



Cineca  
**TRAINING**  
High Performance  
Computing 2017

# Domain specific libraries for PDEs

**Simone Bnà** - [simone.bna@ Cineca.it](mailto:simone.bna@ Cineca.it)  
SuperComputing Applications and Innovation Department





## Outline

- Introduction to Sparse Matrix algebra
- The PETSc toolkit
- Sparse Matrix Computation with PETSc
- Profiling and preliminary tests on KNL



# Introduction to Sparse matrix algebra



## Definition of a Sparse Matrix and a Dense Matrix

- A **sparse matrix** is a matrix in which the number of non-zeroes entries is  $O(n)$  (The average number of non-zeroes entries in each row is bounded independently from  $n$ )

$$\begin{pmatrix} 1.0 & 0 & 5.0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 3.0 & 0 & 0 & 0 & 0 & 11.0 & 0 \\ 0 & 0 & 0 & 0 & 9.0 & 0 & 0 & 0 \\ 0 & 0 & 6.0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 7.0 & 0 & 0 & 0 & 0 \\ 2.0 & 0 & 0 & 0 & 0 & 10.0 & 0 & 0 \\ 0 & 0 & 0 & 8.0 & 0 & 0 & 0 & 0 \\ 0 & 4.0 & 0 & 0 & 0 & 0 & 0 & 12.0 \end{pmatrix}$$

- A **dense matrix** is a non-sparse matrix (The number of non-zeroes elements is  $O(n^2)$ )

$$\begin{pmatrix} 1.0 & 3.4 & 5.0 & 7.5 & 2.3 & 0 & 2.1 & 8.5 \\ 6.5 & 3.5 & 0 & 5.4 & 1.0 & 1 & 0 & 2.1 \\ 0 & 2.8 & 5.7 & 9.2 & 1.1 & 3 & 0 & 2.4 \\ 3.4 & 5.4 & 0 & 4.3 & 3.4 & 2.1 & 1.1 & 4.3 \\ 8.6 & 5.8 & 2.1 & 2.2 & 3.1 & 5.5 & 3.4 & 2.3 \\ 5.4 & 6.7 & 9.8 & 2.1 & 3.4 & 4.3 & 2.1 & 3.5 \\ 4.3 & 3.4 & 1.2 & 5.4 & 0.2 & 3.2 & 0.8 & 1.2 \\ 3.2 & 0 & 1.3 & 4.5 & 0.7 & 9.8 & 0.3 & 1.2 \end{pmatrix}$$



## Sparsity and Density

- The **sparsity** of a matrix is defined as the number of zero-valued elements divided by the total number of elements ( $m \times n$  for an  $m \times n$  matrix)
- The **density** of a matrix is defined as the complementary of the sparsity:  $\text{density} = 1 - \text{sparsity}$
- For Sparse matrices the **sparsity** is  $\approx 1$  and the **density** is  $\ll 1$

Example:

$$m = 8 \quad \text{nnzeros} = 12$$

$$n = 8 \quad \text{nzeros} = m \cdot n - \text{nnzeros}$$

$$\text{sparsity} = 64 - 12 / 64 = \mathbf{0.8125}$$

$$\begin{pmatrix} 1.0 & 0 & 5.0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 3.0 & 0 & 0 & 0 & 0 & 11.0 & 0 \\ 0 & 0 & 0 & 0 & 9.0 & 0 & 0 & 0 \\ 0 & 0 & 6.0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 7.0 & 0 & 0 & 0 & 0 \\ 2.0 & 0 & 0 & 0 & 0 & 10.0 & 0 & 0 \\ 0 & 0 & 0 & 8.0 & 0 & 0 & 0 & 0 \\ 0 & 4.0 & 0 & 0 & 0 & 0 & 0 & 12.0 \end{pmatrix}$$

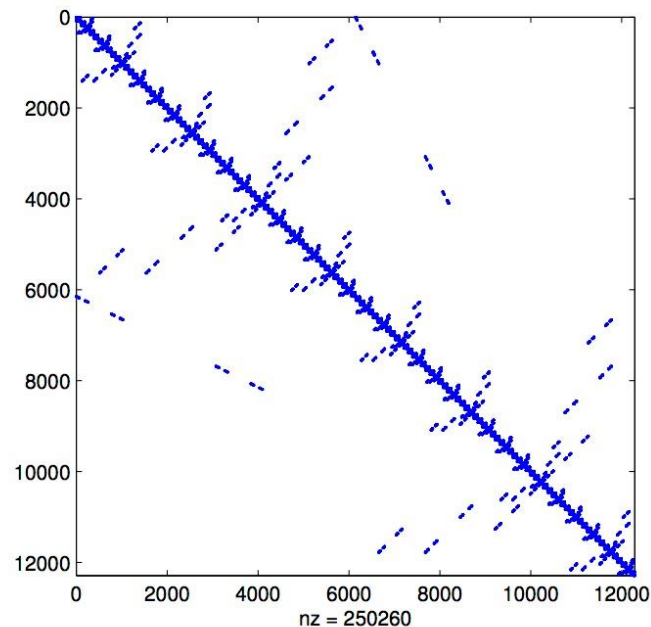
$$\text{density} = 1 - 0.8125 = 0.1875$$



