

User Documentation for CVMODE v2.6.0

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

May 9, 2009



UCRL-SM-208108

DISCLAIMER

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

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

Contents

List of Tables	vii
List of Figures	ix
1 Introduction	1
1.1 Historical Background	1
1.2 Changes from previous versions	2
1.3 Reading this User Guide	3
2 Mathematical Considerations	5
2.1 IVP solution	5
2.2 Preconditioning	9
2.3 BDF stability limit detection	9
2.4 Rootfinding	10
3 Code Organization	13
3.1 SUNDIALS organization	13
3.2 CVODE organization	13
4 Using CVODE for C Applications	17
4.1 Access to library and header files	17
4.2 Data Types	18
4.3 Header files	18
4.4 A skeleton of the user's main program	19
4.5 User-callable functions	21
4.5.1 CVODE initialization and deallocation functions	21
4.5.2 CVODE tolerance specification functions	22
4.5.3 Linear solver specification functions	24
4.5.4 Rootfinding initialization function	27
4.5.5 CVODE solver function	28
4.5.6 Optional input functions	29
4.5.6.1 Main solver optional input functions	30
4.5.6.2 Direct linear solvers optional input functions	35
4.5.6.3 Iterative linear solvers optional input functions	36
4.5.6.4 Rootfinding optional input functions	39
4.5.7 Interpolated output function	39
4.5.8 Optional output functions	40
4.5.8.1 Main solver optional output functions	40
4.5.8.2 Rootfinding optional output functions	47
4.5.8.3 Direct linear solvers optional output functions	48
4.5.8.4 Diagonal linear solver optional output functions	49
4.5.8.5 Iterative linear solvers optional output functions	51
4.5.9 CVODE reinitialization function	54

4.6	User-supplied functions	54
4.6.1	ODE right-hand side	54
4.6.2	Error message handler function	55
4.6.3	Error weight function	55
4.6.4	Rootfinding function	56
4.6.5	Jacobian information (direct method with dense Jacobian)	56
4.6.6	Jacobian information (direct method with banded Jacobian)	57
4.6.7	Jacobian information (matrix-vector product)	59
4.6.8	Preconditioning (linear system solution)	59
4.6.9	Preconditioning (Jacobian data)	60
4.7	Preconditioner modules	61
4.7.1	A serial banded preconditioner module	61
4.7.2	A parallel band-block-diagonal preconditioner module	63
5	FCVODE, an Interface Module for FORTRAN Applications	69
5.1	FCVODE routines	69
5.1.1	Important note on portability	70
5.2	Usage of the FCVODE interface module	71
5.3	FCVODE optional input and output	78
5.4	Usage of the FCVROOT interface to rootfinding	79
5.5	Usage of the FCVBP interface to CVBANDPRE	81
5.6	Usage of the FCVBBD interface to CVBBDPRE	82
6	Description of the NVECTOR module	85
6.1	The NVECTOR_SERIAL implementation	89
6.2	The NVECTOR_PARALLEL implementation	91
6.3	NVECTOR functions used by CVODE	93
7	Providing Alternate Linear Solver Modules	95
7.1	Initialization function	95
7.2	Setup function	96
7.3	Solve function	97
7.4	Memory deallocation function	97
8	Generic Linear Solvers in SUNDIALS	99
8.1	The DLS modules: DENSE and BAND	99
8.1.1	Type DlsMat	100
8.1.2	Accessor macros for the DLS modules	101
8.1.3	Functions in the DENSE module	103
8.1.4	Functions in the BAND module	105
8.2	The SPILS modules: SPGMR, SPBCG, and SPTFQMR	107
8.2.1	The SPGMR module	107
8.2.2	The SPBCG module	108
8.2.3	The SPTFQMR module	108
A	CVODE Installation Procedure	109
A.1	Autotools-based installation	110
A.1.1	Configuration options	111
A.1.2	Configuration examples	114
A.2	CMake-based installation	114
A.2.1	Configuring, building, and installing on Unix-like systems	115
A.2.2	Configuring, building, and installing on Windows	116
A.2.3	Configuration options	116
A.3	Manually building SUNDIALS	119
A.4	Installed libraries and exported header files	120

B CVODE Constants	123
B.1 CVODE input constants	123
B.2 CVODE output constants	123
Bibliography	127
Index	129

List of Tables

4.1	Optional inputs for CVODE, CVDLS, and CVSPILS	30
4.2	Optional outputs from CVODE, CVDLS, CVDIAG, and CVSPILS	41
5.1	Keys for setting FCVODE optional inputs	79
5.2	Description of the FCVODE optional output arrays IOUT and ROUT	80
6.1	Description of the NVECTOR operations	87
6.2	List of vector functions usage by CVODE code modules	94
A.1	SUNDIALS libraries and header files	122

List of Figures

3.1	Organization of the SUNDIALS suite	14
3.2	Overall structure diagram of the CVODE package	15
8.1	Diagram of the storage for a banded matrix of type <code>DlsMat</code>	102

Chapter 1

Introduction

CVODE is part of a software family called SUNDIALS: SUite of Nonlinear and Differential/ALgebraic equation Solvers [14]. This suite consists of CVODE, KINSOL, and IDA, and variants of these with sensitivity analysis capabilities.

1.1 Historical Background

FORTRAN solvers for ODE initial value problems are widespread and heavily used. Two solvers that have been written at LLNL in the past are VODE [1] and VODPK [3]. VODE is a general purpose solver that includes methods for 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 [18]. 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 have been combined in the C-language package CVODE [8].

At present, CVODE contains three Krylov methods that can be used in conjunction with Newton iteration: the GMRES (Generalized Minimal RESidual) [19], Bi-CGStab (Bi-Conjugate Gradient Stabilized) [20], and TFQMR (Transpose-Free Quasi-Minimal Residual) linear iterative methods [9]. 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 CVODE, 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 been 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, 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 minimal impacts on the rest of the solver, resulting in PVODE [6], the parallel variant of CVODE.

Recently, the functionality of CVODE and PVODE has been combined into one single code, simply called CVODE. Development of the new version of CVODE 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 kernels 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 linked into an executable file.

There are several motivations for choosing the C language for CVODE. First, a general movement away from FORTRAN and toward C in scientific computing is apparent. Second, the pointer, structure, and dynamic memory allocation features in C are extremely useful in software of this complexity, with the great variety of method options offered. Finally, we prefer C over C++ for CVODE 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 v2.6.0

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

The user interface has been further refined. Some of the API changes involve: (a) a reorganization of all linear solver modules into two families (besides the existing family of scaled preconditioned iterative linear solvers, the direct solvers, including the new Lapack-based ones, were also organized into a *direct* family); (b) maintaining a single pointer to user data, optionally specified through a **Set**-type function; and (c) a general streamlining of the preconditioner modules distributed with the solver.

Changes in v2.5.0

The main changes in this release involve a rearrangement of the entire SUNDIALS source tree (see §3.1). 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.

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 4). Corresponding additions were made to the FORTRAN interface module FCVODE. At the same time, function type names for Scaled Preconditioned Iterative Linear Solvers were added for the user-supplied Jacobian-times-vector and preconditioner setup and solve functions.

The 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 (**cvode_** and **sundials_**). When using the default installation procedure, the header files are exported under various subdirectories of the target **include** directory. For more details see Appendix A.

Changes in v2.3.0

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

Changes in v2.2.1

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

Changes in v2.2.0

The major changes from the previous version involve a redesign of the user interface across the entire SUNDIALS suite. We have eliminated the mechanism of providing optional inputs and extracting optional statistics from the solver through the `iopt` and `ropt` arrays. Instead, CVODE 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 §4.5.6 and §4.5.8.

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

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

1.3 Reading this User Guide

This user guide is a combination of general usage instructions and specific example programs. We expect that some readers will want to concentrate on the general instructions, while others will refer mostly to the examples, and the organization is intended to accommodate both styles.

There are different possible levels of usage of CVODE. The most casual user, with a small IVP problem only, can get by with reading §2.1, then Chapter 4 through §4.5.5 only, and looking at examples in [15]. In a different direction, a more expert user with an IVP problem may want to (a) use a package preconditioner (§4.7), (b) supply his/her own Jacobian or preconditioner routines (§4.6), (c) do multiple runs of problems of the same size (§4.5.9), (d) supply a new NVECTOR module (Chapter 6), or even (e) supply a different linear solver module (§3.2 and Chapter 7).

The structure of this document is as follows:

- In Chapter 2, we give short descriptions of the numerical methods implemented by CVODE for the solution of initial value problems for systems of ODEs, and continue with short descriptions of preconditioning (§2.2), stability limit detection (§2.3), and rootfinding (§2.4).
- The following chapter describes the structure of the SUNDIALS suite of solvers (§3.1) and the software organization of the CVODE solver (§3.2).
- Chapter 4 is the main usage document for CVODE for C applications. It includes a complete description of the user interface for the integration of ODE initial value problems.
- In Chapter 5, we describe FCVODE, an interface module for the use of CVODE with FORTRAN applications.

- Chapter 6 gives a brief overview of the generic NVECTOR module shared among the various components of SUNDIALS, and details on the two NVECTOR implementations provided with SUNDIALS: a serial implementation (§6.1) and a parallel implementation based on MPI (§6.2).
- Chapter 7 describes the interfaces to the linear solver modules, so that a user can provide his/her own such module.
- Chapter 8 describes in detail the generic linear solvers shared by all SUNDIALS solvers.
- Finally, in the appendices, we provide detailed instructions for the installation of CVODE, within the structure of SUNDIALS (Appendix A), as well as a list of all the constants used for input to and output from CVODE functions (Appendix B).

Finally, the reader should be aware of the following notational conventions in this user guide: program listings and identifiers (such as `CVodeInit`) 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. Usage and installation instructions that constitute important warnings are marked with a triangular symbol in the margin.



Acknowledgments. We wish to acknowledge the contributions to previous versions of the CVODE and PVODE codes and their user guides by Scott D. Cohen [7] and George D. Byrne [5].

Chapter 2

Mathematical Considerations

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

$$\dot{y} = f(t, y), \quad y(t_0) = y_0, \quad (2.1)$$

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. CVODE solves both stiff and nonstiff 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.

2.1 IVP solution

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

$$\sum_{i=0}^{K_1} \alpha_{n,i} y^{n-i} + h_n \sum_{i=0}^{K_2} \beta_{n,i} \dot{y}^{n-i} = 0. \quad (2.2)$$

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 CVODE must choose appropriately one of two multistep methods. For nonstiff problems, CVODE 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, CVODE includes the Backward Differentiation Formulas (BDFs) in so-called fixed-leading coefficient 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 [17].

For either choice of formula, the nonlinear system

$$G(y^n) \equiv y^n - h_n \beta_{n,0} f(t_n, y^n) - a_n = 0, \quad (2.3)$$

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, CVODE offers the choice of either *functional iteration*, suitable only for nonstiff 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

$$M[y^{n(m+1)} - y^{n(m)}] = -G(y^{n(m)}), \quad (2.4)$$

in which

$$M \approx I - \gamma J, \quad J = \partial f / \partial y, \quad \text{and} \quad \gamma = h_n \beta_{n,0}. \quad (2.5)$$

The initial guess for the iteration is a predicted value $y^{n(0)}$ computed explicitly from the available history data.

For the solution of the linear systems within the Newton corrections, CVODE provides several choices, including the option of an user-supplied linear solver module. The linear solver modules distributed with SUNDIALS are organized in two families, a *direct* family comprising direct linear solvers for dense or banded matrices and a *spils* family comprising scaled preconditioned iterative (Krylov) linear solvers. In addition, CVODE also provides a linear solver module which only uses a diagonal approximation of the Jacobian matrix. The methods offered through these modules are as follows:

- dense direct solvers, using either an internal implementation or a Blas/Lapack implementation (serial version only),
- band direct solvers, using either an internal implementation or a Blas/Lapack implementation (serial version only),
- a diagonal approximate Jacobian solver,
- SPGMR, a scaled preconditioned GMRES (Generalized Minimal Residual method) solver without restarts,
- SPBCG, a scaled preconditioned Bi-CGStab (Bi-Conjugate Gradient Stable method) solver, or
- SPTFQMR, a scaled preconditioned TFQMR (Transpose-Free Quasi-Minimal Residual method) solver.

For large stiff systems, where direct methods are not feasible, the combination of a BDF integrator and any of the preconditioned Krylov methods (SPGMR, SPBCG, or SPTFQMR) yields a powerful tool because it combines established methods for stiff integration, nonlinear iteration, and Krylov (linear) iteration with a problem-specific treatment of the dominant source of stiffness, in the form of the user-supplied preconditioner matrix [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, CVODE uses a weighted root-mean-square norm, denoted $\|\cdot\|_{\text{WRMS}}$, for all error-like quantities. The multiplicative weights used are based on the current solution and on the relative and absolute tolerances input by the user, namely

$$W_i = 1/[\text{RTOL} \cdot |y_i| + \text{ATOL}_i]. \quad (2.6)$$

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, or diagonal), the iteration is a Modified Newton iteration, in that 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 (direct cases) 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 or may not involve a reevaluation of J (in M) or of Jacobian data (in P), depending on whether 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 $\bar{\gamma}$ of γ at the last update satisfies $|\gamma/\bar{\gamma} - 1| < 0.2$, or
- a convergence failure occurred that forced a step size reduction.

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 (2.3), 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 (2.6). 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

$$Jv = [f(t, y + \sigma v) - f(t, y)]/\sigma. \quad (2.7)$$

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

A critical part of CVODE — making 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, CVODE returns to the user with a give-up message.

In addition to adjusting the step size to meet the local error test, CVODE periodically adjusts the order, with the goal of maximizing the step size. The integration starts out at order 1 and varies the order 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 in step size or order is done. 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 CVODE described above, as inherited from the solvers VODE and VODPK, are documented in [1, 3, 13]. They are also summarized in [14].

Normally, CVODE 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 CVODE not to integrate past a given stopping point $t = t_{\text{stop}}$.

2.2 Preconditioning

When using a Newton method to solve the nonlinear system (2.3), CVODE 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_LP_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 CVODE 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_LP_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 CVODE are based on approximations to the system Jacobian, $J = \partial f / \partial y$. Since the Newton iteration matrix involved is $M = I - \gamma J$, any approximation \tilde{J} to J yields a matrix that is of potential use as a preconditioner, namely $P = I - \gamma \tilde{J}$. Because the Krylov iteration occurs within a Newton iteration and further also within a time integration, and since each of these iterations has its own test for convergence, the preconditioner may use a very crude approximation, as long as it captures the dominant numerical feature(s) of the system. We have found that the combination of a preconditioner with the Newton-Krylov iteration, using even a fairly poor approximation to the Jacobian, can be surprisingly superior to using the same matrix without Krylov acceleration (i.e., a modified Newton iteration), as well as to using the Newton-Krylov method with no preconditioning.

2.3 BDF stability limit detection

CVODE 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, CVODE 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 in 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, and 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 [11]. The algorithm supplements (but differs greatly from) the existing algorithms in CVODE for choosing step size and order based on estimated local truncation errors. It works directly with history data that is readily available in CVODE. 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 [12], where it works well. The implementation in CVODE has been successfully tested on linear and nonlinear advection-diffusion problems, among others.

This stability limit detection option adds some overhead computational cost to the CVODE 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 CVODE solution with this option turned off appears to take an inordinately large number of steps at orders 3-5 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 the efficiency of the solution.

2.4 Rootfinding

The CVODE solver has been augmented to include a rootfinding feature. This means that, while integrating the Initial Value Problem (2.1), CVODE can also find the roots of a set of user-defined functions $g_i(t, y)$ that depend on t and 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 CVODE. If such a root is desired, the user should reformulate the root function so that it changes sign at the desired root.

The basic scheme used is to check for sign changes of any $g_i(t)$ over each time step taken, and then (when a sign change is found) to home in on the root (or roots) with a modified secant method [10]. In addition, each time g is computed, CVODE 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 , CVODE 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, CVODE stops and reports an error. This way, each time CVODE 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, CVODE has an interval $(t_{lo}, t_{hi}]$ in which roots of the $g_i(t)$ are to be sought, such that t_{hi} is further ahead in the direction of integration, and all $g_i(t_{lo}) \neq 0$. The endpoint t_{hi} is either t_n , the end of the time step last taken, or the next requested output time t_{out} if this comes sooner. The endpoint

t_{lo} is either t_{n-1} , or the last output time t_{out} (if this occurred within the last step), or the last root location (if a root was just located within this step), possibly adjusted slightly toward t_n if an exact zero was found. The algorithm checks g at t_{hi} for zeros and for sign changes in (t_{lo}, t_{hi}) . If no sign changes are found, then either a root is reported (if some $g_i(t_{hi}) = 0$) or we proceed to the next time interval (starting at t_{hi}). If one or more sign changes were found, then a loop is entered to locate the root to within a rather tight tolerance, given by

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

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

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

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

where α a weight parameter. On the first two passes through the loop, α is set to 1, making t_{mid} the secant method value. Thereafter, α is reset according to the side of the subinterval (low vs high, i.e. toward t_{lo} vs toward t_{hi}) in which the sign change was found in the previous two passes. If the two sides were opposite, α is set to 1. If the two sides were the same, α is halved (if on the low side) or doubled (if on the high side). The value of t_{mid} is closer to t_{lo} when $\alpha < 1$ and closer to t_{hi} when $\alpha > 1$. If the above value of t_{mid} is within $\tau/2$ of t_{lo} or t_{hi} , it is adjusted inward, such that its fractional distance from the endpoint (relative to the interval size) is between .1 and .5 (.5 being the midpoint), and the actual distance from the endpoint is at least $\tau/2$.

Chapter 3

Code Organization

3.1 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, SUNDIALS also includes variants of CVODE and IDA with sensitivity analysis capabilities (using either forward or adjoint methods): CVODES and IDAS, respectively.

The various solvers of this family share many subordinate modules. For this reason, it is organized as a family, with a directory structure that exploits that sharing (see Fig. 3.1). The following is a list of the solver packages presently available:

- CVODE, a solver for stiff and nonstiff ODEs $dy/dt = f(t, y)$;
- CVODES, a solver for stiff and nonstiff ODEs with sensitivity analysis capabilities;
- IDA, a solver for differential-algebraic systems $F(t, y, \dot{y}) = 0$;
- IDAS, a solver for differential-algebraic systems with sensitivity analysis capabilities;
- KINSOL, a solver for nonlinear algebraic systems $F(u) = 0$.

3.2 CVODE organization


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

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


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

- CVDENSE: LU factorization and backsolving with dense matrices (using either an internal implementation or Blas/Lapack);
- CVBAND: LU factorization and backsolving with banded matrices (using either an internal implementation or Blas/Lapack);

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



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



(b) Directory structure of the source tree

Figure 3.1: Organization of the SUNDIALS suite




Figure 3.2: Overall structure diagram of the CVODE package. Modules specific to CVODE are distinguished by rounded boxes, while generic solver and auxiliary modules are in rectangular boxes. Note that the direct linear solvers using Lapack implementations are not explicitly represented.

- CVSPGMR: scaled preconditioned GMRES method;
- CVSPBCG: scaled preconditioned Bi-CGStab method;
- CVSPTFQMR: scaled preconditioned TFQMR method.

Additionally, CVODE includes:

- CVDIAG: an internally generated diagonal approximation to the Jacobian;

The set of linear solver modules distributed with CVODE 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 by difference quotients of the product between the Jacobian matrix and a vector of appropriate length. Again, the user has the option of providing a routine for this operation. For the Krylov methods, the preconditioning must be supplied by the user, in two phases: 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 and demonstration programs included with CVODE, offer considerable assistance in building preconditioners.

Each CVODE 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 CVODE module to each of the five associated functions is fixed, thus allowing the central module to be completely independent of the linear system method.

These modules are also decomposed in another way. With the exception of CVDIAG and the modules interfacing to Lapack linear solvers, 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 CVODE 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 CVODE package elsewhere.

CVODE 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 a banded difference-quotient Jacobian-based preconditioner, with corresponding setup and solve 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 CVODE 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 CVODE 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 CVODE memory structure. The reentrancy of CVODE 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 4

Using CVODE for C Applications

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

The listings of the sample programs in the companion document [15] may also be helpful. Those codes may be used as templates (with the removal of some lines used in testing) and are included in the CVODE package.

Users with applications written in FORTRAN77 should see Chapter 5, which describes the FORTRAN/C interface module.

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 CVODE modules can only be used with NVECTOR_SERIAL: CVDENSE, CVBAND (using either the internal or the Lapack implementation) and CVBANDPRE. Also, the preconditioner module CVBBDPRE can only be used with NVECTOR_PARALLEL.

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

4.1 Access to library and header files

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

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

- *libdir/libsundials_cvode.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/cvode*
- *incdir/include/sundials*
- *incdir/include/nvector*

The directories *libdir* and *incdir* are the install library and include directories, resp. For a default installation, these are *instdir/lib* and *instdir/include*, respectively, where *instdir* is the directory where SUNDIALS was installed (see Appendix A).

4.2 Data Types

The `sundials_types.h` file contains the definition of the type `realtype`, which is used by the SUNDIALS solvers for all floating-point data. The type `realtype` can be `float`, `double`, or `long double`, with the default being `double`. The user can change the precision of the SUNDIALS solvers arithmetic at the configuration stage (see §A.1.1).

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 that this usage is consistent with the typedef for `realtype`). Thus, a previously existing piece of ANSI C code can use SUNDIALS without modifying the code to use `realtype`, so long as the SUNDIALS libraries use the correct precision (for details see §A.1.1).

4.3 Header files

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

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

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

The calling program must also include an `NVECTOR` implementation header file (see Chapter 6 for details). For the two `NVECTOR` implementations that are included in the CVODE 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 CVODE are:

- `ccode_dense.h`, which is used with the dense direct linear solver;
- `ccode_band.h`, which is used with the band direct linear solver;

- `cvode_lapack.h`, which is used with Lapack implementations of dense or band direct linear solvers;
- `cvode_diag.h`, which is used with the diagonal linear solver;
- `cvode_spgmr.h`, which is used with the scaled, preconditioned GMRES Krylov linear solver SPGMR;
- `cvode_spgbcs.h`, which is used with the scaled, preconditioned Bi-CGStab Krylov linear solver SPBCG;
- `cvode_sptfqmr.h`, which is used with the scaled, preconditioned TFQMR Krylov solver SPT-FQMR;

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

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

Other headers may be needed, according to the choice of preconditioner, etc. For example, in the `cvDiurnal_kry_p` example (see [15]), preconditioning is done with a block-diagonal matrix. For this, even though the CVSPGMR linear solver is used, the header `sundials_dense.h` is included for access to the underlying generic dense linear solver.

4.4 A skeleton of the user's main program

The following is a skeleton of the user's main program (or calling program) for the integration of 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 CVODE: steps marked [P] correspond to NVECTOR_PARALLEL, while steps marked [S] correspond to NVECTOR_SERIAL.

1. [P] Initialize MPI

Call `MPI_Init(&argc, &argv)` to initialize MPI if used by the user's program. 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 the 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 CVODE object

Call `cvode_mem = CVodeCreate(lmm, iter)` to create the CVODE memory block and to specify the solution method (linear multistep method and nonlinear solver iteration type). `CVodeCreate` returns a pointer to the CVODE memory structure. See §4.5.1 for details.

5. Initialize CVODE solver

Call `CVodeInit(...)` to provide required problem specifications, allocate internal memory for CVODE, and initialize CVODE. `CVodeInit` returns a flag, the value of which indicates either success or an illegal argument value. See §4.5.1 for details.

6. Specify integration tolerances

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

7. Set optional inputs

Call `CVodeSet*` functions to change any optional inputs that control the behavior of CVODE from their default values. See §4.5.6.1 for details.

8. Attach linear solver module

If Newton iteration is chosen, initialize the linear solver module with one of the following calls (for details see §4.5.3):

[S] `ier = CVDense(...);`

[S] `ier = CVBand(...);`

[S] `flag = CVLapackDense(...);`

[S] `flag = CVLapackBand(...);`

`ier = CVDiag(...);`

`ier = CVSpqmr(...);`

`ier = CVSpbcg(...);`

`ier = CVSptfqmr(...);`

9. 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 §4.5.6 for details.

10. Specify rootfinding problem

Optionally, call `CVodeRootInit` to initialize a rootfinding problem to be solved during the integration of the ODE system. See §4.5.4, and see §4.5.6.4 for relevant optional input calls.

11. Advance solution in time

For each point at which output is desired, call `ier = CVode(cvode_mem, tout, yout, &tret, itask)`. Here `itask` specifies the return mode. The vector `y` (which can be the same as the vector `y0` above) will contain $y(t)$. See §4.5.5 for details.

12. Get optional outputs

Call `CV*Get*` functions to obtain optional output. See §4.5.8 for details.

