

Example Programs for CVODE v2.9.0

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

September 26, 2016



UCRL-SM-208110

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

1	Introduction	1
2	Serial example problems	4
2.1	A dense example: cvRoberts_dns	4
2.2	A banded example: cvAdvDiff_bnd	6
2.3	A Krylov example: cvDiurnal_kry	9
3	Parallel example problems	13
3.1	A nonstiff example: cvAdvDiff_non_p	13
3.2	A user preconditioner example: cvDiurnal_kry_p	14
3.3	A CVBBDPRE preconditioner example: cvDiurnal_kry_bbd_p	16
4	hypre example problems	20
4.1	A nonstiff example: cvAdvDiff_non_ph	20
5	Fortran example problems	21
5.1	A serial example: fcvDiurnal_kry	21
5.2	A parallel example: fcvDiag_kry_bbd_p	22
6	Parallel tests	25
	References	27

1 Introduction

This report is intended to serve as a companion document to the User Documentation of CVODE [2]. It provides details, with listings, on the example programs supplied with the CVODE distribution package.

The CVODE distribution contains examples of four types: serial C examples, parallel C examples, and serial and parallel FORTRAN examples. With the exception of "demo"-type example files, the names of all the examples distributed with SUNDIALS are of the form `[slv][PbName]_[ls]_[prec]_[p]`, where

`[slv]` identifies the solver (for CVODE examples this is `cv`, while for FCVODE examples, this is `fcv`);

`[PbName]` identifies the problem;

`[ls]` identifies the linear solver module used (for examples using functional iteration for the nonlinear system solver, `non` specifies that no linear solver was used);

`[prec]` indicates the CVODE preconditioner module used, `bp` for CVBANDPRE or `bbd` for CVBB-DPRE (only if applicable, for examples using a Krylov linear solver);

`[p]` indicates an example using the parallel vector module `NVECTOR_PARALLEL`.

The following lists summarize all examples distributed with CVODE.

Supplied in the `srcdir/examples/cvode/serial` directory are the following ten serial examples (using the `NVECTOR_SERIAL` module):

- `cvRoberts_dns` solves a chemical kinetics problem consisting of three rate equations. This program solves the problem with the BDF method and Newton iteration, with the `CVDENSE` linear solver and a user-supplied Jacobian routine. It also uses the rootfinding feature of CVODE.
- `cvRoberts_dnsL` is the same as `cvRoberts_dns` but uses the Lapack implementation of `CVDENSE`.
- `cvRoberts_dns_uw` is the same as `cvRoberts_dns` but demonstrates the user-supplied error weight function feature of CVODE.
- `cvRoberts_klu` is the same as `cvRoberts_dns` but uses the `KLU` sparse direct linear solver.
- `cvRoberts_sps` is the same as `cvRoberts_dns` but uses the `SuperLUMT` sparse direct linear solver (with one thread).
- `cvAdvDiff_bnd` solves the semi-discrete form of an advection-diffusion equation in 2-D. This program solves the problem with the BDF method and Newton iteration, with the `CVBAND` linear solver and a user-supplied Jacobian routine.
- `cvAdvDiff_bndL` is the same as `cvAdvDiff_bnd` but uses the Lapack implementation of `CVBAND`.

- `cvDiurnal_kry` solves the semi-discrete form of a two-species diurnal kinetics advection-diffusion PDE system in 2-D.
The problem is solved with the BDF/GMRES method (i.e. using the CVSPGMR linear solver) and the block-diagonal part of the Newton matrix as a left preconditioner. A copy of the block-diagonal part of the Jacobian is saved and conditionally reused within the preconditioner setup routine.
- `cvDiurnal_kry_bp` solves the same problem as `cvDiurnal_kry`, with the BDF/GMRES method and a banded preconditioner, generated by difference quotients, using the module CVBANDPRE.
The problem is solved twice: with preconditioning on the left, then on the right.
- `cvDirectDemo_ls` is a demonstration program for CVODE with direct linear solvers. Two separate problems are solved using both the Adams and BDF linear multistep methods in combination with functional and Newton iterations.
The first problem is the Van der Pol oscillator for which the Newton iteration cases use the following types of Jacobian approximations: (1) dense, user-supplied, (2) dense, difference-quotient approximation, (3) diagonal approximation. The second problem is a linear ODE with a banded lower triangular matrix derived from a 2-D advection PDE. In this case, the Newton iteration cases use the following types of Jacobian approximation: (1) banded, user-supplied, (2) banded, difference-quotient approximation, (3) diagonal approximation.
- `cvKrylovDemo_ls` solves the same problem as `cvDiurnal_kry`, with the BDF method, but with three Krylov linear solvers: CVSPGMR, CVSPBCG, and CVSPTFQMR.
- `cvKrylovDemo_prec` is a demonstration program with the GMRES linear solver. This program solves a stiff ODE system that arises from a system of partial differential equations. The PDE system is a six-species food web population model, with predator-prey interaction and diffusion on the unit square in two dimensions.
The ODE system is solved using Newton iteration and the CVSPGMR linear solver (scaled preconditioned GMRES).
The preconditioner matrix used is the product of two matrices: (1) a matrix, only defined implicitly, based on a fixed number of Gauss-Seidel iterations using the diffusion terms only; and (2) a block-diagonal matrix based on the partial derivatives of the interaction terms only, using block-grouping.
Four different runs are made for this problem. The product preconditioner is applied on the left and on the right. In each case, both the modified and classical Gram-Schmidt options are tested.

Supplied in the `srcdir/examples/cvode/parallel` directory are the following four parallel examples (using the NVECTOR_PARALLEL module):

- `cvAdvDiff_non_p` solves the semi-discrete form of a 1-D advection-diffusion equation. This program solves the problem with the option for nonstiff systems, i.e. Adams method and functional iteration.
- `cvAdvDiff_diag_p` solves the same problem as `cvAdvDiff_non_p`, with the Adams method, but with Newton iteration and the CVDiag linear solver.

- `cvDiurnal_kry_p` is a parallel implementation of `cvDiurnal_kry`.
- `cvDiurnal_kry_bbd_p` solves the same problem as `cvDiurnal_kry_p`, with BDF and the GMRES linear solver, using a block-diagonal matrix with banded blocks as a preconditioner, generated by difference quotients, using the module `CVBBDPRE`.

Within the `FCVODE` module, in the directories `srcdir/examples/cvode/fcmix_serial` and `srcdir/examples/cvode/fcmix_parallel`, are the following examples for the FORTRAN-C interface. The first five of these are serial, while the last three are parallel.

- `fcvRoberts_dns` is a serial chemical kinetics example (BDF/CVDENSE) with rootfinding.
- `fcvRoberts_dnsL` is the same as `fcvRoberts_dns` but uses the Lapack implementation of CVDENSE.
- `fcvAdvDiff_bnd` is a serial advection-diffusion example (BDF/CVBAND).
- `fcvDiurnal_kry` is a serial kinetics-transport example (BDF/CVSPGMR).
- `fcvDiurnal_kry_bp` is the `fcvDiurnal_kry` example with `FCVBP`.
- `fcvDiag_non_p` is a nonstiff parallel diagonal ODE example (ADAMS/FUNCTIONAL).
- `fcvDiag_kry_p` is a stiff parallel diagonal ODE example (BDF/CVSPGMR).
- `fcvDiag_kry_bbd_p` is the same as the `fcvDiag_kry_p` example but using the `FCVBBD` module.

In the following sections, we give detailed descriptions of some (but not all) of these examples. We also give our output files for each of these examples, but users should be cautioned that their results may differ slightly from these. Differences in solution values may differ within the tolerances, and differences in cumulative counters, such as numbers of steps or Newton iterations, may differ from one machine environment to another by as much as 10% to 20%.

The final section of this report describes a set of tests done with the parallel version of `CVODE`, using a problem based on the `cvDiurnal_kry/cvDiurnal_kry_p` example.

In the descriptions below, we make frequent references to the `CVODE` User Document [2]. All citations to specific sections (e.g. §4.2) are references to parts of that User Document, unless explicitly stated otherwise.

Note. The examples in the `CVODE` distribution are written in such a way as to compile and run for any combination of configuration options during the installation of `SUNDIALS` (see Appendix A in the User Guide). As a consequence, they contain portions of code that will not be typically present in a user program. For example, all C example programs make use of the variables `SUNDIALS_EXTENDED_PRECISION` and `SUNDIALS_DOUBLE_PRECISION` to test if the solver libraries were built in extended or double precision, and use the appropriate conversion specifiers in `printf` functions. Similarly, the FORTRAN examples in `FCVODE` are automatically pre-processed to generate source code that corresponds to the manner in which the `CVODE` libraries were built (see 5 in this document for more details).

2 Serial example problems

2.1 A dense example: `cvRoberts_dns`

As an initial illustration of the use of the `CVODE` package for the integration of IVP ODEs, we give a sample program called `cvRoberts_dns.c`. It uses the `CVODE` dense linear solver module `CVDENSE` and the `NVECTOR_SERIAL` module (which provides a serial implementation of `NVECTOR`) in the solution of a 3-species chemical kinetics problem.

The problem consists of the following three rate equations:

$$\begin{aligned} \dot{y}_1 &= -0.04 \cdot y_1 + 10^4 \cdot y_2 \cdot y_3 \\ \dot{y}_2 &= 0.04 \cdot y_1 - 10^4 \cdot y_2 \cdot y_3 - 3 \cdot 10^7 \cdot y_2^2 \\ \dot{y}_3 &= 3 \cdot 10^7 \cdot y_2^2 \end{aligned} \tag{1}$$

on the interval $t \in [0, 4 \cdot 10^{10}]$, with initial conditions $y_1(0) = 1.0$, $y_2(0) = y_3(0) = 0.0$. While integrating the system, we also use the rootfinding feature to find the points at which $y_1 = 10^{-4}$ or at which $y_3 = 0.01$.

