



Applications available on FERMI

[e.rossi@cineca.it](mailto:e.rossi@ Cineca.it)



www.cineca.it

What is Fermi best suitable for?

- FERMI has an enormous number of cores
- Fermi cores are not so powerful
- Memory per core quite low (1MB)
- Cores are arranged 16 per node, sharing a 16GB memory
- Cores are dedicated, nodes are dedicated
- Minimum num of nodes/job is 64/128 (1024/2048 cores)
- It is not possible to have a general number of nodes, only given partitions are fine:
 - 64 (only on 2 out of 10 racks)
 - 128, 256,
 - 512 (one midplane)
 - 1024 (two MP = 16Kcores)
 - 1536 (three MP= 24Kcores)
 - 2048 (two racks = 32Kcores)
 - (not in CINECA)
- Also topology is important ...

What is Fermi best suitable for?

- *FERMI is not really a “general purpose” machine*
- **IF your application is well suited** for the architecture you will get unprecedented results
- **IF your application is not**, very poor results, even difficult to install

- **A good application for BGQ must be**
 - **Highly scalable (several K at least)**
 - **Not memory demanding (1GB at most per core)**
 - **Hybrid parallelism can help managing the memory**

Largely scalable applications usually imply LARGE physical problems, as those dealing with at least one of the following aspects:

- large system size
- high resolution (in space)
- long duration of simulations (or high resolution in time)
- large number or set of interconnected systems or degrees of freedom

- For programmers
 - Mathematical libraries
 - I/O libraries
 - Tools for profiling and debugging
- For production
 - Scientific Applications



- **Mathematical libraries**

- `essl/5.1`
- `fftw/2.1.5, 3.3.3`
- `gsl/1.15`
- `lapack/3.4.1`
- `petsc/3.3-p2`
- `scalapack/2.0.2`
- `blas/2007`
- `mass/7.3`
- `papi/4.4.0`

- I/O libraries

- `netcdf/4.1.3`
- `hdf5/1.8.9`
 - ✓ `gzip/2.1`
 - ✓ `zlib/1.2.7`

- Scientific Applications

- `abinit/6.12.3`
- `cpmd/v3.15.3`
- `namd/2.9`
- `amber/12`
- `crystal09/1.01`
- `QuantumEspresso/5.0bgq`
- `bigdft/1.6.0`
- `dl_poly/4.03`
- `siesta/3.1`
- `cp2k/2.3`
- `gromacs/4.5.5`
- `vasp/5.2.12, 5.3.2`
- `cpmd/3.15.3_hfx`
- `lammps/20120816`
- `OpenFoam (to be done)`