13. 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);
```

14. Free solver memory

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

15. [P] Finalize MPI

Call `MPI_Finalize()` to terminate MPI.

4.5 User-callable functions

This section describes the `CVODE` functions that are called by the user to setup and then solve an IVP. Some of these are required. However, starting with §4.5.6, the functions listed involve optional inputs/outputs or restarting, and those paragraphs may be skipped for a casual use of `CVODE`. In any case, refer to §4.4 for the correct order of these calls.

On an error, each user-callable function returns a negative value and sends an error message to the error handler routine, which prints the message on `stderr` by default. However, the user can set a file as error output or can provide his own error handler function (see §4.5.6.1).

4.5.1 CVODE 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 `CVODE` memory block created and allocated by the first two calls.

CVodeCreate	
Call	<code>cvode_mem = CVodeCreate(lmm, iter);</code>
Description	The function <code>CVodeCreate</code> instantiates a <code>CVODE</code> solver object and specifies the solution method.
Arguments	<p><code>lmm</code> (<code>int</code>) specifies the linear multistep method and may be one of two possible values: <code>CV_ADAMS</code> or <code>CV_BDF</code>.</p> <p><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>.</p> <p>The recommended choices for <code>(lmm, iter)</code> are <code>(CV_ADAMS, CV_FUNCTIONAL)</code> for nonstiff problems and <code>(CV_BDF, CV_NEWTON)</code> for stiff problems.</p>
Return value	If successful, <code>CVodeCreate</code> returns a pointer to the newly created <code>CVODE</code> memory block (of type <code>void *</code>). Otherwise, it returns <code>NULL</code> .

CVodeInit

Call	<code>flag = CVodeInit(cvode_mem, f, t0, y0);</code>
Description	The function <code>CVodeInit</code> provides required problem and solution specifications, allocates internal memory, and initializes CVODE.
Arguments	<p><code>cvode_mem</code> (<code>void *</code>) pointer to the CVODE memory block returned by <code>CVodeCreate</code>.</p> <p><code>f</code> (<code>CVRhsFn</code>) is the C function which computes the right-hand side function f in the ODE. This function has the form <code>f(t, y, ydot, user_data)</code> (for full details see §4.6.1).</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>
Return value	The return value <code>flag</code> (of type <code>int</code>) will be one of the following: <p><code>CV_SUCCESS</code> The call to <code>CVodeInit</code> was successful.</p> <p><code>CV_MEM_NULL</code> The CVODE memory block was not initialized through a previous call to <code>CVodeCreate</code>.</p> <p><code>CV_MEM_FAIL</code> A memory allocation request has failed.</p> <p><code>CV_ILL_INPUT</code> An input argument to <code>CVodeInit</code> has an illegal value.</p>
Notes	If an error occurred, <code>CVodeInit</code> also sends an error message to the error handler function.

CVodeFree

Call	<code>CVodeFree(&cvode_mem);</code>
Description	The function <code>CVodeFree</code> frees the memory allocated by a previous call to <code>CVodeCreate</code> .
Arguments	The argument is the pointer to the CVODE memory block (of type <code>void *</code>).
Return value	The function <code>CVodeFree</code> has no return value.

4.5.2 CVODE tolerance specification functions

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

CVodeSStolerances

Call	<code>flag = CVodeSStolerances(cvode_mem, reltol, abstol);</code>
Description	The function <code>CVodeSStolerances</code> specifies scalar relative and absolute tolerances.
Arguments	<p><code>cvode_mem</code> (<code>void *</code>) pointer to the CVODE memory block returned by <code>CVodeCreate</code>.</p> <p><code>reltol</code> (<code>realtype</code>) is the scalar relative error tolerance.</p> <p><code>abstol</code> (<code>realtype</code>) is the scalar absolute error tolerance.</p>
Return value	The return value <code>flag</code> (of type <code>int</code>) will be one of the following: <p><code>CV_SUCCESS</code> The call to <code>CVodeSStolerances</code> was successful.</p> <p><code>CV_MEM_NULL</code> The CVODE memory block was not initialized through a previous call to <code>CVodeCreate</code>.</p> <p><code>CV_NO_MALLOC</code> The allocation function <code>CVodeInit</code> has not been called.</p> <p><code>CV_ILL_INPUT</code> One of the input tolerances was negative.</p>

CVodeSVtolerances

Call	<code>flag = CVodeSVtolerances(cvode_mem, reltol, abstol);</code>
Description	The function <code>CVodeSVtolerances</code> specifies scalar relative tolerance and vector absolute tolerances.
Arguments	<p><code>cvode_mem</code> (void *) pointer to the CVODE memory block returned by <code>CVodeCreate</code>.</p> <p><code>reltol</code> (realtype) is the scalar relative error tolerance.</p> <p><code>abstol</code> (N_Vector) is the vector of absolute error tolerances.</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>CVodeSVtolerances</code> was successful.</p> <p><code>CV_MEM_NULL</code> The CVODE memory block was not initialized through a previous call to <code>CVodeCreate</code>.</p> <p><code>CV_NO_MALLOC</code> The allocation function <code>CVodeInit</code> has not been called.</p> <p><code>CV_ILL_INPUT</code> The relative error tolerance was negative or the absolute tolerance had a negative component.</p>
Notes	This choice of tolerances is important when the absolute error tolerance needs to be different for each component of the state vector y .

CVodeWftolerances

Call	<code>flag = CVodeWftolerances(cvode_mem, efun);</code>
Description	The function <code>CVodeWftolerances</code> specifies a user-supplied function <code>efun</code> that sets the multiplicative error weights W_i for use in the weighted RMS norm, which are normally defined by Eq. (2.6).
Arguments	<p><code>cvode_mem</code> (void *) pointer to the CVODE memory block returned by <code>CVodeCreate</code>.</p> <p><code>efun</code> (CWEwtFn) is the C function which defines the <code>ewt</code> vector (see §4.6.3).</p>
Return value	<p>The return value <code>flag</code> (of type <code>int</code>) will be one of the following:</p> <p><code>CV_SUCCESS</code> The call to <code>CVodeWftolerances</code> was successful.</p> <p><code>CV_MEM_NULL</code> The CVODE memory block was not initialized through a previous call to <code>CVodeCreate</code>.</p> <p><code>CV_NO_MALLOC</code> The allocation function <code>CVodeInit</code> has not been called.</p>

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

(1) The scalar relative tolerance `reltol` is to be set to control relative errors. So `reltol` = 10^{-4} means that errors are controlled to .01%. We do not recommend using `reltol` larger than 10^{-3} . On the other hand, `reltol` should not be so small that it is comparable to the unit roundoff of the machine arithmetic (generally around $1.0\text{E-}15$).

(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 `cvRoberts_dns` in the CVODE package, and the discussion of it in the CVODE Examples document [15]. 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` = 10^{-6} . But in any case, it is a good idea to do a few experiments with the tolerances to see how the computed solution values vary as tolerances are reduced.

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

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

(2) If output plots or tables are being generated, and it is important to avoid having negative numbers appear there (for the sake of avoiding a long explanation of them, if nothing else), then eliminate them, but only in the context of the output medium. Then the internal values carried by the solver are unaffected. Remember that a small negative value in `y` returned by CVODE, 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 $f(t, y)$.

(4) Positivity and non-negativity constraints on components can be enforced by use of the recoverable error return feature in the user-supplied right-hand side function. However, because this option involves some extra overhead cost, it should only be exercised if the use of absolute tolerances to control the computed values is unsuccessful.

4.5.3 Linear solver specification functions

As previously explained, Newton iteration requires the solution of linear systems of the form (2.4). There are six CVODE linear solvers currently available for this task: CVDENSE, CVBAND, CVDIAG, CVSPGMR, CVSPBCG, and CVSPTFQMR.

The first two linear solvers are direct and derive their names from the type of approximation used for the Jacobian $J = \partial f / \partial y$; CVDENSE and CVBAND work with dense and banded approximations to J , respectively. The SUNDIALS suite includes both internal implementations of these two linear solvers and interfaces to Lapack implementations. Together, these linear solvers are referred to as CVDLS (from Direct Linear Solvers).

The CVDIAG linear solver is also a direct linear solver, but it only uses a diagonal approximation to J .

The last three CVODE 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 §4.5.6 and §4.6.

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 (2.5).

To specify a CVODE linear solver, after the call to `CVodeCreate` but before any calls to `CVode`, the user's program must call one of the functions `CVDense/CVLapackDense`, `CVBand/CVLapackBand`, `CVDiag`, `CVSpgrmr`, `CVSpbcg`, or `CVSptfqmr`, as documented below. The first argument passed to these functions is the CVODE memory pointer returned by `CVodeCreate`. A call to one of these functions links the main CVODE 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 and the Lapack direct solvers, the linear solver module used by CVODE 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 8.

CVDense

Call	<code>flag = CVDense(cvode_mem, N);</code>
Description	<p>The function <code>CVDense</code> selects the CVDENSE linear solver and indicates the use of the internal direct dense linear algebra functions.</p> <p>The user's main program must include the <code>cvode_dense.h</code> header file.</p>
Arguments	<p><code>cvode_mem</code> (void *) pointer to the CVODE memory block.</p> <p><code>N</code> (int) problem dimension.</p>
Return value	<p>The return value <code>flag</code> (of type <code>int</code>) is one of:</p> <p><code>CVDLS_SUCCESS</code> The CVDENSE initialization was successful.</p> <p><code>CVDLS_MEM_NULL</code> The <code>cvode_mem</code> pointer is <code>NULL</code>.</p> <p><code>CVDLS_ILL_INPUT</code> The CVDENSE solver is not compatible with the current NVECTOR module.</p> <p><code>CVDLS_MEM_FAIL</code> A memory allocation request failed.</p>
Notes	The CVDENSE linear solver may not be compatible with the particular implementation of the NVECTOR module. Of the two NVECTOR modules provided with SUNDIALS, only NVECTOR_SERIAL is compatible.

CVLapackDense

Call	<code>flag = CVLapackDense(cvode_mem, N);</code>
Description	<p>The function <code>CVLapackDense</code> selects the CVDENSE linear solver and indicates the use of Lapack functions.</p> <p>The user's main function must include the <code>cvode_lapack.h</code> header file.</p>
Arguments	The input arguments are identical to those of <code>CVDense</code> .
Return value	The values of the returned <code>flag</code> (of type <code>int</code>) are identical to those of <code>CVDense</code> .

CVBand

Call	<code>flag = CVBand(cvode_mem, N, mupper, mlower);</code>
Description	<p>The function <code>CVBand</code> selects the CVBAND linear solver and indicates the use of the internal direct band linear algebra functions.</p> <p>The user's main program must include the <code>cvode_band.h</code> header file.</p>
Arguments	<p><code>cvode_mem</code> (void *) pointer to the CVODE memory block.</p> <p><code>N</code> (int) problem dimension.</p> <p><code>mupper</code> (int) upper half-bandwidth of the problem Jacobian (or of the approximation of it).</p> <p><code>mlower</code> (int) lower half-bandwidth of the problem Jacobian (or of the approximation of it).</p>
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of:

CVDLS_SUCCESS The CVBAND initialization was successful.
CVDLS_MEM_NULL The `cvode_mem` pointer is NULL.
CVDLS_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 $(0 \dots N-1)$.
CVDLS_MEM_FAIL A memory allocation request failed.

Notes The CVBAND linear solver may not be compatible with the 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 $-\text{mlower} \leq j - i \leq \text{mupper}$.

CVLapackBand

Call `flag = CVLapackBand(cvode_mem, N, mupper, mlower);`
Description The function CVLapackBand selects the CVBAND linear solver and indicates the use of Lapack functions.
 The user's main function must include the `cvode_lapack.h` header file.
Arguments The input arguments are identical to those of CVBand.
Return value The values of the returned `flag` (of type `int`) are identical to those of CVBand.

CVDiag

Call `flag = CVDiag(cvode_mem);`
Description The function CVDiag selects the CVDIAG linear solver.
 The user's main function must include the `cvode_diag.h` header file.
Arguments `cvode_mem` (`void *`) pointer to the CVODE 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 CVODE 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 `cvode_spgmr.h` header file.
Arguments `cvode_mem` (`void *`) pointer to the CVODE 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 (2.4).

CVSpbcg

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

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

Arguments

- `cvode_mem` (`void *`) pointer to the CVODE 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 (2.4).

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 `cvode_sptfqmr.h` header file.

Arguments

- `cvode_mem` (`void *`) pointer to the CVODE 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 CVSPTFQMR 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 CVSPTFQMR solver uses a scaled preconditioned TFQMR iterative method to solve the linear system (2.4).

4.5.4 Rootfinding initialization function

While solving the IVP, CVODE has the capability to find the roots of a set of user-defined functions. To activate the root finding algorithm, call the following function:

CVodeRootInit

Call	<code>flag = CVodeRootInit(cvode_mem, nrtfn, g);</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> (<code>void *</code>) pointer to the CVODE memory block returned by <code>CVodeCreate</code>.</p> <p><code>nrtfn</code> (<code>int</code>) is the number of root functions g_i.</p> <p><code>g</code> (<code>CVRootFn</code>) is the C function which defines the <code>nrtfn</code> functions $g_i(t, y)$ whose roots are sought. See §4.6.4 for details.</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 <code>NULL</code>.</p> <p><code>CV_MEM_FAIL</code> A memory allocation failed.</p> <p><code>CV_ILL_INPUT</code> The function <code>g</code> is <code>NULL</code>, but <code>nrtfn > 0</code>.</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 = 0</code> .

4.5.5 CVODE solver function

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

CVode

Call	<code>flag = CVode(cvode_mem, tout, yout, &tret, itask);</code>
Description	The function <code>CVode</code> integrates the ODE over an interval in t .
Arguments	<p><code>cvode_mem</code> (<code>void *</code>) pointer to the CVODE memory block.</p> <p><code>tout</code> (<code>realtype</code>) the next time at which a computed solution is desired.</p> <p><code>yout</code> (<code>N_Vector</code>) the computed solution vector.</p> <p><code>tret</code> (<code>realtype</code>) the time reached by the solver (output).</p> <p><code>itask</code> (<code>int</code>) 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 $y(\text{tout})$. 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.</p>
Return value	<p><code>CVode</code> returns a vector <code>yout</code> and a corresponding independent variable value $t = \text{tret}$, 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 <code>flag</code> (of type <code>int</code>) will be one of the following:</p> <p><code>CV_SUCCESS</code> <code>CVode</code> succeeded and no roots were found.</p> <p><code>CV_TSTOP_RETURN</code> <code>CVode</code> succeeded by reaching the stopping point specified through the optional input function <code>CVodeSetStopTime</code> (see §4.5.6.1).</p> <p><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.</p> <p><code>CV_MEM_NULL</code> The <code>cvode_mem</code> argument was <code>NULL</code>.</p> <p><code>CV_NO_MALLOC</code> The CVODE memory was not allocated by a call to <code>CVodeInit</code>.</p>

CV_ILL_INPUT	One of the inputs to <code>CVode</code> was illegal, or some other input to the solver was either illegal or missing. The latter category includes the following situations: (a) The tolerances have not been set. (b) A component of the error weight vector became zero during internal time-stepping. (c) The linear solver initialization function (called by the user after calling <code>CVodeCreate</code>) failed to set the linear solver-specific <code>lsolve</code> field in <code>cvode_mem</code> . (d) A root of one of the root functions was found both at a point t and also very near t . In any case, the user should see the error message for details.
CV_TOO_CLOSE	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.
CV_TOO_MUCH_WORK	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 = 500</code> .
CV_TOO_MUCH_ACC	The solver could not satisfy the accuracy demanded by the user for some internal step.
CV_ERR_FAILURE	Either error test failures occurred too many times (<code>MXNEF = 7</code>) during one internal time step, or with $ h = h_{min}$.
CV_CONV_FAILURE	Either convergence test failures occurred too many times (<code>MXNCF = 10</code>) during one internal time step, or with $ h = h_{min}$.
CV_LINIT_FAIL	The linear solver's initialization function failed.
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_RHSFUNC_FAIL	The right-hand side function failed in an unrecoverable manner.
CV_FIRST_RHSFUNC_FAIL	The right-hand side function had a recoverable error at the first call.
CV_REPTD_RHSFUNC_ERR	Convergence test failures occurred too many times due to repeated recoverable errors in the right-hand side function. This flag will also be returned if the right-hand side function had repeated recoverable errors during the estimation of an initial step size.
CV_UNREC_RHSFUNC_ERR	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.
CV_RTFUNC_FAIL	The rootfinding function failed.

Notes

The vector `yout` can occupy the same space as the vector `y0` of initial conditions that was passed to `CVodeInit`.

In the `CV_ONE_STEP` mode, `tout` is used only on the first call, and only to get the direction and a rough scale of the independent variable.

All failure return values are negative and so the test `ier < 0` will trap all `CVode` failures.

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

4.5.6 Optional input functions

There are numerous optional input parameters that control the behavior of the CVODE solver. CVODE provides functions that can be used to change these optional input parameters from their default values. Table 4.1 lists all optional input functions in CVODE which are then described in detail in the remainder of this section, beginning with those for the main CVODE solver and continuing with those

Table 4.1: Optional inputs for CVODE, CVDLS, and CVSPILS

Optional input	Function name	Default
CVODE main solver		
Pointer to an error file	CVodeSetErrFile	stderr
Error handler function	CVodeSetErrHandlerFn	internal fn.
User data	CVodeSetUserData	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
Direction of zero-crossing	CVodeSetRootDirection	both
Disable rootfinding warnings	CVodeSetNoInactiveRootWarn	none
CVDLS linear solvers		
Dense Jacobian function	CVDlsSetDenseJacFn	DQ
Band Jacobian function	CVDlsSetBandJacFn	DQ
CVSPILS linear solvers		
Preconditioner functions	CVSpilsSetPreconditioner	NULL, NULL
Jacobian-times-vector function	CVSpilsSetJacTimesVecFn	DQ
Preconditioning type	CVSpilsSetPrecType	none
Ratio between linear and nonlinear tolerances	CVSpilsSetEpsLin	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

for the linear solver modules. Note that the diagonal linear solver module has no optional inputs. For the most casual use of CVODE, the reader can skip to §4.6.

We note that, on an 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.

4.5.6.1 Main solver optional input functions

The calls listed here can be executed in any order. However, if either of the functions `CVodeSetErrFile` or `CVodeSetErrHandlerFn` is to be called, that call should be first, in order to take effect for any later error message.

CVodeSetErrFile

Call `flag = CVodeSetErrFile(cvode_mem, errfp);`

Description	The function <code>CVodeSetErrFile</code> specifies a pointer to the file where all CVODE messages should be directed when the default CVODE error handler function is used.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODE 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 CVODE memory pointer is NULL). This use of <code>CVodeSetErrFile</code> is strongly discouraged. 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.



`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 CVODE memory block. <code>ehfun</code> (CErrorHandlerFn) is the C error handler function (see §4.6.2). <code>eh_data</code> (void *) pointer to user data passed to <code>ehfun</code> every time it is called.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>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	Error messages indicating that the CVODE solver memory is NULL will always be directed to <code>stderr</code> .

`CVodeSetUserData`

Call	<code>flag = CVodeSetUserData(cvode_mem, user_data);</code>
Description	The function <code>CVodeSetUserData</code> specifies the user data block <code>user_data</code> and attaches it to the main CVODE memory block.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODE memory block. <code>user_data</code> (void *) 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>cvode_mem</code> pointer is NULL.
Notes	If specified, the pointer to <code>user_data</code> is passed to all user-supplied functions that have it as an argument. Otherwise, a NULL pointer is passed.

`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 CVODE memory block. <code>maxord</code> (int) value of the maximum method order. This must be positive.

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 specified value `maxord` is ≤ 0 , or larger than its previous value.

Notes The default value is `ADAMS_Q_MAX = 12` for the Adams-Moulton method and `BDF_Q_MAX = 5` for the BDF method. Since `maxord` affects the memory requirements for the internal CVODE memory block, its value cannot be increased past its previous value.

An input value greater than the default will result in the default value.

CVodeSetMaxNumSteps

Call `flag = CVodeSetMaxNumSteps(cvode_mem, mxsteps);`

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

Arguments `cvode_mem` (`void *`) pointer to the CVODE memory block.
`mxsteps` (`long int`) maximum allowed number of steps.

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

CVodeSetMaxHnilWarns

Call `flag = CVodeSetMaxHnilWarns(cvode_mem, mxhnil);`

Description The function `CVodeSetMaxHnilWarns` specifies the maximum number of messages issued by the solver warning that $t + h = t$ on the next internal step.

Arguments `cvode_mem` (`void *`) pointer to the CVODE memory block.
`mxhnil` (`int`) maximum number of warning messages (> 0).

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 The default value is 10. A negative value for `mxhnil` indicates that no warning messages should be issued.

CVodeSetStabLimDet

Call `flag = CVodeSetstabLimDet(cvode_mem, stldet);`

Description The function `CVodeSetStabLimDet` indicates if the BDF stability limit detection algorithm should be used. See §2.3 for further details.

Arguments `cvode_mem` (`void *`) pointer to the CVODE memory block.
`stldet` (`boolean` type) 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 CVODE memory block.
`hin` (`realtype`) value of the initial step size to be attempted. Pass 0.0 to use the default value.

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, CVODE 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 t_0 .

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 CVODE memory block.
`hmin` (`realtype`) minimum absolute value of the step size (≥ 0.0).

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 CVODE memory block.
`hmax` (`realtype`) maximum absolute value of the step size (≥ 0.0).

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 = 0.0` 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 CVODE 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. <code>CV_ILL_INPUT</code> The value of <code>tstop</code> is beyond the current t value, t_n .
Notes	The default, if this routine is not called, is that no stop time is imposed.

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 CVODE memory block. <code>maxnef</code> (int) maximum number of error test failures allowed on one step (> 0).
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 CVODE memory block. <code>maxcor</code> (int) maximum number of nonlinear solver iterations allowed per step (> 0).
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>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 CVODE memory block. <code>maxncf</code> (int) maximum number of allowable nonlinear solver convergence failures per step (> 0).
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.

CV_MEM_NULL The `cvode_mem` pointer is NULL.

Notes The default value is 10.

CVodeSetNonlinConvCoef

Call `flag = CVodeSetNonlinConvCoef(cvode_mem, nlscoef);`

Description The function `CVodeSetNonlinConvCoef` specifies the safety factor used in the nonlinear convergence test (see §2.1).

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

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 The default value is 0.1.

CVodeSetIterType

Call `flag = CVodeSetIterType(cvode_mem, iter);`

Description The function `CVodeSetIterType` resets the nonlinear solver iteration type to `iter`.

Arguments `cvode_mem` (void *) pointer to the CVODE memory block.
`iter` (int) specifies the type of nonlinear solver iteration and may be either `CV_NEWTON` or `CV_FUNCTIONAL`.

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 `iter` value passed is neither `CV_NEWTON` nor `CV_FUNCTIONAL`.

Notes The nonlinear solver iteration type is initially specified in the call to `CVodeCreate` (see §4.5.1). This function call is needed only if `iter` is being changed from its value in the prior call to `CVodeCreate`.

4.5.6.2 Direct linear solvers optional input functions

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

CVDlsSetDenseJacFn

Call `flag = CVDlsSetDenseJacFn(cvode_mem, djac);`

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

Arguments `cvode_mem` (void *) pointer to the CVODE memory block.
`djac` (`CVDlsDenseJacFn`) user-defined dense Jacobian approximation function.

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

	<code>CVDLS_SUCCESS</code>	The optional value has been successfully set.
	<code>CVDLS_MEM_NULL</code>	The <code>cvode_mem</code> pointer is NULL.
	<code>CVDLS_LMEM_NULL</code>	The CVDENSE linear solver has not been initialized.
Notes	By default, CVDENSE uses an internal difference quotient function. If NULL is passed to <code>djac</code> , this default function is used.	
	The function type <code>CVDlsDenseJacFn</code> is described in §4.6.5.	

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

<code>CVDlsSetBandJacFn</code>		
Call	<code>flag = CVDlsSetBandJacFn(cvode_mem, bjac);</code>	
Description	The function <code>CVDlsSetBandJacFn</code> specifies the banded Jacobian approximation function to be used.	
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODE memory block. <code>bjac</code> (CVBandJacFn) user-defined banded Jacobian approximation function.	
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of	
	<code>CVDLS_SUCCESS</code>	The optional value has been successfully set.
	<code>CVDLS_MEM_NULL</code>	The <code>cvode_mem</code> pointer is NULL.
	<code>CVDLS_LMEM_NULL</code>	The CVBAND linear solver has not been initialized.
Notes	By default, CVBAND uses an internal difference quotient function. If NULL is passed to <code>bjac</code> , this default function is used.	
	The function type <code>CVBandJacFn</code> is described in §4.6.6.	

4.5.6.3 Iterative linear solvers optional input functions

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 §4.6. If used, the `psetup` function should also be specified in the call to `CVSpilsSetPreconditioner`.

The pointer `user_data` received through `CVodeSetUserData` (or a pointer to NULL if `user_data` was not specified) is passed to the preconditioner `psetup` and `psolve` functions. This allows the user to create an arbitrary structure with relevant problem data and access it during the execution of the user-supplied preconditioner functions without using global data in the program.

Ther 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 default internal difference quotient function 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 §4.6.7 for specification details). As with the preconditioner user-supplied functions, a pointer to the user-defined data structure, `user_data`, specified through `CVodeSetUserData` (or a NULL pointer otherwise) is passed to the Jacobian-times-vector function `jt看` each time it is called.

CVSpilsSetPreconditioner

Call	<code>flag = CVSpilsSetPreconditioner(cvode_mem, psetup, psolve);</code>
Description	The function <code>CVSpilsSetPreconditioner</code> specifies the preconditioner setup and solve functions.
Arguments	<p><code>cvode_mem</code> (void *) pointer to the CVODE memory block.</p> <p><code>psetup</code> (<code>CVSpilsPrecSetupFn</code>) user-defined preconditioner setup function. Pass NULL if no setup is to be done.</p> <p><code>psolve</code> (<code>CVSpilsPrecSolveFn</code>) user-defined preconditioner solve function.</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 values have 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 §4.6.8. The function type <code>CVSpilsPrecSetupFn</code> is described in §4.6.9.

CVSpilsSetJacTimesVecFn

Call	<code>flag = CVSpilsSetJacTimesVecFn(cvode_mem, jtimes);</code>
Description	The function <code>CVSpilsSetJacTimesFn</code> specifies the Jacobian-vector function to be used.
Arguments	<p><code>cvode_mem</code> (void *) pointer to the CVODE memory block.</p> <p><code>jtimes</code> (<code>CVSpilsJacTimesVecFn</code>) user-defined Jacobian-vector product function.</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. If NULL is passed to <code>jtimes</code>, this default function is used.</p> <p>The function type <code>CVSpilsJacTimesVecFn</code> is described in §4.6.7.</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 CVODE 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 §4.5.3). 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 CVOICE 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 <code>NULL</code> . <code>CVSPILS_LMEM_NULL</code> The CVSPILS linear solver has not been initialized. <code>CVSPILS_ILL_INPUT</code> The value of <code>gstype</code> is not valid.
Notes	The default value is <code>MODIFIED_GS</code> . This option is available only for the CVSPGMR linear solver.

**CVSpilsSetEpsLin**

Call	<code>flag = CVSpilsSetEpsLin(cvode_mem, eplifac);</code>
Description	The function <code>CVSpilsSetEpsLin</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 CVOICE memory block. <code>eplifac</code> (realtype) linear convergence safety factor (≥ 0.0).
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 <code>NULL</code> . <code>CVSPILS_LMEM_NULL</code> The CVSPILS linear solver has not been initialized. <code>CVSPILS_ILL_INPUT</code> The factor <code>eplifac</code> is negative.
Notes	The default value is 0.05. Passing a value <code>eplifac</code> = 0.0 also indicates using the default value.

CVSpilsSetMaxl

Call	<code>flag = CVSpilsSetMaxl(cv_mem, maxl);</code>
Description	The function <code>CVSpilsSetMaxl</code> resets the maximum Krylov subspace dimension for the Bi-CGSTab or TFQMR methods.
Arguments	<code>cv_mem</code> (void *) pointer to the CVOICE 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>cvode_mem</code> pointer is <code>NULL</code> . <code>CVSPILS_LMEM_NULL</code> The CVSPILS linear solver has not been initialized. <code>CVSPILS_ILL_INPUT</code> The current linear solver is SPGMR.
Notes	The maximum subspace dimension is initially specified in the call to the linear solver specification function (see §4.5.3). This function call is needed only if <code>maxl</code> is being changed from its previous value. An input value <code>maxl</code> ≤ 0 will result in the default value, 5. This option is available only for the CVSPBCG and CVSPTFQMR linear solvers.



4.5.6.4 Rootfinding optional input functions

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

CVodeSetRootDirection

Call	<code>flag = CVodeSetRootDirection(cvode_mem, rootdir);</code>
Description	The function <code>CVodeSetRootDirection</code> specifies the direction of zero-crossings to be located and returned.
Arguments	<p><code>cvode_mem</code> (void *) pointer to the CVODE memory block.</p> <p><code>rootdir</code> (int *) state array of length <code>nrtfn</code>, the number of root functions g_i, as specified in the call to the function <code>CVodeRootInit</code>. A value of 0 for <code>rootdir[i]</code> indicates that crossing in either direction for g_i should be reported. A value of +1 or -1 indicates that the solver should report only zero-crossings where g_i is increasing or decreasing, respectively.</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_ILL_INPUT</code> rootfinding has not been activated through a call to <code>CVodeRootInit</code>.</p>
Notes	The default behavior is to monitor for both zero-crossing directions.

CVodeSetNoInactiveRootWarn

Call	<code>flag = CVodeSetNoInactiveRootWarn(cvode_mem);</code>
Description	The function <code>CVodeSetNoInactiveRootWarn</code> disables issuing a warning if some root function appears to be identically zero at the beginning of the integration.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODE memory block.
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	CVODE will not report the initial conditions as a possible zero-crossing (assuming that one or more components g_i are zero at the initial time). However, if it appears that some g_i is identically zero at the initial time (i.e., g_i is zero at the initial time and after the first step), CVODE will issue a warning which can be disabled with this optional input function.

4.5.7 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 CVODE.

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

CVodeGetDky

Call	<code>flag = CVodeGetDky(cvode_mem, t, k, dky);</code>
Description	<p>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 (optional output <code>qlast</code>).</p>

Arguments	<p><code>cvode_mem</code> (void *) pointer to the CVODE 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> <code>k</code> is not in the range $0, 1, \dots, q_u$.</p> <p><code>CV_BAD_T</code> <code>t</code> is not in the interval $[t_n - h_u, t_n]$.</p> <p><code>CV_BAD_DKY</code> The <code>dky</code> argument was <code>NULL</code>.</p> <p><code>CV_MEM_NULL</code> The <code>cvode_mem</code> argument was <code>NULL</code>.</p>
Notes	<p>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.</p>

4.5.8 Optional output functions

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

Some of the optional outputs, especially the various counters, can be very useful in determining how successful the CVODE solver is in doing its job. For example, the counters `nsteps` and `nfevals` provide a rough measure of the overall cost of a given run, and can be compared among runs with differing input options to suggest which set of options is most efficient. The ratio `nniters/nsteps` measures the performance of the Newton iteration in solving the nonlinear systems at each time step; typical values for this range from 1.1 to 1.8. The ratio `njevals/nniters` (in the case of a direct linear solver), and the ratio `npevals/nniters` (in the case of an iterative linear solver) measure the overall degree of nonlinearity in these systems, and also the quality of the approximate Jacobian or preconditioner being used. Thus, for example, `njevals/nniters` can indicate if a user-supplied Jacobian is inaccurate, if this ratio is larger than for the case of the corresponding internal Jacobian. The ratio `nliters/nniters` measures the performance of the Krylov iterative linear solver, and thus (indirectly) the quality of the preconditioner.

4.5.8.1 Main solver optional output functions

CVODE 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 CVODE 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 CVODE nonlinear solver used. As a convenience, additional information extraction functions provide the optional outputs in groups. These optional output functions are described next.

CVodeGetWorkSpace

Call	<code>flag = CVodeGetWorkSpace(cvode_mem, &lenrw, &leniw);</code>
Description	The function <code>CVodeGetWorkSpace</code> returns the CVODE real and integer workspace sizes.
Arguments	<p><code>cvode_mem</code> (void *) pointer to the CVODE memory block.</p> <p><code>lenrw</code> (long int) the number of <code>realtype</code> values in the CVODE workspace.</p> <p><code>leniw</code> (long int) the number of integer values in the CVODE workspace.</p>

Table 4.2: Optional outputs from CVODE, CVDLS, CVDIAG, and CVSPILS

Optional output	Function name
CVODE main solver	
Size of CVODE 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 CVODE integrator statistics	CVodeGetIntegratorStats
CVODE 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
CVDLS linear solvers	
Size of real and integer workspaces	CVDlsGetWorkSpace
No. of Jacobian evaluations	CVDlsGetNumJacEvals
No. of r.h.s. calls for finite diff. Jacobian evals.	CVDlsGetNumRhsEvals
Last return from a linear solver function	CVDlsGetLastFlag
Name of constant associated with a return flag	CVDlsGetReturnFlagName
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

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`.

Notes In terms of the problem size N , the maximum method order `maxord`, and the number `nrtfn` of root functions (see §4.5.4), the actual size of the real workspace, in `realtype` words, is given by the following:

- base value: `lenrw` = $96 + (\text{maxord}+5) * N_r + 3*\text{nrtfn}$;
- using `CVodeSVtolerances`: `lenrw` = `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: `leniw` = $40 + (\text{maxord}+5) * N_i + \text{nrtfn}$;
- using `CVodeSVtolerances`: `leniw` = `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`, no rootfinding, and without using `CVodeSVtolerances`, these lengths are given roughly by:

- For the Adams method: `lenrw` = $96 + 17N$ and `leniw` = 57
- For the BDF method: `lenrw` = $96 + 10N$ and `leniw` = 50

CVodeGetNumSteps

Call `flag = CVodeGetNumSteps(cvode_mem, &nsteps);`

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

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

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`.

CVodeGetNumRhsEvals

Call `flag = CVodeGetNumRhsEvals(cvode_mem, &nfevals);`

Description The function `CVodeGetNumRhsEvals` returns the number of calls to the user's right-hand side function.

Arguments `cvode_mem` (`void *`) pointer to the CVODE memory block.
`nfevals` (`long int`) number of calls to the user's `f` 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`.

Notes The `nfevals` value returned by `CVodeGetNumRhsEvals` does not account for calls made to `f` by a linear solver or preconditioner module.

CVodeGetNumLinSolvSetups

Call `flag = CVodeGetNumLinSolvSetups(cvode_mem, &nlinsetups);`

Description The function `CVodeGetNumLinSolvSetups` returns the number of calls made to the linear solver's setup function.

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

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

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

CVodeGetNumErrTestFails

Call `flag = CVodeGetNumErrTestFails(cvode_mem, &netfails);`

Description The function `CVodeGetNumErrTestFails` returns the number of local error test failures that have occurred.

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

CVodeGetLastOrder

Call `flag = CVodeGetLastOrder(cvode_mem, &qlast);`

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

Arguments `cvode_mem` (void *) pointer to the CVODE memory block.
 `qlast` (int) method order used 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.

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 CVODE 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 CVODE 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 CVODE 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 CVODE 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 CVODE 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, &tcur);`

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

Arguments `cvode_mem` (void *) pointer to the CVODE memory block.
 `tcur` (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	<code>flag = CVodeGetNumStabLimOrderReds(cvode_mem, &nsired);</code>
Description	The function <code>CVodeGetNumStabLimOrderReds</code> returns the number of order reductions dictated by the BDF stability limit detection algorithm (see §2.3).
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODE memory block. <code>nsired</code> (long int) number of order reductions due to stability limit detection.
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.
Notes	If the stability limit detection algorithm was not initialized (<code>CVodeSetStabLimDet</code> was not called), then <code>nsired = 0</code> .

CVodeGetTolScaleFactor

Call	<code>flag = CVodeGetTolScaleFactor(cvode_mem, &tolsfac);</code>
Description	The function <code>CVodeGetTolScaleFactor</code> 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	<code>cvode_mem</code> (void *) pointer to the CVODE memory block. <code>tolsfac</code> (realtype) suggested scaling factor for user-supplied tolerances.
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.

CVodeGetErrWeights

Call	<code>flag = CVodeGetErrWeights(cvode_mem, eweight);</code>
Description	The function <code>CVodeGetErrWeights</code> returns the solution error weights at the current time. These are the reciprocals of the W_i given by (2.6).
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODE memory block. <code>eweight</code> (N_Vector) solution error weights at the current time.
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.
Notes	The user must allocate memory for <code>eweight</code> .

**CVodeGetEstLocalErrors**

Call	<code>flag = CVodeGetEstLocalErrors(cvode_mem, ele);</code>
Description	The function <code>CVodeGetEstLocalErrors</code> returns the vector of estimated local errors.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODE memory block. <code>ele</code> (N_Vector) estimated local errors.
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.



Notes The user must allocate memory for `ele`.

The values returned in `ele` are valid only if `CVode` returned a non-negative value.

The `ele` vector, together with the `eweight` vector from `CVodeGetErrWeights`, can be used to determine how the various components of the system contributed to the estimated local error test. Specifically, that error test uses the RMS norm of a vector whose components are the products of the components of these two vectors. Thus, for example, if there were recent error test failures, the components causing the failures are those with largest values for the products, denoted loosely as `eweight[i]*ele[i]`.

CVodeGetIntegratorStats

Call `flag = CVodeGetIntegratorStats(cvode_mem, &nsteps, &nfevals, &nlinsetups, &netfails, &qlast, &qcur, &hinused, &hlast, &hcur, &tcur);`

Description The function `CVodeGetIntegratorStats` returns the CVODE integrator statistics as a group.

Arguments

<code>cvode_mem</code>	(void *) pointer to the CVODE memory block.
<code>nsteps</code>	(long int) number of steps taken by CVODE.
<code>nfevals</code>	(long int) number of calls to the user's <code>f</code> function.
<code>nlinsetups</code>	(long int) number of calls made to the linear solver setup function.
<code>netfails</code>	(long int) number of error test failures.
<code>qlast</code>	(int) method order used on the last internal step.
<code>qcur</code>	(int) method order to be used on the next internal step.
<code>hinused</code>	(realtype) actual value of initial step size.
<code>hlast</code>	(realtype) step size taken on the last internal step.
<code>hcur</code>	(realtype) step size to be attempted on the next internal step.
<code>tcur</code>	(realtype) current internal time reached.

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

CVodeGetNumNonlinSolvIters

Call `flag = CVodeGetNumNonlinSolvIters(cvode_mem, &nniters);`

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

Arguments

<code>cvode_mem</code>	(void *) pointer to the CVODE memory block.
<code>nniters</code>	(long int) number of nonlinear iterations performed.

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

CVodeGetNumNonlinSolvConvFails

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

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

Arguments

<code>cvode_mem</code>	(void *) pointer to the CVODE 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 CVODE nonlinear solver statistics as a group.

Arguments **cvode_mem** (void *) pointer to the CVODE 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.

4.5.8.2 Rootfinding optional output functions

There are two optional output functions associated with rootfinding.

CVodeGetRootInfo

Call **flag = CVodeGetRootInfo(cvode_mem, rootsfound);**

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

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

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

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

CV_SUCCESS The optional output values have been successfully set.

CV_MEM_NULL The **cvode_mem** pointer is NULL.

Notes Note that, for the components g_i for which a root was found, the sign of **rootsfound**[i] indicates the direction of zero-crossing. A value of $+1$ indicates that g_i is increasing, while a value of -1 indicates a decreasing g_i .

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



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 <i>g</i> .
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODE memory block. <code>ngevals</code> (long int) number of calls made to the user's function <i>g</i> 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.

4.5.8.3 Direct linear solvers optional output functions

The following optional outputs are available from the CVDLS modules: 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 CVDLS function. Note that, where the name of an output would otherwise conflict with the name of an optional output from the main solver, a suffix LS (for Linear Solver) has been added here (e.g. `lenrwLS`).

CVDlsGetWorkSpace

Call	<code>flag = CVDlsGetWorkSpace(cvode_mem, &lenrwLS, &leniwLS);</code>
Description	The function <code>CVDlsGetWorkSpace</code> returns the sizes of the real and integer workspaces used by a CVDLS linear solver (CVDENSE or CVBAND).
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODE memory block. <code>lenrwLS</code> (long int) the number of <code>realtype</code> values in the CVDLS workspace. <code>leniwLS</code> (long int) the number of integer values in the CVDLS workspace.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CVDLS_SUCCESS</code> The optional output values have been successfully set. <code>CVDLS_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL. <code>CVDLS_LMEM_NULL</code> The CVDLS linear solver has not been initialized.
Notes	For the CVDENSE linear solver, in terms of the problem size <i>N</i> , the actual size of the real workspace is $2N^2$ <code>realtype</code> words, and the actual size of the integer workspace is <i>N</i> integer words. For the CVBAND linear solver, in terms of <i>N</i> and Jacobian half-bandwidths, the actual size of the real workspace is $(2 \text{ mupper} + 3 \text{ mlower} + 2) N$ <code>realtype</code> words, and the actual size of the integer workspace is <i>N</i> integer words.

CVDlsGetNumJacEvals

Call	<code>flag = CVDlsGetNumJacEvals(cvode_mem, &njevals);</code>
Description	The function <code>CVDlsGetNumJacEvals</code> returns the number of calls made to the CVDLS (dense or band) Jacobian approximation function.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODE memory block. <code>njevals</code> (long int) the number of calls to the Jacobian function.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CVDLS_SUCCESS</code> The optional output value has been successfully set. <code>CVDLS_MEM_NULL</code> The <code>cvode_mem</code> pointer is NULL. <code>CVDLS_LMEM_NULL</code> The CVDLS linear solver has not been initialized.

CVDlsGetNumRhsEvals

Call	<code>flag = CVDlsGetNumRhsEvals(cvode_mem, &nfevalsLS);</code>
Description	The function <code>CVDlsGetNumRhsEvals</code> returns the number of calls made to the user-supplied right-hand side function due to the finite difference (dense or band) Jacobian approximation.
Arguments	<code>cvode_mem</code> (<code>void *</code>) pointer to the CVODE memory block. <code>nfevalsLS</code> (<code>long int</code>) the number of calls made to the user-supplied right-hand side function.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CVDLS_SUCCESS</code> The optional output value has been successfully set. <code>CVDLS_MEM_NULL</code> The <code>cvode_mem</code> pointer is <code>NULL</code> . <code>CVDLS_LMEM_NULL</code> The CVDLS linear solver has not been initialized.
Notes	The value <code>nfevalsLS</code> is incremented only if the default internal difference quotient function is used.

CVDlsGetLastFlag

Call	<code>flag = CVDlsGetLastFlag(cvode_mem, &lsflag);</code>
Description	The function <code>CVDlsGetLastFlag</code> returns the last return value from a CVDLS routine.
Arguments	<code>cvode_mem</code> (<code>void *</code>) pointer to the CVODE memory block. <code>lsflag</code> (<code>int</code>) the value of the last return flag from a CVDLS function.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CVDLS_SUCCESS</code> The optional output value has been successfully set. <code>CVDLS_MEM_NULL</code> The <code>cvode_mem</code> pointer is <code>NULL</code> . <code>CVDLS_LMEM_NULL</code> The CVDLS linear solver has not been initialized.
Notes	If the <code>CVDENSE</code> setup function failed (<code>CVode</code> returned <code>CV_LSETUP_FAIL</code>), then the value of <code>lsflag</code> is equal to the column index (numbered from one) at which a zero diagonal element was encountered during the LU factorization of the (dense or banded) Jacobian matrix. For all other failures, <code>lsflag</code> is negative.

CVDlsGetReturnFlagName

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

4.5.8.4 Diagonal linear solver optional output functions

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. Note that, where the name of an output would otherwise conflict with the name of an optional output from the main solver, a suffix `LS` (for Linear Solver) has been added here (e.g. `lenrwLS`).

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 CVOICE 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.
 `CVDIAG_LMEM_NULL` The CVDIAG linear solver has not been initialized.

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

4.5.8.5 Iterative linear solvers optional output functions

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. Note that, where the name of an output would otherwise conflict with the name of an optional output from the main solver, a suffix LS (for Linear Solver) has been added here (e.g. `lenrwLS`).

CVSpilsGetWorkSpace

Call	<code>flag = CVSpilsGetWorkSpace(cvode_mem, &lenrwLS, &leniwLS);</code>
Description	The function <code>CVSpilsGetWorkSpace</code> returns the global sizes of the CVSPGMR real and integer workspaces.
Arguments	<code>cvode_mem</code> (void *) pointer to the CCODE memory block. <code>lenrwLS</code> (long int) the number of <code>realtype</code> values in the CVSPILS workspace. <code>leniwLS</code> (long int) the number of integer values in the CVSPILS workspace.
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	In terms of the problem size N and maximum subspace size <code>maxl</code> , the actual size of the real workspace is roughly: $(\text{maxl}+5) * N + \text{maxl} * (\text{maxl}+4) + 1$ <code>realtype</code> words for CVSPGMR, $9 * N$ <code>realtype</code> words for CVSPBCG, and $11 * N$ <code>realtype</code> words for CVSPTFQMR. In a parallel setting, the above values are global, summed over all processors.

CVSpilsGetNumLinIters

Call	<code>flag = CVSpilsGetNumLinIters(cvode_mem, &n timers);</code>
Description	The function <code>CVSpilsGetNumLinIters</code> returns the cumulative number of linear iterations.
Arguments	<code>cvode_mem</code> (void *) pointer to the CCODE memory block. <code>n timers</code> (long int) the current number of linear iterations.
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.

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 CVOICE 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 CVOICE 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 CVOICE 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 CVOICE 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	<code>flag = CVSpilsGetNumRhsEvals(cvode_mem, &nfevalsLS);</code>
Description	The function <code>CVSpilsGetNumRhsEvals</code> returns the number of calls to the user right-hand side function for finite difference Jacobian-vector product approximation.
Arguments	<code>cvode_mem</code> (<code>void *</code>) pointer to the CVODE memory block. <code>nfevalsLS</code> (<code>long int</code>) the number of calls to the user right-hand side 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 <code>NULL</code> . <code>CVSPILS_LMEM_NULL</code> The CVSPILS linear solver has not been initialized.
Notes	The value <code>nfevalsLS</code> is incremented only if the default <code>CVSpilsDQJtimes</code> 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> (<code>void *</code>) pointer to the CVODE memory block. <code>lsflag</code> (<code>int</code>) 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 <code>NULL</code> . <code>CVSPILS_LMEM_NULL</code> The CVSPILS linear solver has not been initialized.
Notes	<p>If the CVSPILS setup function failed (Cvode 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>.</p> <p>If the CVSPGMR solve function failed (Cvode 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 <code>NULL</code>; <code>SPGMR_ATIMES_FAIL_UNREC</code>, indicating an unrecoverable failure in the $J*v$ function; <code>SPGMR_PSOLVE_FAIL_UNREC</code>, indicating that the preconditioner solve function <code>psolve</code> failed unrecoverably; <code>SPGMR_GS_FAIL</code>, indicating a failure in the Gram-Schmidt procedure; or <code>SPGMR_QRSOL_FAIL</code>, indicating that the matrix R was found to be singular during the QR solve phase.</p> <p>If the CVSPBCG solve function failed (Cvode 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 <code>NULL</code>; <code>SPBCG_ATIMES_FAIL_UNREC</code>, indicating an unrecoverable failure in the $J*v$ function; or <code>SPBCG_PSOLVE_FAIL_UNREC</code>, indicating that the preconditioner solve function <code>psolve</code> failed unrecoverably.</p> <p>If the CVSPTFQMR solve function failed (Cvode 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 <code>NULL</code>; <code>SPTFQMR_ATIMES_FAIL_UNREC</code>, indicating an unrecoverable failure in the $J*v$ function; or <code>SPTFQMR_PSOLVE_FAIL_UNREC</code>, indicating that the preconditioner solve function <code>psolve</code> failed unrecoverably.</p>

CVSpilsGetReturnFlagName

Call	<code>name = CVSpilsGetReturnFlagName(lsflag);</code>
Description	The function <code>CVSpilsGetReturnFlagName</code> returns the name of the CVSPILS constant corresponding to <code>lsflag</code> .

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

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

4.5.9 CVODE reinitialization function

The function `CVodeReInit` reinitializes the main CVODE solver for the solution of a problem, where a prior call to `CVodeInit` been made. The new problem must have the same size as the previous one. `CVodeReInit` performs the same input checking and initializations that `CVodeInit` 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 §4.5.3

CVodeReInit

Call `flag = CVodeReInit(cvode_mem, t0, y0);`

Description The function `CVodeReInit` provides required problem specifications and reinitializes CVODE.

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

`t0` (`realtype`) is the initial value of t .

`y0` (`N_Vector`) is the initial value of y .

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 CVODE memory block was not initialized through a previous call to `CVodeCreate`.

`CV_NO_MALLOC` Memory space for the CVODE memory block was not allocated through a previous call to `CVodeInit`.

`CV_ILL_INPUT` An input argument to `CVodeReInit` has an illegal value.

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

4.6 User-supplied functions

The user-supplied functions consist of one function defining the 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.

4.6.1 ODE right-hand side

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

CVRhsFn

Definition `typedef int (*CVRhsFn)(realtype t, N_Vector y, N_Vector ydot,
void *user_data);`

Purpose This function computes the ODE right-hand side for a given value of the independent variable t and state vector y .

$\sqrt{(1/N) \sum_1^N (W_i \cdot v_i)^2}$. These weights will be used in place of those defined by Eq. (2.6). The function type `CVEwtFn` is defined as follows:

`CVEwtFn`

Definition	<code>typedef int (*CVEwtFn)(N_Vector y, N_Vector ewt, void *user_data);</code>
Purpose	This function computes the WRMS error weights for the vector y .
Arguments	<p>y is the value of the dependent variable vector at which the weight vector is to be computed.</p> <p>ewt is the output vector containing the error weights.</p> <p>$user_data$ is a pointer to user data, the same as the <code>user_data</code> parameter passed to <code>CVodeSetUserData</code>.</p>
Return value	A <code>CVEwtFn</code> function type must return 0 if it successfully set the error weights and -1 otherwise.
Notes	<p>Allocation of memory for <code>ewt</code> is handled within CVODE.</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>



4.6.4 Rootfinding function

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 *user_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	<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>$gout$ is the output array, of length <code>nrtfn</code>, with components $g_i(t, y)$.</p> <p>$user_data$ is a pointer to user data, the same as the <code>user_data</code> parameter passed to <code>CVodeSetUserData</code>.</p>
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 CVODE.

4.6.5 Jacobian information (direct method with dense Jacobian)

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

`CVDlsDenseJacFn`

Definition	<code>typedef (*CVDlsDenseJacFn)(int N, realtype t, N_Vector y, N_Vector fy, DlsMat Jac, void *user_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>t is the current value of the independent variable.</p>

<code>y</code>	is the current value of the dependent variable vector, namely the predicted value of $y(t)$.
<code>fy</code>	is the current value of the vector $f(t, y)$.
<code>Jac</code>	is the output dense Jacobian matrix (of type <code>DlsMat</code>).
<code>user_data</code>	is a pointer to user data, the same as the <code>user_data</code> parameter passed to <code>CVodeSetUserData</code> .
<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 a <code>CVDlsDenseJacFn</code> as temporary storage or work space.
Return value	A <code>CVDlsDenseJacFn</code> should return 0 if successful, a positive value if a recoverable error occurred (in which case <code>CVODE</code> will attempt to correct, while <code>CVDENSE</code> sets <code>last_flag</code> on <code>CVDLS_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>CVDLS_JACFUNC_UNRECVR</code>).
Notes	<p>A user-supplied dense Jacobian function must load the N by N dense matrix <code>Jac</code> with an approximation to the Jacobian matrix $J(t, y)$ at the point (t, y). Only nonzero elements need to be loaded into <code>Jac</code> because <code>Jac</code> is set to the zero matrix before the call to the Jacobian function. The type of <code>Jac</code> is <code>DlsMat</code>.</p> <p>The accessor macros <code>DENSE_ELEM</code> and <code>DENSE_COL</code> allow the user to read and write dense matrix elements without making explicit references to the underlying representation of the <code>DlsMat</code> type. <code>DENSE_ELEM(J, i, j)</code> references the (i, j)-th element of the dense matrix <code>Jac</code> ($i, j = 0 \dots N - 1$). 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 <code>DENSE_ELEM(J, m-1, n-1) = J_{m,n}</code>. Alternatively, <code>DENSE_COL(J, j)</code> returns a pointer to the first element of the j-th column of <code>Jac</code> ($j = 0 \dots N - 1$), 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 <code>col_n = DENSE_COL(J, n-1); col_n[m-1] = J_{m,n}</code>. For large problems, it is more efficient to use <code>DENSE_COL</code> than to use <code>DENSE_ELEM</code>. Note that both of these macros number rows and columns starting from 0.</p> <p>The <code>DlsMat</code> type and accessor macros <code>DENSE_ELEM</code> and <code>DENSE_COL</code> are documented in §8.1.3.</p> <p>If the user's <code>CVDenseJacFn</code> 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 <code>CVodeGet*</code> functions described in §4.5.8.1. The unit roundoff can be accessed as <code>UNIT_ROUNDOFF</code> defined in <code>sundials_types.h</code>.</p>

4.6.6 Jacobian information (direct method with banded Jacobian)

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

`CVDlsBandJacFn`

Definition	<pre>typedef int (*CVBandJacFn)(int N, int mupper, int mlower, realtype t, N_Vector y, N_Vector fy, DlsMat Jac, void *user_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	N	is the problem size.
	mlower	
	mupper	are the lower and upper half-bandwidths of the Jacobian.
	t	is the current value of the independent variable.
	y	is the current value of the dependent variable vector, namely the predicted value of $y(t)$.
	fy	is the current value of the vector $f(t, y)$.
	Jac	is the output band Jacobian matrix (of type <code>DlsMat</code>).
	user_data	is a pointer to user data, the same as the <code>user_data</code> parameter passed to <code>CVodeSetUserData</code> .
	tmp1	
	tmp2	
	tmp3	are pointers to memory allocated for variables of type <code>N_Vector</code> which can be used by <code>CVDlsBandJacFn</code> as temporary storage or work space.

Return value A `CVDlsBandJacFn` function should return 0 if successful, a positive value if a recoverable error occurred (in which case CVODE will attempt to correct, while CVBAND sets `last_flag` on `CVDLS_JACFUNC_RECVR`), or a negative value if it failed unrecoverably (in which case the integration is halted, CVode returns `CV_LSETUP_FAIL` and CVBAND sets `last_flag` on `CVDLS_JACFUNC_UNRECVR`).

Notes A user-supplied band Jacobian function must load the band matrix `Jac` of type `DlsMat` with the elements of the Jacobian $J(t, y)$ at the point (t, y) . Only nonzero elements need to be loaded into `Jac` because `Jac` 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 `DlsMat` type. `BAND_ELEM(J, i, j)` references the (i, j) -th element of the band matrix `Jac`, counting from 0. This macro is meant for use in small problems for which efficiency of access is not a major concern. Thus, in terms of the indices m and n ranging from 1 to N with (m, n) within the band defined by `mupper` and `mlower`, the Jacobian element $J_{m,n}$ can be loaded using the statement `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 `Jac`, 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 `DlsMat`. 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 `DlsMat` type and the accessor macros `BAND_ELEM`, `BAND_COL` and `BAND_COL_ELEM` are documented in §8.1.4.

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 §4.5.8.1. The unit roundoff can be accessed as `UNIT_ROUNDOFF` defined in `sundials_types.h`.

4.6.7 Jacobian information (matrix-vector product)

If one of the Krylov iterative linear solvers SPGMR, SPBCG, or SPTFQMR is selected (CVSp* is called in step 8 of §4.4), the user may provide a function of type CVSpilsJacTimesVecFn in the following form, to compute matrix-vector products Jv . If such a function is not supplied, the default is a difference quotient approximation to these products.

CVSpilsJacTimesVecFn

Definition	<pre>typedef int (*CVSpilsJacTimesVecFn)(N_Vector v, N_Vector Jv, realtype t, N_Vector y, N_Vector fy, void *user_data, N_Vector tmp);</pre>	
Purpose	This function computes the product $Jv = (\partial f / \partial y)v$ (or an approximation to it).	
Arguments	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 $f(t, y)$.
	user_data	is a pointer to user data, the same as the user_data parameter passed to CNodeSetUserData.
	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 CNodeGet* functions described in §4.5.8.1. The unit roundoff can be accessed as UNIT_ROUNDOFF defined in sundials_types.h.	

4.6.8 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. Here P should approximate (at least crudely) the Newton matrix $M = I - \gamma J$, where $J = \partial f / \partial y$. If preconditioning is done on both sides, the product of the two preconditioner matrices should approximate M . This function must be of type CVSpilsPrecSolveFn, defined as follows:

CVSpilsPrecSolveFn

Definition	<pre>typedef int (*CVSpilsPrecSolveFn)(realtype t, N_Vector y, N_Vector fy, N_Vector r, N_Vector z, realtype gamma, realtype delta, int lr, void *user_data, N_Vector tmp);</pre>	
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 §4.5.8.1).
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);
user_data	is a pointer to user data, the same as the user_data parameter passed to the function CVodeSetUserData .
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).

4.6.9 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 *user_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 CVSpilsPrecSetupFn are as follows:</p> <table> <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, namely the predicted value of $y(t)$.</td></tr> <tr> <td>fy</td><td>is the current value of the vector $f(t, y)$.</td></tr> <tr> <td>jok</td><td>is an input flag indicating whether the Jacobian-related data needs to be updated. The jok argument provides for the reuse of Jacobian data in the preconditioner solve function. jok = FALSE means that the Jacobian-related data must be recomputed from scratch. jok = TRUE means that the Jacobian data, if saved from the previous call to this function, can be reused (with the current value of gamma). A call with jok = TRUE can only occur after a call with jok = FALSE.</td></tr> <tr> <td>jcurPtr</td><td>is a pointer to a flag which should be set to TRUE if Jacobian data was recomputed, or set to FALSE if Jacobian data was not recomputed, but saved data was still reused.</td></tr> <tr> <td>gamma</td><td>is the scalar γ appearing in the Newton matrix $M = I - \gamma J$.</td></tr> <tr> <td>user_data</td><td>is a pointer to user data, the same as the user_data parameter passed to the function CVodeSetUserData.</td></tr> <tr> <td>tmp1</td><td></td></tr> <tr> <td>tmp2</td><td></td></tr> <tr> <td>tmp3</td><td>are pointers to memory allocated for variables of type N_Vector which can be used by CVSpilsPrecSetupFn as temporary storage or work space.</td></tr> </table>	t	is the current value of the independent variable.	y	is the current value of the dependent variable vector, namely the predicted value of $y(t)$.	fy	is the current value of the vector $f(t, y)$.	jok	is an input flag indicating whether the Jacobian-related data needs to be updated. The jok argument provides for the reuse of Jacobian data in the preconditioner solve function. jok = FALSE means that the Jacobian-related data must be recomputed from scratch. jok = TRUE means that the Jacobian data, if saved from the previous call to this function, can be reused (with the current value of gamma). A call with jok = TRUE can only occur after a call with jok = FALSE .	jcurPtr	is a pointer to a flag which should be set to TRUE if Jacobian data was recomputed, or set to FALSE if Jacobian data was not recomputed, but saved data was still reused.	gamma	is the scalar γ appearing in the Newton matrix $M = I - \gamma J$.	user_data	is a pointer to user data, the same as the user_data parameter passed to the function CVodeSetUserData .	tmp1		tmp2		tmp3	are pointers to memory allocated for variables of type N_Vector which can be used by CVSpilsPrecSetupFn as temporary storage or work space.
t	is the current value of the independent variable.																				
y	is the current value of the dependent variable vector, namely the predicted value of $y(t)$.																				
fy	is the current value of the vector $f(t, y)$.																				
jok	is an input flag indicating whether the Jacobian-related data needs to be updated. The jok argument provides for the reuse of Jacobian data in the preconditioner solve function. jok = FALSE means that the Jacobian-related data must be recomputed from scratch. jok = TRUE means that the Jacobian data, if saved from the previous call to this function, can be reused (with the current value of gamma). A call with jok = TRUE can only occur after a call with jok = FALSE .																				
jcurPtr	is a pointer to a flag which should be set to TRUE if Jacobian data was recomputed, or set to FALSE if Jacobian data was not recomputed, but saved data was still reused.																				
gamma	is the scalar γ appearing in the Newton matrix $M = I - \gamma J$.																				
user_data	is a pointer to user data, the same as the user_data parameter passed to the function CVodeSetUserData .																				
tmp1																					
tmp2																					
tmp3	are pointers to memory allocated for variables of type N_Vector which can be used by CVSpilsPrecSetupFn as temporary storage or work space.																				

Return value	The value to be returned by the preconditioner setup function is a flag indicating whether it was successful. This value should be 0 if successful, positive for a recoverable error (in which case the step will be retried), 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 (\mathbf{t}, \mathbf{y}) 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 §4.5.8.1. The unit roundoff can be accessed as <code>UNIT_ROUNDOFF</code> defined in <code>sundials_types.h</code>.</p>

4.7 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, CVODE provides a banded preconditioner in the module CVBANDPRE and a band-block-diagonal preconditioner module CVBBDPRE.

4.7.1 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 §4.3), to use the CVBANDPRE module, the main program must include the header file `cvode_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 §4.4 are grayed out.

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

- (a) `flag = CVSpqr(cvode_mem, pretype, maxl);`
- (b) `flag = CVSpbcg(cvode_mem, pretype, maxl);`
- (c) `flag = CVSptfqmr(cvode_mem, pretype, maxl);`

7. Initialize the CVBANDPRE preconditioner module

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

