

Hands-On Exercises for SLEPc

Apr 2023
Intended for use with version 3.19 of
SLEPc

These exercises introduce application development in SLEPc, the Scalable Library for Eigenvalue Problem Computations. A basic knowledge of PETSc is assumed. The first exercise (exercise 0) is just a guided tour on how to get started compiling and running programs. Exercises 1, 2, 3, and 4 are intended to cover most of the basic SLEPc functionality. The rest of the exercises illustrate more advanced features. All the example programs used in the exercises are included in the SLEPc distribution and are also available at its web site.

- Exercise 0:** [Hello World](#)
- Exercise 1:** [Standard Symmetric Eigenvalue Problem](#)
- Exercise 2:** [Standard Non-Symmetric Eigenvalue Problem](#)
- Exercise 3:** [Generalized Eigenvalue Problem Stored in a File](#)
- Exercise 4:** [Singular Value Decomposition](#)
- Exercise 5:** [Problem without Explicit Matrix Storage](#)
- Exercise 6:** [Parallel Execution](#)
- Exercise 7:** [Use of Deflation Subspaces](#)
- Exercise 8:** [Quadratic Eigenvalue Problem](#)

For reference, detailed information on usage of SLEPc and PETSc may be found at the following links:

- [SLEPc on-line documentation](#)
- [PETSc on-line documentation](#)

Exercise 0: Hello World

This exercise shows how to build and run a simple example program with SLEPc.

Note: The description below related to directories and the use of the PETSC_ARCH variable will be different in the case of a prefix-based installation.

Compiling

SLEPc needs the following environment variables to be set:

SLEPC_DIR - the location of SLEPc
PETSC_DIR - the location of PETSc
PETSC_ARCH - the architecture being used

Make sure that you have them correctly set.

Like in PETSc, a makefile is necessary to compile a SLEPc program. Paste this simple example into a file named `makefile` in your working directory:

```
hello: hello.o
    -${CLINKER} -o hello hello.o ${SLEPC_SYS_LIB}
    ${RM} hello.o
```

```
include ${SLEPC_DIR}/lib/slepc/conf/slepc_common
```

Note: In the above text, the blank space in the 2nd and 3rd lines represents a tab.

Also place the following source code into a file named "hello.c" in the same directory:

```
static char help[] = "Simple Hello World example program in SLEPc\n";

#include

int main( int argc, char **argv )
{
    PetscFunctionBeginUser;
    PetscCall(SlepcInitialize(&argc,&argv,(char*)0,help));
    PetscCall(PetscPrintf(PETSC_COMM_WORLD,"Hello world\n"));
    PetscCall(SlepcFinalize());
    return 0;
}
```

Compile the program with the following command:

```
$ make hello
```

Source Code Details

Examine the source code of the sample program. The following comments will help you understand the code thoroughly.

Header File: All SLEPc programs must include a header file with all the necessary definitions. In this simple example the file `slepcsys.h` (base SLEPc header) is enough because no solver components such as EPS are being used.

Note: SLEPc header files automatically include some PETSc header files.

Library Initialization: All SLEPc programs must begin with a call to `SlepcInitialize`, which in turn initializes PETSc and MPI. Similarly, at the end of the program `SlepcFinalize` must be called for library cleanup.

Input/Output: In this example, we do input/output via a call to a PETSc function, `PetscPrintf`. Remember that in parallel programs input/output cannot be done simply via C standard library functions. Note that in SLEPc programs we can freely use any PETSc function.

Error Checking: All SLEPc routines return an integer indicating whether an error has occurred during the call. The PETSc macro `PetscCall` checks the return value and calls the PETSc error handler upon error detection. All function calls should be wrapped around `PetscCall` to enable a complete error traceback.

Running the Program

SLEPc programs are executed as any other MPI program. Note that this typically differs from one system to another. To run the program with only one process, in some systems you can launch it as a normal program:

```
$ ./hello
```

but in other systems this would not work. Check the documentation of your system. Standard MPI implementations provide the `mpiexec` command to launch the applications

```
$ mpiexec -n 4 ./hello
```

In SLEPc (and PETSc) there are a lot of options (run-time parameters) to control program behavior. These

options are usually equivalent to function calls, so the user can test its effect without changing the source code (this will be illustrated in the next exercises). To show which options are available in a program use:

```
$ ./hello -help
```

Support for Debugging and Complex Numbers

The support for debugging capabilities, complex scalar arithmetic, and other features is managed by SLEPc and PETSc by means of different *architectures*, represented by different values of the PETSC_ARCH variable. In a given system, you can typically find several versions of SLEPc and PETSc, each of them built with different configuration options. For instance, suppose the following values are available:

- arch-linux-gnu-c-opt: built with compiler optimization
- arch-linux-gnu-c-debug: built with debugging support
- arch-linux-gnu-c-opt-complex: optimized with complex scalars
- arch-linux-gnu-c-debug-complex: debug with complex scalars

Note: In order to learn about the particular architectures available in your system, type `ls $SLEPC_DIR`. There should be a subdirectory for each allowed value of PETSC_ARCH.

When using an architecture with support for complex scalars, all scalar values are complex instead of real. Try compiling the example program for complex numbers:

```
$ make PETSC_ARCH=arch-linux-gnu-c-opt-complex hello
```

When using the debug versions some options are available to support debugging. For example

```
$ ./hello -start_in_debugger
```

opens the program in a debugger stopped at the `SlepInitialize` function.

Other useful options are: `-info` to get informative messages about progress of the calculations, `-malloc_info` to print memory usage at end of run, `-log_trace [filename]` to get a full trace of the execution (in a file), `-malloc_dump` to list memory blocks not freed at the end of the program, and `-log_view` to get a summary including performance results.