For the source we give a rather detailed explanation of the parts of the program and their interaction with `CVODE`.

Following the initial comment block, this program has a number of `#include` lines, which allow access to useful items in `CVODE` header files. The `sundials_types.h` file provides the definition of the type `realtype` (see §4.2 for details). For now, it suffices to read `realtype` as `double`. The `cvode.h` file provides prototypes for the `CVODE` functions to be called (excluding the linear solver selection function), and also a number of constants that are to be used in setting input arguments and testing the return value of `CVode`. The `cvode_dense.h` file provides the prototype for the `CVDense` function. The `nvector_serial.h` file is the header file for the serial implementation of the `NVECTOR` module and includes definitions of the `N_Vector` type, a macro to access vector components, and prototypes for the serial implementation specific machine environment memory allocation and freeing functions. The `sundials_dense.h` file provides the definition of the dense matrix type `DlsMat` (`type=1`) and a macro for accessing matrix elements. We have explicitly included `sundials_dense.h`, but this is not necessary because it is included by `cvode_dense.h`.

This program includes two user-defined accessor macros, `Ith` and `IJth`, that are useful in writing the problem functions in a form closely matching the mathematical description of the ODE system, i.e. with components numbered from 1 instead of from 0. The `Ith` macro is used to access components of a vector of type `N_Vector` with a serial implementation. It is defined using the `NVECTOR_SERIAL` accessor macro `NV_Ith_S` which numbers components starting with 0. The `IJth` macro is used to access elements of a dense matrix of type `DlsMat`. It is defined using the `DENSE` accessor macro `DENSE_ELEM` which numbers matrix rows and columns starting with 0. The macro `NV_Ith_S` is fully described in §6.1. The macro `DENSE_ELEM` is fully described in §4.6.5.

Next, the program includes some problem-specific constants, which are isolated to this early location to make it easy to change them as needed. The program prologue ends with prototypes of four private helper functions and the three user-supplied functions that are called by `CVODE`.

The `main` program begins with some dimensions and type declarations, including use of the type `N_Vector`. The next several lines allocate memory for the `y` and `abstol` vectors using `N_VNew_Serial` with a length argument of `NEQ` ($= 3$). The lines following that load

the initial values of the dependent variable vector into `y` and the absolute tolerances into `abstol` using the `Ith` macro.

The calls to `NVNewSerial`, and also later calls to `CVode***` functions, make use of a private function, `check_flag`, which examines the return value and prints a message if there was a failure. The `check_flag` function was written to be used for any serial SUNDIALS application.

The call to `CVodeCreate` creates the CVODE solver memory block, specifying the `CV_BDF` integration method with `CV_NEWTON` iteration. Its return value is a pointer to that memory block for this problem. In the case of failure, the return value is `NULL`. This pointer must be passed in the remaining calls to CVODE functions.

The call to `CVodeInit` allocates and initializes the solver memory block. Its arguments include the name of the C function `f` defining the right-hand side function $f(t, y)$, and the initial values of t and y . The call to `CVodeSVtolerances` specifies a vector of absolute tolerances, and includes the value of the relative tolerance `reltol` and the absolute tolerance vector `abstol`. See §4.5.1 and §4.5.2 for full details of these calls.

The call to `CVodeRootInit` specifies that a rootfinding problem is to be solved along with the integration of the ODE system, that the root functions are specified in the function `g`, and that there are two such functions. Specifically, they are set to $y_1 - 0.0001$ and $y_3 - 0.01$, respectively. See §4.5.4 for a detailed description of this call.

The calls to `CVDense` (see §4.5.3) and `CVDlsSetDenseJacFn` (see §4.5.6) specify the CV-DENSE linear solver with an analytic Jacobian supplied by the user-supplied function `Jac`.

The actual solution of the ODE initial value problem is accomplished in the loop over values of the output time `tout`. In each pass of the loop, the program calls `CVode` in the `CV_NORMAL` mode, meaning that the integrator is to take steps until it overshoots `tout` and then interpolate to $t = \text{tout}$, putting the computed value of $y(\text{tout})$ into `y`, with `t = tout`. The return value in this case is `CV_SUCCESS`. However, if `CVode` finds a root before reaching the next value of `tout`, it returns `CV_ROOT_RETURN` and stores the root location in `t` and the solution there in `y`. In either case, the program prints `t` and `y`. In the case of a root, it calls `CVodeGetRootInfo` to get a length-2 array `rootsfound` of bits showing which root function was found to have a root. If `CVode` returned any negative value (indicating a failure), the program breaks out of the loop. In the case of a `CV_SUCCESS` return, the value of `tout` is advanced (multiplied by 10) and a counter (`iout`) is advanced, so that the loop can be ended when that counter reaches the preset number of output times, `NOUT = 12`. See §4.5.5 for full details of the call to `CVode`.

Finally, the main program calls `PrintFinalStats` to get and print all of the relevant statistical quantities. It then calls `NV_Destroy` to free the vectors `y` and `abstol`, and `CVodeFree` to free the CVODE memory block.

The function `PrintFinalStats` used here is actually suitable for general use in applications of CVODE to any problem with a dense Jacobian. It calls various `CVodeGet***` and `CVDenseGet***` functions to obtain the relevant counters, and then prints them. Specifically, these are: the cumulative number of steps (`nst`), the number of `f` evaluations (`nfe`) (excluding those for difference-quotient Jacobian evaluations), the number of matrix factorizations (`nsetups`), the number of `f` evaluations for Jacobian evaluations (`nfeLS = 0` here), the number of Jacobian evaluations (`nje`), the number of nonlinear (Newton) iterations (`nni`), the number of nonlinear convergence failures (`ncfn`), the number of local error test failures (`netf`), and the number of `g` (root function) evaluations (`nge`). These optional outputs are described in §4.5.8.

The function `f` is a straightforward expression of the ODEs. It uses the user-defined

macro `Ith` to extract the components of `y` and to load the components of `ydot`. See §4.6.1 for a detailed specification of `f`.

Similarly, the function `g` defines the two functions, g_0 and g_1 , whose roots are to be found. See §4.6.4 for a detailed description of the `g` function.

The function `Jac` sets the nonzero elements of the Jacobian as a dense matrix. (Zero elements need not be set because `J` is preset to zero.) It uses the user-defined macro `IJth` to reference the elements of a dense matrix of type `DlsMat`. Here the problem size is small, so we need not worry about the inefficiency of using `NV_Ith_S` and `DENSE_ELEM` to access `N_Vector` and `DlsMat` elements. Note that in this example, `Jac` only accesses the `y` and `J` arguments. See §4.6.5 for a detailed description of the dense `Jac` function.

The output generated by `cvRoberts_dns` is shown below. It shows the output values at the 12 preset values of `tout`. It also shows the two root locations found, first at a root of g_1 , and then at a root of g_0 .

```

cvRoberts_dns sample output

3-species kinetics problem

At t = 2.6391e-01      y =  9.899653e-01      3.470564e-05      1.000000e-02
  rootsfound[] =    0      1
At t = 4.0000e-01      y =  9.851641e-01      3.386242e-05      1.480205e-02
At t = 4.0000e+00      y =  9.055097e-01      2.240338e-05      9.446793e-02
At t = 4.0000e+01      y =  7.158009e-01      9.185098e-06      2.841900e-01
At t = 4.0000e+02      y =  4.505440e-01      3.223217e-06      5.494528e-01
At t = 4.0000e+03      y =  1.831964e-01      8.942051e-07      8.168027e-01
At t = 4.0000e+04      y =  3.898104e-02      1.621656e-07      9.610188e-01
At t = 4.0000e+05      y =  4.938672e-03      1.985172e-08      9.950613e-01
At t = 4.0000e+06      y =  5.166093e-04      2.067499e-09      9.994834e-01
At t = 2.0800e+07      y =  1.000000e-04      4.000395e-10      9.999000e-01
  rootsfound[] =   -1      0
At t = 4.0000e+07      y =  5.206409e-05      2.082671e-10      9.999479e-01
At t = 4.0000e+08      y =  5.211241e-06      2.084507e-11      9.999948e-01
At t = 4.0000e+09      y =  5.200520e-07      2.080209e-12      9.999995e-01
At t = 4.0000e+10      y =  5.699485e-08      2.279794e-13      9.999999e-01

Final Statistics:
nst = 579      nfe = 817      nsetups = 118      nfeLS = 0      nje = 12
nni = 813      ncnfn = 0      netf = 31      nge = 615

```

2.2 A banded example: `cvAdvDiff_bnd`

The example program `cvAdvDiff_bnd.c` solves the semi-discretized form of the 2-D advection-diffusion equation

$$\partial v / \partial t = \partial^2 v / \partial x^2 + .5 \partial v / \partial x + \partial^2 v / \partial y^2 \quad (2)$$

on a rectangle, with zero Dirichlet boundary conditions. The PDE is discretized with standard central finite differences on a $(MX+2) \times (MY+2)$ mesh, giving an ODE system of size $MX*MY$. The discrete value v_{ij} approximates v at $x = i\Delta x$, $y = j\Delta y$. The ODEs are

$$\frac{dv_{ij}}{dt} = f_{ij} = \frac{v_{i-1,j} - 2v_{ij} + v_{i+1,j}}{(\Delta x)^2} + .5 \frac{v_{i+1,j} - v_{i-1,j}}{2\Delta x} + \frac{v_{i,j-1} - 2v_{ij} + v_{i,j+1}}{(\Delta y)^2}, \quad (3)$$

where $1 \leq i \leq \text{MX}$ and $1 \leq j \leq \text{MY}$. The boundary conditions are imposed by taking $v_{ij} = 0$ above if $i = 0$ or $\text{MX}+1$, or if $j = 0$ or $\text{MY}+1$. If we set $u_{(j-1)+i*\text{MY}} = v_{ij}$, so that the ODE system is $\dot{u} = f(u)$, then the system Jacobian $J = \partial f / \partial u$ is a band matrix with upper and lower half-bandwidths both equal to MY . In the example, we take $\text{MX} = 10$ and $\text{MY} = 5$.