```
flag = CVBandPrecInit(cvode_mem, N, mu, m1);
```

to allocate memory and initialize the internal preconditioner data.

8. Set linear solver optional inputs

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

9. Advance solution in time

10. Get optional outputs

Additional optional outputs associated with CVBANDPRE are available by way of two routines described below, `CVBandPrecGetWorkSpace` and `CVBandPrecGetNumRhsEvals`.

11. Deallocate memory for solution vector

12. Free solver memory

The CVBANDPRE preconditioner module is initialized and attached by calling the following function:

CVBandPrecInit

Call	<code>flag = CVBandPrecInit(cvode_mem, N, mu, m1);</code>
Description	The function <code>CVBandPrecInit</code> initializes the CVBANDPRE preconditioner and allocates required (internal) memory for it.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODE memory block. <code>N</code> (int) problem dimension. <code>mu</code> (int) upper half-bandwidth of the Jacobian approximation. <code>m1</code> (int) lower half-bandwidth of the Jacobian approximation.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of <code>CVSPILS_SUCCESS</code> The call to <code>CVBandPrecInit</code> was successful. <code>CVSPILS_MEM_NULL</code> The <code>cvode_mem</code> pointer was NULL. <code>CVSPILS_MEM_FAIL</code> A memory allocation request has failed. <code>CVSPILS_LMEM_NULL</code> A CVSPILS linear solver memory was not attached. <code>CVSPILS_ILL_INPUT</code> The supplied vector implementation was not compatible with block band preconditioner.
Notes	The banded approximate Jacobian will have nonzero elements only in locations (i, j) with $-m1 \leq j - i \leq mu$.

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

CVBandPrecGetWorkSpace

Call	<code>flag = CVBandPrecGetWorkSpace(cvode_mem, &lenrwp, &leniwBP);</code>
Description	The function <code>CVBandPrecGetWorkSpace</code> returns the sizes of the CVBANDPRE real and integer workspaces.
Arguments	<code>cvode_mem</code> (void *) pointer to the CVODE memory block.

	<code>lenrwp</code> (long int) the number of <code>realtype</code> values in the CVBANDPRE workspace.
	<code>leniwp</code> (long int) the number of integer values in the CVBANDPRE workspace.
Return value	The return value <code>flag</code> (of type <code>int</code>) is one of: <code>CVSPILS_SUCCESS</code> The optional output values have been successfully set. <code>CVSPILS_PMEM_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 $(2 \text{ml} + \text{mu} + \text{smu} + 2) N$ <code>realtype</code> 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 function <code>CVSpils***GetWorkSpace</code> .

CVBandPrecGetNumRhsEvals

Call	<code>flag = CVBandPrecGetNumRhsEvals(cvode_mem, &nfevalsBP);</code>
Description	The function <code>CVBandPrecGetNumRhsEvals</code> 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	<code>cvode_mem</code> (void *) pointer to the CVODE memory block. <code>nfevalsBP</code> (long int) the number of calls to the user right-hand side 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_PMEM_NULL</code> The CVBANDPRE preconditioner has not been initialized.
Notes	The counter <code>nfevalsBP</code> is distinct from the counter <code>nfevalsLS</code> returned by the corresponding function <code>CVSpils***GetNumRhsEvals</code> , and also from <code>nfevals</code> , returned by <code>CVodeGetNumRhsEvals</code> . The total number of right-hand side function evaluations is the sum of all three of these counters.

4.7.2 A parallel band-block-diagonal preconditioner module

A principal reason for using a parallel ODE solver such as CVODE lies in the solution of partial differential equations (PDEs). Moreover, the use of a Krylov iterative method for the solution of many such problems is motivated by the nature of the underlying linear system of equations (2.4) 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 [16] and is included in a software module within the CVODE 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 (2.1). 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

$$g(t, y) = [g_1(t, \bar{y}_1), g_2(t, \bar{y}_2), \dots, g_M(t, \bar{y}_M)]^T \quad (4.1)$$

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

$$P = \text{diag}[P_1, P_2, \dots, P_M] \quad (4.2)$$

where

$$P_m \approx I - \gamma J_m \quad (4.3)$$

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 `ml dq` defined as the number of non-zero diagonals above and below the main diagonal, respectively. The difference quotient approximation is computed using `mudq` + `ml dq` + 2 evaluations of g_m , but only a matrix of bandwidth `mukeep` + `mlkeep` + 1 is retained. Neither pair of parameters need be the true half-bandwidths of the Jacobian of the local block of g , if smaller values provide a more efficient preconditioner. The solution of the complete linear system

$$Px = b \quad (4.4)$$

reduces to solving each of the equations

$$P_m x_m = b_m \quad (4.5)$$

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 `cfn` (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 `user_data` that is passed by the user to `CVodeSetUserData` and that was passed to the user's function `f`. The user is responsible for providing space (presumably within `user_data`) for components of y that are communicated between processes by `cfn`, and that are then used by `gloc`, which should not do any communication.

CVLocalFn

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

CVCommFn

Definition	<code>typedef int (*CVCommFn)(int Nlocal, realtype t, N_Vector y, void *user_data);</code>
Purpose	This <code>cfn</code> function performs all interprocess communication necessary for the execution of the <code>gloc</code> function above, using the input vector <code>y</code> .
Arguments	<code>Nlocal</code> is the local vector length. <code>t</code> is the value of the independent variable. <code>y</code> is the dependent variable. <code>user_data</code> is a pointer to user data, the same as the <code>user_data</code> parameter passed to <code>CvodeSetUserData</code> .
Return value	A <code>CVCommFn</code> should return 0 if successful, a positive value if a recoverable error occurred (in which case <code>CVODE</code> will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted and <code>Cvode</code> returns <code>CV_LSETUP_FAIL</code>).
Notes	The <code>cfn</code> function is expected to save communicated data in space defined within the data structure <code>user_data</code> . Each call to the <code>cfn</code> function is preceded by a call to the right-hand side function <code>f</code> with the same <code>(t, y)</code> arguments. Thus, <code>cfn</code> can omit any communication done by <code>f</code> if relevant to the evaluation of <code>glocal</code> . If all necessary communication was done in <code>f</code> , then <code>cfn = NULL</code> can be passed in the call to <code>CVBBDPrecInit</code> (see below).

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

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

- (a) `flag = CVSpgrmr(cvode_mem, pretype, maxl);`
- (b) `flag = CVSpbcg(cvode_mem, pretype, maxl);`
- (c) `flag = CVSptfqmr(cvode_mem, pretype, maxl);`

8. Initialize the `CVBBDPRE` preconditioner module

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

