

DOE ACTS Collection Workshop  
**Robust and High Performance Tools for Scientific Computing**

# ScaLAPACK

**Osni Marques**

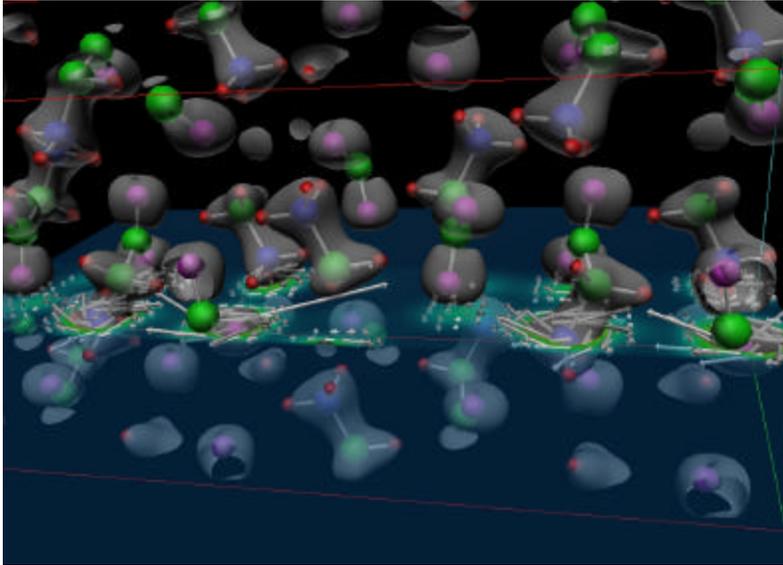
Lawrence Berkeley National Laboratory (LBNL)

(*oamarques@lbl.gov*, *http://www.nersc.gov/~osni*)

# Outline

- Motivation
- ScaLAPACK: *software structure*
  - Basic Linear Algebra Subprograms (BLAS)
  - Linear Algebra PACKage (LAPACK)
  - Basic Linear Algebra Communication Subprograms (BLACS)
  - Parallel BLAS (PBLAS)
- ScaLAPACK: *details*
  - Data layout
  - Array descriptors
  - Error handling
  - Performance
- Hands-on

# ScaLAPACK: Applications



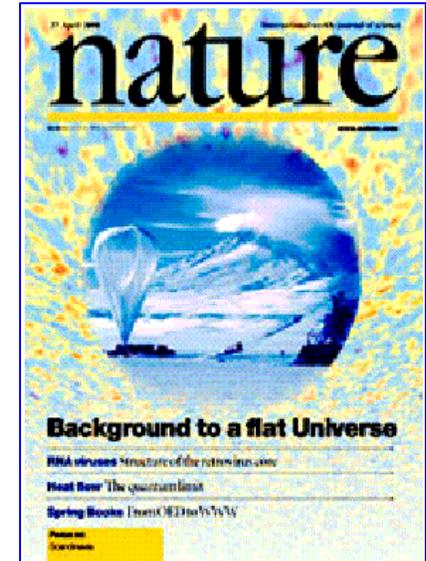
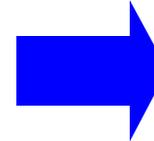
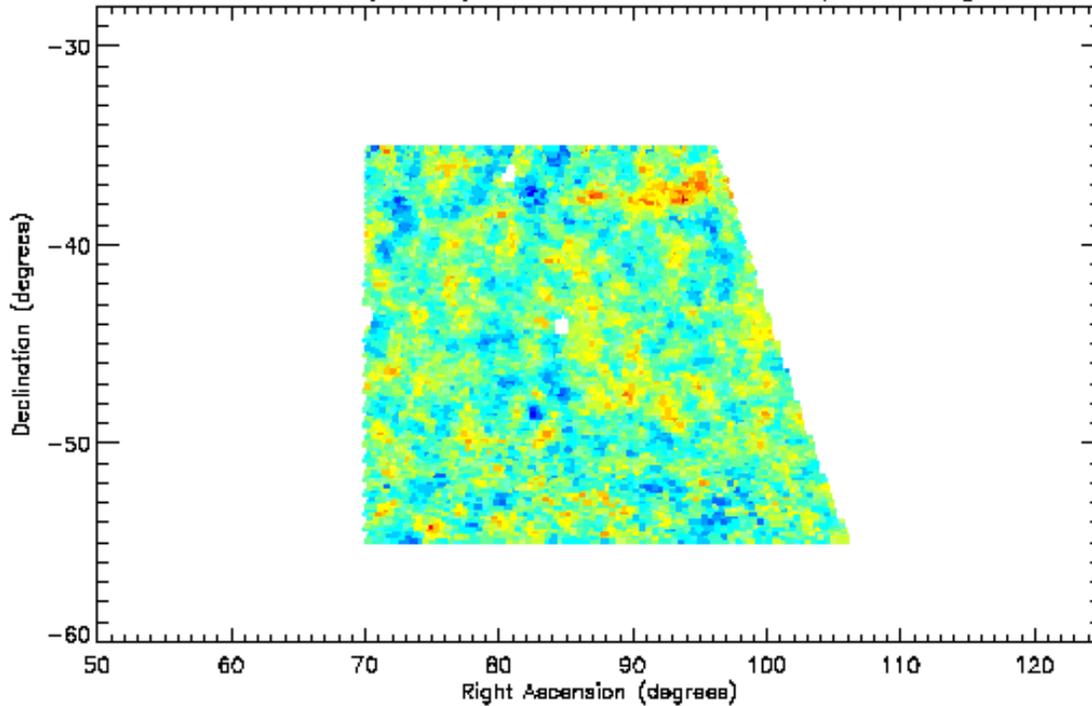
*Induced current (white arrows) and charge density (colored plane and gray surface) in crystallized glycine due to an external field; courtesy of Louie, Yoon, Pfrommer and Canning (UCB and LBNL).*

*Advanced Computational Research in Fusion (SciDAC Project, PI Mitch Pindzola). Point of contact: Dario Mitnik (Dept. of Physics, Rollins College). Mitnik attended the workshop on the ACTS Toolkit in September 2000. Since then he has been actively using some of the ACTS tools, in particular ScaLAPACK, for which he has provided insightful feedback. Dario is currently working on the development, testing and support of new scientific simulation codes related to the study of atomic dynamics using time-dependent close coupling lattice and time-independent methods. He reports that this work could not be carried out in sequential machines and that ScaLAPACK is fundamental for the parallelization of these codes.*

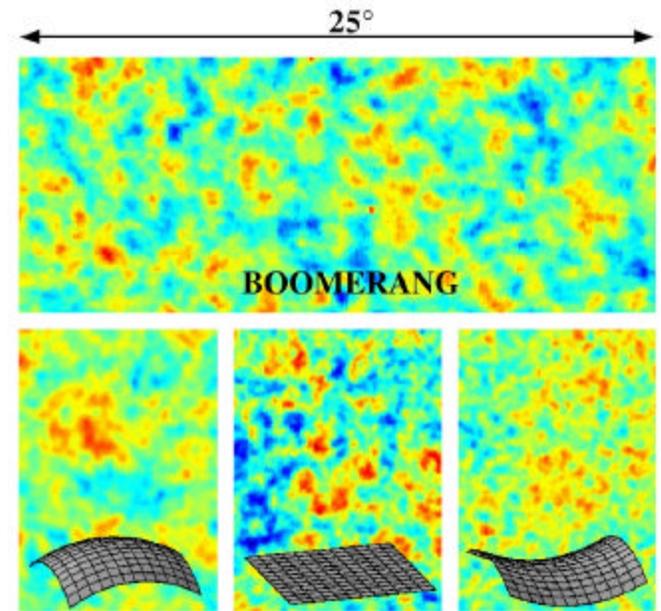
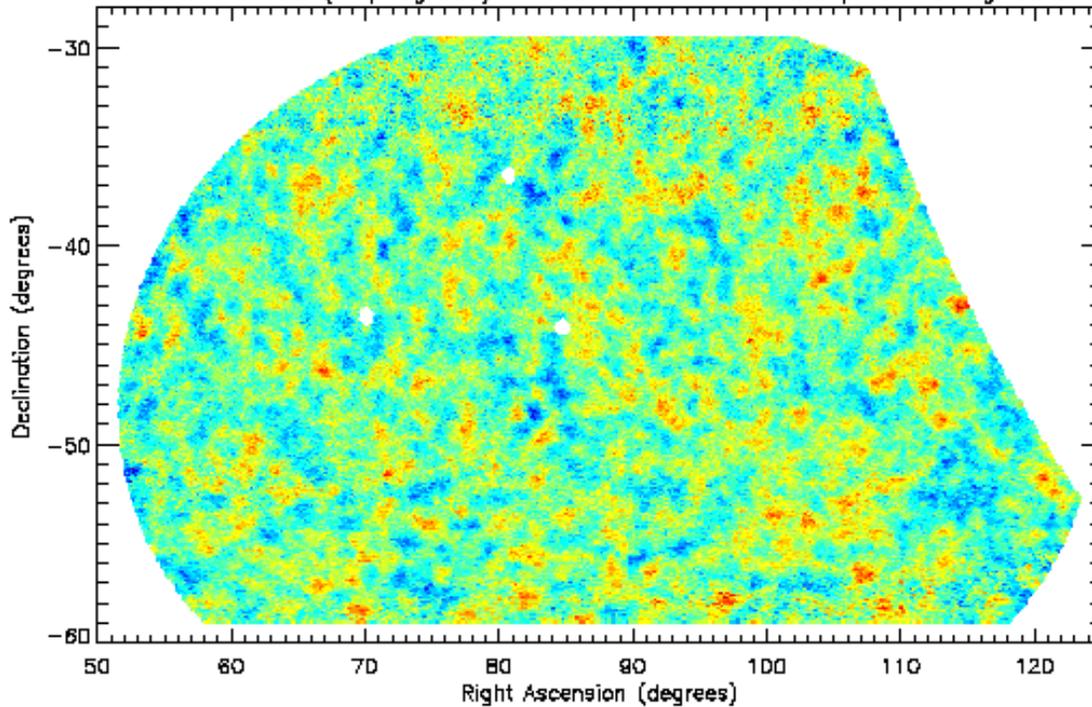
# ScaLAPACK: CMB Analysis

