

Example Programs for KINSOL v7.5.0

Aaron M. Collier and Radu Serban
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory

September 29, 2025



UCRL-SM-208114

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.

CONTRIBUTORS

The SUNDIALS library has been developed over many years by a number of contributors. The current SUNDIALS team consists of Cody J. Balos, David J. Gardner, Alan C. Hindmarsh, Daniel R. Reynolds, and Carol S. Woodward. We thank Radu Serban for significant and critical past contributions.

Other contributors to SUNDIALS include: James Almgren-Bell, Lawrence E. Banks, Peter N. Brown, George Byrne, Rujeko Chinomona, Scott D. Cohen, Aaron Collier, Keith E. Grant, Steven L. Lee, Shelby L. Lockhart, John Loffeld, Daniel McGreer, Yu Pan, Slaven Peles, Cosmin Petra, Steven B. Roberts, H. Hunter Schwartz, Jean M. Sexton, Dan Shumaker, Steve G. Smith, Shahbaj Sohal, Allan G. Taylor, Hilari C. Tiedeman, Chris White, Ting Yan, and Ulrike M. Yang.

Contents

1	Introduction	1
2	C example problems	4
3	C++ example problems	12
	References	16

1 Introduction

This report is intended to serve as a companion document to the User Documentation of KINSOL [1]. It provides details, with listings, on the example programs supplied with the KINSOL distribution package.

The KINSOL distribution contains examples of many types, so as to illustrate the use of various integrator options, various linear solver options, and various NVECTOR kernels. With the exception of "demo"-type example files, the names of all the examples distributed with SUNDIALS are of the form `[slv][PbName]_[strat]_[ls]_[prec]_[p]`, where

`[slv]` identifies the solver (in this case `kin`);

`[PbName]` identifies the problem;

`[strat]` identifies the strategy (absent if "none" or "linesearch");

`[ls]` identifies the linear solver module used;

`[prec]` indicates the KINSOL preconditioner module used (only if applicable, for examples using a Krylov linear solver and the KINBBDPRE module, this will be `bbd`);

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

The following lists summarize all examples distributed with KINSOL.

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

- `kinRoberts_fp` solves the backward Euler time step for a three-species chemical kinetics system, using the fixed point strategy.
- `kinFerTron_dns` solves the Ferraris-Tronconi problem.
This program solves the problem with the `SUNLINSOL_DENSE` linear solver and uses different combinations of globalization and Jacobian update strategies with different initial guesses.
- `kinFerTron_klu` solves the same problem as in `kinFerTron_dns`, but uses the sparse direct solver KLU via the `SUNLINSOL_KLU` module.
- `kinRoboKin_dns` solves a nonlinear system from robot kinematics.
This program solves the problem with the `SUNLINSOL_DENSE` linear solver and a user-supplied Jacobian routine.
- `kinRoboKin_slu` is the same as `kinRoboKin_dns` but uses the SuperLUMT sparse direct linear solver via the `SUNLINSOL_SUPERLUMT` module.
- `kinLaplace_bnd` solves a simple 2-D elliptic PDE on a unit square.
This program solves the problem with the `SUNLINSOL_BAND` linear solver.
- `kinLaplace_picard_bnd` is the same as `kinLaplace_bnd` but uses the Picard strategy.
- `kinLaplace_picard_kry` is the same as `kinLaplace_picard_bnd` but uses the SPGMR linear solver.

- `kinFoodWeb_kry` solves a food web model.
This is a nonlinear system that arises from a system of partial differential equations describing a six-species food web population model, with predator-prey interaction and diffusion on the unit square in two dimensions. This program solves the problem with the `SUNLINSOL_SPGMR` linear solver module and a user-supplied preconditioner. The preconditioner is a block-diagonal matrix based on the partial derivatives of the interaction terms only.
- `kinKrylovDemo_ls` solves the same problem as `kinFoodWeb_kry`, but with three Krylov linear solvers: `SUNLINSOL_SPGMR`, `SUNLINSOL_SPBCGS`, and `SUNLINSOL_SPTFQMR`.
- `kinAnalytic_fp.c` solves a small nonlinear system with known solution.

