



# HPC Numerical Libraries

**Nicola Spallanzani** - [n.spallanzani@cineca.it](mailto:n.spallanzani@ Cineca.it)  
SuperComputing Applications and Innovation Department





# Algorithms and Libraries

Many numerical algorithms are well known and largely available. See for instance:


<http://www.nr.com>

So why should I use libraries?


- They are available on many platforms
- ... and they are usually optimized by vendors
- In case vendor libraries are not installed:

<http://www.netlib.org>

# Numerical Libraries




- Scalable Linear Algebra
- Fast Fourier Transform
- Adaptive Mesh Refinement (AMR)
- Input/Output




# Scalable Linear Algebra

**Nicola Spallanzani** - [n.spallanzani@cineca.it](mailto:n.spallanzani@ Cineca.it)  
SuperComputing Applications and Innovation Department



# Standard Linear Algebra Libraries



PETSc

BLACS

ACML

LAPACK

PLASMA

PSBLAS

MKL

ATLAS

**BLAS**

MAGMA

SLEPc

... but not only

TRILINOS

ARPACK

ESSL

SCALAPACK


CUBLAS

SCOTCH

**which library should I use?**

<http://www.netlib.org/utk/people/JackDongarra/la-sw.html>

# (Parallel) Basic Linear Algebra Subprograms (BLAS and PBLAS)




- **Level 1) : Vector - Vector operations**
- **Level 2) : Vector - Matrix operations**
- **Level 3) : Matrix - Matrix operations**

Linear systems, Eigenvalue equations (LAPACK)

3)  $M \times M$  products

2)  $M \times V$  products

1)  $V \times V$  products



# BLAS subroutines

**matrix multiplication:  $C = A * B$  (level 3)**

**DGEMM( TRANSA, TRANSB, M, N, L, ALPHA, A, LDA, B, LDB, BETA, C, LDC )**  
 'N' or 'T'  $\rightarrow$   $\max(1, M)$

**matrix times vector:  $Y = A * X$  (level 2)**

**DGEMV( TRANS, M, N, ALPHA, A, LDA, X, INCX, BETA, Y, INCY )**  
 1.0d0  $\rightarrow$  0.0d0

**vector swap:  $X \Leftrightarrow Y$  (level 1)**

**DSWAP( N, X, INCX, Y, INCY )**


**scalar product:  $p = X' \cdot Y$  (level 1)**

**$p = \text{DDOT}( N, X, \text{INCX}, Y, \text{INCY} )$**   
 Function  $\leftarrow$  Increment for elements

Quick Reference Guide to the BLAS

<http://www.netlib.org/lapack/lug/node145.html>

# CBLAS subroutines



Instead of calling BLAS routines from a C-language program, you can use the CBLAS interface.

CBLAS is a C-style interface to the BLAS routines. You can call CBLAS routines using regular C-style calls. Use the ***cbblas.h*** header file with the CBLAS interface. The header file specifies enumerated values and prototypes of all the functions.

**matrix multiplication:  $C = A * B$  (level 3)**

`cbblas_dgemm(ORDER, TRANSA, TRANSB, M, N, L, ALPHA, A, LDA, B, LDB, BETA, C, LDC)`


`CblasRowMajor`      `CblasNoTrans`

**matrix times vector:  $Y = A * X$  (level 2)**

`cbblas_dgemv(ORDER, TRANS, M, N, ALPHA, A, LDA, X, INCX, BETA, Y, INCX)`




# **(Scalable) Linear Algebra PACKage** (LAPACK and ScaLAPACK)



- **Matrix Decomposition**
- **Linear Equation Systems**
- **Eigenvalue Equations**
- **Linear Least Square Equations**

**Designed for:**


- **dense, banded, triangular matrices**
- **real and complex matrices**



# LAPACK subroutines

**Eigenvalues and, optionally, eigenvectors of a real symmetric matrix:**

**DSYEV( JOBZ, UPLO, N, A, LDA, W, WORK, LWORK, INFO )**



```

void dsyev_( const char* jobz, const char* uplo, const MKL_INT* n,
             double* a, const MKL_INT* lda, double* w, double* work,
             const MKL_INT* lwork, MKL_INT* info );
  
```


**Index of Driver and Computational Routines:**

<http://www.netlib.org/lapack/lug/node142.html>

**Index of Auxiliary Routines:**

<http://www.netlib.org/lapack/lug/node144.html>

# LAPACKE subroutines



Two-level C interface to LAPACK, consisting of a high-level interface and a middle-level interface.