- All libraries as available as module
`module avail`
- Help on how to use library in the help of the module
`module help <name>`
- Loading the module sets environment variables you have to use in the linking procedure
`-L$(*_LIB)` and `-I$(*_INCLUDE)`
- The show command lists all defined variable
`module show <name>`
- Link sequence important, `<file>.o` always before the libraries, sometimes double linking necessary

```
[erossi00@fen07 ~]$ module avail
----- /cineca/prod/modulefiles/profiles -----
profile/advanced          profile/base(default) profile/front-end
----- /cineca/prod/modulefiles/base/environment -----
autoload/0.1
----- /cineca/prod/modulefiles/base/libraries -----
blas/2007--bgq-xl--1.0(default)      mass/7.3--bgq-xl--1.0
essl/5.1                             mpi4py/1.3--bgq-gnu--4.4.6
fftw/2.1.5--bgq-xl--1.0              netcdf/4.1.3--bgq-xl--1.0
fftw/3.3.2--bgq-xl--1.0              numpy/1.6.2--bgq-gnu--4.4.6
gsl/1.15--bgq-xl--1.0                papi/4.4.0--bgq-gnu--4.4.6
hdf5/1.8.9_par--bgq-xl--1.0          petsc/3.3-p2--bgq-xl--1.0
hdf5/1.8.9_ser--bgq-xl--1.0          scalapack/2.0.2--bgq-xl--1.0(default)
lapack/3.4.1--bgq-xl--1.0(default)    szip/2.1--bgq-xl--1.0
libint/2.0--bgq-xl--1.0(default)      zlib/1.2.7--bgq-gnu--4.4.6
libjpeg/8d--bgq-gnu--4.4.6
----- /cineca/prod/modulefiles/base/compilers -----
bgq-gnu/4.4.6          bgq-xl/1.0          front-end-gnu/4.6.3 front-end-xl/1.0
----- /cineca/prod/modulefiles/base/tools -----
cmake/2.8.8          irods/3.2          tau/2.21.4
globus/5.0.5         scalasca/1.4.2     tgftp/0.5.0
gtransfer/0.0.10a    superc/1.0         uberftp/2.6
----- /cineca/prod/modulefiles/base/applications -----
abinit/6.12.3          cpmd/v3.15.3          namd/2.9
amber/12(default)      crystal09/1.01        qe/5.0bgq
bigdft/1.6.0           dl_poly/4.03(default) siesta/3.1
cp2k/2.3(default)      gromacs/4.5.5(default) vasp/5.2.12
cpmd/3.15.3_hfx(default) lammps/20120816        vasp/5.3.2
```



General Info

```
[erossi00@fen07 ~]$ module help essl
```

```
-----  
Module Specific Help for /cineca/prod/modulefiles/base/libraries/essl/5.1:  
modulefile "essl/5.1"
```

```
essl-5.1  
Collections of state-of-the-art mathematical subroutines
```

```
-----  
License type: commercial  
Web site:      www.ibm.com  
Download url:  www.ibm.com  
-----
```

```
-----  
ESSL and Parallel ESSL are collections of state-of-the-art mathematical subroutines  
specifically designed to improve the performance of engineering and scientific  
applications on the IBM POWER\u002122 processor-based servers and blades.
```

```
*****  
PLEASE NOTE  
- This library is meant to be linked to a program that will be ran on the compute nodes.  
*****
```

```
libesslbg.a (MPI)  
libesslsmpbg.a (MPI & openmp)
```

Example of usage:

```
$ module load essl  
$ module load bgq-xl  
$ mpicc_r -o a.out foo.c -L$ESSL_LIB -lesslbg -lesslsmpbg  
$ mpixlf90_r -o a.out foo.f90 -L$ESSL_LIB -lesslbg -lesslsmpbg  
-----
```

license

descr

example


```
[erossi00@fen07 ~]$ module show essl
```

```
-----  
/cineca/prod/modulefiles/base/libraries/essl/5.1:
```

```
module-whatis    Collections of state-of-the-art mathematical subroutines  
conflict        essl  
setenv  ESSL_HOME      /opt/ibmmath/essl/5.1/  
setenv  ESSL_LIB       /opt/ibmmath/essl/5.1/lib64  
setenv  ESSL_INC       /opt/ibmmath/essl/5.1/include  
setenv  ESSL_INCLUDE   /opt/ibmmath/essl/5.1/include  
prepend-path LIBPATH /opt/ibmmath/essl/5.1/lib64      :  
prepend-path LD_LIBRARY_PATH /opt/ibmmath/essl/5.1/lib64  :
```



- All libraries was compiled with XL compiler for the compute nodes
`-O3 -g -qsimd=auto`
- Additional version compiled with the gnu compilers is sometimes available in the “advanced” profile
- See “module avail” for available versions
- Usually, only the most recent versions are available in the “base” profile

- **Vendor specific libraries**



- ESSL
(Engineering and Scientific Subroutine Library), version 5.1
- MASS (IBM Mathematical Acceleration Subsystem)

- **Public domain Software**



- LAPACK (Linear Algebra PACKage)
- GSL (Gnu Scientific Library)
- BLAS (Basic Linear Algebra Subprograms)
- FFTW (Fast Fourier Transform)

