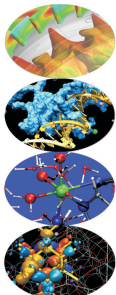


Programmazione Avanzata

Vittorio Ruggiero

(v.ruggiero@cineca.it)

Roma, Marzo 2017



Makefile

Scientific Libraries

- ▶ What do I need to develop my HPC application? At least:
 - ▶ A compiler (e.g. GNU, Intel, PGI, PathScale, ...)
 - ▶ A code editor
- ▶ Several tools may help you (even for non HPC applications)
 - ▶ Debugger: e.g. gdb, TotalView, DDD
 - ▶ Profiler: e.g. gprof, Scalasca, Tau, Vampir
 - ▶ Project management: e.g. make, projects
 - ▶ Revision control: e.g. svn, git, cvs, mercurial
 - ▶ Generating documentation: e.g. doxygen
 - ▶ Source code repository: e.g. sourceforge, github, google.code
 - ▶ Data repository, currently under significant increase
 - ▶ ...

- ▶ You can select the code editor among a very wide range
 - ▶ from the light and fast text editors (e.g. VIM, emacs, ...)
 - ▶ to the more sophisticated Integrated development environment (IDE), (e.g. Eclipse)
 - ▶ or you have intermediate options (e.g. Geany)
- ▶ The choice obviously depends on the complexity and on the software tasks
- ▶ ... but also on your personal taste

- ▶ Non trivial programs are hosted in several source files and link libraries
- ▶ Different types of files require different compilation
 - ▶ different optimization flags
 - ▶ different languages may be mixed, too
 - ▶ compilation and linking require different flags
 - ▶ and the code could work on different platforms
- ▶ During development (and debugging) several recompilations are needed, and we do not want to recompile all the source files but only the modified ones
- ▶ How to deal with it?
 - ▶ use the IDE (with plug-ins) and their project files to manage the content (e.g. Eclipse)
 - ▶ use language-specific compiler features
 - ▶ use external utilities, e.g. Make!

- ▶ “Make is a tool which controls the generation of executables and other non-source files of a program from the program’s source files”
- ▶ Make gets its knowledge from a file called the makefile, which lists each of the non-source files and how to compute it from other files
- ▶ When you write a program, you should write a makefile for it, so that it is possible to use Make to build and install the program and more . . .
- ▶ GNU Make has some powerful features for use in makefiles, beyond what other Make versions have

- ▶ To prepare to use make, you have to write a file that describes:
 - ▶ the relationships among files in your program
 - ▶ commands for updating each file
- ▶ Typically, the executable file is updated from object files, which are created by compiling source files
- ▶ Once a suitable makefile exists, each time you change some source files, the shell command

```
make -f <makefile_name>
```

performs all necessary recompilations

- ▶ If `-f` option is missing, the default names `makefile` or `Makefile` are used

- ▶ A simple makefile consists of “rules”:

```
target ... : prerequisites ...  
    recipe  
    ...  
    ...
```

- ▶ a **target** is usually the name of a file that is generated by a program; examples of targets are executable or object files. A target can also be the name of an action to carry out, such as `clean`
 - ▶ a **prerequisite** is a file that is used as input to create the target. A target often depends on several files.
 - ▶ a **recipe** is an action that make carries out. A recipe may have more than one command, either on the same line or each on its own line. Recipe commands must be preceded by a **tab** character.
- ▶ By default, make starts with the first target (default goal)

- ▶ A simple rule:

```
foo.o : foo.c defs.h
      gcc -c -g foo.c
```

- ▶ This rule says two things
 - ▶ how to decide whether `foo.o` is out of date: it is out of date if it does not exist, or if either `foo.c` or `defs.h` is more recent than it
 - ▶ how to update the file `foo.o`: by running `gcc` as stated. The recipe does not explicitly mention `defs.h`, but we presume that `foo.c` includes it, and that that is why `defs.h` was added to the prerequisites.
- ▶ **WARNING:** Remember the tab character before starting the recipe lines!