## Sparsity pattern

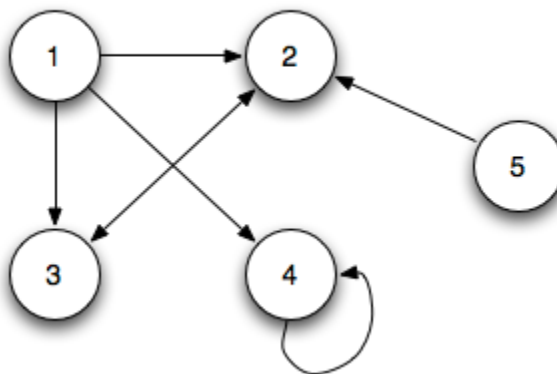
- The distribution of non-zero elements of a sparse matrix can be described by the **sparsity pattern**, which is defined as the set of entries of the matrix different from zero. In symbols:

$$\{ (i, j): A_{ij} \neq 0 \}$$



## Sparsity pattern

- The sparsity pattern can be represented also as a **Graph**, where nodes  $i$  and  $j$  are connected by an edge if and only if  $A_{ij} \neq 0$
- In a Sparse Matrix the **degree of a vertex** in the graph is <<relatively low>>
- Conceptually, sparsity corresponds to a system loosely coupled



|   | 1 | 2 | 3 | 4 | 5 |
|---|---|---|---|---|---|
| 1 | 0 | 1 | 1 | 1 | 0 |
| 2 | 0 | 0 | 1 | 0 | 0 |
| 3 | 0 | 1 | 0 | 0 | 0 |
| 4 | 0 | 0 | 0 | 1 | 0 |
| 5 | 0 | 1 | 0 | 0 | 0 |



## Jacobian of a PDE

- Matrices are used to store the Jacobian of a PDE.
- The following discretizations generates a sparse matrix
  - Finite difference
  - Finite volume
  - Finite element method (FEM)
- Different discretization can lead to a Dense linear matrix:
  - Spectral element method (SEM)

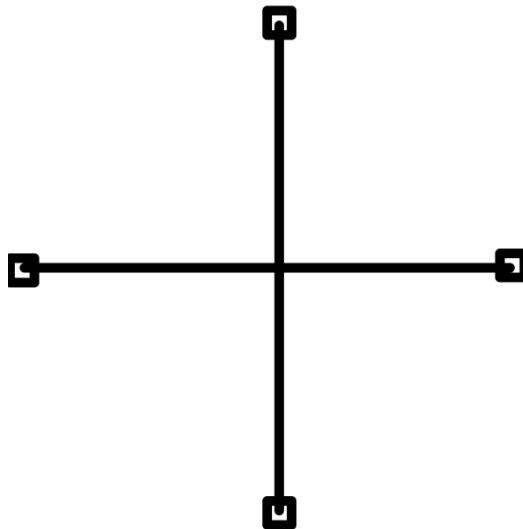




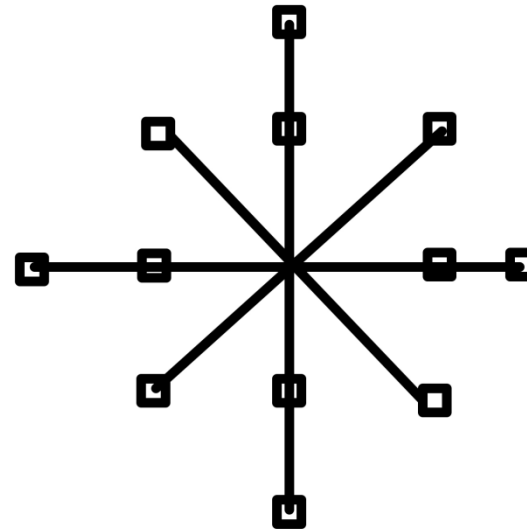
## Sparsity pattern in Finite Difference

- The sparsity pattern in finite difference depends on the topology of the adopted computational grid (e.g. cartesian grid), the indexing of the nodes and the type of stencil

Star stencil



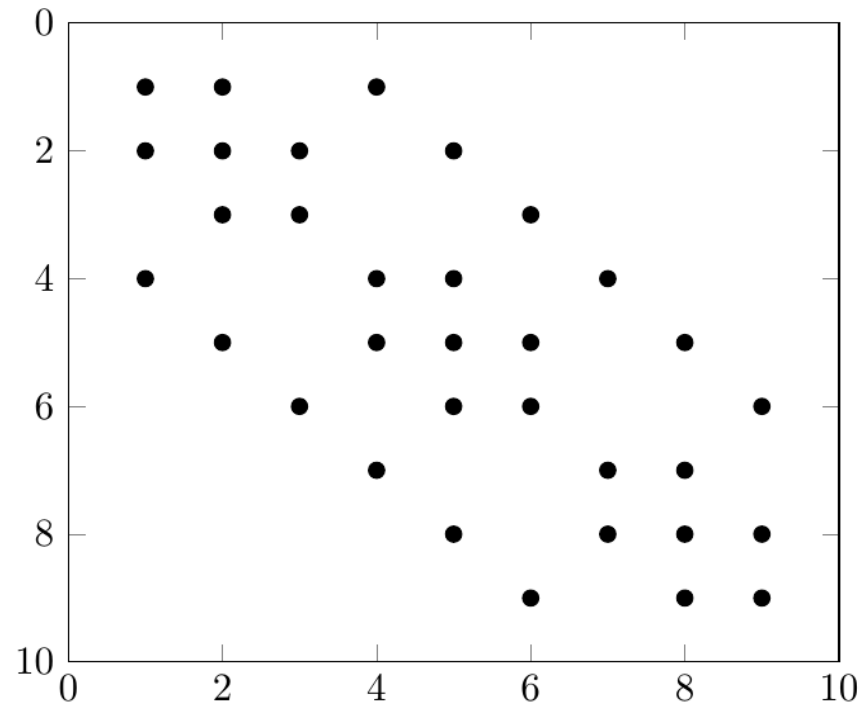
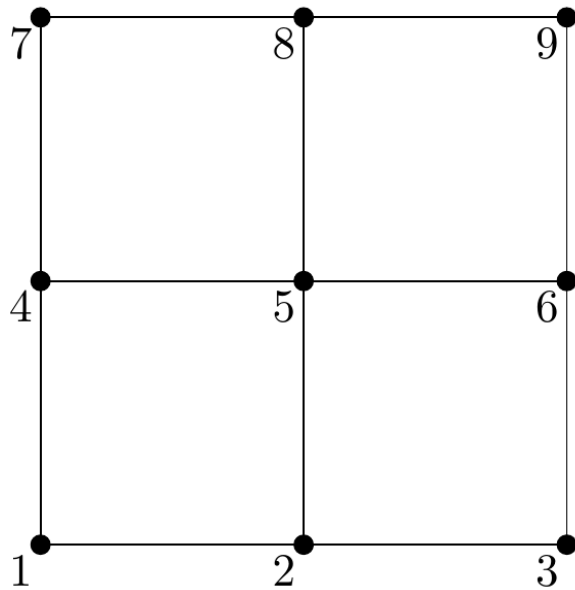
Box stencil





