



# Solving PDEs with PETSc

William Gropp

University of Illinois at Urbana-Champaign

# Introduction

- What and why is PETSc?
  - PETSc is a portable library for solving linear and nonlinear systems of equations in parallel
  - PETSc was originally designed to provide a library for experimentation in domain decomposition algorithms
- What is Domain Decomposition?
  - DD is a algorithmic technique for dividing problems into subproblems and combining the results to solve (or approximate) the solution
  - DD is a natural method for effective parallel algorithms for distributed memory computers

# Overview

- Introduction to PETSc—Hello World
- Building a Poisson Solver in PETSc
  - Using distributed arrays to describe data parallelism
  - Using domain decomposition methods in PETSc
- Solving Nonlinear problems
  - Algorithms for nonlinear problems
  - Bratu example
  - More on distributed arrays in PETSc

# A Few Comments Before We Start

- PETSc is a very large library
  - This tutorial is designed to introduce PETSc without overwhelming you with information
  - Many features will not be covered. PETSc comes with extensive examples and documentation
- PETSc is a freely available and supported research code
  - Available via `http://www.mcs.anl.gov/petsc`
  - Free for everyone, including industrial users
  - Hyperlinked documentation and manual pages for all routines
  - Many tutorial-style examples
  - Support via email: `petsc-maint@mcs.anl.gov`
  - Usable from Fortran 77/90, C, and C++

- Portable to any parallel system supporting MPI, including
  - Tightly coupled systems  
Cray XTn, IBM BlueGene, SGI Altix, SiCortex, Sun, ...
  - Loosely coupled systems, e.g., networks of workstations  
HP, IBM, SGI, Sun, and PCs running Linux or Windows, and Apple Macs
- What is not in PETSc
  - Discretizations
  - Unstructured mesh generation or refinement
  - Load balancing tools
  - Sophisticated visualization support
  - (But PETSc provides ways to interface to other tools)

# A First PETSc Program

- What do PETSc programs look like?
- What do PETSc parallel programs look like?
- How to compile, link, and run PETSc programs?

# Hello World

```
#include "petsc.h"

int main( int argc, char *argv[] )
{
    PetscInitialize( &argc, &argv, 0, 0 );

    PetscPrintf( PETSC_COMM_WORLD, "Hello World\n" );
    PetscFinalize( );
    return 0;
}
```

# Understanding the Code

**PetscInitialize** Initialize PETSc. The arguments allow PETSc to initialize MPI if necessary

**PetscFinalize** Finalize PETSc. Causes PETSc to call `MPI_Finalize` if necessary and also to generate summary reports.

**PetscPrintf** Ensures that only one process prints the data (Try it!)



# Hello World in Fortran

```
integer ierr, rank
#include "include/finclude/petsc.h"
call PetscInitialize( PETSC_NULL_CHARACTER, ierr )
call MPI_Comm_rank( PETSC_COMM_WORLD, rank, ierr )
if (rank .eq. 0) then
    print *, 'Hello World'
endif
call PetscFinalize(ierr)
end
```

# Understanding the Code

- Like the C code, except
  - PetscInitialize has fewer arguments because Fortran has no `argc` or `argv`
  - Must use `MPI_Comm_rank` and `print` because Fortran I/O uses a interface unavailable to libraries
- PETSc 2.1.6 adds a routine that can be used with a single character string (Fortran can't implement its own I/O operations, so PETSc can't provide parallel replacements)

# A Parallel Program

- PETSc uses the distributed memory, shared-nothing model
- Parallel PETSc programs consist of separate communicating processes
- PETSc uses MPI for parallelism
  - You can always access MPI routines
  - You will rarely need to use MPI while using PETSc
  - Many PETSc routines are *collective* in the MPI sense (all processes must call); others are local.
  - Common uses of MPI in PETSc are the routines for communicator size and rank and for processor name.
  - This is illustrated in a revised (and obviously parallel) hello world program.

# Hello World Revisited

```
#include "petsc.h"

int main( int argc, char *argv[] )
{
    int rank;
    PetscInitialize( &argc, &argv, 0, 0 );

    MPI_Comm_rank( PETSC_COMM_WORLD, &rank );
    PetscSynchronizedPrintf( PETSC_COMM_WORLD,
                            "Hello World from rank %d\n", rank );
    PetscSynchronizedFlush( PETSC_COMM_WORLD );
    PetscFinalize( );
    return 0;
}
```

# Understanding the Program

**PetscSynchronizedPrintf** Like `PetscPrintf`, except output comes from all processes in rank order.

**PetscSynchronizedFlush** Indicates that the calling process is done printing.

- Allows the use of multiple `PetscSynchronizedPrintf` calls

**PETSC\_COMM\_WORLD** The PETSc version of `MPI_COMM_WORLD`, they are usually the same set of processes. `PetscSetCommWorld`, used *before* `PetscInitialize`, may be used to give PETSc a subset of processes

# PETSc and PDEs

- PETSc is designed around the mathematics of the problem
  - Specify the data in terms of vectors
  - Specify the problem as linear (using matrices) or nonlinear (using vector-valued functions) equations to be solved
  - Support parallel computing by automatically distributing these objects across all processes
- We'll see a sequence of increasingly sophisticated PDE examples...

# Poisson Problem

Lets solve a simple linear elliptic PDE

$$\begin{aligned}\nabla^2 u &= f \text{ in } [0, 1] \times [0, 1] \\ u &= 0 \text{ on the boundary}\end{aligned}$$

using a simple discretization ( $u_{i,j} = u(x_i, y_j)$ ,  $x_i = ih$ )

$$\begin{aligned}\frac{u_{i+1,j} - 2u_{i,j} + u_{i-1,j}}{h^2} &+ \\ \frac{u_{i,j+1} - 2u_{i,j} + u_{i,j-1}}{h^2} &= f(x_i, y_j).\end{aligned}$$

(We use finite differences for simplicity; finite elements can be used as well.) For simplicity, consider  $f = \sin(\pi x)\sin(\pi y)$ .

We will discretize the interior of the mesh only for this example.

