

User Documentation for KINSOL v2.7.0

Aaron M. Collier, Alan C. Hindmarsh, Radu Serban, and Carol S. Woodward
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory

April 11, 2012



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	v
List of Figures	vii
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
3 Code Organization	11
3.1 SUNDIALS organization	11
3.2 KINSOL organization	11
4 Using KINSOL for C Applications	15
4.1 Access to library and header files	15
4.2 Data types	16
4.3 Header files	16
4.4 A skeleton of the user's main program	17
4.5 User-callable functions	19
4.5.1 KINSOL initialization and deallocation functions	19
4.5.2 Linear solver specification functions	20
4.5.3 KINSOL solver function	22
4.5.4 Optional input functions	24
4.5.4.1 Main solver optional input functions	24
4.5.4.2 Direct linear solvers optional input functions	33
4.5.4.3 Iterative linear solvers optional input functions	34
4.5.5 Optional output functions	35
4.5.5.1 Main solver optional output functions	36
4.5.5.2 Direct linear solvers optional output functions	38
4.5.5.3 Iterative linear solvers optional output functions	39
4.6 User-supplied functions	42
4.6.1 Problem-defining function	42
4.6.2 Error message handler function	42
4.6.3 Informational message handler function	43
4.6.4 Jacobian information (direct method with dense Jacobian)	43
4.6.5 Jacobian information (direct method with banded Jacobian)	44
4.6.6 Jacobian information (matrix-vector product)	45
4.6.7 Preconditioning (linear system solution)	46
4.6.8 Preconditioning (Jacobian data)	47
4.7 A parallel band-block-diagonal preconditioner module	47

5	FKINSOL, an Interface Module for FORTRAN Applications	53
5.1	FKINSOL routines	53
5.2	Important note on portability	54
5.3	Usage of the FKINSOL interface module	54
5.4	FKINSOL optional input and output	59
5.5	Usage of the FKINBBD interface to KINBBDPRE	60
6	Description of the NVECTOR module	65
6.1	The NVECTOR_SERIAL implementation	69
6.2	The NVECTOR_PARALLEL implementation	71
6.3	NVECTOR functions used by KINSOL	73
7	Providing Alternate Linear Solver Modules	75
7.1	Initialization function	75
7.2	Setup function	76
7.3	Solve function	76
7.4	Memory deallocation function	76
8	Generic Linear Solvers in SUNDIALS	79
8.1	The DLS modules: DENSE and BAND	79
8.1.1	Type DlsMat	80
8.1.2	Accessor macros for the DLS modules	81
8.1.3	Functions in the DENSE module	83
8.1.4	Functions in the BAND module	85
8.2	The SPILS modules: SPGMR, SPBCG, and SPTFQMR	87
8.2.1	The SPGMR module	87
8.2.2	The SPBCG module	88
8.2.3	The SPTFQMR module	88
A	KINSOL Installation Procedure	89
A.1	Autotools-based installation	90
A.1.1	Configuration options	91
A.1.2	Configuration examples	94
A.2	CMake-based installation	94
A.2.1	Configuring, building, and installing on Unix-like systems	95
A.2.2	Configuring, building, and installing on Windows	96
A.2.3	Configuration options	96
A.3	Manually building SUNDIALS	99
A.4	Installed libraries and exported header files	100
B	KINSOL Constants	103
B.1	KINSOL input constants	103
B.2	KINSOL output constants	103
	Bibliography	107
	Index	109

List of Tables

4.1	Optional inputs for KINSOL, KINDENSE, and KINSPILS	25
4.2	Optional outputs from KINSOL, KINDLS, and KINSPILS	36
5.1	Keys for setting FKINSOL optional inputs	60
5.2	Description of the FKINSOL optional output arrays IOUT and ROUT	61
6.1	Description of the NVECTOR operations	67
6.2	List of vector functions usage by KINSOL code modules	74
A.1	SUNDIALS libraries and header files	102

List of Figures

3.1	Organization of the SUNDIALS suite	12
3.2	Overall structure diagram of the KINSOL package	13
8.1	Diagram of the storage for a banded matrix of type <code>DlsMat</code>	82

Chapter 1

Introduction

KINSOL is part of a software family called SUNDIALS: SUite of Nonlinear and Differential/ALgebraic equation Solvers. This suite consists of CVODE, KINSOL, and IDA, and variants of these. KINSOL is a general-purpose nonlinear system solver based on Newton-Krylov solver technology.

1.1 Historical Background

The first nonlinear solver packages based on Newton-Krylov methods were written in FORTRAN. In particular, the NKSOL package, written at LLNL, was the first Newton-Krylov solver package written for solution of systems arising in the solution of partial differential equations [2]. This FORTRAN code made use of Newton's method to solve the discrete nonlinear systems and applied a preconditioned Krylov linear solver for solution of the Jacobian system at each nonlinear iteration. The key to the Newton-Krylov method was that the matrix-vector multiplies required by the Krylov method could effectively be approximated by a finite difference of the nonlinear system-defining function, avoiding a requirement for the formation of the actual Jacobian matrix. Significantly less memory was required for the solver as a result.

In the late 1990's, there was a push at LLNL to rewrite the nonlinear solver in C and port it to distributed memory parallel machines. Both Newton and Krylov methods are easily implemented in parallel, and this effort gave rise to the KINSOL package. KINSOL is similar to NKSOL in functionality, except that it provides for more options in the choice of linear system methods and tolerances, and has a more modular design to provide flexibility for future enhancements.

At present, KINSOL contains three Krylov methods: the GMRES (Generalized Minimal RESidual) [9], Bi-CGStab (Bi-Conjugate Gradient Stabilized) [10], and TFQMR (Transpose-Free Quasi-Minimal Residual) [7] linear iterative methods. As Krylov methods, these require almost no matrix storage 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 nonlinear algebraic systems, the Krylov methods are preferable over direct linear solver methods, and are often the only feasible choice. Among the three Krylov methods in KINSOL, 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.

For the sake of completeness in functionality, direct linear systems are now included in KINSOL. These include methods for both dense and banded linear systems, with Jacobians that are either user-supplied or generated internally by difference quotients.

In the process of translating NKSOL into C, the overall KINSOL organization has been changed considerably. One key feature of the KINSOL organization is that a separate module devoted to vector operations has been created. This module facilitated extension to multiprocessor environments with minimal impact on the rest of the solver. The new vector module design is shared across the SUNDIALS suite. This NVECTOR module is written in terms of abstract vector operations with the actual routines

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 KINSOL. 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 KINSOL 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 FORTRAN.

1.2 Changes from previous versions

Changes in v2.7.0

One significant design change was made with this release: The problem size and its relatives, bandwidth parameters, related internal indices, pivot arrays, and the optional output `lsflag` have all been changed from type `int` to type `long int`, except for the problem size and bandwidths in user calls to routines specifying BLAS/LAPACK routines for the dense/band linear solvers. The function `NewIntArray` is replaced by a pair `NewIntArray/NewLintArray`, for `int` and `long int` arrays, respectively.

A large number of errors have been fixed. Three major logic bugs were fixed – involving updating the solution vector, updating the linesearch parameter, and a missing error return. Three minor errors were fixed – involving setting `etachoice` in the Matlab/KINSOL interface, a missing error case in `KINPrintInfo`, and avoiding an exponential overflow in the evaluation of `omega`. In each linear solver interface function, the linear solver memory is freed on an error return, and the `**Free` function now includes a line setting to NULL the main memory pointer to the linear solver memory. In the installation files, we modified the treatment of the macro `SUNDIALS_USE_GENERIC_MATH`, so that the parameter `GENERIC_MATH_LIB` is either defined (with no value) or not defined.

Changes in v2.6.0

This release introduces a new linear solver module, based on Blas and Lapack for both dense and banded matrices.

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

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

KINSPBCG, KINSPTFQMR, KINDENSE, and KINBAND modules have been added to interface with the Scaled Preconditioned Bi-CGStab (SPBCG), Scaled Preconditioned Transpose-Free Quasi-Minimal Residual (SPTFQMR), DENSE, and BAND linear solver modules, respectively. (For details see Chapter 4.) Corresponding additions were made to the FORTRAN interface module FKINSOL. 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.

Regarding the FORTRAN interface module FKINSOL, optional inputs are now set using FKINSETIIN (integer inputs), FKINSETRIN (real inputs), and FKINSETVIN (vector inputs). Optional outputs are still obtained from the IOUT and ROUT arrays which are owned by the user and passed as arguments to FKINMALLOC.

The KINDENSE and KINBAND linear solver modules include support for nonlinear residual monitoring which can be used to control Jacobian updating.

To reduce the possibility of conflicts, the names of all header files have been changed by adding unique prefixes (`kinsol_` 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. Additionally, to resolve potential variable scope issues, all SUNDIALS solvers release user data right after its use. The build system 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, KINSOL now provides a set of routines (with prefix `KINSet`) to change the default values for various quantities controlling the solver and a set of extraction routines (with prefix `KINGet`) 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 Chapter 4.

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

Installation of KINSOL (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 examples. 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 KINSOL. The most casual user, with a small nonlinear system, can get by with reading all of Chapter 2, then Chapter 4 through §4.5.3 only, and looking at examples in [3]. In a different direction, a more expert user with a nonlinear system may want to

(a) use a package preconditioner (§4.7), (b) supply his/her own Jacobian or preconditioner routines (§4.6), (c) supply a new NVECTOR module (Chapter 6), or even (d) supply a different linear solver module (§3.2 and Chapter 7).

The structure of this document is as follows:

- In Chapter 2, we provide short descriptions of the numerical methods implemented by KINSOL for the solution of nonlinear systems.
- The following chapter describes the structure of the SUNDIALS suite of solvers (§3.1) and the software organization of the KINSOL solver (§3.2).
- Chapter 4 is the main usage document for KINSOL for C applications. It includes a complete description of the user interface for the solution of nonlinear algebraic systems.
- In Chapter 5, we describe FKINSOL, an interface module for the use of KINSOL 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 KINSOL, within the structure of SUNDIALS (Appendix A), as well as a list of all the constants used for input to and output from KINSOL functions (Appendix B).

Finally, the reader should be aware of the following notational conventions in this user guide: program listings and identifiers (such as `KINInit`) within textual explanations appear in typewriter type style; fields in C structures (such as *content*) appear in italics; and packages or modules 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 KINSOL code and user guide by Allan G. Taylor.

Chapter 2

Mathematical Considerations

KINSOL solves nonlinear algebraic systems in real N -space,

$$F(u) = 0, \quad F : \mathbf{R}^N \rightarrow \mathbf{R}^N, \quad (2.1)$$

given an initial guess u_0 .

Basic Newton iteration

Depending on the linear solver used, KINSOL can employ either an Inexact Newton method [1, 2, 4, 5, 8], or a Modified Newton method. At the highest level, KINSOL implements the following iteration scheme:

1. Set $u_0 =$ an initial guess
2. For $n = 0, 1, 2, \dots$ until convergence do:
 - (a) Solve $J(u_n)\delta_n = -F(u_n)$
 - (b) Set $u_{n+1} = u_n + \lambda\delta_n$, $0 < \lambda \leq 1$
 - (c) Test for convergence

Here, u_n is the n th iterate to u , and $J(u) = F'(u)$ is the system Jacobian. At each stage in the iteration process, a scalar multiple of the step δ_n , is added to u_n to produce a new iterate, u_{n+1} . A test for convergence is made before the iteration continues.

Newton method variants

For solving the linear system given in step 2(a), KINSOL provides several choices, including the option of a 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. 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),
- 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.

When using one of the direct linear solvers, the linear system in 2(a) is solved exactly, thus resulting in a Modified Newton method (the Jacobian matrix is normally out of date; see below¹). Note that the direct linear solvers (dense and band) can only be used with serial vector representations.

On the other hand, when using any of the iterative linear solvers (GMRES, Bi-CGStab, or TFQMR), the linear system in 2(a) is solved only approximately, thus resulting in an Inexact Newton method. Here right preconditioning is available by way of the preconditioning setup and solve routines supplied by the user, in which case the iterative method is applied to the linear systems $(JP^{-1})(P\delta) = -F$, where P denotes the right preconditioning matrix.

Jacobian information update strategy

In general, unless specified otherwise by the user, KINSOL strives to update Jacobian information (the actual system Jacobian J in the case of direct linear solvers, or the preconditioner matrix P in the case of iterative linear solvers) as infrequently as possible to balance the high costs of matrix operations against other costs. Specifically, these updates occur when:

- the problem is initialized,
- $\|\lambda\delta_{n-1}\|_{D_u, \infty} > 1.5$ (Inexact Newton only),
- `mbset` = 10 nonlinear iterations have passed since the last update,
- the linear solver failed recoverably with outdated Jacobian information,
- the global strategy failed with outdated Jacobian information, or
- $\|\lambda\delta_n\|_{D_u, \infty} < \text{STEPTOL}$ with outdated Jacobian information.

KINSOL allows, through optional solver inputs, changes to the above strategy. Indeed, the user can disable the initial Jacobian information evaluation or change the default value of `mbset`, the number of nonlinear iterations after which a Jacobian information update is enforced.

Scaling

To address the case of ill-conditioned nonlinear systems, KINSOL allows prescribing scaling factors both for the solution vector and for the residual vector. For scaling to be used, the user should supply values D_u , which are diagonal elements of the scaling matrix such that $D_u u_n$ has all components roughly the same magnitude when u_n is close to a solution, and D_F , which are diagonal scaling matrix elements such that $D_F F$ has all components roughly the same magnitude when u_n is not too close to a solution. In the text below, we use the following scaled norms:

$$\|z\|_{D_u} = \|D_u z\|_2, \quad \|z\|_{D_F} = \|D_F z\|_2, \quad \|z\|_{D_u, \infty} = \|D_u z\|_\infty, \quad \text{and} \quad \|z\|_{D_F, \infty} = \|D_F z\|_\infty \quad (2.2)$$

where $\|\cdot\|_\infty$ is the max norm. When scaling values are provided for the solution vector, these values are automatically incorporated into the calculation of the perturbations used for the default difference quotient approximations for Jacobian information; see (2.6) and (2.8) below.

Globalization strategy

Two methods of applying a computed step δ_n to the previously computed solution vector are implemented. The first and simplest is the standard Newton strategy which applies step 2(b) as above with λ always set to 1. The other method is a global strategy, which attempts to use the direction implied by δ_n in the most efficient way for furthering convergence of the nonlinear problem. This technique is implemented in the second strategy, called Linesearch. This option employs both the

¹KINSOL allows the user to enforce a Jacobian evaluation at each iteration thus allowing for an Exact Newton iteration.

α and β conditions of the Goldstein-Armijo linesearch given in [5] for step 2(b), where λ is chosen to guarantee a sufficient decrease in F relative to the step length as well as a minimum step length relative to the initial rate of decrease of F . One property of the algorithm is that the full Newton step tends to be taken close to the solution.

KINSOL implements a backtracking algorithm to first find the value λ such that $u_n + \lambda\delta_n$ satisfies the sufficient decrease condition (or α -condition)

$$F(u_n + \lambda\delta_n) \leq F(u_n) + \alpha \nabla F(u_n) \lambda \delta_n,$$

where $\alpha = 10^{-4}$. Although backtracking in itself guarantees that the step is not too small, KINSOL secondly relaxes λ to satisfy the so-called β -condition (equivalent to Wolfe's curvature condition):

$$F(u_n + \lambda\delta_n) \geq F(u_n) + \beta \nabla F(u_n) \lambda \delta_n,$$

where $\beta = 0.9$. During this second phase, λ is allowed to vary in the interval $[\lambda_{min}, \lambda_{max}]$ where

$$\lambda_{min} = \frac{\text{STEPTOL}}{\|\bar{\delta}_n\|_\infty}, \quad \bar{\delta}_n^j = \frac{\delta_n^j}{1/D_u^j + |u^j|},$$

and λ_{max} corresponds to the maximum feasible step size at the current iteration (typically $\lambda_{max} = \text{STEPMAX}/\|\delta_n\|_{D_u}$). In the above expressions, v^j denotes the j th component of a vector v .

For more details, the reader is referred to [5].

Nonlinear iteration stopping criteria

Stopping criteria for the Newton method are applied to both of the nonlinear residual and the step length. For the former, the Newton iteration must pass a stopping test

$$\|F(u_n)\|_{D_F, \infty} < \text{FTOL},$$

where FTOL is an input scalar tolerance with a default value of $U^{1/3}$. Here U is the machine unit roundoff. For the latter, the Newton method will terminate when the maximum scaled step is below a given tolerance

$$\|\lambda\delta_n\|_{D_u, \infty} < \text{STEPTOL},$$

where STEPTOL is an input scalar tolerance with a default value of $U^{2/3}$. Only the first condition (small residual) is considered a successful completion of KINSOL. The second condition (small step) may indicate that the iteration is stalled near a point for which the residual is still unacceptable.

Additional constraints

As a user option, KINSOL permits the application of inequality constraints, $u^i > 0$ and $u^i < 0$, as well as $u^i \geq 0$ and $u^i \leq 0$, where u^i is the i th component of u . Any such constraint, or no constraint, may be imposed on each component. KINSOL will reduce step lengths in order to ensure that no constraint is violated. Specifically, if a new Newton iterate will violate a constraint, the maximum step length along the Newton direction that will satisfy all constraints is found, and δ_n in Step 2(b) is scaled to take a step of that length.

Residual monitoring for Modified Newton method

When using a Modified Newton method (i.e. when a direct linear solver is used), in addition to the strategy described above for the update of the Jacobian matrix, KINSOL also provides an optional nonlinear residual monitoring scheme to control when the system Jacobian is updated. Specifically, a Jacobian update will also occur when `mbsetsub= 5` nonlinear iterations have passed since the last update and

$$\|F(u_n)\|_{D_F} > \omega \|F(u_m)\|_{D_F},$$

where u_n is the current iterate and u_m is the iterate at the last Jacobian update. The scalar ω is given by

$$\omega = \min \left(\omega_{min} e^{\max(0, \rho-1)}, \omega_{max} \right), \quad (2.3)$$

with ρ defined as

$$\rho = \frac{\|F(u_n)\|_{D_F}}{\text{FTOL}}, \quad (2.4)$$

where FTOL is the input scalar tolerance discussed before. Optionally, a constant value ω_{const} can be used for the parameter ω .

The constants controlling the nonlinear residual monitoring algorithm can be changed from their default values through optional inputs to KINSOL. These include the parameters ω_{min} and ω_{max} , the constant value ω_{const} , and the threshold `mbsetsub`.

Stopping criteria for iterative linear solvers

When using an Inexact Newton method (i.e. when an iterative linear solver is used), the convergence of the overall nonlinear solver is intimately coupled with the accuracy with which the linear solver in 2(a) above is solved. KINSOL provides three options for stopping criteria for the linear system solver, including the two algorithms of Eisenstat and Walker [6]. More precisely, the Krylov iteration must pass a stopping test

$$\|J\delta_n + F\|_{D_F} < (\eta_n + U)\|F\|_{D_F},$$

where η_n is one of:

Eisenstat and Walker Choice 1

$$\eta_n = \frac{|\|F(u_n)\|_{D_F} - \|F(u_{n-1}) + J(u_{n-1})\delta_n\|_{D_F}|}{\|F(u_{n-1})\|_{D_F}},$$

Eisenstat and Walker Choice 2

$$\eta_n = \gamma \left(\frac{\|F(u_n)\|_{D_F}}{\|F(u_{n-1})\|_{D_F}} \right)^\alpha,$$

where default values of γ and α are 0.9 and 2, respectively.

Constant η

$$\eta_n = \text{constant},$$

with 0.1 as the default.