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

- `kinFoodWeb_kry_p` is a parallel implementation of `kinFoodWeb_kry`.
- `kinFoodWeb_kry_bbd_p` solves the same problem as `kinFoodWeb_kry_p`, with a block-diagonal matrix with banded blocks as a preconditioner, generated by difference quotients, using the `KINBBDPRE` module.

Supplied in the `srcdir/examples/kinsol/C_openmp` directory is an example `kinFoodweb_kry_omp`, which solves the same problem as `kinFoodweb_kry`, but using the OpenMP `NVECTOR` module.

Supplied in the `srcdir/examples/kinsol/CUDA_mpi` directory is an example `kin_em_mpicuda`, which solves an expectation-maximization problem for mixture densities, using the `cuda` `NVECTOR` module and MPI.

The following are examples written in C++:

- The directory `srcdir/examples/kinsol/CXX_parallel` contains two examples. `kin_heat2D_nonlin_p` solves a steady state 2D heat equation with an additional nonlinear term. This example is solved with fixed point iteration and illustrates the use of various orthogonalization methods with Anderson acceleration. `kin_em_p` solves the same problem as `kin_em_mpicuda` listed above.
- The directory `srcdir/examples/kinsol/CXX_parhyp` contains two examples, `kin_heat2D_nonlin_hypre_pf` and `kin_bratu2D_hyper_pfm`. These use the `hypre` PFMG preconditioner, fixed point iteration, and Anderson acceleration.

The following are examples in Fortran, using the `SUNDIALS` Fortran interface modules, and based on examples listed earlier:

- The directory `srcdir/examples/kinsol/F2003_serial` contains four examples: `kinDiagon_kry_f2003` solves a simple diagonal test problem. `kinLaplace_bnd_f2003` solves the same problem as `kinLaplace_bnd`. `kinLaplace_picard_kry_f2003` solves the same problem as `kinLaplace_picard_kry`. `kinRoboKin_dns_f2003` solves the same problem as `kinRoboKin_dns`.

- The directory `srdir/examples/kinsol/F2003_parallel` contains one example, `kin_diagon_kry_f2003`, which solves the same problem as `kinDiagon_kry_f2003`, but using MPI.

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 Newton iterations, may differ from one machine environment to another by as much as 10% to 20%.

In the descriptions below, we make frequent references to the KINSOL User Document [1]. All citations to specific sections (e.g. §??) are references to parts of that User Document, unless explicitly stated otherwise.

Note. The examples in the KINSOL distribution are written in such a way as to compile and run for any combination of configuration options used during the installation of SUNDIALS (see Appendix ?? 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.

2 C example problems

2.1 A serial dense example: `kinFerTron_dns`

As an initial illustration of the use of the KINSOL package for the solution of nonlinear systems, we give a sample program called `kinFerTron_dns.c`. It uses the KINSOL dense linear solver module `SUNLINSOL_DENSE` and the `NVECTOR_SERIAL` module (which provides a serial implementation of `NVECTOR`) for the solution of the Ferraris-Tronconi test problem [2].

This problem involves a blend of trigonometric and exponential terms:

$$\begin{aligned}
 0 &= 0.5 \sin(x_1 x_2) - 0.25 x_2 / \pi - 0.5 x_1 \\
 0 &= (1 - 0.25 / \pi)(e^{2x_1} - e) + e x_2 / \pi - 2 e x_1 \\
 &\text{subject to} \\
 x_{1 \min} &= 0.25 \leq x_1 \leq 1 = x_{1 \max} \\
 x_{2 \min} &= 1.5 \leq x_2 \leq 2\pi = x_{2 \max}.
 \end{aligned} \tag{1}$$

The bounds constraints on x_1 and x_2 are treated by introducing four additional variables and using KINSOL's optional constraints feature to enforce non-positivity and non-negativity:

$$\begin{aligned}
 l_1 &= x_1 - x_{1 \min} \geq 0 \\
 L_1 &= x_1 - x_{1 \max} \leq 0 \\
 l_2 &= x_2 - x_{2 \min} \geq 0 \\
 L_2 &= x_2 - x_{2 \max} \leq 0.
 \end{aligned}$$

The Ferraris-Tronconi problem has two known solutions. We solve it with KINSOL using two sets of initial guesses for x_1 and x_2 (first their lower bounds and secondly the middle of their feasible regions), both with an exact and modified Newton method, with and without line search.