- ▶ The main program is in `laplace2d.c` file
 - ▶ includes two header files: `timing.h` and `size.h`
 - ▶ calls functions in two source files: `update_A.c` and `copy_A.c`
- ▶ `update_A.c` and `copy_A.c` includes two header files: `laplace2d.h` and `size.h`
- ▶ A possible (naive) Makefile

```
laplace2d_exe: laplace2d.o update_A.o copy_A.o
    gcc -o laplace2d_exe laplace2d.o update_A.o copy_A.o

laplace2d.o: laplace2d.c timing.h size.h
    gcc -c laplace2d.c

update_A.o: update_A.c laplace2d.h size.h
    gcc -c update_A.c

copy_A.o: copy_A.c laplace2d.h size.h
    gcc -c copy_A.c

.PHONY: clean
clean:
    rm -f laplace2d_exe *.o
```

- ▶ The default goal is (re-)linking `laplace2d_exe`
- ▶ Before `make` can fully process this rule, it must process the rules for the files that `edit` depends on, which in this case are the object files
- ▶ The object files, according to their own rules, are recompiled if the source files, or any of the header files named as prerequisites, is more recent than the object file, or if the object file does not exist
- ▶ Note: in this makefile `.c` and `.h` are not the targets of any rules, but this could happen if they are automatically generated
- ▶ After recompiling whichever object files need it, `make` decides whether to relink `edit` according to the same “updating” rules.
- ▶ Try to follow the path: what happens if, e.g., `laplace2d.h` is modified?

- ▶ The main program is in `laplace2d.f90` file
 - ▶ uses two modules named `prec` and `timing`
 - ▶ calls subroutines in two source files: `update_A.f90` and `copy_A.f90`
- ▶ `update_A.f90` and `copy_A.f90` use only `prec` module
- ▶ sources of `prec` and `timing` modules are in the `prec.f90` and `timing.f90` files
- ▶ Beware of the Fortran modules:
 - ▶ program units using modules require the mod files to exist
 - ▶ a target may be a list of files: e.g., both `timing.o` and `timing.mod` depend on `timing.f90` and are produced compiling `timing.f90`
- ▶ Remember: the order of rules is not significant, except for determining the default goal

```
laplace2d_exe: laplace2d.o update_A.o copy_A.o prec.o timing.o
    gfortran -o laplace2d_exe prec.o timing.o laplace2d.o update_A.o copy_A.o

prec.o prec.mod: prec.f90
    gfortran -c prec.f90

timing.o timing.mod: timing.f90
    gfortran -c timing.f90

laplace2d.o: laplace2d.f90 prec.mod timing.mod
    gfortran -c laplace2d.f90

update_A.o: update_A.f90 prec.mod
    gfortran -c update_A.f90

copy_A.o: copy_A.f90 prec.mod
    gfortran -c copy_A.f90

.PHONY: clean
clean:
    rm -f laplace2d_exe *.o *.mod
```

- ▶ A phony target is one that is not really the name of a file; rather it is just a name for a recipe to be executed when you make an explicit request.
 - ▶ avoid target name clash
 - ▶ improve performance
- ▶ **clean**: an ubiquitous target

```
.PHONY: clean
clean:
    rm *.o temp
```

- ▶ Another common solution: since **FORCE** has no prerequisite, recipe and no corresponding file, make imagines this target to have been updated whenever its rule is run

```
clean: FORCE
    rm *.o temp
FORCE:
```

- ▶ The previous makefiles have several duplications
 - ▶ error-prone and not expressive
- ▶ Use variables!
 - ▶ define
 - objects = laplace2d.o update_A.o copy_A.o**
 - ▶ and use as **\$(objects)**

```
objects := laplace2d.o update_A.o copy_A.o

laplace2d_exe: $(objects)
    gcc -o laplace2d_exe $(objects)

laplace2d.o: laplace2d.c timing.h size.h
    gcc -c laplace2d.c

update_A.o: update_A.c laplace2d.h size.h
    gcc -c update_A.c

copy_A.o: copy_A.c laplace2d.h size.h
    gcc -c copy_A.c

.PHONY: clean
clean:
    rm -f laplace2d_exe *.o
```

