resets maximum Krylov subspace dimension for the Bi-CGStab or TFQMR methods.
Arguments	<code>cv_mem</code> (void *) pointer to the CVODES memory block. <code>maxl</code> (int) maximum dimension of the Krylov subspace.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CVSPILS_SUCCESS</code> The optional value has been successfully set. <code>CVSPILS_MEM_NULL</code> The <code>cv_mem</code> pointer is NULL. <code>CVSPILS_LMEM_NULL</code> The CVSPILS linear solver has not been initialized.
Notes	The maximum subspace dimension is initially specified in the call to the linear solver specification function (see §5.5.3). This function call is needed only if <code>maxl</code> is being changed from its previous value. This option is available only for the CVSPBCG and CVSPTFQMR linear solvers.



5.5.6 Interpolated output function

An optional function `CVodeGetDky` is available to obtain additional output values. This function should only be called after a successful return from `CVode` as it provides interpolated values either of y or of its derivatives (up to the current order of the integration method) interpolated to any value of t in the last internal step taken by CVODES.

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

`CVodeGetDky`

Call	<code>flag = CVodeGetDky(cvode_mem, t, k, dky);</code>
Description	The function <code>CVodeGetDky</code> computes the k -th derivative of the function y at time t , i.e. $d^{(k)}y/dt^{(k)}(t)$, where $t_n - h_u \leq t \leq t_n$, t_n denotes the current internal time reached, and h_u is the last internal step size successfully used by the solver. The user may request $k = 0, 1, \dots, q_u$, where q_u is the current order.
Arguments	<p><code>cvode_mem</code> (void *) pointer to the CVODES memory block.</p> <p><code>t</code> (realtype) the value of the independent variable at which the derivative is to be evaluated.</p> <p><code>k</code> (int) the derivative order requested.</p> <p><code>dky</code> (N_Vector) vector containing the derivative. This vector must be allocated by the user.</p>
Return value	<p>The return value <code>flag</code> (of type <code>int</code>) is one of</p> <p><code>CV_SUCCESS</code> <code>CVodeGetDky</code> succeeded.</p> <p><code>CV_BAD_K</code> k is not in the range $0, 1, \dots, q_u$.</p> <p><code>CV_BAD_T</code> t is not in the interval $[t_n - h_u, t_n]$.</p> <p><code>CV_BAD_DKY</code> The <code>dky</code> argument was NULL.</p> <p><code>CV_MEM_NULL</code> The <code>cvode_mem</code> argument was NULL.</p>
Notes	It is only legal to call the function <code>CVodeGetDky</code> after a successful return from <code>CVode</code> . See <code>CVodeGetCurrentTime</code> , <code>CVodeGetLastOrder</code> , and <code>CVodeGetLastStep</code> in the next section for access to t_n , q_u , and h_u , respectively.

5.5.7 Optional output functions

CVODES provides an extensive set of functions that can be used to obtain solver performance information. Table 5.2 lists all optional output functions in CVODES, which are then described in detail in the remainder of this section, beginning with those for the main CVODES solver and continuing with those for the linear solver modules. Where the name of an output from a linear solver module 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`).

5.5.7.1 Main solver optional output functions

CVODES 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 CVODES memory block (a suggested tolerance scaling factor, the error weight vector, and the vector of estimated local errors). Functions are also provided to extract statistics related to the performance of the CVODES nonlinear solver used. As a convenience, additional information extraction functions provide the optional outputs in groups. These optional output functions are described next.

Table 5.2: Optional outputs from CVODES, CVDENSE, CVBAND, CVDIAG, and CVSPILS

Optional output	Function name
CVODES main solver	
Size of CVODES real and integer workspaces	CVodeGetWorkSpace
Cumulative number of internal steps	CVodeGetNumSteps
No. of calls to r.h.s. function	CVodeGetNumRhsEvals
No. of calls to linear solver setup function	CVodeGetNumLinSolvSetups
No. of local error test failures that have occurred	CVodeGetNumErrTestFails
Order used during the last step	CVodeGetLastOrder
Order to be attempted on the next step	CVodeGetCurrentOrder
No. of order reductions due to stability limit detection	CVodeGetNumStabLimOrderReds
Actual initial step size used	CVodeGetActualInitStep
Step size used for the last step	CVodeGetLastStep
Step size to be attempted on the next step	CVodeGetCurrentStep
Current internal time reached by the solver	CVodeGetCurrentTime
Suggested factor for tolerance scaling	CVodeGetTolScaleFactor
Error weight vector for state variables	CVodeGetErrWeights
Estimated local error vector	CVodeGetEstLocalErrors
No. of nonlinear solver iterations	CVodeGetNumNonlinSolvIters
No. of nonlinear convergence failures	CVodeGetNumNonlinSolvConvFails
All CVODES integrator statistics	CVodeGetIntegratorStats
CVODES nonlinear solver statistics	CVodeGetNonlinSolvStats
Array showing roots found	CVodeGetRootInfo
No. of calls to user root function	CVodeGetNumGEvals
Name of constant associated with a return flag	CVodeGetReturnFlagName
CVDENSE linear solver	
Size of CVDENSE real and integer workspaces	CVDenseGetWorkSpace
No. of Jacobian evaluations	CVDenseGetNumJacEvals
No. of r.h.s. calls for finite diff. Jacobian evals.	CVDenseGetNumRhsEvals
Last return from a CVDENSE function	CVDenseGetLastFlag
Name of constant associated with a return flag	CVDenseGetReturnFlagName
CVBAND linear solver	
Size of CVBAND real and integer workspaces	CVBandGetWorkSpace
No. of Jacobian evaluations	CVBandGetNumJacEvals
No. of r.h.s. calls for finite diff. Jacobian evals.	CVBandGetNumRhsEvals
Last return from a CVBAND function	CVBandGetLastFlag
Name of constant associated with a return flag	CVBandGetReturnFlagName
CVDIAG linear solver	
Size of CVDIAG real and integer workspaces	CVDiagGetWorkSpace
No. of r.h.s. calls for finite diff. Jacobian evals.	CVDiagGetNumRhsEvals
Last return from a CVDIAG function	CVDiagGetLastFlag
Name of constant associated with a return flag	CVDiagGetReturnFlagName
CVSPILS linear solvers	
Size of real and integer workspaces	CVSpilsGetWorkSpace
No. of linear iterations	CVSpilsGetNumLinIters
No. of linear convergence failures	CVSpilsGetNumConvFails
No. of preconditioner evaluations	CVSpilsGetNumPrecEvals
No. of preconditioner solves	CVSpilsGetNumPrecSolves
No. of Jacobian-vector product evaluations	CVSpilsGetNumJtimesEvals
No. of r.h.s. calls for finite diff. Jacobian-vector evals.	CVSpilsGetNumRhsEvals
Last return from a linear solver function	CVSpilsGetLastFlag
Name of constant associated with a return flag	CVSpilsGetReturnFlagName

CVodeGetWorkSpace

Call	<code>flag = CVodeGetWorkSpace(cvode_mem, &lenrw, &leniw);</code>
Description	The function <code>CVodeGetWorkSpace</code> returns the CVODES real and integer workspace sizes.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODES memory block. <code>lenrw</code> (long int) the number of <code>realtype</code> values in the CVODES workspace. <code>leniw</code> (long int) the number of integer values in the CVODES workspace.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CV_SUCCESS</code> The optional output values have been successfully set. <code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL.
Notes	In terms of the problem size N , the maximum method order <code>maxord</code> , and the number <code>nrtfn</code> of root functions (see §5.8), the actual size of the real workspace, in <code>realtype</code> words, is given by the following:

- base value: $\text{lenrw} = 96 + (\text{maxord}+5) * N_r + 3*\text{nrtfn}$;
- if `itol = CV_SV`: $\text{lenrw} = \text{lenrw} + N_r$;

where 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: $\text{leniw} = 40 + (\text{maxord}+5) * N_i + \text{nrtfn}$;
- if `itol = CV_SV`: $\text{leniw} = \text{leniw} + N_i$;

where N_i is the number of integer words in one `N_Vector` (= 1 for `NVECTOR_SERIAL` and $2*\text{npes}$ for `NVECTOR_PARALLEL` and `npes` processors).

For the default value of `maxord`, with no rootfinding, and with `itol` \neq `CV_SV`, these lengths are given roughly by:

- For the Adams method: $\text{lenrw} = 96 + 17N$ and $\text{leniw} = 57$
- For the BDF method: $\text{lenrw} = 96 + 10N$ and $\text{leniw} = 50$

Note that additional memory is allocated if quadratures and/or forward sensitivity integration is enabled. See §5.7.1 and §6.2.1 for more details.

CVodeGetNumSteps

Call	<code>flag = CVodeGetNumSteps(cvode_mem, &nsteps);</code>
Description	The function <code>CVodeGetNumSteps</code> returns the cumulative number of internal steps taken by the solver (total so far).
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODES memory block. <code>nsteps</code> (long int) number of steps taken by CVODES.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CV_SUCCESS</code> The optional output value has been successfully set. <code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL.

CVodeGetNumRhsEvals

Call	<code>flag = CVodeGetNumRhsEvals(cvode_mem, &nfevals);</code>
Description	The function <code>CVodeGetNumRhsEvals</code> returns the number of calls to the user's right-hand side function.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODES memory block. <code>nfevals</code> (long int) number of calls to the user's <code>f</code> function.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CV_SUCCESS</code> The optional output value has been successfully set. <code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is <code>NULL</code> .
Notes	The <code>nfevals</code> value returned by <code>CVodeGetNumRhsEvals</code> does not account for calls made to <code>f</code> by a linear solver or preconditioner module.

CVodeGetNumLinSolvSetups

Call	<code>flag = CVodeGetNumLinSolvSetups(cvode_mem, &nlinsetups);</code>
Description	The function <code>CVodeGetNumLinSolvSetups</code> returns the number of calls made to the linear solver's setup function.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODES memory block. <code>nlinsetups</code> (long int) number of calls made to the linear solver setup function.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CV_SUCCESS</code> The optional output value has been successfully set. <code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is <code>NULL</code> .

CVodeGetNumErrTestFails

Call	<code>flag = CVodeGetNumErrTestFails(cvode_mem, &netfails);</code>
Description	The function <code>CVodeGetNumErrTestFails</code> returns the number of local error test failures that have occurred.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODES memory block. <code>netfails</code> (long int) number of error test failures.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CV_SUCCESS</code> The optional output value has been successfully set. <code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is <code>NULL</code> .

CVodeGetLastOrder

Call	<code>flag = CVodeGetLastOrder(cvode_mem, &qlast);</code>
Description	The function <code>CVodeGetLastOrder</code> returns the integration method order used during the last internal step.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODES memory block. <code>qlast</code> (int) method order used on the last internal step.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CV_SUCCESS</code> The optional output value has been successfully set. <code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is <code>NULL</code> .

CVodeGetCurrentOrder

Call `flag = CVodeGetCurrentOrder(cvode_mem, &qcur);`

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

Arguments `cvode_mem` (`void *`) pointer to the CVODES memory block.
`qcur` (`int`) method order to be used on the next internal step.

Return value The return value `flag` (of type `int`) is one of
`CV_SUCCESS` The optional output value has been successfully set.
`CV_MEM_NULL` The `cvode_mem` pointer is `NULL`.

CVodeGetLastStep

Call `flag = CVodeGetLastStep(cvode_mem, &hlast);`

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

Arguments `cvode_mem` (`void *`) pointer to the CVODES memory block.
`hlast` (`realtype`) step size taken on the last internal step.

Return value The return value `flag` (of type `int`) is one of
`CV_SUCCESS` The optional output value has been successfully set.
`CV_MEM_NULL` The `cvode_mem` pointer is `NULL`.

CVodeGetCurrentStep

Call `flag = CVodeGetCurrentStep(cvode_mem, &hcur);`

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

Arguments `cvode_mem` (`void *`) pointer to the CVODES 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
`CV_SUCCESS` The optional output value has been successfully set.
`CV_MEM_NULL` The `cvode_mem` pointer is `NULL`.

CVodeGetActualInitStep

Call `flag = CVodeGetActualInitStep(cvode_mem, &hinused);`

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

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

Return value The return value `flag` (of type `int`) is one of
`CV_SUCCESS` The optional output value has been successfully set.
`CV_MEM_NULL` The `cvode_mem` pointer is `NULL`.

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

CVodeGetCurrentTime

Call `flag = CVodeGetCurrentTime(cvode_mem, &tcurl);`

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

Arguments `cvode_mem` (void *) pointer to the CVODES memory block.
`tcurl` (realtype) current internal time reached.

Return value The return value `flag` (of type `int`) is one of
`CV_SUCCESS` The optional output value has been successfully set.
`CV_MEM_NULL` The `cvode_mem` pointer is NULL.

CVodeGetNumStabLimOrderReds

Call `flag = CVodeGetNumStabLimOrderReds(cvode_mem, &nsrlred);`

Description The function `CVodeGetNumStabLimOrderReds` returns the number of order reductions dictated by the BDF stability limit detection algorithm (see §3.5).

Arguments `cvode_mem` (void *) pointer to the CVODES memory block.
`nsrlred` (long int) number of order reductions due to stability limit detection.

Return value The return value `flag` (of type `int`) is one of
`CV_SUCCESS` The optional output value has been successfully set.
`CV_MEM_NULL` The `cvode_mem` pointer is NULL.

Notes If the stability limit detection algorithm was not initialized through a call to `CVodeSetStabLimDet`, then `nsrlred=0`.

CVodeGetTolScaleFactor

Call `flag = CVodeGetTolScaleFactor(cvode_mem, &tolsfac);`

Description The function `CVodeGetTolScaleFactor` 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 `cvode_mem` (void *) pointer to the CVODES memory block.
`tolsfac` (realtype) suggested scaling factor for user-supplied tolerances.

Return value The return value `flag` (of type `int`) is one of
`CV_SUCCESS` The optional output value has been successfully set.
`CV_MEM_NULL` The `cvode_mem` pointer is NULL.

CVodeGetErrWeights

Call `flag = CVodeGetErrWeights(cvode_mem, eweight);`

Description The function `CVodeGetErrWeights` returns the solution error weights at the current time. These are the reciprocals of the W_i given by (3.7).

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

Return value The return value `flag` (of type `int`) is one of
`CV_SUCCESS` The optional output value has been successfully set.
`CV_MEM_NULL` The `cvode_mem` pointer is NULL.

Notes The user must allocate memory for `eweight`.



CVodeGetNumNonlinSolvConvFails

Call `flag = CVodeGetNumNonlinSolvConvFails(cvode_mem, &nncfails);`

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

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

Return value The return value `flag` (of type `int`) is one of
 `CV_SUCCESS` The optional output value has been successfully set.
 `CV_MEM_NULL` The `cvode_mem` pointer is NULL.

CVodeGetNonlinSolvStats

Call `flag = CVodeGetNonlinSolvStats(cvode_mem, &nniters, &nncfails);`

Description The function `CVodeGetNonlinSolvStats` returns the CVODES nonlinear solver statistics as a group.

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

Return value The return value `flag` (of type `int`) is one of
 `CV_SUCCESS` The optional output value has been successfully set.
 `CV_MEM_NULL` The `cvode_mem` pointer is NULL.

CVodeGetReturnFlagName

Call `name = CVodeGetReturnFlagName(flag);`

Description The function `CVodeGetReturnFlagName` returns the name of the CVODE constant corresponding to `flag`.

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

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

5.5.7.2 Dense linear solver

The following optional outputs are available from the CVDENSE module: workspace requirements, number of calls to the Jacobian routine, number of calls to the right-hand side routine for finite-difference Jacobian approximation, and last return value from a CVDENSE function.

CVDenseGetWorkSpace

Call `flag = CVDenseGetWorkSpace(cvode_mem, &lenrwLS, &leniwLS);`

Description The function `CVDenseGetWorkSpace` returns the CVDENSE real and integer workspace sizes.

Arguments `cvode_mem` (void *) pointer to the CVODES memory block.
 `lenrwLS` (long int) the number of `realtype` values in the CVDENSE workspace.
 `leniwLS` (long int) the number of integer values in the CVDENSE workspace.

Return value The return value `flag` (of type `int`) is one of
 `CVDENSE_SUCCESS` The optional output values have been successfully set.
 `CVDENSE_MEM_NULL` The `cvode_mem` pointer is NULL.
 `CVDENSE_LMEM_NULL` The CVDENSE linear solver has not been initialized.

Notes In terms of the problem size N , the actual size of the real workspace is $2N^2$ `realtype` words, and the actual size of the integer workspace is N integer words.

CVDenseGetNumJacEvals

Call `flag = CVDenseGetNumJacEvals(cvode_mem, &njevals);`

Description The function `CVDenseGetNumJacEvals` returns the number of calls made to the dense Jacobian approximation function.

Arguments `cvode_mem` (`void *`) pointer to the CVODES memory block.
`njevals` (`long int`) the number of calls to the Jacobian function.

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

`CVDENSE_SUCCESS` The optional output value has been successfully set.
`CVDENSE_MEM_NULL` The `cvode_mem` pointer is NULL.
`CVDENSE_LMEM_NULL` The CVDENSE linear solver has not been initialized.

CVDenseGetNumRhsEvals

Call `flag = CVDenseGetNumRhsEvals(cvode_mem, &nfevalsLS);`

Description The function `CVDenseGetNumRhsEvals` returns the number of calls made to the user-supplied right-hand side function due to the finite difference dense Jacobian approximation.

Arguments `cvode_mem` (`void *`) pointer to the CVODES memory block.
`nfevalsLS` (`long int`) the number of calls made to the user-supplied right-hand side function.

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

`CVDENSE_SUCCESS` The optional output value has been successfully set.
`CVDENSE_MEM_NULL` The `cvode_mem` pointer is NULL.
`CVDENSE_LMEM_NULL` The CVDENSE linear solver has not been initialized.

Notes The value `nfevalsLS` is incremented only if the default `CVDenseDQJac` difference quotient function is used.

CVDenseGetLastFlag

Call `flag = CVDenseGetLastFlag(cvode_mem, &lsflag);`

Description The function `CVDenseGetLastFlag` returns the last return value from a CVDENSE routine.

Arguments `cvode_mem` (`void *`) pointer to the CVODES memory block.
`lsflag` (`int`) the value of the last return flag from a CVDENSE function.

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

`CVDENSE_SUCCESS` The optional output value has been successfully set.
`CVDENSE_MEM_NULL` The `cvode_mem` pointer is NULL.
`CVDENSE_LMEM_NULL` The CVDENSE linear solver has not been initialized.

Notes If the CVDENSE setup function failed (`Cvode` returned `CV_LSETUP_FAIL`), then the value of `lsflag` corresponds to the column index (numbered from one) of a diagonal element with value zero that was encountered during the LU factorization of the dense Jacobian matrix.

CVDenseGetReturnFlagName

Call `name = CVDenseGetReturnFlagName(flag);`

Description The function `CVDenseGetReturnFlagName` returns the name of the CVDENSE constant corresponding to `flag`.

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

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

5.5.7.3 Band linear solver

The following optional outputs are available from the CVBAND module: workspace requirements, number of calls to the Jacobian routine, number of calls to the right-hand side routine for finite-difference Jacobian approximation, and last return value from a CVBAND function.

CVBandGetWorkSpace

Call `flag = CVBandGetWorkSpace(cvode_mem, &lenrwLS, &leniwLS);`

Description The function `CVBandGetWorkSpace` returns the CVBAND real and integer workspace sizes.

Arguments `cvode_mem` (`void *`) pointer to the CVODES memory block.

`lenrwLS` (`long int`) the number of `realtype` values in the CVBAND workspace.

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

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

`CVBAND_SUCCESS` The optional output values have been successfully set.

`CVBAND_MEM_NULL` The `cvode_mem` pointer is `NULL`.

`CVBAND_LMEM_NULL` The CVBAND linear solver has not been initialized.

Notes In terms of the problem size N and Jacobian half-bandwidths, the actual size of the real workspace is $(2 \text{ mupper} + 3 \text{ mlower} + 2) N \text{ realtype}$ words, and the actual size of the integer workspace is N integer words.

CVBandGetNumJacEvals

Call `flag = CVBandGetNumJacEvals(cvode_mem, &njevals);`

Description The function `CVBandGetNumJacEvals` returns the number of calls made to the banded Jacobian approximation function.

Arguments `cvode_mem` (`void *`) pointer to the CVODES memory block.

`njevals` (`long int`) the number of calls to the Jacobian function.

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

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

`CVBAND_MEM_NULL` The `cvode_mem` pointer is `NULL`.

`CVBAND_LMEM_NULL` The CVBAND linear solver has not been initialized.

CVBandGetNumRhsEvals

Call `flag = CVBandGetNumRhsEvals(cvode_mem, &nfevalsLS);`

Description The function `CVBandGetNumRhsEvals` returns the number of calls made to the user-supplied right-hand side function due to the finite difference banded Jacobian approximation.

Arguments `cvode_mem` (`void *`) pointer to the CVODES memory block.

`nfevalsLS` (long int) the number of calls made to the user-supplied right-hand side function.

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

- `CVBAND_SUCCESS` The optional output value has been successfully set.
- `CVBAND_MEM_NULL` The `cvode_mem` pointer is `NULL`.
- `CVBAND_LMEM_NULL` The `CVBAND` linear solver has not been initialized.

Notes The value `nfevalsLS` is incremented only if the default `CVBandDQJac` difference quotient function is used.

CVBandGetLastFlag

Call `flag = CVBandGetLastFlag(cvode_mem, &lsflag);`

Description The function `CVBandGetLastFlag` returns the value of the last return flag from a `CVBAND` routine.

Arguments `cvode_mem` (void *) pointer to the `CVODES` memory block.
`lsflag` (int) the value of the last return flag from a `CVBAND` function.

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

- `CVBAND_SUCCESS` The optional output value has been successfully set.
- `CVBAND_MEM_NULL` The `cvode_mem` pointer is `NULL`.
- `CVBAND_LMEM_NULL` The `CVBAND` linear solver has not been initialized.

Notes If the `CVBAND` setup function failed (`CVode` returned `CV_LSETUP_FAIL`), the value of `lsflag` corresponds to the column index (numbered from one) of a diagonal element with value zero that was encountered during the LU factorization of the banded Jacobian matrix.

CVBandGetReturnFlagName

Call `name = CVBandGetReturnFlagName(flag);`

Description The function `CVBandGetReturnFlagName` returns the name of the `CVBAND` constant corresponding to `flag`.

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

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

5.5.7.4 Diagonal linear solver

The following optional outputs are available from the `CVDIAG` module: workspace requirements, number of calls to the right-hand side routine for finite-difference Jacobian approximation, and last return value from a `CVDIAG` function.

CVDiagGetWorkSpace

Call `flag = CVDiagGetWorkSpace(cvode_mem, &lenrwLS, &leniwLS);`

Description The function `CVDiagGetWorkSpace` returns the `CVDIAG` real and integer workspace sizes.

Arguments `cvode_mem` (void *) pointer to the `CVODES` memory block.
`lenrwLS` (long int) the number of `realtype` values in the `CVDIAG` workspace.
`leniwLS` (long int) the number of integer values in the `CVDIAG` workspace.

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

- `CVDIAG_SUCCESS` The optional output values have been successfully set.
- `CVDIAG_MEM_NULL` The `cvode_mem` pointer is `NULL`.

Notes CVDIAG_LMEM_NULL The CVDIAG linear solver has not been initialized.
 In terms of the problem size N , the actual size of the real workspace is roughly $3N$ `realtype` words.

CVDiagGetNumRhsEvals

Call `flag = CVDiagGetNumRhsEvals(cvode_mem, &nfevalsLS);`

Description The function `CVDiagGetNumRhsEvals` returns the number of calls made to the user-supplied right-hand side function due to the finite difference Jacobian approximation.

Arguments `cvode_mem` (`void *`) pointer to the CVODES memory block.
 `nfevalsLS` (`long int`) the number of calls made to the user-supplied right-hand side function.

Return value The return value `flag` (of type `int`) is one of
 CVDIAG_SUCCESS The optional output value has been successfully set.
 CVDIAG_MEM_NULL The `cvode_mem` pointer is NULL.
 CVDIAG_LMEM_NULL The CVDIAG linear solver has not been initialized.

Notes The number of diagonal approximate Jacobians formed is equal to the number of calls made to the linear solver setup function (see `CVodeGetNumLinSolvSetups`).

CVDiagGetLastFlag

Call `flag = CVDiagGetLastFlag(cvode_mem, &lsflag);`

Description The function `CVDiagGetLastFlag` returns the last return value from a CVDIAG routine.

Arguments `cvode_mem` (`void *`) pointer to the CVODES memory block.
 `lsflag` (`int`) the value of the last return flag from a CVDIAG function.

Return value The return value `flag` (of type `int`) is one of
 CVDIAG_SUCCESS The optional output value has been successfully set.
 CVDIAG_MEM_NULL The `cvode_mem` pointer is NULL.
 CVDIAG_LMEM_NULL The CVDIAG linear solver has not been initialized.

Notes If the CVDIAG setup function failed (`CVode` returned `CV_LSETUP_FAIL`), the value of `lsflag` is equal to `CVDIAG_INV_FAIL`, indicating that a diagonal element with value zero was encountered. The same value is also returned if the CVDIAG solve function failed (`CVode` returned `CV_LSOLVE_FAIL`).

CVDiagGetReturnFlagName

Call `name = CVDiagGetReturnFlagName(flag);`

Description The function `CVDiagGetReturnFlagName` returns the name of the CVDIAG constant corresponding to `flag`.

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

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

5.5.7.5 SPILS linear solvers

The following optional outputs are available from the CVSPILS 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 right-hand side routine for finite-difference Jacobian-vector product approximation, and last return value from a linear solver function.

CVSpilsGetWorkSpace

- Call** `flag = CVSpilsGetWorkSpace(cvode_mem, &lenrwLS, &leniwLS);`
- Description** The function `CVSpilsGetWorkSpace` returns the global sizes of the CVSPGMR real and integer workspaces.
- Arguments** `cvode_mem` (void *) pointer to the CVODES memory block.
`lenrwLS` (long int) the number of **realtype** values in the CVSPILS workspace.
`leniwLS` (long int) the number of integer values in the CVSPILS workspace.
- Return value** The return value `flag` (of type `int`) is one of
`CVSPILS_SUCCESS` The optional output value has been successfully set.
`CVSPILS_MEM_NULL` The `cvode_mem` pointer is NULL.
`CVSPILS_LMEM_NULL` The CVSPILS 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:
 $(\text{maxl}+5) * N + \text{maxl} * (\text{maxl}+4) + 1$ **realtype** words for CVSPGMR,
 $9 * N$ **realtype** words for CVSPBCG,
and $11 * N$ **realtype** words for IDASPTFQMR.
In a parallel setting, the above values are global — summed over all processors.

CVSpilsGetNumLinIters

- Call** `flag = CVSpilsGetNumLinIters(cvode_mem, &nliters);`
- Description** The function `CVSpilsGetNumLinIters` returns the cumulative number of linear iterations.
- Arguments** `cvode_mem` (void *) pointer to the CVODES memory block.
`nliters` (long int) the current number of linear iterations.
- Return value** The return value `flag` (of type `int`) is one of
`CVSPILS_SUCCESS` The optional output value has been successfully set.
`CVSPILS_MEM_NULL` The `cvode_mem` pointer is NULL.
`CVSPILS_LMEM_NULL` The CVSPILS linear solver has not been initialized.

CVSpilsGetNumConvFails

- Call** `flag = CVSpilsGetNumConvFails(cvode_mem, &nlcfails);`
- Description** The function `CVSpilsGetNumConvFails` returns the cumulative number of linear convergence failures.
- Arguments** `cvode_mem` (void *) pointer to the CVODES memory block.
`nlcfails` (long int) the current number of linear convergence failures.
- Return value** The return value `flag` (of type `int`) is one of
`CVSPILS_SUCCESS` The optional output value has been successfully set.
`CVSPILS_MEM_NULL` The `cvode_mem` pointer is NULL.
`CVSPILS_LMEM_NULL` The CVSPILS linear solver has not been initialized.

CVSpilsGetNumPrecEvals

- Call** `flag = CVSpilsGetNumPrecEvals(cvode_mem, &npevals);`
- Description** The function `CVSpilsGetNumPrecEvals` returns the number of preconditioner evaluations, i.e., the number of calls made to `psetup` with `jok = FALSE`.

Arguments `cvode_mem` (void *) pointer to the CVODES memory block.
 `npevals` (long int) the current number of calls to `psetup`.

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

CVSPILS_SUCCESS The optional output value has been successfully set.
 CVSPILS_MEM_NULL The `cvode_mem` pointer is NULL.
 CVSPILS_LMEM_NULL The CVSPILS linear solver has not been initialized.

CVSpilsGetNumPrecSolves

Call `flag = CVSpilsGetNumPrecSolves(cvode_mem, &npsolves);`

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

Arguments `cvode_mem` (void *) pointer to the CVODES memory block.
 `npsolves` (long int) the current number of calls to `psolve`.

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

CVSPILS_SUCCESS The optional output value has been successfully set.
 CVSPILS_MEM_NULL The `cvode_mem` pointer is NULL.
 CVSPILS_LMEM_NULL The CVSPILS linear solver has not been initialized.

CVSpilsGetNumJtimesEvals

Call `flag = CVSpilsGetNumJtimesEvals(cvode_mem, &njvevals);`

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

Arguments `cvode_mem` (void *) pointer to the CVODES memory block.
 `njvevals` (long int) the current number of calls to `jtimes`.

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

CVSPILS_SUCCESS The optional output value has been successfully set.
 CVSPILS_MEM_NULL The `cvode_mem` pointer is NULL.
 CVSPILS_LMEM_NULL The CVSPILS linear solver has not been initialized.

CVSpilsGetNumRhsEvals

Call `flag = CVSpilsGetNumRhsEvals(cvode_mem, &nfevalsLS);`

Description The function `CVSpilsGetNumRhsEvals` returns the number of calls to the user right-hand side function for finite difference Jacobian-vector product approximation.

Arguments `cvode_mem` (void *) pointer to the CVODES memory block.
 `nfevalsLS` (long int) the number of calls to the user right-hand side function.

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

CVSPILS_SUCCESS The optional output value has been successfully set.
 CVSPILS_MEM_NULL The `cvode_mem` pointer is NULL.
 CVSPILS_LMEM_NULL The CVSPILS linear solver has not been initialized.

Notes The value `nfevalsLS` is incremented only if the default `CVSpilsDQJtimes` difference quotient function is used.

CVSpilsGetLastFlag

Call	<code>flag = CVSpilsGetLastFlag(cvode_mem, &lsflag);</code>
Description	The function <code>CVSpilsGetLastFlag</code> returns the last return value from a CVSPILS routine.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVOIDES memory block. <code>flag</code> (int) the value of the last return flag from a CVSPILS function.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CVSPILS_SUCCESS</code> The optional output value has been successfully set. <code>CVSPILS_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL. <code>CVSPILS_LMEM_NULL</code> The CVSPILS linear solver has not been initialized.
Notes	If the CVSPILS setup function failed (<code>CVode</code> returned <code>CV_LSETUP_FAIL</code>), <code>lsflag</code> will be <code>SPGMR_PSET_FAIL_UNREC</code> , <code>SPBCG_PSET_FAIL_UNREC</code> , or <code>SPTFQMR_PSET_FAIL_UNREC</code> . If the CVSPGMR solve function failed (<code>CVode</code> returned <code>CV_LSOLVE_FAIL</code>), <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 Jacobian-times-vector 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. If the CVSPBCG solve function failed (<code>CVode</code> returned <code>CV_LSOLVE_FAIL</code>), <code>lsflag</code> contains the error return flag from <code>SpbcgSolve</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 Jacobian-times-vector function; or <code>SPBCG_PSOLVE_FAIL_UNREC</code> , indicating that the preconditioner solve function <code>psolve</code> failed unrecoverably. If the CVSPTFQMR solve function failed (<code>CVode</code> returned <code>CV_LSOLVE_FAIL</code>), <code>lsflag</code> contains the error return 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 Jacobian-times-vector function; or <code>SPTFQMR_PSOLVE_FAIL_UNREC</code> , indicating that the preconditioner solve function <code>psolve</code> failed unrecoverably.

CVSpilsGetReturnFlagName

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

5.5.8 CVOIDES reinitialization function

The function `CVodeReInit` reinitializes the main CVOIDES solver for the solution of a problem, where a prior call to `CVodeMalloc` has been made. The new problem must have the same size as the previous one. `CVodeReInit` performs the same input checking and initializations that `CVodeMalloc` does, but does no memory allocation as it assumes that the existing internal memory is sufficient for the new problem.

The use of `CVodeReInit` requires that the maximum method order, denoted by `maxord`, be no larger for the new problem than for the previous problem. This condition is automatically fulfilled if the multistep method parameter `lmm` is unchanged (or changed from `CV_ADAMS` to `CV_BDF`) and the default value for `maxord` is specified.

If there are changes to the linear solver specifications, make the appropriate `CV*Set*` calls, as described in §5.5.3

CVodeReInit

Call	<code>flag = CVodeReInit(cvode_mem, f, t0, y0, itol, reltol, abstol);</code>
Description	The function <code>CVodeReInit</code> provides required problem specifications and reinitializes CVODES.
Arguments	<p><code>cvode_mem</code> (<code>void *</code>) pointer to the CVODES memory block.</p> <p><code>f</code> (<code>CVRhsFn</code>) is the C function which computes f in the ODE. This function has the form <code>f(N, t, y, ydot, f_data)</code> (for full details see §5.6).</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>itol</code> (<code>int</code>) is one of <code>CV_SS</code>, <code>CV_SV</code>, or <code>CV_WF</code>, where <code>itol = CV_SS</code> indicates scalar relative error tolerance and scalar absolute error tolerance, while <code>itol = CV_SV</code> indicates scalar relative error tolerance and vector absolute error tolerance. The latter choice is important when the absolute error tolerance needs to be different for each component of the ODE. If <code>itol=CV_WF</code>, the arguments <code>reltol</code> and <code>abstol</code> are ignored and the user is expected to provide a function to evaluate the error weight vector W from (3.7). See <code>CVodeSetEwtFn</code> in §5.5.5.1.</p> <p><code>reltol</code> (<code>realtype</code>) is the relative error tolerance.</p> <p><code>abstol</code> (<code>void *</code>) is a pointer to the absolute error tolerance. If <code>itol=CV_SS</code>, <code>abstol</code> must be a pointer to a <code>realtype</code> variable. If <code>itol=CV_SV</code>, <code>abstol</code> must be an <code>N_Vector</code> variable.</p>
Return value	<p>The return flag <code>flag</code> (of type <code>int</code>) will be one of the following:</p> <p><code>CV_SUCCESS</code> The call to <code>CVodeReInit</code> was successful.</p> <p><code>CV_MEM_NULL</code> The CVODES memory block was not initialized through a previous call to <code>CVodeCreate</code>.</p> <p><code>CV_NO_MALLOC</code> Memory space for the CVODES memory block was not allocated through a previous call to <code>CVodeMalloc</code>.</p> <p><code>CV_ILL_INPUT</code> An input argument to <code>CVodeReInit</code> has an illegal value.</p>
Notes	<p>If an error occurred, <code>CVodeReInit</code> also sends an error message to the error handler function.</p> <p>It is the user's responsibility to provide compatible <code>itol</code> and <code>abstol</code> arguments.</p>



5.6 User-supplied functions

The user-supplied functions consist of one function defining the ODE, (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 iterative algorithms.