# Schematic for Example

In PETSc, your main program remains in control:

main program

PetscInitialize()

A = create the matrix

b = create a vector

any other application code

Use KSP to solve  $Ax = b$

print solution

PetscFinalize()

KSP is the “scalable linear equation solver” component of PETSc (the name is historical and was originally SLES, and the first “S” was for “simple”)



# Creating the Matrix

```
1  #include "petscksp.h"
2
3  /* Form the matrix for the 5-point finite difference 2d Laplacian
4     on the unit square. n is the number of interior points along a si
5  Mat FormLaplacian2d( int n )
6  {
7     Mat      A;
8     int      r, rowStart, rowEnd, i, j;
9     double h, oneByh2;
10
11     h = 1.0 / (n + 1);      oneByh2 = 1.0 / (h * h);
12     MatCreate( PETSC_COMM_WORLD, &A );
13     MatSetSizes( A, PETSC_DECIDE, PETSC_DECIDE, n*n, n*n );
14     MatSetFromOptions( A );
15     MatGetOwnershipRange( A, &rowStart, &rowEnd );
```

# Creating the Matrix II

```
16     /* This is a simple but inefficient way to set the matrix */
17     for (r=rowStart; r<rowEnd; r++) {
18         i = r % n;  j = r / n;
19         if (j - 1 > 0) {
20             MatSetValue( A, r, r - n, oneByh2, INSERT_VALUES ); }
21         if (i - 1 > 0) {
22             MatSetValue( A, r, r - 1, oneByh2, INSERT_VALUES ); }
23         MatSetValue( A, r, r, -4*oneByh2, INSERT_VALUES );
24         if (i + 1 < n - 1) {
25             MatSetValue( A, r, r + 1, oneByh2, INSERT_VALUES ); }
26         if (j + 1 < n - 1) {
27             MatSetValue( A, r, r + n, oneByh2, INSERT_VALUES ); }
28     }
29     MatAssemblyBegin(A, MAT_FINAL_ASSEMBLY);
30     MatAssemblyEnd(A, MAT_FINAL_ASSEMBLY);
31     return A;
32 }
33
```

# Understanding the Code I

**MatCreate** Create a matrix object.

**MatSetSizes** Set the size of the matrix

- $n^2$  equations, so matrix is of size  $n \times n \times n$
- PETSC\_DECIDE tells PETSc to choose the distribution of the matrix across the processes

**MatSetFromOptions** Set basic matrix properties (such as data structure) from command line

**MatGetOwnershipRange** Get the rows of the matrix that PETSc assigned to this process

- PETSc uses a simple assignment of consecutive rows to a process. This simplifies much of the internal structure of PETSc, and, as we shall see, does not reduce the generality
- It is not necessary to set values on the “owning” process
- Returns first row to one + last row on process.
  - Matches common C idiom (`for (i=start; i<end; i++)`)
  - Number of rows is `end-start`

# Understanding the Code II

**MatSetValue** Insert (or optionally add with `ADD_VALUES`) a value to a matrix (*Warning*: This is a macro and needs braces)

**MatAssemblyBegin** and **MatAssemblyEnd** Complete the creation of matrix. The matrix may not be used for any operation (other than `MatSetValue`) until after `MatAssemblyEnd`.

The approach of separating setting values from assembly has several benefits

- Any process may set a value to *any* element of the matrix, even ones not “owned” by the calling process.
- PETSc manages all data communication between processes, *automatically* moving data if necessary
- PETSc can optimize the insertion of matrix elements

# Data Structure Neutral Design

- PETSc matrices are *objects for storing linear operators*
- They allow many types of data structures:
  - Default sparse format MATMPIAIJ and MATSEQAIJ
  - Block sparse MATMPIBAIJ and MATSEQBAIJ
  - Symmetric block sparse MATMPISBAIJ and MATSEQSBAIJ
  - Block diagonal MATMPIBDIAG and MATSEQBDIAG
  - Dense MATMPIDENSE and MATSEQDENSE
  - Many others (see `$PETSC_DIR/include/petscmat.h`)
- Choice of format is made from command line (with `MatSetFromOptions`) or program (with `MatSetType`)
- The *same* routines are used for all choices of data structure
- User-defined data-structures supported with “Shell” objects

# Data Decomposition in PETSc

- How are objects distributed among processes in PETSc?
  - Contiguous rows of a vector or matrix are assigned to processes, starting from the process with rank zero
- The matrix and vector for a  $3 \times 3$  mesh, with two processes, has the following decomposition

$$\begin{array}{c} \text{P0} \\ \hline \text{P1} \end{array} \begin{pmatrix} x_0 \\ x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \\ x_7 \\ x_8 \end{pmatrix} = \begin{pmatrix} 4 & -1 & & -1 & & & & & \\ -1 & 4 & -1 & & -1 & & & & \\ & -1 & 4 & & & -1 & & & \\ -1 & & & 4 & -1 & & -1 & & \\ & -1 & & -1 & 4 & -1 & & -1 & \\ \hline & & -1 & & -1 & 4 & & & -1 \\ & & & -1 & & & 4 & -1 & \\ & & & -1 & & -1 & 4 & -1 & \\ & & & & -1 & & -1 & 4 & \end{pmatrix}$$

# Why Are PETSc Matrices The Way They Are?

- No one data structure is appropriate for all problems
  - Blocked and diagonal formats provide significant performance benefits
  - PETSc provides a large selection of formats and makes it (relatively) easy to extend PETSc by adding new data structures
- Matrix assembly is difficult enough without being forced to worry about data partitioning
  - PETSc provide parallel assembly routines
  - Achieving high performance still requires making most operations local to a process, but this approach allows incremental development of programs
- Matrix decomposition by consecutive rows across processes is simple and makes it easier to work with other codes
  - For applications with other ordering needs, PETSc provides “Application Orderings” (AO)

# Vectors In PETSc

- In order to support the distributed memory “shared nothing” model, as well as single processors and shared memory systems, a PETSc vector is a “handle” to the real vector
  - Allows the vector to be distributed across many processes
  - To access the elements of the vector, we cannot simply do

```
for (i=0; i<n; i++) v[i] = i;
```
  - We do not want to require that the programmer work *only* with the “local” part of the vector; we want to permit operations, such as setting an element of a vector, to be performed by any process.
- The solution is to make vectors an object, just like a parallel matrix



# Creating the Vectors I

```
1  #include "petscvec.h"
2
3  /* Form a vector based on a function for a 2-d regular mesh on the
4     unit square */
5  Vec FormVecFromFunction2d( int n, double (*f)( double, double ) )
6  {
7     Vec      V;
8     int      r, rowStart, rowEnd, i, j;
9     double h;
10
11     h = 1.0 / (n + 1);
12     VecCreate( PETSC_COMM_WORLD, &V );
13     VecSetSizes( V, PETSC_DECIDE, n*n );
14     VecSetFromOptions( V );
```

# Creating the Vectors II

```
15     VecGetOwnershipRange( V, &rowStart, &rowEnd );
16     /* This is a simple but inefficient way to set the vector */
17     for (r=rowStart; r<rowEnd; r++) {
18         i = (r % n) + 1;
19         j = (r / n) + 1;
20         VecSetValue( V, r, (*f)( i * h, j * h ), INSERT_VALUES );
21     }
22     VecAssemblyBegin(V);
23     VecAssemblyEnd(V);
24
25     return V;
26 }
27
```

# Understanding the Code

**VecCreate** Creates the vector.

**VecSetSizes** Sets the global and local size of the vector. Use `PETSC_DECIDE` to have PETSc choose the distribution across processes

**VecSetFromOptions** Like the matrix counterpart. `VecSetType` may be used instead.

**VecGetOwnershipRange** Like the matrix counterpart

**VecSetValue** Sets the value for a vector element. Use `ADD_VALUES` to add to a vector element. Like the matrix routines, elements can be inserted or added by any process.