- The CMB (Cosmic Microwave Background) is the faint echo of the Big Bang.
- The statistics of the tiny variations in the CMB allows the determination of the fundamental parameters of cosmology to the percent level or better.
- MADCAP (Microwave Anisotropy Dataset Computational Analysis Package, see <http://www.nersc.gov/~borrill/cmb/madcap.html>):
  - Makes maps from observations of the CMB and calculates their angular power spectra.
  - Calculations are dominated by the solution of linear systems of the form  $M = A^{-1}B$  for dense  $n \times n$  matrices  $A$  and  $B$  scaling as  $O(n^3)$  in flops.
- On NERSC's Cray T3E (original code):
  - Cholesky factorization and triangular solve, 70-80% peak performance.
  - Solution of systems with  $n \sim 10^4$  using tens of processors.
  - The results demonstrated that the Universe is spatially flat (overall) comprising 70% dark energy, 25% dark matter, and only 5% ordinary matter.
- On NERSC's IBM SP:
  - Porting was trivial but tests showed only 20-30% peak performance.
  - Code rewritten to use Cholesky factorization, triangular matrix inversion and triangular matrix multiplication (one-day work because of the completeness and coherence of ScaLAPACK), performance increased to 50-60% peak.
  - Solution of previously intractable systems with  $n \sim 10^5$  using hundreds of processors.

BOOMERanG 2000 (Nature) – 8147 x 14 arcminute pixels using MCURIE



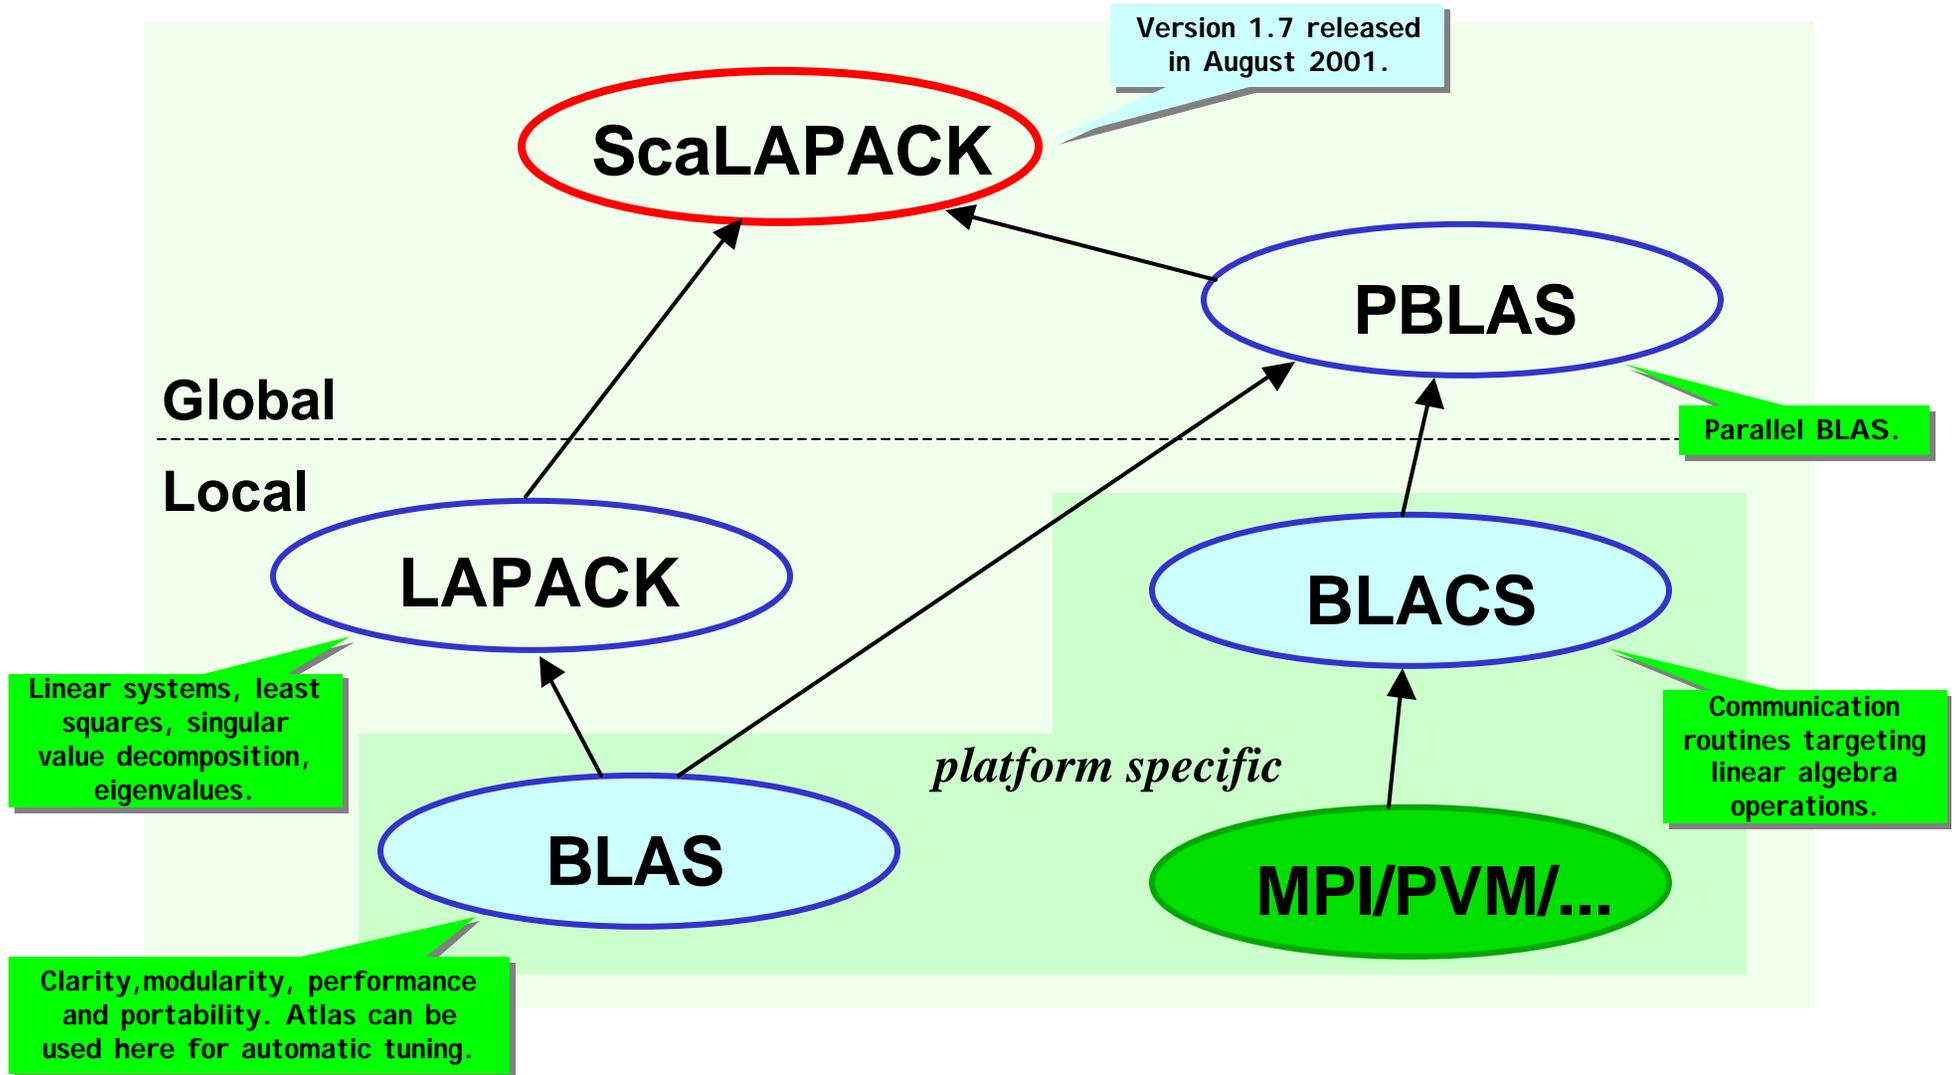
BOOMERanG 2002 (in progress) – 92251 x 7 arcminute pixels using SEABORG