Use the **lapacke.h** header file with the LAPACKE interface. The header file specifies enumerated values and prototypes of all the functions.

system of linear equations:

```
lapack_int LAPACKE_dgesv (int matrix_layout, lapack_int n, lapack_int nrhs,  
                          double *a, lapack_int lda, lapack_int *ipiv,  
                          double *b, lapack_int ldb);
```

LAPACK\_ROW\_MAJOR  
LAPACK\_COL\_MAJOR

```
lapack_int LAPACKE_dgesv_work (int matrix_layout, lapack_int n, lapack_int nrhs,  
                               double *a, lapack_int lda, lapack_int *ipiv,  
                               double *b, lapack_int ldb);
```



# How To Compile

*# load these modules on Galileo*

```
module load gnu/4.9.2
```

```
module load blas/3.5.0--gnu--4.9.2
```

```
module load lapack/3.5.0--gnu--4.9.2
```

*#FORTRAN:*

```
LALIB="-L${LAPACK_LIB} -llapack -L${BLAS_LIB} -lblas"
```

```
gfortran -o program.x program.f90 ${LALIB}
```


*#C: remember to include the header files "cblas.h" and/or "lapacke.h"*

```
LALIB="-I${BLAS_INC} -I${LAPACK_INC} \
```

```
    -L${LAPACK_LIB} -llapacke -llapack \
```

```
    -L${BLAS_LIB} -lcblas -lblas -lgfortran -lm"
```

```
gcc -o program.x program.f90 ${LALIB}
```



# How To Compile

```
module load intel/pe-xe-2016--binary  
module load mkl/11.3.0--binary
```

```
LALIB="-I${MKL_INC} -L${MKL_LIB} -lmkl_intel_lp64 \  
-lmkl_sequential -lmkl_core -lpthread -lm -ldl"
```

*#C: remember to include the header file "mkl.h"*


```
icc -o program.x program.c ${LALIB}  
icc -mkl -o program.x program.c -lpthread -lm -ldl
```

*#FORTRAN:*

```
ifort -o program.x program.f90 ${LALIB}  
ifort -mkl -o program.x program.f90 -lpthread -lm -ldl
```


<http://software.intel.com/en-us/articles/intel-mkl-link-line-advisor>

# Automatically Tuned Linear Algebra Software




- ATLAS is both a research project and a software package.
- ATLAS's purpose is to provide portably optimal linear algebra software (recommended as a way to automatically generate an optimized BLAS library).
- The current version provides a complete BLAS API (for both C and Fortran77), and a very small subset of the LAPACK API.

# (Parallel) ARnoldi PACKage



- ARPACK is a collection of Fortran77 subroutines designed to solve large scale eigenvalue problems.
- It is most appropriate for large sparse or structured matrices. (Where structured means that a matrix-vector product requires order  $n$  rather than the usual order  $n^2$  floating point operations.)
- ARPACK is dependent upon a number of subroutines from LAPACK and the BLAS.
- Main feature: reverse communication interface.
- A parallel version of the ARPACK library is available. The message passing layers currently supported are BLACS and MPI .



The Portable, Extensible Toolkit for Scientific Computation (PETSc, pronounced PET-see; the S is silent), is a suite of data structures and routines developed by Argonne National Laboratory for the scalable (parallel) solution of scientific applications modeled by partial differential equations. It employs the Message Passing Interface (MPI) standard for all message-passing communication. PETSc includes a large suite of parallel linear and nonlinear equation solvers that are easily used in application codes written in C, C++, Fortran and now Python.

*Particularly well suited for Sparse Linear Algebra.*

Language: C, C++, Fortran, Python

Availability: open source


Developers: various

Distributors: Mathematics and Computer Science Division, Argonne National Laboratory

Ref.: Argonne National Laboratory



# TRILINOS




The Trilinos Project is an effort to develop algorithms and enabling technologies within an object-oriented software framework (written in C++) for the solution of large-scale, complex multi-physics engineering and scientific problems. A unique design feature of Trilinos is its focus on packages.