**VecAssemblyBegin** and **VecAssemblyEnd** Like the Matrix counterparts

# Solving a Poisson Problem I

```
1  #include <math.h>
2  #include "petscksp.h"
3  extern Mat FormLaplacian2d( int );
4  extern Vec FormVecFromFunction2d( int, double (*)(double,double) );
5  /* This function is used to define the right-hand side of the
6     Poisson equation to be solved */
7  double func( double x, double y ) {
8     return sin(x*M_PI)*sin(y*M_PI); }
9
10 int main( int argc, char *argv[] )
11 {
12     KSP      sles;
13     Mat      A;
14     Vec      b, x;
15     int      its, n;
16
17     PetscInitialize( &argc, &argv, 0, 0 );
```

# Solving a Poisson Problem II

```
18     n = 10;      /* Get the mesh size.  Use 10 by default */
19     PetscOptionsGetInt( PETSC_NULL, "-n", &n, 0 );
20
21     A = FormLaplacian2d( n );
22     b = FormVecFromFunction2d( n, func );
23     VecDuplicate( b, &x );
24     KSPCreate( PETSC_COMM_WORLD, &sles );
25     KSPSetOperators( sles, A, A, DIFFERENT_NONZERO_PATTERN );
26     KSPSetFromOptions( sles );
27     KSPSolve( sles, b, x );
28     KSPGetIterationNumber( sles, &its );
29     PetscPrintf( PETSC_COMM_WORLD, "Solution in %d iterations is:\n"
30     VecView( x, PETSC_VIEWER_STDOUT_WORLD );
31
32     MatDestroy( A ); VecDestroy( b ); VecDestroy( x );
33     KSPDestroy( sles );
34     PetscFinalize( );
35     return 0;
36 }
```

# Understanding the Code

**KSPCreate** Create a *context* used to solve a linear system. This routine is used for *all* solvers, independent of the choice of algorithm or data structure

**KSPSetOperators** Define the problem.

- The third argument allows the use of a different matrix for preconditioning
- `DIFFERENT_NONZERO_PATTERN` indicates whether the preconditioner has the same nonzero pattern each time a system is solved. This default works with all preconditioners. Other values (e.g., `SAME_NONZERO_PATTERN`) can be used for particular preconditioners. Ignored when solving only one system

**KSPSetFromOptions** Set the algorithm, preconditioner, and the associated parameters, using the command-line

**KSPSolve** Actually solve the system of linear equations.

**KSPGetIterationNumber** The number of iterations is returned. If a direct method is used, one is returned in `its`

**KSPDestroy** Free the KSP context and all storage associated with it

# Objects in PETSc

- How should a matrix be described in a program?

- Old way:

- Dense matrix

```
double precision A(10,10)
```

- Sparse matrix

```
integer ia(11), ja(max_nz)
```

```
double precision a(max_nz)
```

- New way:

```
Mat M
```

- Hides the choice of data structure

- Of course, the library still needs to represent the matrix with some choice of data structure, but this is an *implementation detail*

- Benefit

- Programs become independent of any particular choice of data structure, making it easier to modify and adapt programs.

# Operations in PETSc

- How should operations like “solve linear system” be described in a program?

- Old way

```
mpiaijgmres( ia, ja, a, comm, x, b, nlocal, nglobal,  
            ndir, orthomethod, convtol, &its )
```

- New way

```
KSPSolve( ksp, b, x )  
KSPGetIterationNumber( ksp, &its )
```

- Hides the choice of algorithm
  - Algorithms are to operations as data structures are to objects
- Benefit
  - Programs become independent of a particular choice of algorithm, making it easier to explore algorithmic choices and to adapt to new methods
- In PETSc, operations have their own “handle”, called a “*context variable*”



# Context Variables in PETSc

- Context variables are the key to solver organization
- They contain the complete state of an algorithm, including
  - parameters (e.g., convergence tolerance)
  - functions run by the algorithm (e.g., convergence monitoring routine)
  - information about the current state (e.g., iteration number)

# KSP Structure

- Each KSP object contains two important objects:
  - Krylov Space Method
    - The iterative method
    - The KSP context contains information on the method parameters, e.g. GMRES restart and search directions)
  - PC** Preconditioners
    - Knows how to apply the preconditioner
    - The context contains information on the preconditioner, such as ILU fill level

# Available Methods

KSP		PC	
Name	PETSc option	Name	PETSc option
Conjugate Gradient	cg	Block Jacobi	bjacobi
GMRES	gmres	Overlapping Additive Schwarz	asm
Bi-CG-stab	bicg	ILU	ilu
Transpose-free QMR	tfqmr	SOR	sor
Richardson	richardson	LU (direct solve)	lu
CG-Squared	cgs	Multigrid	mg
SYMMLQ	symmlq	Arbitrary matrix	mat
others		others	

# Using the Command Line Interface

- PETSc makes it easy to try different algorithms

```
mpiexec -n 4 poisson -ksp_type cg
```

```
mpiexec -n 4 poisson -ksp_type gmres
```

```
mpiexec -n 4 poisson -pc_type bjacobi -sub_pc_type ilu \  
-ksp_type bcgs
```

- PETSc make experimentation with different algorithms *easy*
  - Many are already built-in
  - You can add new algorithms and data structures to PETSc; these are then used just like the built-in ones (e.g., a new preconditioner can be used with an existing source code without *any* changes. (However, this is not a one-day project.)
- Many other options available. Use

```
poisson -help | more
```

to get a list of available options

# Monitoring Convergence

- PETSc provides routines to check for and monitor convergence
- The choice of monitor and the output from that monitor can be controlled from the command line
  - ksp\_monitor Print the preconditioned residual norm
  - ksp\_monitor\_draw Plot the preconditioned residual norm
  - ksp\_monitor\_true\_residual\_norm Print the true residual norm  $\|Ax - b\|_2$
  - ksp\_monitor\_draw\_true\_residual Plot the true residual norm
- Custom monitors can be defined by the user

# Accessing the Solution

- *Viewers* are used in PETSc to access and display the contents of an object
- A simple viewer prints data out standard output:

```
VecView( V, PETSC_VIEWER_STDOUT_WORLD );
```

- PETSc provides a wide range of viewers for all major objects
  - Viewers make it easy to send vectors and matrices to Matlab
  - Graphical viewers make it easy to display data
  - Binary viewers make it easy to save and load data

# PETSc Viewers

- PETSc has many viewers

`PETSC_VIEWER_STDOUT_SELF` Sequential, prints to stdout

`PETSC_VIEWER_STDOUT_WORLD` Parallel, prints to stdout

`PETSC_VIEWER_DRAW_WORLD` Parallel, draws using  
X-Windows

- Viewers exist for matrices, vectors, and other objects
  - Matrix viewers provide information and graphical display of matrix sparsity structure and assembly (try `-mat_view_draw`, `-mat_view_info`, or `-mat_view`)
  - Viewers on other objects can print out information about the object