<http://www.physics.ucsb.edu/~boomerang>

# ScaLAPACK: software structure

<http://acts.neresc.gov/scalapack>



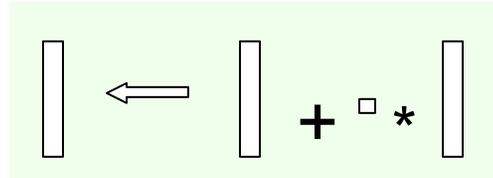
# BLAS

*(Basic Linear Algebra Subroutines)*

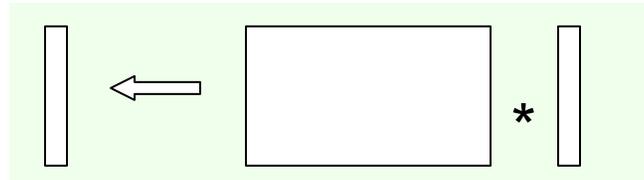
- Clarity: code is shorter and easier to read.
- Modularity: gives programmer larger building blocks.
- Performance: manufacturers (usually) provide tuned machine-specific BLAS.
- Portability: machine dependencies are confined to the BLAS.
- Key to high performance: effective use of memory hierarchy (true on all architectures).

# BLAS: 3 levels

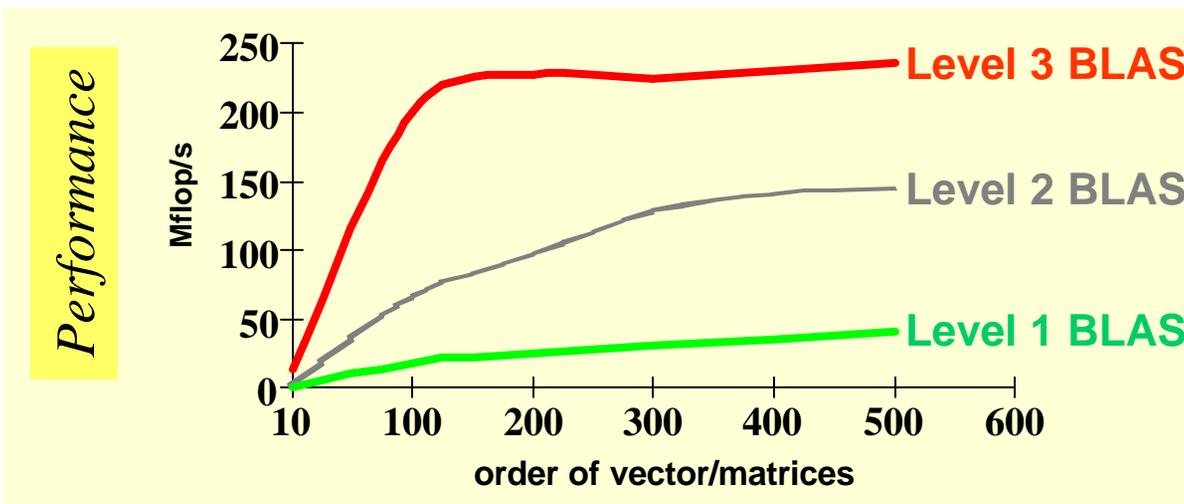
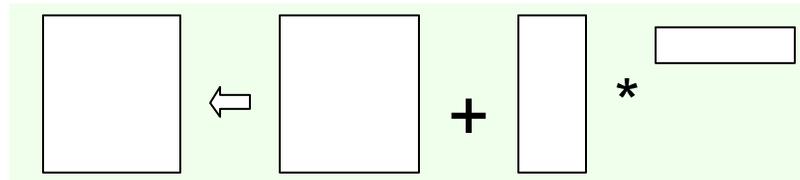
- Level 1 BLAS:  
vector-vector operations.



- Level 2 BLAS:  
matrix-vector operations.



- Level 3 BLAS:  
matrix-matrix operations.



Development of blocked algorithms is important for performance!

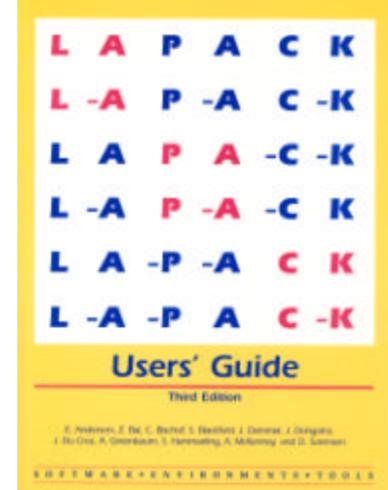
# LAPACK

*(<http://www.netlib.org/lapack>)*

- Linear Algebra library written in Fortran 77 (Fortran 90, C and C++ versions also available).
- Combine algorithms from LINPACK and EISPACK into a single package.
- Efficient on a wide range of computers (RISC, Vector, SMPs).
- User interface similar to LINPACK (Single, Double, Complex, Double Complex).
- Built atop level 1, 2, and 3 BLAS for high performance, clarity, modularity and portability.

# LAPACK: *features*

- Basic problems:
  - Linear systems:  $Ax = b$
  - Least squares:  $\min \|Ax - b\|_2$
  - Singular value decomposition:  $A = U \Sigma V^T$
  - Eigenvalues and eigenvectors:  $Az = \lambda z$ ,  $Az = \lambda Bz$
- LAPACK does not provide routines for structured problems or general sparse matrices (i.e. sparse storage formats such as compressed-row, -column, -diagonal, skyline ...).
- LAPACK Users' Guide, Third Edition (1999)



# BLACS

*(Basic Linear Algebra Communication Subroutines)*

- A design tool, they are a conceptual aid in design and coding.
- Associate widely recognized mnemonic names with communication operations. This improves:
  - program readability
  - self-documenting quality of the code.
- Promote efficiency by identifying frequently occurring operations of linear algebra which can be optimized on various computers.

# BLACS: *basics*

Processes are embedded in a two-dimensional grid.

Example: 3x4 grid

	0	1	2	3
0	0	1	2	3
1	4	5	6	7
2	8	9	10	11

## BLACS: *scoped operations*

An operation which involves more than one sender and one receiver is called a *scoped operation*.

Using a 2D-grid, there are 3 natural scopes:

Scope	Meaning
Row	All processes in a process row participate.
Column	All processes in a process column participate.
All	All processes in the process grid participate.

# BLACS: *communication routines*

## *Send/Receive:*

**\_xxSD2D**( ICTXT , [ UPLO , DIAG ] , M , N , A , LDA , RDEST , CDEST )

**\_xxRV2D**( ICTXT , [ UPLO , DIAG ] , M , N , A , LDA , RSRC , CSRC )

<b>_ (Data type)</b>	<b>xx (Matrix type)</b>
I: Integer, S: Real, D: Double Precision, C: Complex, Z: Double Complex.	GE: General rectangular matrix TR: Trapezoidal matrix

## *Broadcast:*

**\_xxBS2D**( ICTXT , **SCOPE** , **TOP** , [ UPLO , DIAG ] , M , N , A , LDA )

**\_xxBR2D**( ICTXT , **SCOPE** , **TOP** , [ UPLO , DIAG ] , M , N , A , LDA , RSRC , CSRC )

<b>SCOPE</b>	<b>TOP</b>
'Row'	' ' (default)
'Column'	'Increasing Ring'
'All'	'1-tree' ...

# BLACS: *global combine operations*

- Perform element-wise SUM, |MAX|, |MIN|, operations on triangular matrices:
  - `_GSUM2D( ICTXT , SCOPE , TOP , M , N , A , LDA , RDEST , CDEST )`
  - `_GAMX2D( ICTXT , SCOPE , TOP , M , N , A , LDA , RA , CA , RCFLAG , RDEST , CDEST )`
  - `_GAMN2D( ICTXT , SCOPE , TOP , M , N , A , LDA , RA , CA , RCFLAG , RDEST , CDEST )`
- RDEST = -1 indicates that the result of the operation should be left on all processes selected by SCOPE.
- For |MAX|, |MIN|, when RCFLAG = -1, RA and CA are not referenced; otherwise RA and CA are set on output with the coordinates of the process owning the corresponding maximum (or minimum) element in absolute value of A.