# Sparsity pattern in Finite Difference

- The sparsity pattern in finite difference depends on the topology of the adopted computational grid (e.g. cartesian grid), the indexing of the nodes and the type of stencil





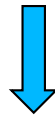
## Sparsity pattern in Finite Element

- The sparsity pattern depends on the topology of the adopted computational grid (e.g. unstructured grid), the kind of the finite element (e.g. Taylor-Hood, Crouzeix-Raviart, Raviart-Thomas, Mini-Element,...) and on the indexing of the nodes.
- In Finite-Element discretizations, the sparsity of the matrix is a direct consequence of the small-support property of the finite element basis
- Finite Volume can be seen as a special case of Finite Element



## Don't reinvent the wheel!

- The use of storage techniques for sparse matrices is fundamental, in particular for large-scale problems
- Standard dense-matrix structures and algorithms are slow and inefficient when applied to large sparse matrices
- There are some available tools to work with Sparse matrices that uses specialised algorithms and data structures to take advantage of the sparse structure of the matrix



- **The PETSc toolkit** (<http://www.mcs.anl.gov/petsc/>)
- **The TRILINOS project** (<https://trilinos.org/>)



# The PETSc toolkit



# PETSc in a nutshell

## PETSc – Portable, Extensible Toolkit for Scientific Computation

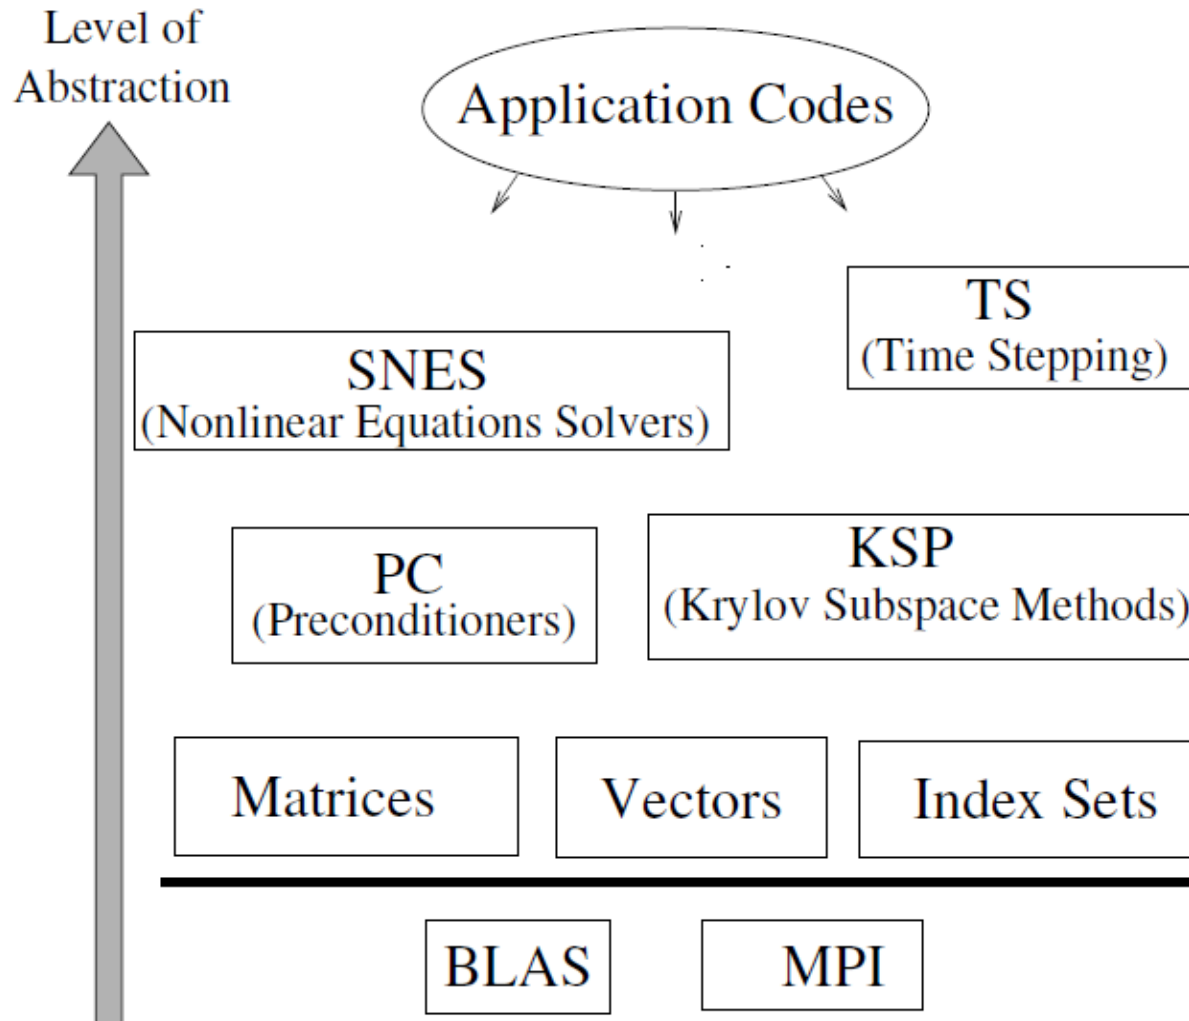
Is a suite of data structures and routines for the scalable (parallel) solution of scientific applications mainly modelled by partial differential equations.