# Working With Vectors

- It is sometimes helpful to have direct access to the storage for the local elements of a vector
- The routines `VecGetArray` and `VecRestoreArray` may be used to get and return the local elements
- The routine `VecGetLocalSize` returns the number of elements in the local part of the vector
- `VecGetArray` returns a pointer to an array that contains the locally-owned values in the vector. Normally, this is just a pointer into the storage that PETSc uses, but for special vector implementations, it may be different storage used just for `VecGetArray`
- `VecRestoreArray` gives the array back to PETSc. Normally, this has no work to do, but if PETSc had to allocate storage for `VecGetArray`, this routine will free that storage



# Example: Computing $\|x - y\|$

- Often need to compute  $\|x - y\|$ , for example, for convergence tests. Also useful in checking a solution
- PETSc does provide routines to compute  $x + \alpha y$  and  $\|x\|$ , but no single routine to compute the norm of the difference of two vectors
- As an example of accessing local elements of a vector, we will implement “mVecNormXPAY” which computes  $\|x + \alpha y\|$ 
  - Accepts all PETSc norm types: NORM\_1, NORM\_2, and NORM\_INFINITY.
- A single routine avoids creating an unneeded temporary vector and avoids extra memory motion needed when using multiple routines

# Computing $\|x - y\|$ I

```
1  #include "petscvec.h"
2
3  /* Comment out this next define if the compiler is
4     not broken and supports C2000 */
5  #define restrict
6  /* This is a new vector routine for PETSc, illustrating the use
7     of several PETSc functions for accessing vector elements */
8
9  int mVecNormXPAY( Vec x, Vec y, const PetscScalar a, NormType ntype,
10                  PetscReal *norm )
11  {
12     const double * restrict xvals, * restrict yvals;
13     int nlocal, i, ierr = 0;
14     MPI_Op normop;
15     double sum = 0.0, totsum;
16
17     /* Get the local arrays and the size */
18     VecGetArray( x, (PetscScalar **)&xvals );
19     VecGetArray( y, (PetscScalar **)&yvals );
```

# Computing $\|x - y\|$ II

```
20     VecGetLocalSize( x, &nlocal );
21
22     if (a == -1) {
23         /* Special case for difference of two vectors */
24         switch (ntype) {
25             case NORM_1:
26                 for (i=0; i<nlocal; i++) {
27                     sum += fabs(xvals[i] - yvals[i]);
28                 }
29                 normop = MPI_SUM;
30                 break;
31             case NORM_2:
32                 for (i=0; i<nlocal; i++) {
33                     register PetscScalar tmp;
34                     tmp = xvals[i] - yvals[i];
35                     sum += tmp*tmp;
36                 }
37                 normop = MPI_SUM;
```

# Computing $\|x - y\|$ III

```
38         break;
39     case NORM_INFINITY:
40         for (i=0; i<nlocal; i++) {
41             register PetscScalar tmp;
42             tmp = fabs(xvals[i] - yvals[i]);
43             if (tmp > sum) sum = tmp;
44         }
45         normop = MPI_MAX;
46         break;
47     default:
48         ierr = 1;
49         break;
50     }
51 }
52 else {
53     /* Unimplemented */
54     ierr = 1;
55 }
```

# Computing $\|x - y\|$ IV

```
56     if (!ierr) {
57         MPI_Comm comm;
58         PetscObjectGetComm( (PetscObject)x, &comm );
59         MPI_Allreduce( &sum, &totsum, 1, MPI_DOUBLE, comm, normop );
60         if (ntype == NORM_2) {
61             totsum = sqrt( totsum );
62         }
63         *norm = totsum;
64     }
65
66     VecRestoreArray( x, (PetscScalar **)&xvals );
67     VecRestoreArray( y, (PetscScalar **)&xvals );
68
69     return ierr;
70 }
```

PetscScalar is just a name for double; using this name allows the PETSc to be rebuilt for float or Complex scalars.

# Distributed Arrays in PETSc

How should a vector be distributed across processes? PETSc's default is a "one-dimensional decomposition"

How can you make use of different data decompositions in PETSc? PETSc provides "Distributed Arrays" (DAs) for this purpose.


For example, consider the layout of a mesh onto this processor mesh:

P2	P3
P0	P1

# Layout Of Distributed Arrays

On this  $2 \times 2$  process grid, the vector elements are numbered like this:

20	21	22	23	24
15	16	17	18	19
10	11	12	13	14
5	6	7	8	9
0	1	2	3	4



18	19	20	23	24
15	16	17	21	22
6	7	8	13	14
3	4	5	11	12
0	1	2	9	10

Natural numbering

PETSc's internal numbering

DAs provide a “logically Cartesian” decomposition. There are no physical coordinates associated with a DA.

# Distributed Arrays

- PETSc's distributed array (DA) provides a way to describe a multidimensional array, distributed across a parallel computer
- DAs provide a way to use more complex data decompositions

```
DACreate2d( PETSC_COMM_WORLD, DA_NONPERIODIC,  
           DA_STENCIL_STAR,  
           nx, ny, px, py, 1, 1, 0, 0, &grid );
```

creates a global  $nx \times ny$  grid, with a  $px \times py$  process decomposition

- The `DA_STENCIL_STAR` and the arguments after `py` have to do with the difference stencil that may be used with this array and will be discussed later.
- `MPI_Dims_create` may be used to determine good values for `px` and `py`.



# Setting the Vector Values I

```
1  #include "petsc.h"
2  #include "petscvec.h"
3  #include "petscda.h"
4
5  /* Form a vector based on a function for a 2-d regular mesh on the
6     unit square */
7  Vec FormVecFromFunctionDA2d( DA grid, int n,
8                               double (*f)( double, double ) )
9  {
10     Vec      V;
11     int      is, ie, js, je, in, jn, i, j;
12     double  h;
13     double  **vval;
14
15     h = 1.0 / (n + 1);
16     DACreateGlobalVector( grid, &V );
17
```

# Setting the Vector Values II

```
18     DAVecGetArray( grid, V, (void **)&vval );
19     /* Get global coordinates of this patch in the DA grid */
20     DAGetCorners( grid, &is, &js, 0, &in, &jn, 0 );
21     ie = is + in - 1;
22     je = js + jn - 1;
23     for (i=is ; i<=ie ; i++) {
24         for (j=js ; j<=je ; j++){
25             vval[j][i] = (*f)( (i + 1) * h, (j + 1) * h );
26         }
27     }
28     DAVecRestoreArray( grid, V, (void **)&vval );
29
30     return V;
31 }
32
```

# Understanding the Code

**DACreateGlobalVector** Creates a PETSc vector that may be used with DAs

**DAVecGetArray** Get a multidimensional array that gives the illusion of a global array (PETSc uses tricks with the array indexing to provide access to the local elements of the vector). Otherwise, like VecGetArray.