# BLACS: *context*

- The BLACS context is the BLACS mechanism for partitioning communication space.
- A message in a context cannot be sent or received in another context.
- The context allows the user to
  - create arbitrary groups of processes
  - create multiple overlapping and/or disjoint grids
  - isolate each process grid so that grids do not interfere with each other
- BLACS context ? MPI communicator

# BLACS: *example*

See <http://www.netlib.org/blacs> for more information.

```
* Get system information
CALL BLACS_PINFO( IAM, NPROCS )
* If underlying system needs additional setup, do it now
IF( NPROCS.LT.1 ) THEN
  IF( IAM.EQ.0 ) NPROCS = 4
  CALL BLACS_SETUP( IAM, NPROCS )
END IF
* Get default system context
CALL BLACS_GET( 0, 0, ICTXT )
□
* Define 1 x (NPROCS/2+1) process grid
NPROW = 1
NPCOL = NPROCS / 2 + 1
CALL BLACS_GRIDINIT( ICTXT, 'Row', NPROW, NPCOL )
CALL BLACS_GRIDINFO( ICTXT, NPROW, NPCOL, MYROW, MYCOL )
* If I'm not in the grid, go to end of program
IF( MYROW.NE.-1 ) THEN
  IF( MYROW.EQ.0 .AND. MYCOL.EQ.0 ) THEN
    CALL DGESD2D( ICTXT, 5, 1, X, 5, 1, 0 )
  ELSE IF( MYROW.EQ.1 .AND. MYCOL.EQ.0 ) THEN
    CALL DGERV2D( ICTXT, 5, 1, Y, 5, 0, 0 )
  END IF
  □
  CALL BLACS_GRIDEXIT( ICTXT )
END IF
CALL BLACS_EXIT( 0 )
END
```

*(out) uniquely identifies each process*  
*(out) number of processes available*

*(in) uniquely identifies each process*  
*(in) number of processes available*  
*(in) integer handle indicating the context*  
*(in) use (default) system context*  
*(out) BLACS context (see next slide)*

*(output) process row and column coordinate*

*send X to process (1,0)*  
*receive X from process (0,0)*

# PBLAS

*(Parallel Basic Linear Algebra Subroutines)*

- Similar to the BLAS in portability, functionality and naming.
- Built atop the BLAS and BLACS
- Provide global view of matrix

```
CALL DGEXXX( M, N, A( IA, JA ), LDA, ... )
```

BLAS

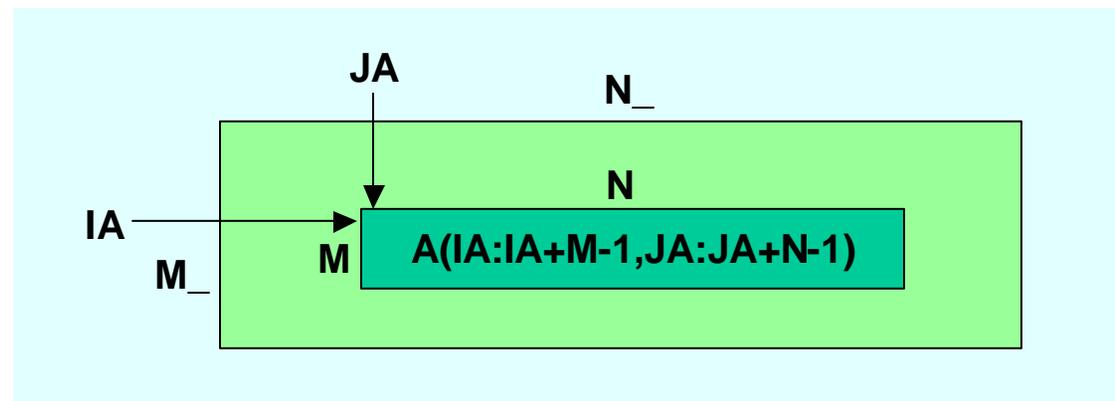
```
CALL PDGEXXX( M, N, A, IA, JA, DESCA, ... )
```

PBLAS

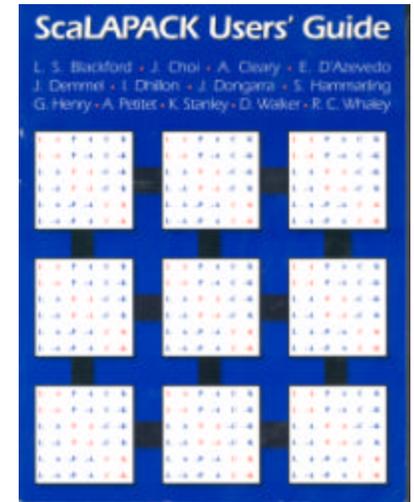
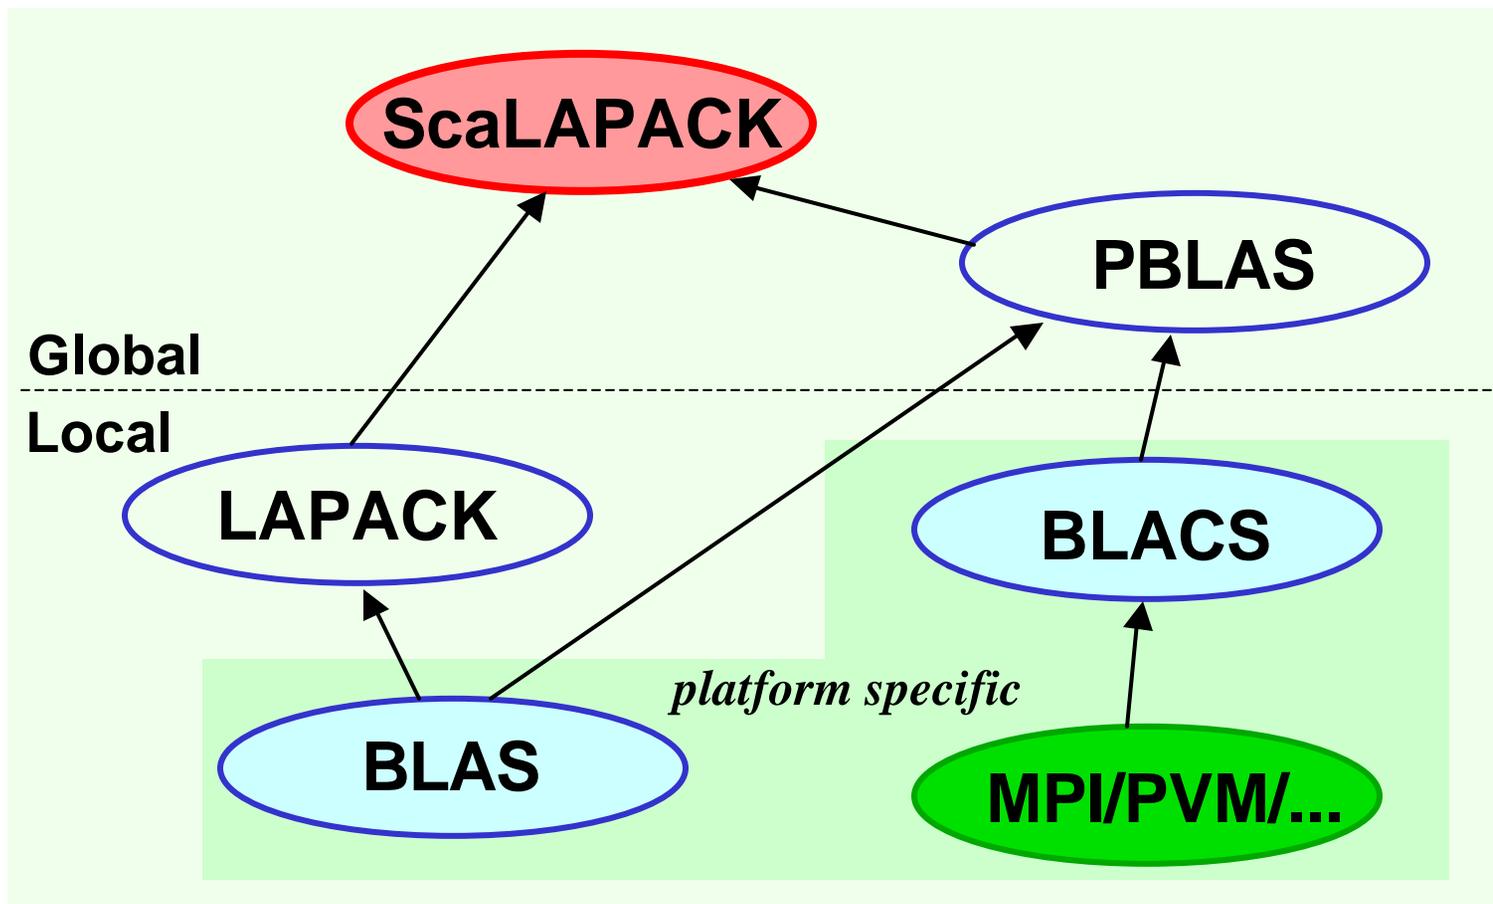
Array descriptor  
(see next slides)