Exercise 1: Standard Symmetric Eigenvalue Problem

This example solves a standard symmetric eigenproblem $Ax=\lambda x$, where A is the matrix resulting from the discretization of the Laplacian operator in 1 dimension by centered finite differences.

$$A = \begin{array}{c|cccccc|} & 2 & -1 & 0 & 0 & 0 & 0 & | \\ & -1 & 2 & -1 & 0 & 0 & 0 & | \\ & 0 & -1 & 2 & -1 & 0 & 0 & | \\ & 0 & 0 & -1 & 2 & -1 & 0 & | \\ & 0 & 0 & 0 & -1 & 2 & -1 & | \\ & 0 & 0 & 0 & 0 & -1 & 2 & | \end{array}$$

Compiling

Copy the file `ex1.c` [plain text] to the working directory and add these lines to the makefile

```
ex1: ex1.o
    -${CLINKER} -o ex1 ex1.o ${SLEPC_EPS_LIB}
    ${RM} ex1.o
```

Note: In the above text, the blank space in the 2nd and 3rd lines represents a tab.

Build the executable with the command

```
$ make ex1
```

Note for Fortran users: Example ex1 is also available in Fortran [ex1f.F \[plain text\]](#).

Running the Program

In order to run the program for a problem of order 50, type the following

```
$ ./ex1 -n 50
```

You should get an output similar to this

```
1-D Laplacian Eigenproblem, n=50

Number of iterations of the method: 8
Solution method: krylovschur

Number of requested eigenvalues: 1
Stopping condition: tol=1e-08, maxit=100
Number of converged eigenpairs: 3

      k          ||Ax-kx||/||kx||
-----
 3.996207      4.30363e-10
 3.984841      2.08631e-09
 3.965946      9.98404e-09
```

Source Code Details

Examine the source code of the sample program and locate the function calls mentioned in the following comments.

The Options Database: All the PETSc functionality related to the options database is available in SLEPc. This allows the user to input control data at run time very easily. In this example, the function `PetscOptionsGetInt` is used to check whether the user has provided a command line option to set the value of `n`, the problem dimension. If so, the variable `n` is set accordingly; otherwise, `n` remains unchanged.

Vectors and Matrices: Usage of matrices and vectors in SLEPc is exactly the same as in PETSc. The user can create a new parallel or sequential matrix, `A`, with subroutine `MatCreate`, where the matrix format can be specified at runtime. The example creates a matrix, sets the nonzero values with `MatSetValues` and then assembles it.

Solving the Eigenvalue Problem: Usage of eigensolvers is very similar to other kinds of solvers provided by PETSc. After creating the matrix, the problem is solved by means of an EPS object (Eigenvalue Problem Solver) via the following sequence of function calls:

```
EPSCreate(MPI_Comm comm, EPS *eps);
EPSSetOperators(EPS eps, Mat A, Mat B);
EPSSetProblemType(EPS eps, EPSProblemType type);
EPSSetFromOptions(EPS eps);
EPSSolve(EPS eps);
EPSGetConverged(EPS eps, int *nconv);
EPSGetEigenpair(EPS eps, int i, PetscScalar *kr, PetscScalar *ki, Vec xr, Vec xi);
EPSTDestroy(EPS eps);
```

First, the eigenproblem solver (EPS) context is created and the operator(s) associated with the

eigensystem are set, as well as the problem type. Then various options are set for customized solution. After that, the program solves the problem, retrieves the solution, and finally destroys the EPS context.

The above function calls are very important and will be present in most SLEPc programs. In the example source code `ex1.c` you will find other functions apart from these. What do they do?

Playing with EPS Options

Now we are going to experiment with different options of the EPS object. A full list of command-line options can be obtained by running the example with the option `-help`.

To show information about the solver object:

```
$ ./ex1 -eps_view
```

Note: This option internally calls the function `EPSView`. Alternatively, we could include a direct call to this function in the source code. Almost all command-line options have a related function call.

Note: All the command-line options related to the EPS object have the `-eps_` prefix.

This time, your output will include something like this

```
EPS Object: 1 MPI processes
  type: krylovschur
    50% of basis vectors kept after restart
    using the locking variant
  problem type: symmetric eigenvalue problem
  selected portion of the spectrum: largest eigenvalues in magnitude
  number of eigenvalues (nev): 1
  number of column vectors (ncv): 16
  maximum dimension of projected problem (mpd): 16
  maximum number of iterations: 100
  tolerance: 1e-08
  convergence test: relative to the eigenvalue
BV Object: 1 MPI processes
  type: svec
  17 columns of global length 30
  vector orthogonalization method: classical Gram-Schmidt
  orthogonalization refinement: if needed (eta: 0.7071)
  block orthogonalization method: GS
  doing matmult as a single matrix-matrix product
DS Object: 1 MPI processes
  type: hep
  solving the problem with: Implicit QR method (_steqr)
ST Object: 1 MPI processes
  type: shift
  shift: 0
  number of matrices: 1
```

This option is very useful to see which solver and options the program is using.

Try solving a much larger problem, for instance with $n=400$. Note that in that case the program does not return a solution. This means that the solver has reached the maximum number of allowed iterations and the convergence criterion was not satisfied. What we can do is either increase the number of iterations or relax the convergence criterion.

```
$ ./ex1 -n 400 -eps_max_it 400
$ ./ex1 -n 400 -eps_tol 1e-3
```

Note that in the latter case the relative error displayed by the program is significantly larger, meaning that the solution has only 3 correct decimal digits, as expected.

It is possible to change the number of requested eigenvalues. Try the following execution

```
$ ./ex1 -n 400 -eps_nev 3 -eps_tol 1e-7
```

In this case, the program did not succeed to compute two of the requested eigenpairs. This is again due to the convergence criterion, which is satisfied by some eigenpairs but not for all. As in the previous case, we could increase further the number of iterations or relax the convergence criterion. Another alternative is to increase the number of column vectors (i.e., the dimension of the subspace with which the eigensolver works). This usually improves the convergence behavior at the expense of larger memory requirements.

```
$ ./ex1 -n 400 -eps_nev 3 -eps_ncv 24
```

Note that the default value of `ncv` depends on the value of `nev`.