- ▶ Use more variables to enhance readability and generality
- ▶ Modifying the first four lines it is easy to modify compilers and flags

```
CC      := gcc
CFLAGS  := -O2
CPPFLAGS :=
LDFLAGS :=

objects := laplace2d.o update_A.o copy_A.o

laplace2d_exe: $(objects)
    $(CC) $(CFLAGS) $(CPPFLAGS) -o laplace2d_exe $(objects) $(LDFLAGS)

laplace2d.o: laplace2d.c timing.h size.h
    $(CC) $(CFLAGS) $(CPPFLAGS) -c laplace2d.c

update_A.o: update_A.c laplace2d.h size.h
    $(CC) $(CFLAGS) $(CPPFLAGS) -c update_A.c

copy_A.o: copy_A.c laplace2d.h size.h
    $(CC) $(CFLAGS) $(CPPFLAGS) -c copy_A.c

.PHONY: clean
clean:
    rm -f laplace2d_exe *.o
```


- ▶ There are still duplications: each compilation needs the same command except for the file name
 - ▶ imagine what happens with hundred/thousand of files!
- ▶ What happens if Make does not find a rule to produce one or more prerequisite (e.g., and object file)?
- ▶ Make searches for an implicit rule, defining default recipes depending on the processed type
 - ▶ C programs: n.o is made automatically from n.c with a recipe of the form

```
$ (CC) $ (CPPFLAGS) $ (CFLAGS) -c
```

- ▶ C++ programs: n.o is made automatically from n.cc, n.cpp or n.C with a recipe of the form

```
$ (CXX) $ (CPPFLAGS) $ (CXXFLAGS) -c
```

- ▶ Fortran programs: n.o is made automatically from n.f, n.F (\$ (CPPFLAGS) only for .F)

```
$ (FC) $ (FFLAGS) $ (CPPFLAGS) -c
```

- ▶ **Implicit rules allow for saving many recipe lines**
 - ▶ but what happens is not clear reading the Makefile
 - ▶ and you are forced to use a predefined structure and variables
 - ▶ to clarify the types to be processed, you may define **.SUFFIXES** variable at the beginning of Makefile

```
.SUFFIXES:
.SUFFIXES: .o .f
```

- ▶ You may use re-define an implicit rule by writing a pattern rule
 - ▶ a pattern rule looks like an ordinary rule, except that its target contains one character %
 - ▶ usually written as first target, does not become the default target

```
%o : %.c
      $(CC) -c $(OPT_FLAGS) $(DEB_FLAGS) $(CPP_FLAGS) $< -o $@
```

- ▶ Automatic variables are usually exploited
 - ▶ \$@ is the target
 - ▶ \$< is the first prerequisite (usually the source code)
 - ▶ \$^ is the list of prerequisites (useful in linking stage)

- ▶ It is possible to select a specific target to be updated, instead of the default goal (remember `clean`)

```
make copy_A.o
```