It is organized in Capability Areas:

- User Experience
- Parallel Programming Environments
- Framework & Tools
- Software Engineering Technologies and Integration
- I/O Support
- Meshes, Geometry, & Load Balancing
- Discretizations
- Scalable Linear Algebra
- Linear & Eigen Solvers
- Embedded Nonlinear Analysis Tools

# MAGMA



**Matrix Algebra for GPU and Multicore Architecture**

<http://icl.cs.utk.edu/magma/>


The MAGMA project aims to develop a dense linear algebra library similar to LAPACK but for heterogeneous/hybrid architectures, starting with current "Multicore+GPU" systems.

## **Methodology: CPU and GPU overlapping**

MAGMA uses HYBRIDIZATION methodology based on

- Representing linear algebra algorithms as collections of TASKS and DATA DEPENDENCIES among them
- Properly SCHEDULING tasks' execution over multicore and GPU hardware components

# MAGMA



## CPU versus GPU interfaces

Why two different interfaces?


If data is already on the GPU

- pointer to GPU memory
- (some) additional memory allocation on CPU side

If data is already on the CPU

- no changes on the prototype
- internal overlap communication/computation (it uses pinned)
- (some) additional memory allocation on GPU side

# MKL on Intel Xeon Phi (MIC)




Intel released a version for Xeon Phi of the MKL mathematical libraries

MKL have three different usage models:


- Automatic offload (AO)
- Compiler assisted offload (CAO)
- Native execution

<https://wiki.u-gov.it/confluence/display/SCAIUS/Guide+for+Intel+Xeon+Phi+%28MIC%29+Usage>




# Fast Fourier Transform

**Nicola Spallanzani** - [n.spallanzani@cineca.it](mailto:n.spallanzani@ Cineca.it)  
SuperComputing Applications and Innovation Department



# Fast Fourier Transform




A fast Fourier transform (FFT) algorithm computes the discrete Fourier transform (DFT) of a sequence, or its inverse. Fourier analysis converts a signal from its original domain (often time or space) to a representation in the frequency domain and vice versa.

An FFT rapidly computes such transformations by factorizing the DFT matrix into a product of sparse (mostly zero) factors.

As a result, it manages to reduce the complexity of computing the DFT from  $O(n^2)$ , which arises if one simply applies the definition of DFT, to  $O(n \log n)$ , where  $n$  is the data size.

# Fast Fourier Transform



## FFTPACK

FFTPACK is a package of Fortran subroutines for the fast Fourier transform. It includes complex, real, sine, cosine, and quarter-wave transforms. It is included in the general-purpose mathematical library SLATEC.

Language: FORTRAN


Availability: public domain

Developers: Paul N. Swarztrauber, National Center for Atmospheric Research, Boulder, CO

Distributors: NETLIB

Ref.: The University of Tennessee at Knoxville and Bell Laboratories

# Fast Fourier Transform



## FFTW

(Fastest Fourier Transform in the West)

FFTW is a C subroutine library for computing the discrete Fourier transform (DFT) in one or more dimensions, of arbitrary input size, and of both real and complex data (as well as of even/odd data, i.e. the discrete cosine/sine transforms or DCT/DST).

The currently most used free library to compute DFTs.

It uses both MPI and OpenMP.

Language: C

Availability: public domain


Developers: Matteo Frigo and Steven G. Johnson.

Distributors: FFTW

Ref.: MIT



# Fast Fourier Transform




## FFTW

(Fastest Fourier Transform in the West)


- Written in C (FORTRAN 77/90/2003 wrappers are included)
- FFTW adapt itself to your system and your problem: It check both total memory, size of your array, cache size register, CPU type, ...
- ... and it chose the best algorithm for you!
- 2 setps is required:
  - Plan creation
  - Execution
- It supports FFTs for any different array size ( $<$  total memory)

# Fast Fourier Transform




## 2Decomp&FFT

The 2DECOMP&FFT library is a software framework in Fortran to build large-scale parallel applications. It is designed for applications using three-dimensional structured mesh and spatially implicit numerical algorithms. At the foundation it implements a general-purpose 2D pencil decomposition for data distribution on distributed-memory platforms. On top it provides a highly scalable and efficient interface to perform three-dimensional distributed FFTs. The library is optimised for supercomputers and scales well to hundreds of thousands of cores. It relies on MPI but provides a user-friendly programming interface that hides communication details from application developers.



VS



Language: FORTRAN

Availability: public domain


Developers: Ning LI.

Distributors: 2decomp.org

Ref.: HECToR Distributed Computational Science and  
 Engineering (CSE) Service operated by NAG Ltd.