- **Threaded Parallelism**

- ESSLsmp 5.1



- **MPI Parallelism**

- ScaLAPACK (Scalable Linear Algebra PACKage)
- FFTW (Fastest Fourier Transform of the West)
- PETSc (scalable solution by partial differential equations)





ESSL - Engineering and Scientific Subroutine Library (IBM)
(collections of state-of-the-art mathematical subroutines specifically designed to improve the performance of engineering and scientific applications on IBM platforms)

- High-performance math libraries
- Tuned for BlueGene/Q
- Can be called from C, C++ and Fortran
- Multi-threaded version (SMP) available
- Can be used together with BLAS or LAPACK libraries

libesslbg.a (serial)
libesslsmpbg.a (OpenMP)

```
$ module load bgq-xl  
$ module load essl  
$ mpicc_r -o a.out foo.c -L$ESSL_LIB -lesslsmpbg  
$ mpixlf90_r -o a.out foo.f90 -L$ESSL_LIB -lesslbg
```



- **Solving dense linear algebra problems efficiently on HPC computers.**
- Developed by Argonne National Laboratory, supported by the National Science Foundation (NSF) and the United States Department of Energy (DOE)
- Performance issues are addressed by incorporating recent algorithmic improvements for linear algebra computation. Algorithms give nearly optimal performance.
- Public domain version 3.4.1
- Must be used together with ESSL (or ESSLsmp)
- Some routines already in ESSL
- Attention, some calling sequences are different!

```
module load essl
module load lapack
mpixlf_r name.f
-L$ESSL_LIB -lesslbg
-L$LAPACK_LIB -llapack
-lesslbg
```


ESSL must be linked also after LAPACK to resolve references

- Scalable Linear Algebra PACKage, or Scalable LAPACK
- Includes a subset of LAPACK routines redesigned for distributed memory MIMD parallel computers.
Matrices are laid out in a two-dimensional block cyclic decomposition.
- Univ. of Knoxville, Oak Ridge National Lab, Univ. of Berkeley
- The BLACS is now part of ScaLAPACK, and is compiled into the ScaLAPACK library. It is no longer necessary to link against BLACS libraries.
- The current release 2.0.2 does no longer include sequential LAPACK routines needed, thus an LAPACK library has to be linked together with ScaLAPACK


```
$ module load scalapack
$ module load essl
$ module load lapack
$ module load bgq-xl

$ mpixlf90_r -o a.out foo.f90
-L$SCALAPACK_LIB -lscalapack -L$ESSL_LIB -lesslbg
-L$LAPACK_LIB -llapack -lesslbg
```

Non-application modules by number of users
Period: from 20140201 to 20140228



NAMD
Scalable Molecular Dynamics



- Applications installed on FERMI if suitable for BGQ architecture
- All available through the “module” environment
- Both licensed and open-source applications. Even if licensed, normally users are allowed to use them freely (site licenses)
- We normally maintain the applications installing new releases as available and stay in touch with providers.

- **Abinit:** periodic systems,
DFT with planewave & pseudopotentials
- **BigDFT:** molecules,
DFT with planewave & pseudopotentials
- **CP2K:** molecular simulation of solids, liquids, bio sys.tem
DFT with GPW
- **Crystal09:** periodic systems,
Quantum Chemistry
- **NwChem:** molecules and periodic solids ,
Quantum Chemistry
- **Octopus:** molecules,
TDDFT
- **VASP** atomic scale material modelling
DFT, HF, hybrid functional, planewave



Available applications today: ab-initio Molecular Dynamics

- **CPMD:** DFT with planewave & pseudopotentials, with Plumed
- **NwChem:** molecules and periodic solids ,
mixed Quantum-Molecular mechanics
- **Quantum ESPRESSO:** molecules,
DFT with planewave & pseudopotentials
- **Siesta:** molecular simulation of solids, liquids, bio sys.