The `cvAdvDiff_bnd.c` program includes files `cvode_band.h` and `sundials_band.h` in order to use the CVBAND linear solver. The `cvode_band.h` file contains the prototype for the CVBand routine. The `sundials_band.h` file contains the definition for band matrix type `DlsMat` (`type=2`) and the `BAND_COL` and `BAND_COL_ELEM` macros for accessing matrix elements (see §8.1.4). We have explicitly included `sundials_band.h`, but this is not necessary because it is included by `cvode_band.h`. The file `nvector_serial.h` is included for the definition of the serial `N_Vector` type.

The include lines at the top of the file are followed by definitions of problem constants which include the x and y mesh dimensions, MX and MY , the number of equations NEQ , the scalar absolute tolerance ATOL , the initial time T0 , and the initial output time T1 .

Spatial discretization of the PDE naturally produces an ODE system in which equations are numbered by mesh coordinates (i, j) . The user-defined macro `IJth` isolates the translation for the mathematical two-dimensional index to the one-dimensional `N_Vector` index and allows the user to write clean, readable code to access components of the dependent variable. The `NV_DATA_S` macro returns the component array for a given `N_Vector`, and this array is passed to `IJth` in order to do the actual `N_Vector` access.

The type `UserData` is a pointer to a structure containing problem data used in the `f` and `Jac` functions. This structure is allocated and initialized at the beginning of `main`. The pointer to it, called `data`, is passed to `CVodeSetUserData`, and as a result it will be passed back to the `f` and `Jac` functions each time they are called. The use of the `data` pointer eliminates the need for global program data.

The `main` program is straightforward. The `CVodeCreate` call specifies the CV_BDF method with a CV_NEWTON iteration. Following the `CVodeInit` call, the call to `CVodeSStolerances` indicates scalar relative and absolute tolerances, and values `reltol` and `abstol` are passed. The call to `CVBand` (see §4.5.3) specifies the CVBAND linear solver, and specifies that both half-bandwidths of the Jacobian are equal to MY . The call to `CVDlsSetBandJacFn` (see §4.5.6) specifies that a user-supplied Jacobian function `Jac` is to be used.

The actual solution of the problem is performed by the call to `CVode` within the loop over the output times `tout`. The max-norm of the solution vector (from a call to `N_VMaxNorm`) and the cumulative number of time steps (from a call to `CVodeGetNumSteps`) are printed at each output time. Finally, the calls to `PrintFinalStats`, `N_VDestroy`, and `CVodeFree` print statistics and free problem memory.

Following the `main` program in the `cvAdvDiff_bnd.c` file are definitions of five functions: `f`, `Jac`, `SetIC`, `PrintHeader`, `PrintOutput`, `PrintFinalStats`, and `check_flag`. The last five functions are called only from within the `cvAdvDiff_bnd.c` file. The `SetIC` function sets the initial dependent variable vector; `PrintHeader` prints the heading of the output page; `PrintOutput` prints a line of solution output; `PrintFinalStats` gets and prints statistics at the end of the run; and `check_flag` aids in checking return values. The statistics printed include counters such as the total number of steps (`nst`), `f` evaluations (excluding those for Jaobian evaluations) (`nfe`), LU decompositions (`nsetups`), `f` evaluations for difference-quotient Jacobians (`nfeLS = 0` here), Jacobian evaluations (`nje`), and nonlinear iterations (`nni`). These optional outputs are described in §4.5.8. Note that `PrintFinalStats` is suitable for general use in applications of CVODE to any problem with a banded Jacobian.

The `f` function implements the central difference approximation (3) with u identically zero

on the boundary. The constant coefficients $(\Delta x)^{-2}$, $.5(2\Delta x)^{-1}$, and $(\Delta y)^{-2}$ are computed only once at the beginning of `main`, and stored in the locations `data->hdcoef`, `data->hacoef`, and `data->vdcoef`, respectively. When `f` receives the `data` pointer (renamed `user_data` here), it pulls out these values from storage in the local variables `hordc`, `horac`, and `verdc`. It then uses these to construct the diffusion and advection terms, which are combined to form `udot`. Note the extra lines setting out-of-bounds values of `u` to zero.

The `Jac` function is an expression of the derivatives

$$\begin{aligned} \partial f_{ij} / \partial v_{ij} &= -2[(\Delta x)^{-2} + (\Delta y)^{-2}] \\ \partial f_{ij} / \partial v_{i\pm 1,j} &= (\Delta x)^{-2} \pm .5(2\Delta x)^{-1}, \quad \partial f_{ij} / \partial v_{i,j\pm 1} = (\Delta y)^{-2}. \end{aligned}$$

This function loads the Jacobian by columns, and like `f` it makes use of the preset coefficients in `data`. It loops over the mesh points (i,j) . For each such mesh point, the one-dimensional index $k = j-1 + (i-1)*MY$ is computed and the k th column of the Jacobian matrix J is set. The row index k' of each component $f_{i',j'}$ that depends on $v_{i,j}$ must be identified in order to load the corresponding element. The elements are loaded with the `BAND_COL_ELEM` macro. Note that the formula for the global index k implies that decreasing (increasing) i by 1 corresponds to decreasing (increasing) k by `MY`, while decreasing (increasing) j by 1 corresponds to decreasing (increasing) k by 1. These statements are reflected in the arguments to `BAND_COL_ELEM`. The first argument passed to the `BAND_COL_ELEM` macro is a pointer to the diagonal element in the column to be accessed. This pointer is obtained via a call to the `BAND_COL` macro and is stored in `kthCol` in the `Jac` function. When setting the components of J we must be careful not to index out of bounds. The guards `(i != 1)` etc. in front of the calls to `BAND_COL_ELEM` prevent illegal indexing. See §4.6.6 for a detailed description of the banded `Jac` function.

The output generated by `cvAdvDiff_bnd` is shown below.

```

----- cvAdvDiff_bnd sample output -----
2-D Advection-Diffusion Equation
Mesh dimensions = 10 X 5
Total system size = 50
Tolerance parameters: reltol = 0    abstol = 1e-05

At t = 0      max.norm(u) = 8.954716e+01
At t = 0.10   max.norm(u) = 4.132889e+00    nst = 85
At t = 0.20   max.norm(u) = 1.039294e+00    nst = 103
At t = 0.30   max.norm(u) = 2.979829e-01    nst = 113
At t = 0.40   max.norm(u) = 8.765774e-02    nst = 120
At t = 0.50   max.norm(u) = 2.625637e-02    nst = 126
At t = 0.60   max.norm(u) = 7.830425e-03    nst = 130
At t = 0.70   max.norm(u) = 2.329387e-03    nst = 134
At t = 0.80   max.norm(u) = 6.953434e-04    nst = 137
At t = 0.90   max.norm(u) = 2.115983e-04    nst = 140
At t = 1.00   max.norm(u) = 6.556853e-05    nst = 142

Final Statistics:
nst = 142    nfe = 174    nsetups = 23    nfeLS = 0    nje = 3
nni = 170    ncfm = 0    netf = 3

```

2.3 A Krylov example: cvDiurnal kry

We give here an example that illustrates the use of CVODE with the Krylov method SPGMR, in the CVSPGMR module, as the linear system solver.

This program solves the semi-discretized form of a pair of kinetics-advection-diffusion partial differential equations, which represent a simplified model for the transport, production, and loss of ozone and the oxygen singlet in the upper atmosphere. The problem includes non-linear diurnal kinetics, horizontal advection and diffusion, and nonuniform vertical diffusion. The PDEs can be written as

$$\frac{\partial c^i}{\partial t} = K_h \frac{\partial^2 c^i}{\partial x^2} + V \frac{\partial c^i}{\partial x} + \frac{\partial}{\partial y} K_v(y) \frac{\partial c^i}{\partial y} + R^i(c^1, c^2, t) \quad (i = 1, 2), \quad (4)$$

where the superscripts i are used to distinguish the two chemical species, and where the reaction terms are given by

$$\begin{aligned} R^1(c^1, c^2, t) &= -q_1 c^1 c^3 - q_2 c^1 c^2 + 2q_3(t) c^3 + q_4(t) c^2, \\ R^2(c^1, c^2, t) &= q_1 c^1 c^3 - q_2 c^1 c^2 - q_4(t) c^2. \end{aligned} \quad (5)$$

The spatial domain is $0 \leq x \leq 20$, $30 \leq y \leq 50$ (in *km*). The various constants and parameters are: $K_h = 4.0 \cdot 10^{-6}$, $V = 10^{-3}$, $K_v = 10^{-8} \exp(y/5)$, $q_1 = 1.63 \cdot 10^{-16}$, $q_2 = 4.66 \cdot 10^{-16}$, $c^3 = 3.7 \cdot 10^{16}$, and the diurnal rate constants are defined as:

$$q_i(t) = \begin{cases} \exp[-a_i / \sin \omega t], & \text{for } \sin \omega t > 0 \\ 0, & \text{for } \sin \omega t \leq 0 \end{cases} \quad (i = 3, 4),$$

where $\omega = \pi/43200$, $a_3 = 22.62$, $a_4 = 7.601$. The time interval of integration is $[0, 86400]$, representing 24 hours measured in seconds.

Homogeneous Neumann boundary conditions are imposed on each boundary, and the initial conditions are

$$\begin{aligned} c^1(x, y, 0) &= 10^6 \alpha(x) \beta(y), \quad c^2(x, y, 0) = 10^{12} \alpha(x) \beta(y), \\ \alpha(x) &= 1 - (0.1x - 1)^2 + (0.1x - 1)^4 / 2, \\ \beta(y) &= 1 - (0.1y - 4)^2 + (0.1y - 4)^4 / 2. \end{aligned} \quad (6)$$

For this example, the equations (4) are discretized spatially with standard central finite differences on a 10×10 mesh, giving an ODE system of size 200.