Following the initial comment block, this program has a number of `#include` lines, which allow access to useful items in CVOID header files. The `kinsol.h` file provides prototypes for the KINSOL 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 `KINSOL`. 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 `sunmatrix_dense.h` file provides the prototype for the `SUNMatrixDense` function. The `sunlinsol_dense.h` file provides the prototype for the `SUNLinSolDense` function. The `sundials_types.h` file provides the definition of the type `sunrealtype` (see §?? for details). For now, it suffices to read `sunrealtype` as `double`. Finally, `sundials_math.h` is included for the definition of the exponential function `RExp`.

Next, the program defines some problem-specific constants, which are isolated to this early location to make it easy to change them as needed. This program includes a user-defined accessor macro, `Ith`, that is useful in writing the problem functions in a form closely matching the mathematical description of the 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 program prologue ends with prototypes of the user-supplied system function `func` and several private helper functions.

The `main` program begins with some dimensions and type declarations, including use of the type `N_Vector`, initializations, and allocation and definitions for the user data structure `data` which contains two arrays with lower and upper bounds for x_1 and x_2 . Next, we create five serial vectors of type `N_Vector` for the two different initial guesses, the solution vector `u`, the scaling factors, and the constraint specifications.

The initial guess vectors `u1` and `u2` are set by the private functions `SetInitialGuess1` and `SetInitialGuess2` and the constraint vector `c` is initialized to $[0, 0, 1, -1, 1, -1]$ indicating that there are no additional constraints on the first two components of `u` (i.e. x_1 and x_2) and that the 3rd and 5th components should be non-negative, while the 4th and 6th should be non-positive.

The calls to `N_VNew_Serial`, and also later calls to various `KIN***` 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 `KINCreate` creates the KINSOL solver memory block. 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 KINSOL functions.

The next four calls to KINSOL optional input functions specify the pointer to the user data structure (to be passed to all subsequent calls to `func`), the vector of additional constraints, and the function and scaled step tolerances, `fnormtol` and `scsteptol`, respectively.

Solver memory is allocated through the call to `KINInit` which specifies the system function `func` and provides the vector `u` which will be used internally as a template for cloning additional necessary vectors of the same type as `u`. The use of the dense linear solver is specified by calling `SUNDenseMatrix` to create the template Jacobian matrix (which also specifies the problem size `NEQ`), then calling `SUNLinSol_Dense` to create the dense-direct linear solver, and finally calling `KINSetLinearSolver` to attach these to KINSOL.

The main program proceeds by solving the nonlinear system eight times, using each of the two initial guesses `u1` and `u2` (which are first copied into the vector `u` using `N_VScale_Serial` from the `NVECTOR_SERIAL` module), with and without globalization through line search (specified by setting `glstr` to `KIN_LINESEARCH` and `KIN_NONE`, respectively), and applying either an exact or a modified Newton method. The switch from exact to modified Newton is done by changing the number of nonlinear iterations after which a Jacobian evaluation is enforced, a value `mset= 1` thus resulting in re-evaluating the Jacobian at every single iteration of the nonlinear solver (exact Newton method). Note that passing `mset= 0` indicates using the default KINSOL value of 10.

The actual problem solution is carried out in the private function `SolveIt` which calls the main solver function `KINSo1` after first setting the optional input `mset`. After a successful return from `KINSo1`, the solution $[x_1, x_2]$ and some solver statistics are printed.

The function `func` is a straightforward expression of the extended nonlinear system. It uses the macro `NV_DATA_S` (defined by the `NVECTOR_SERIAL` module) to extract the pointers to the data arrays of the `N_Vectors` `u` and `f` and sets the components of `fdata` using the current values for the components of `udata`. See §?? for a detailed specification of `f`.