DFT with GPW
- **Vasp:** Quantum-mechanical molecular dynamics from first principle
plane wave, ultrasoft pseudo potentials




Available applications today: classical Molecular Dynamics

- **Namd:** classical molecular dynamics
 - **Amber** classical molecular dynamics
 - **Gromacs** classical molecular dynamics
 - **DL_Poly** classical molecular dynamics
 - **Lammps** classical molecular dynamics code that models an ensemble of particles in a liquid, solid, or gaseous state.
-
- **OpenFOAM** (Open Field Operation and Manipulation) CFD Toolbox

- Data refers to February
- Only 30% of the cpu-h are spent in module-loading jobs

Total CPU-hours by modulename for applications
Period: from 20140201 to 20140228





- We report benchmarks on some application installed on FERMI
- Some of them have been executed on FERMI
- Others refer to the “Unified European Application Benchmark Suite” (UEABS) of PRACE;
 - 12 applications in different scientific fields
 - ✓ Particle Physics (1): QCD
 - ✓ Classical MD (2): NAMD, GROMACS
 - ✓ Quantum MD (3): Quantum Espresso, CP2K, GPAW
 - ✓ CFD (2): Code_Saturne, ALYA
 - ✓ Earth Sciences (2): NEMO, SPECFEM3D
 - ✓ Plasma Physics (1): GENE
 - ✓ Astrophysics (1): GADGET
 - For procurement purposes
 - To assist users when choosing the system
 - To provide conversion between systems

NAMD Scalable molecular dynamics

- CP2K is a freely available (GPL) program to perform atomistic and molecular simulations of solid state, liquid, molecular and biological systems. It provides a general framework for different methods such as e.g. density functional theory (DFT) using a mixed Gaussian and plane waves approach (GPW), and classical pair and many-body potentials.
- It is very well and consistently written, standards-conforming Fortran 95, parallelized with MPI and in some parts with hybrid OpenMP+MPI as an option.
- CP2K provides state-of-the-art methods for efficient and accurate atomistic simulations, sources are freely available and actively improved. It has an active international development team, with the unofficial head quarters in the University of Zürich.
- Test Case B The Tier-0 dataset is a 216 LiH system with Hartree-Fock Exchange (PRACE benchmarks on Juqueen)