Try to set some of the above options directly in the source code by calling the related functions `EPSSetTolerances` and `EPSSetDimensions`. Modify and recompile the program. Use `-eps_view` to check that the values are correctly set. Is it now possible to change these options from the command-line? Does this change whether you place the calls before or after the call to `EPSSetFromOptions`?

Convergence is usually bad when eigenvalues are close to each other, which is the case in this example. In order to see what is happening while the eigensolver iterates, we can use a monitor to display information associated to the convergence of eigenpairs at each iteration:

```
$ ./ex1 -eps_monitor
```

or

```
$ ./ex1 -eps_monitor_all
```

Also, in some SLEPc installations, it is possible to monitor convergence graphically with the `draw` viewer with `draw_lg` format. For example, try this:

```
$ ./ex1 -n 700 -eps_nev 5 -eps_ncv 35 -eps_monitor_all draw::draw_lg -draw_pause .1
```

Note: The plot is drawn in an X11 pop-up window. So this requires that the display is correctly exported.

Changing the Eigensolver

The convergence behavior for a particular problem also depends on the properties of the eigensolver being used. SLEPc provides several eigensolvers which can be selected in the source code with the function `EPSSetType`, or at run time:

```
$ ./ex1 -eps_nev 4 -eps_type lobpcg
```

The following table shows some of the solvers available in SLEPc.

Solver	Command-line Name	Parameter
Krylov-Schur	krylovschur	EPSKRYLOVSCHUR
Generalized Davidson	gd	EPSGD
Jacobi-Davidson	jd	EPSJD
Rayleigh-quotient conjugate gradient	rqcg	EPSRQCG
Locally optimal block preconditioned CG	lobpcg	EPSLOBPCG
Contour integral spectrum slice	ciss	EPSCISS
Lanczos with Explicit Restart	lanczos	EPSLANCZOS

Arnoldi with Explicit Restart	arnoldi	EPSARNOLDI
Subspace Iteration	subspace	EPSSUBSPACE
Power / RQI	power	EPSPOWER
Lapack	lapack	EPSLAPACK
ARPACK	arpack	EPSARPACK

Note: The Lapack solver is not really a full-featured eigensolver but simply an interface to some LAPACK routines. These routines operate sequentially in dense mode and therefore are suitable only for small size problems. This solver should be used only for debugging purposes.

Note: The last one (ARPACK) may or may not be available on your system, depending on whether it was enabled during installation of SLEPc. It consists in an interface to the external [ARPACK library](#). Interfaces to other external libraries may be available as well. These can be used as any other SLEPc native eigensolver.

Note: The default solver is krylovschur for both symmetric and non-symmetric problems.

Exercise 2: Standard Non-Symmetric Eigenvalue Problem

In this exercise we are going to work with a non-symmetric problem. The example solves the eigenvalue problem associated with a Markov model of a random walk on a triangular grid. Although the matrix is non-symmetric, all eigenvalues are real. Eigenvalues come in pairs with the same magnitude and different signs. The values 1 and -1 are eigenvalues for any matrix size. More details about this problem can be found at [Matrix Market](#).

Compiling

Copy the file [ex5.c \[plain text\]](#) to your directory and add these lines to the makefile

```
ex5: ex5.o
    -${CLINKER} -o ex5 ex5.o ${SLEPC_EPS_LIB}
    ${RM} ex5.o
```

Note: In the above text, the blank space in the 2nd and 3rd lines represents a tab.

Build the executable with the command

```
$ make ex5
```

Source Code Details

The example program is very similar to that in Exercise 1. The main difference is that the problem is set to be non-symmetric with [EPSSetProblemType](#):

```
PetscCall(EPSSetProblemType(eps, EPS_NHEP));
```

In this example we also illustrate the use of [EPSSetInitialSpace](#).

Running the Program

Run the program requesting four eigenvalues.

```
$ ./ex5 -eps_nev 4
```

The output will look like this:

```
Markov Model, N=120 (m=15)

Solution method: krylovschur

Number of requested eigenvalues: 4
Linear eigensolve converged (4 eigenpairs) due to CONVERGED_TOL; iterations 5
-----
      k                ||Ax-kx||/||kx||
-----
    -1.000000          7.05086e-10
     1.000000          1.47517e-09
    -0.971367          1.04527e-10
     0.971367          2.23781e-10
-----
```

You can see that the solver returns both positive and negative eigenvalues. This is because largest magnitude eigenvalues are computed by default, that is, internally the solver sorts the eigenvalue approximations according to $|\lambda|$, and the same criterion is used for sorting the finally computed eigenvalues.

Other criteria can be used, see [EPSSetWhichEigenpairs](#) for details. For instance, for computing only the rightmost eigenvalues, try the following.

```
$ ./ex5 -eps_nev 4 -eps_largest_real
```

Similarly, it is possible to request the smallest magnitude eigenvalues with `-eps_smallest_magnitude`. The difference in that case is that the solver needs much more iterations to converge. The justification is that in this problem the smallest magnitude eigenvalues are located in the interior of the spectrum, and computing interior eigenvalues is always harder as explained next.

Computing Interior Eigenvalues

It is well known that computing eigenvalues located at the interior of the spectrum is much more difficult than those in the periphery. We are going to discuss different strategies available in SLEPc.

The general way of computing interior eigenvalues is to specify a target value, around which the eigenvalues must be sought.

```
$ ./ex5 -eps_nev 4 -eps_target 0.75
```

Note that apart from the target value, one should specify a sorting criterion relative to the target (`-eps_target_magnitude`). However, this option can be omitted because it is the default when a target is specified. The output is in this case:

```
Markov Model, N=120 (m=15)

Solution method: krylovschur

Number of requested eigenvalues: 4
Linear eigensolve converged (6 eigenpairs) due to CONVERGED_TOL; iterations 14
-----
      k                ||Ax-kx||/||kx||
-----
     0.771298          1.13746e-09
     0.714286          9.26568e-10
     0.704038          2.49204e-09
     0.702317          8.51587e-10
-----
```

```

0.800066      2.80221e-14
0.660191      7.61205e-09
0.840634      1.94667e-14
0.857143      1.67828e-10
-----

```

We have obtained eigenvalues both on the left and on the right of the target value $\tau=0.75$, and they are sorted according to the distance to τ .