5.6.1 ODE right-hand side

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

CVRhsFn

Definition	<code>typedef int (*CVRhsFn)(realtype t, N_Vector y, N_Vector ydot, void *f_data);</code>
Purpose	This function computes the ODE right-hand side for a given value of the independent variable t and state vector y .

Arguments	<p>t is the current value of the independent variable.</p> <p>y is the current value of the dependent variable vector, $y(t)$.</p> <p>ydot is the output vector $f(t, y)$.</p> <p>f_data is the f_data pointer passed to CVodeSetFdata.</p>
Return value	A CVRhsFn should return 0 if successful, a positive value if a recoverable error occurred (in which case CVODES will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and CV_RHSFUNC_FAIL is returned).
Notes	Allocation of memory for ydot is handled within CVODES.



For efficiency considerations, the right-hand side function is not evaluated at the converged solution of the nonlinear solver. Therefore, a recoverable error in **CVRhsFn** at that point cannot be corrected (as it will occur when the right-hand side function is called the first time during the following integration step and a successful step cannot be undone). However, if the use program also includes quadrature integration, the state variables can be checked for legality in the call to **CVQuadRhsFn** which is called at the converged solution of the nonlinear system and therefore CVODES can be flagged to attempt to recover from such a situation. Also, if sensitivity analysis is performed with one of the staggered methods, the ODE right-hand side function is called at the converged solution of the nonlinear system and a recoverable error at that point will be captured and CVODES will try to correct it.

There are two other situations in which recovery is not possible even if the right-hand side function returns a recoverable error flag. This include the situation when this occurs at the very first call to the **CVRhsFn** (in which case CVODES returns **CV_FIRST_RHSFUNC_ERR**) or if a recoverable error is reported when **CVRhsFn** is called after an error test failure, while the linear multistep method order is equal to 1 (in which case CVODES returns **CV_UNREC_RHSFUNC_ERR**).

5.6.2 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 **CVSetErrFile**), the user may provide a function of type **CVErrorHandlerFn** to process any such messages. The function type **CVErrorHandlerFn** is defined as follows:

CVErrorHandlerFn

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

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

Arguments **error_code** is the error code.

module is the name of the CVODES 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 **CVodeSetErrorHandlerFn**.

Return value A **CVErrorHandlerFn** function has no return value.

Notes **error_code** is negative for errors and positive (**CV_WARNING**) for warnings. If a function returning a pointer to memory (e.g. **CVBBDPAlloc**) encounters an error, it sets **error_code** to 0 before returning **NULL**.

5.6.3 Error weight function

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

`CWEwtFn`

Definition	<code>typedef int (*CWEwtFn)(N_Vector y, N_Vector ewt, void *e_data);</code>
Purpose	This function computes the WRMS error weights for the vector y .
Arguments	<p>y is the value of the vector for which the WRMS norm must be computed.</p> <p>ewt is the output vector containing the error weights.</p> <p>e_data is the <code>e_data</code> pointer passed to <code>CVodeSetEwtFn</code>.</p>
Return value	A <code>CWEwtFn</code> function type must return 0 if it successfully set the error weights and -1 otherwise. In case of failure, a message is printed and the integration stops.
Notes	<p>Allocation of memory for <code>ewt</code> is handled within <code>CVODES</code>.</p> <p>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.</p>



5.6.4 Jacobian information (direct method with dense Jacobian)

If the direct linear solver with dense treatment of the Jacobian is used (i.e., `CVDense` is called in Step 7 of §5.4), the user may provide a function of type `CVDenseJacFn` defined by:

`CVDenseJacFn`

Definition	<code>typedef (*CVDenseJacFn)(long int N, DenseMat J, realtype t, N_Vector y, N_Vector fy, void *jac_data, N_Vector tmp1, N_Vector tmp2, N_Vector tmp3);</code>
Purpose	This function computes the dense Jacobian $J = \partial f / \partial y$ (or an approximation to it).
Arguments	<p>N is the problem size.</p> <p>J is the output Jacobian matrix.</p> <p>t is the current value of the independent variable.</p> <p>y is the current value of the dependent variable vector, namely the predicted value of $y(t)$.</p> <p>fy is the current value of the vector $f(t, y)$.</p> <p>jac_data is the <code>jac_data</code> pointer passed to <code>CVDenseSetJacFn</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>CVDenseJacFn</code> as temporary storage or work space.</p>
Return value	A <code>CVDenseJacFn</code> should return 0 if successful, a positive value if a recoverable error occurred (in which case <code>CVODES</code> will attempt to correct, while <code>CVDENSE</code> sets <code>last_flag</code> on <code>CVDENSE_JACFUNC_RECVR</code>), or a negative value if it failed unrecoverably (in which case the integration is halted, <code>CVode</code> returns <code>CV_LSETUP_FAIL</code> and <code>CVDENSE</code> sets <code>last_flag</code> on <code>CVDENSE_JACFUNC_UNRECVR</code>).
Notes	A user-supplied dense Jacobian function must load the N by N dense matrix J with an approximation to the Jacobian matrix J at the point (t, y) . Only nonzero elements need to be loaded into J because J is set to the zero matrix before the call to the Jacobian function. The type of J is <code>DenseMat</code> .

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 `DenseMat` type. `DENSE_ELEM(J, i, j)` references the (i, j) -th element of the dense matrix J ($i, j = 0 \dots N-1$). This macro is meant for small problems for which efficiency of access is not a major concern. Thus, in terms of the indices m and n ranging from 1 to N , the Jacobian element $J_{m,n}$ can be set using the statement `DENSE_ELEM(J, m-1, n-1) = Jm,n`. Alternatively, `DENSE_COL(J, j)` returns a pointer to the first element of the j -th column of J ($j = 0 \dots N-1$), and the elements of the j -th column can then be accessed using ordinary array indexing. Consequently, $J_{m,n}$ can be loaded using the statements `col_n = DENSE_COL(J, n-1); col_n[m-1] = Jm,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.

The `DenseMat` type and accessor macros `DENSE_ELEM` and `DENSE_COL` are documented in §10.1.

If the user's `CVDenseJacFn` function uses difference quotient approximations, then it may need to access quantities not in the argument list. These include the current step size, the error weights, etc. To obtain these, use the `CVodeGet*` functions described in §5.5.7.1. The unit roundoff can be accessed as `UNIT_ROUNDOFF` defined in `sundials_types.h`.

5.6.5 Jacobian information (direct method with banded Jacobian)

If the direct linear solver with banded treatment of the Jacobian is used (i.e. `CVBand` is called in Step 7 of §5.4), the user may provide a function of type `CVBandJacFn` defined as follows:

<div>CVBandJacFn</div>	
Definition	<pre>typedef int (*CVBandJacFn)(long int N, long int mupper, long int mlower, BandMat J, realtype t, N_Vector y, N_Vector fy, void *jac_data, N_Vector tmp1, N_Vector tmp2, N_Vector tmp3);</pre>
Purpose	This function computes the banded Jacobian $J = \partial f / \partial y$ (or a banded approximation to it).
Arguments	<p><code>N</code> is the problem size.</p> <p><code>mlower</code></p> <p><code>mupper</code> are the lower and upper half-bandwidths of the Jacobian.</p> <p><code>J</code> is the output Jacobian matrix.</p> <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, namely the predicted value of $y(t)$.</p> <p><code>fy</code> is the current value of the vector $f(t, y)$.</p> <p><code>jac_data</code> is the <code>jac_data</code> pointer passed to <code>CVBandSetJacFn</code>.</p> <p><code>tmp1</code></p> <p><code>tmp2</code></p> <p><code>tmp3</code> are pointers to memory allocated for variables of type <code>N_Vector</code> which can be used by <code>CVBandJacFn</code> as temporary storage or work space.</p>
Return value	A <code>CVBandJacFn</code> should return 0 if successful, a positive value if a recoverable error occurred (in which case CVODES will attempt to correct, while CVBAND sets <code>last_flag</code> on <code>CVBAND_JACFUNC_RECVR</code>), or a negative value if it failed unrecoverably (in which case the integration is halted, CVode returns <code>CV_LSETUP_FAIL</code> and CVBAND sets <code>last_flag</code> on <code>CVBAND_JACFUNC_UNRECVR</code>).

Notes A user-supplied band Jacobian function must load the band matrix **J** of type **BandMat** with the elements of the Jacobian $J(t, y)$ at the point (t, y) . Only nonzero elements need to be loaded into **J** because **J** is initialized to the zero matrix before the call to the Jacobian function.

The accessor macros **BAND_ELEM**, **BAND_COL**, and **BAND_COL_ELEM** allow the user to read and write band matrix elements without making specific references to the underlying representation of the **BandMat** type. **BAND_ELEM**(**J**, **i**, **j**) references the (**i**, **j**)-th element of the band matrix **J**, counting from 0. This macro is meant for use in small problems for which efficiency of access is not a major concern. Thus, in terms of the indices m and n ranging from 1 to N with (m, n) within the band defined by **mupper** and **mlower**, the Jacobian element $J_{m,n}$ can be loaded using the statement **BAND_ELEM**(**J**, **m-1**, **n-1**) = $J_{m,n}$. The elements within the band are those with $-\text{mupper} \leq m-n \leq \text{mlower}$. Alternatively, **BAND_COL**(**J**, **j**) returns a pointer to the diagonal element of the **j**-th column of **J**, and if we assign this address to **realtype** ***col_j**, then the **i**-th element of the **j**-th column is given by **BAND_COL_ELEM**(**col_j**, **i**, **j**), counting from 0. Thus, for (m, n) within the band, $J_{m,n}$ can be loaded by setting **col_n** = **BAND_COL**(**J**, **n-1**); **BAND_COL_ELEM**(**col_n**, **m-1**, **n-1**) = $J_{m,n}$. The elements of the **j**-th column can also be accessed via ordinary array indexing, but this approach requires knowledge of the underlying storage for a band matrix of type **BandMat**. The array **col_n** can be indexed from $-\text{mupper}$ to **mlower**. For large problems, it is more efficient to use **BAND_COL** and **BAND_COL_ELEM** than to use the **BAND_ELEM** macro. As in the dense case, these macros all number rows and columns starting from 0.

The **BandMat** type and the accessor macros **BAND_ELEM**, **BAND_COL** and **BAND_COL_ELEM** are documented in §10.2.

If the user's **CVBandJacFn** function uses difference quotient approximations, then it may need to access quantities not in the argument list. These include the current step size, the error weights, etc. To obtain these, use the **CVodeGet*** functions described in §5.5.7.1. The unit roundoff can be accessed as **UNIT_ROUNDOFF** defined in **sundials_types.h**.

5.6.6 Jacobian information (matrix-vector product)

If one of the Krylov iterative linear solvers SPGMR, SPBCG, or SPTFQMR is selected (**CVSp*** is called in step 7 of §5.4), the user may provide a function of type **CVSpilsJacTimesVecFn** in the following form:

CVSpilsJacTimesVecFn

Definition	<pre>typedef int (*CVSpilsJacTimesVecFn)(N_Vector v, N_Vector Jv, realtype t, N_Vector y, N_Vector fy, void *jac_data, N_Vector tmp);</pre>														
Purpose	This function computes the product $Jv = (\partial f / \partial y)v$ (or an approximation to it).														
Arguments	<table border="0"> <tr> <td style="padding-right: 10px;">v</td> <td>is the vector by which the Jacobian must be multiplied.</td> </tr> <tr> <td>Jv</td> <td>is the output vector computed.</td> </tr> <tr> <td>t</td> <td>is the current value of the independent variable.</td> </tr> <tr> <td>y</td> <td>is the current value of the dependent variable vector.</td> </tr> <tr> <td>fy</td> <td>is the current value of the vector $f(t, y)$.</td> </tr> <tr> <td>jac_data</td> <td>is the jac_data pointer passed to CVSp*SetJacTimesVecFn.</td> </tr> <tr> <td>tmp</td> <td>is a pointer to memory allocated for a variable of type N_Vector which can be used for work space.</td> </tr> </table>	v	is the vector by which the Jacobian must be multiplied.	Jv	is the output vector computed.	t	is the current value of the independent variable.	y	is the current value of the dependent variable vector.	fy	is the current value of the vector $f(t, y)$.	jac_data	is the jac_data pointer passed to CVSp*SetJacTimesVecFn .	tmp	is a pointer to memory allocated for a variable of type N_Vector which can be used for work space.
v	is the vector by which the Jacobian must be multiplied.														
Jv	is the output vector computed.														
t	is the current value of the independent variable.														
y	is the current value of the dependent variable vector.														
fy	is the current value of the vector $f(t, y)$.														
jac_data	is the jac_data pointer passed to CVSp*SetJacTimesVecFn .														
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 Jacobian-vector product function should be 0 if successful. Any other return value will result in an unrecoverable error of the SPGMR generic solver, in which case the integration is halted.
- Notes** If the user's `CVSpilsJacTimesVecFn` function uses difference quotient approximations, it may need to access quantities not in the argument list. These include the current step size, the error weights, etc. To obtain these, use the `CVodeGet*` functions described in §5.5.7.1. The unit roundoff can be accessed as `UNIT_ROUNDOFF` defined in `sundials_types.h`.

5.6.7 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 may be either a left or right preconditioner matrix. This function must be of type `CVSpilsPrecSolveFn`, defined as follows:

`CVSpilsPrecSolveFn`

Definition

```
typedef int (*CVSpilsPrecSolveFn)(realtype t, N_Vector y, N_Vector fy,
                                N_Vector r, N_Vector z,
                                realtype gamma, realtype delta,
                                int lr, void *p_data, N_Vector tmp);
```

- Purpose** This function solves the preconditioned system $Pz = r$.
- Arguments**
- t** is the current value of the independent variable.
 - y** is the current value of the dependent variable vector.
 - fy** is the current value of the vector $f(t, y)$.
 - r** is the right-hand side vector of the linear system.
 - z** is the computed output vector.
 - gamma** is the scalar γ appearing in the Newton matrix given by $M = I - \gamma J$.
 - 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 the weighted l_2 norm, i.e., $\sqrt{\sum_i (Res_i \cdot ewt_i)^2} < \text{delta}$. To obtain the `N_Vector ewt` call `CVodeGetErrWeights` (see §5.5.7.1).
 - lr** is an input flag indicating whether the preconditioner solve function is to use the left preconditioner (**lr** = 1) or the right preconditioner (**lr** = 2);
 - p_data** is the **p_data** pointer passed to `CVSpilsSetPreconditioner`.
 - 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), or negative for an unrecoverable error (in which case the integration is halted).

5.6.8 Preconditioning (Jacobian data)

If the user's preconditioner requires that any Jacobian-related data be preprocessed or evaluated, then this needs to be done in a user-supplied C function of type `CVSpilsPrecSetupFn`, defined as follows:

`CVSpilsPrecSetupFn`

Definition	<pre>typedef int (*CVSpilsPrecSetupFn)(realtype t, N_Vector y, N_Vector fy, booleantype jok, booleantype *jcurPtr, realtype gamma, void *p_data, N_Vector tmp1, N_Vector tmp2, N_Vector tmp3);</pre>
Purpose	This function preprocesses and/or evaluates Jacobian-related data needed by the preconditioner.
Arguments	<p>The arguments of a <code>CVSpilsPrecSetupFn</code> are as follows:</p> <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, namely the predicted value of $y(t)$.</p> <p><code>fy</code> is the current value of the vector $f(t, y)$.</p> <p><code>jok</code> is an input flag indicating whether the Jacobian-related data needs to be updated. The <code>jok</code> argument provides for the reuse of Jacobian data in the preconditioner solve function. <code>jok = FALSE</code> means that the Jacobian-related data must be recomputed from scratch. <code>jok = TRUE</code> means that the Jacobian data, if saved from the previous call to this function, can be reused (with the current value of <code>gamma</code>). A call with <code>jok = TRUE</code> can only occur after a call with <code>jok = FALSE</code>.</p> <p><code>jcurPtr</code> is a pointer to a flag which should be set to <code>TRUE</code> if Jacobian data was recomputed, or set to <code>FALSE</code> if Jacobian data was not recomputed, but saved data was still reused.</p> <p><code>gamma</code> is the scalar γ appearing in the Newton matrix $M = I - \gamma P$.</p> <p><code>p_data</code> is the <code>p_data</code> pointer passed to <code>CVSp*SetPreconditioner</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>CVSpilsPrecSetupFn</code> as temporary storage or work space.</p>
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), or negative for an unrecoverable error (in which case the integration is halted).
Notes	<p>The operations performed by this function might include forming a crude approximate Jacobian, and performing an LU factorization of the resulting approximation to $M = I - \gamma J$.</p> <p>Each call to the preconditioner setup function is preceded by a call to the <code>CVRhsFn</code> user function with the same <code>(t,y)</code> arguments. Thus, the preconditioner setup function can use any auxiliary data that is computed and saved during the evaluation of the ODE right-hand side.</p> <p>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.</p> <p>If the user's <code>CVSpilsPrecSetupFn</code> function uses difference quotient approximations, it may need to access quantities not in the call list. These include the current step size, the error weights, etc. To obtain these, use the <code>CVodeGet*</code> functions described in §5.5.7.1. The unit roundoff can be accessed as <code>UNIT_ROUNDOFF</code> defined in <code>sundials_types.h</code>.</p>

5.7 Integration of pure quadrature equations

If the system of ODEs contains *pure quadratures*, it is more efficient to treat them separately by excluding them from the nonlinear solution stage. To do this, begin by excluding the quadrature

variables from the vector `y` and the quadrature equations from within `f`. 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 §5.4 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 vector of initial values**
4. **Create CVODES 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 `CVodeQuadMalloc` to specify the quadrature equation right-hand side function and to allocate internal memory related to quadrature integration. See §5.7.1 for details.
11. **Set optional inputs for quadrature integration**
 - Call `CVodeSetQuadFdata` to specify user data required for the evaluation of the quadrature equation right-hand side. Call `CVodeSetQuadErrCon` to indicate whether or not quadrature variables should be used in the step size control mechanism, and to specify the integration tolerances for quadrature variables. See §5.7.4 for details.
12. **Advance solution in time**
13. **Extract quadrature variables**
 - Call `CVodeGetQuad` to obtain the values of the quadrature variables at the current time. See §5.7.3 for details.
14. **Get optional outputs**
15. **Get quadrature optional outputs**
 - Call `CVodeGetQuad*` functions to obtain optional output related to the integration of quadratures. See §5.7.5 for details.
16. **Deallocate memory for solution vector and for the vector of quadrature variables**
17. **Free solver memory**
18. [P] Finalize MPI

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

5.7.1 Quadrature initialization functions

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

<div style="border: 1px solid black; padding: 2px; display: inline-block;">CVodeQuadMalloc</div>	
Call	<code>flag = CVodeQuadMalloc(cvode_mem, fQ, yQ0);</code>
Description	The function <code>CVodeQuadMalloc</code> provides required problem specifications, allocates internal memory, and initializes quadrature integration.
Arguments	<p><code>cvode_mem</code> (void *) pointer to the CVODES memory block returned by <code>CVodeCreate</code>.</p> <p><code>fQ</code> (<code>CVQuadRhsFn</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, y, yQdot, fQ_data)</code> (for full details see §5.7.6).</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>CV_SUCCESS</code> The call to <code>CVodeQuadMalloc</code> was successful.</p> <p><code>CV_MEM_NULL</code> The CVODES memory was not initialized by a prior call to <code>CVodeCreate</code>.</p> <p><code>CV_MEM_FAIL</code> A memory allocation request failed.</p>
Notes	If an error occurred, <code>CVodeQuadMalloc</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 by:

- Base value: $\text{lenrw} = \text{lenrw} + (\text{maxord}+5)N_q$
- With `itolQ = CV_SV` (see `CVodeSetQuadErrCon`): $\text{lenrw} = \text{lenrw} + N_q$

the size of the integer workspace is increased by:

- Base value: $\text{leniw} = \text{leniw} + (\text{maxord}+5)N_q$
- With `itolQ = CV_SV`: $\text{leniw} = \text{leniw} + N_q$

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

<div style="border: 1px solid black; padding: 2px; display: inline-block;">CVodeQuadReInit</div>	
Call	<code>flag = CVodeQuadReInit(cvode_mem, fQ, yQ0);</code>
Description	The function <code>CVodeQuadReInit</code> provides required problem specifications and reinitializes the quadrature integration.
Arguments	<p><code>cvode_mem</code> (void *) pointer to the CVODES memory block.</p> <p><code>fQ</code> (<code>CVQuadRhsFn</code>) is the C function which computes f_Q, the right-hand side of the quadrature equations.</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>CV_SUCCESS</code> The call to <code>CVodeReInit</code> was successful.</p> <p><code>CV_MEM_NULL</code> The CVODES memory was not initialized by a prior call to <code>CVodeCreate</code>.</p> <p><code>CV_NO_QUAD</code> Memory space for the quadrature integration was not allocated by a prior call to <code>CVodeQuadMalloc</code>.</p>
Notes	If an error occurred, <code>CVodeQuadReInit</code> also sends an error message to the error handler function.

5.7.2 CVODE solver function

Even if quadrature integration was enabled, the call to the main solver function `CVode` is exactly the same as in §5.5.4. However, in this case the return value `flag` can also be one of the following:

<code>CV_QRHSFUNC_FAIL</code>	The quadrature right-hand side function failed in an unrecoverable manner.
<code>CV_FIRST_QRHSFUNC_FAIL</code>	The quadrature right-hand side function failed at the first call.
<code>CV_REPTD_QRHSFUNC_ERR</code>	Convergence tests occurred too many times due to repeated recoverable errors in the quadrature right-hand side function. The <code>CV_REPTD_RHSFUNC_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 quadrature variables are included in the error tests).
<code>CV_UNREC_RHSFUNC_ERR</code>	The quadrature right-hand function had a recoverable error, but no recovery was possible. This failure mode is rare, as it can occur only if the quadrature right-hand side function fails recoverably after an error test failed while at order one.

5.7.3 Quadrature extraction functions

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

<div>CVodeGetQuad</div>	
Call	<code>flag = CVodeGetQuad(cvode_mem, t, yQ);</code>
Description	The function <code>CVodeGetQuad</code> returns the quadrature solution vector after a successful return from <code>CVode</code> .
Arguments	<p><code>cvode_mem</code> (<code>void *</code>) pointer to the memory previously allocated by <code>CVodeMalloc</code>.</p> <p><code>t</code> (<code>realtype</code>) the time at which quadrature information is requested. The time <code>t</code> must fall within the interval defined by the last successful step taken by CVODES.</p> <p><code>yQ</code> (<code>N_Vector</code>) the computed quadrature vector.</p>
Return value	<p>The return value <code>flag</code> of <code>CVodeGetQuad</code> is one of:</p> <p><code>CV_SUCCESS</code> <code>CVodeGetQuad</code> was successful.</p> <p><code>CV_MEM_NULL</code> <code>cvode_mem</code> was NULL.</p> <p><code>CV_NO_QUAD</code> Quadrature integration was not initialized.</p> <p><code>CV_BAD_DKY</code> <code>yQ</code> is NULL.</p> <p><code>CV_BAD_T</code> The time <code>t</code> is not in the allowed range.</p>
Notes	<p>In case of an error return, an error message is also sent to the error handler function.</p> <p>The function <code>CVodeGetQuadDky</code> computes the <code>k</code>-th derivatives of the interpolating polynomials for the quadrature variables at time <code>t</code>. This function is called by <code>CVodeGetQuad</code> with <code>k = 0</code>, but may also be called directly by the user.</p>

<div>CVodeGetQuadDky</div>	
Call	<code>flag = CVodeGetQuadDky(cvode_mem, t, k, dkyQ);</code>
Description	The function <code>CVodeGetQuadDky</code> returns derivatives of the quadrature solution vector after a successful return from <code>CVode</code> .
Arguments	<code>cvode_mem</code> (<code>void *</code>) pointer to the memory previously allocated by <code>CVodeMalloc</code> .

<code>t</code>	(<code>realtype</code>) the time at which quadrature information is requested. The time <code>t</code> must fall within the interval defined by the last successful step taken by CVODES.
<code>k</code>	(<code>int</code>) order of the requested derivative.
<code>dkyQ</code>	(<code>N_Vector</code>) the vector containing the derivative. This vector must be allocated by the user.

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

<code>CV_SUCCESS</code>	<code>CVodeGetQuadDky</code> succeeded.
<code>CV_MEM_NULL</code>	The pointer to <code>cnode_mem</code> was NULL.
<code>CV_NO_QUAD</code>	Quadrature integration was not initialized.
<code>CV_BAD_DKY</code>	The vector <code>dkyQ</code> is NULL.
<code>CV_BAD_K</code>	<code>k</code> is not in the range $0, 1, \dots, q_u$.
<code>CV_BAD_T</code>	The time <code>t</code> is not in the allowed range.

Notes In case of an error return, an error message is also sent to the error handler function.

5.7.4 Optional inputs for quadrature integration

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

CVodeSetQuadFdata

Call	<code>flag = CVodeSetQuadFdata(cnode_mem, fQ_data);</code>
Description	The function <code>CVodeSetQuadFdata</code> specifies the user-defined data block <code>fQ_data</code> and attaches it to the main CVODES memory block.
Arguments	<code>cnode_mem</code> (<code>void *</code>) pointer to the CVODES memory block. <code>fQ_data</code> (<code>void *</code>) pointer to the user data.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of: <code>CV_SUCCESS</code> The optional value has been successfully set. <code>CV_MEM_NULL</code> The <code>cnode_mem</code> pointer is NULL.
Notes	If <code>fQ_data</code> is not specified, a NULL pointer is passed to all user-supplied functions that have it as an argument. Note that <code>fQ_data</code> can be the same as the pointer <code>f_data</code> set through <code>CVodeSetFdata</code> .

CVodeSetQuadErrCon

Call	<code>flag = CVodeSetQuadErrCon(cnode_mem, errconQ, itolQ, reltolQ, abstolQ);</code>
Description	The function <code>CVodeSetQuadErrCon</code> specifies whether or not the quadrature variables should be used in the step size control mechanism, and if so, specifies the integration tolerances for the quadrature variables.
Arguments	<code>cnode_mem</code> (<code>void *</code>) pointer to the CVODES memory block. <code>errconQ</code> (<code>booleantype</code>) specifies whether quadrature variables are included (<code>TRUE</code>) or not (<code>FALSE</code>) in the error control mechanism. If <code>errconQ=FALSE</code> , the following three arguments are ignored. <code>itolQ</code> (<code>int</code>) is either <code>CV_SS</code> or <code>CV_SV</code> , where <code>itolQ = CV_SS</code> indicates scalar relative error tolerance and scalar absolute error tolerance, while <code>itolQ = CV_SV</code> indicates scalar relative error tolerance and vector absolute error tolerance. The latter choice is important when the absolute error tolerance needs to be different for each quadrature variable. <code>reltolQ</code> (<code>realtype *</code>) is a pointer to the relative error tolerance.

abstolQ (void *) is a pointer to the absolute error tolerance. If **itolQ**=CV_SS, **abstolQ** must be a pointer to a **realtype** variable. If **itolQ** = CV_SV, **abstolQ** must be an **N_Vector** variable.

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

CV_SUCCESS The optional value has been successfully set.

CV_MEM_NULL The **cvode_mem** pointer is NULL.

CV_ILL_INPUT An input argument to **CVodeSetQuadErrCon** has an illegal value.

Notes By default, **errconQ** is set to **FALSE**.

It is illegal to call **CVodeSetQuadErrCon** before a call to **CVodeQuadMalloc**.



5.7.5 Optional outputs for quadrature integration

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

CVodeGetQuadNumRhsEvals

Call `flag = CVodeGetQuadNumRhsEvals(cvode_mem, &nfQevals);`

Description The function **CVodeGetQuadNumRhsEvals** returns the number of calls made to the user's quadrature right-hand side function.

Arguments **cvode_mem** (void *) pointer to the CVODES memory block.

nfQevals (long int) number of calls made to the user's **fQ** function.

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

CV_SUCCESS The optional output value has been successfully set.

CV_MEM_NULL The **cvode_mem** pointer is NULL.

CV_NO_QUAD Quadrature integration has not been initialized.

CVodeGetQuadNumErrTestFails

Call `flag = CVodeGetQuadNumErrTestFails(cvode_mem, &nQetfails);`

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

Arguments **cvode_mem** (void *) pointer to the CVODES 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:

CV_SUCCESS The optional output value has been successfully set.

CV_MEM_NULL The **cvode_mem** pointer is NULL.

CV_NO_QUAD Quadrature integration has not been initialized.

CVodeGetQuadErrWeights

Call `flag = CVodeGetQuadErrWeights(cvode_mem, eQweight);`

Description The function **CVodeGetQuadErrWeights** returns the quadrature error weights at the current time.

Arguments **cvode_mem** (void *) pointer to the CVODES memory block.

eQweight (**N_Vector**) quadrature error weights at the current time.

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

	CV_SUCCESS	The optional output value has been successfully set.
	CV_MEM_NULL	The <code>cnode_mem</code> pointer is NULL.
	CV_NO_QUAD	Quadrature integration has not been initialized.
Notes	<p>The user must allocate memory for <code>eQweight</code>.</p> <p>If quadratures were not included in the error control mechanism (through a call to <code>CVodeSetQuadErrCon</code> with <code>errconQ = TRUE</code>), <code>CVodeGetQuadErrWeights</code> does not set the <code>eQweight</code> vector.</p>	



CVodeGetQuadStats

Call	<code>flag = CVodeGetQuadStats(cnode_mem, &nfQevals, &nQetfails);</code>	
Description	The function <code>CVodeGetQuadStats</code> returns the CVODES integrator statistics as a group.	
Arguments	<p><code>cnode_mem</code> (void *) pointer to the CVODES memory block.</p> <p><code>nfQevals</code> (long int) number of calls to the user's <code>fQ</code> function.</p> <p><code>nQetfails</code> (long int) number of error test failures due to quadrature variables.</p>	
Return value	<p>The return value <code>flag</code> (of type <code>int</code>) is one of</p> <p>CV_SUCCESS the optional output values have been successfully set.</p> <p>CV_MEM_NULL the <code>cnode_mem</code> pointer is NULL.</p> <p>CV_NO_QUAD Quadrature integration has not been initialized.</p>	

5.7.6 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. This function must be of type `CVQuadRhsFn` defined as follows:

CVQuadRhsFn

Definition	<pre>typedef int (*CVQuadRhsFn)(realtype t, N_Vector y, N_Vector yQdot, void *fQ_data);</pre>	
Purpose	This function computes the quadrature equation right-hand side for a given value of the independent variable t and state vector y .	
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>yQdot</code> is the output vector $f_Q(t, y)$.</p> <p><code>fQ_data</code> is the <code>fQ_data</code> pointer passed to <code>CVodeSetQuadFdata</code>.</p>	
Return value	A <code>CVQuadRhsFn</code> should return 0 if successful, a positive value if a recoverable error occurred (in which case CVODES will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and <code>CV_QRHSFUNC_FAIL</code> is returned).	
Notes	<p>Allocation of memory for <code>yQdot</code> is automatically handled within CVODES.</p> <p>Both <code>y</code> and <code>yQdot</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 CVODES do not perform any consistency checks with respect to their <code>N_Vector</code> arguments (see §8.1 and §8.2).</p> <p>There are two situations in which recovery is not possible even if <code>CVQuadRhsFn</code> function returns a recoverable error flag. This include the situation when this occurs at the very first call to the <code>CVQuadRhsFn</code> (in which case CVODES returns <code>CV_FIRST_QRHSFUNC_ERR</code>).</p>	

or if a recoverable error is reported when `CVQuadRhsFn` is called after an error test failure, while the linear multistep method order is equal to 1 (in which case CVODES returns `CV_UNREC_QRHSFUNC_ERR`).

5.8 Rootfinding

While solving the IVP, CVODES has the capability to find the roots of a set of user-defined functions. This section describes the user-callable functions used to initialize and define the rootfinding problem and to obtain solution information, and it also describes the required user-supplied function.

5.8.1 User-callable functions for rootfinding

CVodeRootInit

Call	<code>flag = CVodeRootInit(cvode_mem, nrtfn, g, g_data);</code>
Description	The function <code>CVodeRootInit</code> specifies that the roots of a set of functions $g_i(t, y)$ are to be found while the IVP is being solved.
Arguments	<p><code>cvode_mem</code> (void *) pointer to the CVODES memory block returned by <code>CVodeCreate</code>.</p> <p><code>nrtfn</code> (int) is the number of root functions g_i.</p> <p><code>g</code> (CVRtFn) is the C function which defines the <code>nrtfn</code> functions $g_i(t, y)$ whose roots are sought. See §5.8.2 for details.</p> <p><code>g_data</code> (void *) pointer to the user data for use by the user's root function g.</p>
Return value	<p>The return value <code>flag</code> (of type <code>int</code>) is one of</p> <p><code>CV_SUCCESS</code> The call to <code>CVodeRootInit</code> was successful.</p> <p><code>CV_MEM_NULL</code> The <code>cvode_mem</code> argument was NULL.</p> <p><code>CV_MEM_FAIL</code> A memory allocation failed.</p> <p><code>CV_ILL_INPUT</code> The function <code>g</code> is NULL, but <code>nrtfn</code> > 0.</p>
Notes	If a new IVP is to be solved with a call to <code>CVodeReInit</code> , where the new IVP has no rootfinding problem but the prior one did, then call <code>CVodeRootInit</code> with <code>nrtfn</code> =0.

There are two optional output functions associated with rootfinding.

CVodeGetRootInfo

Call	<code>flag = CVodeGetRootInfo(cvode_mem, rootsfound);</code>
Description	The function <code>CVodeGetRootInfo</code> returns an array showing which functions were found to have a root.
Arguments	<p><code>cvode_mem</code> (void *) pointer to the CVODES memory block.</p> <p><code>rootsfound</code> (int *) an <code>int</code> array of length <code>nrtfn</code> showing the indices of the user functions g_i found to have a root. For $i = 0, \dots, \text{nrtfn}-1$, <code>rootsfound[i]</code> = 1 if g_i has a root, and 0 if not.</p>
Return value	<p>The return value <code>flag</code> (of type <code>int</code>) is one of:</p> <p><code>CV_SUCCESS</code> The optional output values have been successfully set.</p> <p><code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL.</p>
Notes	The user must allocate memory for the vector <code>rootsfound</code> .