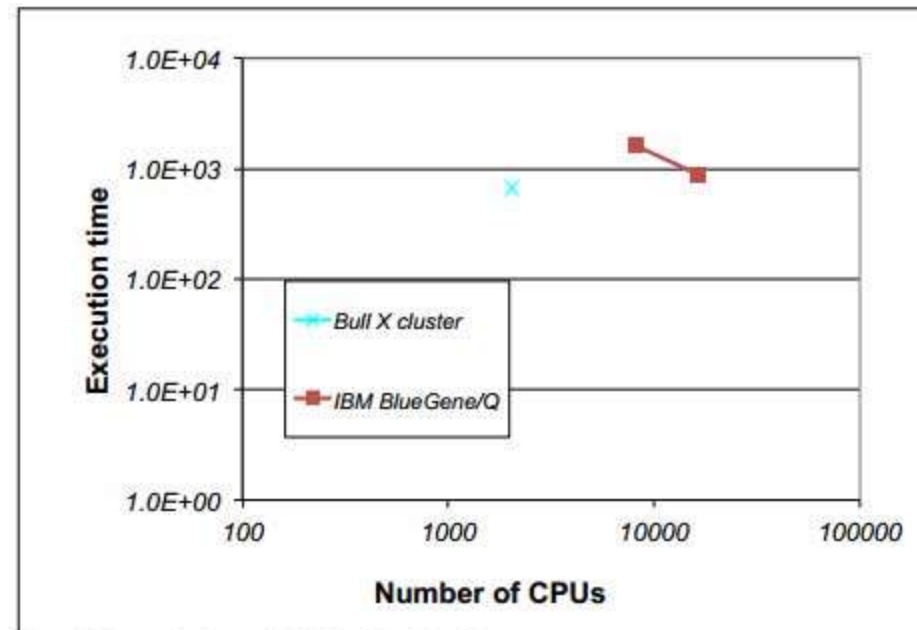


Figure 7 Execution time of CP2K for Test Case B



- GROMACS is a versatile package to perform molecular dynamics, i.e. simulate the Newtonian equations of motion for systems with hundreds to millions of particles.
- It is primarily designed for biochemical molecules such as proteins, lipids and nucleic acids that have a lot of complicated bonded interactions, but since GROMACS is extremely fast at calculating the non-bonded interactions (that usually dominate simulations) many groups also use it for research on non-biological systems, e.g. polymers.
- GROMACS provides extremely high performance compared to all other programs. A lot of algorithmic optimizations have been introduced in the code. The innermost loops are written in assembly with optimized kernels for most common CPU
- GROMACS is written in C and can be run in parallel, using standard MPI communication. Hybrid-MPI/OpenMP implementation is able to push the scaling limits for small to medium systems.

- Test Case B: A model of cellulose and lignocellulosic biomass in an aqueous solution.
- This system of 3.3M atoms is inhomogeneous. This system uses reaction-field electrostatics instead of PME and therefore should scale well.