**DAVecRestoreArray** Like VecRestoreArray, used to allow PETSc to free any storage allocated by DAVecGetArray

**DAGetCorners** Returns the indices of the lower-left corner of the local part of the distributed array relative to the global coordinates, along with the number of points in each direction.

# Setting the Matrix Elements I

```
1  #include "petscksp.h"
2  #include "petscda.h"
3
4  /* Form the matrix for the 5-point finite difference 2d Laplacian
5     on the unit square. n is the number of interior points along a
6     side */
7  Mat FormLaplacianDA2d( DA grid, int n )
8  {
9     Mat      A;
10    int      r, i, j, is, ie, js, je, in, jn, nelm;
11    MatStencil cols[5], row;
12    double    h, oneByh2, vals[5];
13
14    h = 1.0 / (n + 1); oneByh2 = 1.0 / (h*h);
15
16    DAGetMatrix( grid, MATMPIAIJ, &A );
17    /* Get global coordinates of this patch in the DA grid */
```

# Setting the Matrix Elements II

```
18     DAGetCorners( grid, &is, &jns, 0, &in, &jn, 0 );
19     ie = is + in - 1;
20     je = js + jn - 1;
21     /* This is a simple but inefficient way to set the matrix */
22     for (i=is; i<=ie; i++) {
23         for (j=js; j<=je; j++){
24             row.j = j; row.i = i; nelms = 0;
25             if (j - 1 > 0) {
26                 vals[nelms] = oneByh2;
27                 cols[nelms].j = j - 1; cols[nelms+1].i = i;
28             if (i - 1 > 0) {
29                 vals[nelms] = oneByh2;
30                 cols[nelms].j = j;     cols[nelms+1].i = i - 1;
31             vals[nelms] = - 4 * oneByh2;
32             cols[nelms].j = j;     cols[nelms+1].i = i;
33             if (i + 1 < n - 1) {
34                 vals[nelms] = oneByh2;
35                 cols[nelms].j = j;     cols[nelms+1].i = i + 1;}
```

# Setting the Matrix Elements III

```
36         if (j + 1 < n - 1) {
37             vals[nelm] = oneByh2;
38             cols[nelm].j = j + 1; cols[nelm++].i = i;}
39         MatSetValuesStencil( A, 1, &row, nelm, cols, vals,
40                             INSERT_VALUES );
41     }
42 }
43
44 MatAssemblyBegin(A, MAT_FINAL_ASSEMBLY);
45 MatAssemblyEnd(A, MAT_FINAL_ASSEMBLY);
46
47 return A;
48 }
49
```

# Understanding the Code

**DAGetMatrix** Returns a matrix whose elements can be accessed with the coordinates of the distributed array. The type of the matrix must be specified; this chooses a parallel matrix using AIJ format (MATMPIAIJ).

**MatSetValuesStencil** Sets elements of a matrix using grid coordinates

**MatStencil** Data structure that contains the indices of a point in the DA, using the  $i, j, k$  members of the structure

# Poisson Solver Revisited

```
1  #include <math.h>
2  #include "petscksp.h"
3  #include "petscda.h"
4  extern Mat FormLaplacianDA2d( DA, int );
5  extern Vec FormVecFromFunctionDA2d( DA, int, double (*)(double,double) );
6  /* This function is used to define the right-hand side of the
7     Poisson equation to be solved */
8  double func( double x, double y ) {
9     return sin(x*M_PI)*sin(y*M_PI); }
10
11 int main( int argc, char *argv[] )
12 {
13     KSP      ksp;
14     Mat      A;
15     Vec      b, x;
16     DA       grid;
17     int      its, n, px, py, worldSize;
```



# Poisson Solver Revisited II

```
18
19     PetscInitialize( &argc, &argv, 0, 0 );
20
21     /* Get the mesh size.  Use 10 by default */
22     n = 10;
23     PetscOptionsGetInt( PETSC_NULL, "-n", &n, 0 );
24     /* Get the process decomposition.  Default it the same as without
25         DAs */
26     px = 1;
27     PetscOptionsGetInt( PETSC_NULL, "-px", &px, 0 );
28     MPI_Comm_size( PETSC_COMM_WORLD, &worldSize );
29     py = worldSize / px;
30
31     /* Create a distributed array */
32     DACreate2d( PETSC_COMM_WORLD, DA_NONPERIODIC, DA_STENCIL_STAR,
33                n, n, px, py, 1, 1, 0, 0, &grid );
34
35     /* Form the matrix and the vector corresponding to the DA */
```

# Poisson Solver Revisited III

```
36     A = FormLaplacianDA2d( grid, n );
37     b = FormVecFromFunctionDA2d( grid, n, func );
38     VecDuplicate( b, &x );
39     KSPCreate( PETSC_COMM_WORLD, &ksp );
40     KSPSetOperators( ksp, A, A, DIFFERENT_NONZERO_PATTERN );
41     KSPSetFromOptions( ksp );
42     KSPSolve( ksp, b, x );
43     KSPGetIterationNumber( ksp, &its );
44
45     PetscPrintf( PETSC_COMM_WORLD, "Solution is:\n" );
46     VecView( x, PETSC_VIEWER_STDOUT_WORLD );
47     PetscPrintf( PETSC_COMM_WORLD, "Required %d iterations\n", its );
48
49     MatDestroy( A ); VecDestroy( b ); VecDestroy( x );
50     KSPDestroy( ksp ); DADestroy( grid );
51     PetscFinalize( );
52     return 0;
53 }
```

# More Preconditioners

- PETSc provides a large collection of preconditioners, including domain decomposition preconditioners

- Additive Schwarz

```
mpiexec -n 4 poisson -pc_type asm
```

- Control the subdomain solver with `-sub_pc_type`:

```
mpiexec -n 4 poisson -pc_type asm -sub_pc_type ilu
```

(In general, `-sub_pc_<pcparamname>` may be used to change the PC parameter `pcparamname` in the subdomain, and `-sub_ksp_<kspparmname>` for KSP in the subdomain.)