CVodeGetNumGEvals

Call	<code>flag = CVodeGetNumGEvals(cvode_mem, &ngevals);</code>
Description	The function <code>CVodeGetNumGEvals</code> returns the cumulative number of calls made to the user-supplied root function g .
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODES memory block. <code>ngevals</code> (long int) number of calls made to the user's function g thus far.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of: <code>CV_SUCCESS</code> The optional output value has been successfully set. <code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL.

5.8.2 User-supplied function for rootfinding

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

CVRootFn

Definition	<code>typedef int (*CVRootFn)(realtype t, N_Vector y, realtype *gout, void *g_data);</code>
Purpose	This function implements a vector-valued function $g(t, y)$ such that the roots of the <code>nrtfn</code> components $g_i(t, y)$ are sought.
Arguments	<code>t</code> is the current value of the independent variable. <code>y</code> is the current value of the dependent variable vector, $y(t)$. <code>gout</code> is the output array, of length <code>nrtfn</code> , with components $g_i(t, y)$. <code>g_data</code> is the <code>g_data</code> pointer passed to <code>CVodeRootInit</code> .
Return value	A <code>CVRootFn</code> should return 0 if successful or a non-zero value if an error occurred (in which case the integration is halted and <code>CVode</code> returns <code>CV_RTFUNC_FAIL</code>).
Notes	Allocation of memory for <code>gout</code> is automatically handled within CVODES.

5.9 Preconditioner modules

The efficiency of Krylov iterative methods for the solution of linear systems can be greatly enhanced through preconditioning. For problems in which the user cannot define a more effective, problem-specific preconditioner, CVODES provides a banded preconditioner in the module `CVBANDPRE` and a band-block-diagonal preconditioner module `CVBBDPRE`.

5.9.1 A serial banded preconditioner module

This preconditioner provides a band matrix preconditioner for use with any of the Krylov iterative linear solvers, in a serial setting. It uses difference quotients of the ODE right-hand side function \mathbf{f} to generate a band matrix of bandwidth $m_l + m_u + 1$, where the number of super-diagonals (m_u , the upper half-bandwidth) and sub-diagonals (m_l , the lower half-bandwidth) are specified by the user, and uses this to form a preconditioner for use with the Krylov linear solver. Although this matrix is intended to approximate the Jacobian $\partial \mathbf{f} / \partial \mathbf{y}$, it may be a very crude approximation. The true Jacobian need not be banded, or its true bandwidth may be larger than $m_l + m_u + 1$, as long as the banded approximation generated here is sufficiently accurate to speed convergence as a preconditioner.

In order to use the `CVBANDPRE` module, the user need not define any additional functions. Aside from the header files required for the integration of the ODE problem (see §5.3), to use the `CVBANDPRE` module, the main program must include the header file `cvodes_bandpre.h` which declares the needed function prototypes. The following is a summary of the usage of this module. Steps that are unchanged from the skeleton program presented in §5.4 are grayed out.

1. Set problem dimensions
2. Set vector of initial values
3. Create CVODES object
4. Allocate internal memory
5. Set optional inputs

6. **Initialize the CVBANDPRE preconditioner module**

Specify the upper and lower half-bandwidths (`mu` and `m1`, respectively) and call

```
bp_data = CVBandPrecAlloc(cvode_mem, N, mu, m1);
```

to allocate memory for and to initialize a data structure (pointed to by `bp_data`) to be passed to the appropriate CVSp* linear solver.

7. **Attach the Krylov linear solver, one of:**

```
flag = CVBPSpgmr(cvode_mem, pretype, maxl, bp_data);
```

```
flag = CVBPSpbcg(cvode_mem, pretype, maxl, bp_data);
```

```
flag = CVBPSptfqmr(cvode_mem, pretype, maxl, bp_data);
```

Each function CVBPSp* is a wrapper around the corresponding specification function CVSp* and performs the following actions:

- Attaches the CVSPILS linear solver to the main CVODES solver memory;
- Sets the preconditioner data structure for CVBANDPRE;
- Sets the preconditioner setup function for CVBANDPRE;
- Sets the preconditioner solve function for CVBANDPRE;

The arguments `pretype` and `maxl` are described below. The last argument of CVBPSp* is the pointer to the CVBANDPRE data returned by CVBandPrecAlloc.

8. **Set linear solver optional inputs**

Note that the user should not overwrite the preconditioner data, setup function, or solve function through calls to CVSp* optional input functions.

9. **Advance solution in time**

10. **Deallocate memory for solution vector**

11. **Free the CVBANDPRE data structure**

```
CVBandPrecFree(&bp_data);
```

12. **Free solver memory**

The user-callable functions that initialize, attach, and deallocate the CVBANDPRE preconditioner module (steps 6, 7 and 11 above) are described in more detail below.

CVBandPrecAlloc

Call `bp_data = CVBandPrecAlloc(cvode_mem, N, mu, m1);`

Description The function CVBandPrecAlloc initializes and allocates memory for the CVBANDPRE preconditioner.

Arguments `cvode_mem` (void *) pointer to the CVODES memory block.

	N	(long int) problem dimension.
	mu	(long int) upper half-bandwidth of the Jacobian approximation.
	ml	(long int) lower half-bandwidth of the Jacobian approximation.
Return value	If successful, CVBandPrecAlloc returns a pointer to the newly created CVBANDPRE memory block (of type void *). If an error occurred, CVBandPrecAlloc returns NULL .	
Notes	The banded approximate Jacobian will have nonzero elements only in locations (i, j) with $-ml \leq j - i \leq mu$.	

CVBPSpgmr

Call	flag = CVBPSpgmr(cvmem , pretype , maxl , bp_data);	
Description	The function CVBPSpgmr links the CVBANDPRE data to the CVSPGMR linear solver and attaches the latter to the CVOIDES memory block.	
Arguments	cvmem (void *) pointer to the CVOIDES memory block. pretype (int) specifies the preconditioning type and must be either PREC_LEFT or PREC_RIGHT . maxl (int) maximum dimension of the Krylov subspace to be used. Pass 0 to use the default value CVSPILS_MAXL = 5. bp_data (void *) pointer to the CVBANDPRE data structure.	
Return value	The return value flag (of type int) is one of:	
	CVSPILS_SUCCESS	The CVSPGMR initialization was successful.
	CVSPILS_MEM_NULL	The cvmem pointer is NULL .
	CVSPILS_ILL_INPUT	The preconditioner type pretype is not valid.
	CVSPILS_MEM_FAIL	A memory allocation request failed.
	CVBANDPRE_PDATA_NULL	The CVBANDPRE preconditioner has not been initialized.

CVBPSpbcg

Call	flag = CVBPSpbcg(cvmem , pretype , maxl , bp_data);	
Description	The function CVBPSpbcg links the CVBANDPRE data to the CVSPBCG linear solver and attaches the latter to the CVOIDES memory block.	
Arguments	cvmem (void *) pointer to the CVOIDES memory block. pretype (int) preconditioning type. Must be one of PREC_LEFT or PREC_RIGHT . maxl (int) maximum dimension of the Krylov subspace to be used. Pass 0 to use the default value CVSPILS_MAXL = 5. bp_data (void *) pointer to the CVBANDPRE data structure.	
Return value	The return value flag (of type int) is one of	
	CVSPILS_SUCCESS	The CVSPBCG initialization was successful.
	CVSPILS_MEM_NULL	The cvmem pointer is NULL .
	CVSPILS_ILL_INPUT	The preconditioner type pretype is not valid.
	CVSPILS_MEM_FAIL	A memory allocation request failed.
	CVBANDPRE_PDATA_NULL	The CVBANDPRE preconditioner has not been initialized.

CVBPSptfqmr

Call	flag = CVBPSptfqmr(cvmem , pretype , maxl , bp_data);	
Description	The function CVBPSptfqmr links the CVBANDPRE data to the CVSPTFQMR linear solver and attaches the latter to the CVOIDES memory block.	

Arguments `cvode_mem` (void *) pointer to the CVODES memory block.
 `pretype` (int) preconditioning type. Must be one of `PREC_LEFT` or `PREC_RIGHT`.
 `maxl` (int) maximum dimension of the Krylov subspace to be used. Pass 0 to use the default value `CVSPILS_MAXL = 5`.
 `bp_data` (void *) pointer to the CVBANDPRE data structure.

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

<code>CVSPILS_SUCCESS</code>	The CVSPTFQMR initialization was successful.
<code>CVSPILS_MEM_NULL</code>	The <code>cvode_mem</code> pointer is NULL.
<code>CVSPILS_ILL_INPUT</code>	The preconditioner type <code>pretype</code> is not valid.
<code>CVSPILS_MEM_FAIL</code>	A memory allocation request failed.
<code>CVBANDPRE_PDATA_NULL</code>	The CVBANDPRE preconditioner has not been initialized.

CVBandPrecFree

Call `CVBandPrecFree(&bp_data);`

Description The function `CVBandPrecFree` frees the pointer allocated by `CVBandPrecAlloc`.

Arguments The only argument passed to `CVBandPrecFree` is the pointer to the CVBANDPRE data structure (of type `void *`).

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

The following three optional output functions are available for use with the CVBANDPRE module:

CVBandPrecGetWorkSpace

Call `flag = CVBandPrecGetWorkSpace(bp_data, &lenrwBP, &leniwBP);`

Description The function `CVBandPrecGetWorkSpace` returns the sizes of the CVBANDPRE real and integer workspaces.

Arguments `bp_data` (void *) pointer to the CVBANDPRE data structure.
 `lenrwBP` (long int) the number of `realtype` values in the CVBANDPRE workspace.
 `leniwBP` (long int) the number of integer values in the CVBANDPRE workspace.

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

<code>CVBANDPRE_SUCCESS</code>	The optional output values have been successfully set.
<code>CVBANDPRE_PDATA_NULL</code>	The CVBANDPRE preconditioner has not been initialized.

Notes In terms of problem size N and $\text{smu} = \min(N - 1, \text{mu} + \text{ml})$, the actual size of the real workspace is $(2 \text{ml} + \text{mu} + \text{smu} + 2) N$ `realtype` words, and the actual size of the integer workspace is N integer words.

The workspaces referred to here exist in addition to those given by the corresponding `CVSp***GetWorkSpace` function.

CVBandPrecGetNumRhsEvals

Call `flag = CVBandPrecGetNumRhsEvals(bp_data, &nfevalsBP);`

Description The function `CVBandPrecGetNumRhsEvals` returns the number of calls made to the user-supplied right-hand side function for finite difference banded Jacobian approximation used within the preconditioner setup function.

Arguments `bp_data` (void *) pointer to the CVBANDPRE data structure.
 `nfevalsBP` (long int) the number of calls to the user right-hand side function.

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

<code>CVBANDPRE_SUCCESS</code>	The optional output value has been successfully set.
--------------------------------	--

CVBANDPRE_PDATA_NULL The CVBANDPRE preconditioner has not been initialized.

Notes The counter `nfevalsBP` is distinct from the counter `nfevalsLS` returned by the corresponding `CVSp***GetNumRhsEvals` function, and also from `nfevals`, returned by `CVodeGetNumRhsEvals`. The total number of right-hand side function evaluations is the sum of all three of these counters.

CVBandPrecGetReturnFlagName

Call `name = CVBandPrecGetReturnFlagName(flag);`

Description The function `CVBandPrecGetReturnFlagName` returns the name of the CVBANDPRE constant corresponding to `flag`.

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

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

5.9.2 A parallel band-block-diagonal preconditioner module

A principal reason for using a parallel ODE solver such as CVODES 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 (3.5) 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 [20] and is included in a software module within the CVODES package. This module works with the parallel vector module `NVECTOR_PARALLEL` and is usable with any of the Krylov iterative linear solvers. It 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 CVBBDPRE.

One way to envision these preconditioners is to think of the domain of the computational PDE problem as being subdivided into M non-overlapping subdomains. Each of these subdomains is then assigned to one of the M processes to be used to solve the ODE system. The basic idea is to isolate the preconditioning so that it is local to each process, and also to use a (possibly cheaper) approximate right-hand side function. This requires the definition of a new function $g(t, y)$ which approximates the function $f(t, y)$ in the definition of the ODE system (3.1). However, the user may set $g = f$. Corresponding to the domain decomposition, there is a decomposition of the solution vector y into M disjoint blocks y_m , and a decomposition of g into blocks g_m . The block g_m depends both on y_m and on components of blocks $y_{m'}$ associated with neighboring subdomains (so-called ghost-cell data). Let \bar{y}_m denote y_m augmented with those other components on which g_m depends. Then we have

$$g(t, y) = [g_1(t, \bar{y}_1), g_2(t, \bar{y}_2), \dots, g_M(t, \bar{y}_M)]^T \quad (5.1)$$

and each of the blocks $g_m(t, \bar{y}_m)$ is uncoupled from the others.

The preconditioner associated with this decomposition has the form

$$P = \text{diag}[P_1, P_2, \dots, P_M] \quad (5.2)$$

where

$$P_m \approx I - \gamma J_m \quad (5.3)$$

and J_m is a difference quotient approximation to $\partial g_m / \partial y_m$. This matrix is taken to be banded, with upper and lower half-bandwidths `mudq` and `mlldq` defined as the number of non-zero diagonals above and below the main diagonal, respectively. The difference quotient approximation is computed using

$\text{mudq} + \text{mldq} + 2$ evaluations of g_m , but only a matrix of bandwidth $\text{mu} + \text{ml} + 1$ is retained. Neither pair of parameters need be the true half-bandwidths of the Jacobian of the local block of g , if smaller values provide a more efficient preconditioner. The solution of the complete linear system

$$Px = b \quad (5.4)$$

reduces to solving each of the equations

$$P_m x_m = b_m \quad (5.5)$$

and this is done by banded LU factorization of P_m followed by a banded backsolve.

Similar block-diagonal preconditioners could be considered with different treatments of the blocks P_m . For example, incomplete LU factorization or an iterative method could be used instead of banded LU factorization.

The CVBBDPRE module calls two user-provided functions to construct P : a required function **gloc** (of type **CVLocalFn**) which approximates the right-hand side function $g(t, y) \approx f(t, y)$ and which is computed locally, and an optional function **cfm** (of type **CVCommFn**) which performs all interprocess communication necessary to evaluate the approximate right-hand side g . These are in addition to the user-supplied right-hand side function **f**. Both functions take as input the same pointer **f_data** that is passed by the user to **CVodeSetFdata** and that was passed to the user's function **f**, and neither function has a return value. The user is responsible for providing space (presumably within **f_data**) for components of **y** that are communicated between processes by **cfm**, and that are then used by **gloc**, which is not expected to do any communication.

CVLocalFn

Definition	<pre>typedef int (*CVLocalFn)(long int Nlocal, realtype t, N_Vector y, N_Vector glocal, void *f_data);</pre>
Purpose	This function computes $g(t, y)$. It loads the vector glocal as a function of t and y .
Arguments	<p>Nlocal is the local vector length.</p> <p>t is the value of the independent variable.</p> <p>y is the dependent variable.</p> <p>glocal is the output vector.</p> <p>f_data is the f_data pointer passed to CVodeSetFdata.</p>
Return value	A CVLocalFn should return 0 if successful, a positive value if a recoverable error occurred (in which case CVODES will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and CVode returns CV_LSETUP_FAIL).
Notes	<p>This function assumes that all interprocess communication of data needed to calculate glocal has already been done, and that this data is accessible within f_data.</p> <p>The case where g is mathematically identical to f is allowed.</p>

CVCommFn

Definition	<pre>typedef int (*CVCommFn)(long int Nlocal, realtype t, N_Vector y, void *f_data);</pre>
Purpose	This function performs all interprocess communication necessary for the execution of the gloc function above, using the input vector y .
Arguments	<p>Nlocal is the local vector length.</p> <p>t is the value of the independent variable.</p> <p>y is the dependent variable.</p> <p>f_data is the f_data pointer passed to CVodeSetFdata.</p>
Return value	A CVCommFn should return 0 if successful, a positive value if a recoverable error occurred (in which case CVODES will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and CVode returns CV_LSETUP_FAIL).

Notes The `cfn` function is expected to save communicated data in space defined within the data structure `f_data`.

Each call to the `cfn` function is preceded by a call to the right-hand side function `f` with the same `(t, y)` arguments. Thus, `cfn` can omit any communication done by `f` if relevant to the evaluation of `glocal`. If all necessary communication was done in `f`, then `cfn = NULL` can be passed in the call to `CVBBDPrecAlloc` (see below).

Besides the header files required for the integration of the ODE problem (see §5.3), to use the CVBBDPRE module, the main program must include the header file `cvodes.bbdpre.h` which declares the needed function prototypes.

The following is a summary of the proper usage of this module. Steps that are unchanged from the skeleton program presented in §5.4 are grayed out.

1. Initialize MPI
2. Set problem dimensions
3. Set vector of initial values
4. Create CVODES object
5. Allocate internal memory
6. Set optional inputs
7. Initialize the CVBBDPRE preconditioner module

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

```
bbd_data = CVBBDPrecAlloc(cvode_mem, local_N, mudq, mldq,
                          mukeep, mlkeep, dqrely, gloc, cfn);
```

to allocate memory for and to initialize a data structure `bbd_data` (of type `void *`) to be passed to the Krylov linear solver selected (in the next step). The last two arguments passed to `CVBBDPrecAlloc` are the two user-supplied functions described above.

8. Attach the Krylov linear solver, one of:

```
flag = CVBBDSPgmrcvode(cvode_mem, pretype, maxl, bbd_data);
flag = CVBBDSPbcgcvode(cvode_mem, pretype, maxl, bbd_data);
flag = CVBBDSPtfqmr(cvode_mem, pretype, maxl, bbd_data);
```

The function `CVBBDSP*` is a wrapper around the corresponding specification function `CVSp*` and performs the following actions:

- Attaches the CVSPILS linear solver to the main CVODES solver memory;
- Sets the preconditioner data structure for CVBBDPRE;
- Sets the preconditioner setup function for CVBBDPRE;
- Sets the preconditioner solve function for CVBBDPRE;

The arguments `pretype` and `maxl` are described below. The last argument of `CVBBDSP*` is the pointer to the CVBBDPRE data returned by `CVBBDPrecAlloc`.

9. Set linear solver optional inputs

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

10. Advance solution in time

11. Deallocate memory for solution vector

12. Free the CVBBDPRE data structure

```
CVBBDPrecFree(&bbd_data);
```

13. Free solver memory

14. Finalize MPI

The user-callable functions that initialize, attach, and deallocate the CVBBDPRE preconditioner module (steps 7, 8, and 12 above) are described next.

CVBBDPrecAlloc

Call	<code>bbd_data = CVBBDPrecAlloc(cvode_mem, local_N, mudq, mldq, mukeep, mlkeep, dqrely, gloc, cfn);</code>
Description	The function <code>CVBBDPrecAlloc</code> initializes and allocates memory for the CVBBDPRE preconditioner.
Arguments	<p><code>cvode_mem</code> (void *) pointer to the CVODES memory block.</p> <p><code>local_N</code> (long int) local vector length.</p> <p><code>mudq</code> (long int) upper half-bandwidth to be used in the difference quotient Jacobian approximation.</p> <p><code>mldq</code> (long int) lower half-bandwidth to be used in the difference quotient Jacobian approximation.</p> <p><code>mukeep</code> (long int) upper half-bandwidth of the retained banded approximate Jacobian block.</p> <p><code>mlkeep</code> (long int) lower half-bandwidth of the retained banded approximate Jacobian block.</p> <p><code>dqrely</code> (realtype) the relative increment in components of y used in the difference quotient approximations. The default is <code>dqrely = $\sqrt{\text{unit roundoff}}$</code>, which can be specified by passing <code>dqrely = 0.0</code>.</p> <p><code>gloc</code> (CVLocalFn) the C function which computes the approximation $g(t, y) \approx f(t, y)$.</p> <p><code>cfn</code> (CVCommFn) the optional C function which performs all interprocess communication required for the computation of $g(t, y)$.</p>
Return value	If successful, <code>CVBBDPrecAlloc</code> returns a pointer to the newly created CVBBDPRE memory block (of type void *). If an error occurred, <code>CVBBDPrecAlloc</code> returns NULL.
Notes	<p>If one of the half-bandwidths <code>mudq</code> or <code>mldq</code> to be used in the difference quotient calculation of the approximate Jacobian is negative or exceeds the value <code>local_N-1</code>, it is replaced with 0 or <code>local_N-1</code> accordingly.</p> <p>The half-bandwidths <code>mudq</code> and <code>mldq</code> 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 <code>mukeep</code> and <code>mlkeep</code> of the retained banded approximate Jacobian block may be even smaller, to reduce storage and computational costs further.</p> <p>For all four half-bandwidths, the values need not be the same for every process.</p>

CVBBDSpgmr

Call	<code>flag = CVBBDSpgmr(cvode_mem, pretype, maxl, bbd_data);</code>
Description	The function <code>CVBBDSpgmr</code> links the CVBBDPRE data to the CVSPGMR linear solver and attaches the latter to the CVODES memory block.

Arguments `cnode_mem` (void *) pointer to the CVODES memory block.
 `pretype` (int) preconditioning type. Must be either `PREC_LEFT` or `PREC_RIGHT`.
 `maxl` (int) maximum dimension of the Krylov subspace to be used. Pass 0 to use
 the default value `CVSPILS_MAXL = 5`.
 `bbd_data` (void *) pointer to the CVBBDPRE data structure.

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

`CVSPILS_SUCCESS` The CVSPGMR initialization was successful.
`CVSPILS_MEM_NULL` The `cnode_mem` pointer is NULL.
`CVSPILS_ILL_INPUT` The preconditioner type `pretype` is not valid.
`CVSPILS_MEM_FAIL` A memory allocation request failed.
`CVBBDPRE_PDATA_NULL` The CVBBDPRE preconditioner has not been initialized.

CVBBDSpbcg

Call `flag = CVBBDSpbcg(cnode_mem, pretype, maxl, bbd_data);`

Description The function `CVBBDSpbcg` links the CVBBDPRE data to the CVSPBCG linear solver and attaches the latter to the CVODES memory block.

Arguments `cnode_mem` (void *) pointer to the CVODES memory block.
 `pretype` (int) preconditioning type. Must be one of `PREC_LEFT` or `PREC_RIGHT`.
 `maxl` (int) maximum dimension of the Krylov subspace to be used. Pass 0 to use
 the default value `CVSPILS_MAXL = 5`.
 `bbd_data` (void *) pointer to the CVBBDPRE data structure.

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

`CVSPILS_SUCCESS` The CVSPBCG initialization was successful.
`CVSPILS_MEM_NULL` The `cnode_mem` pointer is NULL.
`CVSPILS_ILL_INPUT` The preconditioner type `pretype` is not valid.
`CVSPILS_MEM_FAIL` A memory allocation request failed.
`CVBBDPRE_PDATA_NULL` The CVBBDPRE preconditioner has not been initialized.

CVBBDSpTFqmr

Call `flag = CVBBDSpTFqmr(cnode_mem, pretype, maxl, bbd_data);`

Description The function `CVBBDSpTFqmr` links the CVBBDPRE data to the CVSPTFQMR linear solver and attaches the latter to the CVODES memory block.

Arguments `cnode_mem` (void *) pointer to the CVODES memory block.
 `pretype` (int) preconditioning type. Must be one of `PREC_LEFT` or `PREC_RIGHT`.
 `maxl` (int) maximum dimension of the Krylov subspace to be used. Pass 0 to use
 the default value `CVSPILS_MAXL = 5`.
 `bbd_data` (void *) pointer to the CVBBDPRE data structure.

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

`CVSPILS_SUCCESS` The CVSPTFQMR initialization was successful.
`CVSPILS_MEM_NULL` The `cnode_mem` pointer is NULL.
`CVSPILS_ILL_INPUT` The preconditioner type `pretype` is not valid.
`CVSPILS_MEM_FAIL` A memory allocation request failed.
`CVBBDPRE_PDATA_NULL` The CVBBDPRE preconditioner has not been initialized.

CVBBDPrecFree

Call `CVBBDPrecFree(&bbd_data);`

Description The function `CVBBDPrecFree` frees the memory allocated by `CVBBDPrecAlloc`.

Arguments The only argument passed to `CVBBDPrecFree` is the address of the pointer to the CVBBDPRE data structure (of type `void *`).

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

The CVBBDPRE module also provides a reinitialization function to allow solving a sequence of problems of the same size, with the same linear solver choice, provided there is no change in `local_N`, `mukeep`, or `mlkeep`. After solving one problem, and after calling `CVodeReInit` to re-initialize CVODES for a subsequent problem, a call to `CVBBDPrecReInit` 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 `dqrely`, or one of the user-supplied functions `gloc` and `cfn`. If there is a change in any of the linear solver inputs, an additional call to `CVSpgmr`, `CVSpbcg`, or `CVSptfqmr`, and/or one or more of the corresponding `CVSp***Set***` functions, must also be made.

CVBBDPrecReInit

Call `flag = CVBBDPrecReInit(bbd_data, mudq, mldq, dqrely, gloc, cfn);`

Description The function `CVBBDPrecReInit` re-initializes the CVBBDPRE preconditioner.

Arguments `bbd_data` (`void *`) pointer to the CVBBDPRE data structure.

`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.

`dqrely` (`realtype`) the relative increment in components of y used in the difference quotient approximations. The default is $dqrely = \sqrt{\text{unit roundoff}}$, which can be specified by passing `dqrely = 0.0`.

`gloc` (`CVLocalFn`) the C function which computes the approximation $g(t, y) \approx f(t, y)$.

`cfn` (`CVCommFn`) the optional C function which performs all interprocess communication required for the computation of $g(t, y)$.

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

`CVBBDPRE_SUCCESS` The CVSPBCG re-initialization was successful.

`CVBBDPRE_PDATA_NULL` The CVBBDPRE preconditioner has not been initialized.

Notes If one of the half-bandwidths `mudq` or `mldq` is negative or exceeds the value `local_N-1`, it is replaced with 0 or `local_N-1` accordingly.

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

CVBBDPrecGetWorkSpace

Call `flag = CVBBDPrecGetWorkSpace(bbd_data, &lenrwBBDP, &leniwBBDP);`

Description The function `CVBBDPrecGetWorkSpace` returns the local CVBBDPRE real and integer workspace sizes.

Arguments `bbd_data` (`void *`) pointer to the CVBBDPRE data structure.

`lenrwBBDP` (`long int`) local number of `realtype` values in the CVBBDPRE workspace.

`leniwBBDP` (`long int`) local number of integer values in the CVBBDPRE workspace.

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

`CVBBDPRE_SUCCESS` The optional output values have been successfully set.

`CVBBDPRE_PDATA_NULL` The CVBBDPRE preconditioner has not been initialized.

Notes In terms of `local_N` and `smu = min(local_N - 1, mukeep + mlkeep)`, the actual size of the real workspace is $(2 \text{ mlkeep} + \text{mukeep} + \text{smu} + 2) \text{ local_N realtype}$ words, and the actual size of the integer workspace is `local_N` integer words. These values are local to each process.

The workspaces referred to here exist in addition to those given by the corresponding `CVSp***GetWorkSpace` function.

CVBBDPrecGetNumGfnEvals

Call `flag = CVBBDPrecGetNumGfnEvals(bbd_data, &ngevalsBBDP);`

Description The function `CVBBDPrecGetNumGfnEvals` returns the number of calls made to the user-supplied `gloc` function due to the finite difference approximation of the Jacobian blocks used within the preconditioner setup function.

Arguments `bbd_data` (void *) pointer to the CVBBDPRE data structure.
`ngevalsBBDP` (long int) the number of calls made to the user-supplied `gloc` function.

Return value The return value `flag` (of type `int`) is one of
`CVBBDPRE_SUCCESS` The optional output value has been successfully set.
`CVBBDPRE_PDATA_NULL` The CVBBDPRE preconditioner has not been initialized.

CVBBDPrecGetReturnFlagName

Call `name = CVBBDPrecGetReturnFlagName(flag);`

Description The function `CVBBDPrecGetReturnFlagName` returns the name of the CVBBDPRE constant corresponding to `flag`.

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

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

In addition to the `ngevalsBBDP` `gloc` evaluations, the costs associated with CVBBDPRE also include `nlinsetups` LU factorizations, `nlinsetups` calls to `cfn`, `npsolves` banded backsolve calls, and `nfevalsLS` right-hand side function evaluations, where `nlinsetups` is an optional CVODES output and `npsolves` and `nfevalsLS` are linear solver optional outputs (see §5.5.7).

Chapter 6

Using CVODES for Forward Sensitivity Analysis

This chapter describes the use of CVODES to compute solution sensitivities using forward sensitivity analysis. One of our main guiding principles was to design the CVODES 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 right-hand side of the sensitivity systems (3.9). The only departure from this philosophy is due to the `CVRhsFn` type definition (§5.6). Without changing the definition of this type, the only way to pass values of the problem parameters to the ODE right-hand side function is to require the user data structure `f_data` to contain a pointer to the array of real parameters p .

CVODES uses various constants for both input and output. These are defined as needed in this chapter, but for convenience are also listed separately in Chapter 11.

We begin with a brief overview, in the form of a skeleton user program. Following that are detailed descriptions of the interface to the various user-callable routines and of the user-supplied routines that were not already described in §5.

6.1 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 CVODES. The user program is to have these steps in the order indicated, unless otherwise noted. For the sake of brevity, we defer many of the details to the later sections. As in §5.4, 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 CVODES: steps marked with [P] correspond to NVECTOR_PARALLEL, while steps marked with [S] correspond to NVECTOR_SERIAL. Differences between the user main program in §5.4 and the one below start only at step (9).

First, note that no additional header files need be included for forward sensitivity analysis beyond those for IVP solution (§5.4).

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

6. Set optional inputs

7. Attach linear solver module

8. Set linear solver optional inputs

9. Define the sensitivity problem

•Number of sensitivities (required)

Set `Ns`, the number of parameters with respect to which sensitivities are to be computed.

•Problem parameters (optional)

If CVODES will evaluate the right-hand sides of the sensitivity systems, set `p`, an array of `Np` real parameters upon which the IVP depends. Only parameters with respect to which sensitivities are (potentially) desired need to be included. Attach `p` to the user data structure `f_data`. For example, `f_data->p = p`;

If the user provides a function to evaluate the sensitivity right-hand side, `p` need not be specified.

•Parameter list (optional)

If CVODES will evaluate the right-hand sides of the sensitivity systems, set `plist`, an array of `Ns` integer flags to specify the parameters `p` with respect to which solution sensitivities are to be computed. If sensitivities with respect to the j -th problem parameter are desired, set `plisti = j`, for some $i = 0, \dots, N_s - 1$.

If `plist` is not specified, CVODES will compute sensitivities with respect to the first `Ns` parameters; i.e., `plisti = i`, $i = 0, \dots, N_s - 1$.

If the user provides a function to evaluate the sensitivity right-hand side, `plist` need not be specified.

•Parameter scaling factors (optional)

If CVODES estimates tolerances for the sensitivity solution vectors (based on tolerances for the state solution vector) or if CVODES will evaluate the right-hand sides 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 `Ns` positive scaling factors. Typically, if $p_i \neq 0$, the value $\bar{p}_{\text{plist}_i} = |p_i|$ can be used.

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

If the user provides a function to evaluate the sensitivity right-hand side 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 `f_data` are arbitrary, but they must agree with the arguments passed to `CVodeSetSensParams` below.

10. Set sensitivity initial conditions

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

First, create an array of `Ns` vectors by making the call

```
[S] yS0 = N_VNewVectorArray_Serial(Ns, N);
```

```
[P] yS0 = N_VNewVectorArray_Parallel(Ns, N);
```

and, for each $i = 1, \dots, N_s$, load initial values for the i -th sensitivity vector into the structure defined by:

```
[S] NV_DATA_S(yS0[i])
```

```
[P] NV_DATA_P(yS0[i])
```


If the initial values for the sensitivity variables are already available in `realtype` arrays, create an array of `Ns` “empty” vectors by making the call

```
[S] yS0 = N_VNewVectorArrayEmpty_Serial(Ns, N);
```

```
[P] yS0 = N_VNewVectorArrayEmpty_Parallel(Ns, N);
```

and then attach the `realtype` array `yS0_i` containing the initial values of the i -th sensitivity vector using

```
[S] N_VSetArrayPointer_Serial(yS0_i, yS0[i]);
```

```
[P] N_VSetArrayPointer_Parallel(yS0_i, yS0[i]);
```

11. Activate sensitivity calculations

Call `flag = CVodeSensMalloc(...)` to activate forward sensitivity computations and allocate internal memory for CVODES related to sensitivity calculations (see §6.2.1).

12. Set sensitivity analysis optional inputs

Call `CVodeSetSens*` routines to change from their default values any optional inputs that control the behavior of CVODES in computing forward sensitivities.

13. Advance solution in time

14. Extract sensitivity solution

After each successful return from `CVode`, the solution of the original IVP is available in the `y` argument of `CVode`, while the sensitivity solution can be extracted into `yS` (which can be the same as `yS0`) by calling the routine `flag = CVodeGetSens(cvode_mem, t, yS)`; (see §6.2.3).

15. Deallocate memory for solution vector

16. Deallocate memory for sensitivity vectors

Upon completion of the integration, deallocate memory for the vectors `yS0`:

```
[S] N_VDestroyVectorArray_Serial(yS0, Ns);
```

```
[P] N_VDestroyVectorArray_Parallel(yS0, Ns);
```

If `yS` was created from `realtype` arrays `yS_i`, it is the user’s responsibility to also free the space for the arrays `yS0_i`.

17. Free user data structure

18. Free solver memory

19. Free vector specification memory

6.2 User-callable routines for forward sensitivity analysis

This section describes the CVODES functions, additional to those presented in §5.5, that are called by the user to setup and solve a forward sensitivity problem.

6.2.1 Forward sensitivity initialization and deallocation functions

Activation of forward sensitivity computation is done by calling `CVodeSensMalloc`. The form of the call to this routine is as follows:

CVodeSensMalloc

Call `flag = CVodeSensMalloc(cvode_mem, Ns, ism, yS0);`

Description The routine `CVodeSensMalloc` activates forward sensitivity computations and allocates internal memory related to sensitivity calculations.

Arguments `cvode_mem` (void *) pointer to the CVODES memory block returned by `CVodeCreate`.
`Ns` (int) the number of sensitivities to be computed.
`ism` (int) a flag used to select the sensitivity solution method and can be `CV_SIMULTANEOUS`, `CV_STAGGERED`, or `CV_STAGGERED1`:

- In the `CV_SIMULTANEOUS` approach, the state and sensitivity variables are corrected at the same time. If `CV_NEWTON` was selected as the non-linear system solution method, this amounts to performing a modified Newton iteration on the combined nonlinear system;
- In the `CV_STAGGERED` approach, the correction step for the sensitivity variables takes place at the same time for all sensitivity equations, but only after the correction of the state variables has converged and the state variables have passed the local error test;
- In the `CV_STAGGERED1` approach, all corrections are done sequentially, first for the state variables and then for the sensitivity variables, one parameter at a time. If the sensitivity variables are not included in the error control, this approach is equivalent to `CV_STAGGERED`. Note that the `CV_STAGGERED1` approach can be used only if the user-provided sensitivity right-hand side function is of type `CVSensRhs1Fn` (see §6.3).