The number of iterations is higher than in the default case. The theory says that Krylov methods (and other methods as well) approximate eigenvalues from the periphery to the interior, meaning that before getting eigenvalues closest to 0.75 the solver has to find out the eigenvalues from 0.75 to the rightmost extreme. If we choose a target close to the extreme the number of iterations will be small, and they will increase as τ is moved inside of the spectrum. Therefore, this is not a good strategy because it will not be viable for difficult problems.

Sometimes, an improvement may come from changing the way in which the method extracts the spectral information from the built subspace; see [EPSSetExtraction](#) for details. One such technique is called harmonic extraction. Try the following:

```
$ ./ex5 -eps_nev 4 -eps_target 0.75 -eps_harmonic
```

In this simple problem, harmonic extraction gives no benefit but in difficult problems it may be a significant improvement, especially in combination with preconditioned solvers (discussed later below).

A better solution may be to use a spectral transformation, but with several considerations to take into account regarding cost.

Getting Started with Spectral Transformations

The general idea of the spectral transformation is to substitute the original problem, $Ax=\lambda x$, by another one, $Tx=\theta x$, in which the eigenvalues are mapped to a different position but eigenvectors remain unchanged. With this strategy, one can move interior eigenvalues to the periphery.

Each EPS object uses an ST object internally to manage the spectral transformation. The following table shows the available spectral transformations, which can be selected with the function [STSetType](#) or at run time.

Spectral Transform	Operator	Command-line Name	Parameter
Shift of Origin	$A-\sigma I$	shift	STSHIFT
Shift-and-invert	$(A-\sigma I)^{-1}$	sinvert	STSINVERT
Cayley	$(A-\sigma I)^{-1}(A+\nu I)$	cayley	STCAYLEY
Preconditioner	$K^{-1} \approx (A-\sigma I)^{-1}$	precond	STPRECOND
Polynomial filter	$p(A)$	filter	STFILTER

Note: The default is to do shift of origin with a value $\sigma=0$. This was reported by `-eps_view` in the previous example.

Note: The preconditioner is not really a spectral transformation like the rest. It will be discussed later below.

The shift-and-invert spectral transformation can be used for computing interior eigenvalues:

```
$ ./ex5 -eps_nev 4 -eps_target 0.75 -st_type sinvert
```

Note: All the command-line options related to the ST object have the `-st_` prefix.

With the above execution, the number of iterations is very small, but each iteration is much more costly than in the previous cases because linear systems must be solved to handle the inverted operator (the issue of how to solve linear systems is discussed below). The value of the parameter σ (the shift) is taken to be equal to τ (the target). Run with `-eps_view` to check that it is indeed the case.

Try also with `cayley`, which is nearly equivalent.

Handling the Inverses

In the table of spectral transformations shown above, there are some operators that include the inverse of a certain matrix. These operators are not computed explicitly in order to preserve sparsity. Instead, in the ST object the multiplication by these inverses is replaced by a linear equation solve via a KSP object from PETSc.

SLEPc allows us to pass options to this KSP linear solver object. For instance,

```
$ ./ex5 -eps_nev 4 -eps_target 0.75 -st_type sinvert -st_ksp_type preonly
      -st_pc_type lu
```

Note: In order to specify a command-line option related to the linear solver contained in ST, simply add the `-st_` prefix in front.

The options of the above example specify a direct linear solver (LU factorization). This is what SLEPc does by default. This strategy is usually called *exact shift-and-invert*. Its main drawback is that direct solvers are more costly in terms of flops and storage and are less parallelizable.

An alternative is to do an *inexact shift-and-invert*, that is, to use an iterative linear solver. The following line illustrates how to use an iterative solver

```
$ ./ex5 -eps_nev 4 -eps_target 0.75 -st_type sinvert -st_ksp_type gmres
      -st_pc_type bjacobi -st_ksp_rtol 1e-12
```

Iterative linear solvers may fail to converge if the coefficient matrix is ill-conditioned or close to singular. Also, the accuracy of the eigensolver may be compromised if the iterative linear solver provides a solution far from full working precision.

Note that in SLEPc it is extremely easy to switch between exact and inexact schemes.

Preconditioned Eigensolvers

As mentioned above, the inexact shift-and-invert scheme is very sensitive to the accuracy with which the linear systems are solved. This usually implies using a very stringent tolerance (10^{-12} in the example) and makes it impractical for difficult problems.

An alternative is to use a preconditioned eigensolver, such as those of Davidson type: EPSGD and EPSJD. These solvers try to emulate the idea of shift-and-invert but they are very robust with respect to bad accuracy (i.e., large tolerance) of the iterative linear solve.

Here is an example where Jacobi-Davidson is used:

```
$ ./ex5 -eps_nev 4 -eps_target 0.75 -eps_type jd -st_type precondition
      -st_ksp_type bcgsl -st_pc_type bjacobi -st_ksp_rtol 0.001
```

Note: The `-st_type precondition` key can be omitted in this case, since it is the default in all preconditioned eigensolvers.

Try adding `-eps_harmonic` to the above example. As mentioned before, harmonic extraction is usually better when used in preconditioned solvers.

Exercise 3: Generalized Eigenvalue Problem Stored in a File

In this exercise we are going to work with a generalized eigenvalue problem, $Ax=\lambda Bx$. The example program loads two matrices A and B from a file and then solves the associated eigensystem.

The matrices we are going to work with are BFW62A and BFW62B, which are available at [Matrix Market](#). This particular problem is non-symmetric. Some of the eigenvalues (those of largest magnitude) come in complex conjugate pairs while the rest are real.