Among the initial `#include` lines in this case are lines to include `cvode_spgmr.h` and `sundials_math.h`. The first contains constants and function prototypes associated with the SPGMR method, including the values of the `pretype` argument to `CVSpGmr`. The inclusion of `sundials_math.h` is done to access the `SUNSQR` macro for the square of a `realtyp` number.

The main program calls `CVodeCreate` specifying the `CV_BDF` method and `CV_NEWTON` iteration, and then calls `CVodeInit`, and `CVodeSetSStolerances` specifies the scalar tolerances. It calls `CVSpGmr` (see §4.5.3) to specify the CVSPGMR linear solver with left preconditioning, and the default value (indicated by a zero argument) for `maxl`. The call to `CVSpilsSetJacTimesVecFn` specifies a user-supplied function for Jacobian-vector products. Next, user-supplied preconditioner setup and solve functions, `Precond` and `PSolve`, are specified. See §4.5.6 for details on the `CVSpilsSetPreconditioner` function.

For a sequence of `tout` values, `CVode` is called in the `CV_NORMAL` mode, sampled output is printed, and the return value is tested for error conditions. After that, `PrintFinalStats`

is called to get and print final statistics, and memory is freed by calls to `N_VDestroy`, `FreeUserData`, and `CVodeFree`. The printed statistics include various counters, such as the total numbers of steps (`nst`), of `f` evaluations (excluding those for Jv product evaluations) (`nfe`), of `f` evaluations for Jv evaluations (`nfeLS`), of nonlinear iterations (`nni`), of linear (Krylov) iterations (`nli`), of preconditioner setups (`nsetups`), of preconditioner evaluations (`npe`), and of preconditioner solves (`nps`), among others. Also printed are the lengths of the problem-dependent real and integer workspaces used by the main integrator `CVode`, denoted `lenrw` and `leniw`, and those used by `CVSPGMR`, denoted `lenrwLS` and `leniwLS`. All of these optional outputs are described in §4.5.8. The `PrintFinalStats` function is suitable for general use in applications of `CVODE` to any problem with the `SPGMR` linear solver.

Mathematically, the dependent variable has three dimensions: species number, x mesh point, and y mesh point. But in `NVECTOR_SERIAL`, a vector of type `N_Vector` works with a one-dimensional contiguous array of data components. The macro `IJKth` isolates the translation from three dimensions to one. Its use results in clearer code and makes it easy to change the underlying layout of the three-dimensional data. Here the problem size is 200, so we use the `NV_DATA_S` macro for efficient `N_Vector` access. The `NV_DATA_S` macro gives a pointer to the first component of an `N_Vector` which we pass to the `IJKth` macro to do an `N_Vector` access.

The preconditioner used here is the block-diagonal part of the true Newton matrix. It is generated and factored in the `Precond` routine (see §4.6.10) and backsolved in the `PSolve` routine (see §4.6.9). Its diagonal blocks are 2×2 matrices that include the interaction Jacobian elements and the diagonal contribution of the diffusion Jacobian elements. The block-diagonal part of the Jacobian itself, J_{bd} , is saved in separate storage each time it is generated, on calls to `Precond` with `jok == FALSE`. On calls with `jok == TRUE`, signifying that saved Jacobian data can be reused, the preconditioner $P = I - \gamma J_{bd}$ is formed from the saved matrix J_{bd} and factored. (A call to `Precond` with `jok == TRUE` can only occur after a prior call with `jok == FALSE`.) The `Precond` routine must also set the value of `jcur`, i.e. `*jcurPtr`, to `TRUE` when J_{bd} is re-evaluated, and `FALSE` otherwise, to inform `CVSPGMR` of the status of Jacobian data.

We need to take a brief detour to explain one last important aspect of this program. The generic `DENSE` solver contains two sets of functions: one for “large” matrices and one for “small” matrices. The large dense functions work with the type `DlsMat (type=1)`, while the small dense functions work with `realttype **` as the underlying dense matrix types. The `CVDENSE` linear solver uses the type `DlsMat` for the $N \times N$ dense Jacobian and Newton matrices, and calls the large matrix functions. But to avoid the extra layer of function calls, `cvDiurnal_kry.c` uses the small dense functions for all operations on the 2×2 preconditioner blocks. Thus it includes `sundials_smalldense.h` (by way of `sundials_dense.h`), and calls the small dense matrix functions `newDenseMat`, `newIntArray`, `denseCopy`, `denseScale`, `denseAddI`, `denseGETRF`, and `denseGETRS`. The macro `IJth` defined near the top of the file is used to access individual elements in each preconditioner block, numbered from 1. The small dense functions are available for `CVODE` user programs generally, and are documented in §8.1.3.

In addition to the functions called by `CVODE`, `cvDiurnal_kry.c` includes definitions of several private functions. These are: `AllocUserData` to allocate space for J_{bd} , P , and the pivot arrays; `InitUserData` to load problem constants in the `data` block; `FreeUserData` to free that block; `SetInitialProfiles` to load the initial values in `y`; `PrintOutput` to retrieve and print selected solution values and statistics; `PrintFinalStats` to print statistics; and `check_flag` to check return values for error conditions.

The output generated by `cvDiurnal_kry.c` is shown below. Note that the number of preconditioner evaluations, `npe`, is much smaller than the number of preconditioner setups, `nsetups`, as a result of the Jacobian re-use scheme.

```

cvDiurnal.dns sample output

2-species diurnal advection-diffusion problem

t = 7.20e+03  no. steps = 219  order = 5  stepsize = 1.59e+02
c1 (bot.left/middle/top rt.) = 1.047e+04  2.964e+04  1.119e+04
c2 (bot.left/middle/top rt.) = 2.527e+11  7.154e+11  2.700e+11

t = 1.44e+04  no. steps = 251  order = 5  stepsize = 3.77e+02
c1 (bot.left/middle/top rt.) = 6.659e+06  5.316e+06  7.301e+06
c2 (bot.left/middle/top rt.) = 2.582e+11  2.057e+11  2.833e+11

t = 2.16e+04  no. steps = 277  order = 5  stepsize = 2.75e+02
c1 (bot.left/middle/top rt.) = 2.665e+07  1.036e+07  2.931e+07
c2 (bot.left/middle/top rt.) = 2.993e+11  1.028e+11  3.313e+11

t = 2.88e+04  no. steps = 301  order = 5  stepsize = 3.87e+02
c1 (bot.left/middle/top rt.) = 8.702e+06  1.292e+07  9.650e+06
c2 (bot.left/middle/top rt.) = 3.380e+11  5.029e+11  3.751e+11

t = 3.60e+04  no. steps = 343  order = 3  stepsize = 2.34e+01
c1 (bot.left/middle/top rt.) = 1.404e+04  2.029e+04  1.561e+04
c2 (bot.left/middle/top rt.) = 3.387e+11  4.894e+11  3.765e+11

t = 4.32e+04  no. steps = 421  order = 4  stepsize = 5.26e+02
c1 (bot.left/middle/top rt.) = -4.385e-06  -1.528e-06  -4.905e-06
c2 (bot.left/middle/top rt.) = 3.382e+11  1.355e+11  3.804e+11

t = 5.04e+04  no. steps = 445  order = 3  stepsize = 1.98e+02
c1 (bot.left/middle/top rt.) = 4.461e-07  1.869e-07  4.842e-07
c2 (bot.left/middle/top rt.) = 3.358e+11  4.930e+11  3.864e+11

t = 5.76e+04  no. steps = 462  order = 5  stepsize = 2.35e+02
c1 (bot.left/middle/top rt.) = 3.204e-09  1.203e-09  3.555e-09
c2 (bot.left/middle/top rt.) = 3.320e+11  9.650e+11  3.909e+11

t = 6.48e+04  no. steps = 474  order = 5  stepsize = 6.02e+02
c1 (bot.left/middle/top rt.) = -1.066e-09  -3.409e-10  -1.206e-09
c2 (bot.left/middle/top rt.) = 3.313e+11  8.922e+11  3.963e+11

t = 7.20e+04  no. steps = 486  order = 5  stepsize = 6.02e+02
c1 (bot.left/middle/top rt.) = 2.614e-09  9.722e-10  2.904e-09
c2 (bot.left/middle/top rt.) = 3.330e+11  6.186e+11  4.039e+11

t = 7.92e+04  no. steps = 498  order = 5  stepsize = 6.02e+02
c1 (bot.left/middle/top rt.) = 4.649e-11  1.729e-11  5.161e-11
c2 (bot.left/middle/top rt.) = 3.334e+11  6.669e+11  4.120e+11

t = 8.64e+04  no. steps = 510  order = 5  stepsize = 6.02e+02
c1 (bot.left/middle/top rt.) = -8.856e-14  -3.348e-14  -9.785e-14
c2 (bot.left/middle/top rt.) = 3.352e+11  9.107e+11  4.163e+11

Final Statistics..