`yS0` (`N_Vector *`) a pointer to an array of `Ns` vectors containing the initial values of the sensitivities.

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

`CV_SUCCESS` The call to `CVodeSensMalloc` was successful.

`CV_MEM_NULL` The CVODES memory block was not initialized through a previous call to `CVodeCreate`.

`CV_MEM_FAIL` A memory allocation request has failed.

`CV_ILL_INPUT` An input argument to `CVodeSensMalloc` has an illegal value.

Notes If an error occurred, `CVodeSensMalloc` 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 by:

- Base value: $\text{lenrw} = \text{lenrw} + (\text{maxord}+5)N_sN$
- With `itolS = CV_SV` (see `CVodeSetSensTolerances`): $\text{lenrw} = \text{lenrw} + N_sN$

the size of the integer workspace is increased by:

- Base value: $\text{leniw} = \text{leniw} + (\text{maxord}+5)N_sN$
- With `itolS = CV_SV`: $\text{leniw} = \text{leniw} + N_sN$

The routine `CVodeSensReInit`, useful during the solution of a sequence of problems of same size, reinitializes the sensitivity-related internal memory and must follow a call to `CVodeSensMalloc` (and maybe a call to `CVodeReInit`). The number `Ns` of sensitivities is assumed to be unchanged since the call to `CVodeSensMalloc`. The call to the `CVodeSensReInit` function has the form:

CVodeSensReInit

Call `flag = CVodeSensReInit(cvode_mem, ism, yS0);`

Description The routine `CVodeSensReInit` reinitializes forward sensitivity computations.

Arguments

- `cvode_mem` (`void *`) pointer to the CVODES memory block returned by `CVodeCreate`.
- `ism` (`int`) a flag used to select the sensitivity solution method and can be `CV_SIMULTANEOUS`, `CV_STAGGERED`, or `CV_STAGGERED1`.
- `yS0` (`N_Vector *`) a pointer to an array of `Ns` variables of type `N_Vector` containing the initial values of the sensitivities.

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

- `CV_SUCCESS` The call to `CVodeReInit` was successful.
- `CV_MEM_NULL` The CVODES memory block was not initialized through a previous call to `CVodeCreate`.
- `CV_NO_SENS` Memory space for sensitivity integration was not allocated through a previous call to `CVodeSensMalloc`.
- `CV_ILL_INPUT` An input argument to `CVodeSensReInit` has an illegal value.
- `CV_MEM_FAIL` A memory allocation request has failed.

Notes All arguments of `CVodeSensReInit` are the same as those of `CVodeSensMalloc`.
If an error occurred, `CVodeSensReInit` 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 `CVodeSensMalloc`), the user must call

CVodeSensFree

Call `CVodeSensFree(cvode_mem);`

Description The function `CVodeSensFree` frees the memory allocated for forward sensitivity computations by a previous call to `CVodeSensMalloc`.

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

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

Notes After a call to `CVodeSensFree`, forward sensitivity computations can be reactivated only by calling again `CVodeSensMalloc`.

To activate and deactivate forward sensitivity calculations for successive CVODES runs, without having to allocate and deallocate memory, the following function is provided:

CVodeSensToggleOff

Call `CVodeSensToggleOff(cvode_mem);`

Description The function `CVodeSensToggleOff` deactivates forward sensitivity calculations. It does *not* deallocate sensitivity-related memory.

Arguments `cvode_mem` (`void *`) pointer to the memory previously allocated by `CVodeMalloc`.

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

- `CV_SUCCESS` `CVodeSensToggleOff` was successful.
- `CV_MEM_NULL` `cvode_mem` was `NULL`.

Notes Since sensitivity-related memory is not deallocated, sensitivities can be reactivated at a later time (using `CVodeSensReInit`).

6.2.2 CVODE solver function

Even if forward sensitivity analysis was enabled, the call to the main solver function `CVode` is exactly the same as in §5.5.4. However, in this case the return value `flag` can also be one of the following:

<code>CV_SRHSFUNC_FAIL</code>	The sensitivity right-hand side function failed in an unrecoverable manner.
<code>CV_FIRST_SRHSFUNC_FAIL</code>	The sensitivity right-hand side function failed at the first call.
<code>CV_REPTD_SRHSFUNC_ERR</code>	Convergence tests occurred too many times due to repeated recoverable errors in the sensitivity right-hand side function. The <code>CV_REPTD_SRHSFUNC_ERR</code> will also be returned if the sensitivity right-hand side function had repeated recoverable errors during the estimation of an initial step size.
<code>CV_UNREC_SRHSFUNC_ERR</code>	The sensitivity right-hand function had a recoverable error, but no recovery was possible. This failure mode is rare, as it can occur only if the sensitivity right-hand side function fails recoverably after an error test failed while at order one.

6.2.3 Forward sensitivity extraction functions

If forward sensitivity computations have been initialized by a call to `CVodeSensMalloc`, or reinitialized by a call to `CVSensReInit`, then CVODES computes both a solution and sensitivities at time `t`. However, `CVode` will still return only the solution `y` in `y`. Solution sensitivities can be obtained through one of the following functions:

`CVodeGetSens`

Call	<code>flag = CVodeGetSens(cvode_mem, t, yS);</code>
Description	The function <code>CVodeGetSens</code> returns the sensitivity solution vectors after a successful return from <code>CVode</code> .
Arguments	<p><code>cvode_mem</code> (<code>void *</code>) pointer to the memory previously allocated by <code>CVodeMalloc</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 CVODES.</p> <p><code>yS</code> (<code>N_Vector *</code>) the computed forward sensitivity vectors.</p>
Return value	<p>The return value <code>flag</code> of <code>CVodeGetSens</code> is one of:</p> <p><code>CV_SUCCESS</code> <code>CVodeGetSens</code> was successful.</p> <p><code>CV_MEM_NULL</code> <code>cvode_mem</code> was <code>NULL</code>.</p> <p><code>CV_NO_SENS</code> Forward sensitivity analysis was not initialized.</p> <p><code>CV_BAD_DKY</code> <code>yQ</code> is <code>NULL</code>.</p> <p><code>CV_BAD_T</code> The time <code>t</code> is not in the allowed range.</p>
Notes	In case of an error return, an error message is also printed.

The function `CVodeGetSensDky` computes the `k`-th derivatives of the interpolating polynomials for the sensitivity variables at time `t`. This function is called by `CVodeGetSens` with `k = 0`, but may also be called directly by the user.

`CVodeGetSensDky`

Call	<code>flag = CVodeGetSensDky(cvode_mem, t, k, dkyS);</code>
Description	The function <code>CVodeGetSensDky</code> returns derivatives of the sensitivity solution vectors after a successful return from <code>CVode</code> .
Arguments	<code>cvode_mem</code> (<code>void *</code>) pointer to the memory previously allocated by <code>CVodeMalloc</code> .

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 CVODES.

k (**int**) order of derivatives.

dkyS (**N_Vector ***) the vectors containing the derivatives. The space for **dkyS** must be allocated by the user.

Return value The return value **flag** of **CVodeGetSensDky** is one of:

CV_SUCCESS **CVodeGetSensDky** succeeded.

CV_MEM_NULL The pointer to **cvode_mem** was NULL.

CV_NO_SENS Forward sensitivity analysis was not initialized.

CV_BAD_DKY One of the vectors **dkyS** is NULL.

CV_BAD_K **k** is not in the range $0, 1, \dots, q_u$.

CV_BAD_T The time **t** is not in the allowed range.

Notes In case of an error return, an error message is also printed.

Forward sensitivity solution vectors can also be extracted separately for each parameter in turn through the functions **CVodeGetSens1** and **CVodeGetSensDky1**, defined as follows:

CVodeGetSens1

Call **flag** = **CVodeGetSens1**(**cvode_mem**, **t**, **is**, **yS**);

Description The function **CVodeGetSens1** returns the **is**-th sensitivity solution vector after a successful return from **CVode**.

Arguments **cvode_mem** (**void ***) pointer to the memory previously allocated by **CVodeMalloc**.

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 CVODES.

is (**int**) specifies which sensitivity vector is to be returned ($0 \leq \text{is} < N_s$).

yS (**N_Vector**) the computed forward sensitivity vector.

Return value The return value **flag** of **CVodeGetSens1** is one of:

CV_SUCCESS **CVodeGetSens1** was successful.

CV_MEM_NULL **cvode_mem** was NULL.

CV_NO_SENS Forward sensitivity analysis was not initialized.

CV_BAD_IS The index **is** is not in the allowed range.

CV_BAD_DKY **yQ** is NULL.

CV_BAD_T The time **t** is not in the allowed range.

Notes In case of an error return, an error message is also printed.

CVodeGetSensDky1

Call **flag** = **CVodeGetSensDky1**(**cvode_mem**, **t**, **k**, **is**, **dkyS**);

Description The function **CVodeGetSensDky1** returns the **k**-th derivative of the **is**-th sensitivity solution vector after a successful return from **CVode**.

Arguments **cvode_mem** (**void ***) pointer to the memory previously allocated by **CVodeMalloc**.

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 CVODES.

k (**int**) order of derivative.

is (**int**) specifies the sensitivity derivative vector to be returned ($0 \leq \text{is} < N_s$).

Table 6.1: Forward sensitivity optional inputs

Optional input	Routine name	Default
Sensitivity right-hand side fct. and data	<code>CVodeSetSensRhsFn</code>	internal DQ
Sensitivity right-hand side fct. and data	<code>CVodeSetSensRhs1Fn</code>	internal DQ
Sensitivity scaling factors	<code>CVodeSetSensPbar</code>	NULL
DQ approximation method	<code>CVodeSetSensDQMethod</code>	0.0
Error control strategy	<code>CVodeSetSensErrCon</code>	FALSE
Sensitivity integration tolerances	<code>CVodeSetSensTolerances</code>	estimated
Maximum no. of nonlinear iterations	<code>CVodeSetSensMaxNonlinIters</code>	3

`dkyS` (`N_Vector`) the vector containing the derivative. The space for `dkyS` must be allocated by the user.

Return value The return value `flag` of `CVodeGetSensDky1` is one of:

`CV_SUCCESS` `CVodeGetQuadDky1` succeeded.
`CV_MEM_NULL` The pointer to `cvode_mem` was NULL.
`CV_NO_SENS` Forward sensitivity analysis was not initialized.
`CV_BAD_DKY` One of the vectors `dkyS` is NULL.
`CV_BAD_IS` The index `is` is not in the allowed range.
`CV_BAD_K` `k` is not in the range $0, 1, \dots, q_u$.
`CV_BAD_T` The time `t` is not in the allowed range.

Notes In case of an error return, an error message is also printed.

6.2.4 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 `CVodeSetSens*` functions. Table 6.1 lists all forward sensitivity optional input functions in CVODES which are described in detail in the remainder of this section.

`CVodeSetSensRhsFn`

Call `flag = CVodeSetSensRhsFn(cvode_mem, fS, data);`

Description The function `CVodeSetSensRhsFn` specifies the user-supplied C function used to evaluate the sensitivity right-hand sides (for all parameters at once) and the user data block for use by the user-supplied sensitivity right-hand side function.

Arguments `cvode_mem` (`void *`) pointer to the CVODES memory block.
`fS` (`CVSensRhsFn`) user-defined sensitivity right-hand side function.
`fS_data` (`void *`) pointer to the user data.

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

`CV_SUCCESS` The optional value has been successfully set.
`CV_MEM_NULL` The `cvode_mem` pointer is NULL.



Notes This type of function is not compatible with the `CV_STAGGERED1` sensitivity solution method (argument `ism` to `CVodeSensMalloc`). The compatibility test is performed at the first step in `CVode`.

Passing `fS=NULL` indicates using the default internal difference quotient sensitivity right-hand side routine.

The pointer `fS_data` can be the same as the pointer `f_data`, specified in a prior call to `CVodeSetFdata` (see §5.5.5.1) and passed to the user's right-hand side function `f`.

CVodeSetSensRhs1Fn

Call	<code>flag = CVodeSetSensRhs1Fn(cvode_mem, fS, data);</code>
Description	The function <code>CVodeSetSensRhs1Fn</code> specifies the user-supplied C function used to evaluate the sensitivity right-hand sides (one parameter at a time) and the user data block for use by the user-supplied sensitivity right-hand side function.
Arguments	<p><code>cvode_mem</code> (void *) pointer to the CVODES memory block.</p> <p><code>fS</code> (CVSensRhs1Fn) user-defined sensitivity right-hand side function.</p> <p><code>fS_data</code> (void *) pointer to the user data.</p>
Return value	<p>The return value <code>flag</code> (of type <code>int</code>) is one of:</p> <p><code>CV_SUCCESS</code> The optional value has been successfully set.</p> <p><code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL.</p>
Notes	<p>This type of sensitivity right-hand side function <i>must</i> be used when the <code>CV_STAGGERED1</code> sensitivity solution method is selected through <code>CVodeSensMalloc</code>, but can also be used with the other two choices (<code>CV_SIMULTANEOUS</code> and <code>CV_STAGGERED</code>).</p> <p>Passing <code>fS=NULL</code> indicates using the default internal difference quotient sensitivity right-hand side routine.</p> <p>The pointer <code>fS_data</code> can be the same as the pointer <code>f_data</code>, specified in a prior call to <code>CVodeSetFdata</code> (see §5.5.5.1) and passed to the user's right-hand side function <code>f</code>.</p>

CVodeSetSensParams

Call	<code>flag = CVodeSetSensParams(cvode_mem, p, pbar, plist);</code>
Description	The function <code>CVodeSetSensParams</code> specifies problem parameter information for sensitivity calculations.
Arguments	<p><code>cvode_mem</code> (void *) pointer to the CVODES memory block.</p> <p><code>p</code> (realtype *) a pointer to the array of real problem parameters used to evaluate $f(t, y, p)$. If non-NULL, <code>p</code> must point to a field in the user's data structure <code>f_data</code> passed to the right-hand side function. (See §6.1).</p> <p><code>pbar</code> (realtype *) an array of <code>Ns</code> positive scaling factors. If non-NULL, <code>pbar</code> must have all its components > 0.0. (See §6.1).</p> <p><code>plist</code> (int *) an array of <code>Ns</code> non-negative flags to specify which parameters to use in estimating the sensitivity equations. If non-NULL, <code>plist</code> must have all components ≥ 0. (See §6.1).</p>
Return value	<p>The return value <code>flag</code> (of type <code>int</code>) is one of:</p> <p><code>CV_SUCCESS</code> The optional value has been successfully set.</p> <p><code>CV_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL.</p> <p><code>CV_NO_SENS</code> Forward sensitivity analysis was not initialized.</p> <p><code>CV_ILL_INPUT</code> An argument has an illegal value.</p>
Notes	This function must be preceded by a call to <code>CVodeSensMalloc</code> .

**CVodeSetSensDQMethod**

Call	<code>flag = CVodeSetSensDQMethod(cvode_mem, DQtype, DQrhomax);</code>
Description	The function <code>CVodeSetSensDQMethod</code> specifies the difference quotient strategy in the case in which the right-hand side of the sensitivity equations are to be computed by CVODES.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODES memory block.

DQtype	(int) specifies the difference quotient type and can be one of CV_CENTERED or CV_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 right-hand side.
Return value	The return value flag (of type int) is one of: CV_SUCCESS The optional value has been successfully set. CV_MEM_NULL The cvode_mem pointer is NULL . CV_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 a factor of DQrhomax and the separate approximation is used otherwise. Note that a value DQrhomax < 1.0 will effectively disable switching. See §3.2 for more details. The default value are DQtype = CV_CENTERED and DQrhomax = 0.0.

CVodeSetSensErrCon

Call	flag = CVodeSetSensErrCon (cvode_mem , errconS);
Description	The function CVodeSetSensErrCon specifies the error control strategy for sensitivity variables.
Arguments	cvode_mem (void *) pointer to the CVODES memory block. errconS (boolean) specifies whether sensitivity variables are included (TRUE) or not (FALSE) in the error control mechanism.
Return value	The return value flag (of type int) is one of: CV_SUCCESS The optional value has been successfully set. CV_MEM_NULL The cvode_mem pointer is NULL .
Notes	By default, errconS is set to FALSE . If errconS = TRUE then both state variables and sensitivity variables are included in the error tests. If errconS = FALSE then the sensitivity variables are excluded from the error tests. Note that, in any event, all variables are considered in the convergence tests.

CVodeSetSensTolerances

Call	flag = CVodeSetSensTolerances (cvode_mem , itolS , reltolS , abstolS);
Description	The function CVodeSetSensTolerances specifies the integration tolerances for sensitivity variables.
Arguments	cvode_mem (void *) pointer to the CVODES memory block. itolS (int) is one of CV_SS , CV_SV , or CV_EE , where itolS = CV_SS indicates scalar relative error tolerance and scalar absolute error tolerance, while itolS = CV_SV indicates scalar relative error tolerance and vector absolute error tolerance. If itolS = CV_EE , the arguments reltolS and abstolS are ignored and CVODES will estimate tolerances for the sensitivity variables based on the state tolerances and the scaling factors \bar{p} . reltolS (realtype) is the relative error tolerance. abstolS (void *) is a pointer to the absolute error tolerance. If itolS = CV_SS , then abstolS must be a pointer to an array of realtype variables. If itolS = CV_SV , then abstolS must be an array of Ns variables of type N_Vector . In the latter case, abstolS should be created and set in the same manner as the vectors of initial values for the sensitivity variables (see §6.1).

Table 6.2: Forward sensitivity optional outputs

Optional output	Routine name
No. of calls to sensitivity r.h.s. function	<code>CVodeGetNumSensRhsEvals</code>
No. of calls to r.h.s. function for sensitivity	<code>CVodeGetNumRhsEvalsSens</code>
No. of sensitivity local error test failures	<code>CVodeGetNumSensErrTestFails</code>
No. of calls to lin. solv. setup routine for sens.	<code>CVodeGetNumSensLinSolvSetups</code>
Error weight vector for sensitivity variables	<code>CVodeGetSensErrWeights</code>
No. of sens. nonlinear solver iterations	<code>CVodeGetNumSensNonlinSolvIters</code>
No. of sens. convergence failures	<code>CVodeGetNumSensNonlinSolvConvFails</code>
No. of staggered nonlinear solver iterations	<code>CVodeGetNumStgrSensNonlinSolvIters</code>
No. of staggered convergence failures	<code>CVodeGetNumStgrSensNonlinSolvConvFails</code>

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

`CV_SUCCESS` The optional values have been successfully set.
`CV_MEM_NULL` The `cnode_mem` pointer is `NULL`.
`CV_NO_SENS` Forward sensitivity analysis was not initialized.
`CV_ILL_INPUT` An input argument to `CVodeSetSensTolerances` has an illegal value.

Notes The default behavior is for `CVODES` to estimate appropriate integration tolerances for the sensitivity variables based on the state tolerances and the scaling factors \bar{p} . See §3.2 for details.

It is the user's responsibility to provide compatible `itolS` and `abstolS` arguments.

This function must be preceded by a call to `CVodeSensMalloc`.



`CVodeSetSensMaxNonlinIters`

Call `flag = CVodeSetSensMaxNonlinIters(cnode_mem, maxcorS);`

Description The function `CVodeSetSensMaxNonlinIters` specifies the maximum number of nonlinear solver iterations for sensitivity variables per step.

Arguments `cnode_mem` (`void *`) pointer to the `CVODES` memory block.
`maxcorS` (`int`) maximum number of nonlinear solver iterations allowed per step.

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

`CV_SUCCESS` The optional value has been successfully set.
`CV_MEM_NULL` The `cnode_mem` pointer is `NULL`.

Notes The default value is 3.

6.2.5 Optional outputs for forward sensitivity analysis

Optional output functions that return statistics and solver performance information related to forward sensitivity computations are listed in Table 6.2 and described in detail in the remainder of this section.

`CVodeGetNumSensRhsEvals`

Call `flag = CVodeGetNumSensRhsEvals(cnode_mem, &nfSevals);`

Description The function `CVodeGetNumSensRhsEvals` returns the number of calls to the sensitivity right-hand side function.

Arguments `cnode_mem` (`void *`) pointer to the `CVODES` memory block.
`nfSevals` (`long int`) number of calls to the sensitivity right-hand side function.

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

`CV_SUCCESS` The optional output value has been successfully set.
`CV_MEM_NULL` The `cvode_mem` pointer is `NULL`.
`CV_NO_SENS` Forward sensitivity analysis was not initialized.

Notes In order to accommodate any of the three possible sensitivity solution methods, the default internal finite difference quotient functions evaluate the sensitivity right-hand sides one at a time. Therefore, `nfSevals` will always be a multiple of the number of sensitivity parameters (the same as the case in which the user supplies a routine of type `CVSensRhs1Fn`).

`CVodeGetNumRhsEvalsSens`

Call `flag = CVodeGetNumRhsEvalsSens(cvode_mem, &nfevalsS);`

Description The function `CVodeGetNumRhsEvalsSens` returns the number of calls to the user's right-hand side function due to the internal finite difference approximation of the sensitivity right-hand sides.

Arguments `cvode_mem` (`void *`) pointer to the CVODES memory block.
`nfevalsS` (`long int`) number of calls to the user right-hand side function.

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

`CV_SUCCESS` The optional output value has been successfully set.
`CV_MEM_NULL` The `cvode_mem` pointer is `NULL`.
`CV_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 right-hand sides.

`CVodeGetNumSensErrTestFails`

Call `flag = CVodeGetNumSensErrTestFails(cvode_mem, &nSetfails);`

Description The function `CVodeGetNumSensErrTestFails` returns the number of local error test failures for the sensitivity variables that have occurred.

Arguments `cvode_mem` (`void *`) pointer to the CVODES memory block.
`nSetfails` (`long int`) number of error test failures.

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

`CV_SUCCESS` The optional output value has been successfully set.
`CV_MEM_NULL` The `cvode_mem` pointer is `NULL`.
`CV_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 `CVodeSetSensErrCon` in §6.2.4). Even in that case, this counter is not incremented if the `ism=CV_SIMULTANEOUS` sensitivity solution method has been used.

`CVodeGetNumSensLinSolvSetups`

Call `flag = CVodeGetNumSensLinSolvSetups(cvode_mem, &nlinsetupsS);`

Description The function `CVodeGetNumSensLinSolvSetups` returns the number of calls to the linear solver setup function due to forward sensitivity calculations.

Arguments `cvode_mem` (`void *`) pointer to the CVODES memory block.
`nlinsetupsS` (`long int`) number of calls to the linear solver setup function.

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

CV_SUCCESS The optional output value has been successfully set.

CV_MEM_NULL The `ccode_mem` pointer is NULL.

CV_NO_SENS Forward sensitivity analysis was not initialized.

Notes This counter is incremented only if Newton iteration has been used and if either the `ism=CV_STAGGERED` or the `ism=CV_STAGGERED1` sensitivity solution method has been specified in the call to `CVodeSensMalloc` (see §6.2.1).

CVodeGetSensStats

[illegible]

Description	The function <code>CVodeGetSensStats</code> returns all of the above sensitivity-related solver statistics as a group.
-------------	--

Arguments	<code>ccode_mem</code> (void *) pointer to the CVOIDES memory block.
	<code>nfSevals</code> (long int) number of calls to the sensitivity right-hand side function.
	<code>nfevalsS</code> (long int) number of calls to the user-supplied right-hand side function.
	<code>nSetfails</code> (long int) number of error test failures.
	<code>nlinsetupsS</code> (long int) number of calls to the linear solver setup function.

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

CV_SUCCESS The optional output values have been successfully set.

CV_MEM_NULL The `ccode_mem` pointer is NULL.

CV_NO_SENS Forward sensitivity analysis was not initialized.

CVodeGetSensErrWeights

Call `flag = CNodeGetSensErrWeights(cvode_mem, eSweight);`

Description	The function <code>CVodeGetSensErrWeights</code> returns the sensitivity error weights at the current time. These are the reciprocals of the W_i of (3.7) for the sensitivity variables.
-------------	--

Arguments	<code>cnode_mem</code> (void *) pointer to the CVODES memory block.
	<code>eSweight</code> (N_Vector_S) pointer to the array of error weight vectors.

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

CV_SUCCESS The optional output value has been successfully set.

CV_MEM_NULL The `ccode_mem` pointer is NULL.

CV_NO_SENS Forward sensitivity analysis was not initialized.

Notes	The user must allocate memory for <code>eweightS</code> .
-------	---

CVodeGetNumSensNonlinSolvIters

Call `flag = CVodeGetNumSensNonlinSolvIters(cvode_mem, &nSniters);`

Description	The function <code>CVodeGetNumSensNonlinSolvIters</code> returns the number of nonlinear iterations performed for sensitivity calculations.
-------------	---

Arguments `cvode_mem` (void *) pointer to the CVODES memory block.
 `nSniters` (long int) number of nonlinear iterations performed.

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

CV_SUCCESS The optional output value has been successfully set.

CV_MEM_NULL The `ccode_mem` pointer is NULL.

CV_NO_SENS Forward sensitivity analysis was not initialized.

Notes This counter is incremented only if the `ism` was `CV_STAGGERED` or `CV_STAGGERED1` in the call to `CVodeSensMalloc` (see §6.2.1).

In the `CV_STAGGERED1` case, the value of `nSniters` is the sum of the number of nonlinear iterations performed for each sensitivity equation. These individual counters can be obtained through a call to `CVodeGetNumStgrSensNonlinSolvIters` (see below).

`CVodeGetNumSensNonlinSolvConvFails`

Call `flag = CVodeGetNumSensNonlinSolvConvFails(cvode_mem, &nSncfails);`

Description The function `CVodeGetNumSensNonlinSolvConvFails` returns the number of nonlinear convergence failures that have occurred for sensitivity calculations.

Arguments `cvode_mem` (`void *`) pointer to the CVODES memory block.
`nSncfails` (`long int`) number of nonlinear convergence failures.

Return value The return value `flag` (of type `int`) is one of:
`CV_SUCCESS` The optional output value has been successfully set.
`CV_MEM_NULL` The `cvode_mem` pointer is `NULL`.
`CV_NO_SENS` Forward sensitivity analysis was not initialized.

Notes This counter is incremented only if the `ism` was `CV_STAGGERED` or `CV_STAGGERED1` in the call to `CVodeSensMalloc` (see §6.2.1).

In the `CV_STAGGERED1` case, the value of `nSncfails` is the sum of the number of nonlinear convergence failures that occurred for each sensitivity equation. These individual counters can be obtained through a call to `CVodeGetNumStgrSensNonlinConvFails` (see below).

`CVodeGetSensNonlinSolvStats`

Call `flag = CVodeGetSensNonlinSolvStats(cvode_mem, &nSniters, &nSncfails);`

Description The function `CVodeGetSensNonlinSolvStats` returns the sensitivity-related nonlinear solver statistics as a group.

Arguments `cvode_mem` (`void *`) pointer to the CVODES 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:
`CV_SUCCESS` The optional output values have been successfully set.
`CV_MEM_NULL` The `cvode_mem` pointer is `NULL`.
`CV_NO_SENS` Forward sensitivity analysis was not initialized.

`CVodeGetNumStgrSensNonlinSolvIters`

Call `flag = CVodeGetNumStgrSensNonlinSolvIters(cvode_mem, nSTGR1niters);`

Description The function `CVodeGetNumStgrSensNonlinSolvIters` returns the number of nonlinear (functional or Newton) iterations performed for each sensitivity equation separately, in the `CV_STAGGERED1` case.

Arguments `cvode_mem` (`void *`) pointer to the CVODES memory block.
`nSTGR1niters` (`long int *`) an array (of dimension `Ns`) which will be set with the number of nonlinear iterations performed for each sensitivity system individually.

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

	CV_SUCCESS	The optional output value has been successfully set.
	CV_MEM_NULL	The <code>cnode_mem</code> pointer is NULL.
	CV_NO_SENS	Forward sensitivity analysis was not initialized.
Notes		The user must allocate space for <code>nSTGR1nitters</code> .



CVodeGetNumStgrSensNonlinSolvConvFails

Call	<code>flag = CVodeGetNumStgrSensNonlinSolvConvFails(cnode_mem, nSTGR1ncfails);</code>	
Description	The function <code>CVodeGetNumStgrSensNonlinSolvConvFails</code> returns the number of non-linear convergence failures that have occurred for each sensitivity equation separately, in the CV_STAGGERED1 case.	
Arguments	<code>cnode_mem</code> (void *) pointer to the CVODES memory block. <code>nSTGR1ncfails</code> (long int *) an array (of dimension <code>Ns</code>) which will be set with the number of nonlinear convergence failures for each sensitivity system individually.	
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of:	
	CV_SUCCESS	The optional output value has been successfully set.
	CV_MEM_NULL	The <code>cnode_mem</code> pointer is NULL.
	CV_NO_SENS	Forward sensitivity analysis was not initialized.
Notes	The user must allocate space for <code>nSTGR1ncfails</code> .	



6.3 User-supplied routines for forward sensitivity analysis

In addition to the required and optional user-supplied routines described in §5.6, when using CVODES for forward sensitivity analysis, the user has the option of providing a routine that calculates the right-hand side of the sensitivity equations (3.9).

By default, CVODES uses difference quotient approximation routines for the right-hand sides of the sensitivity equations. However, CVODES allows the option for user-defined sensitivity right-hand side routines (which also provides a mechanism for interfacing CVODES to routines generated by automatic differentiation).

6.3.1 Sensitivity equations right-hand side (all at once)

If the CV_SIMULTANEOUS or CV_STAGGERED approach was selected in the call to `CVodeSensMalloc`, the user may provide the right-hand sides of the sensitivity equations (3.9), for all sensitivity parameters at once, through a function of type `CVSensRhsFn` defined by:

CVSensRhsFn

Definition	<pre>typedef int (*CVSensRhsFn)(int Ns, realtype t, N_Vector y, N_Vector ydot, N_Vector *yS, N_Vector *ySdot, void *fS_data, N_Vector tmp1, N_Vector tmp2);</pre>	
Purpose	This function computes the sensitivity right-hand side for all sensitivity equations at once. It must compute the vectors $(\partial f / \partial y) s_i(t) + (\partial f / \partial p_i)$ and store them in <code>ySdot[i]</code> .	
Arguments	<code>t</code>	is the current value of the independent variable.
	<code>y</code>	is the current value of the state vector, $y(t)$.
	<code>ydot</code>	is the current value of the right-hand side of the state equations.
	<code>yS</code>	contains the current values of the sensitivity vectors.

ySdot is the output of **CVSensRhsFn**. On exit it must contain the sensitivity right-hand side vectors.

f_data is a pointer to user data - the same as the **fS_data** parameter passed to **CVodeSetSensRhsFn**.

tmp1
tmp2 are **N_Vectors** which can be used as temporary storage.

Return value A **CVSensRhsFn** should return 0 if successful, a positive value if a recoverable error occurred (in which case CVODES will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and **CV_SRHSFUNC_FAIL** is returned).



Notes A sensitivity right-hand side function of type **CVSensRhsFn** is not compatible with the **CV_STAGGERED1** approach.



Allocation of memory for **ySdot** is handled within CVODES.

For efficiency considerations, the right-hand side function is not evaluated at the converged solution of the nonlinear solver. Therefore, a recoverable error in **CVSensRhsFn** at that point cannot be corrected (as it will occur when the right-hand side function is called the first time during the following integration step and a successful step cannot be undone).

There are two situations in which recovery is not possible even if **CVSensRhsFn** function returns a recoverable error flag. This include the situation when this occurs at the very first call to the **CVSensRhsFn** (in which case CVODES returns **CV_FIRST_SRHSFUNC_ERR**) or if a recoverable error is reported when **CVSensRhsFn** is called after an error test failure, while the linear multistep method order is equal to 1 (in which case CVODES returns **CV_UNREC_SRHSFUNC_ERR**).

6.3.2 Sensitivity equations right-hand side (one at a time)

Alternatively, the user may provide the sensitivity right-hand sides, one sensitivity parameter at a time, through a function of type **CVSensRhs1Fn**. Note that a sensitivity right-hand side function of type **CVSensRhs1Fn** is compatible with any valid value of the **CVodeSensMalloc** argument **ism**, and is required if **ism=CV_STAGGERED1**. The type **CVSensRhs1Fn** is defined by

CVSensRhs1Fn

Definition

```
typedef int (*CVSensRhs1Fn)(int Ns, realtype t,
                             N_Vector y, N_Vector ydot,
                             int iS, N_Vector yS, N_Vector ySdot,
                             void *fS_data,
                             N_Vector tmp1, N_Vector tmp2);
```

Purpose This function computes the sensitivity right-hand side for one sensitivity equation at a time. It must compute the vector $(\partial f / \partial y) s_i(t) + (\partial f / \partial p_i)$ for $i=iS$ and store it in **ySdot**.

Arguments

- t** is the current value of the independent variable.
- y** is the current value of the state vector, $y(t)$.
- ydot** is the current value of the right-hand side of the state equations.
- iS** is the index of the parameter for which the sensitivity right-hand side must be computed.
- yS** contains the current value of the **iS**-th sensitivity vector.
- ySdot** is the output of **CVSensRhs1Fn**. On exit it must contain the **iS**-th sensitivity right-hand side vector.

	<code>f_data</code> is a pointer to user data - the same as the <code>fS_data</code> parameter passed to <code>CVodeSetSensRhs1Fn</code> .
	<code>tmp1</code> <code>tmp2</code> are <code>N_Vectors</code> which can be used as temporary storage.
Return value	A <code>CVSensRhsFn</code> should return 0 if successful, a positive value if a recoverable error occurred (in which case CVODES will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and <code>CV_SRHSFUNC_FAIL</code> is returned).
Notes	Allocation of memory for <code>ySdot</code> is handled within CVODES. For efficiency considerations, the right-hand side function is not evaluated at the converged solution of the nonlinear solver. Therefore, a recoverable error in <code>CVSensRhs1Fn</code> at that point cannot be corrected (as it will occur when the right-hand side function is called the first time during the following integration step and a successful step cannot be undone). There are two situations in which recovery is not possible even if <code>CVSensRhs1Fn</code> function returns a recoverable error flag. This include the situation when this occurs at the very first call to the <code>CVSensRhs1Fn</code> (in which case CVODES returns <code>CV_FIRST_SRHSFUNC_ERR</code>) or if a recoverable error is reported when <code>CVSensRhs1Fn</code> is called after an error test failure, while the linear multistep method order is equal to 1 (in which case CVODES returns <code>CV_UNREC_SRHSFUNC_ERR</code>).