- Tools for distributed vectors and matrices
- Linear system solvers (sparse/dense, iterative/direct)
- Non linear system solvers
- Serial and parallel computation
- Support for Finite Difference and Finite Elements PDE discretizations
- Structured and Unstructured topologies
- Support for debugging, profiling and graphical output



# PETSc class hierarchy

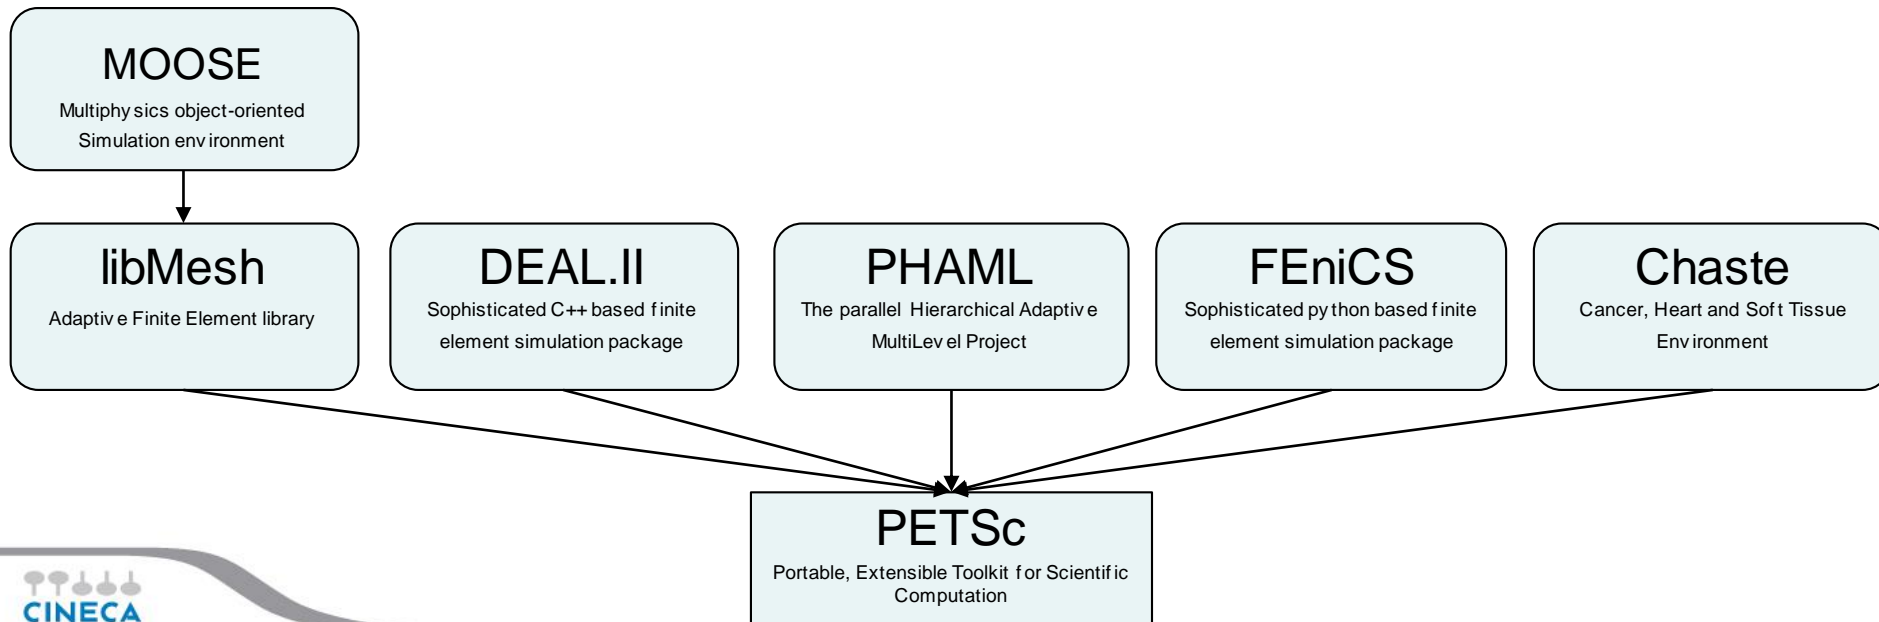
Level of  
Abstraction





# Frameworks built on top of Petsc

- PETSc is a toolkit, not a framework
- PETSc is PDE oriented, but not specific to any kind of PDE
- Alternatives:
  - FEM packages: MOOSE, libMesh, DEAL.II, FEniCS
  - Solvers for classes of problems: CHASTE