Compiling

Copy the file [ex7.c \[plain text\]](#) to your directory and add these lines to the makefile

```
ex7: ex7.o
    -${CLINKER} -o ex7 ex7.o ${SLEPC_EPS_LIB}
    ${RM} ex7.o
```

Note: In the above text, the blank space in the 2nd and 3rd lines represents a tab.

Build the executable with the command

```
$ make ex7
```

Source Code Details

This example uses the PETSc function [MatLoad](#) to load a matrix from a file. The two matrix files are specified in the command line. Note that these files have been converted from Matrix Market format to PETSc binary format.

Compare the source code of the example program with the previous ones. Note that, in this case, two matrix objects are passed in the [EPSSetOperators](#) function call:

```
PetscCall(EPSSetOperators(eps,A,B));
```

Running the Program

Run the program with the following command line:

```
$ ./ex7 -f1 ${SLEPC_DIR}/share/slepc/datafiles/matrices/bfw62a.petsc
    -f2 ${SLEPC_DIR}/share/slepc/datafiles/matrices/bfw62b.petsc
```

Run the program to compute more than one eigenpair. Use the following option to plot the computed eigenvalues:

```
$ ./ex7 -f1 ${SLEPC_DIR}/share/slepc/datafiles/matrices/bfw62a.petsc
    -f2 ${SLEPC_DIR}/share/slepc/datafiles/matrices/bfw62b.petsc
    -eps_type subspace -eps_nev 6 -eps_view_values draw -draw_pause -1
```

Note: The plot is drawn in an X11 pop-up window. So this requires that the display is correctly exported.

Spectral Transformations in Generalized Problems

The following table shows the expressions of the operator in each of the available spectral transformations in the case of generalized problems. Note that both matrices A and B are involved.

Spectral Transform	Operator
Shift of Origin	$B^{-1}A + \sigma I$
Shift-and-invert	$(A - \sigma B)^{-1}B$
Cayley	$(A - \sigma B)^{-1}(A + \nu B)$
Preconditioner	$K^{-1} \approx (A - \sigma B)^{-1}$

In the case of generalized problems, the shift-and-invert transformation does not represent a cost penalty with respect to the simpler shift of origin, since in both cases the inverse of a matrix is required.

```
$ ./ex7 -f1 ${SLEPC_DIR}/share/slepc/datafiles/matrices/bfw62a.petsc
      -f2 ${SLEPC_DIR}/share/slepc/datafiles/matrices/bfw62b.petsc
      -eps_target 0 -st_type sinvert
```

The above execution computes the eigenvalues closest to the origin. Use a target near the left end of the spectrum to compute the largest magnitude eigenvalues

```
$ ./ex7 -f1 ${SLEPC_DIR}/share/slepc/datafiles/matrices/bfw62a.petsc
      -f2 ${SLEPC_DIR}/share/slepc/datafiles/matrices/bfw62b.petsc
      -eps_target -250000 -st_type sinvert
```

Preconditioned Eigensolvers

As hinted above, generalized eigenproblems have the drawback that in the default mode (shift of origin) one of the matrices have to be (implicitly) inverted. However, preconditioned eigensolvers do not have this limitation, and may be able to solve the problem with just a preconditioner or a few iterations of an iterative linear solver.

Here is an example with Generalized Davidson:

```
$ ./ex7 -f1 ${SLEPC_DIR}/share/slepc/datafiles/matrices/bfw62a.petsc
      -f2 ${SLEPC_DIR}/share/slepc/datafiles/matrices/bfw62b.petsc
      -eps_type gd -eps_nev 6
```

Try the above example also with `-eps_target` and `-eps_harmonic`.

Exercise 4: Singular Value Decomposition

In this exercise we turn our attention to the Singular Value Decomposition (SVD). Remember that in real symmetric (or complex Hermitian) matrices, singular values coincide with eigenvalues, but in general this is not the case. The SVD is defined for any matrix, even rectangular ones. Singular values are always non-negative real values.

This example works also by reading a matrix from a file. In particular, the matrix to be used is related to a 2D reaction-diffusion model. More details about this problem can be found at [Matrix Market](#).

Compiling

Copy the file `ex14.c` [\[plain text\]](#) to your directory and add these lines to the makefile

```
ex14: ex14.o
      -${CLINKER} -o ex14 ex14.o ${SLEPC_SVD_LIB}
      ${RM} ex14.o
```

Note: In the above text, the blank space in the 2nd and 3rd lines represents a tab.

Build the executable with the command

```
$ make ex14
```

Running the Program

In order to run the program, type the following

```
$ ./ex14 -file $SLEPC_DIR/share/slepc/datafiles/matrices/rdb200.petsc
```

You should get an output similar to this

```
Singular value problem stored in file.

Reading REAL matrix from a binary file...
Number of iterations of the method: 3
Solution method: cross

Number of requested singular values: 1
Stopping condition: tol=1e-08, maxit=100
SVD solve converged (1 singular triplet) due to CONVERGED_TOL; iterations 3
-----
      sigma          relative error
-----
    35.007519      1.40809e-10
-----
```

Source Code Details

The way in which the SVD object works is very similar to that of EPS. However, some important differences exist. Examine the source code of the example program and pay attention to the differences with respect to EPS. After loading the matrix, the problem is solved by the following sequence of function calls:

```
SVDCreate(MPI_Comm comm,SVD *svd);
SVDSetOperators(SVD svd,Mat A,Mat B);
SVDSetFromOptions(SVD svd);
SVDsolve(SVD svd);
SVDGetConverged(SVD svd, int *nconv);
SVDGetSingularTriplet(SVD svd,int i,PetscReal *sigma,Vec u,Vec v);
SVDDestroy(SVD svd);
```

First, the singular value solver (SVD) context is created and the matrix associated with the problem is specified. Then various options are set for customized solution. After that, the program solves the problem, retrieves the solution, and finally destroys the SVD context.

Note that the singular value, `sigma`, is defined as a `PetscReal`, and that the singular vectors are simple `Vec`'s.

SVD Options

Most of the options available in the EPS object have their equivalent in SVD. A full list of command-line options can be obtained by running the example with the option `-help`.