- Control the subdomain overlap

```
mpiexec -n 4 poisson -pc_type asm -pc_asm_overlap 2
```

# PETSc's Automatic ASM

- PETSc automatically generates overlap by using the structure of the sparse matrix. Control with `-pc_asm_overlap`
- DAs allow you to control the local physical domain
- By using DAs, you can experiment with the effects of different decompositions

```
mpiexec -n 16 poisson -n 64 -pc_type asm
```

```
mpiexec -n 16 poisson2 -n 64 -pc_type asm -mx 8 -my 2
```

```
mpiexec -n 16 poisson2 -n 64 -pc_type asm -mx 4 -my 4
```

- Other ASM types are available with `-pc_asm_type`

**basic** full interpolation and restriction

**restrict** full restriction, local process interpolation

**interpolate** full interpolation, local process restriction

**none** local process restriction and interpolation

# Solving Nonlinear Equations

We would like to solve

$$F(u) = 0$$

for  $u$ . A powerful method for this is *Newton's method*:

$$u^{k+1} = u^k - (F'(u^k))^{-1} F(u^k), \quad k = 0, 1, \dots$$

where  $u^k$  is the approximation to  $u$  at the  $k$ th step. The term  $F'(u^k)$  is a matrix, and this algorithm can be rewritten as

$$\begin{aligned} F'(u^k) \Delta u^k &= -F(u^k) \\ u^{k+1} &= u^k + \Delta u^k \end{aligned}$$

# Newton-based Methods

In practice, various modifications are made to Newton's method. PETSc supports many of the most common:

- Line search strategies
- Trust region strategies
- Pseudo-transient continuation
- Matrix-free variants

PETSc provides a “Simplified Nonlinear Equation Solver” (SNES) for nonlinear problems. SNES is the nonlinear analogue of KSP.

# PDE Jacobian

The matrix  $F'(u)$  is called the *Jacobian*.

For PDE problems, computing the Jacobian can be tricky. Three choices are:

1. Compute  $F'$  analytically, then discretize
2. Discretize  $F$ , then compute  $F'$  by finite difference approximation
3. Discretize  $F$ , then compute  $F'$  by analytically differentiating the discretization of  $F$

PETSc provides additional support for 2, and by interfacing to ADIFOR and ADIC, support for 3

# A Simple Nonlinear PDE

The *Bratu* problem is defined by

$$\begin{aligned} -\nabla^2 u - \lambda e^u &= 0 \text{ in } [0, 1] \times [0, 1] \\ u &= 0 \text{ on the boundary} \end{aligned}$$

We will use the same simple discretization for this problem as for the Poisson problem.




# Evaluating the Function

- Evaluating the function  $F(u) = -\nabla^2 u - \lambda e^u$  is somewhat difficult because it involves a differential operator. This requires information from the neighboring processes. We will use distributed arrays (DAs) to help with this, taking advantage of their support for different *stencils*.
- An alternate approach for *this* example is to use a matrix-vector multiply, using

```
MatMult( A, x, y );
```



to compute  $y = Ax$ . This routine handles all data motion required. However, it is suitable only for relatively simple  $F(u)$ . Thus, we will explore more general techniques

# Stencils




Star Stencil

(DA\_STENCIL\_STAR)




Box Stencil

(DA\_STENCIL\_BOX)





# Stencils




Star Stencil

(DA\_STENCIL\_STAR)




Box Stencil

(DA\_STENCIL\_BOX)




# Global and Local Representations

- A vector associated with a DA has two representations: the *global* and the *local*
- The global representation is nothing more than the natural mesh, distributed across all processes
- The local representation is the local part of the global mesh, *plus* the ghost points



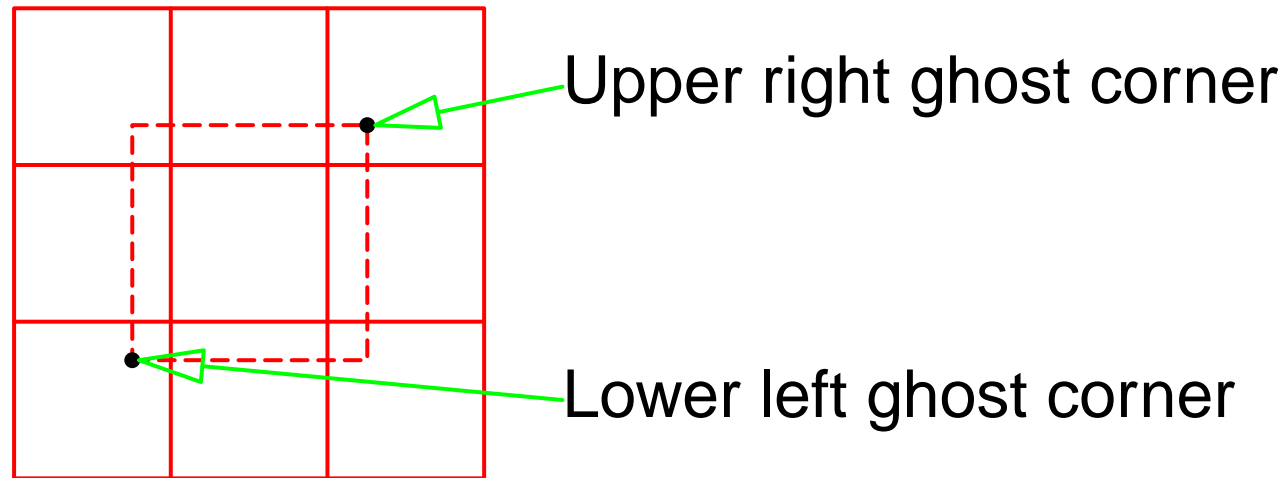
**Global:** each process stores a unique local set of vertices, and each vertex is owned by exactly one process



**Local:** each process stores a unique local set of vertices *as well* as ghost points from neighboring processes

# Using Ghost Points with DAs

A ghost region is defined by the coordinates *in the global representation*:



The routine `DAGetGhostCorners` returns this information, similar to `DAGetCorners`

## Moving Data Between the Global and Local Representations

**DACreateLocalVector** Creates a PETSc vector that can hold the local representation of a DA (the local mesh plus ghost points)

**DAGlobalToLocalBegin** and **DAGlobalToLocalEnd**

Update the ghostpoint values. This involves communication with the neighboring processes. The update may use **INSERT\_VALUES** or **ADD\_VALUES**.

**DALocalToGlobal** Transfers values in the local representation back to the global representation. The ghost points are discarded.

# Parallel Evaluation of the Function

In the Bratu example,

$$F(u) = -\nabla^2 u - \lambda e^u$$

so

$$F'(u)a = -\nabla^2 a - \lambda a e^u,$$

where  $a e^u$  is just  $\{a_i \times e^{u_i}\}$ . Thus the Jacobian  $F'(u)$  is almost the same as the matrix for the Poisson problem, with a diagonal element that depends on  $u$ . Now that we know what these are, how do we provide them to PETSc?

## Providing the Function and Jacobian

We now have functions that evaluate  $F$  and  $F'$ . How can these be used by the SNESolve routine?

- The algorithm needs to evaluate both, under control of the algorithm
- The solution used in PETSc is to pass the functions themselves to the routine that defines the problem, much as the matrix defining a linear problem to solve is passed to KSPSetOperators.
- This is a “callback” method, because the user provides functions to the solver that are called back by the algorithm when their results are needed
- The *calling sequence* for the routine is specified by PETSc.



# Specifying Callbacks

- User provides the routines to perform actions that the library requires. For example

```
SNESSetFunction(snes, f, userfunc, userctx )
```

**snes** SNES context

**f** Vector that will be used to store the function value

**userfunc** Name of (really, pointer to) the function

**userctx** Pointer to data passed that will be passed to the function

- The library can call this function whenever it needs to evaluate the function
- The userctx pointer allows the user to provide an “application context” object. By using this approach, the library need never know the details of data needed only by the application.