```
flag = CVBBDPrecInit(cvode_mem, local_N, mudq, mldq,  
mukeep, mlkeep, dqrely, gloc, cfn);
```

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

9. Set linear solver optional inputs

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

10. Advance solution in time

11. Get optional outputs

Additional optional outputs associated with CVBBDPRE are available by way of two routines described below, CVBBDPrecGetWorkspace and CVBBDPrecGetNumGfnEvals.

12. Deallocate memory for solution vector

13. Free solver memory

14. Finalize MPI

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

CVBBDPrecInit

Call	<code>flag = CVBBDPrecInit(cvode_mem, local_N, mudq, mldq, mukeep, mlkeep, dqrely, gloc, cfn);</code>
Description	The function CVBBDPrecInit initializes and allocates (internal) memory for the CVBBDPRE preconditioner.
Arguments	<p><code>cvode_mem</code> (void *) pointer to the CVODE memory block.</p> <p><code>local_N</code> (int) local vector length.</p> <p><code>mudq</code> (int) upper half-bandwidth to be used in the difference quotient Jacobian approximation.</p> <p><code>mldq</code> (int) lower half-bandwidth to be used in the difference quotient Jacobian approximation.</p> <p><code>mukeep</code> (int) upper half-bandwidth of the retained banded approximate Jacobian block.</p> <p><code>mlkeep</code> (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 $dqrely = \sqrt{\text{unit roundoff}}$, 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	<p>The return value <code>flag</code> (of type <code>int</code>) is one of</p> <p><code>CVSPILS_SUCCESS</code> The call to CVBBDPrecInit was successful.</p> <p><code>CVSPILS_MEM_NULL</code> The <code>cvode_mem</code> pointer was NULL.</p> <p><code>CVSPILS_MEM_FAIL</code> A memory allocation request has failed.</p> <p><code>CVSPILS_LMEM_NULL</code> A CVSPILS linear solver was not attached.</p> <p><code>CVSPILS_ILL_INPUT</code> The supplied vector implementation was not compatible with block band preconditioner.</p>
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 by 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 on every processor.</p>

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 CVODE 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 `CVSpmgr`, `CVSpbcg`, or `CVSptfqmr`, and/or one or more of the corresponding `CVSpils***Set***` functions, must also be made (in the proper order).

CVBBDPrecReInit

Call `flag = CVBBDPrecReInit(cvode_mem, mudq, mldq, dqrely);`

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

Arguments

- `cvode_mem` (void *) pointer to the CVODE memory block.
- `mudq` (int) upper half-bandwidth to be used in the difference quotient Jacobian approximation.
- `mldq` (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`.

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

- `CVSPILS_SUCCESS` The call to `CVBBDPrecReInit` was successful.
- `CVSPILS_MEM_NULL` The `cvode_mem` pointer was NULL.
- `CVSPILS_LMEM_NULL` A CVSPILS linear solver memory was not attached.
- `CVSPILS_PMEM_NULL` The function `CVBBDPrecInit` was not previously called.

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

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

CVBBDPrecGetWorkSpace

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

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

Arguments

- `cvode_mem` (void *) pointer to the CVODE memory block.
- `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

- `CVSPILS_SUCCESS` The optional output value has been successfully set.
- `CVSPILS_MEM_NULL` The `cvode_mem` pointer was NULL.
- `CVSPILS_PMEM_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 function `CVSpils***GetWorkSpace`.

CVBBDPrecGetNumGfnEvals

Call `flag = CVBBDPrecGetNumGfnEvals(cvode_mem, &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 `cvode_mem` (void *) pointer to the CVODE memory block.
 `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

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

`CVSPILS_MEM_NULL` The `cvode_mem` pointer was NULL.

`CVSPILS_PMEM_NULL` The CVBBDPRE preconditioner has not been initialized.

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 CVODE output and `npsolves` and `nfevalsLS` are linear solver optional outputs (see §4.5.8).

Chapter 5

FCVODE, an Interface Module for FORTRAN Applications

The FCVODE interface module is a package of C functions which support the use of the CVODE solver, for the solution of ODE systems $dy/dt = f(t, y)$, in a mixed FORTRAN/C setting. While CVODE is written in C, it is assumed here that the user's calling program and user-supplied problem-defining routines are written in FORTRAN. This package provides the necessary interface to CVODE for both the serial and the parallel NVECTOR implementations.

5.1 FCVODE routines

The user-callable functions, with the corresponding CVODE functions, are as follows:

- Interface to the NVECTOR modules
 - FNVINITS (defined by NVECTOR_SERIAL) interfaces to N_VNewEmpty_Serial.
 - FNVINITP (defined by NVECTOR_PARALLEL) interfaces to N_VNewEmpty_Parallel.
- Interface to the main CVODE module
 - FCVMALLOC interfaces to C_VodeCreate, C_VodeSetUserData, and C_VodeInit, as well as one of C_VodeSStolerances or C_VodeSVtolerances.
 - FCVREINIT interfaces to C_VodeReInit.
 - FCVSETIIN and FCVSETRIN interface to C_VodeSet* functions.
 - FCVEWTSET interfaces to C_VodeWftolerances.
 - FCVODE interfaces to C_Vode, C_VodeGet* functions, and to the optional output functions for the selected linear solver module.
 - FCVDKY interfaces to the interpolated output function C_VodeGetDky.
 - FCVGETERRWEIGHTS interfaces to C_VodeGetErrWeights.
 - FCVGETESTLOCALERR interfaces to C_VodeGetEstLocalErrors.
 - FCVFREE interfaces to C_VodeFree.
- Interface to the linear solver modules
 - FCVDIAG interfaces to CVDiag.
 - FCVDENSE interfaces to CVDense.
 - FCVDENSESETJAC interfaces to CVDlsSetDenseJacFn.
 - FCVLAPACKDENSE interfaces to CVLapackDense.

- FCVLAPACKDENSESETJAC interfaces to CVDlsSetDenseJacFn.
- FCVBAND interfaces to CVBand.
- FCVBANDSETJAC interfaces to CVDlsSetBandJacFn.
- FCVLAPACKBAND interfaces to CVLapackBand.
- FCVLAPACKBANDSETJAC interfaces to CVDlsSetBandJacFn.
- FCVSPGMR interfaces to CVSpgmr and SPGMR optional input functions.
- FCVSPGMRREINIT interfaces to SPGMR optional input functions.
- FCVSPBCG interfaces to CVSpbcg and SPBCG optional input functions.
- FCVSPBCGREINIT interfaces to SPBCG optional input functions.
- FCVSPTFQMR interfaces to CVSp tfqmr and SPTFQMR optional input functions.
- FCVSPTFQMRREINIT interfaces to SPTFQMR optional input functions.
- FCVSPILSSETJAC interfaces to CVSpilsSetJacTimesVecFn.
- FCVSPILSSETPREC interfaces to CVSpilsSetPreconditioner.

The user-supplied functions, each listed with the corresponding internal interface function which calls it (and its type within CVODE), are as follows:

FCVODE routine (FORTRAN, user-supplied)	CVODE function (C, interface)	CVODE type of interface function
FCVFUN	FCVf	CVRhsFn
FCVEWT	FCVEwtSet	CV EwtFn
FCVDJAC	FCVDenseJac	CVDlsDenseJacFn
	FCVLapackDenseJac	CVDlsDenseJacFn
FCVBJAC	FCVBandJac	CVDlsBandJacFn
	FCVLapackBandJac	CVDlsBandJacFn
FCVPSOL	FCVPSol	CVSpilsPrecSolveFn
FCVPSET	FCVPSet	CVSpilsPrecSetupFn
FCVJTIES	FCVJtimes	CVSpilsJacTimesVecFn

In contrast to the case of direct use of CVODE, and of most FORTRAN ODE solvers, the names of all user-supplied routines here are fixed, in order to maximize portability for the resulting mixed-language program.

5.1.1 Important note on portability

In this package, the names of the interface functions, and the names of the FORTRAN user routines called by them, appear as dummy names which are mapped to actual values by a series of definitions in the header files `fcvode.h`, `fcvroot.h`, `fcvbp.h`, and `fcvbbd.h`. By default, those mapping definitions depend in turn on the C macro `F77_FUNC` defined in the header file `sundials.config.h` and decided upon at configuration time (see Appendix A).

The user must also ensure that variables in the user FORTRAN code are declared in a manner consistent with their counterparts in CVODE. All real variables must be declared as `REAL`, `DOUBLE PRECISION`, or perhaps as `REAL*n`, where n denotes the number of bytes, depending on whether CVODE was built in single, double, or extended precision (see Appendix A). Moreover, some of the FORTRAN integer variables must be declared as `INTEGER*4` or `INTEGER*8` according to the C type `long int`. These integer variables include: the array of integer optional outputs (`IOUT`), problem dimensions (`NEQ`, `NLOCAL`, `NGLOBAL`), Jacobian half-bandwidths (`MU`, `ML`, etc.), as well as the array of user integer data, `IPAR`. This is particularly important when using CVODE and the FCVODE package on 64-bit architectures.

5.2 Usage of the FCVODE interface module

The usage of FCVODE requires calls to six or seven interface functions, depending on the method options selected, and one or more user-supplied routines which define the problem to be solved. These function calls and user routines are summarized separately below. Some details are omitted, and the user is referred to the description of the corresponding CVODE functions for information on the arguments of any given user-callable interface routine, or of a given user-supplied function called by an interface function. The usage of FCVODE for rootfinding and with preconditioner modules is described in later subsections.

Steps marked [S] in the instructions below apply to the serial NVECTOR implementation (NVECTOR_SERIAL) only, while those marked [P] apply to NVECTOR_PARALLEL.

1. Right-hand side specification

The user must in all cases supply the following FORTRAN routine

```
SUBROUTINE FCVFUN(T, Y, YDOT, IPAR, RPAR, IER)
  DIMENSION Y(*), YDOT(*), IPAR(*), RPAR(*)
```

It must set the YDOT array to $f(t, y)$, the right-hand side of the ODE system, as function of $T = t$ and the array $Y = y$. The arrays IPAR (of integers) and RPAR (of reals) contain user data and are the same as those passed to FCVMALLOC. IER is an error return flag which should be set to 0 if successful, a positive value if a recoverable error occurred (in which case CVODE will attempt to correct), or a negative value if it failed unrecoverably (in which case the integration is halted).

2. NVECTOR module initialization

[S] To initialize the serial NVECTOR module, the user must make the following call:

```
CALL FNVINITS(KEY, NEQ, IER)
```

where KEY is the solver id (KEY = 1 for CVODE), NEQ is the size of vectors, and IER is a return completion flag which is 0 on success and -1 if a failure occurred.

[P] To initialize the parallel vector module, the user must make the following call:

```
CALL FNVINITP(COMM, KEY, NLOCAL, NGLOBAL, IER)
```

in which the arguments are: COMM = MPI communicator, KEY = 1, NLOCAL = the local size of vectors on this processor, and NGLOBAL = the system size (and the global size of all vectors, equal to the sum of all values of NLOCAL). The return flag IER is set to 0 on a successful return and to -1 otherwise.

If the header file `sundials_config.h` defines `SUNDIALS_MPI_COMM_F2C` to be 1 (meaning the MPI implementation used to build SUNDIALS includes the `MPI_Comm_f2c` function), then COMM can be any valid MPI communicator. Otherwise, `MPI_COMM_WORLD` will be used, so just pass an integer value as a placeholder.



3. Problem specification

To set various problem and solution parameters and allocate internal memory, make the following call:

FCVMALLOC

Call `CALL FCVMALLOC(T0, Y0, METH, ITMETH, IATOL, RTOL, ATOL, & IOUT, ROUT, IPAR, RPAR, IER)`

Description	This function provides required problem and solution specifications, specifies optional inputs, allocates internal memory, and initializes CVODE.
Arguments	<p>T0 is the initial value of t.</p> <p>Y0 is an array of initial conditions.</p> <p>METH specifies the basic integration method: 1 for Adams (nonstiff) or 2 for BDF (stiff).</p> <p>ITMETH specifies the nonlinear iteration method: 1 for functional iteration or 2 for Newton iteration.</p> <p>IATOL specifies the type for absolute tolerance ATOL: 1 for scalar or 2 for array. If IATOL= 3, the arguments RTOL and ATOL are ignored and the user is expected to subsequently call FCVEWTSET and provide the function FCVEWT.</p> <p>RTOL is the relative tolerance (scalar).</p> <p>ATOL is the absolute tolerance (scalar or array).</p> <p>IOUT is an integer array of length 21 for integer optional outputs.</p> <p>ROUT is a real array of length 6 for real optional outputs.</p> <p>IPAR is an integer array of user data which will be passed unmodified to all user-provided routines.</p> <p>RPAR is a real array of user data which will be passed unmodified to all user-provided routines.</p>
Return value	IER is a return completion flag. Values are 0 for successful return and -1 otherwise. See printed message for details in case of failure.
Notes	<p>The user integer data array IPAR must be declared as <code>INTEGER*4</code> or <code>INTEGER*8</code> according to the C type <code>long int</code>.</p> <p>Modifications to the user data arrays IPAR and RPAR inside a user-provided routine will be propagated to all subsequent calls to such routines.</p> <p>The optional outputs associated with the main CVODE integrator are listed in Table 5.2.</p>

As an alternative to providing tolerances in the call to FCVMALLOC, the user may provide a routine to compute the error weights used in the WRMS norm evaluations. If supplied, it must have the following form:

```
SUBROUTINE FCVEWT (Y, EWT, IPAR, RPAR, IER)
  DIMENSION Y(*), EWT(*), IPAR(*), RPAR(*)
```

It must set the positive components of the error weight vector EWT for the calculation of the WRMS norm of Y. On return, set IER = 0 if FCVEWT was successful, and nonzero otherwise. The arrays IPAR (of integers) and RPAR (of reals) contain user data and are the same as those passed to FCVMALLOC.

If the FCVEWT routine is provided, then, following the call to FCVMALOC, the user must make the call:

```
CALL FCVEWTSET (FLAG, IER)
```

with FLAG $\neq 0$ to specify use of the user-supplied error weight routine. The argument IER is an error return flag which is 0 for success or non-zero if an error occurred.

4. Set optional inputs

Call FCVINSETIIN and/or FCVINSETRIN to set desired optional inputs, if any. See §5.3 for details.

5. Linear solver specification

In the case of a stiff system, the implicit BDF method involves the solution of linear systems related to the Jacobian $J = \partial f / \partial y$ of the ODE system. CVODE presently includes eight choices for the

treatment of these systems, and the user of FCVODE must call a routine with a specific name to make the desired choice.

[S] Diagonal approximate Jacobian

This choice is appropriate when the Jacobian can be well approximated by a diagonal matrix. The user must make the call:

```
CALL FCVDIAG(IER)
```

IER is an error return flag set on 0 on success or -1 if a memory failure occurred. There is no additional user-supplied routine. Optional outputs specific to the DIAG case listed in Table 5.2.

[S] Dense treatment of the linear system

To use the direct dense linear solver based on the internal CVODE implementation, the user must make the call:

```
CALL FCVDENSE(NEQ, IER)
```

where NEQ is the size of the ODE system. The argument IER is an error return flag which is 0 for success, -1 if a memory allocation failure occurred, or -2 for illegal input.

Alternatively, to use the Lapack-based direct dense linear solver, the user must make the call:

```
CALL FCVLAPACKDENSE(NEQ, IER)
```

where the arguments have the exact same meanings as for FCVDENSE.

As an option when using the DENSE linear solver, the user may supply a routine that computes a dense approximation of the system Jacobian $J = \partial f / \partial y$. If supplied, it must have the following form:

```
SUBROUTINE FCVDJAC (NEQ, T, Y, FY, DJAC, H, IPAR, RPAR,
&                  WK1, WK2, WK3, IER)
  DIMENSION Y(*), FY(*), DJAC(NEQ,*), IPAR(*), RPAR(*),
&                  WK1(*), WK2(*), WK3(*)
```

Typically this routine will use only NEQ, T, Y, and DJAC. It must compute the Jacobian and store it columnwise in DJAC. The input arguments T, Y, and FY contain the current values of t , y , and $f(t, y)$, respectively. The arrays IPAR (of integers) and RPAR (of reals) contain user data and are the same as those passed to FCVMALLOC. The vectors WK1, WK2, and WK3 of length NEQ are provided as work space for use in FCVDJAC. IER is an error return flag which should be set to 0 if successful, a positive value if a recoverable error occurred (in which case CVODE will attempt to correct), or a negative value if FCVDJAC failed unrecoverably (in which case the integration is halted).

If the user's FCVDJAC uses difference quotient approximations, it may need to use the error weight array EWT and current stepsize H in the calculation of suitable increments. The array EWT can be obtained by calling FCVGETERRWEIGHTS using one of the work arrays as temporary storage for EWT. It may also need the unit roundoff, which can be obtained as the optional output ROUT(6), passed from the calling program to this routine using either RPAR or a common block.

If the FCVDJAC routine is provided, then, following the call to FCVDENSE, the user must make the call:

```
CALL FCVDENSESETJAC (FLAG, IER)
```

with FLAG $\neq 0$ to specify use of the user-supplied Jacobian approximation. The argument IER is an error return flag which is 0 for success or non-zero if an error occurred. If using the Lapack-based

direct dense linear solver, the use of a Jacobian approximation supplied by the user is indicated through the call

```
CALL FCVLAPACKDENSESETJAC (FLAG, IER)
```

Optional outputs specific to the DENSE case are listed in Table 5.2.

[S] Band treatment of the linear system

To use the direct band linear solver based on the internal CVODE implementation, the user must make the call:

```
CALL FCVBAND (NEQ, MU, ML, IER)
```

The arguments are: MU, the upper half-bandwidth; ML, the lower half-bandwidth; and IER an error return flag which is 0 for success, -1 if a memory allocation failure occurred, or -2 in case an input has an illegal value.

Alternatively, to use the Lapack-based direct band linear solver, the user must make the call:

```
CALL FCVLAPACKBAND(NEQ, MU, ML, IER)
```

where the arguments have the exact same meanings as for FCVBAND.

As an option when using the BAND linear solver, the user may supply a routine that computes a band approximation of the system Jacobian $J = \partial f / \partial y$. If supplied, it must have the following form:

```
SUBROUTINE FCVBJAC(NEQ, MU, ML, MDIM, T, Y, FY, BJAC, H, IPAR, RPAR,
&                  WK1, WK2, WK3, IER)
  DIMENSION Y(*), FY(*), BJAC(MDIM,*), IPAR(*), RPAR(*),
&                  WK1(*), WK2(*), WK3(*)
```

Typically this routine will use only NEQ, MU, ML, T, Y, and BJAC. It must load the MDIM by N array BJAC with the Jacobian matrix at the current (t, y) in band form. Store in $BJAC(k, j)$ the Jacobian element $J_{i,j}$ with $k = i - j + MU + 1$ ($k = 1 \cdots ML + MU + 1$) and $j = 1 \cdots N$. The input arguments T, Y, and FY contain the current values of t , y , and $f(t, y)$, respectively. The arrays IPAR (of integers) and RPAR (of reals) contain user data and are the same as those passed to FCVMALLOC. The vectors WK1, WK2, and WK3 of length NEQ are provided as work space for use in FCVBJAC. IER is an error return flag which should be set to 0 if successful, a positive value if a recoverable error occurred (in which case CVODE will attempt to correct), or a negative value if FCVBJAC failed unrecoverably (in which case the integration is halted).

If the user's FCVBJAC uses difference quotient approximations, it may need to use the error weight array EWT and current stepsize H in the calculation of suitable increments. The array EWT can be obtained by calling FCVGETERRWEIGHTS using one of the work arrays as temporary storage for EWT. It may also need the unit roundoff, which can be obtained as the optional output ROUT(6), passed from the calling program to this routine using either RPAR or a common block.

If the FCVBJAC routine is provided, then, following the call to FCVBAND, the user must make the call:

```
CALL FCVBANDSETJAC(FLAG, IER)
```

with FLAG $\neq 0$ to specify use of the user-supplied Jacobian approximation. The argument IER is an error return flag which is 0 for success or non-zero if an error occurred. If using the Lapack-based direct band linear solver, the use of a Jacobian approximation supplied by the user is indicated through the call

```
CALL FCVLAPACKNBANDSETJAC (FLAG, IER)
```

Optional outputs specific to the BAND case are listed in Table 5.2.

[S][P] SPGMR treatment of the linear systems

For the Scaled Preconditioned GMRES solution of the linear systems, the user must make the call

```
CALL FCVSPGMR(IPRETYPE, IGSTYPE, MAXL, DELT, IER)
```

The arguments are as follows. IPRETYPE specifies the preconditioner type: 0 for no preconditioning, 1 for left only, 2 for right only, or 3 for both sides. IGSTYPE indicates the Gram-Schmidt process type: 1 for modified G-S or 2 for classical G-S. MAXL is the maximum Krylov subspace dimension. DELT is the linear convergence tolerance factor. For all of the input arguments, a value of 0 or 0.0 indicates the default. IER is an error return flag which is 0 to indicate success, -1 if a memory allocation failure occurred, or -2 to indicate an illegal input.

Optional outputs specific to the SPGMR case are listed in Table 5.2.

For descriptions of the relevant optional user-supplied routines, see **User-supplied routines for SPGMR/SPBCG/SPTFQMR** below.

[S][P] SPBCG treatment of the linear systems

For the Scaled Preconditioned Bi-CGStab solution of the linear systems, the user must make the call

```
CALL FCVSPBCG(IPRETYPE, MAXL, DELT, IER)
```

Its arguments are the same as those with the same names for FCVSPGMR.

Optional outputs specific to the SPBCG case are listed in Table 5.2.

For descriptions of the relevant optional user-supplied routines, see **User-supplied routines for SPGMR/SPBCG/SPTFQMR** below.

[S][P] SPTFQMR treatment of the linear systems

For the Scaled Preconditioned Transpose-Free Quasi-Minimal Residual solution of the linear systems, the user must make the call

```
CALL FCVSPTFQMR(IPRETYPE, MAXL, DELT, IER)
```

Its arguments are the same as those with the same names for FCVSPGMR.

Optional outputs specific to the SPTFQMR case are listed in Table 5.2.

For descriptions of the relevant optional user-supplied routines, see below.

[S][P] Functions used by SPGMR/SPBCG/SPTFQMR

An optional user-supplied routine, FCVJTIMES (see below), can be provided for Jacobian-vector products. If it is, then, following the call to FCVSPGMR, FCVSPBCG, or FCVSPTFQMR, the user must make the call:

```
CALL FCVSPILSSETJAC(FLAG, IER)
```

with FLAG \neq 0 to specify use of the user-supplied Jacobian-times-vector approximation. The argument IER is an error return flag which is 0 for success or non-zero if an error occurred.

If preconditioning is to be done (IPRETYPE \neq 0), then the user must call

```
CALL FCVSPILSSETPREC(FLAG, IER)
```

with $\text{FLAG} \neq 0$. The return flag IER is 0 if successful, or negative if a memory error occurred. In addition, the user program must include preconditioner routines FCVPSOL and FCVPSET (see below).

[S][P] User-supplied routines for SPGMR/SPBCG/SPTFQMR

With treatment of the linear systems by any of the Krylov iterative solvers, there are three optional user-supplied routines — FCVJTIMES , FCVPSOL , and FCVPSET . The specifications for these routines are given below.

As an option when using the SPGMR, SPBCG, or SPTFQMR linear solvers, the user may supply a routine that computes the product of the system Jacobian $J = \partial f / \partial y$ and a given vector v . If supplied, it must have the following form:

```
SUBROUTINE FCVJTIMES (V, FJV, T, Y, FY, H, IPAR, RPAR, WORK, IER)
  DIMENSION V(*), FJV(*), Y(*), FY(*), IPAR(*), RPAR(*), WORK(*)
```

Typically this routine will use only NEQ , T , Y , V , and FJV . It must compute the product vector Jv , where the vector v is stored in V , and store the product in FJV . The input arguments T , Y , and FY contain the current values of t , y , and $f(t, y)$, respectively. On return, set $\text{IER} = 0$ if FCVJTIMES was successful, and nonzero otherwise. The arrays IPAR (of integers) and RPAR (of reals) contain user data and are the same as those passed to FCVMALLOC . The vector WORK , of length NEQ , is provided as work space for use in FCVJTIMES .

If preconditioning is to be included, the following routine must be supplied, for solution of the preconditioner linear system:

```
SUBROUTINE FCVPSOL(T, Y, FY, R, Z, GAMMA, DELTA, LR, IPAR, RPAR,
&                  WORK, IER)
  DIMENSION Y(*), FY(*), R(*), Z(*), IPAR(*), RPAR(*), WORK(*)
```

It must solve the preconditioner linear system $Pz = r$, where $r = \text{R}$ is input, and store the solution z in Z . Here P is the left preconditioner if $\text{LR}=1$ and the right preconditioner if $\text{LR}=2$. The preconditioner (or the product of the left and right preconditioners if both are nontrivial) should be an approximation to the matrix $I - \gamma J$, where I is the identity matrix, J is the system Jacobian, and $\gamma = \text{GAMMA}$. The input arguments T , Y , and FY contain the current values of t , y , and $f(t, y)$, respectively. On return, set $\text{IER} = 0$ if FCVPSOL was successful, set IER positive if a recoverable error occurred, and set IER negative if a non-recoverable error occurred.

The arrays IPAR (of integers) and RPAR (of reals) contain user data and are the same as those passed to FCVMALLOC . The argument WORK is a work array of length NEQ for use by this routine.

If the user's preconditioner requires that any Jacobian related data be evaluated or preprocessed, then the following routine can be used for the evaluation and preprocessing of the preconditioner:

```
SUBROUTINE FCVPSET(T, Y, FY, JOK, JCUR, GAMMA, H, IPAR, RPAR,
&                  WORK1, WORK2, WORK3, IER)
  DIMENSION Y(*), FY(*), EWT(*), IPAR(*), RPAR(*),
&                  WORK1(*), WORK2(*), WORK3(*)
```

It must perform any evaluation of Jacobian-related data and preprocessing needed for the solution of the preconditioner linear systems by FCVPSOL . The input argument JOK allows for Jacobian data to be saved and reused: If $\text{JOK} = 0$, this data should be recomputed from scratch. If $\text{JOK} = 1$, a saved copy of it may be reused, and the preconditioner constructed from it. The input arguments

T, Y, and FY contain the current values of t , y , and $f(t, y)$, respectively. On return, set JCUR = 1 if Jacobian data was computed, and set JCUR = 0 otherwise. Also on return, set IER = 0 if FCVPSET was successful, set IER positive if a recoverable error occurred, and set IER negative if a non-recoverable error occurred.

The arrays IPAR (of integers) and RPAR (of reals) contain user data and are the same as those passed to FCVMALLOC. The arguments WORK1, WORK2, WORK3 are work arrays of length NEQ for use by this routine.

If the user calls FCVSPILSSETPREC the routine FCVPSET must be provided, even if it is not needed and must return IER=0.



Notes

- (a) If the user's FCVJTIMES or FCVPSET routine uses difference quotient approximations, it may need to use the error weight array EWT, the current stepsize H, and/or the unit roundoff, in the calculation of suitable increments. Also, If FCVPSOL uses an iterative method in its solution, the residual vector $\rho = r - Pz$ of the system should be made less than DELTA in weighted ℓ_2 norm, i.e. $\sqrt{\sum (\rho_i * EWT[i])^2} < DELTA$.
- (b) If needed in FCVJTIMES, FCVPSOL, or FCVPSET, the error weight array EWT can be obtained by calling FCVGETERRWEIGHTS using one of the work arrays as temporary storage for EWT.
- (c) If needed in FCVJTIMES, FCVPSOL, or FCVPSET, the unit roundoff can be obtained as the optional output ROUT(6) (available after the call to FCVMALLOC) and can be passed using either the RPAR user data array or a common block.

6. Problem solution

Carrying out the integration is accomplished by making calls as follows:

```
CALL FCVODE(TOUT, T, Y, ITASK, IER)
```

The arguments are as follows. TOUT specifies the next value of t at which a solution is desired (input). T is the value of t reached by the solver on output. Y is an array containing the computed solution on output. ITASK is a task indicator and should be set to 1 for normal mode (overshoot TOUT and interpolate), to 2 for one-step mode (return after each internal step taken), to 3 for normal mode with the additional tstop constraint, or to 4 for one-step mode with the additional constraint tstop. IER is a completion flag and will be set to a positive value upon successful return or to a negative value if an error occurred. These values correspond to the CVODE returns (see §4.5.5 and §B.2). The current values of the optional outputs are available in IOUT and ROUT (see Table 5.2).

7. Additional solution output

After a successful return from FCVODE, the routine FCVDKY may be used to obtain a derivative of the solution, of order up to the current method order, at any t within the last step taken. For this, make the following call:

```
CALL FCVDKY(T, K, DKY, IER)
```

where T is the value of t at which solution derivative is desired, and K is the derivative order ($0 \leq K \leq QU$). On return, DKY is an array containing the computed K-th derivative of y . The value T must lie between TCUR - HU and TCUR. The return flag IER is set to 0 upon successful return or to a negative value to indicate an illegal input.

8. Problem reinitialization

To re-initialize the CVODE solver for the solution of a new problem of the same size as one already solved, make the following call:

```
CALL FCVREINIT(TO, YO, IATOL, RTOL, ATOL, IER)
```

The arguments have the same names and meanings as those of FCVMALLOC. FCVREINIT performs the same initializations as FCVMALLOC, but does no memory allocation, using instead the existing internal memory created by the previous FCVMALLOC call. The call to specify the linear system solution method may or may not be needed.

Following this call, a call to specify the linear system solver must be made if the choice of linear solver is being changed. Otherwise, a call to reinitialize the linear solver last used may or may not be needed, depending on changes in the inputs to it.

In the case of the BAND solver, for any change in the half-bandwidth parameters, call FCVBAND (or FCVLAPACKBAND) as described above.

In the case of SPGMR, for a change of inputs other than MAXL, make the call

```
CALL FCVSPGMRREINIT (IPRETYPE, IGSTYPE, DELT, IER)
```

which reinitializes SPGMR without reallocating its memory. The arguments have the same names and meanings as those of FCVSPGMR. If MAXL is being changed, then call FCVSPGMR instead.

In the case of SPBCG, for a change in any inputs, make the call

```
CALL FCVSPBCGREINIT (IPRETYPE, MAXL, DELT, IER)
```

which reinitializes SPBCG without reallocating its memory. The arguments have the same names and meanings as those of FCVSPBCG.

In the case of SPTFQMR, for a change in any inputs, make the call

```
CALL FCVSPTFQMRREINIT (IPRETYPE, MAXL, DELT, IER)
```

which reinitializes SPTFQMR without reallocating its memory. The arguments have the same names and meanings as those of FCVSPTFQMR.

9. Memory deallocation

To free the internal memory created by the call to FCVMALLOC, make the call

```
CALL FCVFREE
```

5.3 FCVODE optional input and output

In order to keep the number of user-callable FCVODE interface routines to a minimum, optional inputs to the CVODE solver are passed through only two routines: FCVSETIIN for integer optional inputs and FCVSETRIN for real optional inputs. These functions should be called as follows:

```
CALL FCVSETIIN(KEY, IVAL, IER)
CALL FCVSETRIN(KEY, RVAL, IER)
```

where KEY is a quoted string indicating which optional input is set (see Table 5.1), IVAL is the integer input value to be used, RVAL is the real input value to be used, and IER is an integer return flag which is set to 0 on success and a negative value if a failure occurred.

The optional outputs from the CVODE solver are accessed not through individual functions, but rather through a pair of arrays, IOUT (integer type) of dimension at least 21, and ROUT (real type) of dimension at least 6. These arrays are owned (and allocated) by the user and are passed as arguments

Table 5.1: Keys for setting FCVODE optional inputs

Integer optional inputs (FCVSETIIN)		
Key	Optional input	Default value
MAX_ORD	Maximum LMM method order	5 (BDF), 12 (Adams)
MAX_NSTEPS	Maximum no. of internal steps before t_{out}	500
MAX_ERRFAIL	Maximum no. of error test failures	7
MAX_NITERS	Maximum no. of nonlinear iterations	3
MAX_CONVFAIL	Maximum no. of convergence failures	10
HNIL_WARN	Maximum no. of warnings for $t_n + h = t_n$	10
STAB_LIM	Flag to activate stability limit detection	0

Real optional inputs (FCVSETRIN)		
Key	Optional input	Default value
INIT_STEP	Initial step size	estimated
MAX_STEP	Maximum absolute step size	∞
MIN_STEP	Minimum absolute step size	0.0
STOP_TIME	Value of t_{stop}	undefined
NLCONV_COEF	Coefficient in the nonlinear convergence test	0.1

to FCVMALLOC. Table 5.2 lists the entries in these two arrays and specifies the optional variable as well as the CVODE function which is actually called to extract the optional output.

For more details on the optional inputs and outputs, see §4.5.6 and §4.5.8.

In addition to the optional inputs communicated through FCVSET* calls and the optional outputs extracted from IOUT and ROUT, the following user-callable routines are available:

To obtain the error weight array EWT, containing the multiplicative error weights used the WRMS norms, make the following call:

```
CALL FCVGETERRWEIGHTS (EWT, IER)
```

This computes the EWT array normally defined by Eq. (2.6). The array EWT, of length NEQ or NLOCAL, must already have been declared by the user. The error return flag IER is zero if successful, and negative if there was a memory error.

To obtain the estimated local errors, following a successful call to FCVSOLVE, make the following call:

```
CALL FCVGETESTLOCALERR (ELE, IER)
```

This computes the ELE array of estimated local errors as of the last step taken. The array ELE must already have been declared by the user. The error return flag IER is zero if successful, and negative if there was a memory error.

5.4 Usage of the FCVROOT interface to rootfinding

The FCVROOT interface package allows programs written in FORTRAN to use the rootfinding feature of the CVODE solver module. The user-callable functions in FCVROOT, with the corresponding CVODE functions, are as follows:

- FCVROOTINIT interfaces to CNodeRootInit.
- FCVROOTINFO interfaces to CNodeGetRootInfo.
- FCVROOTFREE interfaces to CNodeRootFree.

Table 5.2: Description of the FCVODE optional output arrays IOUT and ROUT

Integer output array IOUT		
Index	Optional output	CVODE function
CVODE main solver		
1	LENRW	CVodeGetWorkSpace
2	LENIW	CVodeGetWorkSpace
3	NST	CVodeGetNumSteps
4	NFE	CVodeGetNumRhsEvals
5	NETF	CVodeGetNumErrTestFails
6	NCFN	CVodeGetNumNonlinSolvConvFails
7	NNI	CVodeGetNumNonlinSolvIters
8	NSETUPS	CVodeGetNumLinSolvSetups
9	QU	CVodeGetLastOrder
10	QCUR	CVodeGetCurrentOrder
11	NOR	CVodeGetNumStabLimOrderReds
12	NGE	CVodeGetNumGEvals
CVDENSE, CVBAND linear solvers		
13	LENRWLS	CVDlsGetWorkSpace
14	LENIWLS	CVDlsGetWorkSpace
15	LS_FLAG	CVDlsGetLastFlag
16	NFELS	CVDlsGetNumRhsEvals
17	NJE	CVDlsGetNumJacEvals
CVDIAG linear solver		
13	LENRWLS	CVDdiagGetWorkSpace
14	LENIWLS	CVDdiagGetWorkSpace
15	LS_FLAG	CVDdiagGetLastFlag
16	NFELS	CVDdiagGetNumRhsEvals
CVSPGMR, CVSPBCG, CVSPTFQMR linear solvers		
13	LENRWLS	CVSpilsGetWorkSpace
14	LENIWLS	CVSpilsGetWorkSpace
15	LS_FLAG	CVSpilsGetLastFlag
16	NFELS	CVSpilsGetNumRhsEvals
17	NJTV	CVSpilsGetNumJacEvals
18	NPE	CVSpilsGetNumPrecEvals
19	NPS	CVSpilsGetNumPrecSolves
20	NLI	CVSpilsGetNumLinIters
21	NCFL	CVSpilsGetNumConvFails

Real output array ROUT		
Index	Optional output	CVODE function
1	HOU	CVodeGetActualInitStep
2	HU	CVodeGetLastStep
3	HCUR	CVodeGetCurrentStep
4	TCUR	CVodeGetCurrentTime
5	TOLSF	CVodeGetTolScaleFactor
6	UROUND	unit roundoff

Note that at this time, FCVROOT does not provide support to specify the direction of zero-crossing that is to be monitored. Instead, all roots are considered. However, the actual direction of zero-crossing is reported (through the sign of the non-zero elements in the array `INFO` returned by `FCVROOTINFO`).

In order to use the rootfinding feature of `CVODE`, the following call must be made, after calling `FCVMALLOC` but prior to calling `FCVODE`, to allocate and initialize memory for the `FCVROOT` module:

```
CALL FCVROOTINIT (NRTFN, IER)
```

The arguments are as follows: `NRTFN` is the number of root functions. `IER` is a return completion flag; its values are 0 for success, -1 if the `CVODE` memory was `NULL`, and -11 if a memory allocation failed.

To specify the functions whose roots are to be found, the user must define the following routine:

```
SUBROUTINE FCVROOTFN (T, Y, G, IPAR, RPAR, IER)
  DIMENSION Y(*), G(*), IPAR(*), RPAR(*)
```

It must set the `G` array, of length `NRTFN`, with components $g_i(t, y)$, as a function of $T = t$ and the array $Y = y$. The arrays `IPAR` (of integers) and `RPAR` (of reals) contain user data and are the same as those passed to `FCVMALLOC`. Set `IER` on 0 if successful, or on a non-zero value if an error occurred.

When making calls to `FCVODE` to solve the ODE system, the occurrence of a root is flagged by the return value `IER = 2`. In that case, if `NRTFN > 1`, the functions g_i which were found to have a root can be identified by making the following call:

```
CALL FCVROOTINFO (NRTFN, INFO, IER)
```

The arguments are as follows: `NRTFN` is the number of root functions. `INFO` is an integer array of length `NRTFN` with root information. `IER` is a return completion flag; its values are 0 for success, negative if there was a memory failure. The returned values of `INFO(i)` ($i = 1, \dots, \text{NRTFN}$) are 0 or ± 1 , such that `INFO(i) = +1` if g_i was found to have a root and g_i is increasing, `INFO(i) = -1` if g_i was found to have a root and g_i is decreasing, and `INFO(i) = 0` otherwise.

The total number of calls made to the root function `FCVROOTFN`, denoted `NGE`, can be obtained from `IOUT(12)`. If the `FCVODE`/`CVODE` memory block is reinitialized to solve a different problem via a call to `FCVREINIT`, then the counter `NGE` is reset to zero.

To free the memory resources allocated by a prior call to `FCVROOTINIT`, make the following call:

```
CALL FCVROOTFREE
```

5.5 Usage of the FCVBP interface to CVBANDPRE

The `FCVBP` interface sub-module is a package of C functions which, as part of the `FCVODE` interface module, support the use of the `CVODE` solver with the serial `NVECTOR_SERIAL` module, and the combination of the `CVBANDPRE` preconditioner module (see §4.7.1) with any of the Krylov iterative linear solvers.

The two user-callable functions in this package, with the corresponding `CVODE` function around which they wrap, are:

- `FCVBPINIT` interfaces to `CVBandPrecInit`.
- `FCVBPOPT` interfaces to `CVBANDPRE` optional output functions.

As with the rest of the `FCVODE` routines, the names of the user-supplied routines are mapped to actual values through a series of definitions in the header file `fcvbp.h`.

The following is a summary of the usage of this module. Steps that are unchanged from the main program described in §5.2 are grayed-out.

1. **Right-hand side specification**
2. **NVECTOR module initialization**

3. Problem specification

4. Set optional inputs

5. Linear solver specification

First, specify one of the CVSPILS iterative linear solvers, by calling one of FCVSPGMR, FCVSPBCG, or FCVSPTFQMR.

Then, to initialize the CVBANDPRE preconditioner, make the following call:

```
CALL FCVBPINIT(NEQ, MU, ML, IER)
```

The arguments are as follows. `NEQ` is the problem size. `MU` and `ML` are the upper and lower half-bandwidths of the band matrix that is retained as an approximation of the Jacobian. `IER` is a return completion flag. A value of 0 indicates success, while a value of -1 indicates that a memory failure occurred.

Optionally, to specify that SPGMR, SPBCG, or SPTFQMR should use the supplied FCVJTIMES, make the call

```
CALL FCVSPILSSETJAC(FLAG, IER)
```

with `FLAG` $\neq 0$ (see step 5 in §5.2 for details).

6. Problem solution

7. CVBANDPRE Optional outputs

Optional outputs specific to the SPGMR, SPBCG, or SPTFQMR solver are listed in Table 5.2. To obtain the optional outputs associated with the CVBANDPRE module, make the following call:

```
CALL FCVBPOPT(LENRWBP, LENIWBP, NFEBP)
```

The arguments returned are as follows. `LENRWBP` is the length of real preconditioner work space, in `realtype` words. `LENIWBP` is the length of integer preconditioner work space, in integer words. `NFEBP` is the number of $f(t, y)$ evaluations (calls to `FCVFUN`) for difference-quotient banded Jacobian approximations.

8. Memory deallocation

(The memory allocated for the FCVBP module is deallocated automatically by `FCVFREE`.)

5.6 Usage of the FCVBBD interface to CVBBDPRE

The FCVBBD interface sub-module is a package of C functions which, as part of the FCVODE interface module, support the use of the CVODE solver with the parallel NVECTOR_PARALLEL module, and the combination of the CVBBDPRE preconditioner module (see §4.7.2) with any of the Krylov iterative linear solvers.

The user-callable functions in this package, with the corresponding CVODE and CVBBDPRE functions, are as follows:

- FCVBBDINIT interfaces to CVBBDPrecInit.
- FCVBBDREINIT interfaces to CVBBDPrecReInit.
- FCVBBDOPT interfaces to CVBBDPRE optional output functions.

In addition to the FORTRAN right-hand side function `FCVFUN`, the user-supplied functions used by this package, are listed below, each with the corresponding interface function which calls it (and its type within CVBBDPRE or CVODE):

FCVBBD routine (FORTRAN, user-supplied)	CVODE function (C, interface)	CVODE type of interface function
FCVLOCfN	FCVgloc	CVLocalFn
FCVCOMMF	FCVcfn	CVCommFn
FCVJTIMES	FCVJtimes	CVSpilsJacTimesVecFn

As with the rest of the FCVODE routines, the names of all user-supplied routines here are fixed, in order to maximize portability for the resulting mixed-language program. Additionally, based on flags discussed above in §5.1, the names of the user-supplied routines are mapped to actual values through a series of definitions in the header file `fcvbbd.h`.

The following is a summary of the usage of this module. Steps that are unchanged from the main program described in §5.2 are grayed-out.

1. Right-hand side specification
2. NVECTOR module initialization
3. Problem specification
4. Set optional inputs
5. Linear solver specification

First, specify one of the CVSPILS iterative linear solvers, by calling one of `FCVSPGMR`, `FCVSPBCG`, or `FCVSPTFQMR`.

Then, to initialize the CVBBDPRE preconditioner, make the following call:

```
CALL FCVBBDINIT(NLOCAL, MUDQ, MLDQ, MU, ML, DQRELY, IER)
```

The arguments are as follows. `NLOCAL` is the local size of vectors on this processor. `MUDQ` and `MLDQ` are the upper and lower half-bandwidths to be used in the computation of the local Jacobian blocks by difference quotients. These may be smaller than the true half-bandwidths of the Jacobian of the local block of g , when smaller values may provide greater efficiency. `MU` and `ML` are the upper and lower half-bandwidths of the band matrix that is retained as an approximation of the local Jacobian block. These may be smaller than `MUDQ` and `MLDQ`. `DQRELY` is the relative increment factor in y for difference quotients (optional). A value of 0.0 indicates the default, $\sqrt{\text{unit roundoff}}$. `IER` is a return completion flag. A value of 0 indicates success, while a value of -1 indicates that a memory failure occurred or that an input had an illegal value.

Optionally, to specify that SPGMR, SPBCG, or SPTFQMR should use the supplied `FCVJTIMES`, make the call

```
CALL FCVSPILSSETJAC(FLAG, IER)
```

with `FLAG` $\neq 0$ (see step 5 in §5.2 for details).

6. Problem solution
7. CVBBDPRE Optional outputs

Optional outputs specific to the SPGMR, SPBCG, or SPTFQMR solver are listed in Table 5.2. To obtain the optional outputs associated with the CVBBDPRE module, make the following call:

```
CALL FCVBBDOPT(LENRWBBD, LENIWBBBD, NGEBBBD)
```

The arguments returned are as follows. LENRWBBD is the length of real preconditioner work space, in `realtype` words. LENIWBBBD is the length of integer preconditioner work space, in integer words. These sizes are local to the current processor. NGBBBD is the number of $g(t, y)$ evaluations (calls to FCVLOCFN) so far.

8. Problem reinitialization

If a sequence of problems of the same size is being solved using the same linear solver (SPGMR, SPBCG, or SPTFQMR) in combination with the CVBBDPRE preconditioner, then the CVODE package can be re-initialized for the second and subsequent problems by calling FCVREINIT, following which a call to FCVBBDINIT may or may not be needed. If the input arguments are the same, no FCVBBDINIT call is needed. If there is a change in input arguments other than MU or ML, then the user program should make the call

```
CALL FCVBBDREINIT(NLOCAL, MUDQ, MLDQ, DQRELY, IER)
```

This reinitializes the CVBBDPRE preconditioner, but without reallocating its memory. The arguments of the FCVBBDREINIT routine have the same names and meanings as those of FCVBBDINIT. If the value of MU or ML is being changed, then a call to FCVBBDINIT must be made. Finally, if there is a change in any of the linear solver inputs, then a call to FCVSPGMR, FCVSPBCG, or FCVSPTFQMR must also be made; in this case the linear solver memory is reallocated.

9. Memory deallocation

(The memory allocated for the FCVBBD module is deallocated automatically by FCVFREE.)

10. User-supplied routines

The following two routines must be supplied for use with the CVBBDPRE module:

```
SUBROUTINE FCVGLOCFN (NLOC, T, YLOC, GLOC, IPAR, RPAR, IER)
  DIMENSION YLOC(*), GLOC(*), IPAR(*), RPAR(*)
```

This routine is to evaluate the function $g(t, y)$ approximating f (possibly identical to f), in terms of $T = t$, and the array YLOC (of length NLOC), which is the sub-vector of y local to this processor. The resulting (local) sub-vector is to be stored in the array GLOC. The arrays IPAR (of integers) and RPAR (of reals) contain user data and are the same as those passed to FCVMALLOC. IER is an error return flag which should be set to 0 if successful, a positive value if a recoverable error occurred (in which case CVODE will attempt to correct), or a negative value if FCVGLOCFN failed unrecoverably (in which case the integration is halted).

```
SUBROUTINE FCVCOMMFN (NLOC, T, YLOC, IPAR, RPAR, IER)
  DIMENSION YLOC(*), IPAR(*), RPAR(*)
```

This routine is to perform the inter-processor communication necessary for the FCVGLOCFN routine. Each call to FCVCOMMFN is preceded by a call to the right-hand side routine FCVFUN with the same arguments T and YLOC. The arrays IPAR (of integers) and RPAR (of reals) contain user data and are the same as those passed to FCVMALLOC. IER is an error return flag (currently not used; set IER=0). Thus FCVCOMMFN can omit any communications done by FCVFUN if relevant to the evaluation of GLOC. IER is an error return flag which should be set to 0 if successful, a positive value if a recoverable error occurred (in which case CVODE will attempt to correct), or a negative value if FCVCOMMFN failed unrecoverably (in which case the integration is halted).

The subroutine FCVCOMMFN must be supplied even if it is not needed and must return IER=0.

Optionally, the user can supply a routine FCVJTIMES for the evaluation of Jacobian-vector products, as described above in step 5 in §5.2.



Chapter 6

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);
boolean_t   (*nvintest)(N_Vector, N_Vector);
boolean_t   (*nvconstrmask)(N_Vector, N_Vector, N_Vector);
realtype    (*nvminquotient)(N_Vector, N_Vector);
};

```

The generic NVECTOR module defines and implements the vector operations acting on `N_Vector`. These routines are nothing but wrappers for the vector operations defined by a particular NVECTOR implementation, which are accessed through the *ops* field of the `N_Vector` structure. To illustrate this point we show below the implementation of a typical vector operation from the generic NVECTOR module, namely `N_VScale`, which performs the scaling of a vector `x` by a scalar `c`:

```

void N_VScale(realtype c, N_Vector x, N_Vector z)
{
    z->ops->nvscale(c, x, z);
}

```

Table 6.1 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 6.1: 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 a y that is guaranteed to have all nonzero components.

continued on next page

<i>continued from last page</i>	
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.
<i>continued on next page</i>	

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.

6.1 The NVECTOR_SERIAL implementation

The serial implementation of the NVECTOR module provided with SUNDIALS, NVECTOR_SERIAL, defines the *content* field of `N_Vector` to be a structure containing the length of the vector, a pointer to the beginning of a contiguous data array, and a boolean flag *own_data* which specifies the ownership of *data*.

```
struct _N_VectorContent_Serial {
    long int length;
    booleantype own_data;
    realtype *data;
};
```

The following five macros are provided to access the content of an NVECTOR_SERIAL vector. The suffix `_S` in the names denotes 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 6.1. Their names are obtained from those in Table 6.1 by appending the suffix `_Serial`. The module **NVECTOR_SERIAL** provides the following additional user-callable routines:

- **N_VNew_Serial**

This function creates and allocates memory for a serial **N_Vector**. Its only argument is the vector length.

```
N_Vector N_VNew_Serial(long int vec_length);
```

- **N_VNewEmpty_Serial**

This function creates a new serial **N_Vector** with an empty (NULL) data array.

```
N_Vector N_VNewEmpty_Serial(long int vec_length);
```

- **N_VMake_Serial**

This function creates and allocates memory for a serial vector with user-provided data array.

```
N_Vector N_VMake_Serial(long int vec_length, realtype *v_data);
```

- **N_VCloneVectorArray_Serial**

This function creates (by cloning) an array of `count` serial vectors.

```
N_Vector *N_VCloneVectorArray_Serial(int count, N_Vector w);
```

- **N_VCloneEmptyVectorArray_Serial**

This function creates (by cloning) an array of `count` serial vectors, each with an empty (NULL) data array.

```
N_Vector *N_VCloneEmptyVectorArray_Serial(int count, N_Vector w);
```

- **N_VDestroyVectorArray_Serial**

This function frees memory allocated for the array of `count` variables of type **N_Vector** created with **N_VCloneVectorArray_Serial** or with **N_VCloneEmptyVectorArray_Serial**.

```
void N_VDestroyVectorArray_Serial(N_Vector *vs, int count);
```

- **N_VPrint_Serial**

This function prints the content of a serial vector to `stdout`.

```
void N_VPrint_Serial(N_Vector v);
```

Notes

- When looping over the components of an `N_Vector` `v`, it is more efficient to first obtain the component array via `v_data = NV_DATA.S(v)` and then access `v_data[i]` within the loop than it is to use `NV_Ith.S(v,i)` within the loop.
- `N_VNewEmpty_Serial`, `N_VMake_Serial`, and `N_VCloneEmptyVectorArray_Serial` set the field `own_data = FALSE`. `N_VDestroy_Serial` and `N_VDestroyVectorArray_Serial` will not attempt to free the pointer `data` for any `N_Vector` with `own_data` set to `FALSE`. In such a case, it is the user's responsibility to deallocate the `data` pointer.
- To maximize efficiency, vector operations in the `NVECTOR_SERIAL` implementation that have more than one `N_Vector` argument do not check for consistent internal representation of these vectors. It is the user's responsibility to ensure that such routines are called with `N_Vector` arguments that were all created with the same internal representations.



6.2 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 6.1 Their names are obtained from those in Table 6.1 by appending the suffix `_Parallel`. The module NVECTOR_PARALLEL provides the following additional user-callable routines:

- N_VNew_Parallel

This function creates and allocates memory for a parallel vector.

```
N_Vector N_VNew_Parallel(MPI_Comm comm,
                        long int local_length,
                        long int global_length);
```

- N_VNewEmpty_Parallel

This function creates a new parallel N_Vector with an empty (NULL) data array.

```
N_Vector N_VNewEmpty_Parallel(MPI_Comm comm,
                              long int local_length,
                              long int global_length);
```

- N_VMake_Parallel

This function creates and allocates memory for a parallel vector with user-provided data array.

```
N_Vector N_VMake_Parallel(MPI_Comm comm,
                          long int local_length,
                          long int global_length,
                          realtype *v_data);
```

- N_VCloneVectorArray_Parallel

This function creates (by cloning) an array of count parallel vectors.

```
N_Vector *N_VCloneVectorArray_Parallel(int count, N_Vector w);
```

- N_VCloneEmptyVectorArray_Parallel

This function creates (by cloning) an array of count parallel vectors, each with an empty (NULL) data array.

```
N_Vector *N_VCloneEmptyVectorArray_Parallel(int count, N_Vector w);
```

- **N_VDestroyVectorArray_Parallel**

This function frees memory allocated for the array of `count` variables of type `N_Vector` created with `N_VCloneVectorArray_Parallel` or with `N_VCloneEmptyVectorArray_Parallel`.

```
void N_VDestroyVectorArray_Parallel(N_Vector *vs, int count);
```

- **N_VPrint_Parallel**

This function prints the content of a parallel vector to stdout.

```
void N_VPrint_Parallel(N_Vector v);
```

Notes

- When looping over the components of an `N_Vector` `v`, it is more efficient to first obtain the local component array via `v_data = NV_DATA_P(v)` and then access `v_data[i]` within the loop than it is to use `NV_Ith_P(v,i)` within the loop.
- `N_VNewEmpty_Parallel`, `N_VMake_Parallel`, and `N_VCloneEmptyVectorArray_Parallel` set the field `own_data = FALSE`. `N_VDestroy_Parallel` and `N_VDestroyVectorArray_Parallel` will not attempt to free the pointer `data` for any `N_Vector` with `own_data` set to `FALSE`. In such a case, it is the user's responsibility to deallocate the `data` pointer.
- To maximize efficiency, vector operations in the `NVECTOR_PARALLEL` implementation that have more than one `N_Vector` argument do not check for consistent internal representation of these vectors. It is the user's responsibility to ensure that such routines are called with `N_Vector` arguments that were all created with the same internal representations.



6.3 NVECTOR functions used by CVODE

In Table 6.2 below, we list the vector functions in the `NVECTOR` module within the `CVODE` package. The table also shows, for each function, which of the code modules uses the function. The `CVODE` column shows function usage within the main integrator module, while the remaining seven columns show function usage within each of the six `CVODE` linear solvers (CVSPILS stands for any of CVSPGMR, CVSPBCG, or CVSPTFQMR), the CVBANDPRE and CVBBDPRE preconditioner modules, and the FCVODE module.

There is one subtlety in the CVSPILS column hidden by the table, explained here for the case of the CVSPGMR module. The `N_VDotProd` function is called both within the interface file `cvode_spgmr.c` and within the implementation files `sundials_spgmr.c` and `sundials_iterative.c` for the generic SPGMR solver upon which the CVSPGMR solver is built. Also, although `N_VDiv` and `N_VProd` are not called within the interface file `cvode_spgmr.c`, they are called within the implementation file `sundials_spgmr.c`, and so are required by the CVSPGMR solver module. Analogous statements apply to the CVSPBCG and CVSPTFQMR modules, except that they do not use `sundials_iterative.c`. This issue does not arise for the other three `CVODE` 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 `CVODE` user does not need to know anything about the usage of vector functions by the `CVODE` code modules in order to use `CVODE`. The information is presented as an implementation detail for the interested reader.

The vector functions listed in Table 6.1 that are *not* used by `CVODE` are: `N_VWL2Norm`, `N_VL1Norm`, `N_VWrmsNormMask`, `N_VConstrMask`, `N_VCloneEmpty`, and `N_VMinQuotient`. Therefore a user-supplied `NVECTOR` module for `CVODE` could omit these six functions.

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

	CVOICE	CV DENSE	CV BAND	CV DIAG	CV SPILS	CV BANDPRE	CV BBDDPRE	FCVOICE
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 7

Providing Alternate Linear Solver Modules

The central CVODE module interfaces with the linear solver module to be used by way of calls to four routines. 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`: evaluate and preprocess 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 routine (like those described in §4.5.3) which will attach the above four routines to the main CVODE memory block. The CVODE memory block is a structure defined in the header file `cvode_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 routines, only the `lsolve` routine is required. The `lfree` routine must be provided only if the solver specification routine makes any memory allocation. For consistency with the existing CVODE 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 CVODE 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 CVODE memory block can be used to attach a linear solver-specific memory block.

These four routines that interface between CVODE and the linear solver module necessarily have fixed call sequences. Thus, a user wishing to implement another linear solver within the CVODE package must adhere to this set of interfaces. The following is a complete description of the call list for each of these routines. Note that the call list of each routine includes a pointer to the main CVODE memory block, by which the routine can access various data related to the CVODE solution. The contents of this memory block are given in the file `cvode_impl.h` (but not reproduced here, for the sake of space).

7.1 Initialization function

The type definition of `linit` is

linit

Definition `int (*linit)(CNodeMem cv_mem);`

Purpose The purpose of `linit` is to complete initializations for specific linear solver, such as counters and statistics.

Arguments `cv_mem` is the CVODE memory pointer of type `CNodeMem`.

Return value An `linit` function should return 0 if it has successfully initialized the CVODE linear solver and `-1` otherwise.

7.2 Setup function

The type definition of `lsetup` is

lsetup

Definition `int (*lsetup)(CNodeMem 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 re-compute Jacobian-related data if it deems necessary.

Arguments `cv_mem` is the CVODE memory pointer of type `CNodeMem`.

`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 CVODE linear solver needs to be updated or not. Its possible values are:

- `CV_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).
- `CV_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 routine indicated that its Jacobian-related data is not current, or (b) during the previous Newton corrector iteration, the linear solver's solve routine failed in a recoverable manner and the linear solver's setup routine indicated that its Jacobian-related data is not current.
- `CV_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 CVODE 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 re-evaluation of Jacobian data (based on `convfail` and CVODE 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 The `lsetup` routine should return 0 if successful, a positive value for a recoverable error, and a negative value for an unrecoverable error.

The type definition of `lsolve` is

```
Definition      int (*lsolve)(CNodeMem cv_mem, N_Vector b, N_Vector weight,
                             N_Vector ycur, N_Vector fcur);
```

Arguments	<code>cv_mem</code> is the CVODE memory pointer of type <code>CVodeMem</code> .
	<code>b</code> is the right-hand side vector b . The solution is to be returned in the vector <code>b</code> .
	<code>weight</code> is a vector that contains the error weights. These are the W_i of Eq.(2.6).
	<code>ycur</code> is a vector that contains the solver's current approximation to $y(t_n)$.
	<code>fcur</code> is a vector that contains $f(t_n, y_{cur})$.

7.4 Memory deallocation function

```
Definition    void (*lfree)(CNodeMem cv_mem);
```

Arguments The argument `cv_mem` is the CVODE memory pointer of type `CVodeMem`.

Return value This routine has no return value.

Notes	This routine is called once a problem has been completed and the linear solver is no longer needed.
-------	---

Chapter 8

Generic Linear Solvers in SUNDIALS

In this chapter, we describe five generic linear solver code modules that are included in CVODE, but which are of potential use as generic packages in themselves, either in conjunction with the use of CVODE or separately.

These generic linear solver modules in SUNDIALS are organized in two families of solvers, the *dls* family, which includes direct linear solvers appropriate for sequential computations; and the *spils* family, which includes scaled preconditioned iterative (Krylov) linear solvers. The solvers in each family share common data structures and functions.

The *dls* family contains the following two generic linear solvers:

- The DENSE package, a linear solver for dense matrices either specified through a matrix type (defined below) or as simple arrays.
- The BAND package, a linear solver for banded matrices either specified through a matrix type (defined below) or as simple arrays.

Note that this family also includes the Blas/Lapack linear solvers (dense and band) available to the SUNDIALS solvers, but these are not discussed here.

The *spils* family contains the following three generic linear solvers:

- The SPGMR package, a solver for the scaled preconditioned GMRES method.
- The SPBCG package, a solver for the scaled preconditioned Bi-CGStab method.
- The SPTFQMR package, a solver for the scaled preconditioned TFQMR method.

For reasons related to installation, the names of the files involved in these 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 the **dense** and **band** modules that work with a matrix type and the functions in the SPGMR, SPBCG, and SPTFQMR modules are only summarized briefly, since they are less likely to be of direct use in connection with a SUNDIALS solver. However, the functions for dense matrices treated as simple arrays are fully described, because we expect that they will be useful in the implementation of preconditioners used with the combination of one of the SUNDIALS solvers and one of the *spils* linear solvers.

8.1 The DLS modules: DENSE and BAND

The files comprising the DENSE generic linear solver, and their locations in the SUNDIALS *srcdir*, are as follows:

- header files (located in *srcdir/include/sundials*)
`sundials_direct.h` `sundials_dense.h`
`sundials_types.h` `sundials_math.h` `sundials_config.h`
- source files (located in *srcdir/src/sundials*)
`sundials_direct.c` `sundials_dense.c` `sundials_math.c`

The files comprising the BAND generic linear solver are as follows:

- header files (located in *srcdir/include/sundials*)
`sundials_direct.h` `sundials_band.h`
`sundials_types.h` `sundials_math.h` `sundials_config.h`
- source files (located in *srcdir/src/sundials*)
`sundials_direct.c` `sundials_band.c` `sundials_math.c`

Only two of the preprocessing directives in the header file `sundials_config.h` are relevant to the DENSE and BAND packages by themselves (see §A.3 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 `MIN`, `MAX`, and `ABS` macros and `RAbs` function.

The files listed above for either module can be extracted from the SUNDIALS *srcdir* and compiled by themselves into a separate library or into a larger user code.

8.1.1 Type DlsMat

The type `DlsMat`, defined in `sundials_direct.h` is a pointer to a structure defining a generic matrix, and is used with all linear solvers in the *dls* family:

```
typedef struct _DlsMat {
    int type;
    int M;
    int N;
    int ldim;
    int mu;
    int ml;
    int s_mu;
    realtype *data;
    int ldata;
    realtype **cols;
} *DlsMat;
```

For the DENSE module, the relevant fields of this structure are as follows. Note that a dense matrix of type `DlsMat` need not be square.

type - SUNDIALS_DENSE (=1)

M - number of rows

N - number of columns

ldim - leading dimension ($\text{ldim} \geq M$)

data - pointer to a contiguous block of **realtype** variables

ldata - length of the data array ($= \text{ldim} \cdot N$). The (i,j) -th element of a dense matrix **A** of type **DlsMat** (with $0 \leq i < M$ and $0 \leq j < N$) is given by the expression $(A \rightarrow \text{data})[0][j \cdot M + i]$

cols - array of pointers. **cols**[*j*] points to the first element of the *j*-th column of the matrix in the array **data**. The (i,j) -th element of a dense matrix **A** of type **DlsMat** (with $0 \leq i < M$ and $0 \leq j < N$) is given by the expression $(A \rightarrow \text{cols})[j][i]$

For the BAND module, the relevant fields of this structure are as follows (see Figure 8.1 for a diagram of the underlying data representation in a banded matrix of type **DlsMat**). Note that only square band matrices are allowed.

type - **SUNDIALS_BAND** (=2)

M - number of rows

N - number of columns ($N = M$)

mu - upper half-bandwidth, $0 \leq \text{mu} < \min(M, N)$

ml - lower half-bandwidth, $0 \leq \text{ml} < \min(M, N)$

s_mu - storage upper bandwidth, $\text{mu} \leq \text{s_mu} < N$. The LU decomposition routine writes the LU factors into the storage for **A**. The upper triangular factor **U**, however, may have an upper bandwidth as big as $\min(N-1, \text{mu} + \text{ml})$ because of partial pivoting. The **s_mu** field holds the upper half-bandwidth allocated for **A**.

ldim - leading dimension ($\text{ldim} \geq \text{s_mu}$)

data - pointer to a contiguous block of **realtype** variables. The elements of a banded matrix of type **DlsMat** are stored columnwise (i.e. columns are stored one on top of the other in memory). Only elements within the specified half-bandwidths are stored. **data** is a pointer to **ldata** contiguous locations which hold the elements within the band of **A**.

ldata - length of the data array ($= \text{ldim} \cdot (\text{s_mu} + \text{ml} + 1)$)

cols - array of pointers. **cols**[*j*] is a pointer to the uppermost element within the band in the *j*-th column. This pointer may be treated as an array indexed from **s_mu** - **mu** (to access the uppermost element within the band in the *j*-th column) to **s_mu** + **ml** (to access the lowest element within the band in the *j*-th column). Indices from 0 to **s_mu** - **mu** - 1 give access to extra storage elements required by the LU decomposition function. Finally, **cols**[*j*][*i* - *j* + **s_mu**] is the (i, j) -th element, $j - \text{mu} \leq i \leq j + \text{ml}$.

8.1.2 Accessor macros for the DLS modules

The macros below allow a user to efficiently access individual matrix elements without writing out explicit data structure references and without knowing too much about the underlying element storage. The only storage assumption needed is that elements are stored columnwise and that a pointer to the *j*-th column of elements can be obtained via the **DENSE_COL** or **BAND_COL** macros. Users should use these macros whenever possible.

The following two macros are defined by the DENSE module to provide access to data in the **DlsMat** type:

- **DENSE_ELEM**

Usage : **DENSE_ELEM**(**A**, *i*, *j*) = **a_ij**; or **a_ij** = **DENSE_ELEM**(**A**, *i*, *j*);

DENSE_ELEM references the (i,j) -th element of the $M \times N$ **DlsMat** **A**, $0 \leq i < M$, $0 \leq j < N$.




Figure 8.1: Diagram of the storage for a banded matrix of type `DlsMat`. Here A is an $N \times N$ band matrix of type `DlsMat` with upper and lower half-bandwidths μ and m , respectively. The rows and columns of A are numbered from 0 to $N-1$ and the (i, j) -th element of A is denoted $A(i, j)$. The greyed out areas of the underlying component storage are used by the `BandGBTRF` and `BandGBTRS` routines.

- DENSE_COL

Usage : `col_j = DENSE_COL(A,j);`

DENSE_COL references the j -th column of the $M \times N$ `DlsMat` `A`, $0 \leq j < N$. The type of the expression `DENSE_COL(A,j)` is `realtype *`. After the assignment in the usage above, `col_j` may be treated as an array indexed from 0 to $M - 1$. The (i, j) -th element of `A` is referenced by `col_j[i]`.

The following three macros are defined by the BAND module to provide access to data in the `DlsMat` type:

- BAND_ELEM

Usage : `BAND_ELEM(A,i,j) = a_ij; or a_ij = BAND_ELEM(A,i,j);`

BAND_ELEM references the (i,j) -th element of the $N \times N$ band matrix `A`, where $0 \leq i, j \leq N - 1$. The location (i,j) should further satisfy $j - (A \rightarrow \text{mu}) \leq i \leq j + (A \rightarrow \text{ml})$.

- BAND_COL

Usage : `col_j = BAND_COL(A,j);`

BAND_COL references the diagonal element of the j -th column of the $N \times N$ band matrix `A`, $0 \leq j \leq N - 1$. The type of the expression `BAND_COL(A,j)` is `realtype *`. The pointer returned by the call `BAND_COL(A,j)` can be treated as an array which is indexed from $-(A \rightarrow \text{mu})$ to $(A \rightarrow \text{ml})$.

- BAND_COL_ELEM

Usage : `BAND_COL_ELEM(col_j,i,j) = a_ij; or a_ij = BAND_COL_ELEM(col_j,i,j);`

This macro references the (i,j) -th entry of the band matrix `A` when used in conjunction with `BAND_COL` to reference the j -th column through `col_j`. The index (i,j) should satisfy $j - (A \rightarrow \text{mu}) \leq i \leq j + (A \rightarrow \text{ml})$.

8.1.3 Functions in the DENSE module

The DENSE module defines two sets of functions with corresponding names. The first set contains functions (with names starting with a capital letter) that act on dense matrices of type `DlsMat`. The second set contains functions (with names starting with a lower case letter) that act on matrices represented as simple arrays.

The following functions for `DlsMat` dense matrices are available in the DENSE package. For full details, see the header files `sundials_direct.h` and `sundials_dense.h`.