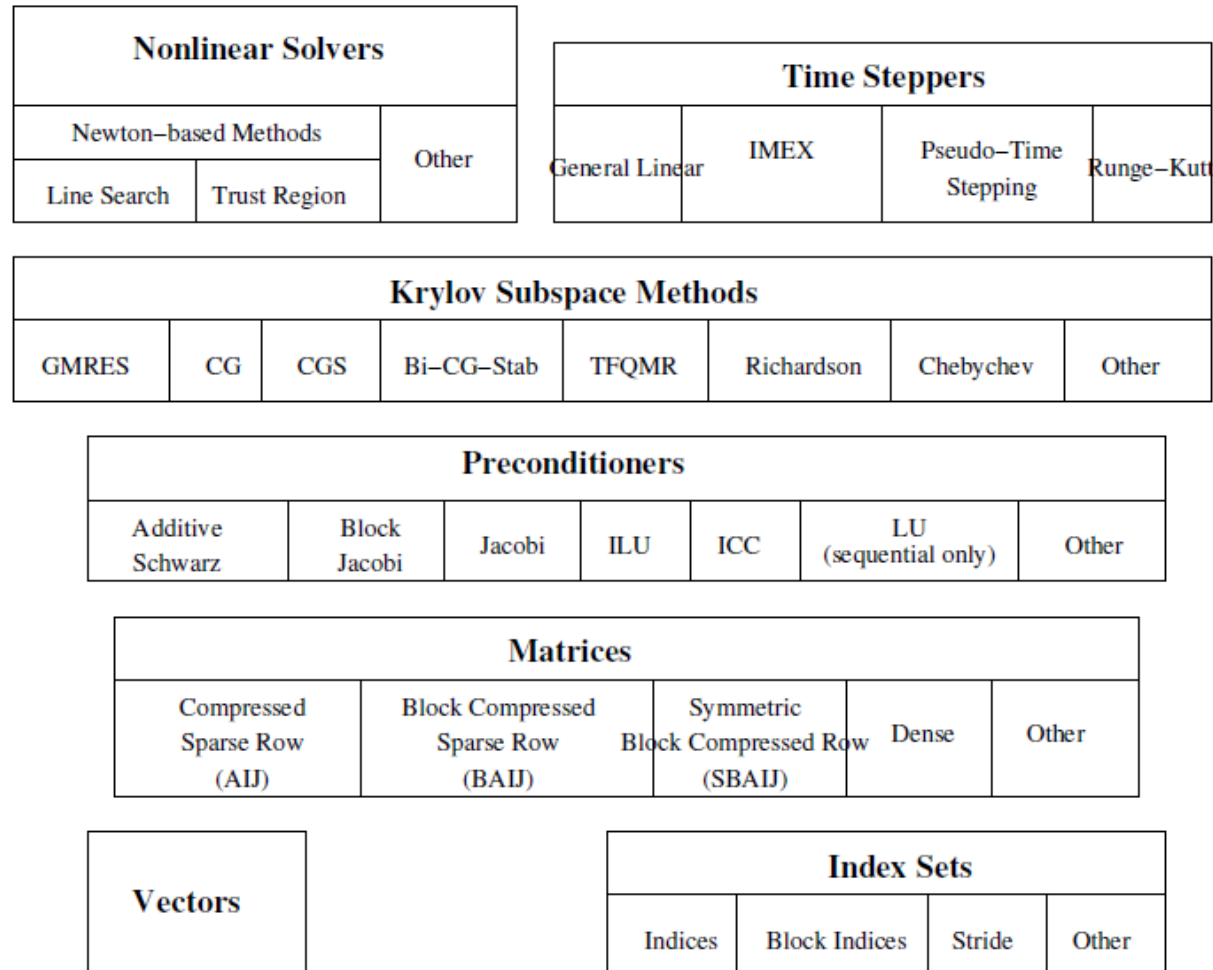






# PETSc numerical components

## Parallel Numerical Components of PETSc





## External Packages

- Dense linear algebra: Scalapack, Plapack
- Sparse direct linear solvers: Mumps, SuperLU, SuperLU\_dist
- Grid partitioning software: Metis, ParMetis, Jostle, Chaco, Party
- ODE solvers: PVODE
- Eigenvalue solvers (including SVD): SLEPc
- Optimization: TAO



# PETSc design concepts

## Goals

- Portability: available on many platforms, basically anything that has MPI
- Performance
- Scalable parallelism
- Flexibility: easy switch among different implementations

## Approach

- Object Oriented Delegation Pattern : many specific implementations of the same object
- Shared interface (overloading):  
`MatMult(A,x,y); // y <- A x`  
same code for sequential, parallel, dense, sparse
- Command line customization

## Drawback

- Nasty details of the implementation hidden



## PETSc and Parallelism

- PETSc is layered on top of **MPI**: you do not need to know much MPI when you use PETSc
- All objects in PETSc are defined on a communicator; they can only interact if on the same communicator
- Parallelism through MPI (**Pure MPI programming model**). Limited support for use with the hybrid MPI-thread model.
  - PETSc supports to have individual threads (OpenMP or others) to each manage their own (sequential) PETSc objects (and each thread can interact only with its own objects).
  - No support for threaded code that made Petsc calls (OpenMP, Pthreads) since PETSc is not «thread-safe».
- Transparent: same code works sequential and parallel.



# Sparse Matrix computation with PETSc



# Vectors

## What are PETSc vectors?

- Represent elements of a vector space over a field (e.g.  $\mathbb{R}^n$ )
- Usually they store field solutions and right-hand sides of PDE
- Vector elements are **PetscScalars** (there are no vectors of integers)
- Each process locally owns a subvector of contiguously numbered global indices

## Features

- Vector types: **STANDARD** (SEQ on one process and **MPI** on several), **VIENNACL**, **CUSP**...
- Supports all vector space operations
  - `VecDot()`, `VecNorm()`, `VecScale()`, ...
- Also unusual ops, like e.g. `VecSqrt()`, `VecReciprocal()`
- Hidden communication of vector values during assembly
- Communications between different parallel vectors