The default strategy is "Eisenstat and Walker Choice 1". For both options 1 and 2, appropriate safeguards are incorporated to ensure that η does not decrease too quickly [6].

Difference quotient Jacobian approximations

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

$$J^{ij} = [F^i(u + \sigma_j e^j) - F^i(u)]/\sigma_j. \quad (2.5)$$

The increments σ_j are given by

$$\sigma_j = \sqrt{U} \max \{|u^j|, 1/D_u^j\}. \quad (2.6)$$

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, Jacobian information is needed only as matrix-vector products Jv . If a routine for Jv is not supplied, these products are approximated by directional difference quotients as

$$J(u)v \approx [F(u + \sigma v) - F(u)]/\sigma, \quad (2.7)$$

where u is the current approximation to a root of (2.1), and σ is a scalar. The choice of σ is taken from [2] and is given by

$$\sigma = \frac{\max\{|u^T v|, u_{typ}^T |v|\}}{\|v\|_2} \text{sign}(u^T v) \sqrt{U}, \quad (2.8)$$

where u_{typ} is a vector of typical values for the absolute values of the solution (and can be taken to be inverses of the scale factors given for u as described below). Convergence of the Newton method is maintained as long as the value of σ remains appropriately small, as shown in [1].

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 KINSOL organization


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

The overall organization of the KINSOL package is shown in Figure 3.2. The central solver module, implemented in the files `kinsol.h`, `kinsol_impl.h` and `kinsol.c`, deals with the solution of a nonlinear algebraic system using either an Inexact Newton method or a line search method for the global strategy. Although this module contains logic for the Newton iteration, it has no knowledge of the method 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.


At present, the package includes the following seven KINSOL 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:

- KINDENSE: LU factorization and backsolving with dense matrices (using either an internal implementation or Blas/Lapack);
- KINBAND: 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 KINSOL package. Modules specific to KINSOL are distinguished by rounded boxes, while generic solver and auxiliary modules are in rectangular boxes. Grayed boxes refer to the encompassing SUNDIALS structure. Note that the direct linear solvers using Lapack implementations are not explicitly represented.

- KINSPGMR: scaled preconditioned GMRES method;
- KINSPBCG: scaled preconditioned Bi-CGStab method;
- KINSPTFQMR: scaled preconditioned TFQMR method.

The set of linear solver modules distributed with KINSOL is intended to be expanded in the future as new algorithms are developed.

In the case of the direct methods KINDENSE and KINBAND 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 methods KINSPGMR, KINSPBCG and KINSPTFQMR, 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.

Each KINSOL 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, as required to achieve convergence. The call list within the central KINSOL module to each of the 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 the modules interfacing to Lapack linear solvers, each of the modules KINDENSE, KINBAND, KINSPGMR, KINSPBCG, and KINSPTFQMR is a set of interface routines built on top of a generic solver module, name DENSE, BAND, SPGMR, SPBCG, and SPTFQMR, respectively. The interface deals with the use of these methods in the KINSOL 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 KINSOL package elsewhere.

KINSOL also provides a preconditioner module called KINBBDPRE for use with any of the Krylov iterative linear solvers. It works in conjunction with NVECTOR_PARALLEL and generates a preconditioner that is a block-diagonal matrix with each block being a band matrix, as further described in §4.7.

All state information used by KINSOL 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 KINSOL 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 KINSOL memory structure. The reentrancy of KINSOL was motivated by the anticipated multi-computer extension.

Chapter 4

Using KINSOL for C Applications

This chapter is concerned with the use of KINSOL for the solution of nonlinear systems. The following subsections treat the header files, the layout of the user's main program, description of the KINSOL user-callable routines, and user-supplied functions. The sample programs described in the companion document [3] may also be helpful. Those codes may be used as templates (with the removal of some lines involved in testing), and are included in the KINSOL 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 KINSOL modules can only be used with NVECTOR_SERIAL: KINDENSE and KINBAND. The preconditioner module KINBDPRE can only be used with NVECTOR_PARALLEL.

KINSOL 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 KINSOL, 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 KINSOL. The relevant library files are

- *libdir/libsundials_kinsol.lib*,
- *libdir/libsundials_nvec*.lib* (one or two files),

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

- *incdir/include*
- *incdir/include/kinsol*
- *incdir/include/sundials*

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

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

`kinsol.h` also includes `sundials_types.h`, which defines the types `realtype` and `booleantype` and 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 KINSOL 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 of these files include in turn the header file `sundials_nvector.h`, which defines the abstract `N_Vector` type.

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

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

- `kinsol_lapack.h`, which is used with Lapack implementations of dense or band direct linear solvers;
- `kinsol_spgmr.h`, which is used with the Krylov solver SPGMR;
- `kinsol_spgcgs.h`, which is used with the Krylov solver SPBCG;
- `kinsol_sptfqmr.h`, which is used with the Krylov solver SPTFQMR;

The header files for the dense and banded linear solvers (both internal and Lapack) include the file `kinsol_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 `kinsol_spils.h` which defined common functions and which in turn includes a header file (`sundials_iterative.h`) which enumerates the kind of preconditioning and for the choices for the Gram-Schmidt process for SPGMR.

Other headers may be needed, according to the choice of preconditioner, etc. For example, in the `kinFoodWeb_kry_p` example (see [3]), preconditioning is done with a block-diagonal matrix. For this, even though the KINSPGMR 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 solution of a nonlinear problem. 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 KINSOL: 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, aside from the internal use in NVECTOR_PARALLEL. Here `argc` and `argv` are the command line argument counter and array received by `main`.

2. Set problem dimensions

[S] Set `N`, the problem size N (of type `int`).

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

3. Set vector with initial guess

To set the vector `u` of initial values, use functions defined by a particular NVECTOR implementation. If a `realtype` array `udata` already exists, containing the initial guess of u_0 , make the call:

```
[S] u = N_VMake_Serial(N, udata);
```

```
[P] u = N_VMake_Parallel(comm, Nlocal, N, udata);
```

Otherwise, make the call:

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

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

and load initial values into the structure defined by:

```
[S] NV_DATA_S(u)
```

```
[P] NV_DATA_P(u)
```

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 KINSOL object

Call `kin_mem = KINCreate()` to create the KINSOL memory block. `KINCreate` returns a pointer to the KINSOL memory structure. See §4.5.1 for details.

5. Set optional inputs

Call `KINSet*` routines to change from their default values any optional inputs that control the behavior of KINSOL. See §4.5.4 for details.

6. Allocate internal memory

Call `KINInit(...)` to specify the problem defining function F , allocate internal memory for KINSOL, and initialize KINSOL. `KINInit` returns a flag to indicate success or an illegal argument value. See §4.5.1 for details.

7. Attach linear solver module

Initialize the linear solver module with one of the following calls (for details see §4.5.2).

```
[S] ier = KINDense(...);
[S] ier = KINBand(...);
[S] flag = KINLapackDense(...);
[S] flag = KINLapackBand(...);
ier = KINSpqmr(...);
ier = KINSpbcg(...);
ier = KINSptfqmr(...);
```

8. Set linear solver optional inputs

Call `KIN*Set*` functions from the selected linear solver module to change optional inputs specific to that linear solver. See §4.5.4 for details.

9. Solve problem

Call `ier = KINSol(...)` to solve the nonlinear problem for a given initial guess. See §4.5.3 for details.

10. Get optional outputs

Call `KINGet*` and `KIN*Get*` functions to obtain optional output. See §4.5.5 for details.

11. Deallocate memory for solution vector

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

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

12. Free solver memory

Call `KINFree(&kin_mem)` to free the memory allocated for KINSOL.

13. [P] Finalize MPI

Call `MPI_Finalize()` to terminate MPI.

4.5 User-callable functions

This section describes the KINSOL functions that are called by the user to set up and solve a nonlinear problem. Some of these are required. However, starting with §4.5.4, the functions listed involve optional inputs/outputs or restarting, and those paragraphs can be skipped for a casual use of KINSOL. In any case, refer to §4.4 for the correct order of these calls.

The return flag (when present) for each of these routines is a negative integer if an error occurred, and non-negative otherwise.

4.5.1 KINSOL 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 problem solution is complete, as it frees the KINSOL memory block created and allocated by the first two calls.

KINCreate

Call `kin_mem = KINCreate();`

Description The function `KINCreate` instantiates a KINSOL solver object.

Arguments This function has no arguments.

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

KINInit

Call `flag = KINInit(kin_mem, func, tmpl);`

Description The function `KINInit` specifies the problem-defining function, allocates internal memory, and initializes KINSOL.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block returned by `KINCreate`.

`func` (`KINSysFn`) is the C function which computes the system function F in the nonlinear problem. This function has the form `func(u, fval, user_data)`. (For full details see §4.6.1.)

`tmpl` (`N_Vector`) is any `N_Vector` (e.g. the initial guess vector `u`) which is used as a template to create (by cloning) necessary vectors in `kin_mem`.

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

`KIN_SUCCESS` The call to `KINInit` was successful.

`KIN_MEM_NULL` The KINSOL memory block was not initialized through a previous call to `KINCreate`.

`KIN_MEM_FAIL` A memory allocation request has failed.

`KIN_ILL_INPUT` An input argument to `KINInit` has an illegal value.

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

KINFree

Call `KINFree(&kin_mem);`

Description The function `KINFree` frees the memory allocated by a previous call to `KINCreate`.

Arguments The argument is the address of the pointer to the KINSOL memory block returned by `KINCreate` (of type `void *`).

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

4.5.2 Linear solver specification functions

As previously explained, Newton iteration requires the solution of linear systems of the form $J\delta = -F$. There are five KINSOL linear solvers currently available for this task: KINDENSE, KINBAND, KINSPGMR, KINSPBCG, and KINSPTFQMR.

The first two linear solvers are direct and derive their names from the type of approximation used for the Jacobian $J = \partial F/\partial u$; KINDENSE and KINBAND 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 KINDLS (from Direct Linear Solvers).

The remaining three KINSOL linear solvers — KINSPGMR, KINSPBCG, and KINSPTFQMR — are Krylov iterative solvers, which use scaled preconditioned GMRES, scaled preconditioned Bi-CGStab, and scaled preconditioned TFQMR, respectively. Together, they are referred to as KINSPILS (from Scaled Preconditioned Iterative Linear Solvers).

With any of the Krylov solvers, only right preconditioning is available. For specification of the preconditioner, see the Krylov solver sections within §4.5.4 and §4.6. If preconditioning is done, user-supplied functions define the right preconditioner matrix P , which should approximate the system Jacobian matrix J .

To specify a KINSOL linear solver, after the call to `KINcreate` but before any calls to `KINsol`, the user's program must call one of the functions `KINDense/KINLapackDense`, `KINBand/KINLapackBand`, `KINSpGmr`, `KINSpbcg`, or `KINSpTfQmr`, as documented below. The first argument passed to these functions is the KINSOL memory pointer returned by `KINcreate`. A call to one of these functions links the main KINSOL nonlinear solver 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 KINBAND 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, with the exception of the Lapack-based direct solvers, the linear solver module used by KINSOL 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.

`KINDense`

Call `flag = KINDense(kin_mem, N);`

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

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

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.

`N` (`long int`) problem dimension.

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

`KINDLS_SUCCESS` The KINDENSE initialization was successful.

`KINDLS_MEM_NULL` The `kin_mem` pointer is NULL.

`KINDLS_ILL_INPUT` The KINDENSE solver is not compatible with the current NVECTOR module.

`KINDLS_MEM_FAIL` A memory allocation request failed.

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

KINLapackDense

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

KINBand

- Call** `flag = KINBand(kin_mem, N, mupper, mlower);`
- Description** The function `KINBand` selects the `KINBAND` linear solver and indicates the use of the internal direct band linear algebra functions.
- The user's main program must include the `kinsol_band.h` header file.
- Arguments** `kin_mem` (`void *`) pointer to the `KINSOL` memory block.
- `N` (`long int`) problem dimension.
- `mupper` (`long int`) upper half-bandwidth of the problem Jacobian (or of the approximation of it).
- `mlower` (`long int`) lower half-bandwidth of the problem Jacobian (or of the approximation of it).
- Return value** The return value `flag` (of type `int`) is one of
- `KINDLS_SUCCESS` The `KINBAND` initialization was successful.
- `KINDLS_MEM_NULL` The `kin_mem` pointer is `NULL`.
- `KINDLS_ILL_INPUT` The `KINBAND` solver is not compatible with the current `NVECTOR` module, or one of the Jacobian half-bandwidths is outside its valid range ($0 \dots N-1$).
- `KINDLS_MEM_FAIL` A memory allocation request failed.
- Notes** The `KINBAND` linear solver may not be compatible with a particular implementation of the `NVECTOR` module. Of the two `NVECTOR` modules provided by `SUNDIALS`, only `NVECTOR_SERIAL` is compatible, while `NVECTOR_PARALLEL` is not. The half-bandwidths are to be set so that the nonzero locations (i, j) in the banded (approximate) Jacobian satisfy $-mlower \leq j - i \leq mupper$.

KINLapackBand

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

KINSpgmr

Call `flag = KINSpgmr(kin_mem, maxl);`

Description The function `KINSpgmr` selects the `KINSPGMR` linear solver.

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

Arguments `kin_mem` (`void *`) pointer to the `KINSOL` memory block.

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

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

`KINSPILS_SUCCESS` The `KINSPGMR` initialization was successful.

`KINSPILS_MEM_NULL` The `kin_mem` pointer is `NULL`.

`KINSPILS_ILL_INPUT` The `NVECTOR` module used does not implement a required operation.

`KINSPILS_MEM_FAIL` A memory allocation request failed.

KINSpbcg

Call `flag = KINSpbcg(kin_mem, maxl);`

Description The function `KINSpbcg` selects the `KINSPBCG` linear solver.

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

Arguments `kin_mem` (`void *`) pointer to the `KINSOL` memory block.

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

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

`KINSPILS_SUCCESS` The `KINSPBCG` initialization was successful.

`KINSPILS_MEM_NULL` The `kin_mem` pointer is `NULL`.

`KINSPILS_ILL_INPUT` The `NVECTOR` module used does not implement a required operation.

`KINSPILS_MEM_FAIL` A memory allocation request failed.

KINSptfqmr

Call `flag = KINSptfqmr(kin_mem, maxl);`

Description The function `KINSptfqmr` selects the `KINSPTFQMR` linear solver.

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

Arguments `kin_mem` (`void *`) pointer to the `KINSOL` memory block.

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

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

`KINSPILS_SUCCESS` The `KINSPTFQMR` initialization was successful.

`KINSPILS_MEM_NULL` The `kin_mem` pointer is `NULL`.

`KINSPILS_ILL_INPUT` The `NVECTOR` module used does not implement a required operation.

`KINSPILS_MEM_FAIL` A memory allocation request failed.

4.5.3 KINSOL solver function

This is the central step in the solution process, the call to solve the nonlinear algebraic system.

KINSol

Call `flag = KINSol(kin_mem, u, strategy, u_scale, f_scale);`

Description The function `KINSol` computes an approximate solution to the nonlinear system.

Arguments

- `kin_mem` (void *) pointer to the KINSOL memory block.
- `u` (N_Vector) vector set to initial guess by user before calling `KINSol`, but which upon return contains an approximate solution of the nonlinear system $F(u) = 0$.
- `strategy` (int) globalization strategy applied to the Newton method. It must be one of `KIN_NONE` or `KIN_LINESEARCH`.
- `u_scale` (N_Vector) vector containing diagonal elements of scaling matrix D_u for vector `u` chosen so that the components of $D_u \cdot u$ (as a matrix multiplication) all have roughly the same magnitude when `u` is close to a root of $F(u)$.
- `f_scale` (N_Vector) vector containing diagonal elements of scaling matrix D_F for $F(u)$ chosen so that the components of $D_F \cdot F(u)$ (as a matrix multiplication) all have roughly the same magnitude when `u` is not too near a root of $F(u)$.

Return value On return, `KINSol` returns the approximate solution in the vector `u` if successful. The return value `flag` (of type `int`) will be one of the following:

`KIN_SUCCESS`

`KINSol` succeeded; the scaled norm of $F(u)$ is less than `fnormtol`.

`KIN_INITIAL_GUESS_OK`

The guess $u = u_0$ satisfied the system $F(u) = 0$ within the tolerances specified.

`KIN_STEP_LT_STPTOL`

`KINSOL` stopped based on scaled step length. This means that the current iterate may be an approximate solution of the given nonlinear system, but it is also quite possible that the algorithm is “stalled” (making insufficient progress) near an invalid solution, or that the scalar `scsteptol` is too large (see `KINSetScaledStepTo1` in §4.5.4 to change `scsteptol` from its default value).

`KIN_MEM_NULL`

The `KINSOL` memory block pointer was `NULL`.

`KIN_ILL_INPUT`

An input parameter was invalid.

`KIN_NO_MALLOC`

The `KINSOL` memory was not allocated by a call to `KINCreate`.

`KIN_LINESEARCH_NONCONV`

The line search algorithm was unable to find an iterate sufficiently distinct from the current iterate, or could not find an iterate satisfying the sufficient decrease condition.

Failure to satisfy the sufficient decrease condition could mean the current iterate is “close” to an approximate solution of the given nonlinear system, the difference approximation of the matrix-vector product $J(u)v$ is inaccurate, or the real scalar `scsteptol` is too large.

`KIN_MAXITER_REACHED`

The maximum number of nonlinear iterations has been reached.

`KIN_MXNEWT_5X_EXCEEDED`

Five consecutive steps have been taken that satisfy the inequality $\|D_u p\|_{L_2} > 0.99 \text{ mxnewtstep}$, where p denotes the current step and `mxnewtstep` is a scalar upper bound on the scaled step length. Such a failure may mean that $\|D_F F(u)\|_{L_2}$ asymptotes from above to a positive value, or the real scalar `mxnewtstep` is too small.

KIN_LINESEARCH_BCFAIL

The line search algorithm was unable to satisfy the “beta-condition” for `MXNBCF + 1` nonlinear iterations (not necessarily consecutive), which may indicate the algorithm is making poor progress.

KIN_LINSOLV_NO_RECOVERY

The user-supplied routine `psolve` encountered a recoverable error, but the preconditioner is already current.

KIN_LINIT_FAIL

The linear solver initialization routine (`linit`) encountered an error.

KIN_LSETUP_FAIL

The user-supplied routine `pset` (used to set up the preconditioner data) encountered an unrecoverable error.

KIN_LSOLVE_FAIL

Either the user-supplied routine `psolve` (used to solve the preconditioned linear system) encountered an unrecoverable error, or the linear solver routine (`lsolve`) encountered an error condition.

KIN_SYSFUNC_FAIL

The system function failed in an unrecoverable manner.

KIN_FIRST_SYSFUNC_ERR

The system function failed recoverably at the first call.

KIN_REPTD_SYSFUNC_ERR

The system function had repeated recoverable errors. No recovery is possible.

Notes

The components of vectors `u_scale` and `f_scale` should be strictly positive.

`KIN_SUCCESS = 0`, `KIN_INITIAL_GUESS_OK = 1`, and `KIN_STEP_LT_STPTOL = 2`. All remaining return values are negative and therefore a test `flag < 0` will trap all KINSOL failures.

4.5.4 Optional input functions

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

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

4.5.4.1 Main solver optional input functions

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

KINSetErrFile

Call `flag = KINSetErrFile(kin_mem, errfp);`

Description The function `KINSetErrFile` specifies the pointer to the file where all KINSOL messages should be directed when the default KINSOL error handler function is used.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`errfp` (`FILE *`) pointer to output file.