The output generated by `kinFerTron_dns` is shown below.

kinFerTron_dns sample output

Ferraris and Tronconi test problem
Tolerance parameters:

```

fnormtol =      1e-05
scstseptol =     1e-05

-----

Initial guess on lower bounds
[x1,x2] =      0.25      1.5

Exact Newton
Solution:
[x1,x2] =  0.299449    2.83693
Final Statistics:
nni =      3    nfe =      4
nje =      3    nfeD =     18

Exact Newton with line search
Solution:
[x1,x2] =  0.299449    2.83693
Final Statistics:
nni =      3    nfe =      4
nje =      3    nfeD =     18

Modified Newton
Solution:
[x1,x2] =  0.299449    2.83693
Final Statistics:
nni =     11    nfe =     12
nje =      2    nfeD =     12

Modified Newton with line search
Solution:
[x1,x2] =  0.299449    2.83693
Final Statistics:
nni =     11    nfe =     12
nje =      2    nfeD =     12

-----

Initial guess in middle of feasible region
[x1,x2] =      0.625    3.89159

Exact Newton
Solution:
[x1,x2] =      0.5    3.14159
Final Statistics:
nni =      5    nfe =      6
nje =      5    nfeD =     30

Exact Newton with line search
Solution:
[x1,x2] =      0.5    3.14159
Final Statistics:
nni =      5    nfe =      6
nje =      5    nfeD =     30

Modified Newton
Solution:
[x1,x2] =  0.500003    3.1416
Final Statistics:

```

```

nni =    12    nfe =    13
nje =     2    nfeD =    12

Modified Newton with line search
Solution:
[x1,x2] =  0.500003    3.1416
Final Statistics:
nni =    12    nfe =    13
nje =     2    nfeD =    12

```

2.2 A serial Krylov example: kinFoodWeb_kry

We give here an example that illustrates the use of KINSOL with the Krylov method SPGMR, via the SUNLINSOL_SPGMR module, as the linear system solver.

This program solves a nonlinear system that arises from a discretized 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. Given the dependent variable vector of species concentrations $c = [c_1, c_2, \dots, c_{n_s}]^T$, where $n_s = 2n_p$ is the number of species and n_p is the number of predators and of prey, then the PDEs can be written as

$$d_i \cdot \left(\frac{\partial^2 c_i}{\partial x^2} + \frac{\partial^2 c_i}{\partial y^2} \right) + f_i(x, y, c) = 0 \quad (i = 1, \dots, n_s), \quad (2)$$

where the subscripts i are used to distinguish the species, and where

$$f_i(x, y, c) = c_i \cdot \left(b_i + \sum_{j=1}^{n_s} a_{i,j} \cdot c_j \right). \quad (3)$$

The problem coefficients are given by

$$a_{ij} = \begin{cases} -1 & i = j \\ -0.5 \cdot 10^{-6} & i \leq n_p, j > n_p \\ 10^4 & i > n_p, j \leq n_p \\ 0 & \text{all other} \end{cases},$$

$$b_i = b_i(x, y) = \begin{cases} 1 + \alpha xy & i \leq n_p \\ -1 - \alpha xy & i > n_p, \end{cases}$$

and

$$d_i = \begin{cases} 1 & i \leq n_p \\ 0.5 & i > n_p. \end{cases}$$

The spatial domain is the unit square $(x, y) \in [0, 1] \times [0, 1]$.

Homogeneous Neumann boundary conditions are imposed and the initial guess is constant in both x and y . For this example, the equations (2) are discretized spatially with standard central finite differences on a 8×8 mesh with $n_s = 6$, giving a system of size 384.

Among the initial `#include` lines in this case are lines to include `sunlinsol_spgmr.h` and `sundials_math.h`. The first contains constants and function prototypes associated with the SPGMR solver module. The inclusion of `sundials_math.h` is done to access the `SUNMAX` and

SUNRabs macros, and the SUNRsqrt function to compute the square root of a `sunrealtype` number.