# PBLAS: *levels and view of the operands*

- Levels:
  - Level 1: vector-vector operations.
  - Level 2: matrix-vector operations.
  - Level 3: matrix-matrix operations.
- Global view of the matrix operands, allowing global addressing of distributed matrices (hiding complex local indexing)



# ScaLAPACK: *structure of the software*

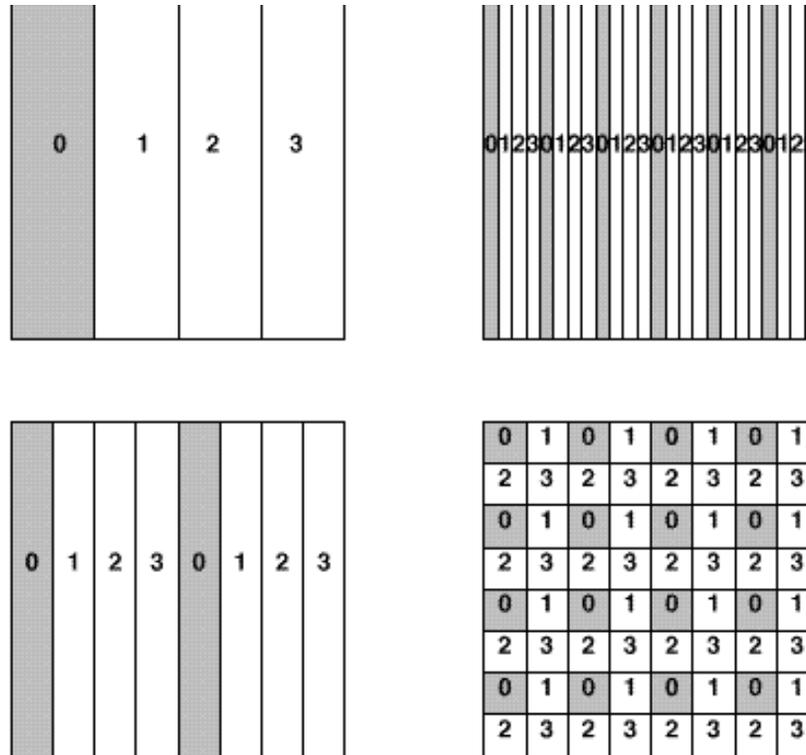


# ScaLAPACK: *goals*

- Efficiency
  - Optimized computation and communication engines
  - Block-partitioned algorithms (Level 3 BLAS) for good node performance
- Reliability
  - Whenever possible, use LAPACK algorithms and error bounds.
- Scalability
  - As the problem size and number of processors grow
  - Replace LAPACK algorithm that did not scale (new ones into LAPACK)
- Portability
  - Isolate machine dependencies to BLAS and the BLACS
- Flexibility
  - Modularity: build rich set of linear algebra tools (BLAS, BLACS, PBLAS)
- Ease-of-Use
  - Calling interface similar to LAPACK

# ScaLAPACK: *data layouts*

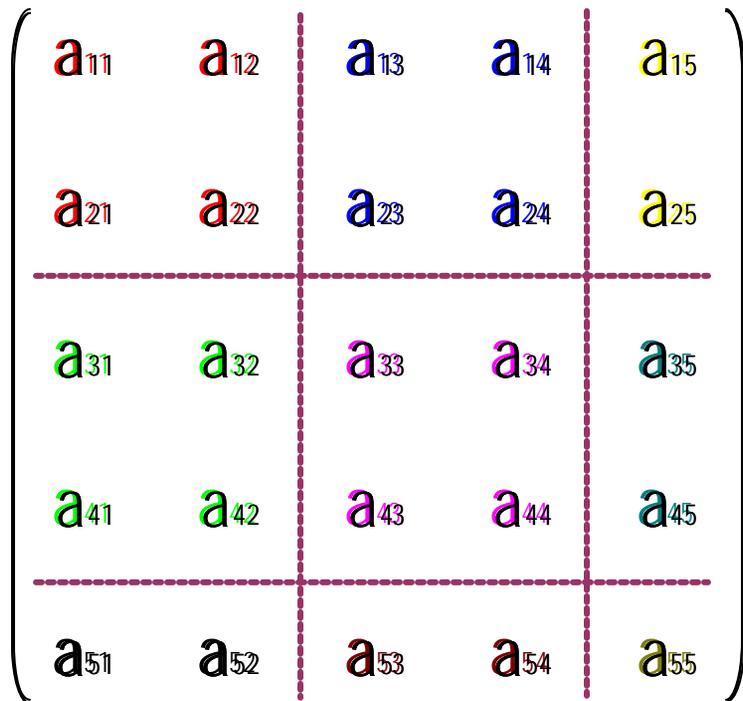
- 1D block and cyclic column distributions



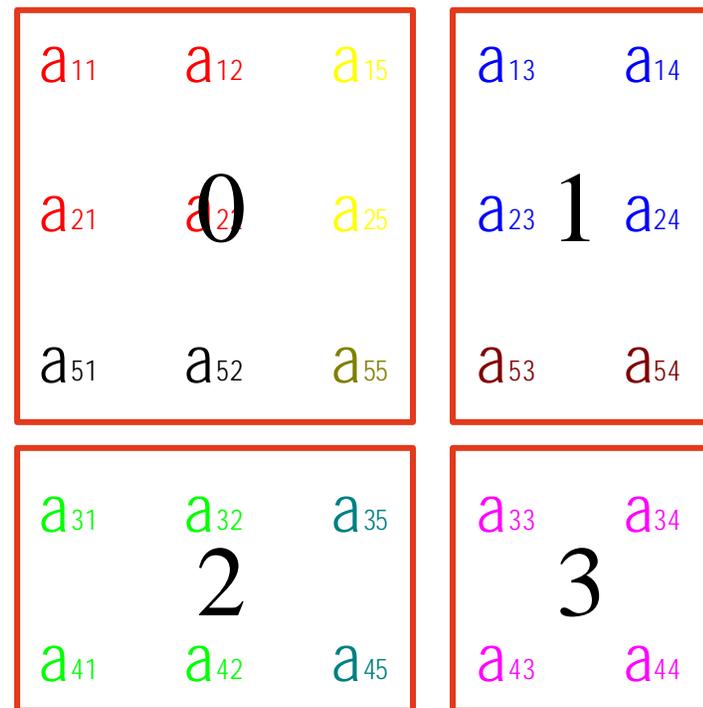
- 1D block-cycle column and 2D block-cyclic distribution
- 2D block-cyclic used in ScaLAPACK for dense matrices

# ScaLAPACK: 2D Block-Cyclic Distribution

5x5 matrix partitioned in 2x2 blocks



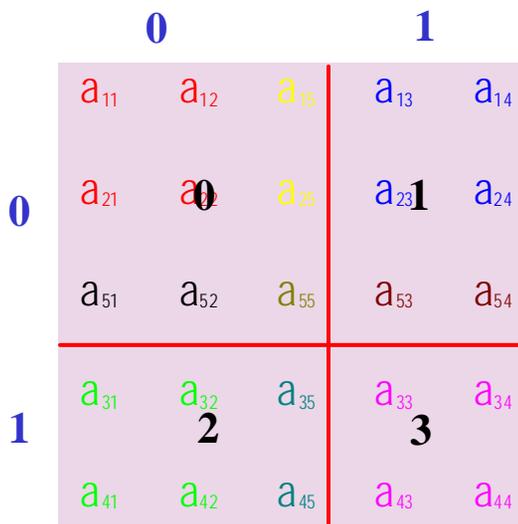
2x2 process grid point of view



# 2D Block-Cyclic Distribution

```

? 1.1  1.2  1.3  1.4  1.5?
?? 2.1  2.2  2.3  2.4  2.5?
?? 3.1  ? 3.2  3.3  3.4  3.5?
?? 4.1  ? 4.2  ? 4.3  4.4  4.5?
?? 5.1  ? 5.2  ? 5.3  ? 5.4  5.5?
    
```



```

□
CALL BLACS_GRIDINFO( ICTXT, NPROW, NPCOL, MYROW, MYCOL )

IF ( MYROW.EQ.0 .AND. MYCOL.EQ.0 ) THEN
  A(1) = 1.1; A(2) = -2.1; A(3) = -5.1;
  A(1+LDA) = 1.2; A(2+LDA) = 2.2; A(3+LDA) = -5.2;
  A(1+2*LDA) = 1.5; A(2+3*LDA) = 2.5; A(3+4*LDA) = -5.5;
ELSE IF ( MYROW.EQ.0 .AND. MYCOL.EQ.1 ) THEN
  A(1) = 1.3; A(2) = 2.3; A(3) = -5.3;
  A(1+LDA) = 1.4; A(2+LDA) = 2.4; A(3+LDA) = -5.4;
ELSE IF ( MYROW.EQ.1 .AND. MYCOL.EQ.0 ) THEN
  A(1) = -3.1; A(2) = -4.1;
  A(1+LDA) = -3.2; A(2+LDA) = -4.2;
  A(1+2*LDA) = 3.5; A(2+3*LDA) = 4.5;
ELSE IF ( MYROW.EQ.1 .AND. MYCOL.EQ.1 ) THEN
  A(1) = 3.3; A(2) = -4.3;
  A(1+LDA) = 3.4; A(2+LDA) = 4.4;
END IF
□
CALL PDGESVD( JOBU, JOBVT, M, N, A, IA, JA, DESCA, S, U, IU,
              JU, DESCU, VT, IVT, JVT, DESCVT, WORK, LWORK,
              INFO )
□
    
```