Table 4.1: Optional inputs for KINSOL, KINDENSE, and KINSPILS

Optional input	Function name	Default
KINSOL main solver		
Error handler function	KINSetErrorHandlerFn	internal fn.
Pointer to an error file	KINSetErrFile	stderr
Info handler function	KINSetInfoHandlerFn	internal fn.
Pointer to an info file	KINSetInfoFile	stdout
Data for problem-defining function	KINSetUserData	NULL
Verbosity level of output	KINSetPrintLevel	0
Max. number of nonlinear iterations	KINSetNumMaxIters	200
No initial preconditioner setup	KINSetNoInitSetup	FALSE
No residual monitoring*	KINSetNoResMon	FALSE
Max. iterations without prec. setup	KINSetMaxSetupCalls	10
Max. iterations without residual check*	KINSetMaxSubSetupCalls	5
Form of η coefficient	KINSetEtaForm	KIN_ETACHOICE1
Constant value of η	KINSetEtaConstValue	0.1
Values of γ and α	KINSetEtaParams	0.9 and 2.0
Values of ω_{min} and ω_{max} *	KINSetResMonParams	0.00001 and 0.9
Constant value of ω^*	KINSetResMonConstValue	0.9
Lower bound on ϵ	KINSetNoMinEps	FALSE
Max. scaled length of Newton step	KINSetMaxNewtonStep	$1000\ D_u u_0\ _2$
Max. number of β -condition failures	KINSetMaxBetaFails	10
Rel. error for D.Q. Jv	KINSetRelErrFunc	$\sqrt{\text{uround}}$
Function-norm stopping tolerance	KINSetFuncNormTol	$\text{uround}^{1/3}$
Scaled-step stopping tolerance	KINSetScaledSteptol	$\text{uround}^{2/3}$
Inequality constraints on solution	KINSetConstraints	NULL
Nonlinear system function	KINSetSysFunc	none
KINDLS linear solvers		
Dense Jacobian function	KINDlsSetDenseJacFn	DQ
Band Jacobian function	KINDlsSetBandJacFn	DQ
KINSPILS linear solvers		
Max. number of restarts**	KINSpilsSetMaxRestarts	0
Preconditioner functions and data	KINSpilsSetPreconditioner	NULL, NULL, NULL
Jacobian-times-vector function and data	KINSpilsSetJacTimesVecFn	internal DQ, NULL

* Only for the KINDLS linear solvers

** Only for KINSPGMR

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

`KIN_SUCCESS` The optional value has been successfully set.

`KIN_MEM_NULL` The `kin_mem` pointer is `NULL`.

Notes The default value for `errfp` is `stderr`.

Passing a value of `NULL` disables all future error message output (except for the case in which the KINSOL memory pointer is `NULL`). This use of `KINSetErrFile` is strongly discouraged.

If `KINSetErrFile` is to be called, it should be called before any other optional input functions, in order to take effect for any later error message.



KINSetErrorHandlerFn

Call `flag = KINSetErrorHandlerFn(kin_mem, ehfun, eh_data);`

Description The function `KINSetErrorHandlerFn` specifies the optional user-defined function to be used in handling error messages.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.

`ehfun` (`KINErrorHandlerFn`) is the user's C error handler function (see §4.6.2).

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

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

`KIN_SUCCESS` The function `ehfun` and data pointer `eh_data` have been successfully set.

`KIN_MEM_NULL` The `kin_mem` pointer is `NULL`.

Notes The default internal error handler function directs error messages to the file specified by the file pointer `errfp` (see `KINSetErrFile` above).

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

KINSetInfoFile

Call `flag = KINSetInfoFile(kin_mem, infofp);`

Description The function `KINSetInfoFile` specifies the pointer to the file where all informative (non-error) messages should be directed.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.

`infofp` (`FILE *`) pointer to output file.

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

`KIN_SUCCESS` The optional value has been successfully set.

`KIN_MEM_NULL` The `kin_mem` pointer is `NULL`.

Notes The default value for `infofp` is `stdout`.

KINSetInfoHandlerFn

Call `flag = KINSetInfoHandlerFn(kin_mem, ihfun, ih_data);`

Description The function `KINSetInfoHandlerFn` specifies the optional user-defined function to be used in handling informative (non-error) messages.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.

`ihfun` (`KINInfoHandlerFn`) is the user's C information handler function (see §4.6.3).

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

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

KIN_SUCCESS The function `ihfun` and data pointer `ih_data` have been successfully set.
KIN_MEM_NULL The `kin_mem` pointer is NULL.
Notes The default internal information handler function directs informative (non-error) messages to the file specified by the file pointer `infofp` (see `KINSetInfoFile` above).

KINSetPrintLevel

Call `flag = KINSetPrintLevel(kin_mem, printf1);`
Description The function `KINSetPrintLevel` specifies the level of verbosity of the output.
Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`printf1` (`int`) flag indicating the level of verbosity. Must be one of:

- 0 no information displayed.
- 1 for each nonlinear iteration display the following information: the scaled Euclidean ℓ_2 norm of the system function evaluated at the current iterate, the scaled norm of the Newton step (only if using `KIN_NONE`), and the number of function evaluations performed so far.
- 2 display level 1 output and the following values for each iteration:
 - $\|F(u)\|_{D_F}$ (only for `KIN_NONE`).
 - $\|F(u)\|_{D_{F,\infty}}$ (for `KIN_NONE` and `KIN_LINESEARCH`).
- 3 display level 2 output plus additional values used by the global strategy (only if using `KIN_LINESEARCH`), and statistical information for the linear solver.

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

- KIN_SUCCESS** The optional value has been successfully set.
- KIN_MEM_NULL** The `kin_mem` pointer is NULL.
- KIN_ILL_INPUT** The argument `printf1` had an illegal value.

Notes The default value for `printf1` is 0.

KINSetUserData

Call `flag = KINSetUserData(kin_mem, user_data);`
Description The function `KINSetUserData` specifies the pointer to user-defined memory that is to be passed to all user-supplied functions.
Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`user_data` (`void *`) pointer to the user-defined memory.
Return value The return value `flag` (of type `int`) is one of:

- KIN_SUCCESS** The optional value has been successfully set.
- KIN_MEM_NULL** The `kin_mem` pointer is NULL.

Notes If specified, the pointer to `user_data` is passed to all user-supplied functions that have it as an argument. Otherwise, a NULL pointer is passed.
 If `user_data` is needed in user preconditioner functions, the call to `KINSetUserData` must be made *before* the call to specify the linear solver.

KINSetNumMaxIters

Call `flag = KINSetNumMaxIters(kin_mem, mxiter);`
Description The function `KINSetNumMaxIters` specifies the maximum number of nonlinear iterations allowed.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`mxiter` (`long int`) maximum number of nonlinear iterations.

Return value The return value `flag` (of type `int`) is one of:
`KIN_SUCCESS` The optional value has been successfully set.
`KIN_MEM_NULL` The `kin_mem` pointer is `NULL`.
`KIN_ILL_INPUT` The maximum number of iterations was non-positive.

Notes The default value for `mxiter` is `MXITER_DEFAULT = 200`.

KINSetNoInitSetup

Call `flag = KINSetNoInitSetup(kin_mem, noInitSetup);`

Description The function `KINSetNoInitSetup` specifies whether an initial call to the preconditioner setup function should be made or not.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`noInitSetup` (`booleantype`) flag controlling whether an initial call to the preconditioner setup function is made (pass `FALSE`) or not made (pass `TRUE`).

Return value The return value `flag` (of type `int`) is one of:
`KIN_SUCCESS` The optional value has been successfully set.
`KIN_MEM_NULL` The `kin_mem` pointer is `NULL`.

Notes The default value for `noInitSetup` is `FALSE`, meaning that an initial call to the preconditioner setup function will be made.

A call to this function is useful when solving a sequence of problems, in which the final preconditioner value from problem is to be used initially for the next problem.

KINSetNoResMon

Call `flag = KINSetNoResMon(kin_mem, noNNIResMon);`

Description The function `KINSetNoResMon` specifies whether or not the nonlinear residual monitoring scheme is used to control Jacobian updating

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`noNNIResMon` (`booleantype`) flag controlling whether residual monitoring is used (pass `FALSE`) or not used (pass `TRUE`).

Return value The return value `flag` (of type `int`) is one of:
`KIN_SUCCESS` The optional value has been successfully set.
`KIN_MEM_NULL` The `kin_mem` pointer is `NULL`.

Notes When using a direct solver, the default value for `noNNIResMon` is `FALSE`, meaning that the nonlinear residual will be monitored.

Residual monitoring is only available for use with the direct linear solver modules (meaning `KINDENSE` and `KINBAND`).



KINSetMaxSetupCalls

Call `flag = KINSetMaxSetupCalls(kin_mem, msbset);`

Description The function `KINSetMaxSetupCalls` specifies the maximum number of nonlinear iterations that can be performed between calls to the preconditioner setup function.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`msbset` (`long int`) maximum number of nonlinear iterations without a call to the preconditioner setup function. Pass 0 to indicate the default.

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

`KIN_SUCCESS` The optional value has been successfully set.
`KIN_MEM_NULL` The `kin_mem` pointer is `NULL`.
`KIN_ILL_INPUT` The argument `msbset` was negative.

Notes The default value for `msbset` is `MSBSET_DEFAULT = 10`.

KINSetMaxSubSetupCalls

Call `flag = KINSetMaxSubSetupCalls(kin_mem, msbsetsub);`

Description The function `KINSetMaxSubSetupCalls` specifies the maximum number of nonlinear iterations between checks by the residual monitoring algorithm.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`msbsetsub` (`long int`) maximum number of nonlinear iterations without checking the nonlinear residual. Pass 0 to indicate the default.

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

`KIN_SUCCESS` The optional value has been successfully set.
`KIN_MEM_NULL` The `kin_mem` pointer is `NULL`.
`KIN_ILL_INPUT` The argument `msbsetsub` was negative.

Notes The default value for `msbsetsub` is `MSBSET_SUB_DEFAULT = 5`.

Residual monitoring is only available for use with the direct linear solver modules (meaning `KINDENSE` and `KINBAND`).



KINSetEtaForm

Call `flag = KINSetEtaForm(kin_mem, etachoice);`

Description The function `KINSetEtaForm` specifies the method for computing the value of the η coefficient used in the calculation of the linear solver convergence tolerance.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`etachoice` (`int`) flag indicating the method for computing η . The value must be one of `KIN_ETACHOICE1`, `KIN_ETACHOICE2`, or `KIN_ETACONSTANT` (see Chapter 2 for details).

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

`KIN_SUCCESS` The optional value has been successfully set.
`KIN_MEM_NULL` The `kin_mem` pointer is `NULL`.
`KIN_ILL_INPUT` The argument `etachoice` had an illegal value.

Notes The default value for `etachoice` is `KIN_ETACHOICE1`.

KINSetEtaConstValue

Call `flag = KINSetEtaConstValue(kin_mem, eta);`

Description The function `KINSetEtaConstValue` specifies the constant value for η in the case `etachoice = KIN_ETACONSTANT`.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`eta` (`realtype`) constant value for η . Pass 0.0 to indicate the default.

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

`KIN_SUCCESS` The optional value has been successfully set.
`KIN_MEM_NULL` The `kin_mem` pointer is `NULL`.
`KIN_ILL_INPUT` The argument `eta` had an illegal value

Notes The default value for `eta` is 0.1. The legal values are $0.0 < \eta \leq 1.0$.

KINSetEtaParams

Call `flag = KINSetEtaParams(kin_mem, egamma, ealpha);`

Description The function `KINSetEtaParams` specifies the parameters γ and α in the formula for η , in the case `etachoice = KIN_ETACHOICE2`.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`egamma` (`realtype`) value of the γ parameter. Pass 0.0 to indicate the default.
`ealpha` (`realtype`) value of the α parameter. Pass 0.0 to indicate the default.

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

- `KIN_SUCCESS` The optional values have been successfully set.
- `KIN_MEM_NULL` The `kin_mem` pointer is `NULL`.
- `KIN_ILL_INPUT` One of the arguments `egamma` or `ealpha` had an illegal value.

Notes The default values for `egamma` and `ealpha` are 0.9 and 2.0, respectively.
The legal values are $0.0 < \text{egamma} \leq 1.0$ and $1.0 < \text{ealpha} \leq 2.0$.

KINSetResMonConstValue

Call `flag = KINSetResMonConstValue(kin_mem, omegaconst);`

Description The function `KINSetResMonConstValue` specifies the constant value for ω when using residual monitoring.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`omegaconst` (`realtype`) constant value for ω . Passing 0.0 results in using Eqn. (2.3).

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

- `KIN_SUCCESS` The optional value has been successfully set.
- `KIN_MEM_NULL` The `kin_mem` pointer is `NULL`.
- `KIN_ILL_INPUT` The argument `omegaconst` had an illegal value

Notes The default value for `omegaconst` is 0.9. The legal values are $0.0 < \text{omegaconst} < 1.0$.

KINSetResMonParams

Call `flag = KINSetResMonParams(kin_mem, omegamin, omegamax);`

Description The function `KINSetResMonParams` specifies the parameters ω_{min} and ω_{max} in the formula (2.3) for ω .

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`omegamin` (`realtype`) value of the ω_{min} parameter. Pass 0.0 to indicate the default.
`omegamax` (`realtype`) value of the ω_{max} parameter. Pass 0.0 to indicate the default.

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

- `KIN_SUCCESS` The optional values have been successfully set.
- `KIN_MEM_NULL` The `kin_mem` pointer is `NULL`.
- `KIN_ILL_INPUT` One of the arguments `omegamin` or `omegamax` had an illegal value.

Notes The default values for `omegamin` and `omegamax` are 0.00001 and 0.9, respectively.
The legal values are $0.0 < \text{omegamin} < \text{omegamax} < 1.0$.

KINSetNoMinEps

- Call `flag = KINSetNoMinEps(kin_mem, noMinEps);`
- Description The function `KINSetNoMinEps` specifies a flag that controls whether or not the value of ϵ , the scaled linear residual tolerance, is bounded from below.
- Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`noMinEps` (`booleantype`) flag controlling the bound on ϵ .
- Return value The return value `flag` (of type `int`) is one of:
`KIN_SUCCESS` The optional value has been successfully set.
`KIN_MEM_NULL` The `kin_mem` pointer is `NULL`.
- Notes The default value for `noMinEps` is `FALSE`, meaning that a positive minimum value, equal to $0.01 * \text{fnormtol}$, is applied to ϵ . (See `KINSetFuncNormTol` below.)

KINSetMaxNewtonStep

- Call `flag = KINSetMaxNewtonStep(kin_mem, mxnewtstep);`
- Description The function `KINSetMaxNewtonStep` specifies the maximum allowable scaled length of the Newton step.
- Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`mxnewtstep` (`realttype`) maximum scaled step length (≥ 0.0). Pass 0.0 to indicate the default.
- Return value The return value `flag` (of type `int`) is one of:
`KIN_SUCCESS` The optional value has been successfully set.
`KIN_MEM_NULL` The `kin_mem` pointer is `NULL`.
`KIN_ILL_INPUT` The input value was negative.
- Notes The default value of `mxnewtstep` is $1000 \|u_0\|_{D_u}$, where u_0 is the initial guess.

KINSetMaxBetaFails

- Call `flag = KINSetMaxBetaFails(kin_mem, mxnbcf);`
- Description The function `KINSetMaxBetaFails` specifies the maximum number of β -condition failures in the linesearch algorithm.
- Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`mxnbcf` (`realttype`) maximum number of β -condition failures. Pass 0.0 to indicate the default.
- Return value The return value `flag` (of type `int`) is one of:
`KIN_SUCCESS` The optional value has been successfully set.
`KIN_MEM_NULL` The `kin_mem` pointer is `NULL`.
`KIN_ILL_INPUT` `mxnbcf` was negative.
- Notes The default value of `mxnbcf` is `MXNBCF_DEFAULT = 10`.

KINSetRelErrFunc

- Call `flag = KINSetRelErrFunc(kin_mem, relfunc);`
- Description The function `KINSetRelErrFunc` specifies the relative error in computing $F(u)$, which is used in the difference quotient approximation of the Jacobian-vector product.
- Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.

relfunc (**realtype**) relative error in $F(u)$ (**relfunc** ≥ 0.0). Pass 0.0 to indicate the default.

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

KIN_SUCCESS The optional value has been successfully set.

KIN_MEM_NULL The **kin_mem** pointer is **NULL**.

KIN_ILL_INPUT The relative error was negative.

Notes The default value for **relfunc** is $\sqrt{\text{unit roundoff}}$.

KINSetFuncNormTol

Call **flag** = **KINSetFuncNormTol**(**kin_mem**, **fnormtol**);

Description The function **KINSetFuncNormTol** specifies the scalar used as a stopping tolerance on the scaled maximum norm of the system function $F(u)$.

Arguments **kin_mem** (**void ***) pointer to the KINSOL memory block.

fnormtol (**realtype**) tolerance for stopping based on scaled function norm (≥ 0.0). Pass 0.0 to indicate the default.

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

KIN_SUCCESS The optional value has been successfully set.

KIN_MEM_NULL The **kin_mem** pointer is **NULL**.

KIN_ILL_INPUT The tolerance was negative.

Notes The default value for **fnormtol** is $(\text{unit roundoff})^{1/3}$.

KINSetScaledStepTol

Call **flag** = **KINSetScaledStepTol**(**kin_mem**, **scsteptol**);

Description The function **KINSetScaledStepTol** specifies the scalar used as a stopping tolerance on the minimum scaled step length.

Arguments **kin_mem** (**void ***) pointer to the KINSOL memory block.

scsteptol (**realtype**) tolerance for stopping based on scaled step length (≥ 0.0). Pass 0.0 to indicate the default.

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

KIN_SUCCESS The optional value has been successfully set.

KIN_MEM_NULL The **kin_mem** pointer is **NULL**.

KIN_ILL_INPUT The tolerance was non-positive.

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

KINSetConstraints

Call **flag** = **KINSetConstraints**(**kin_mem**, **constraints**);

Description The function **KINSetConstraints** specifies a vector that defines inequality constraints for each component of the solution vector u .

Arguments **kin_mem** (**void ***) pointer to the KINSOL memory block.

constraints (**N_Vector**) vector of constraint flags. If **constraints[i]** is

0.0 then no constraint is imposed on u_i .

1.0 then u_i will be constrained to be $u_i \geq 0.0$.

-1.0 then u_i will be constrained to be $u_i \leq 0.0$.

2.0 then u_i will be constrained to be $u_i > 0.0$.

−2.0 then u_i will be constrained to be $u_i < 0.0$.

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

`KIN_SUCCESS` The optional value has been successfully set.
`KIN_MEM_NULL` The `kin_mem` pointer is `NULL`.
`KIN_ILL_INPUT` The constraint vector contains illegal values.

Notes The presence of a non-`NULL` constraints vector that is not 0.0 in all components will cause constraint checking to be performed.

The function creates a private copy of the constraints vector. Consequently, the user-supplied vector can be freed after the function call, and the constraints can only be changed by calling this function.

KINSetSysFunc

Call `flag = KINSetSysFunc(kin_mem, func);`

Description The function `KINSetSysFunc` specifies the user-provided function that evaluates the nonlinear system function $F(u)$.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`func` (`KINSysFn`) user-supplied function that evaluates $F(u)$.

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

`KIN_SUCCESS` The optional value has been successfully set.
`KIN_MEM_NULL` The `kin_mem` pointer is `NULL`.
`KIN_ILL_INPUT` The argument `func` was `NULL`.

Notes The nonlinear system function is initially specified through `KINInit`. The option of changing the system function is provided for a user who wishes to solve several problems of the same size but with different functions.