To show information about the SVD solver, add the `-svd_view` option:

```
$ ./ex14 -file $SLEPC_DIR/share/slepc/datafiles/matrices/rdb200.petsc -svd_view
```

Note: All the command-line options related to the SVD object have the `--svd_` prefix.

Your output will include something like this

```
SVD Object: 1 MPI processes
  type: cross
    implicit matrix
  EPS Object: (svd_cross_) 1 MPI processes
    type: krylovschur
      50% of basis vectors kept after restart
      using the locking variant
    problem type: symmetric eigenvalue problem
    selected portion of the spectrum: largest real parts
    number of eigenvalues (nev): 1
    number of column vectors (ncv): 16
    maximum dimension of projected problem (mpd): 16
    maximum number of iterations: 100
    tolerance: 1e-09
    convergence test: relative to the eigenvalue
  BV Object: (svd_cross_) 1 MPI processes
    type: svec
    17 columns of global length 200
    vector orthogonalization method: classical Gram-Schmidt
    orthogonalization refinement: if needed (eta: 0.7071)
    block orthogonalization method: GS
    doing matmult as a single matrix-matrix product
  DS Object: (svd_cross_) 1 MPI processes
    type: hep
    solving the problem with: Implicit QR method (_steqr)
  ST Object: (svd_cross_) 1 MPI processes
    type: shift
    shift: 0.
    number of matrices: 1
    using a shell matrix
  transpose mode: explicit
  selected portion of the spectrum: largest
  number of singular values (nsv): 1
  number of column vectors (ncv): 16
  maximum dimension of projected problem (mpd): 16
  maximum number of iterations: 100
  tolerance: 1e-08
  convergence test: relative to the singular value
```

The output shows all the options that are susceptible of being changed, either from the command line or from the source code of the program: the method, the portion of the spectrum (largest or smallest singular values), the number of singular values (nsv), etc.

Try to change some of the values, for instance:

```
$ ./ex14 -file $SLEPC_DIR/share/slepc/datafiles/matrices/rdb200.petsc
         -svd_nsv 10 -svd_ncv 40 -svd_smallest
```

The "transpose mode" refers to whether the transpose of matrix A is being built explicitly or not (see [SVDSetImplicitTranspose](#) for an explanation).

Note that in the sample output above, the SVD object contains an EPS object. This only happens in some SVD solver types, as detailed below.

Changing the Singular Value Solver

SLEPc provides several solvers for computing the SVD, which can be selected in the source code with the function [SVDSetType](#), or at run time:

```
$ ./ex14 -file $SLEPC_DIR/share/slepc/datafiles/matrices/rdb200.petsc
```

```
-svd_type trlanczos
```

The following table shows the list of SVD solvers available in SLEPc.

Solver	Command-line Name	Parameter
Cross Product	cross	SVDCROSS
Cyclic Matrix	cyclic	SVDCYCLIC
Lanczos with Explicit Restart	lanczos	SVDLANCZOS
Lanczos with Thick Restart	trlanczos	SVDTRLANCZOS
Lapack	lapack	SVDLAPACK

Note: The Lapack solver is not really a full-featured singular value solver but simply an interface to some LAPACK routines. These routines operate sequentially in dense mode and therefore are suitable only for small size problems. This solver should be used only for debugging purposes.

Note: The default solver is `cross`.

The first two solvers, `cross` and `cyclic`, are not real methods implemented in the SVD module, but are two convenient ways of solving the SVD problem by making use of the eigensolvers available in the EPS module. In those two cases, the SVD object manages an EPS object internally, whose parameters can be set as desired (typically only the method). For example:

```
$ ./ex14 -file $SLEPC_DIR/share/slepc/datafiles/matrices/rdb200.petsc
-svd_type cyclic -svd_cyclic_eps_type lanczos -svd_cyclic_eps_lanczos_reorthog local
```

Exercise 5: Problem without Explicit Matrix Storage

In many applications, it may be better to keep the matrix (or matrices) that define the eigenvalue problem implicit, that is, without storing its nonzero entries explicitly. An example is when we have a matrix-vector routine available. SLEPc allows easy management of this case. This exercise tries to illustrate it by solving a standard symmetric eigenproblem corresponding to the Laplacian operator in 2 dimensions in which the matrix is not built explicitly.

Compiling

Copy the file `ex3.c` [plain text] to your directory and add these lines to the makefile

```
ex3: ex3.o
    -${CLINKER} -o ex3 ex3.o ${SLEPC_EPS_LIB}
    ${RM} ex3.o
```

Note: In the above text, the blank space in the 2nd and 3rd lines represents a tab.

Build the executable with the command

```
$ make ex3
```

Source Code Details

PETSc provides support for matrix-free problems via the *shell* matrix type. This kind of matrices is created with a call to `MatCreateShell`, and their operations are specified with `MatShellSetOperation`. For basic use of these matrices with EPS solvers only the matrix-vector product operation is required. In the

example, this operation is performed by a separate function `MatMult_Laplacian2D`.

Running the Program

Run the program without any spectral transformation options. For instance:

```
$ ./ex3 -eps_type subspace -eps_tol 1e-9 -eps_nev 8
```

Now try running the program with shift-and-invert to get the eigenvalues closest to the origin

```
$ ./ex3 -eps_target 0.0 -st_type sinvert
```

Note that the above command yields a run-time error. Observe the information printed on the screen and try to deduce the reason of the error. In this case, the error is due to the fact that SLEPc tries to use a direct linear solver within the ST object, and this is not possible unless the matrix has been created explicitly as in previous examples.

There are more chances to have success if an inexact shift-and-invert scheme is used. Try using an iterative linear solver without preconditioning:

```
$ ./ex3 -eps_target 0.0 -st_type sinvert -st_ksp_rtol 1e-10
    -st_ksp_type gmres -st_pc_type none
```

You can also try with a nonzero target.