LDA is the leading dimension of the local array (see next slides)

Array descriptor for A (see next slides)

## *2D Block-Cyclic Distribution*

- Ensures good load balance ? performance and scalability (analysis of many algorithms to justify this layout).
- Encompasses a large number of data distribution schemes (but not all).
- Needs redistribution routines to go from one distribution to the other.
- See *<http://acts.nerisc.gov/scalapack/hands-on/datadist.html>*

# ScaLAPACK: *array descriptors*

- Each global data object is assigned an *array descriptor*.
- The *array descriptor*:
  - Contains information required to establish mapping between a global array entry and its corresponding process and memory location (uses concept of BLACS context).
  - Is differentiated by the DTYPE\_ (first entry) in the descriptor.
  - Provides a flexible framework to easily specify additional data distributions or matrix types.
- User must distribute all global arrays prior to the invocation of a ScaLAPACK routine, for example:
  - Each process generates its own submatrix.
  - One processor reads the matrix from a file and send pieces to other processors (may require message-passing for this).

# *Array Descriptor for Dense Matrices*

DESC_()	Symbolic Name	Scope	Definition
1	DTYPE_A	(global)	Descriptor type DTYPE_A=1 for dense matrices.
2	CTXT_A	(global)	BLACS context handle.
3	M_A	(global)	Number of rows in global array A.
4	N_A	(global)	Number of columns in global array A.
5	MB_A	(global)	Blocking factor used to distribute the rows of array A.
6	NB_A	(global)	Blocking factor used to distribute the columns of array A.
7	RSRC_A	(global)	Process row over which the first row of the array A is distributed.
8	CSRC_A	(global)	Process column over which the first column of the array A is distributed.
9	LLD_A	(local)	Leading dimension of the local array.

# *Array Descriptor for Narrow Band Matrices*

DESC_()	Symbolic Name	Scope	Definition
1	DTYPE_A	(global)	Descriptor type DTYPE_A=501 for 1 x Pc process grid for band and tridiagonal matrices block-column distributed.
2	CTXT_A	(global)	BLACS context handle.
3	N_A	(global)	Number of columns in global array A.
4	NB_A	(global)	Blocking factor used to distribute the columns of array A.
5	CSRC_A	(global)	Process column over which the first column of the array A is distributed.
6	LLD_A	(local)	Leading dimension of the local array. For the tridiagonal subroutines, this entry is ignored.
7	?	?	Unused, reserved.

# *Array Descriptor for Right Hand Sides for Narrow Band Linear Solvers*

DESC_()	Symbolic Name	Scope	Definition
1	DTYPE_B	(global)	Descriptor type DTYPE_B=502 for Pr x 1 process grid for block-row distributed matrices .
2	CTXT_B	(global)	BLACS context handle.
3	M_B	(global)	Number of rows in global array B
4	MB_B	(global)	Blocking factor used to distribute the rows of array B.
5	RSRC_B	(global)	Process row over which the first row of the array B is distributed.
6	LLD_B	(local)	Leading dimension of the local array. For the tridiagonal subroutines, this entry is ignored.
7	?	?	Unused, reserved.

# ScaLAPACK: *Functionality*

$Ax = b$	Simple Driver	Expert Driver	Factor	Solve	Inversion	Conditioning Estimator	Iterative Refinement
Triangular				x	x	x	x
SPD	x	x	x	x	x	x	x
SPD Banded	x		x	x			
SPD Tridiagonal	x		x	x			
General	x	x	x	x	x	x	x
General Banded	x		x	x			
General Tridiagonal	x		x	x			
Least Squares	x		x	x			
GQR			x				
GRQ			x				
$Ax = ?x$ or $Ax = ?Bx$	Simple Driver	Expert Driver	Reduction	Solution			
Symmetric	x	x	x	x			
General	x	x	x	x			
Generalized BSPD	x		x	x			
SVD			x	x			

# ScaLAPACK: *error handling*

- Driver and computational routines perform *global* and *local* input error-checking.
  - Global checking ? synchronization
  - Local checking ? validity
- No input error-checking is performed on the auxiliary routines.
- If an error is detected in a PBLAS or BLACS routine program execution stops.