4.5.4.2 Direct linear solvers optional input functions

The `KINDENSE` solver needs a function to compute a dense approximation to the Jacobian matrix $J(u)$. This function must be of type `KINDlsDenseJacFn`. The user can supply his/her own dense Jacobian function, or use the default internal difference quotient approximation that comes with the `KINDENSE` solver. To specify a user-supplied Jacobian function `djac`, `KINDENSE` provides the function `KINDlsSetDenseJacFn`. The `KINDENSE` 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 `use_data` may be specified through `KINSetUserData`.

KINDlsSetDenseJacFn

Call `flag = KINDlsSetDenseJacFn(kin_mem, djac);`

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

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`djac` (`KINDlsDenseJacFn`) user-defined dense Jacobian approximation function.

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

`KINDLS_SUCCESS` The optional value has been successfully set.
`KINDLS_MEM_NULL` The `kin_mem` pointer is `NULL`.
`KINDLS_LMEM_NULL` The `KINDENSE` linear solver has not been initialized.

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

The function type `KINDlsDenseJacFn` is described in §4.6.4.

The KINDENSE solver needs a function to compute a banded approximation to the Jacobian matrix $J(u)$. This function must be of type `KINDlsBandJacFn`. The user can supply his/her own banded Jacobian approximation function, or use the default internal difference quotient approximation that comes with the KINBAND solver. To specify a user-supplied Jacobian function `bjac` KINBAND provides the function `KINDlsSetBandJacFn`. The KINBAND 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 `KINSetUserData`.

`KINDlsSetBandJacFn`

Call `flag = KINDlsSetBandJacFn(kin_mem, bjac);`

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

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`bjac` (`KINDlsBandJacFn`) user-defined banded Jacobian approximation function.

Return value The return value `flag` (of type `int`) is one of
`KINDLS_SUCCESS` The optional value has been successfully set.
`KINDLS_MEM_NULL` The `kin_mem` pointer is NULL.
`KINDLS_LMEM_NULL` The KINBAND linear solver has not been initialized.

Notes By default, KINBAND uses an internal difference quotient approximation. If NULL is passed to `bjac`, this default function is used.

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

4.5.4.3 Iterative linear solvers optional input functions

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

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 `KINSpilsSetPreconditioner`. A KINSPILS solver passes the pointer `user_data` received through `KINSetUserData` to the preconditioner `psetup` and `psolve` functions. This allows the user to create an arbitrary structure with relevant problem data and access it during the execution of the user-supplied preconditioner functions without using global data in the program.

The KINSPILS solvers require a function to compute an approximation to the product between the Jacobian matrix $J(u)$ and a vector v . The user can supply his/her own Jacobian-times-vector approximation function, or use the internal difference quotient approximation that comes with the KINSPILS solvers. A user-defined Jacobian-vector function must be of type `KINSpilsJacTimesVecFn` and can be specified through a call to `KINSpilsSetJacTimesVecFn` (see §4.6.6 for specification details). A KINSPILS solver passes the pointer `user_data` received through `KINSetUserData` to the Jacobian-times-vector function `jt看imes` each time it is called.

`KINSpilsSetPreconditioner`

Call `flag = KINSpilsSetPreconditioner(kin_mem, psetup, psolve);`

Description The function `KINSpilsSetPreconditioner` specifies the preconditioner setup and solve functions.

Arguments `kin_mem` (void *) pointer to the KINSOL memory block.
`psetup` (KINSpilsPrecSetupFn) user-defined preconditioner setup function. Pass NULL if no setup operation is to be done.
`psolve` (KINSpilsPrecSolveFn) user-defined preconditioner solve function.

Return value The return value `flag` (of type `int`) is one of
 KINSPILS_SUCCESS The optional values have been successfully set.
 KINSPILS_MEM_NULL The `kin_mem` pointer is NULL.
 KINSPILS_LMEM_NULL The KINSPILS linear solver has not been initialized.

Notes The function type KINSpilsPrecSolveFn is described in §4.6.7. The function type KINSpilsPrecSetupFn is described in §4.6.8.

KINSpilsSetJacTimesVecFn

Call `flag = KINSpilsSetJacTimesVecFn(kin_mem, jtimes);`

Description The function KINSpilsSetJacTimesFn specifies the Jacobian-vector function to be used.

Arguments `kin_mem` (void *) pointer to the KINSOL memory block.
`jtimes` (KINSpilsJacTimesVecFn) user-defined Jacobian-vector product function.

Return value The return value `flag` (of type `int`) is one of
 KINSPILS_SUCCESS The optional value has been successfully set.
 KINSPILS_MEM_NULL The `kin_mem` pointer is NULL.
 KINSPILS_LMEM_NULL The KINSPILS linear solver has not been initialized.

Notes By default, the KINSPILS linear solvers use an internal difference quotient function KINSpilsDQJtimes. If NULL is passed as `jtimes`, this default function is used.
 The function type KINSpilsJacTimesVecFn is described in §4.6.6.

KINSpilsSetMaxRestarts

Call `flag = KINSpilsSetMaxRestarts(kin_mem, maxrs);`

Description The function KINSpilsSetMaxRestarts specifies the maximum number of times the iterative linear solver can be restarted.

Arguments `kin_mem` (void *) pointer to the KINSOL memory block.
`maxrs` (int) maximum number of restarts (≥ 0).

Return value The return value `flag` (of type `int`) is one of:
 KINSPILS_SUCCESS The optional value has been successfully set.
 KINSPILS_ILL_INPUT The maximum number of restarts specified is negative.
 KINSPILS_MEM_NULL The `kin_mem` pointer is NULL.
 KINSPILS_LMEM_NULL The linear solver has not been initialized.

Notes The default value is 0 (meaning no restarts).
 This option is available only for the KINSPGMR linear solver.



4.5.5 Optional output functions

KINSOL provides an extensive list of functions that can be used to obtain solver performance information. Table 4.2 lists all optional output functions in KINSOL, which are then described in detail in the remainder of this section, beginning with those for the main KINSOL solver and continuing with those for the linear solver modules. Where the name of an output from a linear solver module would otherwise conflict with the name of an optional output from the main solver, a suffix LS (for Linear Solver) has been added here (e.g., `lenrwLS`).

Table 4.2: Optional outputs from KINSOL, KINDLS, and KINSPILS

Optional output	Function name
KINSOL main solver	
Size of KINSOL real and integer workspaces	KINGetWorkSpace
Number of function evaluations	KINGetNumFuncEvals
Number of nonlinear iterations	KINGetNumNolinSolvIters
Number of β -condition failures	KINGetNumBetaCondFails
Number of backtrack operations	KINGetNumBacktrackOps
Scaled norm of F	KINGetFuncNorm
Scaled norm of the step	KINGetStepLength
KINDLS linear solvers	
Size of real and integer workspaces	KINDlsGetWorkSpace
No. of Jacobian evaluations	KINDlsGetNumJacEvals
No. of F calls for D.Q. Jacobian evals.	KINDlsGetNumFuncEvals
Last return from a KINDLS function	KINDlsGetLastFlag
KINSPILS linear solvers	
Size of real and integer workspaces	KINSpilsGetWorkSpace
No. of linear iterations	KINSpilsGetNumLinIters
No. of linear convergence failures	KINSpilsGetNumConvFails
No. of preconditioner evaluations	KINSpilsGetNumPrecEvals
No. of preconditioner solves	KINSpilsGetNumPrecSolves
No. of Jacobian-vector product evaluations	KINSpilsGetNumJtimesEvals
No. of F calls for D.Q. Jacobian-vector evals.	KINSpilsGetNumFuncEvals
Last return from a linear solver function	KINSpilsGetLastFlag

4.5.5.1 Main solver optional output functions

KINSOL 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 and solver performance statistics. These optional output functions are described next.

KINGetWorkSpace

Call `flag = KINGetWorkSpace(kin_mem, &lenrw, &leniw);`

Description The function `KINGetWorkSpace` returns the KINSOL integer and real workspace sizes.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`lenrw` (`long int`) the number of `realtype` values in the KINSOL workspace.
`leniw` (`long int`) the number of integer values in the KINSOL workspace.

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

`KIN_SUCCESS` The optional output values have been successfully set.
`KIN_MEM_NULL` The `kin_mem` pointer is `NULL`.

Notes In terms of the problem size N , the actual size of the real workspace is $17+5N$ `realtype` words. The real workspace is increased by an additional N words if constraint checking is enabled (see `KINSetConstraints`).

The actual size of the integer workspace (without distinction between `int` and `long int`) is $22 + 5N$ (increased by N if constraint checking is enabled).

KINGetNumFuncEvals

Call `flag = KINGetNumFuncEvals(kin_mem, &nfevals);`

Description The function `KINGGetNumFuncEvals` returns the number of evaluations of the system function.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`nfevals` (`long int`) number of calls to the user-supplied function that evaluates $F(u)$.

Return value The return value `flag` (of type `int`) is one of:
`KIN_SUCCESS` The optional output value has been successfully set.
`KIN_MEM_NULL` The `kin_mem` pointer is NULL.

KINGGetNumNonlinSolvIters

Call `flag = KINGGetNumNonlinSolvIters(kin_mem, &nniters);`

Description The function `KINGGetNumNonlinSolvIters` returns the number of nonlinear iterations.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`nniters` (`long int`) number of nonlinear iterations.

Return value The return value `flag` (of type `int`) is one of:
`KIN_SUCCESS` The optional output value has been successfully set.
`KIN_MEM_NULL` The `kin_mem` pointer is NULL.

KINGGetNumBetaCondFails

Call `flag = KINGGetNumBetaCondFails(kin_mem, &nbcfails);`

Description The function `KINGGetNumBetaCondFails` returns the number of β -condition failures.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`nbcfails` (`long int`) number of β -condition failures.

Return value The return value `flag` (of type `int`) is one of:
`KIN_SUCCESS` The optional output value has been successfully set.
`KIN_MEM_NULL` The `kin_mem` pointer is NULL.

KINGGetNumBacktrackOps

Call `flag = KINGGetNumBacktrackOps(kin_mem, &nbacktr);`

Description The function `KINGGetNumBacktrackOps` returns the number of backtrack operations (step length adjustments) performed by the line search algorithm.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`nbacktr` (`long int`) number of backtrack operations.

Return value The return value `flag` (of type `int`) is one of:
`KIN_SUCCESS` The optional output value has been successfully set.
`KIN_MEM_NULL` The `kin_mem` pointer is NULL.

KINGGetFuncNorm

Call `flag = KINGGetFuncNorm(kin_mem, &fnorm);`

Description The function `KINGGetFuncNorm` returns the scaled Euclidean ℓ_2 norm of the nonlinear system function $F(u)$ evaluated at the current iterate.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`fnorm` (`realtype`) current scaled norm of $F(u)$.

Return value The return value `flag` (of type `int`) is one of:
`KIN_SUCCESS` The optional output value has been successfully set.
`KIN_MEM_NULL` The `kin_mem` pointer is NULL.

KINGetStepLength

Call `flag = KINGetStepLength(kin_mem, &steplength);`

Description The function `KINGetStepLength` returns the scaled Euclidean ℓ_2 norm of the step used during the previous iteration.

Arguments `kin_mem` (void *) pointer to the KINSOL memory block.
`steplength` (realtype) scaled norm of the Newton step.

Return value The return value `flag` (of type `int`) is one of:
`KIN_SUCCESS` The optional output value has been successfully set.
`KIN_MEM_NULL` The `kin_mem` pointer is NULL.

4.5.5.2 Direct linear solvers optional output functions

The following optional outputs are available from the KINDLS module: workspace requirements, number of calls to the Jacobian routine, number of calls to the system function routine for difference quotient Jacobian approximation, and last return value from a KINDLS function.

KINDLSGetWorkSpace

Call `flag = KINDLSGetWorkSpace(kin_mem, &lenrwlS, &leniwls);`

Description The function `KINDLSGetWorkSpace` returns the KINDENSE real and integer workspace sizes.

Arguments `kin_mem` (void *) pointer to the KINSOL memory block.
`lenrwlS` (long int) the number of `realtype` values in the KINDLS workspace.
`leniwls` (long int) the number of integer values in the KINDLS workspace.

Return value The return value `flag` (of type `int`) is one of
`KINDLS_SUCCESS` The optional output value has been successfully set.
`KINDLS_MEM_NULL` The `kin_mem` pointer is NULL.
`KINDLS_LMEM_NULL` The KINDENSE linear solver has not been initialized.

Notes For the KINDENSE linear solver, in terms of the problem size N , the actual size of the real workspace is N^2 `realtype` words, and the actual size of the integer workspace is N integer words.
For the KINBAND linear solver, in terms of the problem size N and Jacobian half-bandwidths, the actual size of the real workspace, in `realtype` words, is approximately $(2 \text{ mupper} + 3 \text{ mlower} + 2) N$, and the actual size of the integer workspace is N integer words.

KINDLSGetNumJacEvals

Call `flag = KINDLSGetNumJacEvals(kin_mem, &njevals);`

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

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

Return value The return value `flag` (of type `int`) is one of
`KINDLS_SUCCESS` The optional output value has been successfully set.
`KINDLS_MEM_NULL` The `kin_mem` pointer is NULL.
`KINDLS_LMEM_NULL` The KINDENSE linear solver has not been initialized.

KINDlsGetNumFuncEvals

Call `flag = KINDlsGetNumFuncEvals(kin_mem, &nfevalsLS);`

Description The function `KINDlsGetNumFuncEvals` returns the number of calls to the user system function used to compute the difference quotient approximation to the dense or banded Jacobian.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`nfevalsLS` (`long int`) the number of calls to the user system function.

Return value The return value `flag` (of type `int`) is one of
`KINDLS_SUCCESS` The optional output value has been successfully set.
`KINDLS_MEM_NULL` The `kin_mem` pointer is `NULL`.
`KINDLS_LMEM_NULL` The `KINDENSE` or `KINBAND` linear solver has not been initialized.

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

KINDlsGetLastFlag

Call `flag = KINDlsGetLastFlag(kin_mem, &lsflag);`

Description The function `KINDlsGetLastFlag` returns the last return value from a `KINDENSE` routine.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`lsflag` (`long int`) the value of the last return flag from a `KINDENSE` function.

Return value The return value `flag` (of type `int`) is one of
`KINDLS_SUCCESS` The optional output value has been successfully set.
`KINDLS_MEM_NULL` The `kin_mem` pointer is `NULL`.
`KINDLS_LMEM_NULL` The `KINDENSE` linear solver has not been initialized.

Notes If the `KINDLS` setup function failed (`KINSol` returned `KIN_LSETUP_FAIL`), then `lsflag` is equal to the column index (numbered from one) at which a zero diagonal element was encountered during the LU factorization of the dense Jacobian matrix. For all other failures, `lsflag` is negative.

4.5.5.3 Iterative linear solvers optional output functions

The following optional outputs are available from the `KINSPILS` 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 system function routine for difference quotient Jacobian-vector product approximation, and last return value from a linear solver function.

KINSpilsGetWorkSpace

Call `flag = KINSpilsGetWorkSpace(kin_mem, &lenrwLS, &leniwLS);`

Description The function `KINSpilsGetWorkSpace` returns the global sizes of the linear solver real and integer workspaces.

Arguments `kin_mem` (`void *`) pointer to the KINSOL memory block.
`lenrwLS` (`long int`) the number of `realtype` values in the linear solver workspace.
`leniwLS` (`long int`) the number of integer values in the linear solver workspace.

Return value The return value `flag` (of type `int`) is one of:
`KINSPILS_SUCCESS` The optional output values have been successfully set.

KINSPILS_MEM_NULL The `kin_mem` pointer is NULL.
 KINSPILS_LMEM_NULL The KINSPGMR linear solver has not been initialized.

Notes In terms of the problem size N and maximum subspace size `maxl`, the actual size of the real workspace, in `realttype` words, is roughly:
 $(\text{maxl}+3) * N + \text{maxl} * (\text{maxl}+4) + 1$ for KINSPGMR,
 $7 * N$ for KINSPBCG, and
 $11 * N$ for KINSPTFQMR.
 In a parallel setting, this value is global, summed over all processes.

KINSpilsGetNumLinIters

Call `flag = KINSpilsGetNumLinIters(kin_mem, &nliters);`

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

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

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

KINSPILS_SUCCESS The optional output value has been successfully set.
 KINSPILS_MEM_NULL The `kin_mem` pointer is NULL.
 KINSPILS_LMEM_NULL The linear solver module has not been initialized.

KINSpilsGetNumConvFails

Call `flag = KINSpilsGetNumConvFails(kin_mem, &nlcfails);`

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

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

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

KINSPILS_SUCCESS The optional output value has been successfully set.
 KINSPILS_MEM_NULL The `kin_mem` pointer is NULL.
 KINSPILS_LMEM_NULL The linear solver module has not been initialized.

KINSpilsGetNumPrecEvals

Call `flag = KINSpilsGetNumPrecEvals(kin_mem, &npevals);`

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

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

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

KINSPILS_SUCCESS The optional output value has been successfully set.
 KINSPILS_MEM_NULL The `kin_mem` pointer is NULL.
 KINSPILS_LMEM_NULL The linear solver module has not been initialized.

KINSpilsGetNumPrecSolves

Call `flag = KINSpilsGetNumPrecSolves(kin_mem, &npsolves);`

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

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

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

`KINSPILS_SUCCESS` The optional output value has been successfully set.
`KINSPILS_MEM_NULL` The `kin_mem` pointer is `NULL`.
`KINSPILS_LMEM_NULL` The linear solver module has not been initialized.

KINSpilsGetNumJtimesEvals

Call `flag = KINSpilsGetNumJtimesEvals(kin_mem, &njvevals);`

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

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

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

`KINSPILS_SUCCESS` The optional output value has been successfully set.
`KINSPILS_MEM_NULL` The `kin_mem` pointer is `NULL`.
`KINSPILS_LMEM_NULL` The linear solver module has not been initialized.

KINSpilsGetNumFuncEvals

Call `flag = KINSpilsGetNumFuncEvals(kin_mem, &nfevalsLS);`

Description The function `KINSpilsGetNumFuncEvals` returns the number of calls to the user system function for difference quotient Jacobian-vector product approximations.

Arguments `kin_mem` (void *) pointer to the KINSOL memory block.
`nfevalsLS` (long int) the number of calls to the user system function.

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

`KINSPILS_SUCCESS` The optional output value has been successfully set.
`KINSPILS_MEM_NULL` The `kin_mem` pointer is `NULL`.
`KINSPILS_LMEM_NULL` The linear solver module has not been initialized.

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

KINSpilsGetLastFlag

Call `flag = KINSpilsGetLastFlag(kin_mem, &lsflag);`

Description The function `KINSpilsGetLastFlag` returns the last return value from a `KINSPILS` routine.

Arguments `kin_mem` (void *) pointer to the KINSOL memory block.
`lsflag` (long int) the value of the last return flag from a `KINSPILS` function.

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

`KINSPILS_SUCCESS` The optional output value has been successfully set.
`KINSPILS_MEM_NULL` The `kin_mem` pointer is `NULL`.