The `main` program calls `KINCreate` and then calls `KINInit` with the name of the user-supplied system function `func` and solution vector as arguments. The `main` program then calls a number of `KINSet*` routines to notify KINSOL of the user data pointer, the positivity constraints on the solution, and convergence tolerances on the system function and step size. It calls `SUNLinSol_SPGMR` (see §??) to create the SPGMR linear solver module, supplying the `maxl` value of 15 as the maximum Krylov subspace dimension. It then calls `KINSetLinearSolver` to attach this solver module to KINSOL. Next, a maximum value of `maxlrst = 2` restarts is imposed through a call to `SUNLinSol_SPGMRSetMaxRestarts`. Finally, the user-supplied preconditioner setup and solve functions, `PrecSetupBD` and `PrecSolveBD`, are specified through a call to `KINSetPreconditioner` (see §??). The `data` pointer passed to `KINSetUserData` is passed to `PrecSetupBD` and `PrecSolveBD` whenever these are called.

Next, `KINsol` is called, the return value is tested for error conditions, and the approximate solution vector is printed via a call to `PrintOutput`. After that, `PrintFinalStats` is called to get and print final statistics, and memory is freed by calls to `N_VDestroy_Serial`, `FreeUserData` and `KINFree`. The statistics printed are the total numbers of nonlinear iterations (`nni`), of `func` evaluations (excluding those for Jv product evaluations) (`nfe`), of `func` evaluations for Jv evaluations (`nfeSG`), of linear (Krylov) iterations (`nli`), of preconditioner evaluations (`npe`), and of preconditioner solves (`nps`). All of these optional outputs and others are described in §??.

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 `IJ_Vptr` 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 384, 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 a serial `N_Vector` which is then passed to the `IJ_Vptr` macro.

The preconditioner used here is the block-diagonal part of the true Newton matrix and is based only on the partial derivatives of the interaction terms f in (3) and hence its diagonal blocks are $n_s \times n_s$ matrices ($n_s = 6$). It is generated and factored in the `PrecSetupBD` routine and backsolved in the `PrecSolveBD` routine. See §?? for detailed descriptions of these preconditioner functions.

The program `kinFoodWeb_kry.c` uses the “small” dense functions for all operations on the 6×6 preconditioner blocks. Thus it includes `sundials_dense.h`, and calls the small dense matrix functions `newDenseMat`, `denseGETRF`, and `denseGETRS`. The small dense functions are generally available for KINSOL user programs (for more information, see the comments in the header file `sundials_dense.h`).

In addition to the functions called by KINSOL, `kinFoodWeb_kry.c` includes definitions of several private functions. These are: `AllocUserData` to allocate space for the preconditioner 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 `cc`; `PrintHeader` to print the heading for the output; `PrintOutput` to retrieve and print selected solution values; `PrintFinalStats` to print statistics; and `check_flag` to check return values for error conditions.

The output generated by `kinFoodWeb_kry` is shown below. Note that the solution involved 10 Newton iterations, with an average of about 38 Krylov iterations per Newton iteration.

```

Predator-prey test problem -- KINSol (serial version)

Mesh dimensions = 8 X 8
Number of species = 6
Total system size = 384

Flag globalstrategy = 0 (0 = None, 1 = Linesearch)
Linear solver is SPGMR with maxl = 15, maxlrst = 2
Preconditioning uses interaction-only block-diagonal matrix
Positivity constraints imposed on all components
Tolerance parameters: fnormtol = 1e-07  scsteptol = 1e-13

Initial profile of concentration
At all mesh points:  1 1 1  30000 30000 30000

Computed equilibrium species concentrations:

At bottom left:
  1.16428 1.16428 1.16428 34927.5 34927.5 34927.5

At top right:
  1.25797 1.25797 1.25797 37736.7 37736.7 37736.7

Final Statistics..
nni    =      9      nli    =    329
nfe    =     10      nfeSG  =    338
nps    =    338      npe    =      1      ncfl  =      6

```

2.3 A parallel example: kinFoodWeb_kry_bbd_p

In this example, `kinFoodWeb_kry_bbd_p`, we solve the same problem as with `kinFoodWeb_kry` above, but in parallel, and instead of supplying the preconditioner we use the `KINBBDPRE` module.