<http://www.2decomp.org/>

# Fast Fourier Transform



## P3DFFT

3DFFT is a scalable software library implementing three-dimensional spectral transforms. It has been used in a variety of codes from many areas of computational science. It has been tested and used on many high-end computational system. It uses two-dimensional domain decomposition in order to overcome a scaling bottleneck of one-dimensional decomposition. This allows the programs with this library to scale well on a large number of cores, consistent with bisection bandwidth scaling of interconnect of the underlying hardware system.

Language: C


Availability: public domain

Developers: Dmitry Pekurovsky.

Distributors: 2decomp.org


Ref.: San Diego Supercomputer Center (SDSC) at UC San Diego.

<https://code.google.com/p/p3dfft/>




# Adaptive Mesh Refinement (AMR)

**Nicola Spallanzani** - [n.spallanzani@cineca.it](mailto:n.spallanzani@ Cineca.it)  
SuperComputing Applications and Innovation Department




# AMR libraries



Many problems in numerical analysis, do not require a uniform precision in the numerical grids used for computational simulation.

Adaptive mesh refinement provides a dynamic programming environment for adapting the precision of the numerical computation based on the requirements of a computation problem in specific areas.

# AMR libraries



## METIS

METIS is a software package for graph partitioning that implements various multilevel algorithms. It was written in C, and it can generate, manipulate and partition, graphs grids and matrices. Particularly well suited for partitioning regular grids on a parallel distributed memory system.

Language: C


Availability: free for academy and research

Developers: George Karypis Lab

Distributors: George Karypis Lab

Ref.: Department of Computer Science & Engineering, University of Minnesota

# AMR libraries



## ParMETIS

ParMETIS is an MPI-based parallel library that implements a variety of algorithms for partitioning unstructured graphs, meshes, and for computing fill-reducing orderings of sparse matrices. ParMETIS extends the functionality provided by METIS and includes routines that are especially suited for parallel AMR computations and large scale numerical simulations

Language: C


Availability: free for academy and research

Developers: George Karypis Lab

Distributors: George Karypis Lab

Ref.: Department of Computer Science & Engineering, University of Minnesota

<http://glaros.dtc.umn.edu/gkhome/views/metis>



## PARAMESH

PARAMESH is a package of Fortran 90 subroutines designed to provide an application developer with an easy route to extend an existing serial code which uses a logically Cartesian structured mesh into a parallel code with adaptive mesh refinement(AMR).

Language: FORTRAN


Availability: free for academy and research

Developers: Peter MacNeice, Kevin M. Olson et al.

Distributors: Kevin M. Olson


Ref.: NASA e Drexel University

<http://sourceforge.net/projects/paramesh/>





# AMR libraries



## CHOMBO

Chombo provides a set of tools for implementing finite difference and finite volume methods for the solution of partial differential equations on block-structured adaptively refined rectangular grids.

Language: C


Availability: free for academy and research


Developers: P. Colella, D. T. Graves, J. N. Johnson et al.

Distributors: D. T. Graves

Ref.: Applied Numerical Algorithms Group - Computational Research Division - Lawrence Berkeley National Laboratory- Berkeley, CA, USA


<https://commons.lbl.gov/display/chombo/>






# I/O Libraries

**Nicola Spallanzani** - [n.spallanzani@cineca.it](mailto:n.spallanzani@ Cineca.it)  
SuperComputing Applications and Innovation Department



# I/O libraries



## HDF5

HDF5 is a data model, library, and file format for storing and managing data. It supports an unlimited variety of datatypes, and is designed for flexible and efficient I/O and for high volume and complex data. HDF5 is portable and is extensible, allowing applications to evolve in their use of HDF5. The HDF5 Technology suite includes tools and applications for managing, manipulating, viewing, and analyzing data in the HDF5 format.

Language: Fortran/C/C++/Java


Availability: open source.

Developers: NCSA at the University of Illinois

Distributors: HDF Group

Ref.: <http://www.hdfgroup.org>

# I/O libraries



## NetCDF

NetCDF is a set of software libraries and self-describing, machine-independent data formats that support the creation, access, and sharing of array-oriented scientific data.

The NetCDF source-code is hosted at GitHub, and may be found directly at <http://github.com/Unidata/netcdf-c>.

NetCDF-4 (and pNetCDF) work with MPI version 2 (MPI-2).

Language: Fortran/C/C++/Java

Availability: open source.

Developers: Unidata

Distributors: Unidata

Ref.: <http://www.unidata.ucar.edu/netcdf>