	KINSPILS_LMEM_NULL The linear solver module has not been initialized.
Notes	<p>If the KINSPILS setup function failed (KINSOL returned KIN_LSETUP_FAIL), <code>lsflag</code> will be SPGMR_PSET_FAIL_UNREC, SPBCG_PSET_FAIL_UNREC, or SPTFQMR_PSET_FAIL_UNREC.</p> <p>If the KINSPGMR solve function failed (KINSOL returned KIN_LSOLVE_FAIL), <code>lsflag</code> contains the error return flag from <code>SpgmrSolve</code> and will be one of: SPGMR_MEM_NULL, indicating that the SPGMR memory is NULL; SPGMR_ATIMES_FAIL_UNREC, indicating an unrecoverable failure in the Jacobian-times-vector function; SPGMR_PSOLVE_FAIL_UNREC, indicating that the preconditioner solve function <code>psolve</code> failed unrecoverably; SPGMR_GS_FAIL, indicating a failure in the Gram-Schmidt procedure; or SPGMR_QRSOL_FAIL, indicating that the matrix R was found to be singular during the QR solve phase.</p> <p>If the KINSPBCG solve function failed (KINSOL returned KIN_LSOLVE_FAIL), <code>lsflag</code> contains the error return flag from <code>SpbcgSolve</code> and will be one of: SPBCG_MEM_NULL, indicating that the SPBCG memory is NULL; SPBCG_ATIMES_FAIL_UNREC, indicating an unrecoverable failure in the Jacobian-times-vector function; or SPBCG_PSOLVE_FAIL_UNREC, indicating that the preconditioner solve function <code>psolve</code> failed unrecoverably.</p> <p>If the KINSPTFQMR solve function failed (KINSOL returned KIN_LSOLVE_FAIL), <code>lsflag</code> contains the error return flag from <code>SptfqmrSolve</code> and will be one of: SPTFQMR_MEM_NULL, indicating that the SPTFQMR memory is NULL; SPTFQMR_ATIMES_FAIL_UNREC, indicating an unrecoverable failure in the $J*v$ function; or SPTFQMR_PSOLVE_FAIL_UNREC, indicating that the preconditioner solve function <code>psolve</code> failed unrecoverably.</p>

4.6 User-supplied functions

The user-supplied functions consist of one function defining the nonlinear system, (optionally) a function that handles error and warning messages, (optionally) a function that handles informational messages, (optionally) a function that provides Jacobian-related information for the linear solver, and (optionally) one or two functions that define the preconditioner for use in any of the Krylov iterative algorithms.

4.6.1 Problem-defining function

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

<code>KINSysFn</code>

Definition	<code>typedef void (*KINSysFn)(N_Vector u, N_Vector fval, void *user_data);</code>
Purpose	This function computes $F(u)$ for a given value of the vector u .
Arguments	<p><code>u</code> is the current value of the variable vector, u.</p> <p><code>fval</code> is the output vector $F(u)$.</p> <p><code>user_data</code> is a pointer to user data, the pointer <code>user_data</code> passed to <code>KINSetUserData</code>.</p>
Return value	A <code>KINSysFn</code> function should return 0 if successful, a positive value if a recoverable error occurred (in which case KINSOL will attempt to correct), or a negative value if it failed unrecoverably (in which case the solution process is halted and <code>KIN_SYSFUNC_FAIL</code> is returned).
Notes	Allocation of memory for <code>fval</code> is handled within KINSOL.

4.6.2 Error message handler function

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

KINErrorHandlerFn

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

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

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

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

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

4.6.3 Informational message handler function

As an alternative to the default behavior of directing informational (meaning non-error) messages to the file pointed to by `infofp` (see `KINSetInfoFile`), the user may provide a function of type `KINInfoHandlerFn` to process any such messages. The function type `KINInfoHandlerFn` is defined as follows:

KINInfoHandlerFn

Definition `typedef void (*KINInfoHandlerFn)(const char *module, const char *function, char *msg, void *ih_data);`

Purpose This function processes informational messages from KINSOL and its sub-modules.

Arguments `module` is the name of the KINSOL module reporting the information.
`function` is the name of the function reporting the information.
`msg` is the message.
`ih_data` is a pointer to user data, the same as the `ih_data` parameter passed to `KINSetInfoHandlerFn`.

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

4.6.4 Jacobian information (direct method with dense Jacobian)

If the direct linear solver with dense treatment of the Jacobian is used (`KINDense` or `KINLapackDense` is called in Step 7 of §4.4), the user may provide a function of type `KINDlsDenseJacFn` defined by

KINDlsDenseJacFn

Definition `typedef int (*KINDlsDenseJacFn)(long int N, N_Vector u, N_Vector fu, DlsMat J, void *user_data, N_Vector tmp1, N_Vector tmp2);`

Purpose This function computes the dense Jacobian $J(u)$ or an approximation to it.

Arguments `N` is the problem size.
`u` is the current (unscaled) iterate.
`fu` is the current value of the vector $F(u)$.
`J` is the output approximate Jacobian matrix, $J = \partial F / \partial u$.
`user_data` is a pointer to user data, the same as the `user_data` parameter passed to `KINSetUserData`.

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

Return value A function of type `KINDlsDenseJacFn` should return 0 if successful or a non-zero value otherwise.

Notes A user-supplied dense Jacobian function must load the N by N dense matrix J with an approximation to the Jacobian matrix $J(u)$ at u . Only nonzero elements need to be loaded into J because J is set to the zero matrix before the call to the Jacobian function. The type of J is `DlsMat`.

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

The `DlsMat` type and the accessor macros `DENSE_ELEM` and `DENSE_COL` are documented in §8.1.3.

If the user's `KINDlsDenseJacFn` function uses difference quotient approximations, it may need to access quantities not in the call list. To obtain these, use the `KINGet*` functions described in §4.5.5.1. The unit roundoff can be accessed as `UNIT_ROUNDOFF` defined in `sundials_types.h`.

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

4.6.5 Jacobian information (direct method with banded Jacobian)

If the direct linear solver with banded treatment of the Jacobian is used (`KINBand` or `KINLapackBand`) is called in Step 7 of §4.4), the user may provide a function of type `KINDlsBandJacFn` defined by:

`KINDlsBandJacFn`

```

Definition  typedef int (*KINDlsBandJacFn)(long int N, long int mupper,
                                           long int mlower, N_Vector u, N_Vector fu,
                                           DlsMat J, void *user_data,
                                           N_Vector tmp1, N_Vector tmp2);
  
```

Purpose This function computes the banded Jacobian $J(u)$ or a banded approximation to it.

Arguments N is the problem size.

`mlower`

`mupper` are the lower and upper half-bandwidths of the Jacobian.

`u` is the current (unscaled) iterate.

`fu` is the current value of the vector $F(u)$.

`J` is the output approximate Jacobian matrix, $J = \partial F / \partial u$.

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

`tmp1`

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

Return value A function of type `KINDlsBandJacFn` should return 0 if successful or a non-zero value otherwise.

Notes A user-supplied band Jacobian function must load the band matrix `J` of type `DlsMat` with the elements of the Jacobian $J(u)$ at `u`. Only nonzero elements need to be loaded into `J` because `J` is preset to zero 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 `J`, counting from 0. This macro is for use in small problems in which efficiency of access is not a major concern. Thus, in terms of indices m and n running from 1 to N with (m, n) within the band defined by `mupper` and `mlower`, the Jacobian element $J_{m,n}$ can be loaded with the statement `BAND_ELEM(J, m-1, n-1) = J_{m,n}`. The elements within the band are those with $-\text{mupper} \leq m-n \leq \text{mlower}$. Alternatively, `BAND_COL(J, j)` returns a pointer to the diagonal element of the j th column of `J`, and if we assign this address to `realtype *col_j`, then the i th element of the j th column is given by `BAND_COL_ELEM(col_j, i, j)`, counting from 0. Thus for (m, n) within the band, $J_{m,n}$ can be loaded by setting `col_n = BAND_COL(J, n-1)`; `BAND_COL_ELEM(col_n, m-1, n-1) = J_{m,n}`. The elements of the j th column can also be accessed via ordinary array indexing, but this approach requires knowledge of the underlying storage for a band matrix of type `DlsMat`. The array `col_n` can be indexed from $-\text{mupper}$ to `mlower`. For large problems, it is more efficient to use the combination of `BAND_COL` and `BAND_COL_ELEM` than to use the `BAND_ELEM`. As in the dense case, these macros all number rows and columns starting from 0, not 1.

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 `KINDlsBandJacFn` function uses difference quotient approximations, it may need to access quantities not in the call list. To obtain these, use the `KINGet*` functions described in §4.5.5.1. The unit roundoff can be accessed as `UNIT_ROUNDOFF` defined in `sundials.types.h`.

For the sake of uniformity, the arguments `N`, `mlower`, and `mupper` are of type `long int`, even in the case that the Lapack band solver is to be used.

4.6.6 Jacobian information (matrix-vector product)

If one of the Krylov iterative linear solvers `SPGMR`, `SPBCG`, or `SPTFQMR` is selected (`KINSp*` is called in step 7 of §4.4), the user may provide a `jtimes` function of type `KINSpilsJacTimesVecFn` to compute products Jv . If such a function is not supplied, the default is a difference quotient approximation of these products.

KINSpilsJacTimesVecFn

Definition `typedef int (*KINSpilsJacTimesVecFn)(N_Vector v, N_Vector Jv, N_Vector u, booleantype new_u, void *user_data);`

Purpose This `jtimes` function computes the product Jv (or an approximation to it).

Arguments `v` is the vector by which the Jacobian must be multiplied to the right.
`Jv` is the computed output vector.
`u` is the current value of the dependent variable vector.

new_u is a flag, input from KINSOL and possibly reset by the user's `jtimes` function, indicating whether the iterate vector `u` has been updated since the last call to `jtimes`. This is useful if the `jtimes` function computes and saves Jacobian data that depends on `u` for use in computing $J(u)v$. The input value of `new_u` is `TRUE` following an update by KINSOL, and in that case any saved Jacobian data depending on `u` should be recomputed. The `jtimes` routine should then set `new_u` to `FALSE`, so that on subsequent calls to `jtimes` with the same `u`, the saved data can be reused.

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

Return value The value to be returned by the Jacobian-times-vector function should be 0 if successful. If a recoverable failure occurred, the return value should be positive. In this case, KINSOL will attempt to correct by calling the preconditioner setup function if the preconditioner information is current, KINSOL halts. If the Jacobian-times-vector function encounters an unrecoverable error, it should return a negative value, prompting KINSOL to halt.

Notes If a user-defined routine is not given, then an internal `KINSPGMR` function, using difference quotient approximations, is used.

If the user-provided `KINSpilsJacTimesVecFn` function needs the unit roundoff, this can be accessed as `UNIT_ROUNDOFF` defined in `sundials_types.h`.

4.6.7 Preconditioning (linear system solution)

If preconditioning is used, then the user must provide a C function to solve the linear system $Pz = r$ where P is the preconditioner matrix, approximating (at least crudely) the system Jacobian $J = \partial F / \partial u$. This function must be of type `KINSpilsPrecSolveFn`, defined as follows:

```
KINSpilsPrecSolveFn
```

```
Definition    typedef int (*KINSpilsPrecSolveFn)(N_Vector u, N_Vector uscale,
                                                N_Vector fval, N_Vector fscale,
                                                N_Vector v, void *user_data,
                                                N_Vector tmp);
```

Purpose This function solves the preconditioning system $Pz = r$.

Arguments

- u** is the current (unscaled) value of the iterate.
- uscale** is a vector containing diagonal elements of the scaling matrix for `u`.
- fval** is the vector $F(u)$ evaluated at `u`.
- fscale** is a vector containing diagonal elements of the scaling matrix for `fval`.
- v** on input, `v` is set to the right-hand side vector of the linear system, `r`. On output, `v` must contain the solution `z` of the linear system $Pz = r$.
- user_data** is a pointer to user data, the same as the `user_data` parameter passed to the function `KINSetUserData`.
- 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, and negative for an unrecoverable error.

Notes If the preconditioner solve function fails recoverably and if the preconditioner information (set by the preconditioner setup function) is out of date, KINSOL attempts to correct by calling the setup function. If the preconditioner data is current, KINSOL halts.

4.6.8 Preconditioning (Jacobian data)

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

```
KINSpilsPrecSetupFn
```

Definition `typedef int (*KINSpilsPrecSetupFn)(N_Vector u, N_Vector uscale,
 N_Vector fval, N_Vector fscale,
 void *user_data, N_Vector tmp1,
 N_Vector tmp2);`

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

Arguments The arguments of a `KINSpilsPrecSetupFn` are as follows:

`u` is the current (unscaled) value of the iterate.

`uscale` is a vector containing diagonal elements of the scaling matrix for `u`.

`fval` is the vector $F(u)$ evaluated at `u`.

`fscale` is a vector containing diagonal elements of the scaling matrix for `fval`.

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

`tmp1`
`tmp2` are pointers to memory allocated for variables of type `N_Vector` which can be used by `KINSpilsPrecSetupFn` 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, any other value resulting in halting the KINSOL solver.

Notes The user-supplied preconditioner setup subroutine should compute the right preconditioner matrix P (stored in the memory block referenced by the `user_data` pointer) used to form the scaled preconditioned linear system

$$(D_F J(u) P^{-1} D_u^{-1}) \cdot (D_u P x) = -D_F F(u),$$

where D_u and D_F denote the diagonal scaling matrices whose diagonal elements are stored in the vectors `uscale` and `fscale`, respectively.

The preconditioner setup routine will not be called prior to every call made to the preconditioner solve function, but will instead be called only as often as necessary to achieve convergence of the Newton iteration.

If the user's `KINSpilsPrecSetupFn` function uses difference quotient approximations, it may need to access quantities not in the call list. To obtain these, use the `KINGet*` functions described in §4.5.5.1. The unit roundoff can be accessed as `UNIT_ROUNDOFF` defined in `sundials_types.h`.

If the preconditioner solve routine requires no preparation, then a preconditioner setup function need not be given.

4.7 A parallel band-block-diagonal preconditioner module

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, KINSOL provides a band-block-diagonal preconditioner module `KINBBDPRE`, to be used with the parallel `N_Vector` module described in §6.2.

This module provides a preconditioner matrix for KINSOL that is block-diagonal with banded blocks. The blocking corresponds to the distribution of the dependent variable vector u amongst the processes. Each preconditioner block is generated from the Jacobian of the local part (associated with the current process) of a given function $G(u)$ approximating $F(u)$ ($G = F$ is allowed). The blocks are generated by each process via a difference quotient scheme, utilizing a specified band structure. This structure is given by upper and lower half-bandwidths, `mudq` and `mldq`, defined as the number of non-zero diagonals above and below the main diagonal, respectively. However, from the resulting approximate Jacobian blocks, 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. Such an efficiency gain may occur if the couplings in the system outside a certain bandwidth are considerably weaker than those within the band. Reducing `mukeep` and `mlkeep` while keeping `mudq` and `mldq` at their true values, discards the elements outside the narrower band. Reducing both pairs has the additional effect of lumping the outer Jacobian elements into the computed elements within the band, and requires more caution and experimentation to see whether the lower cost of narrower band matrices offsets the loss of accuracy in the blocks.

The KINBBDPRE module calls two user-provided functions to construct P : a required function `Gloc` (of type `KINLocalFn`) which approximates the nonlinear system function $G(u) \approx F(u)$ and which is computed locally, and an optional function `Gcomm` (of type `KINCommFn`) which performs all interprocess communication necessary to evaluate the approximate function G . These are in addition to the user-supplied nonlinear system function that evaluates $F(u)$. Both functions take as input the same pointer `user_data` as that passed by the user to `KINSetUserData` and passed to the user's function `func`, and neither function has a return value. The user is responsible for providing space (presumably within `user_data`) for components of u that are communicated by `Gcomm` from the other processes, and that are then used by `Gloc`, which should not do any communication.

KINLocalFn

Definition `typedef void (*KINLocalFn)(long int Nlocal, N_Vector u,
N_Vector gval, void *user_data);`

Purpose This `Gloc` function computes $G(u)$, and outputs the resulting vector as `gval`.

Arguments `Nlocal` is the local vector length.

`u` is the current value of the iterate.

`gval` is the output vector.

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

Return value A `KINLocalFn` function type does not have a return value.

Notes This function must assume that all interprocess communication of data needed to calculate `gval` has already been done, and this data is accessible within `user_data`.

Memory for `u` and `gval` is handled within the preconditioner module.

The case where G is mathematically identical to F is allowed.

KINCommFn

Definition `typedef void (*KINCommFn)(long int Nlocal, N_Vector u, void *user_data);`

Purpose This `Gcomm` function performs all interprocess communications necessary for the execution of the `Gloc` function above, using the input vector `u`.

Arguments `Nlocal` is the local vector length.

`u` is the current value of the iterate.

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

Return value A `KINCommFn` function type does not have a return value.

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

Each call to the `Gcomm` function is preceded by a call to the system function `func` with the same `u` argument. Thus `Gcomm` can omit any communications done by `func` if relevant to the evaluation of `Gloc`. If all necessary communication was done in `func`, then `Gcomm = NULL` can be passed in the call to `KINBBDPrecInit` (see below).

Besides the header files required for the solution of a nonlinear problem (see §4.3), to use the `KINBBDPRE` module, the main program must include the header file `kinbbdpre.h` which declares the needed function prototypes.

The following is a summary of the usage of this module and describes the sequence of calls in the user main program. Steps that are unchanged from the user main program presented in §4.4 are grayed out.

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

```
flag = KINSpgmr(kin_mem, maxl);
flag = KINSpbcg(kin_mem, maxl);
flag = KINSptfqmr(kin_mem, maxl);
```

8. Initialize the `KINBBDPRE` preconditioner module

Specify the upper and lower half-bandwidth pairs (`mudq`, `mldq`) and (`mukeep`, `mlkeep`), and call

```
flag = KINBBDPrecInit(kin_mem, Nlocal, mudq, mldq,
                    mukeep, mlkeep, dq_rel_u, Gloc, Gcomm);
```

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

9. Set linear solver optional inputs

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

10. Solve problem
11. Get optional output

Additional optional outputs associated with `KINBBDPRE` are available by way of two routines described below, `KINBBDPrecGetWorkSpace` and `KINBBDPrecGetNumGfnEvals`.

12. Deallocate memory for solution vector
13. Free solver memory
14. Finalize MPI

The user-callable function that initializes `KINBBDPRE` (step 8), is described in more detail below.

KINBBDPrecInit

Call `flag = KINBBDPrecInit(kin_mem, Nlocal, mudq, mldq, mukeep, mlkeep, dq_rel_u, Gloc, Gcomm);`

Description The function `KINBBDPrecInit` initializes and allocates memory for the KINBBDPRE preconditioner.

Arguments

- `kin_mem` (void *) pointer to the KINSOL memory block.
- `Nlocal` (long int) local vector length.
- `mudq` (long int) upper half-bandwidth to be used in the difference-quotient Jacobian approximation.
- `mldq` (long int) lower half-bandwidth to be used in the difference-quotient Jacobian approximation.
- `mukeep` (long int) upper half-bandwidth of the retained banded approximate Jacobian block.
- `mlkeep` (long int) lower half-bandwidth of the retained banded approximate Jacobian block.
- `dq_rel_u` (realtype) the relative increment in components of `u` used in the difference quotient approximations. The default is `dq_rel_u = $\sqrt{\text{unit roundoff}}$` , which can be specified by passing `dq_rel_u = 0.0`.
- `Gloc` (KINLocalFn) the C function which computes the approximation $G(u) \approx F(u)$.
- `Gcomm` (KINCommFn) the optional C function which performs all interprocess communication required for the computation of $G(u)$.

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

- `KINSPILS_SUCCESS` The call to `KINBBDPrecInit` was successful.
- `KINSPILS_MEM_NULL` The `kin_mem` pointer was `NULL`.
- `KINSPILS_MEM_FAIL` A memory allocation request has failed.
- `KINSPILS_LMEM_NULL` A KINSPILS linear solver was not attached.
- `KINSPILS_ILL_INPUT` The supplied vector implementation was not compatible with block band preconditioner.

Notes If one of the half-bandwidths `mudq` or `mldq` to be used in the difference-quotient calculation of the approximate Jacobian is negative or exceeds the value `Nlocal-1`, it is replaced with 0 or `Nlocal-1` accordingly.

The half-bandwidths `mudq` and `mldq` need not be the true half-bandwidths of the Jacobian of the local block of G , when smaller values may provide greater efficiency.