- ▶ of course, it will update the chain of its prerequisite
- ▶ useful during development when the full target has not been programmed, yet
- ▶ And it is possible to set target-specific variables as (repeated) target prerequisites
- ▶ Consider you want to write a Makefile considering both GNU and Intel compilers
- ▶ Use a default goal which is a help to inform that compiler must be specified as target

```
CPPFLAGS :=
LDLFLAGS :=

objects := laplace2d.o update_A.o copy_A.o

.SUFFIXES :=
.SUFFIXES := .c .o

%.o: %.c
    $(CC) $(CFLAGS) $(CPPFLAGS) -c $<

help:
    @echo "Please select gnu or intel compilers as targets"

gnu: CC      := gcc
gnu: CFLAGS  := -O3
gnu: $(objects)
    $(CC) $(CFLAGS) $(CPPFLAGS) -o laplace2d_gnu $(objects) $(LDLFLAGS)

intel: CC     := icc
intel: CFLAGS := -fast
intel: $(objects)
    $(CC) $(CFLAGS) $(CPPFLAGS) -o laplace2d_intel $(objects) $(LDLFLAGS)

laplace2d.o: laplace2d.c timing.h size.h
update_A.o : update_A.c laplace2d.h size.h
copy_A.o   : copy_A.c laplace2d.h size.h

.PHONY: clean
clean:
    rm -f laplace2d_gnu laplace2d_intel $(objects)
```

```

LDFLAGS :=
objects := prec.o timing.o laplace2d.o update_A.o copy_A.o
.SUFFIXES:
.SUFFIXES: .f90 .o .mod

%.o: %.f90
    $(FC) $(FFLAGS) -c $<
%.o %.mod: %.f90
    $(FC) $(FFLAGS) -c $<

help:
    @echo "Please select gnu or intel compilers as targets"
gnu: FC          := gfortran
gnu: FCFLAGS    := -O3
gnu: $(objects)
    $(FC) $(FCFLAGS) -o laplace2d_gnu $^ $(LDFLAGS)
intel: FC       := ifort
intel: FCFLAGS  := -fast
intel: $(objects)
    $(FC) $(CFLAGS) -o laplace2d_intel $^ $(LDFLAGS)

prec.o prec.mod:    prec.f90
timing.o timing.mod: timing.f90
laplace2d.o:       laplace2d.f90 prec.mod timing.mod
update_A.o:        update_A.f90 prec.mod
copy_A.o:          copy_A.f90 prec.mod

.PHONY: clean
clean:
    rm -f laplace2d_gnu laplace2d_intel $(objects) *.mod
  
```

- ▶ Another way to support different compilers or platforms is to include a platform specific file (e.g., `make.inc`) containing the needed definition of variables

```
include make.inc
```

- ▶ Common applications feature many `make.inc.<platform_name>` which you have to select and copy to `make.inc` before compiling
- ▶ When invoking `make`, it is also possible to set a variable

```
make OPTFLAGS=-O3
```

- ▶ this value will override the value inside the Makefile
- ▶ unless `override` directive is used
- ▶ but `override` is useful when you want to add options to the user defined options, e.g.

```
override CFLAGS += -g
```

- ▶ The variables considered until now are called *simply expanded* variables, are assigned using `:=` and work like variables in most programming languages.
- ▶ The other flavour of variables is called *recursively expanded*, and is assigned by the simple `=`
 - ▶ recursive expansion allows to make the next assignments working as expected

```
CFLAGS          = $(include_dirs) -O  
include_dirs    = -Ifoo -Ibar
```

- ▶ but may lead to unpredictable substitutions or even impossible circular dependencies

```
CFLAGS          = $(CFLAGS) -g
```

- ▶ You may use `+=` to add more text to the value of a variable
 - ▶ acts just like normal `=` if the variable is still undefined
 - ▶ otherwise, exactly what `+=` does depends on what flavor of variable you defined originally
- ▶ Use recursive variables only if needed

- ▶ A single file name can specify many files using wildcard characters: `*`, `?` and `[...]`
- ▶ Wildcard expansion depends on the context
 - ▶ performed by make automatically in targets and in prerequisites
 - ▶ in recipes, the shell is responsible for
 - ▶ what happens typing `make print` in the example below?
(The automatic variable `$?` stands for files that have changed)

```
print: *.c
    lpr -p $?
    touch print
```


- ▶ Environment variables are automatically transformed into make variables
- ▶ Variables could be not enough to generalize rules
 - ▶ e.g., you may need non-trivial variable dependencies
- ▶ Imagine your application has to be compiled using GNU on your local machine `mac_loc`, and Intel on the cluster `mac_clus`
- ▶ You can catch the hostname from shell and use a conditional statement (`$SHELL` is not exported)

```
SHELL := /bin/sh
HOST := $(shell hostname)
ifeq ($(HOST),mac_loc)
    CC      := gcc
    CFLAGS := -O3
endif
ifeq ($(HOST),mac_clus)
    CC      := icc
    CFLAGS := -fast
endif
```

- ▶ Be careful on Windows systems!