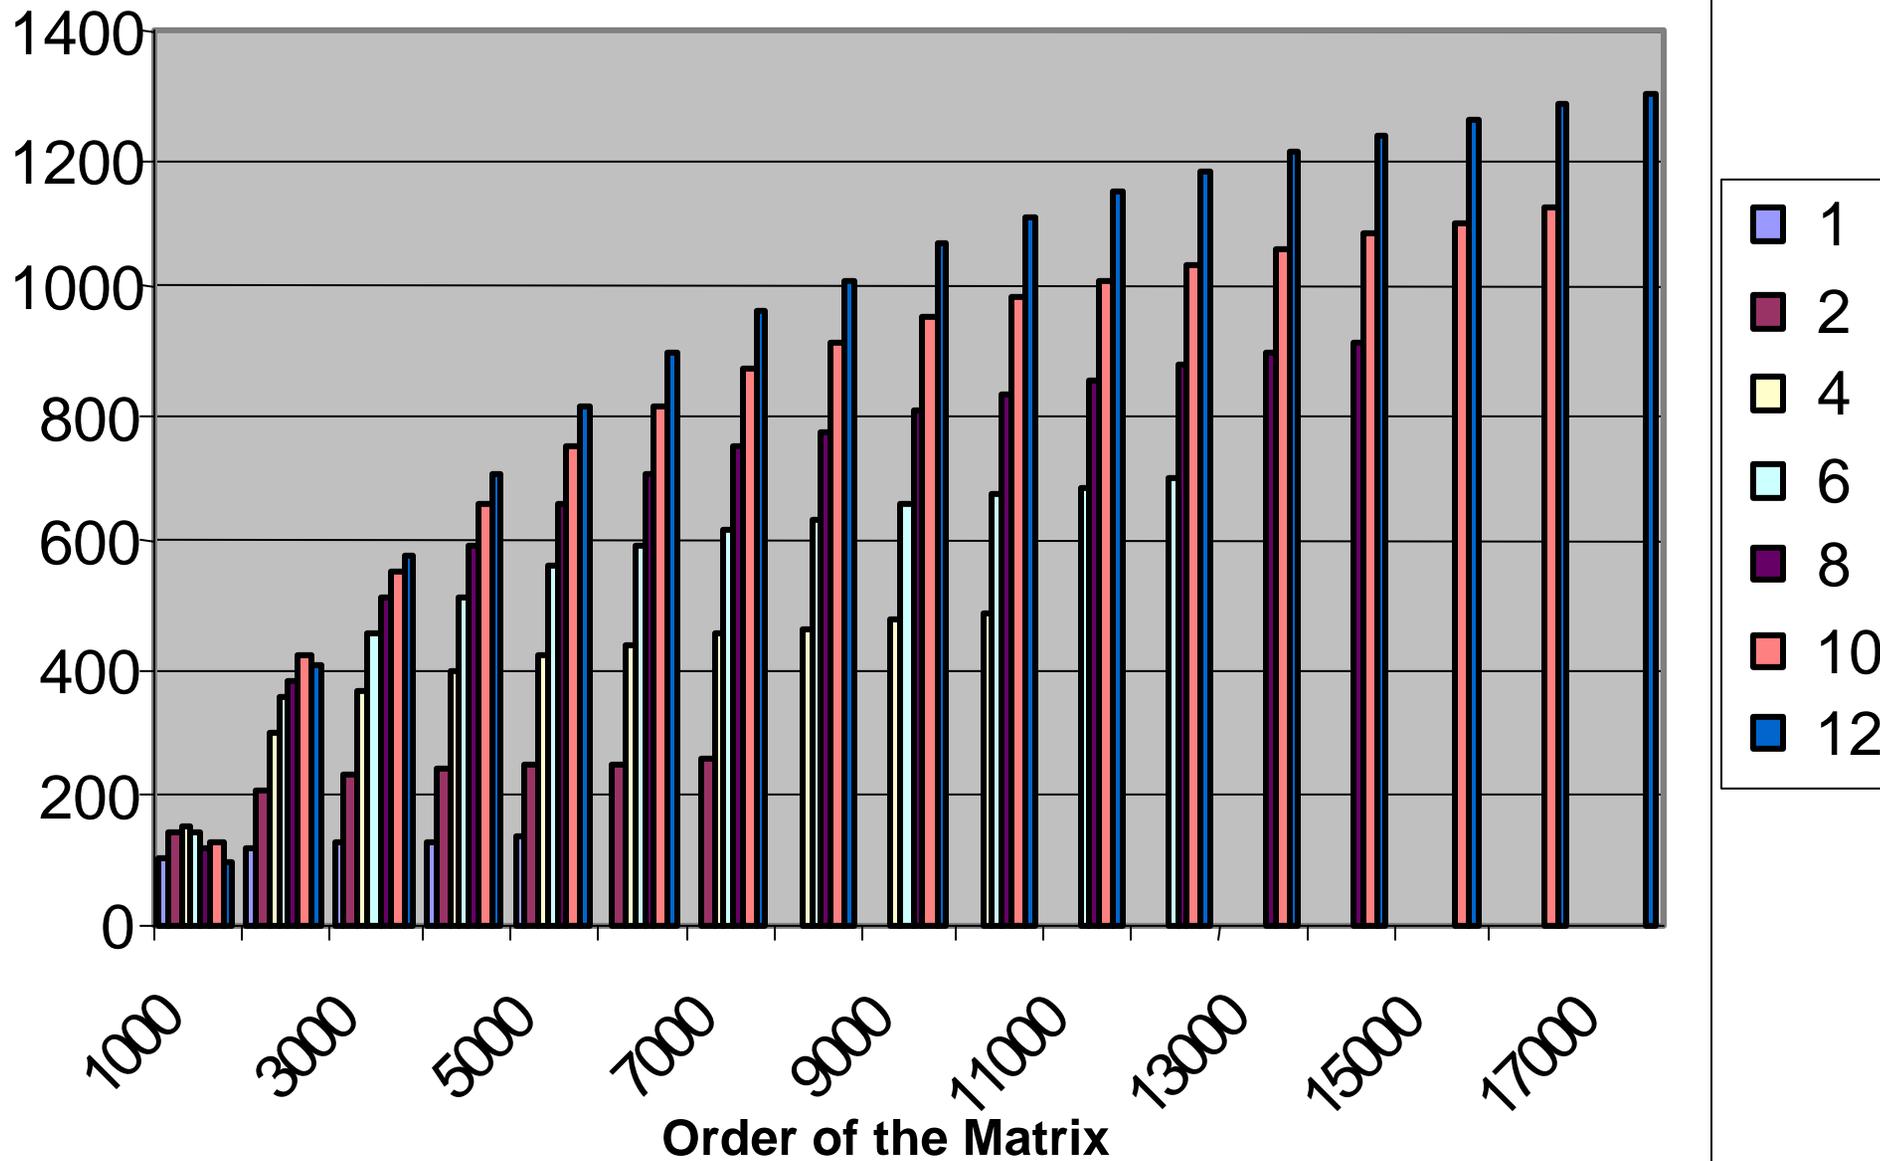
# ScaLAPACK: *debugging hints*

- Look at ScaLAPACK example programs.
- Always check the value of INFO on exit from a ScaLAPACK routine.
- Query for size of workspace, LWORK = -1.
- Link to the Debug Level 1 BLACS (specified by BLACSDBGVL=1 in Bmake.inc).
- Consult errata files on *netlib*:  
<http://www.netlib.org/scalapack/errata.scalapack>  
<http://www.netlib.org/blacs/errata.blacs>

# ScaLAPACK: Performance

- The algorithms implemented in ScaLAPACK are scalable in the sense that the parallel efficiency is an increasing function of  $N^2/P$  (problem size per node).
- Maintaining memory use per node constant allows efficiency to be maintained (in practice, a slight degradation is acceptable).
- Use efficient machine-specific BLAS (not the Fortran 77 source code available in <http://www.netlib.gov>) and BLACS (nondebug installation).
- On a distributed-memory computer:
  - Use the right number of processors
    - Rule of thumb:  $P=M \times N / 1,000,000$  for an  $M \times N$  matrix, which provides a local matrix of size approximately 1000-by-1000.
    - Do not try to solve a small problem on too many processors.
    - Do not exceed the physical memory.
  - Use an efficient data distribution.
    - Block size (i.e.,  $MB, NB$ ) = 64.
    - Square processor grid:  $P_{row} = P_{column}$ .

# Mflop/s LU/Solve on Pentium II 300 MHz Cluster



# ScaLAPACK: *Commercial Use*

ScaLAPACK has been incorporated in the following commercial packages:

- Fujitsu
- Hewlett-Packard/Convex
- Hitachi
- IBM Parallel ESSL
- NAG Numerical PVM (and MPI) Library
- Cray LIBSCI
- NEC Scientific Software Library
- Sun Scientific Software Library
- Visual Numerics (IMSL)

# ScaLAPACK: *Development Team*

- Susan Blackford, UTK
- Jaeyoung Choi, Soongsil University
- Andy Cleary, LLNL
- Ed D'Azevedo, ORNL
- Jim Demmel, UCB
- Inderjit Dhillon, UT Austin
- Jack Dongarra, UTK
- Ray Fellers, LLNL
- Sven Hammarling, NAG
- Greg Henry, Intel
- Osni Marques, LBNL/NERSC
- Caroline Papadopoulos, UCSD
- Antoine Petitet, UTK
- Ken Stanley, UCB
- Francoise Tisseur, Manchester
- David Walker, Cardiff
- Clint Whaley, UTK

# On line tutorial: <http://acts.neresc.gov/scalapack/hands-on/main.html>

## Hands-On Exercises for ScaLAPACK

ScaLAPACK Team  
August 2002

### Introduction

These exercises provide basic and more advanced programming instruction for writing parallel programs calling the BLACS, PBLAS, and ScaLAPACK. A basic knowledge of Fortran, parallel programming with message-passing, and MPI are assumed. Some of the exercises also require an understanding of two-dimensional block cyclic data distribution.

Detailed information on the BLACS, PBLAS, and ScaLAPACK may be found at the respective URLs:

- ♦ <http://www.netlib.org/blacs>
- ♦ <http://www.netlib.org/scalapack>
- ♦ [http://www.netlib.org/scalapack/pblas\\_qref.html](http://www.netlib.org/scalapack/pblas_qref.html)

Exercises 1 and 2 give an introduction to parallel programming with the Basic Linear Algebra Communication Subprograms (BLACS). Exercises 3, 4, and 5 provide a range of simplistic to more complex parallel programs calling ScaLAPACK and PBLAS. More example programs for ScaLAPACK can be found at <http://www.netlib.org/scalapack/examples>.

The instructions for the exercises assume that the underlying system is an IBM SP or a Cray T3E; using up to six processes that do message-passing. These example programs use MPI as the underlying message-passing layer. The version of MPI used in these examples is the version 3.0, and we assume the user has this version installed.

These hands-on exercises were prepared in collaboration with the Joint Institute for Computational Science at the University of Tennessee, based on contributions from A. YarKhan, C. Hastings, S. Blackford, C. Whaley, A. Petitet and O. Marques.

Exercise 1: [BLACS - Hello World Example](#)

Exercise 2: [BLACS - Pi Example](#)

Exercise 3: [ScaLAPACK - Example Program 1](#)

Exercise 4: [ScaLAPACK - Example Program 2](#)

Exercise 5: [PBLAS Example](#)

### Addition Resources:

- ♦ [Block Cyclic Data Distribution](#)
- ♦ [Useful calling sequences](#)
- ♦ [Download all exercises](#)

[ScaLAPACK](#)

[Tools](#)

[Project](#)

[Home](#)

[Search](#)

# Instructions for Hands-on

- Do a “*cp -r ~osni/work/hands-on hands-on*”.
- There are five subdirectories under *hands-on*:
  - Example 1: BLACS, “hello world” example
  - Example 2: BLACS, “pi” example
  - Example 3: ScaLAPACK small problem 1
  - Example 4: ScaLAPACK small problem 2
  - Example 5: PBLAS example
- The examples are written in Fortran. For a successful compilation and execution, you will have to correct some lines in the examples, in particular the lines starting with *\*\*\** (commented lines).
- The examples can be compiled with “*make*”, which will generate an executable file whose name begins with “*x*”.

### Exercise 3: ScaLAPACK - Example Program 1

**Information** This is a very simplistic example program that solves a linear system by calling the ScaLAPACK routine PSGESV. More complete details of this example program can be found in Chapter 2 of the [ScaLAPACK Users' Guide](#). This example program demonstrates the basic requirements to call a ScaLAPACK routine: initializing the process grid, assigning the matrix to the processes, calling the ScaLAPACK routine, and releasing the process grid.