Also, the half-bandwidths `mukeep` and `mlkeep` of the retained banded approximate Jacobian block may be even smaller, to reduce storage and computation costs further.

For all four half-bandwidths, the values need not be the same for every process.

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

KINBBDPrecGetWorkSpace

Call `flag = KINBBDPrecGetWorkSpace(kin_mem, &lenrwBBDP, &leniwBBDP);`

Description The function `KINBBDPrecGetWorkSpace` returns the local KINBBDPRE real and integer workspace sizes.

Arguments

- `kin_mem` (void *) pointer to the KINSOL memory block.
- `lenrwBBDP` (long int) local number of `realtype` values in the KINBBDPRE workspace.
- `leniwBBDP` (long int) local number of integer values in the KINBBDPRE workspace.

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

KINSPILS_SUCCESS The optional output value has been successfully set.
 KINSPILS_MEM_NULL The `kin_mem` pointer was NULL.
 KINSPILS_PMEM_NULL The KINBBDPRE preconditioner has not been initialized.

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

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

KINBBDPrecGetNumGfnEvals

Call `flag = KINBBDPrecGetNumGfnEvals(kin_mem, &ngevalsBBDP);`

Description The function `KINBBDPrecGetNumGfnEvals` returns the number of calls to the user `Gloc` function due to the difference quotient approximation of the Jacobian blocks used within KINBBDPRE's preconditioner setup function.

Arguments `kin_mem` (void *) pointer to the KINSOL memory block.
`ngevalsBBDP` (long int) the number of calls to the user `Gloc` function.

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

KINSPILS_SUCCESS The optional output value has been successfully set.
 KINSPILS_MEM_NULL The `kin_mem` pointer was NULL.
 KINSPILS_PMEM_NULL The KINBBDPRE preconditioner has not been initialized.

In addition to the `ngevalsBBDP Gloc` evaluations, the costs associated with KINBBDPRE also include `nlinsetups` LU factorizations, `nlinsetups` calls to `Gcomm`, `npsolves` banded backsolve calls, and `nfevalsLS` right-hand side function evaluations, where `nlinsetups` is an optional KINSOL output and `npsolves` and `nfevalsLS` are linear solver optional outputs (see §4.5.5).

Chapter 5

FKINSOL, an Interface Module for FORTRAN Applications

The FKINSOL interface module is a package of C functions which support the use of the KINSOL solver, for the solution nonlinear systems $F(u) = 0$, in a mixed FORTRAN/C setting. While KINSOL 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 KINSOL for both the serial and the parallel NVECTOR implementations.

5.1 FKINSOL routines

The user-callable functions, with the corresponding KINSOL 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 KINSOL module
 - FKINMALLOC interfaces to KINCreate, KINSetUserData, and KINInit.
 - FKINSETIIN and FKINSETRIN interface to KINSet* functions.
 - FKINSETVIN interfaces to KINSetConstraints.
 - FKINSOL interfaces to KINSol, KINGet* functions, and to the optional output functions for the selected linear solver module.
 - FKINFREE interfaces to KINFree.
- Interface to the linear solver modules
 - FKINDENSE interfaces to KINDense.
 - FKINDENSESETJAC interfaces to KINDlsSetDenseJacFn.
 - FKINLAPACKDENSE interfaces to KINLapackDense.
 - FKINLAPACKDENSESETJAC interfaces to KINDlsSetDenseJacFn.
 - FKINBAND interfaces to KINBand.
 - FKINBANDSETJAC interfaces to KINDlsSetBandJacFn.
 - FKINLAPACKBAND interfaces to KINLapackBand.
 - FKINLAPACKBANDSETJAC interfaces to KINDlsSetBandJacFn.

- FKINSPGMR interfaces to KINSpgmr and SPGMR optional input functions.
- FKINSPBCG interfaces to KINSpbcg and SPBCG optional input functions.
- FKINSPTFQMR interfaces to KINSptfqmr and SPTFQMR optional input functions.
- FKINSPILSSETJAC interfaces to KINSpilsSetJacTimesVecFn.
- FKINSPILSSETPREC interfaces to KINSpilsSetPreconditioner.

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

FKINSOL routine (FORTRAN, user-supplied)	KINSOL function (C, interface)	KINSOL type of interface function
FKFUN	FKINfunc	KINSysFn
FKDJAC	FKINDenseJac	KINDlsDenseJacFn
	FKINLapackDenseJac	KINDlsDenseJacFn
FKBJAC	FKINBandJac	KINDlsBandJacFn
	FKINLapackBandJac	KINDlsBandJacFn
FKPSET	FKINPSet	KINSpilsPrecSetupFn
FKPSOL	FKINPSol	KINSpilsPrecSolveFn
FKJTICES	FKINJtimes	KINSpilsJacTimesVecFn

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

5.2 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 `fkinsol.h` and `fkinbdd.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 KINSOL. 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 KINSOL 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 optional input (`IVAL`), the problem size and bandwidth parameters (`NEQ`, `NLOCAL`, `NGLOBAL`, `ML`, `MU`, etc.), and the array of integer optional outputs (`IOUT`). EXCEPTION: In the case that Lapack linear solvers are to be used, the input arguments to `FKINLAPACKDENSE` or `FKINLAPACKBAND` (`NEQ`, `ML`, and `MU`) must be declared to be consistent with C type `int`. This type consistency is particularly important when using KINSOL and the FKINSOL package on 64-bit architectures.

5.3 Usage of the FKINSOL interface module

The usage of FKINSOL requires calls to a few different 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 KINSOL 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.

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. Nonlinear system function specification

The user must in all cases supply the following FORTRAN routine

```
SUBROUTINE FKFUN (U, FVAL, IER)
  DIMENSION U(*), FVAL(*)
```

It must set the FVAL array to $F(u)$, the system function, as a function of $U = u$. 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 KINSOL will attempt to correct), or a negative value if it failed unrecoverably (in which case the solution process 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 = 3 for KINSOL), 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 = 3, 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.

NOTE: The integers NEQ, NLOCAL, and NGLOBAL should be declared so as to match C type long int.

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:

```
FKINMALLOC
```

Call	CALL FKINMALLOC (IOUT, ROUT, IER)
Description	This function specifies the optional output arrays, allocates internal memory, and initializes KINSOL.
Arguments	IOUT is an integer array for integer optional outputs. ROUT is a real array for real optional outputs.
Return value	IER is the return completion flag. Values are 0 for successful return and -1 otherwise. See printed message for details in case of failure.
Notes	The user integer data array IOUT must be declared as INTEGER*4 or INTEGER*8 according to the C type long int. The optional outputs associated with the main KINSOL integrator are listed in Table 5.2.

4. Set optional inputs

Call FKINSETIIN and/or FKINSETRIN to set desired optional inputs, if any. See §5.4 for details.

5. Linear solver specification

The solution method in KINSOL involves the solution of linear systems related to the Jacobian of the nonlinear system. KINSOL presently includes five choices for the treatment of these systems, and the user of FKINSOL must call a routine with a specific name to make the desired choice.

[S] Dense treatment of the linear system

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

```
CALL FKINDENSE (NEQ, IER)
```

where NEQ is the size of the nonlinear 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 FKINLAPACKDENSE(NEQ, IER)
```

where the arguments have the same meanings as for FKINDENSE, except that here NEQ must be declared so as to match C type `int`.

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

```
SUBROUTINE FKDJAC (NEQ, U, FVAL, DJAC, WK1, WK2, IER)
  DIMENSION U(*), FVAL(*), DJAC(NEQ,*), WK1(*), WK2(*)
```

Typically this routine will use only NEQ, U, and DJAC. It must compute the Jacobian and store it columnwise in DJAC. The input arguments U and FVAL contain the current values of u and $F(u)$, respectively. The vectors WK1 and WK2, of length NEQ, are provided as work space for use in FKDJAC. 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 KINSOL will attempt to correct), or a negative value if FKDJAC failed unrecoverably (in which case the solution process is halted). NOTE: The argument NEQ has a type consistent with C type `long int` even in the case when the Lapack dense solver is to be used.

If the FKDJAC routine is provided, then, following the call to FKINDENSE, the user must make the call:

```
CALL FKINDENSESETJAC (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 FKINLAPACKDENSESETJAC (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 KINSOL implementation, the user must make the call:

```
CALL FKINBAND (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 FKINLAPACKBAND(NEQ, MU, ML, IER)
```

where the arguments have the same meanings as for FKINBAND, except that here NEQ, MU, and ML must be declared so as to match C type `int`.

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

```
SUBROUTINE FKBJAC (NEQ, MU, ML, MDIM, U, FVAL, BJAC, WK1, WK2, IER)
DIMENSION U(*), FVAL(*), BJAC(MDIM,*), WK1(*), WK2(*)
```

Typically this routine will use only NEQ, MU, ML, U, and BJAC. It must load the MDIM by N array BJAC with the Jacobian matrix at the current u 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 U and FVAL contain the current values of u , and $F(u)$, respectively. The vectors WK1 and WK2 of length NEQ are provided as work space for use in FKBJAC. 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 KINSOL will attempt to correct), or a negative value if FKBJAC failed unrecoverably (in which case the solution process is halted). NOTE: The arguments NEQ, MU, ML, and MDIM have a type consistent with C type `long int` even in the case when the Lapack band solver is to be used.

If the FKBJAC routine is provided, then, following the call to FKINBAND, the user must make the call:

```
CALL FKINBANDSETJAC (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 FKINLAPACKNBANDSETJAC (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 FKINSPGMR (MAXL, MAXLRST, IER)
```

The arguments are as follows. MAXL is the maximum Krylov subspace dimension. MAXLRST is the maximum number of restarts. 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 FKINSPBCG (MAXL, IER)
```

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

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 FKINSPTFQMR (MAXL, IER)
```

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

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, FKINJTIMES (see below), can be provided for Jacobian-vector products. If it is, then, following the call to FKINSPGMR, FKINSPBCG, or FKINSPTFQMR, the user must make the call:

```
CALL FKINSPILSSETJAC (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, then the user must call

```
CALL FKINSPILSSETPREC (FLAG, IER)
```

with $FLAG \neq 0$. The return flag IER is 0 if successful, or negative if a memory error occurred. In addition, the user program must include preconditioner routines FKPSOL and FKPSET (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 — FKINJTIMES, FKPSOL, and FKPSET. 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 $J = \partial F / \partial u$ and a given vector v . If supplied, it must have the following form:

```
SUBROUTINE FKINJTIMES (V, FJV, NEWU, U, IER)
  DIMENSION V(*), FJV(*), U(*)
```

Typically this routine will use only U, 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 argument U contains the current value of u . On return, set IER = 0 if FKINJTIMES was successful, and nonzero otherwise.

NEWU is a flag to indicate if U has been changed since the last call; if it has, then NEWU = 1, and FKINJTIMES should recompute any saved Jacobian data it uses and reset NEWU to 0. (See §4.6.6.)

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

```
SUBROUTINE FKPSOL (U, USCALE, FVAL, FSCALE, VTEM, FTEM, IER)
  DIMENSION U(*), USCALE(*), FVAL(*), FSCALE(*), VTEM(*), FTEM(*)
```

Typically this routine will use only U, FVAL, VTEM and FTEM. It must solve the preconditioned linear system $Pz = r$, where $r = VTEM$ is input, and store the solution z in VTEM as well. Here P is the right preconditioner. If scaling is being used, the routine supplied must also account for scaling on either coordinate or function value, as given in the arrays USCALE and FSCALE, respectively.

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 FKPSET (U, USCALE, FVAL, FSCALE, VTEMP1, VTEMP2, IER)
  DIMENSION U(*), USCALE(*), FVAL(*), FSCALE(*), VTEMP1(*), VTEMP2(*)
```

It must perform any evaluation of Jacobian-related data and preprocessing needed for the solution of the preconditioned linear systems by FKPSOL. The variables U through FSCALE are for use in the preconditioning setup process. Typically, the system function FKFUN is called before any calls to FKPSET, so that FVAL will have been updated. U is the current solution iterate. The arrays VTEMP1 and VTEMP2 are available for work space. If scaling is being used, USCALE and FSCALE are available for those operations requiring scaling.

On return, set IER = 0 if FKPSET was successful, or set IER = 1 if an error occurred.

If the user calls FKINSPILSSETPREC, the routine FKPSET must be provided, even if it is not needed, and then it should return IER = 0.



6. Problem solution

Solving the nonlinear system is accomplished by making the following call:

```
CALL FKINSOL (U, GLOBALSTRAT, USCALE, FSCALE, IER)
```

The arguments are as follows. U is an array containing the initial guess on input, and the solution on return. GLOBALSTRAT is an integer (type INTEGER) defining the global strategy choice (1 specifies Inexact Newton, while 2 indicates line search). USCALE is an array of scaling factors for the U vector. FSCALE is an array of scaling factors for the FVAL vector. IER is an integer completion flag and will have one of the following values: 0 to indicate success, 1 to indicate that the initial guess satisfies $F(u) = 0$ within tolerances, 2 to indicate apparent stalling (small step), or a negative value to indicate an error or failure. These values correspond to the KINSOL returns (see §4.5.3 and §B.2). The values of the optional outputs are available in IOPT and ROPT (see Table 5.2).

7. Memory deallocation

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

```
CALL FKINFREE
```

5.4 FKINSOL optional input and output

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

Table 5.1: Keys for setting FKINSOL optional inputs

Integer optional inputs FKINSETIIN		
Key	Optional input	Default value
PRNT_LEVEL	Verbosity level of output	0
MAX_NITER	Maximum no. of nonlinear iterations	200
ETA_FORM	Form of η coefficient	1 (KIN_ETACHOICE1)
MAX_SETUPS	Maximum no. of iterations without prec. setup	10
MAX_SP_SETUPS	Maximum no. of iterations without residual check	5
NO_INIT_SETUP	No initial preconditioner setup	FALSE
NO_MIN_EPS	Lower bound on ϵ	FALSE
NO_RES_MON	No residual monitoring	FALSE

Real optional inputs (FKINSETRIN)		
Key	Optional input	Default value
FNORM_TOL	Function-norm stopping tolerance	$\text{uround}^{1/3}$
SSTEP_TOL	Scaled-step stopping tolerance	$\text{uround}^{2/3}$
MAX_STEP	Max. scaled length of Newton step	$1000 \ D_u u_0\ _2$
RERR_FUNC	Relative error for F.D. Jv	$\sqrt{\text{uround}}$
ETA_CONST	Constant value of η	0.1
ETA_PARAMS	Values of γ and α	0.9 and 2.0
RMON_CONST	Constant value of ω	0.9
RMON_PARAMS	Values of ω_{min} and ω_{max}	0.00001 and 0.9

CALL FKINSETIIN (KEY, IVAL, IER)

CALL FKINSETRIN (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 integer **IVAL** should be declared in a manner consistent with C type `long int`.

The optional outputs from the KINSOL solver are accessed not through individual functions, but rather through a pair of arrays, **IOUT** (integer type) of dimension at least 15, and **ROUT** (real type) of dimension at least 2. These arrays are owned (and allocated) by the user and are passed as arguments to **FKINMALLOC**. Table 5.2 lists the entries in these two arrays and specifies the optional variable as well as the KINSOL function which is actually called to extract the optional output.

For more details on the optional inputs and outputs, see §4.5.4 and §4.5.5.

5.5 Usage of the FKINBBDD interface to KINBBDDPRE

The FKINBBDD interface sub-module is a package of C functions which, as part of the FKINSOL interface module, support the use of the KINSOL solver with the parallel NVECTOR_PARALLEL module and the KINBBDDPRE preconditioner module (see §4.7), for the solution of nonlinear problems in a mixed FORTRAN/C setting.

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

- FKINBBDDINIT interfaces to KINBBDDPrecInit.
- FKINBBDDOPT interfaces to KINBBDDPRE optional output functions.

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

Table 5.2: Description of the FKINSOL optional output arrays IOU and ROUT

Integer output array IOU		
Index	Optional output	KINSOL function
KINSOL main solver		
1	LENRW	KINGGetWorkSpace
2	LENIW	KINGGetWorkSpace
3	NNI	KINGGetNumNonlinSolvIters
4	NFE	KINGGetNumFuncEvals
5	NBCF	KINGGetNumBetaCondFails
6	NBKTRK	KINGGetNumBacktrackOps
KINDENSE, KINBAND linear solver		
7	LENRWLS	KINDlsGetWorkSpace
8	LENIWLS	KINDlsGetWorkSpace
9	LS_FLAG	KINDlsGetLastFlag
10	NFELS	KINDlsGetNumFuncEvals
11	NJE	KINDlsGetNumJacEvals
KINSPGMR, KINSPBCG, KINSPTFQMR linear solvers		
7	LENRWLS	KINSpilsGetWorkSpace
8	LENIWLS	KINSpilsGetWorkSpace
9	LS_FLAG	KINSpilsGetLastFlag
10	NFELS	KINSpilsGetNumFuncEvals
11	NJTV	KINSpilsGetNumJacEvals
12	NPE	KINSpilsGetNumPrecEvals
13	NPS	KINSpilsGetNumPrecSolves
14	NLI	KINSpilsGetNumLinIters
15	NCFL	KINSpilsGetNumConvFails
Real output array ROUT		
Index	Optional output	KINSOL function
1	FNORM	KINGGetFuncNorm
2	SSTEP	KINGGetStepLength

FKINBBB routine (FORTRAN, user-supplied)	CVODE function (C, interface)	CVODE type of interface function
FKLOCFN	FKINgloc	KINLocalFn
FKCOMMF	FKINgcomm	KINCommFn
FKJTIMES	FKINJtimes	KINSpilsJacTimesVecFn

As with the rest of the FKINSOL 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 `fkinbbd.h`.

The following is a summary of the usage of this module. Steps that are unchanged from the main program described in §5.3 are grayed-out.

1. Nonlinear system function specification
2. NVECTOR module initialization
3. Problem specification
4. Set optional inputs
5. Linear solver specification

First, specify one of the KINSPILS iterative linear solvers, by calling one of FKINSPGMR, FKINSPBCG, or FKINSPTFQMR.

To initialize the KINBBDPRE preconditioner, make the following call:

```
CALL FKINBBDINIT (NLOCAL, MUDQ, MLDQ, MU, ML, IER)
```

The arguments are as follows. `NLOCAL` is the local size of vectors for this process. `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`. `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 FKJTIMES, make the call

```
CALL FKINSPILSSETJAC (FLAG, IER)
```

with `FLAG` $\neq 0$. (See step 5 in §5.3).

6. Problem solution
7. KINBBDPRE **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 KINBBDPRE module, make the following call:

```
CALL FKINBBDOPT (LENRBBB, LENIBBB, NGEBBB)
```

The arguments returned are as follows. `LENRBBB` is the length of real preconditioner work space, in `realtype` words. `LENIBBB` is the length of integer preconditioner work space, in integer words. These sizes are local to the current process. `NGEBBB` is the cumulative number of $G(u)$ evaluations (calls to FKLOCFN) so far.

8. Memory deallocation

(The memory allocated for the FKINBBB module is deallocated automatically by FKINFREE.)

9. User-supplied routines

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

```
SUBROUTINE FKLOCFN (NLOC, ULOC, GLOC, IER)
  DIMENSION ULOC(*), GLOC(*)
```

This routine is to evaluate the function $G(u)$ approximating F (possibly identical to F), in terms of the array `ULOC` (of length `NLOC`), which is the sub-vector of u local to this processor. The resulting (local) sub-vector is to be stored in the array `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 `KINSOL` will attempt to correct), or a negative value if `FKLOCFN` failed unrecoverably (in which case the solution process is halted).

```
SUBROUTINE FKCOMMFN (NLOC, ULOC, IER)
  DIMENSION ULOC(*)
```

This routine is to perform the inter-processor communication necessary for the `FKLOCFN` routine. Each call to `FKCOMMFN` is preceded by a call to the system function routine `FKFUN` with the same argument `ULOC`. `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 `KINSOL` will attempt to correct), or a negative value if `FKCOMMFN` failed recoverably (in which case the solution process is halted).