- ▶ For large systems, it is often desirable to put sources and headers in separate directories from the binaries
- ▶ Using Make, you do not need to change the individual prerequisites, just the search paths
- ▶ A **vpath** pattern is a string containing a % character.
 - ▶ **%.h** matches files that end in **.h**

```
vpath %.c foo  
vpath %.c blish  
vpath %.c bar
```

will look for a file ending in `.c` in `foo`, then `blish`, then `bar`

- ▶ using **vpath** without specifying directory clears all search paths associated with patterns



- ▶ When using directory searching, recipe generalizing is mandatory

```
vpath %.c src
vpath %.h ../headers
foo.o : foo.c defs.h hack.h
      gcc -c $< -o $@
```

- ▶ Again, automatic variables solve the problem
- ▶ And implicit or pattern rules may be used, too
- ▶ Directory search also works for linking libraries using prerequisites of the form `-lname`
- ▶ `make` will search for the file `libname.so` and, if not found, for `libname.a` first searching in `vpath` and then in system directory

```
foo : foo.c -lcurses
      gcc $^ -o $@
```

- ▶ **all** → Compile the entire program. This should be the default target
- ▶ **install** → Compile the program and copy the executables, libraries, and so on to the file names where they should reside for actual use.
- ▶ **uninstall** → Delete all the installed files
- ▶ **clean** → Delete all files in the current directory that are normally created by building the program.
- ▶ **distclean** → Delete all files in the current directory (or created by this makefile) that are created by configuring or building the program.
- ▶ **check** → Perform self-tests (if any).
- ▶ **install-html/install-dvi/install-pdf/install-ps** → Install documentation
- ▶ **html/dvi/pdf/ps** → Create documentation

- ▶ Compiling a large application may require several hours
- ▶ Running `make` in parallel can be very helpful, e.g. to use 8 processes

```
make -j8
```

- ▶ but not completely safe (e.g., recursive `make` compilation)
- ▶ There is much more you could know about **`make`**
 - ▶ this should be enough for your in-house application
 - ▶ but probably not enough for understanding large projects you could encounter

```
http://www.gnu.org/software/make/manual/make.html
```

Makefile

Scientific Libraries

- ▶ you have to link with
`-L<library_directory> -l<library_name>`
- ▶ Static library:
 - ▶ `*.a`
 - ▶ all symbols are included in the executable at linking
 - ▶ if you built a new library that use an other external library it doesn't contains the other symbols: you have to explicitly linking the library
- ▶ Dynamic Library:
 - ▶ `*.so`
 - ▶ Symbols are resolved at run-time
 - ▶ you have to set-up where find the requested library at run-time (i.e. setting `LD_LIBRARY_PATH` environment variable)
 - ▶ `ldd <exe_name>` gives you info about dynamic library needed

- ▶ A (complete?) set of function implementing different numeric algorithms
- ▶ A set of basic function (e.g. Fast Fourier Transform, . . .)
- ▶ A set of low level function (e.g. scalar products or random number generator), or more complex algorithms (Fourier Transform or Matrix diagonalization)
- ▶ (Usually) Faster than hand made code (i.e. sometimes it is written in assembler)
- ▶ Proprietary or Open Source
- ▶ Sometimes developed for a particular compiler/architecture . . .

▶ Pros:

- ▶ Helps to modularize the code
- ▶ Portability
- ▶ Efficient
- ▶ Ready to use

▶ Cons:

- ▶ Some details are hidden (e.g. Memory requirements)
- ▶ You don't have the complete control
- ▶ You have to read carefully the documentation
- ▶ ...

- ▶ It is hard to have a complete overview of Scientific libraries
 - ▶ many different libraries
 - ▶ still evolving . . .
 - ▶ . . . especially for “new architectures” (e.g GPU, MIC)

- ▶ Main libraries used in HPC
 - ▶ Linear Algebra
 - ▶ FFT
 - ▶ I/O libraries
 - ▶ Message Passing
 - ▶ Mesh decomposition
 - ▶ . . .

- ▶ Different parallelization paradigm
 - ▶ Shared memory (i.e. multi-threaded) or/and Distributed Memory
- ▶ Shared memory
 - ▶ BLAS
 - ▶ GOTOBLAS
 - ▶ LAPACK/CLAPACK/LAPACK++
 - ▶ ATLAS
 - ▶ PLASMA
 - ▶ SuiteSparse
 - ▶ ...
- ▶ Distributed Memory
 - ▶ Blacs (only decomposition)
 - ▶ ScaLAPACK
 - ▶ PSBLAS
 - ▶ Elemental
 - ▶ ...