This example program solves the 9-by-9 system of linear equations given by:

$$\begin{array}{cccccccc|cccc} / & 19 & 3 & 1 & 12 & 1 & 16 & 1 & 3 & 11 & \backslash & / & x1 & \backslash & / & 0 \\ | & -19 & 3 & 1 & 12 & 1 & 16 & 1 & 3 & 11 & | & | & x2 & | & | & 0 \\ | & -19 & -3 & 1 & 12 & 1 & 16 & 1 & 3 & 11 & | & | & x3 & | & | & 1 \\ | & -19 & -3 & -1 & 12 & 1 & 16 & 1 & 3 & 11 & | & | & x4 & | & | & 0 \\ | & -19 & -3 & -1 & -12 & 1 & 16 & 1 & 3 & 11 & | & | & x5 & | & | & 0 \\ | & -19 & -3 & -1 & -12 & -1 & 16 & 1 & 3 & 11 & | & | & x6 & | & | & 0 \\ | & -19 & -3 & -1 & -12 & -1 & -16 & 1 & 3 & 11 & | & | & x7 & | & | & 0 \\ | & -19 & 3 & -1 & -12 & -1 & -16 & -1 & 3 & 11 & | & | & x8 & | & | & 0 \\ \backslash & -19 & -3 & -1 & -12 & -1 & -16 & -1 & -3 & 11 & / & \backslash & x9 & / & \backslash & 0 \end{array}$$

using the ScaLAPACK driver routine PSGESV. The ScaLAPACK routine PSGESV solves a system of linear equations  $A \cdot X = B$ , where the coefficient matrix (denoted by  $A$ ) and the right-hand-side matrix (denoted by  $B$ ) are real, general distributed matrices. The coefficient matrix  $A$  is distributed as depicted below and, for simplicity, we shall solve the system for one right-hand side (NRHS=1), that is, the matrix  $B$  is a vector. The third element of the matrix  $B$  is equal to 1, and all other elements are equal to 0. After solving this system of equations, the solution vector  $X$  is given by

$$\begin{array}{cc|cc} / & x1 & \backslash & / & 0 \\ | & x2 & | & -1/6 & | \\ | & x3 & | & 1/2 & | \\ | & x4 & | & 0 & | \\ | & x5 & | & 0 & | \\ | & x6 & | & 0 & | \\ | & x7 & | & -1/2 & | \\ | & x8 & | & 1/6 & | \\ \backslash & x9 & / & \backslash & 0 \end{array}$$

Let us assume that the matrix  $A$  is partitioned and distributed such that we have chosen the row and column block sizes as  $MB=NB=2$ , and the matrix is distributed on a 2-by-3 process grid ( $P_r=2, P_c=3$ ). The partitioning and distribution of our example matrix  $A$  is represented in the two figures below, where, to aid visualization, we use the notation  $s=19, c=3, s=1, l=12, p=16$ , and  $k=11$ .

Partitioning of global matrix  $A$ :

s	c	a	l	a	p	a	c	k
-s	c	a	l	a	p	a	c	k
-s	-c	a	l	a	p	a	c	k
-s	-c	-a	l	a	p	a	c	k
-s	-c	-a	-l	a	p	a	c	k
-s	-c	-a	-l	-a	p	a	c	k
-s	-c	-a	-l	-a	-p	a	c	k
-s	c	-a	-l	-a	-p	-a	c	k
-s	-c	-a	-l	-a	-p	-a	-c	k

Mapping of matrix  $A$  onto process grid ( $P_r=2, P_c=3$ ). Note, for example, that process (0,0) contains a local array of size  $A(3,4)$ .

	0	1	2					
s	c	a	l	k	a	p		
-s	c	a	l	k	a	p		
-s	-c	a	l	k	a	p	0	
-s	-c	a	c	-a	-l	k	a	p
-s	-c	a	c	-a	-l	k	-a	p
-s	-c	-a	-l	k	-a	-p		
-s	-c	-a	-l	k	-a	-p		
-s	-c	-a	-l	k	-a	-p		
-s	-c	-a	-l	k	-a	-p	1	
-s	-c	-a	-l	k	-a	-p		
-s	-c	-a	-l	k	-a	-p		

### PROGRAM PSGESVDRIVER

```

*
*   Example Program solving Ax=b via ScaLAPACK routine PSGESV
*
*   .. Parameters ..
*
*   ..
*
*   .. Local Scalars ..
*
*   ..
*
*   .. Local Arrays ..
*
*   .. Executable Statements ..
*
*
*   INITIALIZE THE PROCESS GRID
*
*
***  CALL SL_INIT( ICTXT, NPROW, NPCOL )
***  CALL BLACS_***** ( ICTXT, NPROW, NPCOL, MYROW, MYCOL )
*
*   If I'm not in the process grid, go to the end of the program
*
*
*   IF( MYROW.EQ.-1 )
*     GO TO 10
*
*
*   DISTRIBUTE THE MATRIX ON THE PROCESS GRID
*
*   Initialize the array descriptors for the matrices A and B
*
*
*   CALL DESCINIT( DESCA, M, N, MB, NB, RSRC, CSRC, ICTXT, MXLLDA,
* $             INFO )
*   CALL DESCINIT( DESCB, N, NRHS, NB, NBRHS, RSRC, CSRC, ICTXT,
* $             MXLLDB, INFO )
*
*   Generate matrices A and B and distribute to the process grid
*
*
*   CALL MATINIT( A, DESCA, B, DESCB )
*
*
*   Make a copy of A and B for checking purposes
*
*
*   CALL PSLACPY( 'All', N, N, A, 1, 1, DESCA, A0, 1, 1, DESCA )
*   CALL PSLACPY( 'All', N, NRHS, B, 1, 1, DESCB, B0, 1, 1, DESCB )
*
*
*   CALL THE SCALAPACK ROUTINE
*
*   Solve the linear system A * X = B
*
***  CALL ***** ( N, NRHS, A, IA, JA, DESCA, IPIV, B, IB, JB, DESCB,
* $             INFO )

```



Consider the 12-by-10 matrix:

$$A \begin{matrix} ? \\ ? \\ ? \\ ? \\ ? \\ ? \\ ? \\ ? \\ ? \\ ? \\ ? \\ ? \end{matrix} \begin{matrix} a_{1,1} & a_{1,2} & a_{1,3} & a_{1,4} & a_{1,5} & a_{1,6} & a_{1,7} & a_{1,8} & a_{1,9} & a_{1,10} \\ a_{2,1} & a_{2,2} & a_{2,3} & a_{2,4} & a_{2,5} & a_{2,6} & a_{2,7} & a_{2,8} & a_{2,9} & a_{2,10} \\ a_{3,1} & a_{3,2} & a_{3,3} & a_{3,4} & a_{3,5} & a_{3,6} & a_{3,7} & a_{3,8} & a_{3,9} & a_{3,10} \\ a_{4,1} & a_{4,2} & a_{4,3} & a_{4,4} & a_{4,5} & a_{4,6} & a_{4,7} & a_{4,8} & a_{4,9} & a_{4,10} \\ a_{5,1} & a_{5,2} & a_{5,3} & a_{5,4} & a_{5,5} & a_{5,6} & a_{5,7} & a_{5,8} & a_{5,9} & a_{5,10} \\ a_{6,1} & a_{6,2} & a_{6,3} & a_{6,4} & a_{6,5} & a_{6,6} & a_{6,7} & a_{6,8} & a_{6,9} & a_{6,10} \\ a_{7,1} & a_{7,2} & a_{7,3} & a_{7,4} & a_{7,5} & a_{7,6} & a_{7,7} & a_{7,8} & a_{7,9} & a_{7,10} \\ a_{8,1} & a_{8,2} & a_{8,3} & a_{8,4} & a_{8,5} & a_{8,6} & a_{8,7} & a_{8,8} & a_{8,9} & a_{8,10} \\ a_{9,1} & a_{9,2} & a_{9,3} & a_{9,4} & a_{9,5} & a_{9,6} & a_{9,7} & a_{9,8} & a_{9,9} & a_{9,10} \\ a_{10,1} & a_{10,2} & a_{10,3} & a_{10,4} & a_{10,5} & a_{10,6} & a_{10,7} & a_{10,8} & a_{10,9} & a_{10,10} \\ a_{11,1} & a_{11,2} & a_{11,3} & a_{11,4} & a_{11,5} & a_{11,6} & a_{11,7} & a_{11,8} & a_{11,9} & a_{11,10} \\ a_{12,1} & a_{12,2} & a_{12,3} & a_{12,4} & a_{12,5} & a_{12,6} & a_{12,7} & a_{12,8} & a_{12,9} & a_{12,10} \end{matrix} \begin{matrix} ? \\ ? \\ ? \\ ? \\ ? \\ ? \\ ? \\ ? \\ ? \\ ? \\ ? \\ ? \end{matrix}$$

Do the following block cyclic distributions:

- 3-by-3 blocking on a 3-by-2 process grid
- 4-by-4 blocking on a 2-by-3 process grid

Use <http://acts.nersc.gov/scalapack/hands-on/datadist.html> to compare