- **NewDenseMat**: allocation of a `DlsMat` dense matrix;
- **DestroyMatrix**: free memory for a `DlsMat` matrix;
- **PrintMat**: print a `DlsMat` matrix to standard output.
- **NewIntArray**: allocation of an array of `int` for use as pivots with `DenseGETRF/DenseGETRS`;
- **NewRealArray**: allocation of an array of `realtype` for use as right-hand side with `DenseGETRS`;
- **DestroyArray**: free memory for an array;
- **SetToZero**: load a matrix with zeros;
- **AddIdentity**: increment a square matrix by the identity matrix;
- **DenseCopy**: copy one matrix to another;
- **DenseScale**: scale a matrix by a scalar;

- **DenseGETRF**: LU factorization with partial pivoting;
- **DenseGETRS**: solution of $Ax = b$ using LU factorization (for square matrices A);
- **DensePOTRF**: Cholesky factorization of a real symmetric positive matrix;
- **DensePOTRS**: solution of $Ax = b$ using the Cholesky factorization of A ;
- **DenseGEQRF**: QR factorization of an $m \times n$ matrix, with $m \geq n$;
- **DenseORMQR**: compute the product $w = Qv$, with Q calculated using **DenseGEQRF**;

The following functions for small dense matrices are available in the DENSE package:

- **newDenseMat**
newDenseMat(m,n) allocates storage for an m by n dense matrix. It returns a pointer to the newly allocated storage if successful. If the memory request cannot be satisfied, then **newDenseMat** returns NULL. The underlying type of the dense matrix returned is **realtype****. If we allocate a dense matrix **realtype** a** by **a = newDenseMat(m,n)**, then **a[j][i]** references the (i,j) -th element of the matrix **a**, $0 \leq i < m$, $0 \leq j < n$, and **a[j]** is a pointer to the first element in the j -th column of **a**. The location **a[0]** contains a pointer to $m \times n$ contiguous locations which contain the elements of **a**.
- **destroyMat**
destroyMat(a) frees the dense matrix **a** allocated by **newDenseMat**;
- **newIntArray**
newIntArray(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.
- **newRealArray**
newRealArray(n) allocates an array of n **realtype** values. It returns a pointer to the first element in the array if successful. It returns NULL if the memory request could not be satisfied.
- **destroyArray**
destroyArray(p) frees the array **p** allocated by **newIntArray** or **newRealArray**;
- **denseCopy**
denseCopy(a,b,m,n) copies the m by n dense matrix **a** into the m by n dense matrix **b**;
- **denseScale**
denseScale(c,a,m,n) scales every element in the m by n dense matrix **a** by the scalar **c**;
- **denseAddIdentity**
denseAddIdentity(a,n) increments the *square* n by n dense matrix **a** by the identity matrix I_n ;
- **denseGETRF**
denseGETRF(a,m,n,p) factors the m by n dense matrix **a**, using Gaussian elimination with row pivoting. It overwrites the elements of **a** with its LU factors and keeps track of the pivot rows chosen in the pivot array **p**.

A successful LU factorization leaves the matrix **a** and the pivot array **p** with the following information:

1. **p[k]** contains the row number of the pivot element chosen at the beginning of elimination step k , $k = 0, 1, \dots, n-1$.

2. If the unique LU factorization of \mathbf{a} is given by $\mathbf{Pa} = \mathbf{LU}$, where \mathbf{P} is a permutation matrix, \mathbf{L} is an \mathbf{m} by \mathbf{n} lower trapezoidal matrix with all diagonal elements equal to 1, and \mathbf{U} is an \mathbf{n} by \mathbf{n} upper triangular matrix, then the upper triangular part of \mathbf{a} (including its diagonal) contains \mathbf{U} and the strictly lower trapezoidal part of \mathbf{a} contains the multipliers, $\mathbf{I} - \mathbf{L}$. If \mathbf{a} is square, \mathbf{L} is a unit lower triangular matrix.

`denseGETRF` returns 0 if successful. Otherwise it encountered a zero diagonal element during the factorization, indicating that the matrix \mathbf{a} does not have full column rank. In this case it returns the column index (numbered from one) at which it encountered the zero.

- `denseGETRS`

`denseGETRS(a,n,p,b)` solves the \mathbf{n} by \mathbf{n} linear system $\mathbf{ax} = \mathbf{b}$. It assumes that \mathbf{a} (of size $\mathbf{n} \times \mathbf{n}$) has been LU-factored and the pivot array \mathbf{p} has been set by a successful call to `denseGETRF(a,n,n,p)`. The solution \mathbf{x} is written into the \mathbf{b} array.

- `densePOTRF`

`densePOTRF(a,m)` calculates the Cholesky decomposition of the \mathbf{m} by \mathbf{m} dense matrix \mathbf{a} , assumed to be symmetric positive definite. Only the lower triangle of \mathbf{a} is accessed and overwritten with the Cholesky factor.

- `densePOTRS`

`densePOTRS(a,m,b)` solves the \mathbf{m} by \mathbf{m} linear system $\mathbf{ax} = \mathbf{b}$. It assumes that the Cholesky factorization of \mathbf{a} has been calculated in the lower triangular part of \mathbf{a} by a successful call to `densePOTRF(a,m)`.

- `denseGEQRF`

`denseGEQRF(a,m,n,beta,wrk)` calculates the QR decomposition of the \mathbf{m} by \mathbf{n} matrix \mathbf{a} ($\mathbf{m} \geq \mathbf{n}$) using Householder reflections. On exit, the elements on and above the diagonal of \mathbf{a} contain the \mathbf{n} by \mathbf{n} upper triangular matrix \mathbf{R} ; the elements below the diagonal, with the array `beta`, represent the orthogonal matrix \mathbf{Q} as a product of elementary reflectors. The real array `wrk`, of length \mathbf{m} , must be provided as temporary workspace.

- `denseORMQR`

`denseORMQR(a,m,n,beta,v,w,wrk)` calculates the product $\mathbf{w} = \mathbf{Qv}$ for a given vector \mathbf{v} of length \mathbf{n} , where the orthogonal matrix \mathbf{Q} is encoded in the \mathbf{m} by \mathbf{n} matrix \mathbf{a} and the vector `beta` of length \mathbf{n} , after a successful call to `denseGEQRF(a,m,n,beta,wrk)`. The real array `wrk`, of length \mathbf{m} , must be provided as temporary workspace.

8.1.4 Functions in the BAND module

The BAND module defines two sets of functions with corresponding names. The first set contains functions (with names starting with a capital letter) that act on band matrices of type `DlsMat`. The second set contains functions (with names starting with a lower case letter) that act on matrices represented as simple arrays.

The following functions for `DlsMat` banded matrices are available in the BAND package. For full details, see the header files `sundials_direct.h` and `sundials_band.h`.

- `NewBandMat`: allocation of a `DlsMat` band matrix;
- `DestroyMatrix`: free memory for a `DlsMat` matrix;
- `PrintMat`: print a `DlsMat` matrix to standard output.
- `NewIntArray`: allocation of an array of `int` for use as pivots with `BandGBRF/BandGBRS`;
- `NewRealArray`: allocation of an array of `realtype` for use as right-hand side with `BandGBRS`;

- **DestroyArray**: free memory for an array;
- **SetToZero**: load a matrix with zeros;
- **AddIdentity**: increment a square matrix by the identity matrix;
- **BandCopy**: copy one matrix to another;
- **BandScale**: scale a matrix by a scalar;
- **BandGBTRF**: LU factorization with partial pivoting;
- **BandGBTRS**: solution of $Ax = b$ using LU factorization;

The following functions for small band matrices are available in the BAND package:

- **newBandMat**
`newBandMat(n, smu, ml)` allocates storage for an n by n band matrix with lower half-bandwidth ml .
- **destroyMat**
`destroyMat(a)` frees the band matrix `a` allocated by `newBandMat`;
- **newIntArray**
`newIntArray(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.
- **newRealArray**
`newRealArray(n)` allocates an array of n `realtype` values. It returns a pointer to the first element in the array if successful. It returns NULL if the memory request could not be satisfied.
- **destroyArray**
`destroyArray(p)` frees the array `p` allocated by `newIntArray` or `newRealArray`;
- **bandCopy**
`bandCopy(a,b,n,a_smu, b_smu,copymu, copym1)` copies the n by n band matrix `a` into the n by n band matrix `b`;
- **bandScale**
`bandScale(c,a,n,mu,ml,smu)` scales every element in the n by n band matrix `a` by `c`;
- **bandAddIdentity**
`bandAddIdentity(a,n,smu)` increments the n by n band matrix `a` by the identity matrix;
- **bandGETRF**
`bandGETRF(a,n,mu,ml,smu,p)` factors the n by n band matrix `a`, using Gaussian elimination with row pivoting. It overwrites the elements of `a` with its LU factors and keeps track of the pivot rows chosen in the pivot array `p`.
- **bandGETRS**
`bandGETRS(a,n,smu,ml,p,b)` solves the n by n linear system $ax = b$. It assumes that `a` (of size $n \times n$) has been LU-factored and the pivot array `p` has been set by a successful call to `bandGETRF(a,n,mu,ml,smu,p)`. The solution x is written into the `b` array.



8.2 The SPILS modules: SPGMR, SPBCG, and SPTFQMR

A linear solver module from the *spils* family can only be used in conjunction with an actual NVECTOR implementation library, such as the NVECTOR_SERIAL or NVECTOR_PARALLEL provided with SUNDIALS.

8.2.1 The SPGMR module

The SPGMR package, in the files `sundials_spgmr.h` and `sundials_spgmr.c`, includes an implementation of the scaled preconditioned GMRES method. A separate code module, implemented in `sundials_iterative.(h,c)`, contains auxiliary functions that support SPGMR, as well as the other Krylov solvers in SUNDIALS (SPBCG and SPTFQMR). For full details, including usage instructions, see the header files `sundials_spgmr.h` and `sundials_iterative.h`.

The files comprising the SPGMR generic linear solver, and their locations in the SUNDIALS *srcdir*, are as follows:

- header files (located in *srcdir/include/sundials*)
`sundials_spgmr.h` `sundials_iterative.h` `sundials_nvector.h`
`sundials_types.h` `sundials_math.h` `sundials_config.h`
- source files (located in *srcdir/src/sundials*)
`sundials_spgmr.c` `sundials_iterative.c` `sundials_nvector.c`

Only two of the preprocessing directives in the header file `sundials_config.h` are required to use the SPGMR package by itself (see §A.3 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 nine files listed above can be extracted from the SUNDIALS *srcdir* and compiled by themselves into an SPGMR library or into a larger user code.

The following functions are available in the SPGMR package:

- `SpgmrMalloc`: allocation of memory for `SpgmrSolve`;
- `SpgmrSolve`: solution of $Ax = b$ by the SPGMR method;
- `SpgmrFree`: free memory allocated by `SpgmrMalloc`.

The following functions are available in the support package `sundials_iterative.(h,c)`:

- `ModifiedGS`: performs modified Gram-Schmidt procedure;
- `ClassicalGS`: performs classical Gram-Schmidt procedure;
- `QRfact`: performs QR factorization of Hessenberg matrix;
- `QRsol`: solves a least squares problem with a Hessenberg matrix factored by `QRfact`.

8.2.2 The SPBCG module

The SPBCG package, in the files `sundials_spgm.h` and `sundials_spgm.c`, includes an implementation of the scaled preconditioned Bi-CGStab method. For full details, including usage instructions, see the file `sundials_spgm.h`.

The files needed to use the SPBCG module by itself are the same as for the SPGMR module, but with `sundials_spgm.h(c)` in place of `sundials_spgm.h(c)`.

The following functions are available in the SPBCG package:

- `SpgmMalloc`: allocation of memory for `SpgmSolve`;
- `SpgmSolve`: solution of $Ax = b$ by the SPBCG method;
- `SpgmFree`: free memory allocated by `SpgmMalloc`.

8.2.3 The SPTFQMR module

The SPTFQMR package, in the files `sundials_sptfqr.h` and `sundials_sptfqr.c`, includes an implementation of the scaled preconditioned TFQMR method. For full details, including usage instructions, see the file `sundials_sptfqr.h`.

The files needed to use the SPTFQMR module by itself are the same as for the SPGMR module, but with `sundials_sptfqr.h(c)` in place of `sundials_spgm.h(c)`.

The following functions are available in the SPTFQMR package:

- `SptfqrMalloc`: allocation of memory for `SptfqrSolve`;
- `SptfqrSolve`: solution of $Ax = b$ by the SPTFQMR method;
- `SptfqrFree`: free memory allocated by `SptfqrMalloc`.

Appendix A

CVODE Installation Procedure

The installation of CVODE 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 CVODE.¹

The SUNDIALS suite (or individual solvers) are distributed as compressed archives (`.tar.gz`). The name of the distribution archive is of the form *solver-x.y.z.tar.gz*, where *solver* is one of: **sundials**, **cvode**, **cvodes**, **ida**, **idas**, or **kinsol**, and *x.y.z* represents the version number (of the SUNDIALS suite or of the individual solver). To begin the installation, first uncompress and expand the sources, by issuing

```
% tar xzf solver-x.y.z.tar.gz
```

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

Starting with version 2.4.0 of SUNDIALS, two installation methods are provided: in addition to the previous autotools-based method, SUNDIALS now provides a method based on CMake. Before providing detailed explanations on the installation procedure for the two approaches, we begin with a few common observations:

- 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 (temporary) directory under which SUNDIALS is built.

instdir is the directory under which the SUNDIALS exported header files and libraries will be installed. Typically, header files are exported under a directory *instdir/include* while libraries are installed under *instdir/lib*, with *instdir* specified at configuration time.

- For the CMake-based installation, in-source builds are prohibited; in other words, the build directory *builddir* can **not** be the same as *srcdir* and such an attempt will lead to an error. For autotools-based installation, in-source builds are allowed, although even in that case we recommend using a separate *builddir*. Indeed, this prevents “polluting” the source tree and allows efficient builds for different configurations and/or options.
- The installation directory *instdir* can **not** be the same as the source directory *srcdir*.
- By default, only the libraries and header files are exported to the installation directory *instdir*. If enabled by the user (with the appropriate option to **configure** or toggle for CMake), the



¹Files for both the serial and parallel versions of CVODE 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/`; `nvector_parallel.h` (in `include/nvector/`); `cvode_bbdpre.c`, `cvode_bbdpre_impl.h` (in `src/cvode/`); `cvode_bbdpre.h` (in `include/cvode/`); `fcvbbd.c`, `fcvbbd.h` (in `src/cvode/fcmix/`); all files in `examples/cvode/parallel/`; all files in `examples/cvode/fcmix_parallel/`. (By “serial version” of CVODE we mean the CVODE solver with the serial NVECTOR module attached, and similarly for “parallel version”.)

examples distributed with SUNDIALS will be built together with the solver libraries but the installation step will result in exporting (by default in a subdirectory of the installation directory) the example sources and sample outputs together with automatically generated configuration files that reference the *installed* SUNDIALS headers and libraries. As such, these configuration files for the SUNDIALS examples can be used as "templates" for your own problems. The `configure` script will install makefiles. CMake installs `CMakeLists.txt` files and also (as an option available only under Unix/Linux) makefiles. Note that both installation approaches also allow the option of building the SUNDIALS examples without having to install them. (This can be used as a sanity check for the freshly built libraries.)

- Even if generation of shared libraries is enabled, only static libraries are created for the FCMIX modules. (Because of the use of fixed names for the Fortran user-provided subroutines, FCMIX shared libraries would result in "undefined symbol" errors at link time.)

A.1 Autotools-based installation

The installation procedure outlined below will work on commodity LINUX/UNIX systems without modification. However, users are still encouraged to carefully read this entire section before attempting to install the SUNDIALS suite, in case non-default choices are desired for compilers, compilation options, installation location, etc. The user may invoke the configuration script with the help flag to view a complete listing of available options, by issuing the command

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

from within *srcdir*.

The installation steps for SUNDIALS can be as simple as the following:

```
% cd (...)/srcdir
% ./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, the example programs are not built and installed. 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
```

The above steps are for an "in-source" build. For an "out-of-source" build (recommended), the procedure is simply:

```
% cd (...)/builddir
% (...)/srcdir/configure
% make
% make install
```

Note that, in this case, `make clean` and `make distclean` are irrelevant. Indeed, if disk space is a priority, the entire *builddir* can be purged after the installation completes. For a new install, a new *builddir* directory can be created and used.

A.1.1 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`, `idas`, 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** defined instead 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.

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 Blas/Lapack support

The configure script will attempt to automatically determine the proper libraries to be linked for support of the new Blas/Lapack linear solver module. If these are not found, or if Blas and/or Lapack libraries are installed in a non-standard location, the following options can be used:

`--with-blas`

Specify the Blas library.

Default: none

`--with-lapack`

Specify the Lapack library.

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.

A.1.2 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, and install libraries, headers, and example sources under appropriate subdirectories of `/home/myname/sundials/`, use

```
% configure --prefix=/home/myname/sundials --enable-examples
```

To disable installation of the examples, use:

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

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 `--with-mpi-libs` option is required so that the configure script can check if `gcc` can link with the appropriate MPI library. The `--disable-lapack` option explicitly disables support for Blas/Lapack, while the `--disable-fcmix` explicitly disables building the FCMIX interfaces. Note that, because of the last two options, no Fortran-related settings are checked for.

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

Finally, a minimal configuration and installation of SUNDIALS in `/home/myname/sundials/` (serial only, no Fortran support, no examples) can be obtained with:

```
% configure --prefix=/home/myname/sundials \
--disable-mpi --disable-lapack --disable-fcmix
```

A.2 CMake-based installation

Support for CMake-based installation has been added to SUNDIALS primarily to provide a platform-independent build system. Like autotools, CMake can generate a Unix Makefile. Unlike autotools, CMake can also create KDevelop, Visual Studio, and (Apple) XCode project files from the same configuration file. In addition, CMake provides a GUI front end and therefore the installation process is more interactive than when using autotools.

The installation options are very similar to the options mentioned above (although their default values may differ slightly). Practically, all configurations supported by the autotools-based installation

approach are also possible with CMake, the only notable exception being cross-compilation, which is currently not implemented in the CMake approach.

The SUNDIALS build process requires CMake version 2.4.x or higher and a working compiler. On Unix-like operating systems, it also requires Make (and `curses`, including its development libraries, for the GUI front end to CMake, `ccmake`), while on Windows it requires Visual Studio. While many Linux distributions offer CMake, the version included is probably out of date. Many new CMake features have been added recently, and you should download the latest version from <http://www.cmake.org/HTML/Download.html>. Build instructions for Cmake (only necessary for Unix-like systems) can be found on the CMake website. Once CMake is installed, Linux/Unix user will be able to use `ccmake`, while Windows user will be able to use `CMakeSetup`.

As noted above, when using CMake to configure, build and install SUNDIALS, it is always required to use a separate build directory. While in-source builds are possible, they are explicitly prohibited by the SUNDIALS CMake scripts (one of the reasons being that, unlike autotools, CMake does not provide a `make distclean` procedure and it is therefore difficult to clean-up the source tree after an in-source build).

A.2.1 Configuring, building, and installing on Unix-like systems

Use `ccmake` from the CMake installed location. `ccmake` is a Curses based GUI for CMake. To run it go to the build directory and specify as an argument the build directory:

```
% mkdir (...)/builddir
% cd (...)/builddir
% ccmake (...)/srcdir
```

About `ccmake`:

- Iterative process
 - Select values, run configure (`c` key)
 - Set the settings, run configure, set the settings, run configure, etc.
- Repeat until all values are set and the generate option is available (`g` key)
- Some variables (advanced variables) are not visible right away
- To see advanced variables, toggle to advanced mode (`t` key)
- To set a variable, move the cursor to the variable and press enter
 - If it is a boolean (ON/OFF) it will flip the value
 - If it is string or file, it will allow editing of the string
 - For file and directories, the `<tab>` key can be used to complete
- To search for a variable press `/` key, and to repeat the search, press the `n` key

CMake will now generate makefiles including all dependencies and all rules to build SUNDIALS on this system. You should not, however, try to move the build directory to another location on this system or to another system. Once you have makefiles you should be able to just type:

```
% make
```

To install SUNDIALS in the installation directory specified at configuration time, simply run

```
% make install
```

A.2.2 Configuring, building, and installing on Windows

Use **CMakeSetup** from the CMake install location. Make sure to select the appropriate source and the build directory. Also, make sure to pick the appropriate generator (on Visual Studio 6, pick the Visual Studio 6 generator). Some CMake versions will ask you to select the generator the first time you press Configure instead of having a drop-down menu in the main dialog.

About **CMakeSetup**:

- Iterative process
 - Select values, press the Configure button
 - Set the settings, run configure, set the settings, run configure, etc.
- Repeat until all values are set and the OK button becomes available.
- Some variables (advanced variables) are not visible right away
- To see advanced variables, toggle to advanced mode ("Show Advanced Values" toggle).
- To set the value of a variable, click on that value.
 - If it is boolean (ON/OFF), a drop-down menu will appear for changing the value.
 - If it is file or directory, an ellipsis button will appear ("...") on the far right of the entry. Clicking this button will bring up the file or directory selection dialog.
 - If it is a string, it will become an editable string.

CMake will now create Visual Studio project files. You should now be able to open the SUNDIALS project (or workspace) file. Make sure to select the appropriate build type (Debug, Release, ...). To build SUNDIALS, simply build the **ALL_BUILD** target. To install SUNDIALS, simply run the **INSTALL** target within the build system.

A.2.3 Configuration options

A complete list of all available options for a CMake-based SUNDIALS configuration is provide below. Note that the default values shown are for a typical configuration on a Linux system and are provided as illustration only. Some of them will be different on different systems.

BUILD_CVODE - Build the CVODE library
Default: ON

BUILD_CVODES - Build the CVODES library
Default: ON

BUILD_IDA - Build the IDA library
Default: ON

BUILD_IDAS - Build the IDAS library
Default: ON

BUILD_KINSOL - Build the KINSOL library
Default: ON

BUILD_SHARED_LIBS - Build shared libraries
Default: OFF

BUILD_STATIC_LIBS - Build static libraries
Default: ON

CMAKE_BUILD_TYPE - Choose the type of build, options are: None (CMAKE_C_FLAGS used) Debug Release RelWithDebInfo MinSizeRel
Default:

CMAKE_C_COMPILER - C compiler
Default: /usr/bin/gcc

CMAKE_C_FLAGS - Flags for C compiler
Default:

CMAKE_C_FLAGS_DEBUG - Flags used by the compiler during debug builds
Default: -g

CMAKE_C_FLAGS_MINSIZEREL - Flags used by the compiler during release minsize builds
Default: -Os -DNDEBUG

CMAKE_C_FLAGS_RELEASE - Flags used by the compiler during release builds
Default: -O3 -DNDEBUG

CMAKE_BACKWARDS_COMPATIBILITY - For backwards compatibility, what version of CMake commands and syntax should this version of CMake allow.
Default: 2.4

CMAKE_Fortran_COMPILER - Fortran compiler
Default: /usr/bin/g77
Note: Fortran support (and all related options) are triggered only if either Fortran-C support is enabled (FCMIX_ENABLE is ON) or Blas/Lapack support is enabled (LAPACK_ENABLE is ON).

CMAKE_Fortran_FLAGS - Flags for Fortran compiler
Default:

CMAKE_Fortran_FLAGS_DEBUG - Flags used by the compiler during debug builds
Default:

CMAKE_Fortran_FLAGS_MINSIZEREL - Flags used by the compiler during release minsize builds
Default:

CMAKE_Fortran_FLAGS_RELEASE - Flags used by the compiler during release builds
Default:

CMAKE_INSTALL_PREFIX - Install path prefix, prepended onto install directories
Default: /usr/local
Note: The user must have write access to the location specified through this option. Exported SUNDIALS header files and libraries will be installed under subdirectories **include** and **lib** of CMAKE_INSTALL_PREFIX, respectively.

EXAMPLES_ENABLE - Build the SUNDIALS examples
Default: OFF
Note: setting this option to ON will trigger additional options related to how and where example programs will be installed.

EXAMPLES_GENERATE_MAKEFILES - Create Makefiles for building the examples
Default: ON
Note: This option is triggered only if enabling the building and installing of the example programs (i.e., both **EXAMPLES_ENABLE** and **EXAMPLES_INSTALL** are set to ON) and if configuration is done on a Unix-like system. If enabled, makefiles for the compilation of the example programs (using the installed SUNDIALS libraries) will be automatically generated and exported to the directory specified by **EXAMPLES_INSTALL_PATH**.

EXAMPLES_INSTALL - Install example files

Default: ON

Note: This option is triggered only if building example programs is enabled (**EXAMPLES_ENABLE** ON). If the user requires installation of example programs then the sources and sample output files for all SUNDIALS modules that are currently enabled will be exported to the directory specified by **EXAMPLES_INSTALL_PATH**. A CMake configuration script will also be automatically generated and exported to the same directory. Additionally, if the configuration is done under a Unix-like system, an additional option (**EXAMPLES_GENERATE_MAKEFILES**) will be triggered.

EXAMPLES_INSTALL_PATH - Output directory for installing example files

Default: /usr/local/examples

Note: The actual default value for this option will be an **examples** subdirectory created under **CMAKE_INSTALL_PREFIX**.

EXAMPLES_USE_STATIC_LIBS - Link examples using the static libraries

Default: OFF

Note: This option is triggered only if building shared libraries is enabled (**BUILD_SHARED_LIBS** is ON).

FCMIX_ENABLE - Enable Fortran-C support

Default: OFF

LAPACK_ENABLE - Enable Lapack support

Default: OFF

Note: Setting this option to ON will trigger the two additional options see below.

LAPACK_LIBRARIES - Lapack (and Blas) libraries

Default: /usr/lib/liblapack.so;/usr/lib/libblas.so

LAPACK_LINKER_FLAGS - Lapack (and Blas) required linker flags

Default: -lg2c

MPI_ENABLE - Enable MPI support

Default: OFF

Note: Setting this option to ON will trigger several additional options related to MPI.

MPI_MPICC - mpicc program

Default: /home/radu/apps/mpich1/gcc/bin/mpicc

Note: This option is triggered only if using MPI compiler scripts (**MPI_USE_MPISCRIPTS** is ON).

MPI_MPIF77 - mpif77 program

Default: /home/radu/apps/mpich1/gcc/bin/mpif77

Note: This option is triggered only if using MPI compiler scripts (**MPI_USE_MPISCRIPTS** is ON) and Fortran-C support is enabled (**FCMIX_ENABLE** is ON).

MPI_INCLUDE_PATH - Path to MPI header files

Default: /home/radu/apps/mpich1/gcc/include

Note: This option is triggered only if not using MPI compiler scripts (**MPI_USE_MPISCRIPTS** is ON).

MPI_LIBRARIES - MPI libraries

Default: /home/radu/apps/mpich1/gcc/lib/libmpich.a

Note: This option is triggered only if not using MPI compiler scripts (**MPI_USE_MPISCRIPTS** is ON).

MPI_USE_MPISCRIPTS - Use MPI compiler scripts

Default: ON

SUNDIALS_PRECISION - Precision used in SUNDIALS, options are: double, single or extended
Default: double

USE_GENERIC_MATH - Use generic (stdc) math libraries
Default: ON

A.3 Manually building SUNDIALS

With the addition of CMake support, the installation of the SUNDIALS package on almost any platform was greatly simplified. However, if for whatever reason, neither of the two procedures described above is convenient (for example for users who prefer to own the build process or otherwise incorporate SUNDIALS or one of its solvers in a larger project with its own build system), we provide here a few directions for a completely manual installation.