6.4 Note on using partial error control

For some problems, when sensitivities are excluded from the error control test, the behavior of CVODES 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. A comparison of the solver diagnostics reported for `cvsdex` and the second run of the `cvsfwddex` example in [18] indicates that this may not always be the case.

The short explanation of this behavior is that the step size selection implemented by the error control mechanism in CVODES is based on the magnitude of the correction calculated by the nonlinear solver. As mentioned in §6.2.1, even with partial error control selected in the call to `CVodeSensMalloc`, the sensitivity variables are included in the convergence tests of the nonlinear solver.

When using the simultaneous corrector method (§3.2), 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. After all, in this case (`ism=CV_STAGGERED` or `ism=CV_STAGGERED1` in the call to `CVodeSensMalloc`), 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, CVODES 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, CVODES may trigger a call to the linear solver's setup routine which typically involves reevaluation of Jacobian information (Jacobian approximation in the case of `CVDENSE` and `CVBAND`, or preconditioner data in the case of `CVSPGMR`). 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 ODEs 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 CVODES 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 final correction), 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 CVODES. 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 7

Using CVODES for Adjoint Sensitivity Analysis

This chapter describes the use of CVODES to compute sensitivities of derived functions using adjoint sensitivity analysis. As mentioned before, the adjoint sensitivity module of CVODES provides the infrastructure for integrating backward in time any system of ODEs that depends on the solution of the original IVP, by providing various interfaces to the main CVODES 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 ODEs that are integrated backward in time. The backward problem can be the adjoint problem (3.17) or (3.20), and can be augmented with some quadrature differential equations.

CVODES uses various constants for both input and output. These are defined as needed in this chapter, but for convenience are also listed separately in Chapter 11.

We begin with a brief overview, in the form of a skeleton user program. Following that are detailed descriptions of the interface to the various user-callable functions and of the user-supplied functions that were not already described in §5.

7.1 A skeleton of the user's main program

The following is a skeleton of the user's main program as an application of CVODES. The user program is to have these steps in the order indicated, unless otherwise noted. For the sake of brevity, we defer many of the details to the later sections. As in §5.4, 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 CVODES: steps marked with [P] correspond to NVECTOR_PARALLEL, while steps marked with [S] correspond to NVECTOR_SERIAL.

1. Include necessary header files

The `cvodes.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` or `nvector_parallel.h` for the two implementations provided with CVODES) 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 CVODES object for the forward problem**
6. **Allocate internal memory for the forward problem**
7. **Set optional inputs for the forward problem**
8. **Attach linear solver module for the forward problem**
9. **Set linear solver optional inputs for the forward problem**
10. **Allocate space for the adjoint computation**

Call `cvadj_mem = CVadjMalloc()` to allocate memory for the combined forward-backward problem (see §7.2.1 for more details). This call requires `Nd`, the number of steps between two consecutive checkpoints. `CVadjMalloc` also specifies the type of interpolation used (see §3.3.1).

11. **Integrate forward problem**

Call `CVodeF`, a wrapper for the CVODES main integration function `CVode`, either in `CV_NORMAL` mode to the time `tout` or in `CV_ONE_STEP` mode inside a loop (if intermediate solutions of the forward problem are desired (see §7.2.2)). The final value of `tret`, denoted `tfinal`, is then the maximum allowable value for the endpoint t_1 .

Backward problem

12. **Set problem dimensions for the backward problem**

[S] set NB, the number of variables in the backward problem
 [P] set NB and NBlocal

13. **Set final values for the backward problem**

Set the vector `yB0` of final values for the backward problem.

14. **Create CVODES object for the backward problem**

Call `CVodeCreateB`, a wrapper for `CVodeCreate`, to create the CVODES memory block and specify the solution method (linear multistep method and nonlinear solver iteration type) for the backward problem. Unlike `CVodeCreate`, the function `CVodeCreateB` does not return a pointer to the newly created memory block. Instead, this pointer is attached to the adjoint memory block (returned by `CVadjMalloc` and passed as the first argument to `CVodeCreateB`).

15. **Allocate memory for the backward problem**

Call `CVodeMallocB`, a wrapper for `CVodeMalloc`, to allocate internal memory and initialize CVODES at `tB0` for the backward problem (see §7.2.3).

16. **Set optional inputs for the backward problem**

Call `CVodeSet*B` functions to change from their default values any optional inputs that control the behavior of CVODES. Unlike their counterparts for the forward problem, these functions take as their first argument the adjoint memory block returned by `CVadjMalloc`.

17. **Attach linear solver module for the backward problem**

If Newton iteration is chosen, initialize the linear solver module for the backward problem by calling the appropriate wrapper function: `CVDenseB`, `CVBandB`, `CVDiagB`, `CVSpqrB`, `CVSpbcgB`, or `CVSptfqmr` (see §7.2.4). 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 `CVDENSE` linear solver and the backward problem with `CVSPGMR`.

18. **Initialize quadrature calculation**

If additional quadrature equations must be evaluated, call `CVodeQuadMallocB`, a wrapper around `CVodeQuadMalloc`, to initialize and allocate memory for quadrature integration. Optionally, call `CVodeSetQuad*B` functions to change from their default values optional inputs that control the integration of quadratures during the backward phase.

19. Integrate backward problem

Call `CVodeB`, a second wrapper around the CVODES main integration function `CVode`, to integrate the backward problem from `tB0` (see §7.2.5). This function can be called either in `CV_NORMAL` or `CV_ONE_STEP` mode. Typically, `CVodeB` will be called in `CV_NORMAL` mode with an end time equal to the initial time of the forward problem.

20. Extract quadrature variables

If applicable, call `CVodeGetQuadB`, a wrapper around `CVodeGetQuad`, to extract the values of the quadrature variables at the time returned by the last call to `CVodeB`.

21. 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 `CVodeFree` to free the CVODES memory block for the forward problem, and a call to `CVadjFree` (see §7.2.1) to free the memory allocated for the combined problem. Note that `CVadjFree` also deallocates the CVODES memory for the backward problem.

22. 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 CVODES was motivated by the desire to keep it as close as possible in look and feel to the one for ODE IVP integration. Note that if steps (12)-(20) are not present, a program with the above structure will have the same functionality as one described in §5.4 for integration of ODEs, albeit with some overhead due to the checkpointing scheme.

7.2 User-callable functions for adjoint sensitivity analysis

7.2.1 Adjoint sensitivity allocation and deallocation functions

After the setup phase for the forward problem, but before the call to `CVodeF`, memory for the combined forward-backward problem must be allocated by a call to the function `CVadjMalloc`. The form of the call to this function is

`CVadjMalloc`

Call	<code>cvadj_mem = CVadjMalloc(cvode_mem, Nd, interpType);</code>		
Description	The function <code>CVadjMalloc</code> allocates internal memory for the combined forward and backward integration, other than the CVODES memory block. Space is allocated for the N_d interpolation data points and a linked list of checkpoints is initialized.		
Arguments	<code>cvode_mem</code>	(<code>void *</code>) is the CVODES memory block for the forward problem returned by a previous call to <code>CVodeCreate</code> .	
	<code>Nd</code>	(<code>long int</code>) is the number of integration steps between two consecutive checkpoints.	
	<code>interpType</code>	(<code>int</code>) specifies the type of interpolation used and can be <code>CV_POLYNOMIAL</code> or <code>CV_HERMITE</code> , indicating variable-degree polynomial and cubic Hermite interpolation, respectively (see §3.3.1).	

Return value If successful, `CVadjMalloc` returns a pointer of type `void *`. The user does not need to access this memory block but must pass it to other adjoint module user-callable functions. In case of failure (`cvode_mem` is `NULL`, an input has an illegal value, or a memory request fails), `CVadjMalloc` prints an error message to the standard output stream `stderr` and returns `NULL`.

Notes The user must set `Nd` so that all data needed for interpolation of the forward problem solution between two checkpoints fits in memory. `CVadjMalloc` attempts to allocate space for $(2Nd+3)$ variables of type `N_Vector`.

CVadjFree

Call `CVadjFree(&cvadj_mem);`

Description The function `CVadjFree` frees the memory allocated by a previous call to `CVadjMalloc`.

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

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

Notes This function frees all memory allocated by `CVadjMalloc`. This includes workspace memory, the linked list of checkpoints, memory for the interpolation data, as well as the CVODES memory for the backward integration phase.

If the user wishes to change the interpolation method after the call to `CVadjMalloc`, this can be done through a call to the function `CVadjSetInterpType`. The form of the call to this function is

CVadjSetInterpType

Call `flag = CVadjSetInterpType(cvadj_mem, interpType);`

Description The function `CVadjSetInterpType` resets the interpolation method to `interpType`.

Arguments `cvadj_mem` (`void *`) pointer to the adjoint memory block.
`interpType` (`int`) specifies the type of interpolation used and can be `CV_POLYNOMIAL` or `CV_HERMITE`, indicating variable-degree polynomial and cubic Hermite interpolation, respectively (see §3.3.1).

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

`CV_SUCCESS` The call to `CVadjSetInterpType` was successful.
`CV_ADJMEM_NULL` The `cvadj_mem` argument was `NULL`.
`CV_ILL_INPUT` The input argument `interpType` has an illegal value.
`CV_MEM_FAIL` An error occurred while trying to allocate memory for the interpolation module.

Notes The function `CVadjSetInterpType` performs memory deallocation and allocation.

7.2.2 Forward integration function

The function `CVodeF` is very similar to the CVODES function `CVode` (see §5.5.4) in that it integrates the solution of the forward problem and returns the solution in `y`. At the same time, however, `CVodeF` stores checkpoint data every `Nd` integration steps. `CVodeF` can be called repeatedly by the user. The call to this function has the form

CVodeF

Call `flag = CVodeF(cvadj_mem, tout, yout, tret, itask, ncheck);`

Description The function `CVodeF` integrates the forward problem over an interval in t and saves checkpointing data.

Arguments `cvadj_mem` (void *) pointer to the adjoint memory block.
 `tout` (realtype) the next time at which a computed solution is desired.
 `yout` (N_Vector) the computed solution vector.
 `tret` (realtype *) the time reached by the solver.
 `itask` (int) a flag indicating the job of the solver for the next step. The `CV_NORMAL` task is to have the solver take internal steps until it has reached or just passed the user-specified `tout` parameter. The solver then interpolates in order to return an approximate value of $y(tout)$. The `CV_ONE_STEP` option tells the solver to just take one internal step and return the solution at the point reached by that step. The `CV_NORMAL_TSTOP` and `CV_ONE_STEP_TSTOP` modes are similar to `CV_NORMAL` and `CV_ONE_STEP`, respectively, except that the integration never proceeds past the value `tstop` (specified through the function `CVodeSetStopTime`).

Return value On return, `CVodeF` returns a vector `yout` and a corresponding independent variable value $t = *tret$, such that `yout` is the computed value of $y(t)$. Additionally, it returns in `ncheck` the number of checkpoints saved. The return value `flag` (of type `int`) will be one of the following. For more details see §5.5.4.

`CV_SUCCESS` `CVodeF` succeeded.
`CV_TSTOP_RETURN` `CVodeF` succeeded by reaching the optional stopping point.
`CV_NO_MALLOC` The function `CVodeMalloc` has not been previously called.
`CV_ILL_INPUT` One of the inputs to `CVodeF` is illegal.
`CV_TOO_MUCH_WORK` The solver took `mxstep` internal steps but could not reach `tout`.
`CV_TOO_MUCH_ACC` The solver could not satisfy the accuracy demanded by the user for some internal step.
`CV_ERR_FAILURE` Error test failures occurred too many times during one internal time step or occurred with $|h| = h_{min}$.
`CV_CONV_FAILURE` Convergence test failures occurred too many times during one internal time step or occurred with $|h| = h_{min}$.
`CV_LSETUP_FAIL` The linear solver's setup function failed in an unrecoverable manner.
`CV_LSOLVE_FAIL` The linear solver's solve function failed in an unrecoverable manner.
`CV_ADJMEM_NULL` The `cvadj_mem` argument was `NULL`.
`CV_MEM_FAIL` A memory allocation request has failed (in an attempt to allocate space for a new checkpoint).

Notes All failure return values are negative and therefore a test `flag < 0` will trap all `CVodeF` failures.

At this time, `CVodeF` 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 `CVODES` internal memory block and contains enough information to restart the integration from that time and to proceed with the same step size and method order sequence as during the forward integration.

In addition, `CVodeF` 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 check points were necessary, there is no need for the second forward integration phase.

Although it is legal to define a value for `tstop` and then call `CVodeF` in `CV_NORMAL_TSTOP` or `CV_ONE_STEP_TSTOP` modes, after a return with `flag = TSTOP_RETURN`, the integration should not be continued (no `tstop` information is stored at checkpoints).

It is illegal to change the integration tolerances between consecutive calls to `CVodeF`, as this information is not captured in the checkpoints data.



7.2.3 Backward problem initialization functions

The functions `CVodeCreateB` and `CVodeMallocB` must be called in the order listed. They instantiate a CVODES solver object, provide problem and solution specifications, and allocate internal memory for the backward problem.

CVodeCreateB

Call	<code>flag = CVodeCreateB(cvadj_mem, lmm, iter);</code>
Description	The function <code>CVodeCreateB</code> instantiates a CVODES solver object and specifies the solution method for the backward problem.
Arguments	<p><code>cvadj_mem</code> (void *) pointer to the adjoint memory block returned by <code>CVadjMalloc</code>.</p> <p><code>lmm</code> (int) specifies the linear multistep method and must be one of two possible values: <code>CV_ADAMS</code> or <code>CV_BDF</code>.</p> <p><code>iter</code> (int) specifies the type of nonlinear solver iteration and must be either <code>CV_NEWTON</code> or <code>CV_FUNCTIONAL</code>.</p>
Return value	<p>If successful, <code>CVodeCreateB</code> stores a pointer to the newly created CVODES memory block (of type void *) for the backward problem. The return <code>flag</code> (of type int) is one of:</p> <p><code>CV_SUCCESS</code> The call to <code>CVodeCreateB</code> was successful.</p> <p><code>CV_ADJMEM_NULL</code> The <code>cvadj_mem</code> argument was NULL.</p> <p><code>CV_MEM_FAIL</code> An error occurred while trying to create the CVODES memory block for the backward problem.</p>

The function `CVodeMallocB` is essentially a call to `CVodeMalloc` with some particularization for backward integration as described below.

CVodeMallocB

Call	<code>flag = CVodeMallocB(cvadj_mem, fB, tB0, yB0, itolB, reltolB, abstolB);</code>
Description	The function <code>CVodeMallocB</code> provides required problem and solution specifications, allocates internal memory, and initializes CVODES for the backward problem.
Arguments	<p><code>cvadj_mem</code> (void *) pointer to the adjoint memory block returned by <code>CVadjMalloc</code>.</p> <p><code>fB</code> (<code>CVRhsFnB</code>) is the C function which computes fB, the right-hand side of the backward ODE problem. This function has the form <code>fB(t, y, yB, yBdot, f_dataB)</code> (for full details see §7.3).</p> <p><code>tB0</code> (<code>realtype</code>) specifies the endpoint where final conditions are provided for the backward problem.</p> <p><code>yB0</code> (<code>N_Vector</code>) is the final value of the backward problem.</p> <p><code>itolB</code> (int) is one of <code>CV_SS</code> or <code>CV_SV</code> where <code>itol=CV_SS</code> indicates scalar relative error tolerance and scalar absolute error tolerance, while <code>itol=CV_SV</code> indicates scalar relative error tolerance and vector absolute error tolerance.</p> <p><code>reltolB</code> (<code>realtype</code>) is the relative error tolerance.</p> <p><code>abstolB</code> (void *) is a pointer to the absolute error tolerance. If <code>itolB = CV_SS</code>, <code>abstolB</code> must be a pointer to a <code>realtype</code> variable. If <code>itolB = CV_SV</code>, <code>abstolB</code> must be an <code>N_Vector</code> variable.</p>
Return value	<p>The return <code>flag</code> (of type int) will be one of the following:</p> <p><code>CV_SUCCESS</code> The call to <code>CVodeMallocB</code> was successful.</p> <p><code>CV_ADJMEM_NULL</code> The <code>cvadj_mem</code> argument was NULL.</p> <p><code>CV_BAD_TB0</code> The final time <code>tB0</code> is outside the interval over which the forward problem was solved.</p>

CV_MEM_NULL	The CVODES memory block for the backward problem was not initialized through a previous call to <code>CVodeCreateB</code> .
CV_MEM_FAIL	A memory allocation request has failed.
CV_ILL_INPUT	An input argument to <code>CVodeMallocB</code> has an illegal value.

Notes It is the user's responsibility to provide compatible `itolB` and `abstolB` arguments.

The memory allocated by `CVodeMallocB` is deallocated by the function `CVadjFree`.

For the case when it is needed to solve several different backward problems corresponding to the same original problem, CVODES provides a mechanism to reuse the existing checkpoints. The function `CVodeReInitB` reinitializes the CVODES memory block for the backward problem, where a prior call to `CVodeMallocB` has been made with the same problem size `NB`. `CVodeReInitB` performs the same input checking and initializations that `CVodeMallocB` does, but does no memory allocation, assuming that the existing internal memory is sufficient for the new problem. Note that `CVReInitB` is essentially a wrapper for `CVodeReInit` and so all details given for `CVodeReInit` in §5.5.8 apply. The call to the `CVodeReInitB` function has the form

`CVodeReInitB`

Call	<code>flag = CVodeReInitB(cvadj_mem, fB, tB0, yB0, itolB, reltolB, abstolB);</code>
Description	The function <code>CVodeReInitB</code> provides required problem specifications and reinitializes CVODES for the backward problem.
Arguments	<p><code>cvadj_mem</code> (void *) pointer to the adjoint memory block returned by <code>CVadjMalloc</code>.</p> <p><code>fB</code> (<code>CVRhsFnB</code>) is the C function which computes fB, the right-hand side of the backward ODE problem. This function has the form <code>fB(t, y, yB, yBdot, f_dataB)</code> (for full details see §7.3).</p> <p><code>tB0</code> (<code>realtype</code>) specifies the endpoint where final conditions are provided for the backward problem.</p> <p><code>yB0</code> (<code>N_Vector</code>) is the final value of the backward problem.</p> <p><code>itolB</code> (<code>int</code>) is either <code>CV_SS</code> or <code>CV_SV</code>, where <code>itol=CV_SS</code> indicates scalar relative error tolerance and scalar absolute error tolerance, while <code>itol=CV_SV</code> indicates scalar relative error tolerance and vector absolute error tolerance.</p> <p><code>reltolB</code> (<code>realtype</code>) is the relative error tolerance.</p> <p><code>abstolB</code> (void *) is a pointer to the absolute error tolerance. If <code>itolB = CV_SS</code>, <code>abstolB</code> must be a pointer to a <code>realtype</code> variable. If <code>itolB = CV_SV</code>, <code>abstolB</code> must be an <code>N_Vector</code> variable.</p>
Return value	The return value <code>flag</code> (of type <code>int</code>) will be one of the following: <p>CV_SUCCESS The call to <code>CVodeReInitB</code> was successful.</p> <p>CV_ADJMEM_NULL The <code>cvadj_mem</code> argument was NULL.</p> <p>CV_BAD_TB0 The final time <code>tB0</code> is outside the interval over which the forward problem was solved.</p> <p>CV_MEM_NULL The CVODES memory block for the backward problem was not initialized through a previous call to <code>CVodeCreateB</code>.</p> <p>CV_NO_MALLOC Memory space for the CVODES memory block for the backward problem was not allocated through a previous call to <code>CVodeMallocB</code>.</p> <p>CV_ILL_INPUT An input argument to <code>CVodeReInitB</code> has an illegal value.</p>
Notes	It is the user's responsibility to provide compatible <code>itolB</code> and <code>abstolB</code> arguments.

7.2.4 Linear solver initialization functions for backward problem

All linear solver modules in CVODES provide additional specification functions for the case in which Newton iteration is selected for the solution of the backward problem. The initialization functions

described in §5.5.3 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 §7.3).

The following six 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 §5.5.3 with the exception of their first argument which must be the pointer to the adjoint memory block returned by `CVadjMalloc`.

```
flag = CVDenseB(cvadj_mem, nB);
flag = CVDiagB(cvadj_mem);
flag = CVBandB(cvadj_mem, nB, mupperB, mlowerB);
flag = CVSpgrB(cvadj_mem, pretypeB, maxlB);
flag = CVSpbcgB(cvadj_mem, pretypeB, maxlB);
flag = CVSptfqmrB(cvadj_mem, pretypeB, maxlB);
flag = CVSpbcgB(cvadj_mem, pretypeB, maxlB);
flag = CVSptfqmrB(cvadj_mem, pretypeB, maxlB);
```

Their return value `flag` (of type `int`) can have any of the return values of their counterparts. If the `cvadj_mem` argument was `NULL`, `flag` will be `CVDENSE_ADMEM_NULL`, `CVDIAG_ADMEM_NULL`, `CVBAND_ADMEM_NULL`, or `CVSPILS_ADMEM_NULL`.

7.2.5 Backward integration function

The function `CVodeB` performs the integration of the backward problem. It is essentially a wrapper for the CVODES main integration function `CVode` and, in the case in which checkpoints were needed, it evolves the solution of the backward problem through a sequence of forward-backward integrations between consecutive checkpoints. 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 call to this function has the form

CVodeB

Call `flag = CVodeB(cvadj_mem, tBout, yBout, tBret, itaskB);`

Description The function `CVodeB` integrates the backward ODE problem over an interval in t .

Arguments `cvadj_mem` (`void *`) pointer to the adjoint memory block returned by `CVadjMalloc`.

`tBout` (`realtype`) the next time at which a computed solution is desired.

`yBout` (`N_Vector`) the computed solution vector of the backward problem.

`tBret` (`realtype *`) the time reached by the solver.

`itaskB` (`int`) a flag indicating the job of the solver for the next step. The `CV_NORMAL` task is to have the solver take internal steps until it has reached or just passed the user specified `tBout` parameter. The solver then interpolates in order to return an approximate value of $yB(tBout)$. The `CV_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 On return, `CVodeB` returns a vector `yBout` and a corresponding independent variable value $t = *tBret$, such that `yBout` is the computed value of the solution of the backward problem.

The return value `flag` (of type `int`) will be one of the following. For more details see §5.5.4.

`CV_SUCCESS` `CVodeB` succeeded.

`CV_NO_MALLOC` The CVODES memory for the backward problem was `NULL`.

`CV_ILL_INPUT` One of the inputs to `CVode` is illegal.

`CV_BAD_ITASK` The `itaskB` argument has an illegal value.

CV_TOO_MUCH_WORK	The solver took <code>mxstep</code> internal steps but could not reach <code>tBout</code> .
CV_TOO_MUCH_ACC	The solver could not satisfy the accuracy demanded by the user for some internal step.
CV_ERR_FAILURE	Error test failures occurred too many times during one internal time step.
CV_CONV_FAILURE	Convergence test failures occurred too many times during one internal time step.
CV_LSETUP_FAIL	The linear solver's setup function failed in an unrecoverable manner.
CV_SOLVE_FAIL	The linear solver's solve function failed in an unrecoverable manner.
CV_ADJMEM_NULL	The <code>cvadj_mem</code> argument was <code>NULL</code> .
CV_BCKMEM_NULL	The <code>cvodes</code> memory for the backward problem was not created through a call to <code>CVodeCreateB</code> .
CV_BAD_TBOUT	The desired output time <code>tBout</code> is outside the interval over which the forward problem was solved.
CV_REIFWD_FAIL	Reinitialization of the forward problem failed at the first checkpoint (corresponding to the initial time of the forward problem).
CV_FWD_FAIL	An error occurred during the integration of the forward problem.
Notes	All failure return values are negative and therefore a test <code>flag < 0</code> will trap all <code>CVodeB</code> failures.

7.2.6 Optional input functions for the backward problem

7.2.6.1 Main solver optional input functions

The adjoint module in `CVODES` provides wrappers for most of the optional input functions defined in §5.5.5.1. The only difference is that the first argument of the optional input functions for the backward problem is the pointer to the adjoint memory block, `cvadj_mem`, of type `void *`, returned by `CVadjMalloc`. The optional input functions defined for the backward problem are:

```
flag = CVodeSetErrHandlerB(cvadj_mem, ehfunB, eh\_dataB);
flag = CVodeSetErrFileB(cvadj_mem, errfpB);
flag = CVodeSetIterTypeB(cvadj_mem, iterB);
flag = CVodeSetFdataB(cvadj_mem, f_dataB);
flag = CVodeSetMaxOrdB(cvadj_mem, maxordB);
flag = CVodeSetMaxNumStepsB(cvadj_mem, mxstepsB);
flag = CVodeSetStabLimDetB(cvadj_mem, stldetB);
flag = CVodeSetInitStepB(cvadj_mem, hinB);
flag = CVodeSetMinStepB(cvadj_mem, hminB);
flag = CVodeSetMaxStepB(cvadj_mem, hmaxB);
```

Their return value `flag` (of type `int`) can have any of the return values of their counterparts, but it can also be `CV_ADJMEM_NULL` if the `cvadj_mem` argument was `NULL`.

7.2.6.2 Dense linear solver

Optional inputs for the `CVDENSE` linear solver module can be set for the backward problem through the following function:

CVDenseSetJacFnB

Call `flag = CVDenseSetJacFnB(cvadj_mem, djacB, jac_dataB);`

Description The function `CVDenseSetJacFnB` specifies the dense Jacobian approximation function to be used for the backward problem and the pointer to user data.

Arguments `cvadj_mem` (`void *`) pointer to the adjoint memory block.

	<code>djacB</code>	(<code>CVDenseJacFnB</code>) user-defined dense Jacobian approximation function.
	<code>jac_dataB</code>	(<code>void *</code>) pointer to the user-defined data structure.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of:	
	<code>CVDENSE_SUCCESS</code>	The optional value has been successfully set.
	<code>CVDENSE_MEM_NULL</code>	The CVODES solver memory block was not created through a call to <code>CVodeCreateB</code> .
	<code>CVDENSE_LMEM_NULL</code>	The CVDENSE linear solver has not been initialized through a call to <code>CVDenseB</code> .
	<code>CVDENSE_ADJMEM_NULL</code>	The <code>cvadj_mem</code> argument was NULL.
Notes	The function type <code>CVDenseJacFnB</code> is described in §7.3.	

7.2.6.3 Band linear solver

Optional inputs for the CVBAND linear solver module can be set for the backward problem through the following function:

CVBandSetJacFnB

Call	<code>flag = CVBandSetJacFnB(cvadj_mem, bjacB, jac_dataB);</code>	
Description	The function <code>CVBandSetJacFnB</code> specifies the banded Jacobian approximation function to be used for the backward problem and the pointer to user data.	
Arguments	<code>cvadj_mem</code> (<code>void *</code>) pointer to the adjoint memory block. <code>bjacB</code> (<code>CVBandJacFnB</code>) user-defined banded Jacobian approximation function. <code>jac_dataB</code> (<code>void *</code>) pointer to the user-defined data structure.	
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of:	
	<code>CVBAND_SUCCESS</code>	The optional value has been successfully set.
	<code>CVBAND_MEM_NULL</code>	The CVODES solver memory block was not created through a call to <code>CVodeCreateB</code> .
	<code>CVBAND_LMEM_NULL</code>	The CVBAND linear solver has not been initialized through a call to <code>CVBandB</code> .
	<code>CVBAND_ADJMEM_NULL</code>	The <code>cvadj_mem</code> argument was NULL.
Notes	The function type <code>CVBandJacFnB</code> is described in §7.3.	

7.2.6.4 SPILS linear solvers

Optional inputs for the CVSPILS linear solver module can be set for the backward problem through the following functions:

CVSpilsSetPreconditionerB

Call	<code>flag = CVSpilsSetPreconditionerB(cvadj_mem, psolveB, psetupB, p_dataB);</code>	
Description	The function <code>CVSpilsSetPrecSolveFnB</code> specifies the preconditioner setup and solve functions and the pointer to user data for the backward integration.	
Arguments	<code>cvadj_mem</code> (<code>void *</code>) pointer to the adjoint memory block. <code>psolveB</code> (<code>CVSpilsPrecSolveFnB</code>) user-defined preconditioner solve function. <code>psetupB</code> (<code>CVSpilsPrecSetupFnB</code>) user-defined preconditioner setup function. <code>p_dataB</code> (<code>void *</code>) pointer to the user-defined data structure.	
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of:	
	<code>CVSPILS_SUCCESS</code>	The optional value has been successfully set.

	CVSPILS_MEM_NULL	The CVODES solver memory block was not created through a call to <code>CVodeCreateB</code> .
	CVSPILS_LMEM_NULL	The CVSPGMR linear solver has not been initialized.
	CVSPILS_ADJMEM_NULL	The <code>cvadj_mem</code> argument was NULL.
Notes	The function types <code>CVSpilsPrecSolveFnB</code> and <code>CVSpilsPrecSetupFnB</code> are described in §7.3.	

CVSpilsSetJacTimesVecFnB

Call	<code>flag = CVSpilsSetJacTimesVecFnB(cvadj_mem, jtimesB, jac_data);</code>	
Description	The function <code>CVSpilsSetJacTimesFnB</code> specifies the Jacobian-vector product function to be used and te pointer to user data.	
Arguments	<code>cvadj_mem</code> (void *) pointer to the adjoint memory block. <code>jtimesB</code> (<code>CVSpilsJacTimesVecFnB</code>) user-defined Jacobian-vector product function. <code>jac_dataB</code> (void *) pointer to the user-defined data structure.	
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of:	
	CVSPILS_SUCCESS	The optional value has been successfully set.
	CVSPILS_MEM_NULL	The CVODES solver memory block was not created through a call to <code>CVodeCreateB</code> .
	CVSPILS_LMEM_NULL	The CVSPGMR linear solver has not been initialized.
	CVSPILS_ADJMEM_NULL	The <code>cvadj_mem</code> argument was NULL.
Notes	The function type <code>CVSpilsJacTimesVecFnB</code> is described in §7.3.	

CVSpilsSetGSTypeB

Call	<code>flag = CVSpilsSetGSType(cvadj_mem, gstypeB);</code>	
Description	The function <code>CVSpilsSetGSTypeB</code> specifies the type of Gram-Schmidt orthogonalization to be used with CVSPGMR. 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>cvadj_mem</code> (void *) pointer to the adjoint memory block. <code>gstypeB</code> (int) type of Gram-Schmidt orthogonalization.	
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of:	
	CVSPILS_SUCCESS	The optional value has been successfully set.
	CVSPILS_MEM_NULL	The CVODES solver memory block was not created through a call to <code>CVodeCreateB</code> .
	CVSPILS_LMEM_NULL	The CVSPGMR linear solver has not been initialized.
	CVSPILS_ILL_INPUT	The Gram-Schmidt orthogonalization type <code>gstypeB</code> is not valid.
	CVSPILS_ADJMEM_NULL	The <code>cvadj_mem</code> argument was NULL.
Notes	The default value is <code>MODIFIED_GS</code> . This option is available only with CVSPGMR.	



CVSpilsSetDeltB

Call	<code>flag = CVSpilsSetDeltB(cvadj_mem, deltB);</code>	
Description	The function <code>CVSpilsSetDeltB</code> specifies the factor by which the Krylov linear solver's convergence test constant is reduced from the Newton iteration test constant.	
Arguments	<code>cvadj_mem</code> (void *) pointer to the adjoint memory block.	

	<code>deltB</code> (<code>realtype</code>) the value of the convergence test constant reduction factor.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of:
	<code>CVSPILS_SUCCESS</code> The optional value has been successfully set.
	<code>CVSPILS_MEM_NULL</code> The CVODES solver memory block was not created through a call to <code>CVodeCreateB</code> .
	<code>CVSPILS_LMEM_NULL</code> The CVSPGMR linear solver has not been initialized.
	<code>CVSPILS_ILL_INPUT</code> The factor <code>deltB</code> is negative.
	<code>CVSPILS_ADJMEM_NULL</code> The <code>cvadj_mem</code> argument was <code>NULL</code> .
Notes	The default value is 0.05. Passing a value <code>deltB= 0.0</code> also indicates using the default value.

CVSpilsSetMaxlB

Call	<code>flag = CVSpilsSetMaxlB(cvadj_mem, maxlB);</code>
Description	The function <code>CVSpilsSetMaxlB</code> resets maximum Krylov subspace dimension for the Bi-CGStab or TFQMR methods.
Arguments	<code>cvadj_mem</code> (<code>void *</code>) pointer to the adjoint memory block. <code>maxlB</code> (<code>realtype</code>) maximum dimension of the Krylov subspace.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of:
	<code>CVSPILS_SUCCESS</code> The optional value has been successfully set.
	<code>CVSPILS_MEM_NULL</code> The CVODES solver memory block was not created through a call to <code>CVodeCreateB</code> .
	<code>CVSPILS_LMEM_NULL</code> The CVSPGMR linear solver has not been initialized.
	<code>CVSPILS_ILL_INPUT</code> The factor <code>deltB</code> is negative.
	<code>CVSPILS_ADJMEM_NULL</code> The <code>cvadj_mem</code> argument was <code>NULL</code> .
Notes	The maximum subspace dimension is initially specified in the call to <code>CVSpbcgB</code> or <code>CVSptfqmrB</code> . The call to <code>CVSpilsSetMaxlB</code> is needed only if <code>maxl</code> is being changed from its previous value.
	This option is available only for the CVSPBCG and CVSPTFQMR linear solvers.



CVSpilsSetPrecTypeB

Call	<code>flag = CVSpilsSetPrecTypeB(cvadj_mem, pretypeB);</code>
Description	The function <code>CVSpilsSetPrecTypeB</code> resets the type of preconditioning to be used.
Arguments	<code>cvadj_mem</code> (<code>void *</code>) pointer to the adjoint memory block. <code>pretypeB</code> (<code>int</code>) specifies the type of preconditioning and must be one of: <code>PREC_NONE</code> , <code>PREC_LEFT</code> , <code>PREC_RIGHT</code> , or <code>PREC_BOTH</code> .
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of:
	<code>CVSPILS_SUCCESS</code> The optional value has been successfully set.
	<code>CVSPILS_MEM_NULL</code> The CVODES solver memory block was not created through a call to <code>CVodeCreateB</code> .
	<code>CVSPILS_LMEM_NULL</code> The CVSPGMR linear solver has not been initialized.
	<code>CVSPILS_ILL_INPUT</code> The preconditioner type <code>pretype</code> is not valid.
	<code>CVSPILS_ADJMEM_NULL</code> The <code>cvadj_mem</code> argument was <code>NULL</code> .
Notes	The preconditioning type is initially specified in the call to the linear solver specification function (see §7.2.4). The call to <code>CVSpilsSetPrecTypeB</code> is needed only if <code>pretypeB</code> is being changed from its previous value.