```

lenrw	=	2089	leniw	=	50
lenrwLS	=	2046	leniwLS	=	10
nst	=	510			
nfe	=	675	nfeLS	=	641
nni	=	671	nli	=	641
nsetups	=	94	netf	=	36
npe	=	9	nps	=	1243
ncfn	=	0	ncfl	=	0

3 Parallel example problems

3.1 A nonstiff example: `cvAdvDiff_non_p`

This problem begins with a simple diffusion-advection equation for $u = u(t, x)$

$$\frac{\partial u}{\partial t} = \frac{\partial^2 u}{\partial x^2} + 0.5 \frac{\partial u}{\partial x} \quad (7)$$

for $0 \leq t \leq 5$, $0 \leq x \leq 2$, and subject to homogeneous Dirichlet boundary conditions and initial values given by

$$\begin{aligned} u(t, 0) &= 0, & u(t, 2) &= 0, \\ u(0, x) &= x(2-x)e^{2x}. \end{aligned} \quad (8)$$

A system of MX ODEs is obtained by discretizing the x -axis with $\text{MX}+2$ grid points and replacing the first and second order spatial derivatives with their central difference approximations. Since the value of u is constant at the two endpoints, the semi-discrete equations for those points can be eliminated. With u_i as the approximation to $u(t, x_i)$, $x_i = i(\Delta x)$, and $\Delta x = 2/(\text{MX}+1)$, the resulting system of ODEs, $\dot{u} = f(t, u)$, can now be written:

$$\dot{u}_i = \frac{u_{i+1} - 2u_i + u_{i-1}}{(\Delta x)^2} + 0.5 \frac{u_{i+1} - u_{i-1}}{2(\Delta x)}. \quad (9)$$

This equation holds for $i = 1, 2, \dots, \text{MX}$, with the understanding that $u_0 = u_{\text{MX}+1} = 0$.

In the parallel processing environment, we may think of the several processors as being laid out on a straight line with each processor to compute its contiguous subset of the solution vector. Consequently the computation of the right hand side of Eq. (9) requires that each interior processor must pass the first component of its block of the solution vector to its left-hand neighbor, acquire the last component of that neighbor's block, pass the last component of its block of the solution vector to its right-hand neighbor, and acquire the first component of that neighbor's block. If the processor is the first (0th) or last processor, then communication to the left or right (respectively) is not required.

This problem uses the Adams (non-stiff) integration formula and functional iteration. It is unrealistically simple, but serves to illustrate use of the parallel version of CVODE.

The `cvAdvDiff_non_p.c` file begins with `#include` declarations for various required header files, including lines for `nvector_parallel` to access the parallel `N_Vector` type and related macros, and for `mpi.h` to access MPI types and constants. Following that are definitions of problem constants and a data block for communication with the `f` routine. That block includes the number of PEs, the index of the local PE, and the MPI communicator.

The `main` program begins with MPI calls to initialize MPI and to set multi-processor environment parameters `npes` (number of PEs) and `my_pe` (local PE index). The local vector length is set according to `npes` and the problem size `NEQ` (which may or may not be multiple of `npes`). The value `my_base` is the base value for computing global indices (from 1 to `NEQ`) for the local vectors. The solution vector `u` is created with a call to `N_VNew_Parallel` and loaded with a call to `SetIC`. The calls to `CVodeCreate`, `CVodeInit`, and `CVodeSStolerances` specify a CVODE solution with the nonstiff method and scalar tolerances. The call to `CVodeSetUserData` insures that the pointer `data` is passed to the `f` routine whenever it is called. A heading is printed (if on processor 0). In a loop over `tout` values, `CVode` is called, and the return value checked for errors. The max-norm of the solution and the total number

of time steps so far are printed at each output point. Finally, some statistical counters are printed, memory is freed, and MPI is finalized.

The `SetIC` routine uses the last two arguments passed to it to compute the set of global indices (`my_base+1` to `my_base+my_length`) corresponding to the local part of the solution vector `u`, and then to load the corresponding initial values. The `PrintFinalStats` routine uses `CVodeGet***` calls to get various counters, and then prints these. The counters are: `nst` (number of steps), `nfe` (number of `f` evaluations), `nni` (number of nonlinear iterations), `netf` (number of error test failures), and `ncfn` (number of nonlinear convergence failures). This routine is suitable for general use with CVODE applications to nonstiff problems.

The `f` function is an implementation of Eq. (9), but preceded by communication operations appropriate for the parallel setting. It copies the local vector `u` into a larger array `z`, shifted by 1 to allow for the storage of immediate neighbor components. The first and last components of `u` are sent to neighboring processors with `MPI_Send` calls, and the immediate neighbor solution values are received from the neighbor processors with `MPI_Recv` calls, except that zero is loaded into `z[0]` or `z[my_length+1]` instead if at the actual boundary. Then the central difference expressions are easily formed from the `z` array, and loaded into the data array of the `udot` vector.

The `cvAdvDiff_non_p.c` file includes a routine `check_flag` that checks the return values from calls in `main`. This routine was written to be used by any parallel SUNDIALS application.

The output below is for `cvAdvDiff_non_p` with `MX = 10` and four processors. Varying the number of processors will alter the output, only because of roundoff-level differences in various vector operations. The fairly high value of `ncfn` indicates that this problem is on the borderline of being stiff.

```

                                cvAdvDiff_non_p sample output
-----

1-D advection-diffusion equation, mesh size = 10

Number of PEs =    4

At t = 0.00  max.norm(u) =  1.569909e+01  nst =    0
At t = 0.50  max.norm(u) =  3.052881e+00  nst =  113
At t = 1.00  max.norm(u) =  8.753188e-01  nst =  191
At t = 1.50  max.norm(u) =  2.494926e-01  nst =  265
At t = 2.00  max.norm(u) =  7.109707e-02  nst =  339
At t = 2.50  max.norm(u) =  2.026223e-02  nst =  418
At t = 3.00  max.norm(u) =  5.777332e-03  nst =  486
At t = 3.50  max.norm(u) =  1.650483e-03  nst =  563
At t = 4.00  max.norm(u) =  4.754357e-04  nst =  646
At t = 4.50  max.norm(u) =  1.374222e-04  nst =  715
At t = 5.00  max.norm(u) =  3.937469e-05  nst =  795

Final Statistics:

nst = 795      nfe = 1465      nni = 1461      ncfn = 146      netf = 5

```

3.2 A user preconditioner example: `cvDiurnal_kry_p`

As an example of using CVODE with the Krylov linear solver CVSPGMR and the parallel MPI `NVECTOR_PARALLEL` module, we describe a test problem based on the system PDEs given

above for the `cvDiurnal_kry` example. As before, we discretize the PDE system with central differencing, to obtain an ODE system $\dot{u} = f(t, u)$ representing (4). But in this case, the discrete solution vector is distributed over many processors. Specifically, we may think of the processors as being laid out in a rectangle, and each processor being assigned a subgrid of size $\text{MXSUB} \times \text{MYSUB}$ of the $x - y$ grid. If there are NPEX processors in the x direction and NPEY processors in the y direction, then the overall grid size is $\text{MX} \times \text{MY}$ with $\text{MX} = \text{NPEX} \times \text{MXSUB}$ and $\text{MY} = \text{NPEY} \times \text{MYSUB}$, and the size of the ODE system is $2 \cdot \text{MX} \cdot \text{MY}$.

To compute f in this setting, the processors pass and receive information as follows. The solution components for the bottom row of grid points in the current processor are passed to the processor below it and the solution for the top row of grid points is received from the processor below the current processor. The solution for the top row of grid points for the current processor is sent to the processor above the current processor, while the solution for the bottom row of grid points is received from that processor by the current processor. Similarly the solution for the first column of grid points is sent from the current processor to the processor to its left and the last column of grid points is received from that processor by the current processor. The communication for the solution at the right edge of the processor is similar. If this is the last processor in a particular direction, then message passing and receiving are bypassed for that direction.

This code is intended to provide a more realistic example than that in `cvAdvDiff_non_p`, and to provide a template for a stiff ODE system arising from a PDE system. The solution method is BDF with Newton iteration and SPGMR. The left preconditioner is the block-diagonal part of the Newton matrix, with 2×2 blocks, and the corresponding diagonal blocks of the Jacobian are saved each time the preconditioner is generated, for re-use later under certain conditions.

The organization of the `cvDiurnal_kry_p` program deserves some comments. The right-hand side routine `f` calls two other routines: `ucomm`, which carries out inter-processor communication; and `fcalc`, which operates on local data only and contains the actual calculation of $f(t, u)$. The `ucomm` function in turn calls three routines which do, respectively, non-blocking receive operations, blocking send operations, and receive-waiting. All three use MPI, and transmit data from the local `u` vector into a local working array `uext`, an extended copy of `u`. The `fcalc` function copies `u` into `uext`, so that the calculation of $f(t, u)$ can be done conveniently by operations on `uext` only. Most other features of `cvDiurnal_kry_p.c` are the same as in `cvDiurnal_kry.c`, except for extra logic involved with distributed vectors.

The following is a sample output from `cvDiurnal_kry_p`, for four processors (in a 2×2 array) with a 5×5 subgrid on each. The output will vary slightly if the number of processors is changed.

```

cvDiurnal_kry_p sample output

2-species diurnal advection-diffusion problem

t = 7.20e+03   no. steps = 219   order = 5   stepsize = 1.59e+02
At bottom left:  c1, c2 =    1.047e+04    2.527e+11
At top right:    c1, c2 =    1.119e+04    2.700e+11

t = 1.44e+04   no. steps = 251   order = 5   stepsize = 3.77e+02
At bottom left:  c1, c2 =    6.659e+06    2.582e+11
At top right:    c1, c2 =    7.301e+06    2.833e+11

t = 2.16e+04   no. steps = 277   order = 5   stepsize = 2.75e+02