In this case, we think of the parallel MPI processes as being laid out in a rectangle, and each process being assigned a subgrid of size $\text{MXSUB} \times \text{MYSUB}$ of the $x-y$ grid. If there are NPEX processes in the x direction and NPEY processes 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 nonlinear system is $\text{NUM_SPECIES} \cdot \text{MX} \cdot \text{MY}$.

The evaluation of the nonlinear system function is performed in `func`. In this parallel setting, the processes first communicate the subgrid boundary data and then compute the local components of the nonlinear system function. The MPI communication is isolated in the private function `ccomm` (which in turn calls `BRecvPost`, `BSend`, and `BRecvWait`) and the subgrid boundary data received from neighboring processes is loaded into the work array `cext`. The computation of the nonlinear system function is done in `func_local` which starts by copying the local segment of the `cc` vector into `cext`, and then by imposing the boundary conditions by copying the first interior mesh line from `cc` into `cext`. After this, the nonlinear system function is evaluated by using central finite-difference approximations using the data in `cext` exclusively.

`KINBBDPRE` uses a band-block-diagonal preconditioner, generated by difference quotients.

The upper and lower half-bandwidths of the Jacobian block generated on each process are both equal to $2 \cdot n_s - 1$, and that is the value passed as `mudq` and `mldq` in the call to `KINBBDPrecInit`. These values are much less than the true half-bandwidths of the Jacobian blocks, which are $n_s \cdot \text{MXSUB}$. However, an even narrower band matrix is retained as the preconditioner, with half-bandwidths equal to n_s , and this is the value passed to `KINBBDPrecInit` for `mukeep` and `mlkeep`.

The function `func_local` is also passed as the `gloc` argument to `KINBBDPrecInit`. Since all communication needed for the evaluation of the local approximation of f used in building the band-block-diagonal preconditioner is already done for the evaluation of f in `func`, a `NULL` pointer is passed as the `gcomm` argument to `KINBBDPrecInit`.

The `main` program resembles closely that of the `kinFoodWeb_kry` example, with particularization arising from the use of the parallel MPI `NVECTOR_PARALLEL` module. It begins by initializing MPI and obtaining the total number of processes and the rank of the local process. The local length of the solution vector is then computed as `NUM_SPECIES * MXSUB * MYSUB`. Distributed vectors are created by calling the constructor defined in `NVECTOR_PARALLEL` with the MPI communicator and the local and global problem sizes as arguments. All output is performed only from the process with `id` equal to 0. Finally, after `KINSol` is called and the results are printed, all memory is deallocated, and the MPI environment is terminated by calling `MPI_Finalize`.

The output generated by `kinFoodWeb_kry_bbd_p` is shown below. Note that 9 Newton iterations were required, with an average of about 51.6 Krylov iterations per Newton iteration.

```

kinFoodWeb_kry_bbd_p sample output

Predator-prey test problem--  KINSol (parallel-BBD version)

Mesh dimensions = 20 X 20
Number of species = 6
Total system size = 2400

Subgrid dimensions = 10 X 10
Processor array is 2 X 2

Flag globalstrategy = 0 (0 = None, 1 = Linesearch)
Linear solver is SPGMR with maxl = 20, maxlrst = 2
Preconditioning uses band-block-diagonal matrix from KINBBDPRE
  Difference quotient half-bandwidths: mudq = 11, mldq = 11
  Retained band block half-bandwidths: mukeep = 6, mlkeep = 6
Tolerance parameters:  fnormtol = 1e-07   scsteptol = 1e-13

Initial profile of concentration
At all mesh points:  1 1 1   30000 30000 30000

Computed equilibrium species concentrations:

At bottom left:
  1.165 1.165 1.165 34949 34949 34949

At top right:
  1.25552 1.25552 1.25552 37663.2 37663.2 37663.2

Final Statistics..
nni      =      9      nli      =    464

```


nfe	=	10	nfeSG	=	473		
nps	=	473	npe	=	1	ncfl	= 6

3 C++ example problems

3.1 A parallel matrix-free example: `kin_heat2D_nonlin_p`

As an illustration of the use of the KINSOL package for the solution of nonlinear systems in parallel, we give a sample program called `kin_heat2D_nonlin_p.cpp`. It uses the KINSOL KINFP iteration with Anderson Acceleration and the NVECTOR_PARALLEL module (which provides a parallel implementation of NVECTOR) for the solution of the following test problem. This example highlights the use of the various orthogonalization routine options within Anderson Acceleration, passed to the example problem via the `--orthaa` flag. Available options include 0 (KIN_ORTH_MGS), 1 (KIN_ORTH_ICWY), 2 (KIN_ORTH_CGS2), and 3 (KIN_ORTH_DCGS2).