- ▶ **BLAS: Basic Linear Algebra Subprograms**
 - ▶ it is one of the first library developed for HPC (1979, vector machine)
 - ▶ it includes basic operations between vectors, matrix and vector, matrix and matrix
 - ▶ it is used by many other high level libraries
- ▶ It is divided into 3 different levels
 - ▶ BLAS lev. 1: basic subroutines for scalar-vector operations (1977-79, vector machine)
 - ▶ BLAS lev. 2: basic subroutines for vector-matrix operations (1984-86)
 - ▶ BLAS lev. 3: subroutine for matrix-matrix operations (1988)

- ▶ It apply to real/complex data, in single/double precision
- ▶ Old Fortran77 style
- ▶ Level 1: scalar-vector operations ($O(n)$)
 - ▶ *SWAP vector swap
 - ▶ *COPY vector copy
 - ▶ *SCAL scaling
 - ▶ *NRM2 L2-norm
 - ▶ *AXPY sum: $a*X+Y$ (X, Y are vectors)
- ▶ Level 2: vector-matrix operations ($O(n^2)$)
 - ▶ *GEMV product vector/matrix (generic)
 - ▶ *HEMV product vector/matrix (hermitian)
 - ▶ *SYMV product vector/matrix (simmetric)

- ▶ Level 3: matrix-matrix operations ($O(n^3)$)
 - ▶ *GEMM product matrix/matrix (generic)
 - ▶ *HEMM product matrix/matrix (hermitian)
 - ▶ *SYMM product matrix/matrix (simmetric)

- ▶ GOTOBLAS
 - ▶ optimized (using assembler) BLAS library for different supercomputers. Developed by Kazushige Goto, now at Texas Advanced Computing Center, University of Texas at Austin.

- ▶ **LAPACK: Linear Algebra PACKage**
 - ▶ Linear algebraic solvers (linear systems of equations, Ordinary Least Square, eigenvalues, . . .)
 - ▶ evolution of LINPACK e EISPACK
- ▶ **ATLAS: Automatically Tuned Linear Algebra Software**
 - ▶ BLAS and LAPACK (but only some subroutine) implementations
 - ▶ Automatic optimization of Software paradigm
- ▶ **PLASMA: Parallel Linear Algebra Software for Multi-core Architectures**
 - ▶ Similare to LAPACK (less subroutines) developed to be efficient on multicore systems.
- ▶ **SuiteSparse**
 - ▶ Sparse Matrix

- ▶ Eigenvalues/Eigenvectors
 - ▶ EISPACK: with specialized version for matrix fo different kind (real/complex, hermitia, simmetrich, tridiagonal, ...)
 - ▶ ARPACK: eigenvalus for big size problems. Parallel version use BLACs and MPI libraries.
- ▶ Distributed Memory Linear Algebra
 - ▶ BLACS: linear algebra oriented message passing interface
 - ▶ ScaLAPACK: Scalable Linear Algebra PACKage
 - ▶ Elemental: framework for dense linear algebra
 - ▶ PSBLAS: Parallel Sparse Basic Linear Algebra Subroutines
 - ▶ ...

- ▶ I/O Libraries are extremely important for
 - ▶ Interoperability: C/Fortran, Little Endian/Big Endian, ...
 - ▶ Visualizzazione
 - ▶ Sub-set data analysis
 - ▶ Metadata
 - ▶ Parallel I/O

- ▶ HDF5: “is a data model, library, and file format for storing and managing data”

- ▶ 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”

- ▶ VTK: “open-source, freely available software system for 3D computer graphics, image processing and visualization”

- ▶ **MPI: Message Passing Interface**
 - ▶ De facto standard for Distributed Memory Parallelization (MPICH/OpenMPI)
- ▶ **Mesh decomposition**
 - ▶ METIS e ParMETIS: “can partition a graph, partition a finite element mesh, or reorder a sparse matrix”
 - ▶ Scotch e PT-Scotch: “sequential and parallel graph partitioning, static mapping and clustering, sequential mesh and hypergraph partitioning, and sequential and parallel sparse matrix block ordering”