7.2.7 Optional output functions for the backward problem

The user of the adjoint module in CVODES has access to any of the optional output functions described in §5.5.7, both for the main solver and for the linear solver modules. The first argument of these `CVodeGet*` and `CV*Get*` functions is the CVODES memory block for the backward problem. In order to call any of these functions, the user must first call the following function to obtain a pointer to this memory block:

CVadjGetCVodeBmem

Call	<code>cvode_memB = CVadjGetCVodeBmem(cvadj_mem);</code>
Description	The function <code>CVadjGetCVodeBmem</code> returns a pointer to the CVODES memory block for the backward problem.
Arguments	The argument <code>cvadj_mem</code> (of type <code>void *</code>) is a pointer to the adjoint memory block returned by <code>CVadjMalloc</code> .
Return value	The return value, <code>cvode_memB</code> (of type <code>void *</code>), is a pointer to the CVODES memory for the backward problem.
Notes	The user should not modify in any way <code>cvode_memB</code> .



7.2.8 Backward integration of pure quadrature equations

7.2.8.1 Backward quadrature initialization functions

The function `CVodeQuadMallocB` initializes and allocates memory for the backward integration of quadrature equations. It has the following form:

CVodeQuadMallocB

Call	<code>flag = CVodeQuadMallocB(cvadj_mem, fQB, yQB0);</code>
Description	The function <code>CVodeQuadMallocB</code> provides required problem specifications, allocates internal memory, and initializes backward quadrature integration.
Arguments	<p><code>cvadj_mem</code> (<code>void *</code>) pointer to the adjoint memory block returned by <code>CVadjMalloc</code>.</p> <p><code>fQB</code> (<code>CVQuadRhsFnB</code>) is the C function which computes fQB, the right-hand side of the backward quadrature equations. This function has the form <code>fQB(t, y, yB, qBdot, fQ_dataB)</code> (for full details see below).</p> <p><code>yQB0</code> (<code>N_Vector</code>) is the value of the quadrature variables at <code>tB0</code>.</p>
Return value	<p>The return value <code>flag</code> (of type <code>int</code>) will be one of the following:</p> <p><code>CV_SUCCESS</code> The call to <code>CVodeQuadMallocB</code> was successful.</p> <p><code>CV_MEM_NULL</code> The CVODES solver memory block was not created through a previous call to <code>CVodeCreateB</code>.</p> <p><code>CV_MEM_FAIL</code> A memory allocation request has failed.</p> <p><code>CV_ADJMEM_NULL</code> The <code>cvadj_mem</code> argument was <code>NULL</code>.</p>

The integration of quadrature equations during the backward phase can be re-initialized by calling

CVodeQuadReInitB

Call	<code>flag = CVodeQuadReInitB(cvadj_mem, fQB, yQB0);</code>
Description	The function <code>CVodeQuadReInitB</code> re-initializes the backward quadrature integration.
Arguments	<p><code>cvadj_mem</code> (<code>void *</code>) pointer to the adjoint memory block.</p> <p><code>fQB</code> (<code>CVQuadRhsFnB</code>) is the C function which computes fQB, the right-hand side of the backward quadrature equations.</p> <p><code>yQB0</code> (<code>N_Vector</code>) is the value of the quadrature variables at <code>tB0</code>.</p>

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

<code>CV_SUCCESS</code>	The call to <code>CVodeReInitB</code> was successful.
<code>CV_MEM_NULL</code>	The CVODES solver memory block was not created through a previous call to <code>CVodeCreateB</code> .
<code>CV_NO_QUAD</code>	Quadrature integration was not activated through a previous call to <code>CVodeQuadMallocB</code> .
<code>CV_ADJMEM_NULL</code>	The <code>cvadj_mem</code> argument was <code>NULL</code> .

7.2.8.2 Backward quadrature extraction function

To extract the values of the quadrature variables at the last return time of `CVodeB`, CVODES provides a wrapper for the function `CVodeGetQuad` (see §5.7.3). The call to this function has the form

`CVodeGetQuadB`

Call `flag = CVodeGetQuadB(cvadj_mem, yQB);`

Description The function `CVodeGetQuadB` returns the quadrature solution vector after a successful return from `CVodeB`.

Arguments `cvadj_mem` (`void *`) pointer to the adjoint memory returned by `CVadjMalloc`.
`yQB` (`N_Vector`) the computed backward quadrature vector.

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

<code>CV_SUCCESS</code>	<code>CVodeGetQuadB</code> was successful.
<code>CV_MEM_NULL</code>	The CVODES solver memory block was not created through a previous call to <code>CVodeCreateB</code> .
<code>CV_NO_QUAD</code>	Quadrature integration was not initialized through a previous call to <code>CVodeQuadMallocB</code> .
<code>Cv_BAD_DKY</code>	<code>yQB</code> is <code>NULL</code> .
<code>CV_ADJMEM_NULL</code>	The <code>cvadj_mem</code> argument was <code>NULL</code> .

7.2.8.3 Optional input/output functions for backward quadrature integration

Optional values controlling the backward integration of quadrature equations can be changed from their default values through calls to one of the following functions which are wrappers for the corresponding optional input functions defined in §5.7.4:

```
flag = CVodeSetQuadFdataB(cvadj_mem, fQ_dataB);
flag = CVodeSetQuadErrConB(cvadj_mem, errconQB, itolQB, reltolQB, abstolQB);
```

Their return value `flag` (of type `int`) can have any of the return values of their counterparts, but it can also be `CV_ADJMEM_NULL` if the `cvadj_mem` argument was `NULL`.

Access to optional outputs related to backward quadrature integration can be obtained by calling the corresponding `CVodeGetQuad*` functions (see §5.7.5). A pointer to the CVODES memory block for the backward problem, required as the first argument of these functions, can be obtained through a call to the functions `CVadjGetCVodeBmem` (see §7.2.7).

7.2.9 Optional output from the adjoint module

7.2.9.1 Checkpoint information function

For debugging purposes, CVODES provides a function `CVadjGetCheckPointsInfo` which returns partial information from the linked list of checkpoints generated by `CVodeF`. The call to this function has the form:

CVadjGetCheckPointsInfo

- Call** `flag = CVadjGetCheckPointsList(cvadj_mem, ckpnt);`
- Description** The function `CVadjGetCheckPointsList` returns a structure array with checkpoint information.
- Arguments** `cvadj_mem` (`void *`) pointer to the adjoint memory returned by `CVadjMalloc`.
`ckpnt` (`CVadjCheckPointRec *`) an array of `ncheck+1` structures with checkpoint information, where `ncheck` is the numebr of checkpoints returned by `CVodeF`.
- Return value** The return value `flag` of `CVadjGetCheckPointsInfo` is one of:
`CV_SUCCESS` `CVadjGetCheckPointsInfo` was successful.
`CV_ADJMEM_NULL` The `cvadj_mem` argument was `NULL`.
- Notes** For an example of using `CVadjGetCheckPointsInfo`, see the `cvadjdenx` example.
- The type `CVadjCheckPointRec` is defined in the header file `cvodes.h`:

```
typedef struct {
    void *my_addr;
    void *next_addr;
    realtype t0;
    realtype t1;
    long int nstep;
    int order;
    realtype step;
} CVadjCheckPointRec;
```

The fields in this structure have the following meanings:

- `my_addr` Address of current checkpoint.
- `next_addr` Address of next checkpoint.
- `t0`
- `t1` Time interval between current and next checkpoint.
- `nstep` Step number at which the current checkpoint was saved.
- `order` Linear multistep method order at the current checkpoint.
- `step` Integration stepsize at current checkpoint.

7.2.9.2 Interpolation data

For debugging purposes, CVODEA provides two extraction functions which return the data stored for interpolation purposes.

CVadjGetDataPointHermite

- Call** `int = CVadjGetDataPointHermite(cvadj_mem, which, &t, y, yd);`
- Description** The function `CVadjGetDataPointHermite` returns the time and two vectors associated with the `which` interpolation data point.
- Arguments** `cvadj_mem` (`void *`) pointer to the adjoint memory returned by `CVadjMalloc`.
`which` (`long int`) index of the interpolation data point.
`t` (`realtype *`)
`y` (`N_Vector`)
`yd` (`N_Vector`) time, solution, and solution derivative for the forward problem stored for interpolation purposes at the `which` data point.
- Return value** The return value `flag` is one of:

CV_SUCCESS CVadjGetDataPointHermite was successful.
 CV_ILL_INPUT The interpolation type was not cubic Hermite.
 CV_ADJMEM_NULL The `cvadj_mem` argument was NULL.

Notes It is the user's responsibility to allocate space for `y` and `yd`.

CVadjGetDataPointPolynomial

Call `int = CVadjGetDataPointPolynomial(cvadj_mem, which, &t, order, y);`

Description The function `CVadjGetDataPointPolynomial` returns the time and two vectors associated with the `which` interpolation data point.

Arguments `cvadj_mem` (void *) pointer to the adjoint memory returned by `CVadjMalloc`.

`which` (long int) index of the interpolation data point.

`t` (realtype *)

`order` (int)

`yd` (N_Vector) time, method order, and solution of the forward problem stored for interpolation purposes at the `which` data point.

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

CV_SUCCESS CVadjGetDataPointPolynomial was successful.

CV_ILL_INPUT The interpolation type was not variable-order polynomial.

CV_ADJMEM_NULL The `cvadj_mem` argument was NULL.

Notes It is the user's responsibility to allocate space for `y`.

7.2.9.3 Return flag name

The names of constants associated with CVODEA-specific return values can be obtained by calling the following function:

CVadjGetReturnFlagName

Call `name = CVadjGetReturnFlagName(flag);`

Description The function `CVadjGetReturnFlagName` returns the name of the CVODEA constant corresponding to `flag`.

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

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

7.3 User-supplied functions for adjoint sensitivity analysis

In addition to the required ODE right-hand side function and any optional functions for the forward problem, when using the adjoint sensitivity module in CVODES, the user must supply one function defining the backward problem ODE and, optionally, functions to supply Jacobian-related information (if Newton iteration is chosen) and one or two functions that define the preconditioner (if one of the CVSPILS solvers is selected) for the backward problem. Type definitions for all these user-supplied functions are given below.

7.3.1 ODE right-hand side for the backward problem

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

CVRhsFnB

Definition	<code>typedef int (*CVRhsFnB)(realtype t, N_Vector y, N_Vector yB, N_Vector yBdot, void *f_dataB);</code>
Purpose	This function computes the right-hand side of the backward problem ODE system. This could be (3.17) or (3.20).
Arguments	<p><code>t</code> is the current value of the independent variable.</p> <p><code>y</code> is the current value of the forward solution vector.</p> <p><code>yB</code> is the current value of the dependent variable vector.</p> <p><code>yBdot</code> is the output vector containing the right-hand side of the backward ODE problem.</p> <p><code>f_dataB</code> is a pointer to user data - the same as the <code>f_dataB</code> parameter passed to <code>CVodeSetFdataB</code>.</p>
Return value	A <code>CVRhsFnB</code> should return 0 if successful, a positive value if a recoverable error occurred (in which case <code>CVODES</code> will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and <code>CVodeB</code> returns <code>CV_RHSFUNC_FAIL</code>).
Notes	<p>Allocation of memory for <code>yBdot</code> is handled within <code>CVODES</code>.</p> <p>The <code>y</code>, <code>yB</code>, and <code>yBdot</code> arguments are all of type <code>N_Vector</code>, but <code>yB</code> and <code>yBdot</code> typically have different internal representations from <code>y</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 <code>CVODES</code> do not perform any consistency checks with respect to their <code>N_Vector</code> arguments (see §8.1 and §8.2).</p> <p>The <code>f_dataB</code> pointer is passed to the user's <code>fB</code> function every time it is called and can be the same as the <code>f_data</code> pointer used for the forward problem.</p> <p>Before calling the user's <code>CVRhsFnB</code>, <code>CVODEA</code> needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, <code>CVODEA</code> triggers an unrecoverable failure in the right-hand side function which will halt the integration and <code>CVodeB</code> will return <code>CV_RHSFUNC_FAIL</code>.</p>



7.3.2 Quadrature right-hand side for the backward problem

The user must provide a function of type `CVQuadRhsFnB` defined by

CVQuadRhsFnB

Definition	<code>typedef int (*CVQuadRhsFnB)(realtype t, N_Vector y, N_Vector yB, N_Vector qBdot, void *fQ_dataB);</code>
Purpose	This function computes the quadrature equation right-hand side for the backward problem.
Arguments	<p><code>t</code> is the current value of the independent variable.</p> <p><code>y</code> is the current value of the forward solution vector.</p> <p><code>yB</code> is the current value of the dependent variable vector.</p> <p><code>qBdot</code> is the output vector containing the right-hand side of the backward quadrature equations.</p> <p><code>fQ_dataB</code> is a pointer to user data - the same as the <code>fQ_dataB</code> parameter passed to <code>CVodeSetQuadFdataB</code>.</p>
Return value	A <code>CVQuadRhsFnB</code> should return 0 if successful, a positive value if a recoverable error occurred (in which case <code>CVODES</code> will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and <code>CVodeB</code> returns <code>CV_QRHSFUNC_FAIL</code>).

Notes	<p>Allocation of memory for <code>qBdot</code> is handled within CVODES.</p> <p>The <code>y</code>, <code>yB</code>, and <code>yQBdot</code> 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 CVODES do not perform any consistency checks with respect to their <code>N_Vector</code> arguments (see §8.1 and §8.2).</p> <p>The <code>fQ_dataB</code> pointer is passed to the user's <code>fQB</code> function every time it is called and can be the same as the <code>f_data</code> pointer used for the forward problem.</p> <p>Before calling the user's <code>CVQuadRhsFnB</code>, CVODEA needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, CVODEA triggers an unrecoverable failure in the quadrature right-hand side function which will halt the integration and <code>CVodeB</code> will return <code>CV_QRHSFUNC_FAIL</code>.</p>
-------	---



7.3.3 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. `CVDenseB` is called in step 17 of §7.1), the user may provide, through a call to `CVDenseSetJacFnB` (see §7.2.6), a function of the following type:

CVDenseJacFnB

Definition	<pre>typedef int (*CVDenseJacFnB)(long int nB, DenseMat JB, realtype t, N_Vector y, N_Vector yB, N_Vector fyB, void *jac_dataB, N_Vector tmp1B, N_Vector tmp2B, N_Vector tmp3B);</pre>
Purpose	<p>This function computes the dense Jacobian of the backward problem (or an approximation to it). If the backward problem is the adjoint of the original IVP, then this Jacobian is just the transpose of $J = \partial f / \partial y$ with a change in sign.</p>
Arguments	<p><code>nB</code> is the backward problem size.</p> <p><code>J</code> is the output Jacobian matrix.</p> <p><code>t</code> is the current value of the independent variable.</p> <p><code>y</code> is the current value of the forward solution vector.</p> <p><code>yB</code> is the current value of the dependent variable vector.</p> <p><code>fyB</code> is the current value of the right-hand side of the backward problem.</p> <p><code>jac_dataB</code> is a pointer to user data - the same as the <code>jac_dataB</code> parameter passed to <code>CVDenseSetJacDataB</code>.</p> <p><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>CVDenseJacFnB</code> as temporary storage or work space.</p>
Return value	<p>A <code>CVDenseJacFnB</code> should return 0 if successful, a positive value if a recoverable error occurred (in which case CVODES will attempt to correct, while <code>CVDENSE</code> sets <code>last_flag</code> on <code>CVDENSE_JACFUNC_RECVR</code>), or a negative value if it failed unrecoverably (in which case the integration is halted, <code>CVodeB</code> returns <code>CV_LSETUP_FAIL</code> and <code>CVDENSE</code> sets <code>last_flag</code> on <code>CVDENSE_JACFUNC_UNRECVR</code>).</p>
Notes	<p>A user-supplied dense Jacobian function must load the <code>nB</code> by <code>nB</code> dense matrix <code>JB</code> with an approximation to the Jacobian matrix at the point (t, y, yB), where <code>y</code> is the solution of the original IVP at time <code>t</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>JB</code> as this matrix is set to</p>

zero before the call to the Jacobian function. The type of JB is `DenseMat`. The user is referred to §5.6.4 for details regarding accessing a `DenseMat` object.

Before calling the user's `CVDenseJacFnB`, CVODEA needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, CVODEA triggers an unrecoverable failure in the Jacobian function which will halt the integration (`CVodeB` returns `CV_LSETUP_FAIL` and `CVDENSE` sets `last_flag` on `CVDENSE_JACFUNC_UNRECVR`).



7.3.4 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. `CVBandB` is called in step 17 of §7.1), the user may provide, through a call to `CVBandSetJacFnB` (see §7.2.6), a function the following type:

`CVBandJacFnB`

Definition	<pre>typedef int (*CVBandJacFnB)(long int nB, long int mupperB, long int mlowerB, BandMat JB, realtype t, N_Vector y, N_Vector yB, N_Vector fyB, void *jac_dataB, N_Vector tmp1B, N_Vector tmp2B, N_Vector tmp3B);</pre>	
Purpose	This function computes the banded Jacobian of the backward problem (or a banded approximation to it).	
Arguments	nB	is the backward problem size.
	mlowerB	
	mupperB	are the lower and upper half-bandwidth of the Jacobian.
	JB	is the output Jacobian matrix.
	t	is the current value of the independent variable.
	y	is the current value of the forward solution vector.
	yB	is the current value of the dependent variable vector.
	fyB	is the current value of the right-hand side of the backward problem.
	jac_dataB	is a pointer to user data - the same as the <code>jac_dataB</code> parameter passed to <code>CVBandSetJacDataB</code> .
	tmp1B	
	tmp2B	
	tmp3B	are pointers to memory allocated for variables of type <code>N_Vector</code> which can be used by <code>CVBandJacFnB</code> as temporary storage or work space.
Return value	A <code>CVBandJacFnB</code> should return 0 if successful, a positive value if a recoverable error occurred (in which case CVODES will attempt to correct, while <code>CVBAND</code> sets <code>last_flag</code> on <code>CVBAND_JACFUNC_RECVR</code>), or a negative value if it failed unrecoverably (in which case the integration is halted, <code>CVodeB</code> returns <code>CV_LSETUP_FAIL</code> and <code>CVDENSE</code> sets <code>last_flag</code> on <code>CVBAND_JACFUNC_UNRECVR</code>).	
Notes	A user-supplied band Jacobian function must load the band matrix JB (of type <code>BandMat</code>) with the elements of the Jacobian at the point (t, y, yB) , where y is the solution of the original IVP at time t and yB is the solution of the backward problem at the same time. Only nonzero elements need to be loaded into JB because JB is preset to zero before the call to the Jacobian function. More details on the accessor macros provided for a <code>BandMat</code> object and on the rest of the arguments passed to a function of type <code>CVBandJacFnB</code> are given in §5.6.5.	



Before calling the user's `CVBandJacFnB`, CVODEA needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, CVODEA triggers an unrecoverable failure in the Jacobian function which will halt the integration (`CVodeB` returns `CV_LSETUP_FAIL` and `CVBAND` sets `last_flag` on `CVBAND_JACFUNC_UNRECVR`).

7.3.5 Jacobian information for the backward problem (matrix-vector product)

If one of the Krylov iterative linear solvers SPGMR, SPBCG, or SPTFQMR is selected (`CVSp*B` is called in step 17 of §7.1), the user may provide a function of type `CVSpilsJacTimesVecFnB` in the following form:

`CVSpilsJacTimesVecFnB`

Definition	<pre>typedef int (*CVSpilsJacTimesVecFnB)(N_Vector vB, N_Vector JvB, realtype t, N_Vector y, N_Vector yB, N_Vector fyB, void *jac_dataB, N_Vector tmpB);</pre>		
Purpose	This function computes the action of the Jacobian on a given vector <code>vB</code> for the backward problem (or an approximation to it).		
Arguments	<code>vB</code>	is the vector by which the Jacobian must be multiplied to the right.	
	<code>JvB</code>	is the output vector computed.	
	<code>t</code>	is the current value of the independent variable.	
	<code>y</code>	is the current value of the forward solution vector.	
	<code>yB</code>	is the current value of the dependent variable vector.	
	<code>fyB</code>	is the current value of the right-hand side of the backward problem.	
	<code>jac_dataB</code>	is a pointer to user data - the same as the <code>jac_dataB</code> parameter passed to <code>CVSp*SetJacTimesVecFnB</code> .	
	<code>tmpB</code>	is a pointer to memory allocated for a vector which can be used for work space.	
Return value	The return value of a function of type <code>CVSpilsJtimesFnB</code> should be 0 if successful or nonzero if an error was encountered, in which case the integration is halted.		
Notes	<p>A user-supplied Jacobian-vector product function must load the vector <code>JvB</code> with the result of the product between 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>CVSpilsJacTimesVecFn</code> (see §5.6.6). If the backward problem is the adjoint of $\dot{y} = f(t, y)$, then this function is to compute $-(\partial f / \partial y)^T v_B$.</p>		

7.3.6 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 may be either a left or a right preconditioner matrix. This function must be of type `CVSpilsPrecSolveFnB` defined by

`CVSpilsPrecSolveFnB`

Definition	<pre>typedef int (*CVSpilsPrecSolveFnB)(realtype t, N_Vector y, N_Vector yB, N_Vector fyB, N_Vector rB, N_Vector zB, realtype gammaB, realtype deltaB, int lrB, void *P_dataB, N_Vector tmpB);</pre>
Purpose	This function solves the preconditioning system $Pz = r$ for the backward problem.
Arguments	<p>t is the current value of the independent variable.</p> <p>y is the current value of the forward solution vector.</p> <p>yB is the current value of the dependent variable vector.</p> <p>fyB is the current value of the right-hand side of the backward problem.</p> <p>rB is the right-hand side vector of the linear system.</p> <p>zB is the output vector computed.</p> <p>gammaB is the scalar γ appearing in the Newton matrix.</p> <p>deltaB is an input tolerance to be used if an iterative method is employed in the solution.</p> <p>lrB is an input flag indicating whether the preconditioner solve function is to use the left preconditioner (lr=1) or the right preconditioner (lr=2);</p> <p>P_dataB is a pointer to user data - the same as the P_dataB parameter passed to the function CVSp*SetPreconditionerB.</p> <p>tmpB is a pointer to memory allocated for a vector which can be used for work space.</p>
Return value	The return value of a preconditioner solve 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).

7.3.7 Preconditioning for the backward problem (Jacobian data)

If the user's preconditioner requires that any Jacobian-related data be preprocessed or evaluated, then this needs to be done in a user-supplied C function of type **CVSpilsPrecSetupFnB** defined by

CVSpilsPrecSetupFnB

Definition	<pre>typedef int (*CVSpilsPrecSetupFnB)(realtype t, N_Vector y, N_Vector yB, N_Vector fyB, booleantype jokB, booleantype *jcurPtrB, realtype gammaB, void *P_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 a CVSpilsPrecSetupFnB are as follows:</p> <p>t is the current value of the independent variable.</p> <p>y is the current value of the forward solution vector.</p> <p>yB is the current value of the dependent variable vector.</p> <p>fyB is the current value of the right-hand side of the backward problem.</p> <p>jokB is an input flag indicating whether Jacobian-related data needs to be recomputed.</p> <p>jcurPtrB is a pointer to an output integer flag which is to be set to TRUE if Jacobian data was recomputed, or to FALSE if Jacobian data was not recomputed but saved data was still reused.</p> <p>gammaB is the scalar γ appearing in the Newton matrix.</p>

P_dataB is a pointer to user data - the same as the **P_dataB** parameter passed to **CVSpilsSetPreconditionerB**.

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).

7.4 Using CVODES preconditioner modules for the backward problem

As on the forward integration phase, the efficiency of Krylov iterative methods for the solution of linear systems can be greatly enhanced through preconditioning. Both preconditioner modules provided with SUNDIALS, the banded preconditioner **CVBANDPRE** and the band-block-diagonal preconditioner module **CVBBDPRE**, provide interface functions through which they can be used on the backward integration phase.

7.4.1 Using the banded preconditioner CVBANDPRE

The adjoint module in CVODES offers an interface to the banded preconditioner module **CVBANDPRE** described in section §5.9.1. This preconditioner provides a band matrix preconditioner based on difference quotients of the backward problem right-hand side function **fB**. It generates a banded approximation to the Jacobian with m_{lB} sub-diagonals and m_{uB} super-diagonals to be used with one of the Krylov linear solvers.

In order to use the **CVBANDPRE** module in the solution of the backward problem, the user need not define any additional functions. First, the user must initialize the **CVBANDPRE** module by calling

CVBandPrecAllocB

Call `flag = CVBandPrecAlloc(cvadj_mem, nB, muB, mlB);`

Description The function **CVBandPrecAllocB** initializes and allocates memory for the **CVBANDPRE** preconditioner for the backward problem.

Arguments

- cvadj_mem** (void *) pointer to the adjoint memory block returned by **CVadjMalloc**.
- nB** (long int) backward problem dimension.
- muB** (long int) upper half-bandwidth of the backward problem Jacobian approximation.
- mlB** (long int) lower half-bandwidth of the backward problem Jacobian approximation.

Return value If successful, **CVBandPrecAlloc** stores a pointer to the newly created **CVBANDPRE** memory block. The return value **flag** (of type **int**) is one of:

- CVBANDPRE_SUCCESS** The call to **CVBandPrecAllocB** was successful.
- CVBANDPRE_MEM_FAIL** An error occurred while trying to create the **CVBANDPRE** memory block.
- CVBANDPRE_ADJMEM_NULL** The **cvadj_mem** argument was **NULL**.

To specify the use of the **CVSPGMR** linear solver module with the **CVBANDPRE** preconditioner module, make the following call:

CVBPSpgmrB

Call `flag = CVBPSpgmrB(cvadj_mem, pretypeB, maxlB);`

Description The function `CVBPSpgmrB` links the `CVBANDPRE` data to the `CVSPGMR` linear solver and attaches the latter to the `CVODES` memory block for the backward problem.

Arguments

- `cvadj_mem` (`void *`) pointer to the adjoint memory block returned by `CVadjMalloc`.
- `pretypeB` (`int`) preconditioning type. Can be one of `PREC_LEFT` or `PREC_RIGHT`.
- `maxlB` (`int`) maximum dimension of the Krylov subspace to be used. Pass 0 to use the default value `CVSPILS_MAXL=5`.

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

- `CVSPILS_SUCCESS` The `CVSPGMR` initialization was successful.
- `CVSPILS_MEM_NULL` The `CVODES` memory block for the backward problem was not initialized through a previous call to `CVodeCreateB`.
- `CVSPILS_ILL_INPUT` The preconditioner type `pretypeB` is not valid.
- `CVSPILS_MEM_FAIL` A memory allocation request failed.
- `CVBANDPRE_PDATA_NULL` The `CVBANDPRE` preconditioner has not been initialized.
- `CVBANDPRE_ADJMEM_NULL` The `cvadj_mem` argument was `NULL`.

To specify the use of the `CVSPBCG` linear solver module with the `CVBANDPRE` preconditioner module, make the following call:

CVBPSpbcgB

Call `flag = CVBPSpbcgB(cvadj_mem, pretypeB, maxlB);`

Description The function `CVBPSpbcgB` links the `CVBANDPRE` data to the `CVSPBCG` linear solver and attaches the latter to the `CVODES` memory block for the backward problem.

Arguments

- `cvadj_mem` (`void *`) pointer to the adjoint memory block returned by `CVadjMalloc`.
- `pretypeB` (`int`) preconditioning type. Can be one of `PREC_LEFT` or `PREC_RIGHT`.
- `maxlB` (`int`) maximum dimension of the Krylov subspace to be used. Pass 0 to use the default value `CVSPILS_MAXL=5`.

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

- `CVSPILS_SUCCESS` The `CVSPBCG` initialization was successful.
- `CVSPILS_MEM_NULL` The `CVODES` memory block for the backward problem was not initialized through a previous call to `CVodeCreateB`.
- `CVSPILS_ILL_INPUT` The preconditioner type `pretypeB` is not valid.
- `CVSPILS_MEM_FAIL` A memory allocation request failed.
- `CVBANDPRE_PDATA_NULL` The `CVBANDPRE` preconditioner has not been initialized.
- `CVBANDPRE_ADJMEM_NULL` The `cvadj_mem` argument was `NULL`.

To specify the use of the `CVSPTFQMR` linear solver module with the `CVBANDPRE` preconditioner module, make the following call:

CVBPSptfqmrB

Call `flag = CVBPSptfqmrB(cvadj_mem, pretypeB, maxlB);`

Description The function `CVBPSptfqmrB` links the `CVBANDPRE` data to the `CVSPTFQMR` linear solver and attaches the latter to the `CVODES` memory block for the backward problem.

Arguments

- `cvadj_mem` (`void *`) pointer to the adjoint memory block returned by `CVadjMalloc`.
- `pretypeB` (`int`) preconditioning type. Can be one of `PREC_LEFT` or `PREC_RIGHT`.
- `maxlB` (`int`) maximum dimension of the Krylov subspace to be used. Pass 0 to use the default value `CVSPILS_MAXL=5`.

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

`CVBANDPRE_ADJMEM_NULL` The `CVSPTFQMR` initialization was successful.
`CVSPILS_MEM_NULL` The `CVODES` memory block for the backward problem was not initialized through a previous call to `CVodeCreateB`.
`CVSPILS_ILL_INPUT` The preconditioner type `pretypeB` is not valid.
`CVSPILS_MEM_FAIL` A memory allocation request failed.
`CVBANDPRE_PDATA_NULL` The `CVBANDPRE` preconditioner has not been initialized.
`CVBANDPRE_ADJMEM_NULL` The `cvadj_mem` argument was `NULL`.

To deallocate the `CVBANDPRE` preconditioner module memory used for the backward integration, make the following call:

CVBandPrecFreeB

Call `CVBandPrecFreeB(cvadj_mem);`

Description The function `CVBandPrecFreeB` frees the memory allocated by `CVBandPrecAllocB`.

Arguments The only argument passed to `CVBandPrecFreeB` is the pointer to the `CVODEA` memory block (of type `void *`).

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

Notes The call to `CVBandPrecFreeB` must precede that to `CVadjFree`.

For more details on `CVBANDPRE` see §5.9.1.

7.4.2 Using the band-block-diagonal preconditioner CVBBDPRE

The adjoint module in `CVODES` offers an interface to the band-block-diagonal preconditioner module `CVBBDPRE` described in section §5.9.2. 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 parallel vector module `NVECTOR_PARALLEL`.

In order to use the `CVBBDPRE` module in the solution of the backward problem, the user must define one or two additional functions, described at the end of this section.

7.4.2.1 Usage of CVBBDPRE for the backward problem

The `CVBBDPRE` module is initialized by calling

CVBBDPrecAllocB

Call `flag = CVBBDPrecAllocB(cvadj_mem, NlocalB, mudqB, mldqB, mukeepB, mlkeepB, dqrelyB, glocB, cfnB);`

Description The function `CVBBDPrecAllocB` initializes and allocates memory for the `CVBBDPRE` preconditioner for the backward problem.

Arguments `cvadj_mem` (`void *`) pointer to the adjoint memory block returned by `CVadjMalloc`.
`NlocalB` (`long int`) local vector dimension for 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.
`mukeepB` (`long int`) upper half-bandwidth of the retained banded approximate Jacobian block.
`mlkeepB` (`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 **dqrelyB** = $\sqrt{\text{unit roundoff}}$, which can be specified by passing **dqrely** = 0.0.
- glocB** (**CVLocalFnB**) the C function which computes the approximation $g_B(t, y)$ to the right-hand side of the backward problem.
- cfnB** (**CVCommFnB**) the optional C function which performs all interprocess communication required for the computation of $g_B(t, y)$.
- Return value** If successful, **CVBBDPrecAlloc** stores a pointer to the newly created **CVBBDPRE** memory block. The return value **flag** (of type **int**) is one of:
- CVBBDPRE_SUCCESS** The call to **CVBBDPrecAllocB** was successful.
- CVBBDPRE_PDATA_NULL** An error occurred while trying to create the **CVBBDPRE** memory block.
- CVBBDPRE_ADJMEM_NULL** The **cvadj_mem** argument was **NULL**.

To specify the use of the **CVSPGMR** linear solver module with the **CVBBDPRE** preconditioner module, make the following call:

CVBBDSpgmrB

- Call** `flag = CVBBDSpgmrB(cvadj_mem, pretypeB, maxlB);`
- Description** The function **CVBBDSpgmrB** links the **CVBBDPRE** data to the **CVSPGMR** linear solver and attaches the latter to the **CVODES** memory block for the backward problem.
- Arguments** **cvadj_mem** (**void ***) pointer to the adjoint memory block returned by **CVadjMalloc**.
pretypeB (**int**) preconditioning type. Can be one of **PREC_LEFT** or **PREC_RIGHT**.
maxlB (**int**) maximum dimension of the Krylov subspace to be used. Pass 0 to use the default value **CVSPILS_MAXL** = 5.
- Return value** The return value **flag** (of type **int**) is one of:
- CVSPILS_SUCCESS** The **CVSPGMR** initialization was successful.
- CVSPILS_MEM_NULL** The **CVODES** memory block for the backward problem was not initialized through a previous call to **CVodeCreateB**.
- CVSPILS_ILL_INPUT** The preconditioner type **pretypeB** is not valid.
- CVSPILS_MEM_FAIL** A memory allocation request failed.
- CVBBDPRE_PDATA_NULL** The **CVBBDPRE** preconditioner has not been initialized.
- CVBBDPRE_ADJMEM_NULL** The **cvadj_mem** argument was **NULL**.

To specify the use of the **CVSPBCG** linear solver module with the **CVBBDPRE** preconditioner module, make the following call:

CVBBDSpbcgB

- Call** `flag = CVBBDSpbcgB(cvadj_mem, pretypeB, maxlB);`
- Description** The function **CVBBDSpbcgB** links the **CVBBDPRE** data to the **CVSPBCG** linear solver and attaches the latter to the **CVODES** memory block for the backward problem.
- Arguments** **cvadj_mem** (**void ***) pointer to the adjoint memory block returned by **CVadjMalloc**.
pretypeB (**int**) preconditioning type. Can be one of **PREC_LEFT** or **PREC_RIGHT**.
maxlB (**int**) maximum dimension of the Krylov subspace to be used. Pass 0 to use the default value **CVSPILS_MAXL** = 5.
- Return value** The return value **flag** (of type **int**) is one of:
- CVSPILS_SUCCESS** The **CVSPBCG** initialization was successful.
- CVSPILS_MEM_NULL** The **CVODES** memory block for the backward problem was not initialized through a previous call to **CVodeCreateB**.

CVSPILS_ILL_INPUT	The preconditioner type <code>pretypeB</code> is not valid.
CVSPILS_MEM_FAIL	A memory allocation request failed.
CVBBDPRE_PDATA_NULL	The CVBBDPRE preconditioner has not been initialized.
CVBBDPRE_ADJMEM_NULL	The <code>cvadj_mem</code> argument was NULL.