This problem involves solving a steady-state 2D heat equation with an additional nonlinear term defined by $c(u)$:

$$b = \nabla \cdot (D \nabla u) + c(u) \quad \text{in } \mathcal{D} = [0, 1] \times [0, 1] \quad (4)$$

where D is a diagonal matrix with entries k_x and k_y for the diffusivity in the x and y directions respectively. The boundary conditions are

$$u(0, y) = u(1, y) = u(x, 0) = u(x, 1) = 0. \quad (5)$$

We chose the analytical solution to be

$$u_{\text{exact}} = u(x, y) = \sin^2(\pi x) \sin^2(\pi y), \quad (6)$$

hence, we define the static term b as follows

$$\begin{aligned} b = & k_x 2\pi^2 (\cos^2(\pi x) - \sin^2(\pi x)) \sin^2(\pi y) \\ & + k_y 2\pi^2 (\cos^2(\pi y) - \sin^2(\pi y)) \sin^2(\pi x) + c(u_{\text{exact}}) \end{aligned} \quad (7)$$

The spatial derivatives are computed using second-order centered differences, with the data distributed over $n_x \times n_y$ points on a uniform spatial grid. The problem is set up to use spatial grid parameters $n_x = 64$ and $n_y = 64$, with heat conductivity parameters $k_x = 1.0$ and $k_y = 1.0$.

This problem is solved via a fixed point iteration with Anderson acceleration, where the fixed point is function formed by adding u to both sides of the equation i.e.,

$$b + u = \nabla \cdot (D \nabla u) + c(u) + u, \quad (8)$$

so that the fixed point function is

$$G(u) = \nabla \cdot (D \nabla u) + c(u) + u - b. \quad (9)$$

The problem is run using a tolerance of 10^{-8} , and a starting vector containing all ones. The following tables contains all available input parameters when running the example problem.

Table 1: Optional input parameter flags.

Flag	Description
--mesh <nx> <ny>	mesh points in the x and y directions
--np <npx> <npy>	number of MPI processes in the x and y directions
--domain <xu> <yu>	domain upper bound in the x and y direction
--k <kx> <ky>	diffusion coefficients
--rtol <rtol>	relative tolerance
--maa <maa>	number of previous residuals for Anderson Acceleration
--damping <damping>	damping parameter for Anderson Acceleration
--orthaa <orthaa>	orthogonalization routine used in Anderson Acceleration
--maxits <maxits>	max number of iterations
--c <cu>	nonlinear function choice (integer between 1 - 17)
--timing	print timing data
--help	print available input parameters and exit

Table 2: Input parameter flags for setting the nonlinear function $c(u)$.

Flag	Function
--c 1	$c(u) = u$
--c 2	$c(u) = u^3 - u$
--c 3	$c(u) = u - u^2$
--c 4	$c(u) = e^u$
--c 5	$c(u) = u^4$
--c 6	$c(u) = \cos^2(u) - \sin^2(u)$
--c 7	$c(u) = \cos^2(u) - \sin^2(u) - e^u$
--c 8	$c(u) = e^u u^4 - u e^{\cos(u)}$
--c 9	$c(u) = e^{(\cos^2(u))}$
--c 10	$c(u) = 10(u - u^2)$
--c 11	$c(u) = -13 + u + ((5 - u)u - 2)u$
--c 12	$c(u) = \sqrt{5}(u - u^2)$
--c 13	$c(u) = (u - e^u)^2 + (u + u \sin(u) - \cos(u))^2$
--c 14	$c(u) = u + u e^u + u e^{-u}$
--c 15	$c(u) = u + u e^u + u e^{-u} + (u - e^u)^2$
--c 16	$c(u) = u + u e^u + u e^{-u} + (u - e^u)^2 + (u + u \sin(u) - \cos(u))^2$
--c 17	$c(u) = u + u e^{-u} + e^u(u + \sin(u) - \cos(u))^3$