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** or CMake scripts) is provided below.

```

1  /* SUNDIALS configuration header file */
2
3  #define SUNDIALS_PACKAGE_VERSION "2.4.0"
4
5  #define F77_FUNC(name,NAME) name ## _
6  #define F77_FUNC_(name,NAME) name ## _
7
8  #define SUNDIALS_DOUBLE_PRECISION 1
9
10 #define SUNDIALS_USE_GENERIC_MATH 1
11
12 #define SUNDIALS_MPLCOMM_F2C 1
13
14 #define SUNDIALS_EXPORT

```

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.

- Use of generic math functions

If `SUNDIALS_USE_GENERIC_MATH` is defined, then the functions in `sundials_math.(h,c)` 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 is specified by appropriately defining the following parameterized macros (using the stringization operator, `##`, if necessary):

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

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

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

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.

- Mark SUNDIALS API functions for export/import. When building shared SUNDIALS libraries under Windows, use

```
#define SUNDIALS_EXPORT __declspec(dllexport)
```

When linking to shared SUNDIALS libraries under Windows, use

```
#define SUNDIALS_EXPORT __declspec(dllimport)
```

In all other cases (other platforms or static libraries under Windows), the `SUNDIALS_EXPORT` macro is empty.

A.4 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 §A.1) or the appropriate CMake options (see §A.2). 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 A.1. The file extension *.lib* is typically *.so* for shared libraries and *.a* for static libraries. Note that, in Table A.1, names are relative to *libdir* for libraries and to *includedir* for header files.

A typical user program need not explicitly include any of the shared SUNDIALS header files from under the *includedir/sundials* directory since they are explicitly included by the appropriate solver header files (*e.g.*, *cnode_dense.h* includes *sundials_dense.h*). However, it is both legal and safe to do so (*e.g.*, the functions declared in *sundials_dense.h* could be used in building a preconditioner).

Table A.1: SUNDIALS libraries and header files

SHARED	Libraries	n/a	
	Header files	sundials/sundials_config.h sundials/sundials_math.h sundials/sundials_nvector.h sundials/sundials_direct.h sundials/sundials_dense.h sundials/sundials_iterative.h sundials/sundials_spgmrs.h sundials/sundials_spgmrs.h	sundials/sundials_types.h sundials/sundials_fnvector.h sundials/sundials_lapack.h sundials/sundials_band.h sundials/sundials_spgmrs.h sundials/sundials_sptfqmr.h
NVECTOR_SERIAL	Libraries	libsundials_nvecserial.lib	libsundials_fnvecserial.a
	Header files	nvector/nvector_serial.h	
NVECTOR_PARALLEL	Libraries	libsundials_nvecparallel.lib	libsundials_fnvecparallel.a
	Header files	nvector/nvector_parallel.h	
CVODE	Libraries	libsundials_cvode.lib	libsundials_fcvcde.a
	Header files	cvode/cvode.h cvode/cvode_direct.h cvode/cvode_dense.h cvode/cvode_diag.h cvode/cvode_spils.h cvode/cvode_sptfqmr.h cvode/cvode_bandpre.h	cvode/cvode_impl.h cvode/cvode_lapack.h cvode/cvode_band.h cvode/cvode_spgmrs.h cvode/cvode_spgmrs.h cvode/cvode_spgmrs.h cvode/cvode_spgmrs.h
CVODES	Libraries	libsundials_cvodes.lib	
	Header files	cvodes/cvodes.h cvodes/cvodes_direct.h cvodes/cvodes_dense.h cvodes/cvodes_diag.h cvodes/cvodes_spils.h cvodes/cvodes_sptfqmr.h cvodes/cvodes_bandpre.h	cvodes/cvodes_impl.h cvodes/cvodes_lapack.h cvodes/cvodes_band.h cvodes/cvodes_spgmrs.h cvodes/cvodes_spgmrs.h cvodes/cvodes_spgmrs.h cvodes/cvodes_spgmrs.h
IDA	Libraries	libsundials_ida.lib	libsundials_fida.a
	Header files	ida/ida.h ida/ida_direct.h ida/ida_dense.h ida/ida_spils.h ida/ida_spgmrs.h ida/ida_sptfqmr.h	ida/ida_impl.h ida/ida_lapack.h ida/ida_band.h ida/ida_spgmrs.h ida/ida_spgmrs.h ida/ida_sptfqmr.h
IDAS	Libraries	libsundials_idas.lib	
	Header files	idas/idas.h idas/idas_direct.h idas/idas_dense.h idas/idas_spils.h idas/idas_spgmrs.h idas/idas_sptfqmr.h	idas/idas_impl.h idas/idas_lapack.h idas/idas_band.h idas/idas_spgmrs.h idas/idas_spgmrs.h idas/idas_sptfqmr.h
KINSOL	Libraries	libsundials_kinsol.lib	libsundials_fkinsol.a
	Header files	kinsol/kinsol.h kinsol/kinsol_direct.h kinsol/kinsol_dense.h kinsol/kinsol_spils.h kinsol/kinsol_spgmrs.h kinsol/kinsol_sptfqmr.h	kinsol/kinsol_impl.h kinsol/kinsol_lapack.h kinsol/kinsol_band.h kinsol/kinsol_spgmrs.h kinsol/kinsol_spgmrs.h kinsol/kinsol_sptfqmr.h

Appendix B

CVODE Constants

Below we list all input and output constants used by the main solver and linear solver modules, together with their numerical values and a short description of their meaning.

B.1 CVODE input constants

CVODE 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_NORMAL	1	Solver returns at specified output time.
CV_ONE_STEP	2	Solver returns after each successful step.
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.

B.2 CVODE output constants

CVODE 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_WARNING	99	CVode succeeded but an unusual situation occurred.
CV_TOO_MUCH_WORK	-1	The solver took mxstep 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.

CVDSL linear solver modules

CVDSL_SUCCESS	0	Successful function return.
CVDSL_MEM_NULL	-1	The <code>cvode_mem</code> argument was NULL.
CVDSL_LMEM_NULL	-2	The CVDSL linear solver has not been initialized.
CVDSL_ILL_INPUT	-3	The CVDSL solver is not compatible with the current NVECTOR module.
CVDSL_MEM_FAIL	-4	A memory allocation request failed.
CVDSL_JACFUNC_UNRECV	-5	The Jacobian function failed in an unrecoverable manner.
CVDSL_JACFUNC_RECV	-6	The Jacobian function had a recoverable error.

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_INV_FAIL	-5	A diagonal element of the Jacobian was 0.
CVDIAG_RHSFUNC_UNRECV	-6	The right-hand side function failed in an unrecoverable manner.

CVDIAG_RHSFUNC_RECVR	-7	The right-hand side function had a recoverable error.
<hr/>		
CVSPILS linear solver modules		
<hr/>		
CVSPILS_SUCCESS	0	Successful function return.
CVSPILS_MEM_NULL	-1	The <code>cvode_mem</code> argument was NULL.
CVSPILS_LMEM_NULL	-2	The CVSPILS linear solver has not been initialized.
CVSPILS_ILL_INPUT	-3	The CVSPILS 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_PMEM_NULL	-5	The preconditioner module has not been initialized.
<hr/>		
SPGMR generic linear solver module		
<hr/>		
SPGMR_SUCCESS	0	Converged.
SPGMR_RES_REDUCED	1	No convergence, but the residual norm was reduced.
SPGMR_CONV_FAIL	2	Failure to converge.
SPGMR_QRFACT_FAIL	3	A singular matrix was found during the QR factorization.
SPGMR_PSOLVE_FAIL_REC	4	The preconditioner solve function failed recoverably.
SPGMR_ATIMES_FAIL_REC	5	The Jacobian-times-vector function failed recoverably.
SPGMR_PSET_FAIL_REC	6	The preconditioner setup 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.

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. User Documentation for PVODE, An ODE Solver for Parallel Computers. Technical Report UCRL-ID-130884, LLNL, May 1998.
- [6] 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.
- [7] S. D. Cohen and A. C. Hindmarsh. CVODE User Guide. Technical Report UCRL-MA-118618, LLNL, September 1994.
- [8] S. D. Cohen and A. C. Hindmarsh. CVODE, a Stiff/Nonstiff ODE Solver in C. *Computers in Physics*, 10(2):138–143, 1996.
- [9] R. W. Freund. A Transpose-Free Quasi-Minimal Residual Algorithm for Non-Hermitian Linear Systems. *SIAM J. Sci. Comp.*, 14:470–482, 1993.
- [10] K. L. Hiebert and L. F. Shampine. Implicitly Defined Output Points for Solutions of ODEs. Technical Report SAND80-0180, Sandia National Laboratories, February 1980.
- [11] 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.
- [12] A. C. Hindmarsh. Avoiding BDF Stability Barriers in the MOL Solution of Advection-Dominated Problems. *Appl. Num. Math.*, 17:311–318, 1995.
- [13] A. C. Hindmarsh. The PVODE and IDA Algorithms. Technical Report UCRL-ID-141558, LLNL, December 2000.
- [14] 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.
- [15] A. C. Hindmarsh and R. Serban. Example Programs for CVODE v2.6.0. Technical report, LLNL, 2008. UCRL-SM-208110.

-
- [16] 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.
 - [17] 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.
 - [18] 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.
 - [19] 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.
 - [20] 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, 5
- BAND generic linear solver
 - functions, 105–106
 - small matrix, 106
 - macros, 103
 - type DlsMat, 100–101
- BAND_COL, 58, 103
- BAND_COL_ELEM, 58, 103
- BAND_ELEM, 58, 103
- bandAddIdentity, 106
- bandCopy, 106
- bandGETRF, 106
- bandGETRS, 106
- bandScale, 106
- BDF method, 5
- Bi-CGSTab method, 27, 38, 108
- BIG_REAL, 18, 89
- CLASSICAL_GS, 38
- CV_ADAMS, 21, 54
- CV_BAD_DKY, 40
- CV_BAD_K, 40
- CV_BAD_T, 40
- CV_BDF, 21, 54
- CV_CONV_FAILURE, 29
- CV_ERR_FAILURE, 29
- CV_FIRST_RHSFUNC_ERR, 55
- CV_FIRST_RHSFUNC_FAIL, 29
- CV_FUNCTIONAL, 21, 35
- CV_ILL_INPUT, 22, 23, 29, 32–35, 39, 54
- CV_LINIT_FAIL, 29
- CV_LSETUP_FAIL, 29, 57, 58, 64, 65
- CV_LSOLVE_FAIL, 29
- CV_MEM_FAIL, 22
- CV_MEM_NULL, 22, 23, 28, 31–35, 39, 40, 42–48, 54
- CV_NEWTON, 21, 35
- CV_NO_MALLOC, 22, 23, 28, 54
- CV_NORMAL, 28
- CV_ONE_STEP, 28
- CV_REPTD_RHSFUNC_ERR, 29
- CV_RHSFUNC_FAIL, 29, 55
- CV_ROOT_RETURN, 28
- CV_RTFUNC_FAIL, 29, 56
- CV_SUCCESS, 22, 23, 28, 31–35, 39, 40, 42–48, 54
- CV_TOO_CLOSE, 29
- CV_TOO_MUCH_ACC, 29
- CV_TOO_MUCH_WORK, 29
- CV_TSTOP_RETURN, 28
- CV_UNREC_RHSFUNC_ERR, 29, 55
- CV_WARNING, 55
- CVBAND linear solver
 - Jacobian approximation used by, 36
 - memory requirements, 48
 - NVECTOR compatibility, 25
 - optional input, 35–36
 - optional output, 48–49
 - selection of, 25
 - use in FCVODE, 74
- CVBand, 20, 24, 25, 57
- CVBANDPRE preconditioner
 - description, 61
 - optional output, 62–63
 - usage, 61–62
 - user-callable functions, 62
- CVBandPrecGetNumRhsEvals, 63
- CVBandPrecGetWorkSpace, 62
- CVBandPrecInit, 62
- CVBBDPRE preconditioner
 - description, 63–64
 - optional output, 67–68
 - usage, 65–66
 - user-callable functions, 66–67
 - user-supplied functions, 64–65
- CVBBDPrecGetNumGfnEvals, 68
- CVBBDPrecGetWorkSpace, 67
- CVBBDPrecInit, 66
- CVBBDPrecReInit, 67
- CVDENSE linear solver
 - Jacobian approximation used by, 35
 - memory requirements, 48
 - NVECTOR compatibility, 25
 - optional input, 35–36
 - optional output, 48–49
 - selection of, 25
 - use in FCVODE, 73
- CVDense, 20, 24, 25, 56

- CVDIAG linear solver
 - Jacobian approximation used by, 26
 - memory requirements, 49
 - optional output, 49–51
 - selection of, 26
 - use in FCVODE, 73
- CVDiag, 20, 24, 26
- CVDIAG_ILL_INPUT, 26
- CVDIAG_LMEM_NULL, 50
- CVDIAG_MEM_FAIL, 26
- CVDIAG_MEM_NULL, 26, 50
- CVDIAG_SUCCESS, 26, 50
- CVDiagGetLastFlag, 50
- CVDiagGetNumRhsEvals, 50
- CVDiagGetReturnFlagName, 51
- CVDiagGetWorkSpace, 50
- CVDLS_ILL_INPUT, 25, 26
- CVDLS_JACFUNC_RECVR, 57, 58
- CVDLS_JACFUNC_UNRECVR, 57, 58
- CVDLS_LMEM_NULL, 36, 48, 49
- CVDLS_MEM_FAIL, 25, 26
- CVDLS_MEM_NULL, 25, 26, 36, 48, 49
- CVDLS_SUCCESS, 25, 26, 36, 48, 49
- CVDlsBandJacFn, 57
- CVDlsDenseJacFn, 56
- CVDlsGetLastFlag, 49
- CVDlsGetNumJacEvals, 48
- CVDlsGetNumRhsEvals, 49
- CVDlsGetReturnFlagName, 49
- CVDlsGetWorkSpace, 48
- CVDlsSetBandJacFn, 36
- CVDlsSetDenseJacFn, 35
- CVErrorHandlerFn, 55
- CVewtFn, 55
- CVLapackBand, 20, 24, 26, 57
- CVLapackDense, 20, 24, 25, 56
- CVODE, 1
 - motivation for writing in C, 2
 - package structure, 13
 - relationship to CVODE, PODE, 1–2
 - relationship to VODE, VODPK, 1
- CVODE linear solvers
 - built on generic solvers, 25
 - CVBAND, 25
 - CVDENSE, 25
 - CVDIAG, 26
 - CVSPBCG, 27
 - CVSPGMR, 26
 - CVSPTFQMR, 27
 - header files, 18
 - implementation details, 15
 - list of, 13–15
 - NVECTOR compatibility, 17
 - selecting one, 24–25
 - CVode, 21, 28
 - cvode.h, 18
 - cvode_band.h, 18
 - cvode_dense.h, 18
 - cvode_diag.h, 19
 - cvode_lapack.h, 19
 - cvode_spgmr.h, 19
 - cvode_sptfqmr.h, 19
 - CVodeCreate, 21
 - CVodeFree, 21, 22
 - CVodeGetActualInitStep, 44
 - CVodeGetCurrentOrder, 43
 - CVodeGetCurrentStep, 44
 - CVodeGetCurrentTime, 44
 - CVodeGetDky, 39
 - CVodeGetErrWeights, 45
 - CVodeGetEstLocalErrors, 45
 - CVodeGetIntegratorStats, 46
 - CVodeGetLastOrder, 43
 - CVodeGetLastStep, 44
 - CVodeGetNonlinSolvStats, 47
 - CVodeGetNumErrTestFails, 43
 - CVodeGetNumGEvals, 48
 - CVodeGetNumLinSolvSetups, 43
 - CVodeGetNumNonlinSolvConvFails, 46
 - CVodeGetNumNonlinSolvIters, 46
 - CVodeGetNumRhsEvals, 42
 - CVodeGetNumStabLimOrderReds, 45
 - CVodeGetNumSteps, 42
 - CVodeGetReturnFlagName, 47
 - CVodeGetRootInfo, 47
 - CVodeGetTolScaleFactor, 45
 - CVodeGetWorkSpace, 40
 - CVodeInit, 22, 54
 - CVodeReInit, 54
 - CVodeRootInit, 28
 - CVodeSetErrFile, 31
 - CVodeSetErrorHandlerFn, 31
 - CVodeSetInitStep, 33
 - CVodeSetIterType, 35
 - CVodeSetMaxConvFails, 34
 - CVodeSetMaxErrTestFails, 34
 - CVodeSetMaxHnilWarns, 32
 - CVodeSetMaxNonlinIters, 34
 - CVodeSetMaxNumSteps, 32
 - CVodeSetMaxOrder, 31
 - CVodeSetMaxStep, 33
 - CVodeSetMinStep, 33
 - CVodeSetNoInactiveRootWarn, 39
 - CVodeSetNonlinConvCoef, 35
 - CVodeSetRootDirection, 39
 - CVodeSetStabLimDet, 32
 - CVodeSetStopTime, 34

- CVodeSetUserData, 31
- CVodeSStolerances, 22
- CVodeSVtolerances, 23
- CVodeWftolerances, 23
- CVRhsFn, 22, 54
- CVRootFn, 56
- CVSPBCG linear solver
 - Jacobian approximation used by, 36
 - memory requirements, 51
 - optional input, 36–38
 - optional output, 51–54
 - preconditioner setup function, 36, 60
 - preconditioner solve function, 36, 59
 - selection of, 27
 - use in FCVODE, 75
- CVSpbcg, 20, 24, 27
- CVSPGMR linear solver
 - Jacobian approximation used by, 36
 - memory requirements, 51
 - optional input, 36–38
 - optional output, 51–54
 - preconditioner setup function, 36, 60
 - preconditioner solve function, 36, 59
 - selection of, 26
 - use in FCVODE, 75
- CVSpgrmr, 20, 24, 26
- CVSPILS_ILL_INPUT, 27, 37, 38, 62, 66
- CVSPILS_LMEM_NULL, 37, 38, 51–53, 62, 66, 67
- CVSPILS_MEM_FAIL, 27, 62, 66
- CVSPILS_MEM_NULL, 27, 37, 38, 51–53
- CVSPILS_PMEM_NULL, 63, 67, 68
- CVSPILS_SUCCESS, 27, 37, 38, 51–53, 63
- CVSpilsGetLastFlag, 53
- CVSpilsGetNumConvFails, 52
- CVSpilsGetNumJtimesEvals, 52
- CVSpilsGetNumLinIters, 51
- CVSpilsGetNumPrecEvals, 52
- CVSpilsGetNumPrecSolves, 52
- CVSpilsGetNumRhsEvals, 53
- CVSpilsGetReturnFlagName, 53
- CVSpilsGetWorkSpace, 51
- CVSpilsJacTimesVecFn, 59
- CVSpilsPrecSetupFn, 60
- CVSpilsPrecSolveFn, 59
- CVSpilsSetEpsLin, 38
- CVSpilsSetGSType, 38
- CVSpilsSetJacTimesFn, 37
- CVSpilsSetMaxl, 38
- CVSpilsSetPreconditioner, 37
- CVSpilsSetPrecType, 37
- CVSPTFQMR linear solver
 - Jacobian approximation used by, 36
 - memory requirements, 51
 - optional input, 36–38
 - optional output, 51
 - preconditioner setup function, 36, 60
 - preconditioner solve function, 36, 59
 - selection of, 27
 - use in FCVODE, 75
- CVSptfqmr, 20, 24, 27
- DENSE generic linear solver
 - functions
 - large matrix, 103–104
 - small matrix, 104–105
 - macros, 101–103
 - type DlsMat, 100–101
- DENSE_COL, 57, 103
- DENSE_ELEM, 57, 101
- denseAddIdentity, 104
- denseCopy, 104
- denseGEQRF, 105
- denseGETRF, 104
- denseGETRS, 105
- denseORMQR, 105
- densePOTRF, 105
- densePOTRS, 105
- denseScale, 104
- destroyArray, 104, 106
- destroyMat, 104, 106
- DlsMat, 57, 58, 100
- eh_data, 55
- error control
 - order selection, 8
 - step size selection, 7–8
- error messages, 30
 - redirecting, 30
 - user-defined handler, 31, 55
- FCVBAND, 74
- FCVBANDSETJAC, 74
- FCVBBDINIT, 83
- FCVBBDOPT, 83
- FCVBBDREINIT, 84
- FCVBJAC, 74
- FCVBPINIT, 82
- FCVBPOPT, 82
- FCVCOMMFN, 84
- FCVDENSE, 73
- FCVDENSESETJAC, 73
- FCVDIAG, 73
- FCVDJAC, 73
- FCVDKY, 77
- FCVEWT, 72
- FCVEWTSET, 72
- FCVFREE, 78
- FCVFUN, 71
- FCVGETERRWEIGHTS, 79

- FCVGETESTLOCALERR, 79
- FCVGLOCFN, 84
- FCVJTIMES, 76, 84
- FCVMALLOC, 72
- FCVMALLOC, 71
- FCVODE, 77
- FCVODE interface module
 - interface to the CVBANDPRE module, 81–82
 - interface to the CVBBDPRE module, 82–84
 - optional input and output, 78
 - rootfinding, 79–81
 - usage, 71–78
 - user-callable functions, 69–70
 - user-supplied functions, 70
- FCVPSET, 76
- FCVPSOL, 76
- FCVREINIT, 77
- FCVSETIIN, 78
- FCVSETRIN, 78
- FCVSPBCG, 75
- FCVSPBCGREINIT, 78
- FCVSPGMR, 75
- FCVSPGMRREINIT, 78
- FCVSPILSSETJAC, 75, 82, 83
- FCVSPILSSETPREC, 75
- FCVSPTFQMR, 75
- FCVSPTFQMRREINIT, 78
- FNVINITP, 71
- FNVINITS, 71
- generic linear solvers
 - BAND, 99
 - DENSE, 99
 - SPBCG, 108
 - SPGMR, 107
 - SPTFQMR, 108
 - use in CVODE, 15–16
- GMRES method, 27, 107
- Gram-Schmidt procedure, 38
- half-bandwidths, 25, 57–58, 62, 66
- header files, 18, 61, 65
- HNIL_WARNINGS, 79
- INIT_STEP, 79
- IOUT, 78, 80
- itask, 21, 28
- iter, 21, 35
- Jacobian approximation function
 - band
 - difference quotient, 36
 - use in FCVODE, 74
 - user-supplied, 36, 57–58
 - dense
 - difference quotient, 35
 - use in FCVODE, 73
 - user-supplied, 35, 56–57
 - diagonal
 - difference quotient, 26
 - Jacobian times vector
 - difference quotient, 36
 - use in FCVODE, 76
 - user-supplied, 37
 - Jacobian-vector product
 - user-supplied, 59
- linit, 95
- lmm, 21, 54
- LSODE, 1
- MAX_CONVFAIL, 79
- MAX_ERRFAIL, 79
- MAX_NITERS, 79
- MAX_NSTEPS, 79
- MAX_ORD, 79
- MAX_STEP, 79
- maxl, 26, 27
- maxord, 32, 54
- memory requirements
 - CVBAND linear solver, 48
 - CVBANDPRE preconditioner, 63
 - CVBBDPRE preconditioner, 67
 - CVDENSE linear solver, 48
 - CVDIAG linear solver, 49
 - CVODE solver, 42
 - CVSPGMR linear solver, 51
- MIN_STEP, 79
- MODIFIED_GS, 38
- MPI, 4
- N_VCloneEmptyVectorArray, 86
- N_VCloneEmptyVectorArray_Parallel, 92
- N_VCloneEmptyVectorArray_Serial, 90
- N_VCloneVectorArray, 86
- N_VCloneVectorArray_Parallel, 92
- N_VCloneVectorArray_Serial, 90
- N_VDestroyVectorArray, 86
- N_VDestroyVectorArray_Parallel, 93
- N_VDestroyVectorArray_Serial, 90
- N_Vector, 18, 85
- N_VMake_Parallel, 92
- N_VMake_Serial, 90
- N_VNew_Parallel, 92
- N_VNew_Serial, 90
- N_VNewEmpty_Parallel, 92
- N_VNewEmpty_Serial, 90
- N_VPrint_Parallel, 93
- N_VPrint_Serial, 90
- newBandMat, 106

- `newDenseMat`, 104
- `newIntArray`, 104, 106
- `newRealArray`, 104, 106
- `NLCONV_COEF`, 79
- nonlinear system
 - definition, 5
 - Newton convergence test, 7
 - Newton iteration, 6–7
- `NV_COMM_P`, 92
- `NV_CONTENT_P`, 91
- `NV_CONTENT_S`, 89
- `NV_DATA_P`, 91
- `NV_DATA_S`, 89
- `NV_GLOBLENGTH_P`, 91
- `NV_Ith_P`, 92
- `NV_Ith_S`, 90
- `NV_LENGTH_S`, 89
- `NV_LOCLENGTH_P`, 91
- `NV_OWN_DATA_P`, 91
- `NV_OWN_DATA_S`, 89
- `NVECTOR` module, 85
- `nvector_parallel.h`, 18
- `nvector_serial.h`, 18
- optional input
 - band linear solver, 35–36
 - dense linear solver, 35–36
 - iterative linear solver, 36–38
 - rootfinding, 39
 - solver, 30–35
- optional output
 - band linear solver, 48–49
 - band-block-diagonal preconditioner, 67–68
 - banded preconditioner, 62–63
 - dense linear solver, 48–49
 - diagonal linear solver, 49–51
 - interpolated solution, 39
 - iterative linear solver, 51–54
 - solver, 40–47
- output mode, 8, 28
- portability, 18
 - Fortran, 70
- `PREC_BOTH`, 26, 27, 37
- `PREC_LEFT`, 26, 27, 37
- `PREC_NONE`, 26, 27, 37
- `PREC_RIGHT`, 26, 27, 37
- preconditioning
 - advice on, 9, 15
 - band-block diagonal, 63
 - banded, 61
 - setup and solve phases, 15
 - user-supplied, 36–37, 59, 60
- `pretype`, 26, 27, 37
- `PVODE`, 1
- `RCONST`, 18
- `realtype`, 18
- reinitialization, 54
- right-hand side function, 54
- Rootfinding, 10, 20, 27, 79
- `ROUT`, 78, 80
- `SMALL_REAL`, 18
- SPBCG generic linear solver
 - description of, 108
 - functions, 108
- SPGMR generic linear solver
 - description of, 107
 - functions, 107
 - support functions, 107
- SPTFQMR generic linear solver
 - description of, 108
 - functions, 108
- `STAB_LIM`, 79
- Stability limit detection, 9
- step size bounds, 33
- `STOP_TIME`, 79
- `sundials_nvector.h`, 18
- `sundials_types.h`, 18
- TFQMR method, 27, 38, 108
- tolerances, 6, 23, 24, 55
- `UNIT_ROUNDOFF`, 18
- User main program
 - CVBANDPRE usage, 61
 - CVBBDPRE usage, 65
 - FCVBBD usage, 83
 - FCVBP usage, 81
 - FCVODE usage, 71
 - IVP solution, 19
- `user_data`, 31, 55–58, 64, 65
- `VODE`, 1
- `VODPK`, 1
- weighted root-mean-square norm, 6