This example is slow, because the iterative linear solver takes a lot of iterations to reach the required precision (add `-st_ksp_converged_reason` to monitor the convergence of the linear solver). In order to alleviate this problem, a preconditioner should be used. However, a powerful preconditioner such as ILU cannot be used in this case, for the same reason a direct solver is not available. The only possibility is to use a simple preconditioner such as Jacobi. Try running the last example again with `-st_pc_type jacobi`. Note that this works because the `MATOP_GET_DIAGONAL` has been defined in our program. In this particular example the Jacobi preconditioner does not help reduce the linear iterations since the matrix is not strongly diagonally dominant.

Some of the above difficulties can be avoided by using a preconditioned eigensolver, as in the examples shown in previous exercises.

Exercise 6: Parallel Execution

The objective of this exercise is to run an example program with different number of processors to see how execution time is reduced. This time, we are going to solve a standard eigensystem $Ax=\lambda x$ with the matrix loaded from a file. In particular, the matrix we are going to use is QC2534. It is a complex matrix of order 2534 arising from a quantum chemistry application (more details can be found at [Matrix Market](#)).

Compiling

Copy the file `ex4.c` [\[plain text\]](#) to your directory and add these lines to the makefile

```
ex4: ex4.o
    -${CLINKER} -o ex4 ex4.o ${SLEPC_EPS_LIB}
    ${RM} ex4.o
```

Note: In the above text, the blank space in the 2nd and 3rd lines represents a tab.

Build the executable with the command (optimized complex version)

```
$ make PETSC_ARCH=arch-linux-gnu-c-opt-complex ex4
```

Source Code Details

This example program is very similar to that of exercise 3. It uses the PETSc function `MatLoad` to load a matrix from a file. The matrix file is specified in the command line.

Running the Program

In order to run this example, you will need the file `qc2534.petsc`. Locate it in the file system and then run the program with the command

```
$ ./ex4 -file qc2534.petsc
```

For execution with more than one process:

```
$ mpiexec -n 2 ./ex4 -file qc2534.petsc
```

Check the output of the program. It should be the same as with one process.

Try using the `-log_view` option to have a look at the profiling information collected by PETSc and SLEPc. For instance, check the size and number of MPI messages.

```
$ mpiexec -n 2 ./ex4 -file qc2534.petsc -log_view
```

Try to find out how much time was spent is solving the eigenvalue problem. Is there a significant reduction when we increase the number of processors?

Instrumenting the Source Code

If we are just interested in knowing the time used by the eigensolver, then it may be better to let our example program inform us. With the function `PetscTime`, it is possible to obtain the current time of day (wall-clock time) in seconds. Edit the source code and add two calls to this function just before and after the `EPSSolve` call, as in the following fragment of code

```
PetscCall(PetscTime(&t1));
PetscCall(EPSSolve(eps));
PetscCall(PetscTime(&t2));
PetscCall(PetscPrintf(PETSC_COMM_WORLD," Elapsed Time: %f\n",t2-t1));
```

Also you must add the definition of the two new variables

```
PetscLogDouble t1,t2;
```

as well as the header

```
#include <petsctime.h>
```

Recompile the program with

```
$ make PETSC_ARCH=arch-linux-gnu-c-opt-complex ex4
```

Run it with one, two and four processors checking the time spent by the solver

```
$ mpiexec -n 1 ./ex4 -file qc2534.petsc
$ mpiexec -n 2 ./ex4 -file qc2534.petsc
$ mpiexec -n 4 ./ex4 -file qc2534.petsc
```

Exercise 7: Use of Deflation Subspaces

The term deflation refers to the use of the knowledge of one or more eigenpairs to find other eigenpairs. For instance, most eigensolvers try to approximate a number of eigenpairs and, as soon as one of them has converged, they deflate it for better approximating the other ones. Another case is when one eigenpair is known a priori and one wants to use this knowledge to compute other eigenpairs. SLEPc supports this by means of *deflation subspaces*.

This example illustrates the use of deflation subspaces to compute the smallest nonzero eigenvalue of the Laplacian of a graph corresponding to a 2-D regular mesh. The problem is a standard symmetric eigenproblem $Ax = \lambda x$, where $A = L(G)$ is the Laplacian of graph G , defined as follows: $A_{ii} = \text{degree of node } i$, $A_{ij} = -1$ if edge (i,j) exists in G , zero otherwise. This matrix is symmetric positive semidefinite and singular, and $[1 \ 1 \ \dots \ 1]^T$ is the eigenvector associated with the zero eigenvalue. In graph theory, one is usually interested in computing the eigenvector associated with the next eigenvalue (the so-called Fiedler vector).

Compiling

Copy the file `ex11.c` [\[plain text\]](#) to your directory and add these lines to the makefile

```
ex11: ex11.o
    -${CLINKER} -o ex11 ex11.o ${SLEPC_EPS_LIB}
    ${RM} ex11.o
```

Note: In the above text, the blank space in the 2nd and 3rd lines represents a tab.

Build the executable with the command

```
$ make ex11
```

Source Code Details

This example computes the smallest eigenvalue by setting `EPS_SMALLEST_REAL` in `EPSSetWhichEigenpairs`. An alternative would be to use a shift-and-invert spectral transformation with a zero target to compute the eigenvalues closest to the origin, or to use harmonic extraction with a zero target.

By specifying a deflation subspace (the one associated to the eigenvector $[1 \ 1 \ \dots \ 1]^T$) with the function `EPSSetDeflationSpace`, the convergence to the zero eigenvalue is avoided. Thus, the program should compute the smallest nonzero eigenvalues.

Running the Program

Run the program simply with

```
$ ./ex11
```

For the case of using an inexact spectral transformation, the command line would be:

```
$ ./ex11 -eps_target 0.0 -eps_target_real -st_type sinvert
    -st_ksp_rtol 1e-10 -st_ksp_type gmres -st_pc_type jacobi
```

Note that a shift-and-invert spectral transformation should always be used in combination with EPS_TARGET_MAGNITUDE or EPS_TARGET_REAL.