# Numerical vector operations

| Function Name  | Operation                   |
|--|-----------------------------|
| <code>VecAXPY(Vec y, PetscScalar a, Vec x);</code>                 | $y = y + a * x$             |
| <code>VecAYPX(Vec y, PetscScalar a, Vec x);</code>                 | $y = x + a * y$             |
| <code>VecWAXPY(Vec w, PetscScalar a, Vec x, Vec y);</code>         | $w = a * x + y$             |
| <code>VecAXPBY(Vec y, PetscScalar a, PetscScalar b, Vec x);</code> | $y = a * x + b * y$         |
| <code>VecScale(Vec x, PetscScalar a);</code>                       | $x = a * x$                 |
| <code>VecDot(Vec x, Vec y, PetscScalar *r);</code>                 | $r = \bar{x}' * y$          |
| <code>VecTDot(Vec x, Vec y, PetscScalar *r);</code>                | $r = x' * y$                |
| <code>VecNorm(Vec x, NormType type, PetscReal *r);</code>          | $r =   x  _{type}$          |
| <code>VecSum(Vec x, PetscScalar *r);</code>                        | $r = \sum x_i$              |
| <code>VecCopy(Vec x, Vec y);</code>                                | $y = x$                     |
| <code>VecSwap(Vec x, Vec y);</code>                                | $y = x$ while $x = y$       |
| <code>VecPointwiseMult(Vec w, Vec x, Vec y);</code>                | $w_i = x_i * y_i$           |
| <code>VecPointwiseDivide(Vec w, Vec x, Vec y);</code>              | $w_i = x_i / y_i$           |
| <code>VecMDot(Vec x, int n, Vec y[], PetscScalar *r);</code>       | $r[i] = \bar{x}' * y[i]$    |
| <code>VecMTDot(Vec x, int n, Vec y[], PetscScalar *r);</code>      | $r[i] = x' * y[i]$          |
| <code>VecMAXPY(Vec y, int n, PetscScalar *a, Vec x[]);</code>      | $y = y + \sum_i a_i * x[i]$ |
| <code>VecMax(Vec x, int *idx, PetscReal *r);</code>                | $r = \max x_i$              |
| <code>VecMin(Vec x, int *idx, PetscReal *r);</code>                | $r = \min x_i$              |
| <code>VecAbs(Vec x);</code>  | $x_i =  x_i $               |
| <code>VecReciprocal(Vec x);</code>                                 | $x_i = 1/x_i$               |
| <code>VecShift(Vec x, PetscScalar s);</code>                       | $x_i = s + x_i$             |
| <code>VecSet(Vec x, PetscScalar alpha);</code>                     | $x_i = \alpha$              |



# Matrices

## What are PETSc matrices?

- Roughly represent linear operators that belong to the dual of a vector space over a field (e.g.  $\mathbb{R}^n$ )
- In most of the PETSc low-level implementations, each process logically owns a submatrix of contiguous rows

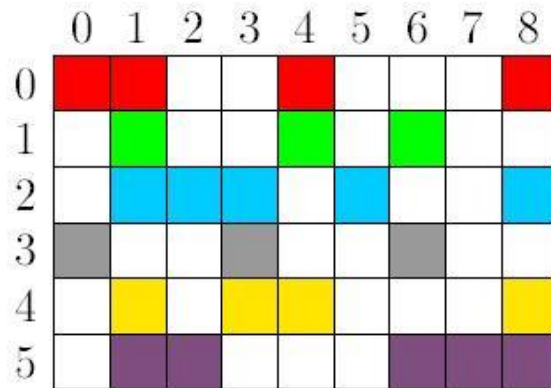
## Features

- Supports many storage formats
  - AIJ, BAIJ, SBAIJ, DENSE, VIENNACL, CUSP (on GPU) ...
- Data structures for many external packages
  - MUMPS (parallel), SuperLU\_dist (parallel), SuperLU, UMFPack
- Hidden communications in parallel matrix assembly
- Matrix operations are defined from a common interface
- Shell matrices via user defined MatMult and other ops





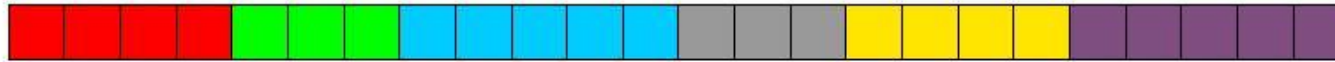
# Matrices



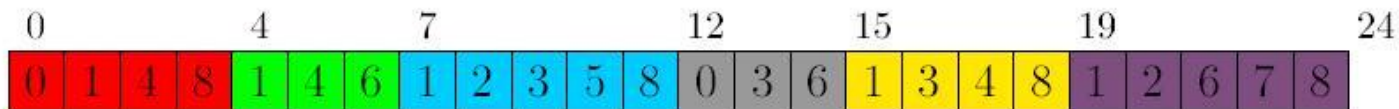
The default matrix representation within PETSc is the general sparse **AIJ format** (Yale sparse matrix or Compressed Sparse Row, CSR)

- The nonzero elements are stored by rows
- Array of corresponding column numbers
- Array of pointers to the beginning of each row

value



index



row pointer





## Matrix memory preallocation

- PETSc matrix creation is very flexible: No preset sparsity pattern
- Memory **preallocation** is critical for achieving **good performance** during matrix assembly, as this reduces the number of allocations and copies required during the assembling process. Remember: malloc is very expensive (run your code with `-memory_info`, `-malloc_log`)
- Private representations of PETSc sparse matrices are dynamic data structures: **additional nonzeros can be freely added** (if no preallocation has been explicitly provided).
- No preset sparsity pattern, any processor can set any element: potential for lots of malloc calls
- Dynamically adding many nonzeros
  - requires additional memory allocations
  - requires copies
  - **kills performances!**