# Forming the Function I

```
#include "petscsnes.h"
#include "petscda.h"
#include "bratu.h"
#include <math.h>

/* Evaluate the function for the Bratu nonlinear problem on the local
   mesh points */
int FormBratuFunction( SNES snes, Vec v, Vec f, void *ctx )
{
    UserBratuCtx *bratu = (UserBratuCtx *)ctx;
    DA           da      = bratu->da;
    double       lambda  = bratu->lambda;
    double       h       = bratu->h;
    Vec          lv;
    int          i, j;
    int          lli, llj, ni, nj; /* lower left i,j and size for local
                                     part of mesh */
```

# Forming the Function II

```
const double **varr;  
double      **fvarr;  
  
/* Get the coordinates of our part of the global mesh */  
DAGetCorners( da, &lli, &llj, 0, &ni, &nj, 0 );  
  
DAGetLocalVector( da, &lv );  
  
/* Scatter the ghost points to the other processes, using  
   the values in the input vector v */  
DAGlobalToLocalBegin( da, v, INSERT_VALUES, lv );  
DAGlobalToLocalEnd( da, v, INSERT_VALUES, lv );  
  
DAVecGetArray( da, lv, (void **)&varr );  
DAVecGetArray( da, f, (void **)&fvarr );  
  
for (j=llj ; j<llj+nj ; j++)  
    for (i=lli ; i<lli+ni ; i++) {
```

# Forming the Function III

```
if (i == 0 || j == 0 ||
    i == bratu->n + 1 || j == bratu->n + 1) {
    fvarr[j][i] = 0.0;
}
else {
    fvarr[j][i] = -( varr[j-1][i] + varr[j][i-1] +
                    varr[j+1][i] + varr[j][i+1] -
                    4 * varr[j][i] ) / (h*h) -
                lambda * exp(varr[j][i]);
}
}
```

```
DAVecRestoreArray( da, f, (void **)&fvarr );
```

```
DAVecRestoreArray( da, lv, (void **)&varr );
```

```
DARestoreLocalVector( da, &lv );
```

```
return 0;
```

```
}
```

# Understanding the Code

- One key feature of this routine is the use of the fourth argument, “ctx”, to pass additional information to the Function. In this case, we use a user-defined structure define in bratu.h:

```
/* This typedef defines a struct that contains the
   data that we need to have when evaluating the
   function or the Jacobian for the Bratu problem */
typedef struct {
    DA      da;          /* DA for grid */
    double h;          /* Mesh spacing */
    double lambda;     /* parameter in problem */
    int     n;          /* interior grid is n x n */
} UserBratuCtx;
```

- The rest of the code uses the DA to provide ghost values for the the evaluation of the finite difference scheme
  - Boundary conditions, as always, add complexity

# Forming the Jacobian I

```
#include "petscsnes.h"
#include "petscda.h"
#include "bratu.h"
#include <math.h>

/* Form the matrix for the Jacobian of the Bratu problem, where the
   function uses a 5-point finite difference 2d Laplacian on the
   unit square. n is the number of interior points along a side */
Mat FormBratuJacobian( SNES snes, Vec u, Mat *A, Mat *B,
                      MatStructure *flag, void *ctx )
{
    Mat      jac = *A;
    UserBratuCtx *bratu = (UserBratuCtx *)ctx;
    DA      da = bratu->da;
    int     r, i, j, n = bratu->n;
    double  h = bratu->h, lambda = bratu->lambda;
    double  oneByh2 = 1.0 / (h*h), **uvals, v[5];
```

# Forming the Jacobian II

```
int    lli, llj, ni, nj; /* lower left i,j and size for local
                        part of mesh */

MatStencil row, col[5];

DAGetCorners( da, &lli, &llj, 0, &ni, &nj, 0 );
DAVecGetArray( da, u, (void **)&uvals );

/* This is a simple but inefficient way to set the matrix */
for (j=llj; j<llj+nj; j++) {
    for (i=lli; i<lli+ni; i++) {
        row.i = i; row.j = j;
        if (i == 0 || j == 0 ||
            i == n + 1 || j == n + 1) {
            v[0] = 1.0;
            MatSetValuesStencil( jac, 1, &row, 1, &row, v, INSERT_VALUES );
        }
        else {
            col[0].i = i; col[0].j = j - 1; v[0] = - oneByh2;
```

# Forming the Jacobian III

```
        col[1].i = i; col[1].j = j + 1; v[1] = - oneByh2;
        col[2].i = i - 1; col[2].j = j; v[2] = - oneByh2;
        col[3].i = i + 1; col[3].j = j; v[3] = - oneByh2;
        col[4].i = i; col[4].j = j;
        v[4] = 4.0 * oneByh2 - lambda * exp( uvals[j][i] );
        MatSetValuesStencil( jac, 1, &row, 5, col, v, INSERT_VALUES );
    }
}
}
MatAssemblyBegin( jac, MAT_FINAL_ASSEMBLY );
DAVecRestoreArray( da, u, (void **)&uvals );

*flag = SAME_NONZERO_PATTERN; /* preconditioner has same structure
MatAssemblyEnd( jac, MAT_FINAL_ASSEMBLY );

return 0;
}
```



# Bratu Example I

```
#include "petscsnes.h"
#include "petscda.h"
#include "bratu.h"

extern int FormBratuJacobian( SNES,Vec,Mat *,Mat *,MatStructure *,void * );
extern int FormBratuFunction( SNES, Vec, Vec, void * );

int main( int argc, char *argv[] )
{
    UserBratuCtx bratu;
    SNES          snes;
    Vec           x, r;
    Mat           J;
    int           its;

    PetscInitialize( &argc, &argv, 0, 0 );
```

# Bratu Example II

```
/* Get the problem parameters */
bratu.lambda = 6.0;
PetscOptionsGetReal( 0, "-lambda", &bratu.lambda, 0 );
if (bratu.lambda >= 6.81 || bratu.lambda < 0) {
    SETERRQ(1, "Lambda must be between 0 and 6.81");
}
bratu.n = 10; /* Get the mesh size. Use 10 by default */
PetscOptionsGetInt( PETSC_NULL, "-n", &bratu.n, 0 );
bratu.h = 1.0 / (bratu.n + 1);

SNESCreate( PETSC_COMM_WORLD, &snes );

/* Create the mesh and decomposition */
DACreate2d( PETSC_COMM_WORLD, DA_NONPERIODIC, DA_STENCIL_STAR,
            bratu.n + 2, bratu.n + 2, PETSC_DECIDE, PETSC_DECIDE,
            1, 1, 0, 0, &bratu.da );

DACreateGlobalVector( bratu.da, &x );
```