And for the case of harmonic extraction:

```
$ ./ex11 -eps_target 0.0 -eps_target_real -eps_harmonic
```

Exercise 8: Quadratic Eigenvalue Problem

Now we are going to focus on the solution of quadratic eigenvalue problems with PEP solvers. In this case, the problem to be solved is formulated as $(\lambda^2 M + \lambda C + K)x = 0$. In our simple example, M is a diagonal matrix, C is tridiagonal, and K is the 2-D Laplacian.

Compiling

Copy the file `ex16.c` [plain text] to your directory and add these lines to the makefile

```
ex16: ex16.o
    -${CLINKER} -o ex16 ex16.o ${SLEPC_PEP_LIB}
    ${RM} ex16.o
```

Note: In the above text, the blank space in the 2nd and 3rd lines represents a tab.

Build the executable with the command

```
$ make ex16
```

Running the Program

Run the program without arguments and see the output:

```
Quadratic Eigenproblem, N=100 (10x10 grid)
Number of requested eigenvalues: 1
      k          ||P(k)x||/||kx||
-----
-1.164037+1.653625i  2.02021e-12
-1.164037-1.653625i  2.02021e-12
```

Source Code Details

The PEP object is used very much like EPS or SVD, as can be seen in the source code. Here is a summary of the main function calls:

```
PEPCreate(MPI_Comm comm, PEP *pep);
PEPSetOperators(PEP pep, PetscInt nmat, Mat A[]);
PEPSetProblemType(PEP pep, PEPProblemType type);
PEPSetFromOptions(PEP pep);
PEPSolve(PEP pep);
PEPGetConverged(PEP pep, int *nconv);
PEPGetEigenpair(PEP pep, int i, PetscScalar *kr, PetscScalar *ki, Vec xr, Vec xi);
PEPDestroy(PEP pep);
```

First, the solver context (PEP) is created and the three problem matrices are specified. Then various

options are set for customized solution. After that, the program solves the problem, retrieves the solution, and finally destroys the context.

PEP Options

Most of the options available in the EPS object have their equivalent in PEP. A full list of command-line options can be obtained by running the example with the option `-help`.

To show information about the PEP solver, add the `-pep_view` option:

```
PEP Object: 1 MPI processes
  type: toar
    50% of basis vectors kept after restart
    using the locking variant
  problem type: symmetric polynomial eigenvalue problem
  polynomial represented in MONOMIAL basis
  selected portion of the spectrum: largest eigenvalues in magnitude
  number of eigenvalues (nev): 1
  number of column vectors (ncv): 16
  maximum dimension of projected problem (mpd): 16
  maximum number of iterations: 100
  tolerance: 1e-08
  convergence test: relative to the eigenvalue
  extraction type: NORM
BV Object: 1 MPI processes
  type: svec
  18 columns of global length 100
  vector orthogonalization method: classical Gram-Schmidt
  orthogonalization refinement: if needed (eta: 0.7071)
  block orthogonalization method: GS
  doing matmult as a single matrix-matrix product
DS Object: 1 MPI processes
  type: nhsp
ST Object: 1 MPI processes
  type: shift
  shift: 0.
  number of matrices: 3
  all matrices have different nonzero pattern
KSP Object: (st_) 1 MPI processes
  type: preonly
  maximum iterations=10000, initial guess is zero
  tolerances: relative=1e-08, absolute=1e-50, divergence=10000.
  left preconditioning
  using NONE norm type for convergence test
PC Object: (st_) 1 MPI processes
  type: lu
  out-of-place factorization
  tolerance for zero pivot 2.22045e-14
  matrix ordering: nd
  factor fill ratio given 5., needed 1.
  Factored matrix follows:
    Mat Object: 1 MPI processes
      type: seqaij
      rows=100, cols=100
      package used to perform factorization: petsc
      total: nonzeros=100, allocated nonzeros=100
      total number of mallocs used during MatSetValues calls =0
      not using I-node routines
  linear system matrix = precond matrix:
  Mat Object: 1 MPI processes
    type: seqaij
    rows=100, cols=100
    total: nonzeros=100, allocated nonzeros=500
    total number of mallocs used during MatSetValues calls =0
    not using I-node routines
```

Note: All the command-line options related to the PEP object have the `-pep_` prefix.

Try changing some of the values, for example:

```
$ ./ex16 -pep_nev 4 -pep_ncv 24 -pep_smallest_magnitude -pep_tol 1e-5
```

Choosing the Solver Method

Several polynomial eigensolvers are available, which can be selected in the source code with the function `PEPSetType`, or at run time:

```
$ ./ex16 -pep_type qarnoldi
```

The following table shows the list of PEP solvers available in SLEPc.

Solver	Command-line Name	Parameter
Linearization	linear	PEPLINEAR
Quadratic Arnoldi	qarnoldi	PEPQARNOLDI
Two-level orthogonal Arnoldi	toar	PEPTOAR
Symmetric TOAR	stoar	PEPSTOAR
Jacobi-Davidson	jd	PEPJD
Contour integral	ciss	PEPCISS

Note: The default solver is `toar`.

The linear solver performs an explicit linearization of the quadratic eigenproblem, resulting in a generalized eigenproblem. This linearization can be customized, see the Users Manual for details. For instance:

```
$ ./ex16 -pep_type linear -pep_linear_linearization 1,0 -pep_linear_explicitmatrix
```

Since in this problem all matrices are symmetric, the problem type is set to `PEP_HERMITIAN` in the source code with `PEPSetProblemType`, and this obliges us to set the explicit matrix flag, see `PEPLinearSetExplicitMatrix`. It is also possible to use a non-symmetric linearization by choosing the corresponding problem type:

```
$ ./ex16 -pep_type linear -pep_general
```

In the linear solver it is also possible to tune any of the EPS options, including those corresponding to ST and the linear solvers. For instance:

```
$ ./ex16 -pep_type linear -pep_general -pep_linear_st_ksp_type bcgs
          -pep_linear_st_pc_type jacobi
```