To specify the use of the CVSPTFQMR linear solver module with the CVBBDPRE preconditioner module, make the following call:

CVBBDSPtfqmrB

Call	<code>flag = CVBBDSPtfqmrB(cvadj_mem, pretypeB, maxlB);</code>												
Description	The function <code>CVBBDSPtfqmrB</code> links the CVBBDPRE data to the CVSPTFQMR linear solver and attaches the latter to the CVODES memory block for the backward problem.												
Arguments	<p><code>cvadj_mem</code> (void *) pointer to the adjoint memory block returned by <code>CVadjMalloc</code>.</p> <p><code>pretypeB</code> (int) preconditioning type. Can be one of <code>PREC_LEFT</code> or <code>PREC_RIGHT</code>.</p> <p><code>maxlB</code> (int) maximum dimension of the Krylov subspace to be used. Pass 0 to use the default value <code>CVSPILS_MAXL=5</code>.</p>												
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of:												
	<table> <tr> <td>CVSPILS_SUCCESS</td><td>The CVSPTFQMR initialization was successful.</td></tr> <tr> <td>CVSPILS_MEM_NULL</td><td>The CVODES memory block for the backward problem was not initialized through a previous call to <code>CVodeCreateB</code>.</td></tr> <tr> <td>CVSPILS_ILL_INPUT</td><td>The preconditioner type <code>pretypeB</code> is not valid.</td></tr> <tr> <td>CVSPILS_MEM_FAIL</td><td>A memory allocation request failed.</td></tr> <tr> <td>CVBBDPRE_PDATA_NULL</td><td>The CVBBDPRE preconditioner has not been initialized.</td></tr> <tr> <td>CVBBDPRE_ADJMEM_NULL</td><td>The <code>cvadj_mem</code> argument was NULL.</td></tr> </table>	CVSPILS_SUCCESS	The CVSPTFQMR initialization was successful.	CVSPILS_MEM_NULL	The CVODES memory block for the backward problem was not initialized through a previous call to <code>CVodeCreateB</code> .	CVSPILS_ILL_INPUT	The preconditioner type <code>pretypeB</code> is not valid.	CVSPILS_MEM_FAIL	A memory allocation request failed.	CVBBDPRE_PDATA_NULL	The CVBBDPRE preconditioner has not been initialized.	CVBBDPRE_ADJMEM_NULL	The <code>cvadj_mem</code> argument was NULL.
CVSPILS_SUCCESS	The CVSPTFQMR initialization was successful.												
CVSPILS_MEM_NULL	The CVODES memory block for the backward problem was not initialized through a previous call to <code>CVodeCreateB</code> .												
CVSPILS_ILL_INPUT	The preconditioner type <code>pretypeB</code> is not valid.												
CVSPILS_MEM_FAIL	A memory allocation request failed.												
CVBBDPRE_PDATA_NULL	The CVBBDPRE preconditioner has not been initialized.												
CVBBDPRE_ADJMEM_NULL	The <code>cvadj_mem</code> argument was NULL.												

To reinitialize the CVBBDPRE preconditioner module for the backward problem call the following function:

CVBBDPrecReInitB

Call	<code>flag = CVBBDPrecReInitB(cvadj_mem, mudqB, mldqB, dqrelyB, glocB, cfnB);</code>						
Description	The function <code>CVBBDPrecReInitB</code> reinitializes the CVBBDPRE preconditioner for the backward problem.						
Arguments	<p><code>cvadj_mem</code> (void *) pointer to the adjoint memory block returned by <code>CVadjMalloc</code>.</p> <p><code>mudqB</code> (long int) upper half-bandwidth to be used in the difference-quotient Jacobian approximation.</p> <p><code>mldqB</code> (long int) lower half-bandwidth to be used in the difference-quotient Jacobian approximation.</p> <p><code>dqrelyB</code> (realtype) the relative increment in components of <code>yB</code> used in the difference quotient approximations.</p> <p><code>glocB</code> (CVLocalFnB) the C function which computes the approximation $g_B(t, y)$ to the right-hand side of the backward problem.</p> <p><code>cfnB</code> (CVCommFnB) the optional C function which performs all interprocess communication required for the computation of $g_B(t, y)$.</p>						
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of:						
	<table> <tr> <td>CVBBDPRE_SUCCESS</td><td>The CVBBDPRE initialization was successful.</td></tr> <tr> <td>CVBBDPRE_PDATA_NULL</td><td>The CVBBDPRE preconditioner has not been initialized.</td></tr> <tr> <td>CVBBDPRE_ADJMEM_NULL</td><td>The <code>cvadj_mem</code> argument was NULL.</td></tr> </table>	CVBBDPRE_SUCCESS	The CVBBDPRE initialization was successful.	CVBBDPRE_PDATA_NULL	The CVBBDPRE preconditioner has not been initialized.	CVBBDPRE_ADJMEM_NULL	The <code>cvadj_mem</code> argument was NULL.
CVBBDPRE_SUCCESS	The CVBBDPRE initialization was successful.						
CVBBDPRE_PDATA_NULL	The CVBBDPRE preconditioner has not been initialized.						
CVBBDPRE_ADJMEM_NULL	The <code>cvadj_mem</code> argument was NULL.						

To deallocate the CVBBDPRE preconditioner module memory used for the backward integration, make the following call:

CVBBDPrecFreeB

Call	<code>CVBBDPrecFreeB(cvadj_mem);</code>
Description	The function <code>CVBBDPrecFreeB</code> frees the memory allocated by <code>CVBBDPrecAllocB</code> .
Arguments	The only argument passed to <code>CVBBDPrecFreeB</code> is the pointer to the CVODEA memory block (of type <code>void *</code>).
Return value	The function <code>CVBBDPrecFreeB</code> has no return value.
Notes	The call to <code>CVBBDPrecFreeB</code> must precede that to <code>CVadjFree</code> .
For more details on CVBBDPRE see §5.9.2.	

7.4.2.2 User-supplied functions for CVBBDPRE

To use the CVBBDPRE module, the user must supply one or two functions which the module calls to construct the preconditioner: a required function `glocB` (of type `CVLocalFnB`) which approximates the right-hand side of the backward problem and which is computed locally, and an optional function `cfnB` (of type `CVCommFnB`) which performs all interprocess communication necessary to evaluate this approximate right-hand side (see §5.9.2). The prototypes for these two functions are described below.

CVLocalFnB

Definition	<pre>typedef int (*CVLocalFnB)(long int NlocalB, realtype t, N_Vector y, N_Vector yB, N_Vector gB, void *f_dataB);</pre>	
Purpose	This function loads the vector <code>gB</code> as a function of <code>t</code> , <code>y</code> , and <code>yB</code> .	
Arguments	<p><code>NlocalB</code> is the local vector length for the backward problem.</p> <p><code>t</code> is the value of the independent variable.</p> <p><code>y</code> is the current value of the forward solution vector.</p> <p><code>yB</code> is the current value of the dependent variable vector.</p> <p><code>gB</code> is the output vector.</p> <p><code>f_dataB</code> is a pointer to user data - the same as the <code>f_dataB</code> parameter passed to <code>CVodeSetFdataB</code>.</p>	
Return value	A <code>CVLocalFnB</code> should return 0 if successful, a positive value if a recoverable error occurred (in which case CVODES will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and <code>CVodeB</code> returns <code>CV_LSETUP_FAIL</code>).	
Notes	<p>This routine assumes that all interprocess communication of data needed to calculate <code>gB</code> has already been done, and this data is accessible within <code>f_dataB</code>.</p> <p>Before calling the user's <code>CVLocalFnB</code>, CVODEA needs to evaluate (through interpolation) the values of the states from the forward integration. If an error occurs in the interpolation, CVODEA triggers an unrecoverable failure in the preconditioner setup function which will halt the integration (<code>CVodeB</code> returns <code>CV_LSETUP_FAIL</code>).</p>	

**CVCommFnB**

Definition	<pre>typedef int (*CVCommFnB)(long int NlocalB, realtype t, N_Vector y, N_Vector yB, void *f_dataB);</pre>	
Purpose	This function performs all interprocess communications necessary for the execution of the <code>glocB</code> function above, using the input vectors <code>y</code> and <code>yB</code> .	
Arguments	<p><code>NlocalB</code> is the local vector length.</p> <p><code>t</code> is the value of the independent variable.</p> <p><code>y</code> is the current value of the forward solution vector.</p> <p><code>yB</code> is the current value of the dependent variable vector.</p>	

`f_dataB` is a pointer to user data - the same as the `f_dataB` parameter passed to `CVodeSetFdataB`.

Return value A `CVCommFn` should return 0 if successful, a positive value if a recoverable error occurred (in which case CVODES will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and `CVodeB` returns `CV_LSETUP_FAIL`).

Notes The `cfnB` function is expected to save communicated data in space defined within the structure `f_dataB`.

Each call to the `cfnB` function is preceded by a call to the function that evaluates the right-hand side of the backward problem with the same `t`, `y`, and `yB` arguments. If there is no additional communication needed, then pass `cfnB = NULL` to `CVBBDPrecAllocB`.

Chapter 8

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 two provided within SUNDIALS, a serial and an MPI parallel implementations.

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    (*nvlinorm)(N_Vector);
void        (*nvcompare)(realtype, N_Vector, N_Vector);
booleantype (*nvinvtest)(N_Vector, N_Vector);
booleantype (*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 8.1 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 8.1: 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 the data array.
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.
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 linear solvers, as well as the interfaces to the banded preconditioners 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 linear solver.
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 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 c : $z_i = c$, $i = 0, \dots, n-1$.
N_VProd	<code>N_VProd(x, y, z);</code> Sets the N_Vector z to be the component-wise product of the N_Vector inputs x and y : $z_i = x_i y_i$, $i = 0, \dots, n-1$.
N_VDiv	<code>N_VDiv(x, y, z);</code> Sets the N_Vector z to be the component-wise ratio of the N_Vector inputs x and y : $z_i = x_i / y_i$, $i = 0, \dots, n-1$. The y_i may not be tested for 0 values. It should only be called with an x that is guaranteed to have all nonzero components.
<i>continued on next page</i>	

continued from last page	
Name	Usage and Description
N_VScale	<code>N_VScale(c, x, z);</code> Scales the <code>N_Vector</code> <code>x</code> by the scalar <code>c</code> and returns the result in <code>z</code> : $z_i = cx_i$, $i = 0, \dots, n-1$.
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 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 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 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> : $m = \min_i x_i$.
N_VWL2Norm	<code>m = N_VWL2Norm(x, w);</code> Returns the weighted Euclidean ℓ_2 norm of the <code>N_Vector</code> <code>x</code> with 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 <code>N_Vector</code> <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 <code>N_Vector</code> <code>x</code> to the scalar <code>c</code> and returns an <code>N_Vector</code> <code>z</code> such that: $z_i = 1.0$ if $ x_i \geq c$ and $z_i = 0.0$ otherwise.
continued on next page	

continued from last page	
Name	Usage and Description
N_VInvTest	<code>t = N_VInvTest(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> , with prior testing for zero values: $z_i = 1.0/x_i$, $i = 0, \dots, n-1$. This routine returns <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 <code>FALSE</code> if any element failed the constraint test, <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.

8.1 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;
    boolean_t 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 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 8.1. Their names are obtained from those in Table 8.1 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_VCloneVectorArrayEmpty_Serial

This function creates (by cloning) an array of `count` serial vectors, each with an empty (`NULL`) data array.

```
N_Vector *N_VCloneVectorArrayEmpty_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_VCloneVectorArrayEmpty_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_VCloneVectorArrayEmpty_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.



8.2 The NVECTOR_PARALLEL implementation

The parallel implementation of the `NVECTOR` module provided with `SUNDIALS`, `NVECTOR_PARALLEL`, 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, an a boolean flag `own_data` indicating ownership of the data array `data`.

```
struct _N_VectorContent_Parallel {
    long int local_length;
    long int global_length;
    boolean_t 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 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 8.1 Their names are obtained from those in Table 8.1 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_VCloneVectorArrayEmpty_Parallel

This function creates (by cloning) an array of count parallel vectors, each with an empty (NULL) data array.

```
N_Vector *N_VCloneVectorArrayEmpty_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_VCloneVectorArrayEmpty_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_VCloneVectorArrayEmpty_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.



8.3 NVECTOR functions used by CVODES

In Table 8.2 below, we list the vector functions in the `NVECTOR` module within the `CVODES` package. The table also shows, for each function, which of the code modules uses the function. The `CVODES` column shows function usage within the main integrator module, while the remaining seven columns show function usage within each of the six `CVODES` linear solvers (`CVSPILS` stands for any of `CVSPGMR`, `CVSPBCG`, or `CVSPTFQMR`), the `CVBANDPRE` and `CVBBDPRE` preconditioner modules, and the `CVODEA` adjoint sensitivity module.

There is one subtlety in the `CVSPILS` column hidden by the table, explained here for the case of the `CVSPGMR` module). The dot product function `N_VDotProd` is called both within the implementation file `cvodes_spgmr.c` for the `CVSPGMR` solver and within the implementation files `sundials_spgmr.c` and `sundials_iterative.c` for the generic `SPGMR` solver upon which the `CVSPGMR` solver is implemented. Also, although `N_VDiv` and `N_VProd` are not called within the implementation file `cvodes_spgmr.c`, they are called within the implementation file `sundials_spgmr.c` and so are required by the `CVSPGMR` solver module. This issue does not arise for the other three `CVODES` linear solvers because the generic `DENSE` and `BAND` solvers (used in the implementation of `CVDENSE` and `CVBAND`) do not make calls to any vector functions and `CVDIAG` is not implemented using a generic diagonal solver.

At this point, we should emphasize that the `CVODES` user does not need to know anything about the usage of vector functions by the `CVODES` code modules in order to use `CVODES`. The information is presented as an implementation detail for the interested reader.

The vector functions listed in Table 8.1 that are *not* used by `CVODES` are: `N_VWL2Norm`, `N_VL1Norm`, `N_VWrmsNormMask`, `N_VConstrMask`, `N_VCloneEmpty`, and `N_VMinQuotient`. Therefore a user-supplied `NVECTOR` module for `CVODES` could omit these six kernels.

Table 8.2: List of vector functions usage by CVODES code modules

	CVODES	CVDENSE	CVBAND	CVDIAG	CVSPILS	CVBANDPRE	CVBBDPRE	CVODEA
N_VClone	✓			✓	✓			✓
N_VDestroy	✓			✓	✓			✓
N_VSpace	✓							
N_VGetArrayPointer		✓	✓			✓	✓	
N_VSetArrayPointer		✓						
N_VLinearSum	✓	✓		✓	✓			✓
N_VConst	✓				✓			
N_VProd	✓			✓	✓			
N_VDiv	✓			✓	✓			
N_VScale	✓	✓	✓	✓	✓	✓	✓	✓
N_VAbs	✓							
N_VInv	✓			✓				
N_VAddConst	✓			✓				
N_VDotProd					✓			
N_VMaxNorm	✓							
N_VWrmsNorm	✓	✓	✓		✓	✓	✓	
N_VMin	✓							
N_VCompare				✓				
N_VInvTest				✓				

Chapter 9

Providing Alternate Linear Solver Modules

The central CVODES module interfaces with the linear solver module to be used by way of calls to four functions. These are denoted here by `linit`, `lsetup`, `lsolve`, and `lfree`. Briefly, their purposes are as follows:

- `linit`: initialize and allocate memory specific to the linear solver;
- `lsetup`: preprocess and evaluate the Jacobian or preconditioner;
- `lsolve`: solve the linear system;
- `lfree`: free the linear solver memory.

A linear solver module must also provide a user-callable specification function (like those described in §5.5.3) which will attach the above four functions to the main CVODES memory block. The CVODES memory block is a structure defined in the header file `cvodes_impl.h`. A pointer to such a structure is defined as the type `CVodeMem`. The four fields in a `CVodeMem` structure that must point to the linear solver's functions are `cv_linit`, `cv_lsetup`, `cv_lsolve`, and `cv_lfree`, respectively. Note that of the four interface functions, only the `lsolve` function is required. The `lfree` function must be provided only if the solver specification function makes any memory allocation. For consistency with the existing CVODES linear solver modules, we recommend that the return value of the specification function be 0 for a successful return or a negative value if an error occurs (the pointer to the main CVODES memory block is NULL, an input is illegal, the NVECTOR implementation is not compatible, a memory allocation fails, etc.)

To facilitate data exchange between the four interface functions, the field `cv_lmem` in the CVODES memory block can be used to attach a linear solver-specific memory block.

To be used during the backward integration with the CVODEA module, a linear solver module must also provide an additional user-callable specification function (like those described in §7.2.4) which will attach the four functions to the CVODES memory block for the backward integration. Note that this block (of type `struct CVodeMemRec`) is not directly accessible to the user, but rather is itself a field (`cvb_mem`) in the CVODEA memory block. The CVODEA memory block is a structure defined in the header file `cvodea_impl.h`. A pointer to such a structure is defined as the type `CVadjMem`. The specification function for backward integration should also return a negative value if the adjoint CVODEA memory block is NULL.

An additional field (`ca_lmemB`) in the CVODEA memory block provides a hook-up for optionally attaching a linear solver-specific memory block.

The four functions that interface between CVODES and the linear solver module necessarily have fixed call sequences. Thus, a user wishing to implement another linear solver within the CVODES package must adhere to this set of interfaces. The following is a complete description of the argument list for each of these functions. Note that the argument list of each function includes a pointer to the

main CVODES memory block, by which the function can access various data related to the CVODES solution. The contents of this memory block (of type `CVodeMem`) are given in the file `cvodes_impl.h` (but not reproduced here, for the sake of space).

9.1 Initialization function

The type definition of `linit` is

`linit`

Definition `int (*linit)(CVodeMem cv_mem);`

Purpose The purpose of `linit` is to complete linear solver-specific initializations, such as counters and statistics.

Arguments `cv_mem` is the CVODES memory pointer of type `CVodeMem`.

Return value An `linit` function should return 0 if it has successfully initialized the CVODES linear solver and `-1` otherwise.

9.2 Setup function

The type definition of `lsetup` is

`lsetup`

Definition `int (*lsetup)(CVodeMem cv_mem, int convfail, N_Vector ypred,
N_Vector fpred, booleantype *jcurPtr,
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`. It may recompute Jacobian-related data if it is deemed necessary.

Arguments `cv_mem` is the CVODES memory pointer of type `CVodeMem`.
`convfail` is an input flag used to indicate any problem that occurred during the solution of the nonlinear equation on the current time step for which the linear solver is being used. This flag can be used to help decide whether the Jacobian data kept by a CVODES linear solver needs to be updated or not. Its possible values are:

- `NO_FAILURES`: this value is passed to `lsetup` if either this is the first call for this step, or the local error test failed on the previous attempt at this step (but the Newton iteration converged).
- `FAIL_BAD_J`: this value is passed to `lsetup` if (a) the previous Newton corrector iteration did not converge and the linear solver's setup function indicated that its Jacobian-related data is not current, or (b) during the previous Newton corrector iteration, the linear solver's solve function failed in a recoverable manner and the linear solver's setup function indicated that its Jacobian-related data is not current.
- `FAIL_OTHER`: this value is passed to `lsetup` if during the current internal step try, the previous Newton iteration failed to converge even though the linear solver was using current Jacobian-related data.

`ypred` is the predicted `y` vector for the current CVODES internal step.

`fpred` is the value of the right-hand side at `ypred`, i.e. $f(t_n, y_{pred})$.

`jcurPtr` is a pointer to a boolean to be filled in by `lsetup`. The function should set `*jcurPtr = TRUE` if its Jacobian data is current after the call, and should set `*jcurPtr = FALSE` if its Jacobian data is not current. If `lsetup` calls for

reevaluation of Jacobian data (based on `convfail` and CVODES state data), it should return `*jcurPtr = TRUE` unconditionally; otherwise an infinite loop can result.

`vtemp1`

`vtemp2`

`vtemp3` are temporary variables of type `N_Vector` provided for use by `lsetup`.

Return value An `lsetup` function should return 0 if successful, a positive value for a recoverable error, and a negative value for an unrecoverable error.

9.3 Solve function

The type definition of `lsolve` is

lsolve

Definition `int (*lsolve)(CNodeMem cv_mem, N_Vector b, N_Vector weight, N_Vector ycur, N_Vector fcur);`

Purpose The function `lsolve` must solve the linear equation $Mx = b$, where M is some approximation to $I - \gamma J$, $J = (\partial f / \partial y)(t_n, y_{cur})$ (see Eq.(3.6)), and the right-hand side vector b is input.

Arguments `cv_mem` is the CVODES memory pointer of type `CNodeMem`.

`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 Eq.(3.7).

`ycur` is a vector that contains the solver's current approximation to $y(t_n)$.

`fcur` is a vector that contains $f(t_n, y_{cur})$.

Return value An `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.

9.4 Memory deallocation function

The type definition of `lfree` is

lfree

Definition `void (*lfree)(CNodeMem cv_mem);`

Purpose The function `lfree` should free up any memory allocated by the linear solver.

Arguments The argument `cv_mem` is the CVODES memory pointer of type `CNodeMem`.

Return value An `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 10

Generic Linear Solvers in SUNDIALS

In this chapter, we describe five generic linear solver code modules that are included in CVODES, but which are of potential use as generic packages in themselves, either in conjunction with the use of CVODES or separately. These modules are:

- The DENSE matrix package, which includes the matrix type `DenseMat`, macros and functions for `DenseMat` matrices, and functions for small dense matrices treated as simple array types.
- The BAND matrix package, which includes the matrix type `BandMat`, macros and functions for `BandMat` matrices.
- The SPGMR package, which includes a solver for the scaled preconditioned GMRES method.
- The SPBCG package, which includes a solver for the scaled preconditioned Bi-CGStab method.
- The SPTFQMR package, which includes a solver for the scaled preconditioned TFQMR method.

For reasons related to installation, the names of the files involved in these generic solvers begin with the prefix `sundials_`. But despite this, each of the solvers is in fact generic, in that it is usable completely independently of SUNDIALS.

For the sake of space, the functions for `DenseMat` and `BandMat` matrices and the functions in SPGMR, SPBCG and SPTFQMR are only summarized briefly, since they are less likely to be of direct use in connection with CVODES. The functions for small dense matrices are fully described, because we expect that they will be useful in the implementation of preconditioners used with the combination of CVODES and the CVSPGMR, CVSPBCG or CVSPTFQMR solver.

10.1 The DENSE module

Relative to the SUNDIALS *srcdir*, the files comprising the DENSE generic linear solver are as follows:

- header files (located in *srcdir/include/sundials*)
`sundials_dense.h` `sundials_smallldense.h`
`sundials_types.h` `sundials_math.h` `sundials_config.h`
- source files (located in *srcdir/src/sundials*)
`sundials_dense.c` `sundials_smallldense.c` `sundials_math.c`

Only two of the preprocessing directives in the header file `sundials_config.h` are relevant to the DENSE package by itself (see §2.4 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 `ABS` macro and `Rabs` function.

The eight files listed above can be extracted from the SUNDIALS *srcdir* and compiled by themselves into a DENSE library or into a larger user code.

10.1.1 Type DenseMat

The type `DenseMat` is defined to be a pointer to a structure with the number of rows, number of columns, and a data field:

```
typedef struct {
    long int M;
    long int N;
    realtype **data;
} *DenseMat;
```

The M and N fields indicates the number of columns and rows, respectively, of a dense matrix, while the *data* field is a two dimensional array used for component storage. The elements of a dense matrix are stored columnwise (i.e columns are stored one on top of the other in memory). If A is of type `DenseMat`, then the (i,j) -th element of A (with $0 \leq i < M$ and $0 \leq j < N$) is given by the expression `(A->data)[j][i]` or by the expression `(A->data)[0][j*M+i]`. 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` macro. Users should use these macros whenever possible.

10.1.2 Accessor Macros

The following two macros are defined by the DENSE module to provide access to data in the `DenseMat` 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 $M \times N$ `DenseMat` 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 $M \times N$ `DenseMat` 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]`.

10.1.3 Functions

The following functions for `DenseMat` matrices are available in the DENSE package. For full details, see the header file `sundials_dense.h`.

- `DenseAllocMat`: allocation of a `DenseMat` matrix;
- `DenseAllocPiv`: allocation of a pivot array for use with `DenseGETRF`/`DenseGETRS`;
- `DenseGETRF`: LU factorization with partial pivoting;
- `DenseGETRS`: solution of $Ax = b$ using LU factorization (for square matrices A);
- `DenseZero`: load a matrix with zeros;
- `DenseCopy`: copy one matrix to another;
- `DenseScale`: scale a matrix by a scalar;
- `DenseAddI`: increment a square matrix by the identity matrix;
- `DenseFreeMat`: free memory for a `DenseMat` matrix;
- `DenseFreePiv`: free memory for a pivot array;
- `DensePrint`: print a `DenseMat` matrix to standard output.

10.1.4 Small Dense Matrix Functions

The following functions for small dense matrices are available in the DENSE package:

- `denalloc`
`denalloc(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 `denalloc` returns `NULL`. The underlying type of the dense matrix returned is `realtype**`. If we allocate a dense matrix `realtype** a` by `a = denalloc(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`.
- `denallocpiv`
`denallocpiv(n)` allocates an array of n integers. It returns a pointer to the first element in the array if successful. It returns `NULL` if the memory request could not be satisfied.
- `denGETRF`
`denGETRF(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.`denGETRF` 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.

- **denGETRS**
`denGETRS(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 `denGETRF(a,n,n,p)`. The solution x is written into the **b** array.
- **denzero**
`denzero(a,m,n)` sets all the elements of the m by n dense matrix **a** to be 0.0;
- **dencopy**
`dencopy(a,b,m,n)` copies the m by n dense matrix **a** into the m by n dense matrix **b**;
- **denscale**
`denscale(c,a,m,n)` scales every element in the m by n dense matrix **a** by **c**;
- **denaddI**
`denaddI(a,n)` increments the n by n dense matrix **a** by the identity matrix;
- **denfreepiv**
`denfreepiv(p)` frees the pivot array **p** allocated by `denallocpiv`;
- **denfree**
`denfree(a)` frees the dense matrix **a** allocated by `denalloc`;
- **denprint**
`denprint(a,m,n)` prints the m by n dense matrix **a** to standard output as it would normally appear on paper. It is intended as a debugging tool with small values of **n**. The elements are printed using the `%g` option. A blank line is printed before and after the matrix.

10.2 The BAND module

Relative to the SUNDIALS *srcdir*, the files comprising the BAND generic linear solver are as follows:

- header files (located in *srcdir/include/sundials*)
`sundials_band.h`
`sundials_types.h` `sundials_math.h` `sundials_config.h`
- source files (located in *srcdir/src/sundials*)
`sundials_band.c` `sundials_math.c`

Only two of the preprocessing directives in the header file `sundials_config.h` are required to use the BAND package by itself (see §2.4 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 of the SUNDIALS `realtype` and `booleantype` types and the macro `RCONST`, while the `sundials_math.h` header file is needed for the `MIN`, `MAX`, and `ABS` macros and `RAbs` function.

The six files listed above can be extracted from the SUNDIALS *srcdir* and compiled by themselves into a BAND library or into a larger user code.

10.2.1 Type BandMat

The type **BandMat** is the type of a large band matrix **A** (possibly distributed). It is defined to be a pointer to a structure defined by:

```
typedef struct {
    long int size;
    long int mu, ml, smu;
    realtype **data;
} *BandMat;
```

The fields in the above structure are:

- *size* is the number of columns (which is the same as the number of rows);
- *mu* is the upper half-bandwidth, $0 \leq mu \leq size-1$;
- *ml* is the lower half-bandwidth, $0 \leq ml \leq size-1$;
- *smu* is the storage upper half-bandwidth, $mu \leq smu \leq size-1$. The **BandGBTRF** routine writes the LU factors into the storage for **A**. The upper triangular factor **U**, however, may have an upper half-bandwidth as big as $\min(size-1, mu+ml)$ because of partial pivoting. The *smu* field holds the upper half-bandwidth allocated for **A**.
- *data* is a two dimensional array used for component storage. The elements of a band matrix of type **BandMat** 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.

If we number rows and columns in the band matrix starting from 0, then

- *data[0]* is a pointer to $(smu+ml+1)*size$ contiguous locations which hold the elements within the band of **A**
- *data[j]* is a pointer to the uppermost element within the band in the *j*-th column. This pointer may be treated as an array indexed from $smu-mu$ (to access the uppermost element within the band in the *j*-th column) to $smu+ml$ (to access the lowest element within the band in the *j*-th column). Indices from 0 to $smu-mu-1$ give access to extra storage elements required by **BandGBTRF**.
- *data[j][i-j+smu]* is the (i, j) -th element, $j-mu \leq i \leq j+ml$.

The macros below allow a user to 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 into the *j*-th column of elements can be obtained via the **BAND_COL** macro. Users should use these macros whenever possible.

See Figure 10.1 for a diagram of the **BandMat** type.

10.2.2 Accessor Macros

The following three macros are defined by the **BAND** module to provide access to data in the **BandMat** 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 $j-(A->mu) \leq i \leq j+(A->ml)$.




Figure 10.1: Diagram of the storage for a band matrix of type **BandMat**. Here A is an $N \times N$ band matrix of type **BandMat** with upper and lower half-bandwidths μ and $m1$, 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.

- **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 $j - (A \rightarrow \text{mu}) \leq i \leq j + (A \rightarrow \text{ml})$.

10.2.3 Functions

The following functions for **BandMat** matrices are available in the **BAND** package. For full details, see the header file `sundials_band.h`.

- **BandAllocMat**: allocation of a **BandMat** matrix;
- **BandAllocPiv**: allocation of a pivot array for use with **BandGBTRF**/**BandGBTRS**;
- **BandGBTRF**: LU factorization with partial pivoting;
- **BandGBTRS**: solution of $Ax = b$ using LU factorization;
- **BandZero**: load a matrix with zeros;
- **BandCopy**: copy one matrix to another;
- **BandScale**: scale a matrix by a scalar;
- **BandAddI**: increment a matrix by the identity matrix;
- **BandFreeMat**: free memory for a **BandMat** matrix;
- **BandFreePiv**: free memory for a pivot array;
- **BandPrint**: print a **BandMat** matrix to standard output.

10.3 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` and `sundials_iterative.c`, contains auxiliary functions that support SPGMR, as well as the other Krylov solvers in SUNDIALS (SPBCG and SPTFQMR). For full details, including usage instructions, see the header files `sundials_spgmr.h` and `sundials_iterative.h`.

Relative to the SUNDIALS *srcdir*, the files comprising the SPGMR generic linear solver 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 §2.4 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 `MAX` and `ABS` macros and `RAbs` and `RSqrt` functions.

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 SPGMR package can only be used in conjunction with an actual NVECTOR implementation library, such as the `NVECTOR_SERIAL` or `NVECTOR_PARALLEL` provided with SUNDIALS.

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.

10.3.1 Functions

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`.

10.4 The SPBCG module

The SPBCG package, in the files `sundials_spbcgs.h` and `sundials_spbcgs.c`, includes an implementation of the scaled preconditioned Bi-CGStab method. For full details, including usage instructions, see the file `sundials_spbcgs.h`.

The SPBCG package can only be used in conjunction with an actual NVECTOR implementation library, such as the `NVECTOR_SERIAL` or `NVECTOR_PARALLEL` provided with SUNDIALS.

The files needed to use the SPBCG module by itself are the same as for the SPGMR module, with `sundials_spbcgs.(h,c)` replacing `sundials_spgmr.(h,c)`.

10.4.1 Functions

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`.

10.5 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 SPTFQMR package can only be used in conjunction with an actual NVECTOR implementation library, such as the NVECTOR_SERIAL or NVECTOR_PARALLEL provided with SUNDIALS.

The files needed to use the SPTFQMR module by itself are the same as for the SPGMR module, with `sundials_sptfqmr.(h,c)` replacing `sundials_spgmr.(h,c)`.



10.5.1 Functions

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`.

Chapter 11

CVODES 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.

11.1 CVODES input constants

CVODES main solver module		
CV_ADAMS	1	Adams-Moulton linear multistep method.
CV_BDF	2	BDF linear multistep method.
CV_FUNCTIONAL	1	Nonlinear system solution through functional iterations.
CV_NEWTON	2	Nonlinear system solution through Newton iterations.
CV_SS	1	Scalar relative tolerance, scalar absolute tolerance.
CV_SV	2	Scalar relative tolerance, vector absolute tolerance.
CV_EE	2	Estimated relative tolerance and absolute tolerance for sensitivity variables.
CV_NORMAL	1	Solver returns at specified output time.
CV_ONE_STEP	2	Solver returns after each successful step.
CV_NORMAL_TSTOP	3	Solver returns at specified output time, but does not proceed past the specified stopping time.
CV_ONE_STEP_TSTOP	4	Solver returns after each successful step, but does not proceed past the specified stopping time.
CV_SIMULTANEOUS	1	Simultaneous corrector forward sensitivity method.
CV_STAGGERED	2	Staggered corrector forward sensitivity method.
CV_STAGGERED1	3	Staggered (variant) corrector forward sensitivity method.
CV_CENTERED	1	Central difference quotient approximation (2^{nd} order) of the sensitivity RHS.
CV_FORWARD	2	Forward difference quotient approximation (1^{st} order) of the sensitivity RHS.
CVODEA adjoint solver module		
CV_HERMITE	1	Use Hermite interpolation.
CV_POLYNOMIAL	2	Use variable-degree polynomial interpolation.

Iterative linear solver module		
PREC_NONE	0	No preconditioning
PREC_LEFT	1	Preconditioning on the left only.
PREC_RIGHT	2	Preconditioning on the right only.
PREC_BOTH	3	Preconditioning on both the left and the right.
MODIFIED_GS	1	Use modified Gram-Schmidt procedure.
CLASSICAL_GS	2	Use classical Gram-Schmidt procedure.

11.2 CVODES output constants

CVODES main solver module		
CV_SUCCESS	0	Successful function return.
CV_TSTOP_RETURN	1	CVode succeeded by reaching the specified stopping point.
CV_ROOT_RETURN	2	CVode succeeded and found one or more roots.
CV_TOO_MUCH_WORK	-1	The solver took <code>mxstep</code> internal steps but could not reach tout.
CV_TOO_MUCH_ACC	-2	The solver could not satisfy the accuracy demanded by the user for some internal step.
CV_ERR_FAILURE	-3	Error test failures occurred too many times during one internal time step or minimum step size was reached.
CV_CONV_FAILURE	-4	Convergence test failures occurred too many times during one internal time step or minimum step size was reached.
CV_LINIT_FAIL	-5	The linear solver's initialization function failed.
CV_LSETUP_FAIL	-6	The linear solver's setup function failed in an unrecoverable manner.
CV_LSOLVE_FAIL	-7	The linear solver's solve function failed in an unrecoverable manner.
CV_RHSFUNC_FAIL	-8	The right-hand side function failed in an unrecoverable manner.
CV_FIRST_RHSFUNC_ERR	-9	The right-hand side function failed at the first call.
CV_REPTD_RHSFUNC_ERR	-10	The right-hand side function had repeated recoverable errors.
CV_UNREC_RHSFUNC_ERR	-11	The right-hand side function had a recoverable error, but no recovery is possible.
CV_RTFUNC_FAIL	-12	The rootfinding function failed in an unrecoverable manner.
CV_MEM_FAIL	-20	A memory allocation failed.
CV_MEM_NULL	-21	The <code>cvode_mem</code> argument was NULL.
CV_ILL_INPUT	-22	One of the function inputs is illegal.
CV_NO_MALLOC	-23	The CVODE memory block was not allocated by a call to <code>CVodeMalloc</code> .
CV_BAD_K	-24	The derivative order k is larger than the order used.
CV_BAD_T	-25	The time t is outside the last step taken.
CV_BAD_DKY	-26	The output derivative vector is NULL.
CV_TOO_CLOSE	-27	The output and initial times are too close to each other.
CV_NO_QUAD	-30	Quadrature integration was not activated.
CV_QRHSFUNC_FAIL	-31	The quadrature right-hand side function failed in an unrecoverable manner.

CV_FIRST_QRHSFUNC_ERR	-32	The quadrature right-hand side function failed at the first call.
CV_REPTD_QRHSFUNC_ERR	-33	The quadrature right-hand side function had repeated recoverable errors.
CV_UNREC_QRHSFUNC_ERR	-34	The quadrature right-hand side function had a recoverable error, but no recovery is possible.
CV_BAD_IS	-40	The sensitivity index is larger than the number of sensitivities computed.
CV_NO_SENS	-41	Forward sensitivity integration was not activated.
CV_SRHSFUNC_FAIL	-42	The sensitivity right-hand side function failed in an unrecoverable manner.
CV_FIRST_SRHSFUNC_ERR	-43	The sensitivity right-hand side function failed at the first call.
CV_REPTD_SRHSFUNC_ERR	-44	The sensitivity right-hand side function had repeated recoverable errors.
CV_UNREC_SRHSFUNC_ERR	-45	The sensitivity right-hand side function had a recoverable error, but no recovery is possible.
<hr/> CVODEA adjoint solver module <hr/>		
CV_ADJMEM_NULL	-101	The <code>cvadj_mem</code> argument was NULL.
CV_BAD_TBO	-103	The final time for the adjoint problem is outside the interval over which the forward problem was solved.
CV_BCKMEM_NULL	-104	The <code>cvodes</code> memory for the backward problem was not created.
CV_REIFWD_FAIL	-105	Reinitialization of the forward problem failed at the first checkpoint.
CV_FWD_FAIL	-106	An error occurred during the integration of the forward problem.
CV_BAD_ITASK	-107	Wrong task for backward integration.
CV_BAD_TBOUT	-108	The desired output time is outside the interval over which the forward problem was solved.
CV_GETY_BADT	-109	Wrong time in interpolation function.
<hr/> CVDENSE linear solver module <hr/>		
CVDENSE_SUCCESS	0	Successful function return.
CVDENSE_MEM_NULL	-1	The <code>cvode_mem</code> argument was NULL.
CVDENSE_LMEM_NULL	-2	The CVDENSE linear solver has not been initialized.
CVDENSE_ILL_INPUT	-3	The CVDENSE solver is not compatible with the current NVECTOR module.
CVDENSE_MEM_FAIL	-4	A memory allocation request failed.
CVDENSE_JACFUNC_UNRECV	-5	The Jacobian function failed in an unrecoverable manner.
CVDENSE_JACFUNC_RECV	-6	The Jacobian function had a recoverable error.
CVDENSE_ADJMEM_NULL	-101	The <code>cvadj_mem</code> argument was NULL.
CVDENSE_LMEMB_NULL	-102	The CVDENSE linear solver has not been initialized for the backward integration.
<hr/> CVBAND linear solver module <hr/>		

CVBAND_SUCCESS	0	Successful function return.
CVBAND_MEM_NULL	-1	The <code>cvode_mem</code> argument was NULL.
CVBAND_LMEM_NULL	-2	The CVBAND linear solver has not been initialized.
CVBAND_ILL_INPUT	-3	The CVBAND solver is not compatible with the current NVECTOR module, or an input value was illegal.
CVBAND_MEM_FAIL	-4	A memory allocation request failed.
CVBAND_JACFUNC_UNRECVR	-5	The Jacobian function failed in an unrecoverable manner.
CVBAND_JACFUNC_RECVR	-6	The Jacobian function had a recoverable error.
CVBAND_ADJMEM_NULL	-101	The <code>cvadj_mem</code> argument was NULL.
CVBAND_LMEMB_NULL	-102	The CVBAND linear solver has not been initialized for the backward integration.

CVDIAG linear solver module		
CVDIAG_SUCCESS	0	Successful function return.
CVDIAG_MEM_NULL	-1	The <code>cvode_mem</code> argument was NULL.
CVDIAG_LMEM_NULL	-2	The CVDIAG linear solver has not been initialized.
CVDIAG_ILL_INPUT	-3	The CVDIAG solver is not compatible with the current NVECTOR module.
CVDIAG_MEM_FAIL	-4	A memory allocation request failed.
CVDIAG_ADJMEM_NULL	-101	The <code>cvadj_mem</code> argument was NULL.

CVSPILS linear solver modules		
CVSPILS_SUCCESS	0	Successful function return.
CVSPILS_MEM_NULL	-1	The <code>cvode_mem</code> argument was NULL.
CVSPILS_LMEM_NULL	-2	The linear solver has not been initialized.
CVSPILS_ILL_INPUT	-3	The solver is not compatible with the current NVECTOR module, or an input value was illegal.
CVSPILS_MEM_FAIL	-4	A memory allocation request failed.
CVSPILS_ADJMEM_NULL	-101	The <code>cvadj_mem</code> argument was NULL.
CVSPILS_LMEMB_NULL	-102	The linear solver has not been initialized for the backward integration.

SPGMR generic linear solver module		
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 function 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 function 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 function 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 function 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 function 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 function failed unrecoverably.
<hr/> CVBANDPRE preconditioner module <hr/>		
CVBANDPRE_SUCCESS	0	Successful function return.
CVBANDPRE_PDATA_NULL	-11	The preconditioner module has not been initialized.
CVBANDPRE_RHSFUNC_UNRECV	-12	The right-hand side function failed unrecoverably.
CVBANDPRE_ADJMEM_NULL	-111	The <code>cvadj_mem</code> argument was NULL.
CVBANDPRE_MEM_FAIL	-112	A memory allocation failed.
<hr/> CVBBDPRE preconditioner module <hr/>		
CVBBDPRE_SUCCESS	0	Successful function return.
CVBBDPRE_PDATA_NULL	-11	The preconditioner module has not been initialized.
CVBBDPRE_FUNC_UNRECV	-12	A user supplied function failed unrecoverably.
CVBBDPRE_ADJMEM_NULL	-111	The <code>cvadj_mem</code> argument was NULL.
CVBBDPRE_PDATAB_NULL	-112	The CVBBDPRE preconditioner module has not been initialized for the backward integration.
CVBBDPRE_MEM_FAIL	-113	A memory allocation failed.

Bibliography

- [1] P. N. Brown, G. D. Byrne, and A. C. Hindmarsh. VODE, a Variable-Coefficient ODE Solver. *SIAM J. Sci. Stat. Comput.*, 10:1038–1051, 1989.
- [2] P. N. Brown and A. C. Hindmarsh. Reduced Storage Matrix Methods in Stiff ODE Systems. *J. Appl. Math. & Comp.*, 31:49–91, 1989.
- [3] 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.
- [4] G. D. Byrne and A. C. Hindmarsh. A Polyalgorithm for the Numerical Solution of Ordinary Differential Equations. *ACM Trans. Math. Softw.*, 1:71–96, 1975.
- [5] 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.
- [6] 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.
- [7] 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.
- [8] S. D. Cohen and A. C. Hindmarsh. CVODE User Guide. Technical Report UCRL-MA-118618, LLNL, September 1994.
- [9] S. D. Cohen and A. C. Hindmarsh. CVODE, a Stiff/Nonstiff ODE Solver in C. *Computers in Physics*, 10(2):138–143, 1996.
- [10] 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.
- [11] R. W. Freund. A Transpose-Free Quasi-Minimal Residual Algorithm for Non-Hermitian Linear Systems. *SIAM J. Sci. Comp.*, 14:470–482, 1993.
- [12] K. L. Hiebert and L. F. Shampine. Implicitly Defined Output Points for Solutions of ODEs. Technical Report SAND80-0180, Sandia National Laboratories, February 1980.
- [13] A. C. Hindmarsh. Detecting Stability Barriers in BDF Solvers. In J.R. Cash and I. Gladwell, editor, *Computational Ordinary Differential Equations*, pages 87–96, Oxford, 1992. Oxford University Press.
- [14] A. C. Hindmarsh. Avoiding BDF Stability Barriers in the MOL Solution of Advection-Dominated Problems. *Appl. Num. Math.*, 17:311–318, 1995.
- [15] A. C. Hindmarsh. The PVODE and IDA Algorithms. Technical Report UCRL-ID-141558, LLNL, December 2000.

-
- [16] 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.
 - [17] A. C. Hindmarsh and R. Serban. Example Programs for CVODE v2.4.0. Technical report, LLNL, 2005. UCRL-SM-208110.
 - [18] A. C. Hindmarsh and R. Serban. Example Programs for CVODES v2.3.0. Technical Report UCRL-SM-208115, LLNL, 2005.
 - [19] A. C. Hindmarsh and R. Serban. User Documentation for CVODE v2.4.0. Technical Report UCRL-SM-208108, LLNL, 2005.
 - [20] 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.
 - [21] K. R. Jackson and R. Sacks-Davis. An Alternative Implementation of Variable Step-Size Multistep Formulas for Stiff ODEs. *ACM Trans. Math. Softw.*, 6:295–318, 1980.
 - [22] 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.
 - [23] T. Maly and L. R. Petzold. Numerical Methods and Software for Sensitivity Analysis of Differential-Algebraic Systems. *Applied Numerical Mathematics*, 20:57–79, 1997.
 - [24] K. Radhakrishnan and A. C. Hindmarsh. Description and Use of LSODE, the Livermore Solver for Ordinary Differential Equations. Technical Report UCRL-ID-113855, LLNL, march 1994.
 - [25] 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.
 - [26] 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.
 - [27] 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

- Adams method, 15
- adjoint sensitivity analysis
 - checkpointing, 23
 - implementation in CVODES, 24, 29
 - mathematical background, 21–24
 - quadrature evaluation, 127
 - right-hand side evaluation, 126
- BAND generic linear solver
 - functions, 159
 - macros, 157–159
 - type BandMat, 157
- BAND_COL, 71, 159
- BAND_COL_ELEM, 71, 159
- BAND_ELEM, 71, 157
- BandMat, 33, 71, 129, 157
- BDF method, 15
- Bi-CGStab method, 39, 51, 122, 160
- BIG_REAL, 32, 143
- CLASSICAL_GS, 51, 121
- CV_ADAMS, 35, 66, 116
- CV_ADJMEM_NULL, 114–117, 119, 123–126
- CV_BAD_DKY, 52, 77, 98–100
- Cv_BAD_DKY, 124
- CV_BAD_IS, 99, 100
- CV_BAD_ITASK, 118
- CV_BAD_K, 52, 77, 99, 100
- CV_BAD_T, 52, 77, 98–100
- CV_BAD_TBO, 116, 117
- CV_BAD_TBOUT, 119
- CV_BCKMEM_NULL, 119
- CV_BDF, 35, 66, 116
- CV_CENTERED, 102
- CV_CONV_FAILURE, 41, 115, 119
- CV_EE, 102
- CV_ERR_FAILURE, 41, 115, 119
- CV_FIRST_QRHSFUNC_ERR, 79
- CV_FIRST_QRHSFUNC_FAIL, 76
- CV_FIRST_RHSFUNC_ERR, 68
- CV_FIRST_RHSFUNC_FAIL, 41
- CV_FIRST_SRHSFUNC_ERR, 108, 109
- CV_FIRST_SRHSFUNC_FAIL, 98
- CV_FORWARD, 102
- CV_FUNCTIONAL, 35, 47, 116
- CV_FWD_FAIL, 119
- CV_HERMITE, 113, 114
- CV_ILL_INPUT, 36, 40, 44, 45, 47, 67, 78, 96, 97, 101–103, 114, 117, 118, 126
- CV_ILL_IPUT, 115
- CV_LINIT_FAIL, 41
- CV_LSETUP_FAIL, 41, 69, 70, 86, 115, 119, 128, 129, 137, 138
- CV_LSOLVE_FAIL, 41, 115
- CV_MEM_FAIL, 36, 75, 96, 97, 114–117, 123
- CV_MEM_NULL, 36, 40, 43–48, 52, 54–59, 67, 75, 77–81, 96–107, 117, 123, 124
- CV_NEWTON, 35, 47, 116
- CV_NO_MALLOC, 40, 67, 115, 117, 118
- CV_NO_QUAD, 75, 77–79, 124
- CV_NO_SENS, 97–101, 103–107
- CV_NORMAL, 40, 112, 115, 118
- CV_NORMAL_TSTOP, 40, 115
- CV_ONE_STEP, 40, 112, 115, 118
- CV_ONE_STEP_TSTOP, 40, 115
- CV_POLYNOMIAL, 113, 114
- CV_QRHSFUNC_FAIL, 76, 79, 127
- CV_REIFWD_FAIL, 119
- CV_REPTD_QRHSFUNC_ERR, 76
- CV_REPTD_RHSFUNC_ERR, 41
- CV_REPTD_SRHSFUNC_ERR, 98
- CV_RHSFUNC_FAIL, 41, 68, 127
- CV_ROOT_RETURN, 40
- CV_RTFUNC_FAIL, 41, 81
- CV_SIMULTANEOUS, 27, 96, 107
- CV_SOLVE_FAIL, 119
- CV_SRHSFUNC_FAIL, 98, 108, 109
- CV_SS, 35, 47, 67, 77, 102, 116, 117
- CV_STAGGERED, 27, 96, 107
- CV_STAGGERED1, 29, 96, 108
- CV_SUCCESS, 36, 40, 43–48, 52, 54–59, 67, 75–81, 96–107, 114–118, 123–126
- CV_SV, 35, 47, 67, 77, 102, 116, 117
- CV_TOO_CLOSE, 41
- CV_TOO_MUCH_ACC, 41, 115, 119
- CV_TOO_MUCH_WORK, 41, 115, 119
- CV_TSTOP_RETURN, 40, 115
- CV_UNREC_QRHSFUNC_ERR, 80
- CV_UNREC_RHSFUNC_ERR, 41, 68, 76
- CV_UNREC_SRHSFUNC_ERR, 98, 108, 109

- CV_WARNING, 68
- CV_WF, 35, 67
- CVadjCheckPointRec, 125
- CVadjFree, 114
- CVadjGetCheckPointsInfo, 124
- CVadjGetCheckPointsList, 125
- CvadjGetCvodeBmem, 123
- CVadjGetDataPointHermite, 125
- CVadjGetDataPointPolynomial, 126
- CVadjGetReturnFlagName, 126
- CVadjMalloc, 112, 113
- CVadjSetInterpType, 114
- CVBAND linear solver
 - Jacobian approximation used by, 49
 - memory requirements, 61
 - NVECTOR compatibility, 38
 - optional input, 49, 120
 - optional output, 61–62
 - selection of, 38
- CVBand, 34, 37, 38, 70
- CVBAND_ADJMEM_NULL, 120
- CVBAND_ILL_INPUT, 38
- CVBAND_JACFUNC_RECVR, 70, 129
- CVBAND_JACFUNC_UNRECVR, 70, 129, 130
- CVBAND_LMEM_NULL, 49, 61, 62, 120
- CVBAND_MEM_FAIL, 38
- CVBAND_MEM_NULL, 38, 49, 61, 62, 120
- CVBAND_SUCCESS, 38, 49, 61, 62, 120
- CVBandB, 129
- CVBandDQJac, 49
- CVBandGetLastFlag, 62
- CVBandGetNumJacEvals, 61
- CVBandGetNumRhsEvals, 61
- CVBandGetReturnFlagName, 62
- CVBandGetWorkSpace, 61
- CVBandJacFn, 70
- CVBANDPRE preconditioner
 - description, 81
 - optional output, 84–85
 - usage, 81–82
 - usage with adjoint module, 132–134
 - user-callable functions, 82–84, 132–134
- CVBANDPRE_ADJMEM_NULL, 132–134
- CVBANDPRE_MEM_FAIL, 132
- CVBANDPRE_PDATA_NULL, 83–85, 133, 134
- CVBANDPRE_SUCCESS, 84, 132
- CVBandPrecAlloc, 82
- CVBandPrecAllocB, 132
- CVBandPrecFree, 84
- CVBandPrecFreeB, 134
- CVBandPrecGetNumRhsEvals, 84
- CVBandPrecGetReturnFlagName, 85
- CVBandPrecGetWorkSpace, 84
- CVBandSetJacFn, 49
- CVBandSetJacFnB, 120
- CVBBDPRE preconditioner
 - description, 85–86
 - optional output, 90–91
 - usage, 87–88
 - usage with adjoint module, 134–138
 - user-callable functions, 88–90, 134–137
 - user-supplied functions, 86–87, 137–138
- CVBBDPRE_ADJMEM_NULL, 135, 136
- CVBBDPRE_PDATA_NULL, 89–91, 135, 136
- CVBBDPRE_SUCCESS, 90, 91, 135, 136
- CVBBDPrecAlloc, 88
- CVBBDPrecAllocB, 134
- CVBBDPrecFree, 90
- CVBBDPrecFreeB, 137
- CVBBDPrecGetNumGfnEvals, 91
- CVBBDPrecGetReturnFlagName, 91
- CVBBDPrecGetWorkSpace, 90
- CVBBDPrecReInit, 90
- CVBBDPrecReInitB, 136
- CVBBDSPbcg, 89
- CVBBDSPbcgB, 135
- CVBBDSPgmr, 88
- CVBBDSPgmrB, 135
- CVBBDSPtfqmr, 89
- CVBBDSPtfqmrB, 136
- CVBPSpbcg, 83
- CVBPSpbcgB, 133
- CVBPSpgmr, 83
- CVBPSpgmrB, 133
- CVBPSptfqmr, 83
- CVBPSptfqmrB, 133
- CVDENSE linear solver
 - Jacobian approximation used by, 48
 - memory requirements, 59
 - NVECTOR compatibility, 37
 - optional input, 48, 119–120
 - optional output, 59–61
 - selection of, 37
- CVDense, 34, 37, 69
- CVDENSE_ADJMEM_NULL, 120
- CVDENSE_ILL_INPUT, 38
- CVDENSE_JACFUNC_RECVR, 69, 128
- CVDENSE_JACFUNC_UNRECVR, 69, 128, 129
- CVDENSE_LMEM_NULL, 48, 59, 60, 120
- CVDENSE_MEM_FAIL, 38
- CVDENSE_MEM_NULL, 38, 48, 59, 60, 120
- CVDENSE_SUCCESS, 38, 48, 59, 60, 120
- CVDenseB, 128
- CVDenseDQJac, 48
- CVDenseGetLastFlag, 60
- CVDenseGetNumJacEvals, 60
- CVDenseGetNumRhsEvals, 60
- CVDenseGetReturnFlagName, 61

- CVDenseGetWorkSpace, 59
- CVDenseJacFn, 69
- CVDenseSetJacFn, 48
- CVDenseSetJacFnB, 119
- CVDIAG linear solver
 - Jacobian approximation used by, 38
 - memory requirements, 62
 - optional output, 62–63
 - selection of, 38
- CVDiag, 34, 37, 38
- CVDIAG_ILL_INPUT, 38
- CVDIAG_LMEM_NULL, 63
- CVDIAG_MEM_FAIL, 38
- CVDIAG_MEM_NULL, 38, 62, 63
- CVDIAG_SUCCESS, 38, 62, 63
- CVDiagGetLastFlag, 63
- CVDiagGetNumRhsEvals, 63
- CVDiagGetReturnFlagName, 63
- CVDiagGetWorkSpace, 62
- CVErrHandlerFn, 68
- CVewtFn, 69
- CVODE, 1
- CVode, 34, 40
- CVodeB, 113, 118
- CVodeCreate, 35
- CVodeCreateB, 112, 116
- CVodeF, 112, 114
- CVodeFree, 34, 36
- CVodeGetActualInitStep, 56
- CVodeGetCurrentOrder, 56
- CVodeGetCurrentStep, 56
- CVodeGetCurrentTime, 57
- CVodeGetDky, 52
- CVodeGetErrWeights, 57
- CVodeGetEstLocalErrors, 58
- CVodeGetIntegratorStats, 58
- CVodeGetLastOrder, 55
- CVodeGetLastStep, 56
- CVodeGetNonlinSolvStats, 59
- CVodeGetNumErrTestFails, 55
- CVodeGetNumGEvals, 81
- CVodeGetNumLinSolvSetups, 55
- CVodeGetNumNonlinSolvConvFails, 59
- CVodeGetNumNonlinSolvIters, 58
- CVodeGetNumRhsEvals, 55
- CVodeGetNumRhsEvalsSEns, 104
- CVodeGetNumSensErrTestFails, 104
- CVodeGetNumSensLinSolvSetups, 104
- CVodeGetNumSensNonlinSolvConvFails, 106
- CVodeGetNumSensNonlinSolvIters, 105
- CVodeGetNumSensRhsEvals, 103
- CVodeGetNumStabLimOrderReds, 57
- CVodeGetNumSteps, 54
- CVodeGetNumStgrSensNonlinSolvConvFails, 107
- CVodeGetNumStgrSensNonlinSolvIters, 106
- CVodeGetQuad, 124
- CVodeGetQuadB, 113
- CVodeGetQuadDky, 76
- CVodeGetQuadErrWeights, 78
- CVodeGetQuadNumErrTestFails, 78
- CVodeGetQuadNumRhsEvals, 78
- CVodeGetQuadStats, 79
- CVodeGetReturnFlagName, 59
- CVodeGetRootInfo, 80
- CVodeGetSens, 95
- CVodeGetSensDky, 98
- CVodeGetSensErrWeights, 105
- CVodeGetSensNonlinSolvStats, 106
- CVodeGetSensStats, 105
- CVodeGetTolScaleFactor, 57
- CVodeGetWorkSpace, 54
- CVodeMalloc, 35, 66
- CVodeMallocB, 112, 116, 117
- CVodeQuadMalloc, 75
- CVodeQuadMallocB, 123
- CVodeReInit, 66
- CVodeReInitB, 117
- CVodeRootInit, 80
- CVODES
 - brief description of, 1
 - motivation for writing in C, 2
 - package structure, 27
 - relationship to CVODE, PVODE, 1–2
 - relationship to VODE, VODPK, 1
- CVODES linear solvers
 - built on generic solvers, 37
 - CVBAND, 38
 - CVDENSE, 37
 - CVDIAG, 38
 - CVSPBCG, 39
 - CVSPGMR, 39
 - CVSPTFQMR, 39
 - header files, 32
 - implementation details, 30
 - list of, 29
 - NVECTOR compatibility, 31
 - selecting one, 37
 - usage with adjoint module, 117–118
- cvodes.h, 32
- cvodes_band.h, 33
- cvodes_dense.h, 32
- cvodes_diag.h, 33
- cvodes_spgmr.h, 33
- cvodes_spgmr.h, 33
- cvodes_sptfqmr.h, 33
- CVodeSensFree, 97
- CVodeSensMalloc, 95, 96
- CVodeSensReInit, 96, 97

- CVodeSensToggleOff, 97
- CVodeSetErrFile, 43
- CVodeSetErrHandlerFn, 43
- CVodeSetEwtFn, 48
- CVodeSetFdata, 43
- CVodeSetInitStep, 45
- CVodeSetIterType, 47
- CVodeSetMaxConvFails, 46
- CVodeSetMaxErrTestFails, 46
- CVodeSetMaxHnilWarns, 44
- CVodeSetMaxNonlinIters, 46
- CVodeSetMaxNumSteps, 44
- CVodeSetMaxOrder, 44
- CVodeSetMaxStep, 45
- CVodeSetMinStep, 45
- CVodeSetNonlinConvCoef, 47
- CVodeSetQuadErrCon, 77
- CVodeSetQuadFdata, 77
- CVodeSetSensDQMethod, 101
- CVodeSetSensErrCon, 102
- CVodeSetSensMaxNonlinIters, 103
- CVodeSetSensParams, 101
- CVodeSetSensRhs1Fn, 101
- CVodeSetSensRhsFn, 100
- CVodeSetSensTolerances, 102
- CVodeSetStabLimDet, 44
- CVodeSetStopTime, 46
- CVodeSetTolerances, 47
- CVQuadRhsFn, 75, 79
- CVQuadRhsFnB, 123, 127
- CVRhsFn, 35, 67
- CVRhsFnB, 116, 117, 126
- CVRootFn, 81
- CVSensRhs1Fn, 108
- CVSensRhsFn, 107
- CVSPBCG linear solver
 - Jacobian approximation used by, 49
 - memory requirements, 63
 - optional input, 49–51, 120–122
 - optional output, 63–66
 - preconditioner setup function, 49, 72, 131
 - preconditioner solve function, 49, 72, 130
 - selection of, 39
- CVSpgbcg, 34, 37, 39
- CVSPGMR linear solver
 - Jacobian approximation used by, 49
 - memory requirements, 63
 - optional input, 49–51, 120–122
 - optional output, 63–66
 - preconditioner setup function, 49, 72, 131
 - preconditioner solve function, 49, 72, 130
 - selection of, 39
- CVSPILS_ADJMEM_NULL, 121, 122
- CVSPILS_ILL_INPUT, 39, 40, 50, 51, 83, 84, 89, 121, 122, 133–136
- CVSPILS_LMEM_NULL, 50, 51, 64–66, 121, 122
- CVSPILS_MEM_FAIL, 39, 40, 83, 84, 89, 133–136
- CVSPILS_MEM_NULL, 39, 40, 50, 51, 64–66, 83, 84, 89, 121, 122, 133–136
- CVSPILS_SUCCESS, 39, 40, 50, 51, 64–66, 83, 84, 89, 120–122, 133, 135, 136
- CVSpilsDQJtimes, 49
- CVSpilsGetLastFlag, 66
- CVSpilsGetNumConvFails, 64
- CVSpilsGetNumJtimesEvals, 65
- CVSpilsGetNumLinIters, 64
- CVSpilsGetNumPrecEvals, 64
- CVSpilsGetNumPrecSolves, 65
- CVSpilsGetNumRhsEvals, 65
- CVSpilsGetReturnFlagName, 66
- CVSpilsGetWorkSpace, 64
- CVSpilsJacTimesVecFn, 71
- CVSpilsJacTimesVecFnB, 130
- CVSpilsPrecSetupFn, 72
- CVSpilsPrecSetupFnB, 131
- CVSpilsPrecSolveFn, 72
- CVSpilsPrecSolveFnB, 130
- CVSpilsSetDelt, 51
- CVSpilsSetDeltB, 121
- CVSpilsSetGSType, 51
- CVSpilsSetGSTypeB, 121
- CVSpilsSetJacTimesFn, 50
- CVSpilsSetJacTimesFnB, 121
- CVSpilsSetMax1, 51
- CVSpilsSetMax1B, 122
- CVSpilsSetPreconditioner, 50
- CVSpilsSetPrecSolveFnB, 120
- CVSpilsSetPrecType, 50
- CVSpilsSetPrecTypeB, 122
- CVSPTFQMR linear solver
 - Jacobian approximation used by, 49
 - memory requirements, 63
 - optional input, 49–51, 120–122
 - optional output, 63
 - preconditioner setup function, 49, 72, 131
 - preconditioner solve function, 49, 72, 130
 - selection of, 39
- CVSptfqmr, 34, 37, 39
- denaddI, 156
- denalloc, 155
- denallocpiv, 155
- dencopy, 156
- denfree, 156
- denfreepiv, 156
- denGETRF, 155
- denGETRS, 156

- denprint, 156
- denscale, 156
- DENSE generic linear solver
 - functions
 - large matrix, 154–155
 - small matrix, 155–156
 - macros, 154
 - type DenseMat, 154
- DENSE_COL, 70, 154
- DENSE_ELEM, 70, 154
- DenseMat, 32, 69, 129, 154
- denzero, 156

- e_data, 69
- eh_data, 68
- error control
 - order selection, 18
 - sensitivity variables, 20
 - step size selection, 17–18
- error messages, 41, 43
 - user-defined handler, 43, 68

- f_data, 43, 68, 86
- f_dataB, 127, 137, 138
- forward sensitivity analysis
 - absolute tolerance selection, 20–21
 - correction strategies, 19–20, 27, 96, 97
 - mathematical background, 18–21
 - right hand side evaluation, 21
 - right-hand side evaluation, 21, 107–109
- fQ_data, 77, 79
- fQ_dataB, 127
- fS_data, 108, 109

- g_data, 81
- generic linear solvers
 - BAND, 156
 - DENSE, 153
 - SPBCG, 160
 - SPGMR, 159
 - SPTFQMR, 161
 - use in CVODES, 30
- GMRES method, 39, 159
- Gram-Schmidt procedure, 51, 121

- half-bandwidths, 38, 70–71, 82, 88
- header files, 32, 81, 87

- itask, 34, 40, 115
- iter, 35, 47
- itol, 35, 116, 117
- itolQ, 77
- itolS, 102

- Jacobian approximation function
 - band
 - difference quotient, 49
 - user-supplied, 49, 70–71
 - user-supplied (backward), 120, 129
 - dense
 - difference quotient, 48
 - user-supplied, 48, 69–70
 - user-supplied (backward), 119, 128
 - diagonal
 - difference quotient, 38
 - Jacobian times vector
 - difference quotient, 49
 - user-supplied, 50
 - Jacobian-vector product
 - user-supplied, 71–72
 - user-supplied (backward), 121, 130

- linit, 150
- lmm, 35, 66
- LSODE, 1

- maxl, 39, 40, 83, 84, 89
- maxord, 44, 66
- memory requirements
 - CVBAND linear solver, 61
 - CVBANDPRE preconditioner, 84
 - CVBBDPRE preconditioner, 91
 - CVDENSE linear solver, 59
 - CVDIAG linear solver, 62
 - CVODES solver, 75, 96
 - CVODES solver, 54
 - CVSPGMR linear solver, 63
- MODIFIED_GS, 51, 121
- MPI, 4

- N_VCloneEmptyVectorArray, 140
- N_VCloneVectorArray, 140
- N_VCloneVectorArray_Parallel, 146
- N_VCloneVectorArray_Serial, 144
- N_VCloneVectorArrayEmpty_Parallel, 146
- N_VCloneVectorArrayEmpty_Serial, 144
- N_VDestroyVectorArray, 140
- N_VDestroyVectorArray_Parallel, 147
- N_VDestroyVectorArray_Serial, 144
- N_Vector, 32, 139
- N_VMake_Parallel, 146
- N_VMake_Serial, 144
- N_VNew_Parallel, 146
- N_VNew_Serial, 144
- N_VNewEmpty_Parallel, 146
- N_VNewEmpty_Serial, 144
- N_VPrint_Parallel, 147
- N_VPrint_Serial, 144
- nonlinear system
 - definition, 15–16

- Newton convergence test, 17
- Newton iteration, 16–17
- NV_COMM_P, 146
- NV_CONTENT_P, 145
- NV_CONTENT_S, 143
- NV_DATA_P, 145
- NV_DATA_S, 143
- NV_GLOBLENGTH_P, 145
- NV_Ith_P, 146
- NV_Ith_S, 144
- NV_LENGTH_S, 143
- NV_LOCLENGTH_P, 145
- NV_OWN_DATA_P, 145
- NV_OWN_DATA_S, 143
- NVECTOR module, 139
- nvector_parallel.h, 32
- nvector_serial.h, 32
- optional input
 - backward solver, 119
 - band linear solver, 49, 120
 - dense linear solver, 48, 119–120
 - forward sensitivity, 100–103
 - iterative linear solver, 49–51, 120–122
 - quadrature integration, 77–78, 124
 - solver, 43–48
- optional output
 - backward solver, 123
 - band linear solver, 61–62
 - band-block-diagonal preconditioner, 90–91
 - banded preconditioner, 84–85
 - checkpoint information, 124
 - dense linear solver, 59–61
 - diagonal linear solver, 62–63
 - forward sensitivity, 103–107
 - interpolated quadratures, 76
 - interpolated sensitivities, 98
 - interpolated solution, 52
 - interpolation data, 125
 - iterative linear solver, 63–66
 - quadrature integration, 78–79, 124
 - solver, 52–59
- output mode, 18, 40, 115, 118
- partial error control
 - explanation of CVODES behavior, 109
- portability, 32
- PREC_BOTH, 39, 40, 50, 122
- PREC_LEFT, 39, 40, 50, 83, 84, 89, 122, 133, 135, 136
- PREC_NONE, 39, 40, 50, 122
- PREC_RIGHT, 39, 40, 50, 83, 84, 89, 122, 133, 135, 136
- preconditioning
 - advice on, 24, 29
 - band-block diagonal, 85
 - banded, 81
 - setup and solve phases, 29
 - user-supplied, 49–50, 72, 120–121, 130, 131
- pretype, 39, 40, 50, 83, 84, 89
- pretypeB, 122, 133, 135, 136
- PVODE, 1
- RCONST, 32
- realtype, 32
- reinitialization, 66, 117
- right-hand side function, 67
 - backward problem, 126
 - forward sensitivity, 107–109
 - quadrature backward problem, 127
 - quadrature equations, 79
- Rootfinding, 25, 34, 80
- SMALL_REAL, 32
- SPBCG generic linear solver
 - description of, 160
 - functions, 160
- SPGMR generic linear solver
 - description of, 159
 - functions, 160
 - support functions, 160
- SPTFQMR generic linear solver
 - description of, 161
 - functions, 161
- Stability limit detection, 24
- step size bounds, 45
- sundials_nvector.h, 32
- sundials_types.h, 32
- TFQMR method, 40, 51, 122, 161
- tolerances, 16, 35, 36, 47, 69, 77, 102, 116, 117
- UNIT_ROUNDOFF, 32
- User main program
 - Adjoint sensitivity analysis, 111
 - CVBANDPRE usage, 81
 - CVBBDPRE usage, 87
 - forward sensitivity analysis, 93
 - integration of quadratures, 74
 - IVP solution, 33
- VODE, 1
- VODPK, 1
- weighted root-mean-square norm, 16