# Bratu Example III

```
VecDuplicate( x, &r ); /* Use this as the vector to give SetFunction
SNESSetFunction( snes, r, FormBratuFunction, &bratu );

DAGetMatrix( bratu.da, MATMPIAIJ, &J );
SNESSetJacobian( snes, J, J, FormBratuJacobian, &bratu );

SNESSetFromOptions( snes );

FormBratuInitialGuess( &bratu, x );
SNESsolve( snes, x, &its );

PetscPrintf( PETSC_COMM_WORLD,
             "Number of Newton iterations = %d\n", its );

VecDestroy(r); DADestroy(bratu.da);
SNESDestroy(snes);
PetscFinalize( );
return 0;
}
```

# Understanding the Code

**SNESCreate** Creates the SNES context

**SNESSetFunction** Specify the function to be called to evaluate the function  $F(u)$

**SNESSetJacobian** Specify the function to be called to create the Jacobian matrix.

**SNESSetFromOptions** Set SNES parameters from the commandline

**VecSet** Set all elements of a vector to the same value

**SNESolve** Solve the system of nonlinear equations. Return the number of iterations in `its`

**SNESDestroy** Free the SNES context and recover space

**SETERRQ** The counterpart to `CHKERRQ`, it sets the error and returns a message

# Using the Command Line Interface

- Easy to control Newton features
  - -snes\_type ls
  - -snes\_type tr
  - -snes\_rtol num (relative convergence tolerance)
- Complete control over solution of Jacobian problem—just use the same commandline parameters
  - -ksp\_type cgs
  - -pc\_type asm

# Convenience Functions

- PETSc's design makes it relatively easy to layer functionality
- One example is the support for function and Jacobian evaluation on DAs
  - DASetLocalFunction** Attach a function to a DA
  - DASetLocalJacobian** Attach a Jacobian to a DA
  - SNESDAFormFunction** Tell SNES that the function evaluation should use the function on a DA. to provide the function values
  - SNESDAComputeJacobian** Tell SNES that the Jacobian evaluation should use the Jacobian function on a DA
- The functions provide just the computation applied to the local vector (from the DA, which includes the ghost points)
- *Wrapper* functions provided by **DASetLocalFunction** and **DASetLocalJacobian** handle all of the details of setting up the local vectors and arrays.
- The function passed to **DASetLocalFunction** has the calling sequence:

```
FormFunctionLocal(DALocalInfo *info, PetscScalar **x,  
                 PetscScalar **f, AppCtx *user)
```

# Example Local Function I

```
int FormFunctionLocal(DALocalInfo *info, PetscScalar **x,
                    PetscScalar **f, AppCtx *user)
{
    int          ierr, i, j;
    PetscReal    two = 2.0, lambda, hx, hy, hxdhy, hydhx, sc;
    PetscScalar  u, uxx, uyy;

    PetscFunctionBegin;

    lambda = user->param;
    hx      = 1.0 / (PetscReal)(info->mx-1);
    hy      = 1.0 / (PetscReal)(info->my-1);
    sc      = hx*hy*lambda;
    hxdhy   = hx/hy;
    hydhx   = hy/hx;
```

# Example Local Function II

```
/*
   Compute function over the locally owned part of the grid
*/
for (j=info->ys; j<info->ys+info->ym; j++) {
  for (i=info->xs; i<info->xs+info->xm; i++) {
    if (i == 0 || j == 0 || i == info->mx-1 || j == info->my-1) {
      f[j][i] = x[j][i];
    } else {
      u          = x[j][i];
      uxx        = (two*u - x[j][i-1] - x[j][i+1])*hydhx;
      uyy        = (two*u - x[j-1][i] - x[j+1][i])*hxdhy;
      f[j][i] = uxx + uyy - sc*PetscExpScalar(u);
    }
  }
}

ierr = PetscLogFlops(11*info->ym*info->xm);CHKERRQ(ierr);
PetscFunctionReturn(0);
}
```



# Conclusion

- PETSc provides a powerful framework for
  - Developing applications
  - Experimenting with different algorithms
  - Using abstractions to simplify parallel programming
- PETSc continues to grow and develop
  - New routines added as needed *and understood*
  - PETSc 3 will provide a more powerful framework for combining tools written in different programming languages

# References

- Documentation `www.mcs.anl.gov/petsc/docs`
  - PETSc Users Manual
  - Manual pages (the most up-to-date)
  - Many hyperlinked examples
  - FAQ, Troubleshooting info, installation info, etc.
- Publications `www.mcs.anl.gov/petsc/publications`
  - Research and publications that make use of PETSc
- MPI information `www.mpi-forum.org`
- ***Using MPI*** (2<sup>nd</sup> Edition), by Gropp, Lusk, and Skjellum
- ***Domain Decomposition***, by Smith, Björstad, and Gropp

# Topics Not Covered

- PETSc contains many features, each introduced to provide a necessary feature for an application or researcher
  - Unstructured Meshes
  - Matrix free methods
  - Access to other packages
  - Using different preconditioner matrices
  - Others

# Using PETSc with Other Packages

- Linear solvers

- AMG

[www.mgnet.org/mgnet-codes-gmd.html](http://www.mgnet.org/mgnet-codes-gmd.html)

- BlockSolve95

[www.mcs.anl.gov/BlockSolve95](http://www.mcs.anl.gov/BlockSolve95)

- Hypre

[www.llnl.gov/casc/hypre](http://www.llnl.gov/casc/hypre)

- ILUTP [www.cs.umn.edu/~saad](http://www.cs.umn.edu/~saad)

- LUSOL

[www.sbsi-sol-optimize.com](http://www.sbsi-sol-optimize.com)

- SPAI

[www.sam.math.ethz.ch/~grote/spai](http://www.sam.math.ethz.ch/~grote/spai)

- SuperLU

[www.nersc.gov/~xiaoye/SuperLU](http://www.nersc.gov/~xiaoye/SuperLU)

- ODE solvers

- PVODE

[www.llnl.gov/CASC/PVODE](http://www.llnl.gov/CASC/PVODE)

- Mesh and discretization tools

- Overture

[www.llnl.gov/CASC/Overture](http://www.llnl.gov/CASC/Overture)

- SAMRAI

[www.llnl.gov/CASC/SAMRAI](http://www.llnl.gov/CASC/SAMRAI)

- SUMAA3d

[www.mcs.anl.gov/sumaa3d](http://www.mcs.anl.gov/sumaa3d)

- Optimization software

- TAO [www.mcs.anl.gov/tao](http://www.mcs.anl.gov/tao)

- Veltisto

[www.cs.nyu.edu/~biros/veltisto](http://www.cs.nyu.edu/~biros/veltisto)

- Others

- Matlab [www.mathworks.com](http://www.mathworks.com)

- ParMETIS

[www.cs.umn.edu/~karypis/metis/parmetis](http://www.cs.umn.edu/~karypis/metis/parmetis)

- SLEPc

[www.grycap.upv.es/slepc](http://www.grycap.upv.es/slepc)