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

Optionally, the user can supply a routine `FKINJTIMES` for the evaluation of Jacobian-vector products, as described above in step 5 in §5.3.



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    (*nvl1norm)(N_Vector);
void        (*nvcompare)(realtype, N_Vector, N_Vector);
booleantype (*nvinvtest)(N_Vector, N_Vector);
booleantype (*nvconstrmask)(N_Vector, N_Vector, N_Vector);
realtype    (*nvminquotient)(N_Vector, N_Vector);
};

```

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

```

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

```

Table 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	$v = \text{N_VClone}(w);$ Creates a new <code>N_Vector</code> of the same type as an existing vector <code>w</code> and sets the <code>ops</code> field. It does not copy the vector, but rather allocates storage for the new vector.
N_VCloneEmpty	$v = \text{N_VCloneEmpty}(w);$ Creates a new <code>N_Vector</code> of the same type as an existing vector <code>w</code> and sets the <code>ops</code> field. It does not allocate storage for the data array.
N_VDestroy	$\text{N_VDestroy}(v);$ Destroys the <code>N_Vector</code> <code>v</code> and frees memory allocated for its internal data.
N_VSpace	$\text{N_VSpace}(nvSpec, \&lrw, \&liw);$ Returns storage requirements for one <code>N_Vector</code> . <code>lrw</code> contains the number of realtype words and <code>liw</code> contains the number of integer words. This function is advisory only, for use in determining a user's total space requirements; it could be a dummy function in a user-supplied NVECTOR module if that information is not of interest.
N_VGetArrayPointer	$vdata = \text{N_VGetArrayPointer}(v);$ Returns a pointer to a <code>realtype</code> array from the <code>N_Vector</code> <code>v</code> . Note that this assumes that the internal data in <code>N_Vector</code> is a contiguous array of <code>realtype</code> . This routine is only used in the solver-specific interfaces to the dense and banded (serial) linear solvers, and in the interfaces to the banded (serial) and band-block-diagonal (parallel) preconditioner modules provided with SUNDIALS.
N_VSetArrayPointer	$\text{N_VSetArrayPointer}(vdata, v);$ Overwrites the data in an <code>N_Vector</code> with a given array of <code>realtype</code> . Note that this assumes that the internal data in <code>N_Vector</code> is a contiguous array of <code>realtype</code> . This routine is only used in the interfaces to the dense (serial) linear solver, hence need not exist in a user-supplied NVECTOR module for a parallel environment.
N_VLinearSum	$\text{N_VLinearSum}(a, x, b, y, z);$ Performs the operation $z = ax + by$, where a and b are scalars and x and y are of type <code>N_Vector</code> : $z_i = ax_i + by_i, i = 0, \dots, n - 1$.
N_VConst	$\text{N_VConst}(c, z);$ Sets all components of the <code>N_Vector</code> <code>z</code> to <code>c</code> : $z_i = c, i = 0, \dots, n - 1$.
N_VProd	$\text{N_VProd}(x, y, z);$ Sets the <code>N_Vector</code> <code>z</code> to be the component-wise product of the <code>N_Vector</code> inputs <code>x</code> and <code>y</code> : $z_i = x_i y_i, i = 0, \dots, n - 1$.

continued on next page

<i>continued from last page</i>	
Name	Usage and Description
N_VDiv	<code>N_VDiv(x, y, z);</code> Sets the <code>N_Vector</code> <code>z</code> to be the component-wise ratio of the <code>N_Vector</code> inputs <code>x</code> and <code>y</code> : $z_i = x_i/y_i$, $i = 0, \dots, n-1$. The y_i may not be tested for 0 values. It should only be called with a <code>y</code> that is guaranteed to have all nonzero components.
N_VScale	<code>N_VScale(c, x, z);</code> Scales the <code>N_Vector</code> <code>x</code> by the 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> : $m = \sqrt{\sum_{i=0}^{n-1} (x_i w_i)^2}$.

continued on next page

<i>continued from last page</i>	
Name	Usage and Description
N_VL1Norm	<code>m = N_VL1Norm(x);</code> Returns the ℓ_1 norm of the N_Vector <code>x</code> : $m = \sum_{i=0}^{n-1} x_i $.
N_VCompare	<code>N_VCompare(c, x, z);</code> Compares the components of the N_Vector <code>x</code> to the scalar <code>c</code> and returns an N_Vector <code>z</code> such that: $z_i = 1.0$ if $ x_i \geq c$ and $z_i = 0.0$ otherwise.
N_VInvTest	<code>t = N_VInvTest(x, z);</code> Sets the components of the N_Vector <code>z</code> to be the inverses of the components of the N_Vector <code>x</code> , with prior testing for zero values: $z_i = 1.0/x_i$, $i = 0, \dots, n-1$. This routine returns <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;
    booleantype own_data;
    realtype *data;
    MPI_Comm comm;
};
```

The following seven macros are provided to access the content of a `NVECTOR_PARALLEL` vector. The suffix `_P` in the names denotes 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 KINSOL

In Table 6.2 below, we list the vector functions in the `NVECTOR` module within the `KINSOL` package. The table also shows, for each function, which of the code modules uses the function. The `KINSOL` column shows function usage within the main solver module, while the remaining five columns show function usage within each of the four `KINSOL` linear solvers (`KINSPILS` stands for any of `KINSPGMR`, `KINSPBCG`, or `KINSPTFQMR`), the `KINBBDPRE` preconditioner module, and the `FKINSOL` module.

There is one subtlety in the `KINSPILS` column hidden by the table, explained here for the case of the `KINSPGMR` module. The `N_VDotProd` function is called both within the interface file `kinsol_spgmr.c` and within the implementation files `sundials_spgmr.c` and `sundials_iterative.c` for the generic `SPGMR` solver upon which the `KINSPGMR` solver is built. Also, although `N_VDiv` and `N_VProd` are not called within the interface file `kinsol_spgmr.c`, they are called within the implementation file `sundials_spgmr.c`, and so are required by the `KINSPGMR` solver module. Analogous statements apply to the `KINSPBCG` and `KINSPTFQMR` modules, except that they do not use `sundials_iterative.c`. This issue does not arise for the direct `KINSOL` linear solvers because the generic `DENSE` and `BAND` solvers (used in the implementation of `KINDENSE` and `KINBAND`) do not make calls to any vector functions.

At this point, we should emphasize that the `KINSOL` user does not need to know anything about the usage of vector functions by the `KINSOL` code modules in order to use `KINSOL`. 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 `KINSOL` are: `N_VAddConst`, `N_VWrmsNorm`, `N_VWrmsNormMask`, `N_VCompare`, and `N_VInvTest`. Therefore a user-supplied `NVECTOR` module for `KINSOL` could omit these five functions.

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

	KINSOL	KINDENSE	KINBAND	KINSPILS	KINBDDPRE	FKINSOL
N_VClone	✓			✓	✓	
N_VCloneEmpty						✓
N_VDestroy	✓			✓	✓	✓
N_VSpace	✓					
N_VGetArrayPointer		✓	✓		✓	✓
N_VSetArrayPointer		✓				✓
N_VLinearSum	✓	✓		✓		
N_VConst				✓		
N_VProd	✓	✓	✓	✓		
N_VDiv	✓			✓		
N_VScale	✓	✓	✓	✓	✓	
N_VAbs	✓					
N_VInv	✓					
N_VDotProd		✓	✓	✓		
N_VMaxNorm	✓					
N_VMin	✓					
N_VWL2Norm	✓	✓	✓	✓		
N_VL1Norm				✓		
N_VConstrMask	✓					
N_VMinQuotient	✓					

Chapter 7

Providing Alternate Linear Solver Modules

The central KINSOL 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 that described in §4.5.2) which will attach the above four routines to the main KINSOL memory block. The KINSOL memory block is a structure defined in the header file `kinsol_impl.h`. A pointer to such a structure is defined as the type `KINMem`. The four fields in a `KINMem` structure that must point to the linear solver's functions are `kin_linit`, `kin_lsetup`, `kin_lsolve`, and `kin_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. The linear solver specification function must also set the value of the field `kin_setupNonNull` in the KINSOL memory block — to `TRUE` if `lsetup` is used, or `FALSE` otherwise.

For consistency with the existing KINSOL 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 KINSOL 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 `kin_lmem` in the KINSOL memory block can be used to attach a linear solver-specific memory block. That memory should be allocated in the linear solver specification function.

These four routines, which interface between KINSOL and the linear solver module, necessarily have fixed call sequences. Thus, a user wishing to implement another linear solver within the KINSOL 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 KINSOL memory block, by which the routine can access various data related to the KINSOL solution. The contents of this memory block are given in the file `kinsol_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)(KINMem kin_mem);`

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

Arguments `kin_mem` is the KINSOL memory pointer of type `KINMem`.

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

7.2 Setup function

The type definition of `lsetup` is

lsetup

Definition `int (*lsetup)(KINMem kin_mem);`

Purpose The job of `lsetup` is to prepare the linear solver for subsequent calls to `lsolve`. It may recompute Jacobian-related data if it deems necessary.

Arguments `kin_mem` is the KINSOL memory pointer of type `KINMem`.

Return value The `lsetup` routine should return 0 if successful and a non-zero value otherwise.

7.3 Solve function

The type definition of `lsolve` is

lsolve

Definition `int (*lsolve)(KINMem kin_mem, N_Vector x,
N_Vector b, realtype *res_norm);`

Purpose The routine `lsolve` must solve the linear equation $Jx = b$, where $J = \partial F/\partial u$ is evaluated at the current iterate and the right-hand side vector b is input.

Arguments `kin_mem` is the KINSOL memory pointer of type `KINMem`.

`x` is a vector set to an initial guess prior to calling `lsolve`. On return it should contain the solution to $Jx = b$.

`b` is the right-hand side vector b , set to $-F(u)$, evaluated at the current iterate.
`res_norm` holds the value of the L_2 norm of the residual vector upon return.

Return value `lsolve` should return 0 if successful. If an error occurs and recovery could be possible by calling again the `lsetup` function, then it should return a positive value. Otherwise, `lsolve` should return a negative value.

7.4 Memory deallocation function

The type definition of `lfree` is

lfree

Definition `void (*lfree)(KINMem kin_mem);`

Purpose The routine `lfree` should free any linear solver memory allocated by the `linit` routine.

Arguments `kin_mem` is the KINSOL memory pointer of type `KINMem`.