# Preallocation of a parallel sparse matrix

Each process **logically owns** a matrix subset of contiguously numbered global rows. Each subset consists of two sequential matrices corresponding to diagonal and **off-diagonal** parts.

|    |    |    |    |    |    |    |    |    |
|----|----|----|----|----|----|----|----|----|
| P0 | 1  | 2  | 0  | 0  | 3  | 0  | 0  | 4  |
|    | 0  | 5  | 6  | 7  | 0  | 0  | 8  | 0  |
|    | 9  | 0  | 10 | 11 | 0  | 0  | 12 | 0  |
| P1 | 13 | 0  | 14 | 15 | 16 | 17 | 0  | 0  |
|    | 0  | 18 | 0  | 19 | 20 | 21 | 0  | 0  |
|    | 0  | 0  | 0  | 22 | 23 | 0  | 24 | 0  |
| P2 | 25 | 26 | 27 | 0  | 0  | 28 | 29 | 0  |
|    | 30 | 0  | 0  | 31 | 32 | 33 | 0  | 34 |

## Process 0

$dnz=2, onz=2$

$dnnz[0]=2, onnz[0]=2$

$dnnz[1]=2, onnz[1]=2$

$dnnz[2]=2, onnz[2]=2$

## Process 1

$dnz=3, onz=2$

$dnnz[0]=3, onnz[0]=2$

$dnnz[1]=3, onnz[1]=1$

$dnnz[2]=2, onnz[2]=1$

## Process 2

$dnz=1, onz=4$

$dnnz[0]=1, onnz[0]=4$

$dnnz[1]=1, onnz[1]=4$



# Numerical Matrix Operations

| Function Name  | Operation                                 |
|--|---|
| <code>MatAXPY(Mat Y, PetscScalar a, Mat X, MatStructure);</code> | $Y = Y + a * X$                           |
| <code>MatMult(Mat A, Vec x, Vec y);</code>                       | $y = A * x$                               |
| <code>MatMultAdd(Mat A, Vec x, Vec y, Vec z);</code>             | $z = y + A * x$                           |
| <code>MatMultTranspose(Mat A, Vec x, Vec y);</code>              | $y = A^T * x$                             |
| <code>MatMultTransposeAdd(Mat A, Vec x, Vec y, Vec z);</code>    | $z = y + A^T * x$                         |
| <code>MatNorm(Mat A, NormType type, double *r);</code>           | $r = \ A\ _{type}$                        |
| <code>MatDiagonalScale(Mat A, Vec l, Vec r);</code>              | $A = \text{diag}(l) * A * \text{diag}(r)$ |
| <code>MatScale(Mat A, PetscScalar a);</code>                     | $A = a * A$                               |
| <code>MatConvert(Mat A, MatType type, Mat *B);</code>            | $B = A$                                   |
| <code>MatCopy(Mat A, Mat B, MatStructure);</code>                | $B = A$                                   |
| <code>MatGetDiagonal(Mat A, Vec x);</code>                       | $x = \text{diag}(A)$                      |
| <code>MatTranspose(Mat A, MatReuse, Mat* B);</code>              | $B = A^T$                                 |
| <code>MatZeroEntries(Mat A);</code>                              | $A = 0$                                   |
| <code>MatShift(Mat Y, PetscScalar a);</code>                     | $Y = Y + a * I$                           |



# Matrix multiplication (MatMult)

$$y \leftarrow A * x_A + B * x_B$$

- $x_B$  needs to be communicated
- $A * x_A$  can be computed in the meantime

## Algorithm

- Initiate asynchronous sends/receives for  $x_B$
- compute  $A * x_A$
- make sure  $x_B$  is in
- compute  $B * x_B$

Due to the splitting of the matrix storage into A (diag) and B (off-diag) part, code for the sequential case can be reused.

Off-diagonal block  
has off-processor connections

A

B

Diagonal block has on-processor  
connections



## Sparse Matrices and Linear Solvers

- Solve a linear system  $A x = b$  using the Gauss Elimination method can be very time-resource consuming
- Alternatives to direct solvers are iterative solvers
- Convergence of the succession is not always guaranteed
- Possibly much faster and less memory consuming
- Basic iteration:  $y \leftarrow A x$  executed once  $x$  iteration
- Also needed a good preconditioner:  $B \approx A^{-1}$



## Iterative solver basics

- **KSP** (*Krylov SPace Methods*) objects are used for solving linear systems by means of iterative methods.
- Convergence can be improved by using a suitable **PC** object (preconditioner).
- Almost all iterative methods are implemented.
- Classical iterative methods (not belonging to KSP solvers) are classified as preconditioners
- Direct solution for parallel square matrices available through external solvers (MUMPS, SuperLU\_dist). Petsc provides a built-in LU serial solver.
- Many KSP options can be controlled by command line
- Tolerances, convergence and divergence reason
- Custom monitors and convergence tests



# Solver Types

| Method  | KSPType       | Options Database Name |
|---|---------------|-----------------------|
| Richardson                                    | KSPRICHARDSON | richardson            |
| Chebyshev                                     | KSPCHEBYSHEV  | chebyshev             |
| Conjugate Gradient [12]                       | KSPCG         | cg                    |
| BiConjugate Gradient                          | KSPBICG       | bicg                  |
| Generalized Minimal Residual [16]             | KSPGMRES      | gmres                 |
| Flexible Generalized Minimal Residual         | KSPFGMRES     | fgmres                |
| Deflated Generalized Minimal Residual         | KSPDGMRES     | dgmres                |
| Generalized Conjugate Residual                | KSPGCR        | gcr                   |
| BiCGSTAB [19]                                 | KSPBCGS       | bcgs                  |
| Conjugate Gradient Squared [18]               | KSPCGS        | cgs                   |
| Transpose-Free Quasi-Minimal Residual (1) [8] | KSPTFQMR      | tfqmr                 |
| Transpose-Free Quasi-Minimal Residual (2)     | KSPTCQMR      | tcqmr                 |
| Conjugate Residual                            | KSPCR         | cr                    |
| Least Squares Method                          | KSPLSQR       | lsqr                  |
| Shell for no KSP method                       | KSPPREONLY    | preonly               |