3.2 A parallel example using *hypre*: `kin_heat2D_nonlin_hypr_pfm`

As an illustration of the use of the KINSOL package for the solution of nonlinear systems in parallel and using *hypre* linear solvers, we give a sample program called `kin_heat2D_nonlin_hypr_pfm.cpp`. It uses the KINSOL KINFP iteration and the NVECTOR_PARALLEL module (which provides a parallel implementation of NVECTOR) for the solution of the following test problem. This example highlights the use of the various orthogonalization routine options within Anderson Acceleration, passed to the example problem via the `--orthaa` flag. Available options include 0 (KIN_ORTH_MGS), 1 (KIN_ORTH_ICWY), 2 (KIN_ORTH_CGS2), and 3 (KIN_ORTH_DCGS2).

This problem involves solving a steady-state 2D heat equation with an additional nonlinear term defined by $c(u)$:

$$b = \nabla \cdot (D \nabla u) + c(u) \quad \text{in } \mathcal{D} = [0, 1] \times [0, 1] \quad (10)$$

where D is a diagonal matrix with entries k_x and k_y for the diffusivity in the x and y directions respectively. The boundary conditions are

$$u(0, y) = u(1, y) = u(x, 0) = u(x, 1) = 0 \quad (11)$$

We chose the analytical solution to be

$$u_{\text{exact}} = u(x, y) = \sin^2(\pi x) \sin^2(\pi y), \quad (12)$$

hence, we define the static term b as follows

$$\begin{aligned} b = & k_x 2\pi^2 (\cos^2(\pi x) - \sin^2(\pi x)) \sin^2(\pi y) \\ & + k_y 2\pi^2 (\cos^2(\pi y) - \sin^2(\pi y)) \sin^2(\pi x) + c(u_{\text{exact}}) \end{aligned} \quad (13)$$

The spatial derivatives are computed using second-order centered differences, with the data distributed over $n_x \times n_y$ points on a uniform spatial grid. The problem is set up to use spatial grid parameters $n_x = 64$ and $n_y = 64$, with heat conductivity parameters $k_x = 1.0$ and $k_y = 1.0$.

This problem is solved via a fixed point iteration with Anderson acceleration, where the fixed point function formed by implementing the Laplacian as a matrix-vector product,

$$b = Au + c(u), \quad (14)$$

and solving for u results to get the fixed point function

$$G(u) = A^{-1}(b - c(u)). \quad (15)$$

The problem is run using a tolerance of 10^{-8} , and a starting vector containing all ones. The linear system solve is executed using the SUNLINSOL_PCG linear solver with the *hypr* PFMG preconditioner. The setup of the linear solver can be found in the `Setup_LS` function, and setup of the *hypr* preconditioner can be found in the `Setup_Hypr` function within the main file.

All input parameter flags available for Example 3.1 are also available for this problem. In addition, all runtime flags controlling the linear solver and *hypr* related parameters are set using the flags in the following table.

Table 3: Optional input parameter flags for setting *hypr* related parameters.

Flag	Description
<code>--lsinfo</code>	output residual history for PCG
<code>--lin_iters <lin_iters></code>	max number of iterations for PCG
<code>--eps_lin <eps_lin></code>	linear tolerance for PCG
<code>--pfmg_relax <pfmg_relax></code>	relaxation type in PFMG
<code>--pfmg_nrelax <pfmg_nrelax></code>	pre/post relaxation sweeps in PFMG

References

- [1] A. M. Collier, A. C. Hindmarsh, R. Serban, and C.S. Woodward. User Documentation for KINSOL v7.5.0. Technical Report UCRL-SM-208116, LLNL, 2025.
- [2] C. Floudas, P. Pardalos, C. Adjiman, W. Esposito, Z. Gumus, S. Harding, J. Klepeis, C. Meyer, and C. Schweiger. *Handbook of Test Problems in Local and Global Optimization*. Kluwer Academic Publishers, Dordrecht, 1999.