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 SUNDIALS, but which are of potential use as generic packages in themselves, either in conjunction with the use of KINSOL 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 `RCNST`, 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;
    long int M;
    long int N;
    long int ldim;
    long int mu;
    long int ml;
    long int s_mu;
    realtype *data;
    long int ldata;
    realtype **cols;
} *DlsMat;
```

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

`type` - SUNDIALS_DENSE (=1)

`M` - number of rows

`N` - number of columns

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

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

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

cols - array of pointers. $\text{cols}[j]$ 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 $(\mathbf{A} \rightarrow \text{cols})[j][i]$

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, $0 \leq \text{mu} < \min(M,N)$

m1 - lower half-bandwidth, $0 \leq \text{m1} < \min(M,N)$

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 $\min(N-1, \text{mu} + \text{m1})$ 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 ($= \text{ldim} \cdot (\text{s_mu} + \text{m1} + 1)$)

cols - array of pointers. $\text{cols}[j]$ is a pointer to the uppermost element within the band in the j -th column. This pointer may be treated as an array indexed from $\text{s_mu} - \text{mu}$ (to access the uppermost element within the band in the j -th column) to $\text{s_mu} + \text{m1}$ (to access the lowest element within the band in the j -th column). Indices from 0 to $\text{s_mu} - \text{mu} - 1$ give access to extra storage elements required by the LU decomposition function. Finally, $\text{cols}[j][i - j + \text{s_mu}]$ is the (i, j) -th element, $j - \text{mu} \leq i \leq j + \text{m1}$.

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: $\text{DENSE_ELEM}(\mathbf{A}, i, j) = \mathbf{a}_{ij}$; or $\mathbf{a}_{ij} = \text{DENSE_ELEM}(\mathbf{A}, i, j)$;

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 ml , respectively. The rows and columns of A are numbered from 0 to $N - 1$ and the (i, j) -th element of A is denoted $A(i, j)$. The greyed out areas of the underlying component storage are used by the `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;
- `DestroyMat`: free memory for a `DlsMat` matrix;
- `PrintMat`: print a `DlsMat` matrix to standard output.
- `NewLintArray`: allocation of an array of `long int` integers for use as pivots with `DenseGETRF` and `DenseGETRS`;
- `NewIntArray`: allocation of an array of `int` integers for use as pivots with the Lapack dense solvers;
- `NewRealArray`: allocation of an array of `realtype` for use as right-hand side with `DenseGETRS`;
- `DestroyArray`: free memory for an array;
- `SetToZero`: load a matrix with zeros;
- `AddIdentity`: increment a square matrix by the identity matrix;

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

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

- `newDenseMat`

`newDenseMat(m,n)` allocates storage for an m by n dense matrix. It returns a pointer to the newly allocated storage if successful. If the memory request cannot be satisfied, then `newDenseMat` returns NULL. The underlying type of the dense matrix returned is `realtype**`. If we allocate a dense matrix `realtype** a` by `a = newDenseMat(m,n)`, then `a[j][i]` references the (i,j) -th element of the matrix `a`, $0 \leq i < m$, $0 \leq j < n$, and `a[j]` is a pointer to the first element in the j -th column of `a`. The location `a[0]` contains a pointer to $m \times n$ contiguous locations which contain the elements of `a`.

- `destroyMat`

`destroyMat(a)` frees the dense matrix `a` allocated by `newDenseMat`;

- `newLintArray`

`newLintArray(n)` allocates an array of n integers, all `long int`. It returns a pointer to the first element in the array if successful. It returns NULL if the memory request could not be satisfied.

- `newIntArray`

`newIntArray(n)` allocates an array of n integers, all `int`. It returns a pointer to the first element in the array if successful. It returns NULL if the memory request could not be satisfied.

- `newRealArray`

`newRealArray(n)` allocates an array of n `realtype` values. It returns a pointer to the first element in the array if successful. It returns NULL if the memory request could not be satisfied.

- `destroyArray`

`destroyArray(p)` frees the array `p` allocated by `newLintArray`, `newIntArray`, or `newRealArray`;

- `denseCopy`

`denseCopy(a,b,m,n)` copies the m by n dense matrix `a` into the m by n dense matrix `b`;

- `denseScale`

`denseScale(c,a,m,n)` scales every element in the m by n dense matrix `a` by the scalar `c`;

- `denseAddIdentity`

`denseAddIdentity(a,n)` increments the *square* n by n dense matrix `a` by the identity matrix I_n ;

- **denseGETRF**

denseGETRF(*a*,*m*,*n*,*p*) factors the *m* by *n* dense matrix *a*, using Gaussian elimination with row pivoting. It overwrites the elements of *a* with its LU factors and keeps track of the pivot rows chosen in the pivot array *p*.

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

1. *p*[*k*] contains the row number of the pivot element chosen at the beginning of elimination step *k*, *k* = 0, 1, ..., *n*−1.
2. If the unique LU factorization of *a* is given by $Pa = LU$, where *P* is a permutation matrix, *L* is an *m* by *n* lower trapezoidal matrix with all diagonal elements equal to 1, and *U* is an *n* by *n* upper triangular matrix, then the upper triangular part of *a* (including its diagonal) contains *U* and the strictly lower trapezoidal part of *a* contains the multipliers, $I - L$. If *a* is square, *L* is a unit lower triangular matrix.

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

- **denseGETRS**

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

- **densePOTRF**

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

- **densePOTRS**

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

- **denseGEQRF**

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

- **denseORMQR**

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

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;
- `DestroyMat`: free memory for a `DlsMat` matrix;
- `PrintMat`: print a `DlsMat` matrix to standard output.
- `NewLintArray`: allocation of an array of `int` integers for use as pivots with `BandGBRF` and `BandGBRS`;
- `NewIntArray`: allocation of an array of `int` integers for use as pivots with the Lapack band solvers;
- `NewRealArray`: allocation of an array of `realtype` for use as right-hand side with `BandGBRS`;
- `DestroyArray`: free memory for an array;
- `SetToZero`: load a matrix with zeros;
- `AddIdentity`: increment a square matrix by the identity matrix;
- `BandCopy`: copy one matrix to another;
- `BandScale`: scale a matrix by a scalar;
- `BandGBTRF`: LU factorization with partial pivoting;
- `BandGBTRS`: solution of $Ax = b$ using LU factorization;

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

- `newBandMat`
`newBandMat(n, smu, ml)` allocates storage for an `n` by `n` band matrix with lower half-bandwidth `ml`.
- `destroyMat`
`destroyMat(a)` frees the band matrix `a` allocated by `newBandMat`;
- `newLintArray`
`newLintArray(n)` allocates an array of `n` integers, all `long int`. It returns a pointer to the first element in the array if successful. It returns `NULL` if the memory request could not be satisfied.
- `newIntArray`
`newIntArray(n)` allocates an array of `n` integers, all `int`. It returns a pointer to the first element in the array if successful. It returns `NULL` if the memory request could not be satisfied.
- `newRealArray`
`newRealArray(n)` allocates an array of `n` `realtype` values. It returns a pointer to the first element in the array if successful. It returns `NULL` if the memory request could not be satisfied.
- `destroyArray`
`destroyArray(p)` frees the array `p` allocated by `newLintArray`, `newIntArray`, or `newRealArray`;
- `bandCopy`
`bandCopy(a,b,n,a_smu, b_smu,copymu, 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 `RCNST`, 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_spgmgs.h` and `sundials_spgmgs.c`, includes an implementation of the scaled preconditioned Bi-CGStab method. For full details, including usage instructions, see the file `sundials_spgmgs.h`.

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

The following functions are available in the SPBCG package:

- `SpgmgsMalloc`: allocation of memory for `SpgmgsSolve`;
- `SpgmgsSolve`: solution of $Ax = b$ by the SPBCG method;
- `SpgmgsFree`: free memory allocated by `SpgmgsMalloc`.

8.2.3 The SPTFQMR module

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

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

The following functions are available in the SPTFQMR package:

- `SptfqmrMalloc`: allocation of memory for `SptfqmrSolve`;
- `SptfqmrSolve`: solution of $Ax = b$ by the SPTFQMR method;
- `SptfqmrFree`: free memory allocated by `SptfqmrMalloc`.

Appendix A

KINSOL Installation Procedure

The installation of KINSOL 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 KINSOL.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 KINSOL 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/`); `kinbbdpre.c`, `kinbbdpre_impl.h` (in `src/kinsol/`); `kinbbdpre.h` (in `include/kinsol/`); `fkinbbd.c`, `fkinbbd.h` (in `src/kinsol/fcmix/`); all files in `examples/kinsol/parallel/`; all files in `examples/kinsol/fcmix_parallel/`. (By “serial version” of KINSOL we mean the KINSOL 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=BLASDIR`

Specify the Blas library.

Default: none

`--with-lapack=LAPACKDIR`

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` compiler 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 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
11
12 #define SUNDIALS_MPICOMM_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.*, *cvode_dense.h* includes *sundials_dense.h*). However, it is both legal and safe to do so (*e.g.*, the functions declared in *sundials_dense.h* could be used in building a preconditioner).

Table A.1: SUNDIALS libraries and header files

SHARED	Libraries	n/a	
	Header files	sundials/sundials_config.h sundials/sundials_math.h sundials/sundials_nvector.h sundials/sundials_direct.h sundials/sundials_dense.h sundials/sundials_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_fcvsode.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_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_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_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

KINSOL 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 KINSOL input constants

KINSOL main solver module		
KIN_ETACHOICE1	1	Use Eisenstat and Walker Choice 1 for η .
KIN_ETACHOICE2	2	Use Eisenstat and Walker Choice 2 for η .
KIN_ETACONSTANT	3	Use constant value for η .
KIN_NONE	0	Use inexact Newton globalization.
KIN_LINESEARCH	1	Use linesearch globalization.

Iterative linear solver module		
PREC_NONE	0	No preconditioning
PREC_RIGHT	2	Preconditioning on the right.
MODIFIED_GS	1	Use modified Gram-Schmidt procedure.
CLASSICAL_GS	2	Use classical Gram-Schmidt procedure.

B.2 KINSOL output constants

KINSOL main solver module		
KIN_SUCCESS	0	Successful function return.
KIN_INITIAL_GUESS_OK	1	The initial user-supplied guess already satisfies the stopping criterion.
KIN_STEP_LT_STPTOL	2	The stopping tolerance on scaled step length was satisfied.
KIN_WARNING	99	A non-fatal warning. The solver will continue.
KIN_MEM_NULL	-1	The <code>kin_mem</code> argument was NULL.
KIN_ILL_INPUT	-2	One of the function inputs is illegal.
KIN_NO_MALLOC	-3	The KINSOL memory was not allocated by a call to <code>KINMalloc</code> .
KIN_MEM_FAIL	-4	A memory allocation failed.
KIN_LINESEARCH_NONCONV	-5	The linesearch algorithm was unable to find an iterate sufficiently distinct from the current iterate.

KIN_MAXITER_REACHED	-6	The maximum number of nonlinear iterations has been reached.
KIN_MXNEWT_5X_EXCEEDED	-7	Five consecutive steps have been taken that satisfy a scaled step length test.
KIN_LINESEARCH_BCFAIL	-8	The linesearch algorithm was unable to satisfy the β -condition for <code>nbcfails</code> iterations.
KIN_LINSOLV_NO_RECOVERY	-9	The user-supplied routine preconditioner slve function failed recoverably, but the preconditioner is already current.
KIN_LINIT_FAIL	-10	The linear solver's initialization function failed.
KIN_LSETUP_FAIL	-11	The linear solver's setup function failed in an unrecoverable manner.
KIN_LSOLVE_FAIL	-12	The linear solver's solve function failed in an unrecoverable manner.
KIN_SYSFUNC_FAIL	-13	The system function failed in an unrecoverable manner.
KIN_FIRST_SYSFUNC_ERR	-14	The system function failed recoverably at the first call.
KIN_REPTD_SYSFUNC_ERR	-15	The system function had repeated recoverable errors.

KINDLS linear solver module

KINDLS_SUCCESS	0	Successful function return.
KINDLS_MEM_NULL	-1	The <code>kin_mem</code> argument was NULL.
KINDLS_LMEM_NULL	-2	The KINDLS linear solver has not been initialized.
KINDLS_ILL_INPUT	-3	The KINDLS solver is not compatible with the current NVECTOR module.
KINDLS_MEM_FAIL	-4	A memory allocation request failed.
KINDLS_JACFUNC_UNRECVR	-5	The Jacobian function failed in an unrecoverable manner.
KINDLS_JACFUNC_RECVR	-6	The Jacobian function had a recoverable error.

KINSPILS linear solver modules

KINSPILS_SUCCESS	0	Successful function return.
KINSPILS_MEM_NULL	-1	The <code>kin_mem</code> argument was NULL.
KINSPILS_LMEM_NULL	-2	The KINSPILS linear solver has not been initialized.
KINSPILS_ILL_INPUT	-3	The KINSPILS solver is not compatible with the current NVECTOR module, or an input value was illegal.
KINSPILS_MEM_FAIL	-4	A memory allocation request failed.
KINSPILS_PMEM_NULL	-5	The preconditioner module has not been initialized.

SPGMR generic linear solver module

SPGMR_SUCCESS	0	Converged.
SPGMR_RES_REDUCED	1	No convergence, but the residual norm was reduced.
SPGMR_CONV_FAIL	2	Failure to converge.
SPGMR_QRFACT_FAIL	3	A singular matrix was found during the QR factorization.
SPGMR_PSOLVE_FAIL_REC	4	The preconditioner solve function failed recoverably.
SPGMR_ATIMES_FAIL_REC	5	The Jacobian-times-vector function failed recoverably.
SPGMR_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.

SPBCG generic linear solver module

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_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.

SPTFQMR generic linear solver module

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_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.

Bibliography

- [1] P. N. Brown. A local convergence theory for combined inexact-Newton/finite difference projection methods. *SIAM J. Numer. Anal.*, 24(2):407–434, 1987.
- [2] P. N. Brown and Y. Saad. Hybrid Krylov Methods for Nonlinear Systems of Equations. *SIAM J. Sci. Stat. Comput.*, 11:450–481, 1990.
- [3] A. M. Collier and R. Serban. Example Programs for KINSOL v2.7.0. Technical Report UCRL-SM-208114, LLNL, 2011.
- [4] R. S. Dembo, S. C. Eisenstat, and T. Steihaug. Inexact Newton Methods. *SIAM J. Numer. Anal.*, 19:400–408, 1982.
- [5] J. E. Dennis and R. B. Schnabel. *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. SIAM, Philadelphia, 1996.
- [6] S. C. Eisenstat and H. F. Walker. Choosing the Forcing Terms in an Inexact Newton Method. *SIAM J. Sci. Comput.*, 17:16–32, 1996.
- [7] R. W. Freund. A Transpose-Free Quasi-Minimal Residual Algorithm for Non-Hermitian Linear Systems. *SIAM J. Sci. Comp.*, 14:470–482, 1993.
- [8] C. T. Kelley. *Iterative Methods for Solving Linear and Nonlinear Equations*. SIAM, Philadelphia, 1995.
- [9] 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.
- [10] 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

- BAND generic linear solver
 - functions, 85–86
 - small matrix, 86–87
 - macros, 83
 - type DlsMat, 80–81
- BAND_COL, 45, 83
- BAND_COL_ELEM, 45, 83
- BAND_ELEM, 45, 83
- bandAddIdentity, 87
- bandCopy, 86
- bandGETRF, 87
- bandGETRS, 87
- bandScale, 86
- Bi-CGStab method, 88
- BIG_REAL, 16, 69

- DENSE generic linear solver
 - functions
 - large matrix, 83–84
 - small matrix, 84–85
 - macros, 81–83
 - type DlsMat, 80–81
- DENSE_COL, 44, 83
- DENSE_ELEM, 44, 81
- denseAddIdentity, 84
- denseCopy, 84
- denseGEQRF, 85
- denseGETRF, 85
- denseGETRS, 85
- denseORMQR, 85
- densePOTRF, 85
- densePOTRS, 85
- denseScale, 84
- destroyArray, 84, 86
- destroyMat, 84, 86
- DlsMat, 44, 45, 80

- eh_data, 43
- error message
 - user-defined handler, 26
- error messages, 24
 - redirecting, 24
 - user-defined handler, 42
- ETA_CONST, 60

- ETA_FORM, 60
- ETA_PARAMS, 60

- FKBJAC, 57
- FKCOMMFN, 63
- FKDJAC, 56
- FKFUN, 55
- FKINBAND, 56
- FKINBANDSETJAC, 57
- FKINBBD interface module
 - interface to the KINBBDPRE module, 60
- FKINBBDINIT, 62
- FKINBBDOPT, 62
- FKINDENSE, 56
- FKINDENSESETJAC, 56
- FKINFREE, 59
- FKINJTIMES, 58, 63
- FKINMALLOC, 55
- FKINSETIIN, 59
- FKINSETRIN, 59
- FKINSOL, 59
- FKINSOL interface module
 - interface to the KINBBDPRE module, 63
 - optional input and output, 59
 - usage, 54–59
 - user-callable functions, 53–54
 - user-supplied functions, 54
- FKINSPBCG, 58
- FKINSPGMR, 57
- FKINSPILSSETJAC, 58, 62
- FKINSPILSSETPREC, 58
- FKINSPTFQMR, 58
- FKLOCFN, 63
- FKPSET, 59
- FKPSOL, 59
- FNORM_TOL, 60
- FNVINITP, 55
- FNVINITS, 55

- generic linear solvers
 - BAND, 79
 - DENSE, 79
 - SPBCG, 88
 - SPGMR, 87

- SPTFQMR, 88
 - use in KINSOL, 13–14
- GMRES method, 87
- half-bandwidths, 21, 44–45, 49
- header files, 16, 49
- ih_data, 43
- Inexact Newton iteration
 - definition, 5
- info message
 - user-defined handler, 26
- info messages
 - redirecting, 26
- informational messages
 - user-defined handler, 43
- IOUT, 60, 61
- Jacobian approximation function
 - band
 - difference quotient, 34
 - use in FKINSOL, 57
 - user-supplied, 34, 44–45
 - dense
 - difference quotient, 33
 - use in FKINSOL, 56
 - user-supplied, 33, 43–44
 - Jacobian times vector
 - difference quotient, 34
 - use in FKINSOL, 58
 - user-supplied, 35, 45–46
- KIN_ETACHOICE1, 29
- KIN_ETACHOICE2, 29
- KIN_ETACONSTANT, 29
- KIN_FIRST_SYSFUNC_ERR, 24
- KIN_ILL_INPUT, 19, 23, 27–33
- KIN_INITIAL_GUESS_OK, 23
- KIN_LINESEARCH, 23
- KIN_LINESEARCH_BCFAIL, 24
- KIN_LINESEARCH_NONCONV, 23
- KIN_LINIT_FAIL, 24
- KIN_LINSOLV_NO_RECOVERY, 24
- KIN_LSETUP_FAIL, 24
- KIN_LSOLVE_FAIL, 24
- KIN_MAXITER_REACHED, 23
- KIN_MEM_FAIL, 19
- KIN_MEM_NULL, 19, 23, 26–33, 36–38
- KIN_MXNEWT_5X_EXCEEDED, 23
- KIN_NO_MALLOC, 23
- KIN_NONE, 23
- KIN_REPTD_SYSFUNC_ERR, 24
- KIN_STEP_LT_STPTOL, 23
- KIN_SUCCESS, 19, 23, 26–33, 36–38
- KIN_SYSFUNC_FAIL, 24
- KIN_WARNING, 43
- KINBAND linear solver
 - Jacobian approximation used by, 34
 - memory requirements, 38
 - NVECTOR compatibility, 21
 - optional input, 33–34
 - optional output, 38–39
 - selection of, 21
 - use in FKINSOL, 56
- KINBand, 18, 20, 21, 44
- KINBandSetJacFn, 34
- KINBBDPRE preconditioner
 - optional output, 50–51
 - usage, 49
 - user-callable functions, 49–50
 - user-supplied functions, 48–49
- KINBBDPrecGetNumGfnEvals, 51
- KINBBDPrecGetWorkSpace, 50
- KINBBDPrecInit, 50
- KINCreate, 19
- KINDENSE linear solver
 - Jacobian approximation used by, 33
 - memory requirements, 38
 - NVECTOR compatibility, 20
 - optional input, 33–34
 - optional output, 38–39
 - selection of, 20
 - use in FKINSOL, 56
- KINDense, 18, 20, 43
- KINDLS_ILL_INPUT, 20, 21
- KINDLS_LMEM_NULL, 33, 34, 38, 39
- KINDLS_MEM_FAIL, 20, 21
- KINDLS_MEM_NULL, 20, 21, 33, 34, 38, 39
- KINDLS_SUCCESS, 20, 21, 33, 34, 38, 39
- KINDlsBandJacFn, 44
- KINDlsDenseJacFn, 43
- KINDlsGetLastFlag, 39
- KINDlsGetNumFuncEvals, 39
- KINDlsGetNumJacEvals, 38
- KINDlsGetWorkSpace, 38
- KINDlsSetDenseJacFn, 33
- KINErrorHandlerFn, 42
- KINFree, 19
- KINGetFuncNorm, 37
- KINGetNumBacktrackOps, 37
- KINGetNumBetaCondFails, 37
- KINGetNumFuncEvals, 37
- KINGetNumNonlinSolvIters, 37
- KINGetStepLength, 38
- KINGetWorkSpace, 36
- KINInfoHandlerFn, 43
- KINInit, 19
- KINLapackBand, 18, 20, 21
- KINLapackDense, 18, 20, 21

- KINSetConstraints, 32
- KINSetErrFile, 24
- KINSetErrHandlerFn, 26
- KINSetEtaConstValue, 29
- KINSetEtaForm, 29
- KINSetEtaParams, 30
- KINSetFuncNormTol, 32
- KINSetInfoFile, 26
- KINSetInfoHandlerFn, 26
- KINSetMaxBetaFails, 31
- KINSetMaxNewtonStep, 31
- KINSetMaxSetupCalls, 28
- KINSetMaxSubSetupCalls, 29
- KINSetNoInitSetup, 28
- KINSetNoMinEps, 31
- KINSetNoResMon, 28
- KINSetNumMaxIters, 27
- KINSetPrintLevel, 27
- KINSetRelErrFunc, 31
- KINSetResMonConstValue, 30
- KINSetResMonParams, 30
- KINSetScaledStepTol, 32
- KINSetSysFunc, 33
- KINSetUserData, 27
- KINSOL
 - brief description of, 1
 - motivation for writing in C, 2
 - package structure, 11
 - relationship to NKSOL, 1
- KINSOL linear solvers
 - built on generic solvers, 20
 - header files, 16
 - implementation details, 13
 - KINBAND, 21
 - KINDENSE, 20
 - KINSPBCG, 22
 - KINSPGMR, 21
 - KINSPTFQMR, 22
 - list of, 11–13
 - NVECTOR compatibility, 15
 - selecting one, 20
- KINSol, 18, 23
- kinsol.h, 16
- kinsol_band.h, 16
- kinsol_dense.h, 16
- kinsol_lapack.h, 17
- kinsol_spgmrs.h, 17
- kinsol_spgmr.h, 17
- kinsol_sptfqmr.h, 17
- KINSOLKINSOL linear solvers
 - selecting one, 20
- KINSPBCG linear solver
 - Jacobian approximation used by, 34
 - memory requirements, 39
 - optional input, 34–35
 - optional output, 39–42
 - preconditioner setup function, 34, 47
 - preconditioner solve function, 34, 46
 - selection of, 22
 - use in FKINSOL, 58
- KINSPbcg, 18, 20, 22
- KINSPGMR linear solver
 - Jacobian approximation used by, 34
 - memory requirements, 39
 - optional input, 34–35
 - optional output, 39–42
 - preconditioner setup function, 34, 47
 - preconditioner solve function, 34, 46
 - selection of, 21
 - use in FKINSOL, 57
- KINSPgmr, 18, 20, 22
- KINSPILS_ILL_INPUT, 22, 35, 50
- KINSPILS_LMEM_NULL, 35, 40–42, 50
- KINSPILS_MEM_FAIL, 22, 50
- KINSPILS_MEM_NULL, 22, 35, 40, 41
- KINSPILS_PMEM_NULL, 51
- KINSPILS_SUCCESS, 22, 35, 39–41
- KINSpilsGetLastFlag, 41
- KINSpilsGetNumConvFails, 40
- KINSpilsGetNumFuncEvals, 41
- KINSpilsGetNumJtimesEvals, 41
- KINSpilsGetNumLinIters, 40
- KINSpilsGetNumPrecEvals, 40
- KINSpilsGetNumPrecSolves, 41
- KINSpilsGetWorkSpace, 39
- KINSpilsJacTimesVecFn, 45
- KINSpilsPrecSetupFn, 47
- KINSpilsPrecSolveFn, 46
- KINSpilsSetJacTimesFn, 35
- KINSpilsSetMaxRestarts, 35
- KINSpilsSetPreconditioner, 34
- KINSPLIS_LMEM_NULL, 40
- KINSPTFQMR linear solver
 - Jacobian approximation used by, 34
 - memory requirements, 39
 - optional input, 34–35
 - optional output, 39–42
 - preconditioner setup function, 34, 47
 - preconditioner solve function, 34, 46
 - selection of, 22
 - use in FKINSOL, 58
- KINSptfqmr, 18, 20, 22
- KINSysFn, 19, 42
- linit, 75
- MAX_NITER, 60
- MAX_SETUPS, 60

- MAX_SP_SETUPS, 60
- MAX_STEP, 60
- memory requirements
 - KINBAND linear solver, 38
 - KINBBDPRE preconditioner, 51
 - KINDENSE linear solver, 38
 - KINSOL solver, 36
 - KINSPGMR linear solver, 39
- Modified Newton iteration
 - definition, 5
- MPI, 4

- N_VCloneEmptyVectorArray, 66
- N_VCloneEmptyVectorArray_Parallel, 73
- N_VCloneEmptyVectorArray_Serial, 70
- N_VCloneVectorArray, 66
- N_VCloneVectorArray_Parallel, 72
- N_VCloneVectorArray_Serial, 70
- N_VDestroyVectorArray, 66
- N_VDestroyVectorArray_Parallel, 73
- N_VDestroyVectorArray_Serial, 70
- N_Vector, 16, 65
- N_VMake_Parallel, 72
- N_VMake_Serial, 70
- N_VNew_Parallel, 72
- N_VNew_Serial, 70
- N_VNewEmpty_Parallel, 72
- N_VNewEmpty_Serial, 70
- N_VPrint_Parallel, 73
- N_VPrint_Serial, 71
- newBandMat, 86
- newDenseMat, 84
- newIntArray, 84, 86
- newLintArray, 84, 86
- newRealArray, 84, 86
- NO_INIT_SETUP, 60
- NO_MIN_EPS, 60
- NO_RES_MON, 60
- nonlinear system
 - definition, 5
- NV_COMM_P, 72
- NV_CONTENT_P, 71
- NV_CONTENT_S, 69
- NV_DATA_P, 71
- NV_DATA_S, 69
- NV_GLOBLENGTH_P, 71
- NV_Ith_P, 72
- NV_Ith_S, 70
- NV_LENGTH_S, 69
- NV_LOCLENGTH_P, 71
- NV_OWN_DATA_P, 71
- NV_OWN_DATA_S, 69
- NVECTOR module, 65
- nvector_parallel.h, 16
- nvector_serial.h, 16

- optional input
 - band linear solver, 33–34
 - dense linear solver, 33–34
 - iterative linear solver, 34–35
 - solver, 24–33
- optional output
 - band linear solver, 38–39
 - band-block-diagonal preconditioner, 50–51
 - dense linear solver, 38–39
 - iterative linear solver, 39–42
 - solver, 36–38

- portability, 16
 - Fortran, 54
- preconditioning
 - setup and solve phases, 13
 - user-supplied, 34–35, 46, 47
- PRNT_LEVEL, 60
- problem-defining function, 42

- RCONST, 16
- realtype, 16
- RERR_FUNC, 60
- RMON_CONST, 60
- RMON_PARAMS, 60
- ROUT, 60, 61

- SMALL_REAL, 16
- SPBCG generic linear solver
 - description of, 88
 - functions, 88
- SPGMR generic linear solver
 - description of, 87
 - functions, 88
 - support functions, 88
- SPTFQMR generic linear solver
 - description of, 88
 - functions, 88
- SSTEP_TOL, 60
- sundials_nvector.h, 16
- sundials_types.h, 16

- TFQMR method, 88

- UNIT_ROUNDOFF, 16
- User main program
 - FKINBBD usage, 62
 - FKINSOL usage, 54
 - KINBBDPRE usage, 49
 - KINSOL usage, 17
- user_data, 42, 48