```

```

At bottom left:  c1, c2 =    2.665e+07    2.993e+11
At top right:    c1, c2 =    2.931e+07    3.313e+11

t = 2.88e+04    no. steps = 307    order = 4    stepsize = 1.98e+02
At bottom left:  c1, c2 =    8.702e+06    3.380e+11
At top right:    c1, c2 =    9.650e+06    3.751e+11

t = 3.60e+04    no. steps = 335    order = 5    stepsize = 1.17e+02
At bottom left:  c1, c2 =    1.404e+04    3.387e+11
At top right:    c1, c2 =    1.561e+04    3.765e+11

t = 4.32e+04    no. steps = 388    order = 4    stepsize = 4.48e+02
At bottom left:  c1, c2 =   -5.732e-07    3.382e+11
At top right:    c1, c2 =   -6.367e-07    3.804e+11

t = 5.04e+04    no. steps = 406    order = 5    stepsize = 3.97e+02
At bottom left:  c1, c2 =   -4.317e-09    3.358e+11
At top right:    c1, c2 =   -8.233e-09    3.864e+11

t = 5.76e+04    no. steps = 418    order = 5    stepsize = 4.74e+02
At bottom left:  c1, c2 =   -2.576e-09    3.320e+11
At top right:    c1, c2 =   -1.259e-09    3.909e+11

t = 6.48e+04    no. steps = 428    order = 5    stepsize = 7.70e+02
At bottom left:  c1, c2 =    3.451e-09    3.313e+11
At top right:    c1, c2 =    2.081e-09    3.963e+11

t = 7.20e+04    no. steps = 437    order = 5    stepsize = 7.70e+02
At bottom left:  c1, c2 =    1.630e-11    3.330e+11
At top right:    c1, c2 =    1.843e-11    4.039e+11

t = 7.92e+04    no. steps = 447    order = 5    stepsize = 7.70e+02
At bottom left:  c1, c2 =   -1.704e-11    3.334e+11
At top right:    c1, c2 =   -1.131e-11    4.120e+11

t = 8.64e+04    no. steps = 456    order = 5    stepsize = 7.70e+02
At bottom left:  c1, c2 =    1.496e-12    3.352e+11
At top right:    c1, c2 =    8.085e-13    4.163e+11

Final Statistics:

lenrw  = 2089    leniw  = 120
lenrwls = 2046    leniwls = 80
nst    = 456
nfe    = 586    nfels  = 619
nni    = 582    nli    = 619
nsetups = 73    netf   = 25
npe    = 8      nps    = 1149
ncfn   = 0      ncfl   = 0

```

3.3 A CVBBDPRE preconditioner example: cvDiurnal_kry_bbd_p

In this example, `cvDiurnal_kry_bbd_p`, we solve the same problem as in `cvDiurnal_kry_p` above, but instead of supplying the preconditioner, we use the CVBBDPRE module, which generates and uses a band-block-diagonal preconditioner. The half-bandwidths of the Jaco-

bian block on each processor are both equal to $2 \cdot \text{MXSUB}$, and that is the value supplied as `mudq` and `mldq` in the call to `CVBBDPrecInit`. But in order to reduce storage and computation costs for preconditioning, we supply the values `mukeep = mlkeep = 2` ($= \text{NVAR}$) as the half-bandwidths of the retained band matrix blocks. This means that the Jacobian elements are computed with a difference quotient scheme using the true bandwidth of the block, but only a narrow band matrix (bandwidth 5) is kept as the preconditioner.

As in `cvDiurnal_kry_p.c`, the `f` routine in `cvDiurnal_kry_bbd_p.c` simply calls a communication routine, `fucomm`, and then a strictly computational routine, `flocal`. However, the call to `CVBBDPrecInit` specifies the pair of routines to be called as `ucomm` and `flocal`, where `ucomm` is `NULL`. This is because each call by the solver to `ucomm` is preceded by a call to `f` with the same (\mathbf{t}, \mathbf{u}) arguments, and therefore the communication needed for `flocal` in the solver's calls to it have already been done.

In `cvDiurnal_kry_bbd_p.c`, the problem is solved twice — first with preconditioning on the left, and then on the right. Thus prior to the second solution, calls are made to reset the initial values (`SetInitialProfiles`), the main solver memory (`CVodeReInit`), the CVBBD-PRE memory (`CVBBDPrecReInit`), as well as the preconditioner type (`CVSpilsSetPrecType`).

Sample output from `cvDiurnal_kry_bbd_p` follows, again using 5×5 subgrids on a 2×2 processor grid. The performance of the preconditioner, as measured by the number of Krylov iterations per Newton iteration, `nli/nni`, is very close to that of `cvDiurnal_kry_p` when preconditioning is on the left, but slightly poorer when it is on the right.

cvDiurnal_kry_bbd_p sample output

```

2-species diurnal advection-diffusion problem
 10 by 10 mesh on 4 processors
Using CVBBDPRE preconditioner module
  Difference-quotient half-bandwidths are mudq = 10,  mldq = 10
  Retained band block half-bandwidths are mukeep = 2,  mlkeep = 2

Preconditioner type is:  jpre = PREC_LEFT

t = 7.20e+03  no. steps = 190  order = 5  stepsize = 1.61e+02
At bottom left:  c1, c2 = 1.047e+04  2.527e+11
At top right:   c1, c2 = 1.119e+04  2.700e+11

t = 1.44e+04  no. steps = 221  order = 5  stepsize = 3.85e+02
At bottom left:  c1, c2 = 6.659e+06  2.582e+11
At top right:   c1, c2 = 7.301e+06  2.833e+11

t = 2.16e+04  no. steps = 247  order = 5  stepsize = 3.00e+02
At bottom left:  c1, c2 = 2.665e+07  2.993e+11
At top right:   c1, c2 = 2.931e+07  3.313e+11

t = 2.88e+04  no. steps = 272  order = 4  stepsize = 4.05e+02
At bottom left:  c1, c2 = 8.702e+06  3.380e+11
At top right:   c1, c2 = 9.650e+06  3.751e+11

t = 3.60e+04  no. steps = 309  order = 4  stepsize = 7.53e+01
At bottom left:  c1, c2 = 1.404e+04  3.387e+11
At top right:   c1, c2 = 1.561e+04  3.765e+11

t = 4.32e+04  no. steps = 377  order = 4  stepsize = 4.02e+02
At bottom left:  c1, c2 = 1.908e-07  3.382e+11
At top right:   c1, c2 = 2.345e-07  3.804e+11

```

```

t = 5.04e+04   no. steps = 392   order = 5   stepsize = 3.67e+02
At bottom left:  c1, c2 =   -6.408e-10   3.358e+11
At top right:    c1, c2 =   -6.654e-10   3.864e+11

t = 5.76e+04   no. steps = 403   order = 5   stepsize = 4.72e+02
At bottom left:  c1, c2 =    2.017e-08   3.320e+11
At top right:    c1, c2 =    3.353e-08   3.909e+11

t = 6.48e+04   no. steps = 415   order = 5   stepsize = 7.47e+02
At bottom left:  c1, c2 =   -2.502e-10   3.313e+11
At top right:    c1, c2 =    2.005e-10   3.963e+11

t = 7.20e+04   no. steps = 424   order = 5   stepsize = 7.47e+02
At bottom left:  c1, c2 =    4.217e-12   3.330e+11
At top right:    c1, c2 =   -2.693e-12   4.039e+11

t = 7.92e+04   no. steps = 434   order = 5   stepsize = 7.47e+02
At bottom left:  c1, c2 =    2.779e-12   3.334e+11
At top right:    c1, c2 =   -1.865e-12   4.120e+11

t = 8.64e+04   no. steps = 444   order = 5   stepsize = 7.47e+02
At bottom left:  c1, c2 =    2.331e-13   3.352e+11
At top right:    c1, c2 =   -1.599e-13   4.163e+11

```

Final Statistics:

```

lenrw  = 2089      leniw  = 120
lenrwls = 2046     leniwls = 80
nst     = 444
nfe     = 581      nfels  = 526
nni     = 577      nli    = 526
nsetups = 75       netf   = 28
npe     = 8        nps    = 1057
ncfn    = 0        ncfl   = 0

```

```

In CVBBDPRE: real/integer local work space sizes = 600, 50
              no. floccal evals. = 176

```

Preconditioner type is: jpre = PREC_RIGHT

```

t = 7.20e+03   no. steps = 191   order = 5   stepsize = 1.22e+02
At bottom left:  c1, c2 =    1.047e+04   2.527e+11
At top right:    c1, c2 =    1.119e+04   2.700e+11

t = 1.44e+04   no. steps = 223   order = 5   stepsize = 2.79e+02
At bottom left:  c1, c2 =    6.659e+06   2.582e+11
At top right:    c1, c2 =    7.301e+06   2.833e+11

t = 2.16e+04   no. steps = 249   order = 5   stepsize = 4.31e+02
At bottom left:  c1, c2 =    2.665e+07   2.993e+11
At top right:    c1, c2 =    2.931e+07   3.313e+11

t = 2.88e+04   no. steps = 314   order = 3   stepsize = 9.38e+01

```

```

At bottom left:  c1, c2 =    8.702e+06    3.380e+11
At top right:    c1, c2 =    9.650e+06    3.751e+11

t = 3.60e+04    no. steps = 350    order = 5    stepsize = 9.78e+01
At bottom left:  c1, c2 =    1.404e+04    3.387e+11
At top right:    c1, c2 =    1.561e+04    3.765e+11

t = 4.32e+04    no. steps = 403    order = 4    stepsize = 3.87e+02
At bottom left:  c1, c2 =    1.504e-09    3.382e+11
At top right:    c1, c2 =    1.683e-09    3.804e+11

t = 5.04e+04    no. steps = 416    order = 5    stepsize = 5.91e+02
At bottom left:  c1, c2 =   -1.137e-11    3.358e+11
At top right:    c1, c2 =   -1.439e-11    3.864e+11

t = 5.76e+04    no. steps = 432    order = 5    stepsize = 1.73e+02
At bottom left:  c1, c2 =    1.293e-09    3.320e+11
At top right:    c1, c2 =    2.448e-10    3.909e+11

t = 6.48e+04    no. steps = 447    order = 5    stepsize = 6.38e+02
At bottom left:  c1, c2 =    7.963e-13    3.313e+11
At top right:    c1, c2 =   -2.943e-13    3.963e+11

t = 7.20e+04    no. steps = 459    order = 5    stepsize = 6.38e+02
At bottom left:  c1, c2 =   -2.414e-12    3.330e+11
At top right:    c1, c2 =    2.797e-13    4.039e+11

t = 7.92e+04    no. steps = 470    order = 5    stepsize = 6.38e+02
At bottom left:  c1, c2 =   -1.059e-13    3.334e+11
At top right:    c1, c2 =    3.557e-14    4.120e+11

t = 8.64e+04    no. steps = 481    order = 5    stepsize = 6.38e+02
At bottom left:  c1, c2 =    6.045e-15    3.352e+11
At top right:    c1, c2 =   -2.016e-15    4.163e+11

```

Final Statistics:

```

lenrw  = 2089    leniw  = 120
lenrwls = 2046    leniwls = 80
nst    = 481
nfe    = 622    nfels  = 769
nni    = 618    nli    = 769
nsetups = 104    netf   = 33
npe    = 9      nps    = 1281
ncfn   = 0      ncfl   = 0

```

```

In CVBBDPRE: real/integer local work space sizes = 600, 50
              no. flocal evals. = 198

```

4 *hypr*e example problems

4.1 A nonstiff example: `cvAdvDiff_non_ph`

This example is same as 3.1, except it uses *hypr*e vector instead of SUNDIALS native parallel vector implementation. The output of the two examples is identical. In the following, we will point out only the differences between the two. Familiarity with *hypr*e library [1] is helpful.