## Preconditioner types

| Method                         | PCType      | Options Database Name |
|--------------------------------|-------------|-----------------------|
| Jacobi                         | PCJACOBI    | jacobi                |
| Block Jacobi                   | PCBJACOBI   | bjacobi               |
| SOR (and SSOR)                 | PCSOR       | sor                   |
| SOR with Eisenstat trick       | PCEISENSTAT | eisenstat             |
| Incomplete Cholesky            | PCICC       | icc                   |
| Incomplete LU                  | PCILU       | ilu                   |
| Additive Schwarz               | PCASM       | asm                   |
| Algebraic Multigrid            | PCGAMG      | gamg                  |
| Linear solver                  | PCKSP       | ksp                   |
| Combination of preconditioners | PCCOMPOSITE | composite             |
| LU                             | PCLU        | lu                    |
| Cholesky                       | PCCHOLESKY  | cholesky              |
| No preconditioning             | PCNONE      | none                  |
| Shell for user-defined PC      | PCSHELL     | shell                 |



## Factorization preconditioner

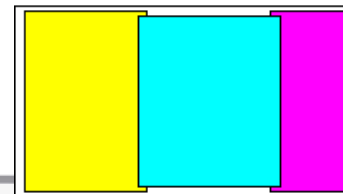
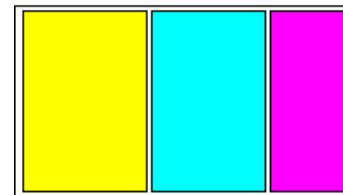
- Exact factorization:  $A = LU$
- Inexact factorization:  $A \approx M = \underline{L} \underline{U}$  where  $\underline{L}$ ,  $\underline{U}$  obtained by throwing away the 'fill-in' during the factorization process (sparsity pattern of  $M$  is the same as  $A$ )
- Application of the preconditioner (that is, solve  $Mx = y$ ) approx same cost as matrix-vector product  $y \leftarrow Ax$
- Factorization preconditioners are sequential
- PCICC: symmetric matrix, PCILU: nonsymmetric matrix



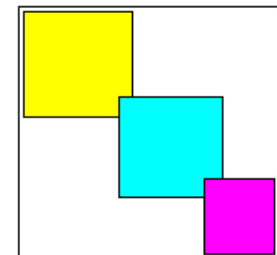
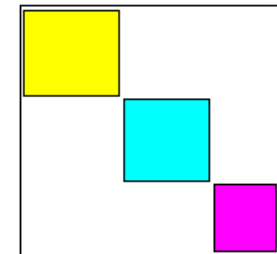
# Parallel preconditioners

- Factorization preconditioners are sequential
- We can use them in parallel as a subpreconditioner of a parallel preconditioner as Block Jacobi or Additive Schwarz Methods (ASM)
- Each processor has its own block(s) to work with
- Block Jacobi is fully parallel, ASM requires communications between neighbours
- ASM can be more robust than Block Jacobi and have better convergence properties

Domain partitioning



Matrix blocks





# Profiling and preliminary tests on KNL



## Profiling and performance tuning

- Integrated profiling of:
  - time
  - floating-point performance
  - memory usage
  - communication
- User-defined events
- Profiling by stages of an application

**`-log_view`** - Prints an ASCII version of performance data at program's conclusion. These statistics are comprehensive and concise and require little overhead; thus, `-log_view` is intended as the primary means of monitoring the performance of PETSc codes.





## Petsc benchmark: ex56 (3D linear elasticity)

- 3D, tri-linear quadrilateral (Q1), displacement finite element formulation of linear elasticity.  $E=1.0$ ,  $\nu=0.25$ .
- Unit box domain with Dirichlet boundary condition on the  $y=0$  side only.
- Load of 1.0 in  $x + 2y$  direction on all nodes (not a true uniform load).
- $np$  = number of processes;  $np^{1/3}$  must be integer
- $ne$  = number of elements in the  $x,y,z$  direction;  $(ne+1)\%(np^{1/3})$  must equal zero
- Default solver: GMRES + BLOCK\_JACOBI + ILU(0)



## Petsc benchmark: ex56 (3D linear elasticity)

|  | Command  | Time    |
|--|--|---------|
| Broadwell (ne=80, np=27)                           | <code>mpirun -np 27 ./ex56 -ne 80 -log_view</code>                       | 14.2 s  |
| KNL (ne=80, np=27) + DRAM                          | <code>mpirun -np 64 numactl --membind=0,1 ./ex56 -ne 79 -log_view</code> | 38.61 s |
| KNL (ne=80, np=27) +<br>MCDRAM=FLAT +<br>NUMA=SNC2 | <code>mpirun -np 27 numactl --membind=2,3 ./ex56 -ne 80 -log_view</code> | 12.12 s |
| KNL (ne=79, np=64) +<br>MCDRAM=FLAT +<br>NUMA=SNC2 | <code>mpirun -np 64 numactl --membind=2,3 ./ex56 -ne 79 -log_view</code> | 10.90 s |
| KNL (ne=80, np=27) +<br>MCDRAM=CACHE               | <code>mpirun -np 27 ./ex56 -ne 80 -log_view</code>                       | 14.12 s |
| KNL (ne=80, np=64) +<br>MCDRAM=CACHE               | <code>mpirun -np 64 ./ex56 -ne 79 -log_view</code>                       | 12.50 s |





Thank you for the attention