▶ Trilinos

- ▶ object oriented software framework for the solution of large-scale, complex multi-physics engineering and scientific problems
- ▶ A two-level software structure designed around collections of packages
- ▶ A package is an integral unit developed by a team of experts in a particular algorithms area

▶ PETSc

- ▶ It is a suite of data structures and routines for the (parallel) solution of applications modeled by partial differential equations.
- ▶ It supports MPI, shared memory pthreads, and GPUs through CUDA or OpenCL, as well as hybrid MPI-shared memory pthreads or MPI-GPU parallelism.

- ▶ **MKL: Intel Math Kernel Library**
 - ▶ Major functional categories include Linear Algebra, Fast Fourier Transforms (FFT), Vector Math and Statistics. Cluster-based versions of LAPACK and FFT are also included to support MPI-based distributed memory computing.
- ▶ **ACML: AMD Core Math Library**
 - ▶ Optimized functions for AMD processors. It includes BLAS, LAPACK, FFT, Random Generators ...
- ▶ **GSL: GNU Scientific Library**
 - ▶ The library provides a wide range of mathematical routines such as random number generators, special functions and least-squares fitting. There are over 1000 functions in total with an extensive test suite.
- ▶ **ESSL (IBM): Engineering and Scientific Subroutine library**
 - ▶ BLAS, LAPACK, ScaLAPACK, Sparse Solvers, FFT e may other. The Parallel version uses MPI

- ▶ first of all the syntax should be correct (read the manual!!!)
- ▶ always check for the right version
- ▶ sometimes for proprietary libraries linking could be “complicated”
- ▶ e.g. Intel ScaLAPACK

```
mpif77 <program> -L$MKLRROOT/lib/intel64 \  
-lmkl_scalapack_lp64 -lmkl_blacs_openmpi \  
-lmkl_intel_lp64 -lmkl_intel_thread -lmkl_core \  
-liomp5 -lpthread
```

- ▶ Many libraries are written using C, many others using Fortran
- ▶ This can produce some problems
 - ▶ type matching: C **int** is not granted to be the same with Fortran **integer**
 - ▶ symboli Match: Fortran e C++ “alter” symbol’s name producing object file (e.g. Fortran put an extra `_`)
- ▶ Brute force approach
 - ▶ hand-made match all types and add `_` to match all librerie’s objects.
 - ▶ **nm <object_file>** lists all symbols
- ▶ Standard Fortran 2003 (module **iso_c_binding**)
 - ▶ The most important library gives you Fortran2003 interface
- ▶ In C++ command **extern "C"**

- ▶ To call libraries from C to Fortran and viceversa
- ▶ Example: `mpi` written using C/C++:
 - ▶ old style: `include "mpif.h"`
 - ▶ new style: `use mpi`
 - ▶ the two approach are not fully equivalent: using the module implies also a compile-time check type!
- ▶ Example: `BLAS` written using Fortran
 - ▶ legacy: call `dgemm_` instead of `dgemm`
 - ▶ modern: call `cblas_dgemm`
- ▶ Standardization still lacking...
 - ▶ Read the manual ...

- ▶ Take a look at “netlib” web site

```
http://www.netlib.org/blas/
```

- ▶ BLAS was written in Fortran 77, some compilers gives you interfaces (types check, F95 features)
 - ▶ Using Intel e MKL

```
use mk195_blas
```


- ▶ C (legacy):
 - ▶ add underscore to function's name
 - ▶ Fortran: arguments by reference, it is mandatory to pass pointers
 - ▶ Type matching (compiler dependent): probably `double`, `int`, `char` → `double precision`, `integer`, `character`
- ▶ C (modern)
 - ▶ use interface `cbblas`: GSL (GNU) or MKL (Intel)
 - ▶ include header file `#include <gsl.h>` OR `#include<mk1.h>`

http://www.gnu.org/software/gsl/manual/html_node/GSL-CBLAS-Library.htm