We use *hypr*e IJ vector interface to allocate the template vector and create parallel partitioning:

```
HYPRE_IJVectorCreate(comm, my_base, my_base + local_N - 1, &Uij);
HYPRE_IJVectorSetObjectType(Uij, HYPRE_PARCSR);
HYPRE_IJVectorInitialize(Uij);
```

The initialize call means that vector elements are ready to be set using the IJ interface. We choose initial condition vector $x_0 = x(t_0)$ as the template vector and we set its values in `SetIC(...)` function. We complete *hypr*e vector assembly by:

```
HYPRE_IJVectorAssemble(Uij);
HYPRE_IJVectorGetObject(Uij, (void**) &Upar);
```

The assemble call is collective and it makes *hypr*e vector ready to use. The sets the handle `Upar` to the actual *hypr*e vector. The handle is then passed to the `N_VMake` function, which creates the template `N_Vector` as a wrapper around the *hypr*e vector. All other vectors in the computation are created by cloning the template vector. The template vector does not own the underlying *hypr*e vector, and it is user's responsibility to destroy it using `HYPRE_IJVectorDestroy(Uij)` call after the template vector has been destroyed. This function will destroy both, the *hypr*e vector and its IJ interface.

To access individual elements of solution vectors `u` and `u \dot` in the residual function, user needs to extract the *hypr*e vector first by calling `N_VGetVector_ParHyp`, and then use *hypr*e methods from that point on.

Notes



- At this point interfaces to *hypr*e solvers and preconditioners are not available. They will be provided in subsequent SUNDIALS releases. The interface to *hypr*e vector is included in this release mainly for testing purposes and as a preview of functionality to come.

5 Fortran example problems

The FORTRAN example problem programs supplied with the CVODE package are all written in standard FORTRAN77 and use double precision arithmetic. However, when the FORTRAN examples are built, the source code is automatically modified according to the configure options supplied by the user and the system type. Integer variables are declared as `INTEGER*n`, where n denotes the number of bytes in the corresponding C type (`long int` or `int`). Floating-point variable declarations remain unchanged if double precision is used, but are changed to `REAL*n`, where n denotes the number of bytes in the SUNDIALS type `realtype`, if using single precision. Also, if using single precision, then declarations of floating-point constants are appropriately modified; e.g. `0.5D-4` is changed to `0.5E-4`.

5.1 A serial example: `fcvDiurnal kry`

The `fcvDiurnal kry` example is a FORTRAN equivalent of the `cvDiurnal kry` problem. (In fact, it was derived from an earlier FORTRAN example program for `VODPK`.)

The main program begins with a call to `INITKX`, which sets problem parameters, loads these into arrays `IPAR` and `RPAR` for use by other routines, and loads `Y` (here called `U0`) with its initial values. `Main` calls `FNVINITS`, `FCVMALLOC`, `FCVSPGMR`, and `FCVSPILSSETPREC`, to initialize the `NVECTOR_SERIAL` module, the main solver memory, and the `CVSPGMR` module, and to specify user-supplied preconditioner setup and solve routines. It calls `FCVODE` in a loop over `TOUT` values, with printing of selected solution values and performance data (from the `IOUT` and `ROUT` arrays). At the end, it prints a number of performance counters, and frees memory with calls to `FCVFREE`.

In `fcvDiurnal kry.f`, the `FCVFUN` routine is a straightforward implementation of the discretized form of Eqns. (4). In `FCVPSET`, the block-diagonal part of the Jacobian, J_{bd} , is computed (and copied to `P`) if `JOK = 0`, but is simply copied from `BD` to `P` if `JOK = 1`. In both cases, the preconditioner matrix P is formed from J_{bd} and its 2×2 blocks are LU-factored. In `FCVPSOL`, the solution of a linear system $Px = z$ is solved by doing backsolve operations on the blocks. Subordinate routines are used to isolate these evaluation, factorization, and backsolve operations. The remainder of `fcvDiurnal kry.f` consists of routines from `LINPACK` and the `BLAS` needed for matrix and vector operations.

The following is sample output from `fcvDiurnal kry`, using a 10×10 mesh. The performance of `FCVODE` here is quite similar to that of `CVODE` on the `cvDiurnal kry` problem, as expected.

```
fcvDiurnal kry sample output
Krylov example problem:

Kinetics-transport, NEQ = 200

t = 0.720E+04 nst = 219 q = 5 h = 0.158696E+03
c1 (bot.left/middle/top rt.) = 0.104683E+05 0.296373E+05 0.111853E+05
c2 (bot.left/middle/top rt.) = 0.252672E+12 0.715376E+12 0.269977E+12

t = 0.144E+05 nst = 251 q = 5 h = 0.377205E+03
c1 (bot.left/middle/top rt.) = 0.665902E+07 0.531602E+07 0.730081E+07
c2 (bot.left/middle/top rt.) = 0.258192E+12 0.205680E+12 0.283286E+12

t = 0.216E+05 nst = 277 q = 5 h = 0.274587E+03
```

```

c1 (bot.left/middle/top rt.) = 0.266498E+08 0.103636E+08 0.293077E+08
c2 (bot.left/middle/top rt.) = 0.299279E+12 0.102810E+12 0.331344E+12

t = 0.288E+05 nst = 312 q = 4 h = 0.367517E+03
c1 (bot.left/middle/top rt.) = 0.870209E+07 0.129197E+08 0.965002E+07
c2 (bot.left/middle/top rt.) = 0.338035E+12 0.502929E+12 0.375096E+12

t = 0.360E+05 nst = 350 q = 4 h = 0.683836E+02
c1 (bot.left/middle/top rt.) = 0.140404E+05 0.202903E+05 0.156090E+05
c2 (bot.left/middle/top rt.) = 0.338677E+12 0.489443E+12 0.376517E+12

t = 0.432E+05 nst = 407 q = 4 h = 0.383863E+03
c1 (bot.left/middle/top rt.) = 0.803367E-06 0.363858E-06 0.889797E-06
c2 (bot.left/middle/top rt.) = 0.338233E+12 0.135487E+12 0.380352E+12

t = 0.504E+05 nst = 436 q = 3 h = 0.215343E+03
c1 (bot.left/middle/top rt.) = 0.375001E-05 0.665499E-06 0.454113E-05
c2 (bot.left/middle/top rt.) = 0.335816E+12 0.493028E+12 0.386445E+12

t = 0.576E+05 nst = 454 q = 5 h = 0.428080E+03
c1 (bot.left/middle/top rt.) = 0.112301E-08 0.194567E-09 0.136087E-08
c2 (bot.left/middle/top rt.) = 0.332031E+12 0.964985E+12 0.390900E+12

t = 0.648E+05 nst = 466 q = 5 h = 0.690422E+03
c1 (bot.left/middle/top rt.) = 0.353041E-08 0.590752E-09 0.428410E-08
c2 (bot.left/middle/top rt.) = 0.331303E+12 0.892184E+12 0.396342E+12

t = 0.720E+05 nst = 476 q = 5 h = 0.690422E+03
c1 (bot.left/middle/top rt.) = -0.121418E-09 -0.206756E-10 -0.147240E-09
c2 (bot.left/middle/top rt.) = 0.332972E+12 0.618620E+12 0.403885E+12

t = 0.792E+05 nst = 487 q = 5 h = 0.690422E+03
c1 (bot.left/middle/top rt.) = -0.341376E-11 -0.568210E-12 -0.414339E-11
c2 (bot.left/middle/top rt.) = 0.333441E+12 0.666893E+12 0.412026E+12

t = 0.864E+05 nst = 497 q = 5 h = 0.690422E+03
c1 (bot.left/middle/top rt.) = 0.309841E-12 0.526192E-13 0.375773E-12
c2 (bot.left/middle/top rt.) = 0.335178E+12 0.910652E+12 0.416251E+12

```

Final statistics:

```

number of steps      = 497          number of f evals.    = 651
number of prec. setups = 91
number of prec. evals. = 9          number of prec. solves = 1233
number of nonl. iters. = 647        number of lin. iters. = 652
average Krylov subspace dimension (NLI/NNI) = 0.100773E+01
number of conv. failures.. nonlinear = 0 linear = 0
number of error test failures = 34

```

5.2 A parallel example: fcvDiag_kry_bbd_p

This example, `fcvDiag_kry_bbd_p`, uses a simple diagonal ODE system to illustrate the use of FCVODE in a parallel setting. The system is

$$\dot{y}_i = -\alpha i y_i \quad (i = 1, \dots, N) \quad (10)$$

on the time interval $0 \leq t \leq 1$. In this case, we use $\alpha = 10$ and $N = 10 \cdot \text{NPES}$, where NPES is the number of processors and is specified at run time. The linear solver to be used is SPGMR with the CVBBDPRE (band-block-diagonal) preconditioner. Since the system Jacobian is diagonal, the half-bandwidths specified are all zero. The problem is solved twice — with preconditioning on the left, then on the right.

The source file for this problem begins with MPI calls to initialize MPI and to get the number of processors and local processor index. Following the call to FCVMALLOC, the linear solver specification is done with calls to FCVSPGMR and FCVBBDINIT. In a loop over TOUT values, it calls FCVODE and prints the step and f evaluation counters. After that, it computes and prints the maximum global error, and all the relevant performance counters. Those specific to CVBBDPRE are obtained by a call to FCVBBDOPT. To prepare for the second run, the program calls FCVREINIT, FCVBBDREINIT, and FCVSPGMRREINIT, in addition to resetting the initial conditions. Finally, it frees memory and terminates MPI. Notice that in the FCVFUN routine, the local processor index MYPE and the local vector size NLOCAL are used to form the global index values needed to evaluate the right-hand side of Eq. (10).

The following is a sample output from fcvDiag_kry_bbd_p, with NPES = 4. As expected, the performance is identical for left vs right preconditioning.

```

----- fcvDiag_kry_bbd_p sample output -----
Diagonal test problem:

NEQ = 40
parameter alpha = 10.000
ydot_i = -alpha*i * y_i (i = 1,...,NEQ)
RTOL, ATOL = 0.1E-04 0.1E-09
Method is BDF/NEWTON/SPGMR
Preconditioner is band-block-diagonal, using CVBBDPRE
Number of processors = 4

Preconditioning on left

t = 0.10E+00    no. steps = 221    no. f-s = 262
t = 0.20E+00    no. steps = 265    no. f-s = 308
t = 0.30E+00    no. steps = 290    no. f-s = 334
t = 0.40E+00    no. steps = 306    no. f-s = 351
t = 0.50E+00    no. steps = 319    no. f-s = 365
t = 0.60E+00    no. steps = 329    no. f-s = 375
t = 0.70E+00    no. steps = 339    no. f-s = 386
t = 0.80E+00    no. steps = 345    no. f-s = 392
t = 0.90E+00    no. steps = 352    no. f-s = 399
t = 0.10E+01    no. steps = 359    no. f-s = 406

Max. absolute error is 0.28E-08

Final statistics:

number of steps          = 359      number of f evals.      = 406
number of prec. setups   = 38
number of prec. evals.   = 7        number of prec. solves  = 728
number of nonl. iters.   = 402      number of lin. iters.   = 364
average Krylov subspace dimension (NLI/NNI) = 0.9055
number of conv. failures.. nonlinear = 0 linear = 0
number of error test failures = 5

```

```
main solver real/int workspace sizes = 489 120
linear solver real/int workspace sizes = 446 80
In CVBBDPRE:
real/int local workspace = 20 10
number of g evals. = 14
```

Preconditioning on right

t =	0.10E+00	no. steps =	221	no. f-s =	262
t =	0.20E+00	no. steps =	265	no. f-s =	308
t =	0.30E+00	no. steps =	290	no. f-s =	334
t =	0.40E+00	no. steps =	306	no. f-s =	351
t =	0.50E+00	no. steps =	319	no. f-s =	365
t =	0.60E+00	no. steps =	329	no. f-s =	375
t =	0.70E+00	no. steps =	339	no. f-s =	386
t =	0.80E+00	no. steps =	345	no. f-s =	392
t =	0.90E+00	no. steps =	352	no. f-s =	399
t =	0.10E+01	no. steps =	359	no. f-s =	406

Max. absolute error is 0.28E-08

Final statistics:

number of steps	=	359	number of f evals.	=	406
number of prec. setups	=	38			
number of prec. evals.	=	7	number of prec. solves	=	728
number of nonl. iters.	=	402	number of lin. iters.	=	364
average Krylov subspace dimension (NLI/NNI)	=	0.9055			
number of conv. failures.. nonlinear	=	0	linear	=	0
number of error test failures	=	5			
main solver real/int workspace sizes	=	489 120			
linear solver real/int workspace sizes	=	446 80			
In CVBBDPRE:					
real/int local workspace	=	20 10			
number of g evals.	=	14			

6 Parallel tests

The stiff example problem `cvDiurnal_kry` described above, or rather its parallel version `cvDiurnal_kry_p`, has been modified and expanded to form a test problem for the parallel version of `CVODE`. This work was largely carried out by M. Wittman and reported in [3].

To start with, in order to add realistic complexity to the solution, the initial profile for this problem was altered to include a rather steep front in the vertical direction. Specifically, the function $\beta(y)$ in Eq. (6) has been replaced by:

$$\beta(y) = .75 + .25 \tanh(10y - 400) . \quad (11)$$

This function rises from about .5 to about 1.0 over a y interval of about .2 (i.e. 1/100 of the total span in y). This vertical variation, together with the horizontal advection and diffusion in the problem, demands a fairly fine spatial mesh to achieve acceptable resolution.

In addition, an alternate choice of differencing is used in order to control spurious oscillations resulting from the horizontal advection. In place of central differencing for that term, a biased upwind approximation is applied to each of the terms $\partial c^i / \partial x$, namely:

$$\partial c / \partial x|_{x_j} \approx \left[\frac{3}{2} c_{j+1} - c_j - \frac{1}{2} c_{j-1} \right] / (2\Delta x) . \quad (12)$$

With this modified form of the problem, we performed tests similar to those described above for the example. Here we fix the subgrid dimensions at `MXSUB` = `MYSUB` = 50, so that the local (per-processor) problem size is 5000, while the processor array dimensions, `NPEX` and `NPEY`, are varied. In one (typical) sequence of tests, we fix `NPEY` = 8 (for a vertical mesh size of `MY` = 400), and set `NPEX` = 8 (`MX` = 400), `NPEX` = 16 (`MX` = 800), and `NPEX` = 32 (`MX` = 1600). Thus the largest problem size N is $2 \cdot 400 \cdot 1600 = 1,280,000$. For these tests, we also raise the maximum Krylov dimension, `max1`, to 10 (from its default value of 5).

For each of the three test cases, the test program was run on a Cray-T3D (256 processors) with each of three different message-passing libraries:

- `MPICH`: an implementation of MPI on top of the Chameleon library
- `EPCC`: an implementation of MPI by the Edinburgh Parallel Computer Centre
- `SHMEM`: Cray's Shared Memory Library

The following table gives the run time and selected performance counters for these 9 runs. In all cases, the solutions agreed well with each other, showing expected small variations with grid size. In the table, M-P denotes the message-passing library, RT is the reported run time in CPU seconds, `nst` is the number of time steps, `nfe` is the number of f evaluations, `nni` is the number of nonlinear (Newton) iterations, `nli` is the number of linear (Krylov) iterations, and `npe` is the number of evaluations of the preconditioner.

Some of the results were as expected, and some were surprising. For a given mesh size, variations in performance counts were small or absent, except for moderate (but still acceptable) variations for `SHMEM` in the smallest case. The increase in costs with mesh size can be attributed to a decline in the quality of the preconditioner, which neglects most of the spatial coupling. The preconditioner quality can be inferred from the ratio `nli/nni`, which is the average number of Krylov iterations per Newton iteration. The most interesting (and unexpected) result is the variation of run time with library: `SHMEM` is the most efficient,

NPEX	M-P	RT	nst	nfe	nni	nli	npe
8	MPICH	436.	1391	9907	1512	8392	24
8	EPCC	355.	1391	9907	1512	8392	24
8	SHMEM	349.	1999	10,326	2096	8227	34
16	MPICH	676.	2513	14,159	2583	11,573	42
16	EPCC	494.	2513	14,159	2583	11,573	42
16	SHMEM	471.	2513	14,160	2581	11,576	42
32	MPICH	1367.	2536	20,153	2696	17,454	43
32	EPCC	737.	2536	20,153	2696	17,454	43
32	SHMEM	695.	2536	20,121	2694	17,424	43

Table 1: Parallel CVODE test results vs problem size and message-passing library

but EPCC is a very close second, and MPICH loses considerable efficiency by comparison, as the problem size grows. This means that the highly portable MPI version of CVODE, with an appropriate choice of MPI implementation, is fully competitive with the Cray-specific version using the SHMEM library. While the overall costs do not represent a well-scaled parallel algorithm (because of the preconditioner choice), the cost per function evaluation is quite flat for EPCC and SHMEM, at .033 to .037 (for MPICH it ranges from .044 to .068).

For tests that demonstrate speedup from parallelism, we consider runs with fixed problem size: $MX = 800$, $MY = 400$. Here we also fix the vertical subgrid dimension at $MYSUB = 50$ and the vertical processor array dimension at $NPEY = 8$, but vary the corresponding horizontal sizes. We take $NPEX = 8, 16, \text{ and } 32$, with $MXSUB = 100, 50, \text{ and } 25$, respectively. The runs for the three cases and three message-passing libraries all show very good agreement in solution values and performance counts. The run times for EPCC are 947, 494, and 278, showing speedups of 1.92 and 1.78 as the number of processors is doubled (twice). For the SHMEM runs, the times were slightly lower and the ratios were 1.98 and 1.91. For MPICH, consistent with the earlier runs, the run times were considerably higher, and in fact show speedup ratios of only 1.54 and 1.03.

References

- [1] R Falgout and UM Yang. Hypre user's manual. Technical report, LLNL, 2015.
- [2] A. C. Hindmarsh and R. Serban. User Documentation for CVODE v2.7.0. Technical Report UCRL-SM-208108, LLNL, 2011.
- [3] M. R. Wittman. Testing of PVODE, a Parallel ODE Solver. Technical Report UCRL-ID-125562, LLNL, August 1996.