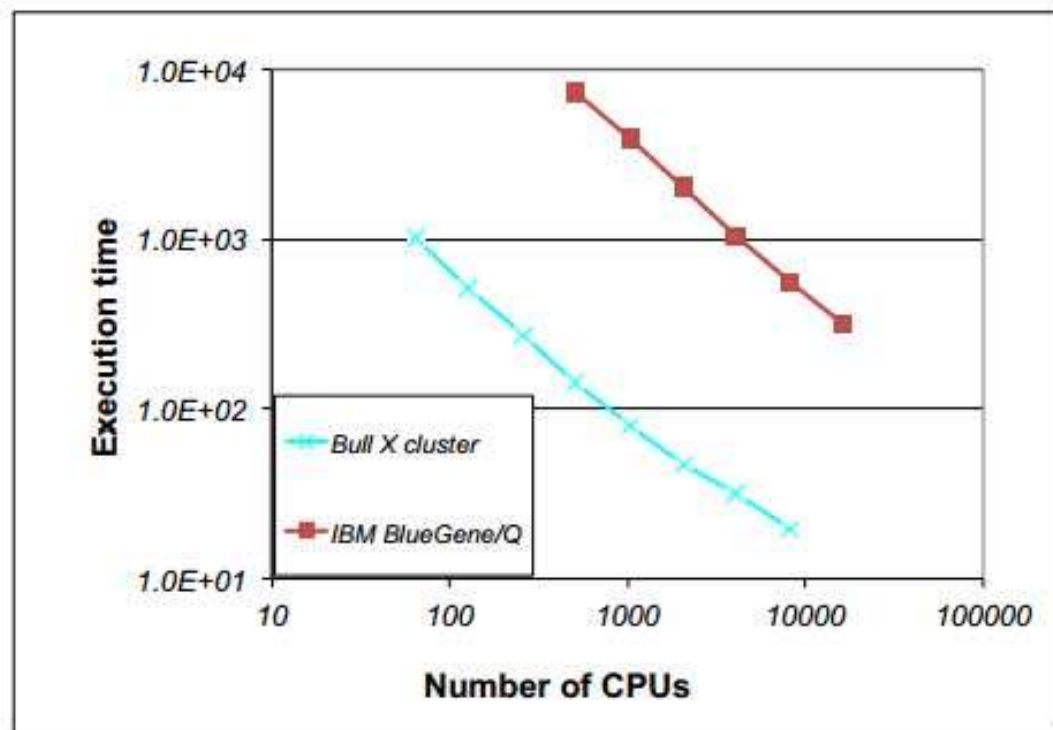


Figure 21 Execution time of GROMACS for Test Case B



NAMD Scalable molecular dynamics


- parallel molecular dynamics code designed for high-performance simulation of large biomolecular systems.
- scales to hundreds of processors on high-end parallel platforms
- Compiled with XL compilers, mixed parallelism MPI + multi-threads, uses FFTW2 (single-precision), BG/Q optimised by developers
- Benchmarks use STMV (virus) (1,066,628 atoms, periodic, PME), as defined on the NAMD web-site

Satellite Tobacco Mosaic Virus
NAMD 2.9, 1.2 MAtoms, BlueGene/Q

- **Graphic 1:** the NAMD performance (simulation time in ns/day) is reported vs. the increasing number of nodes (fixed ranks-per-node=4 and ppn=4)

- bg_size = from 64 to 2048
- ranks-per-node = 4
- ppn =4

- num. of core =
from 1.024 to 32.768



The datasets are based on the original "Satellite Tobacco Mosaic Virus (STMV)" dataset from official NAMD site.

Test Case B: This is a 3x3x3 replication of the STMV dataset.

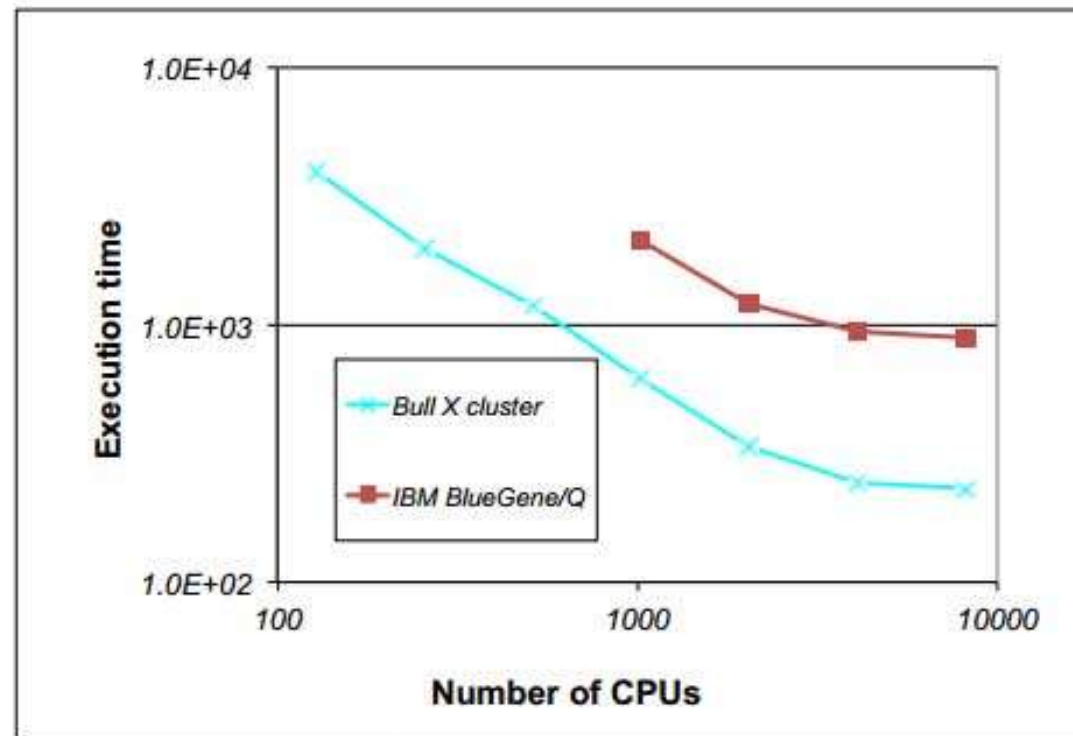
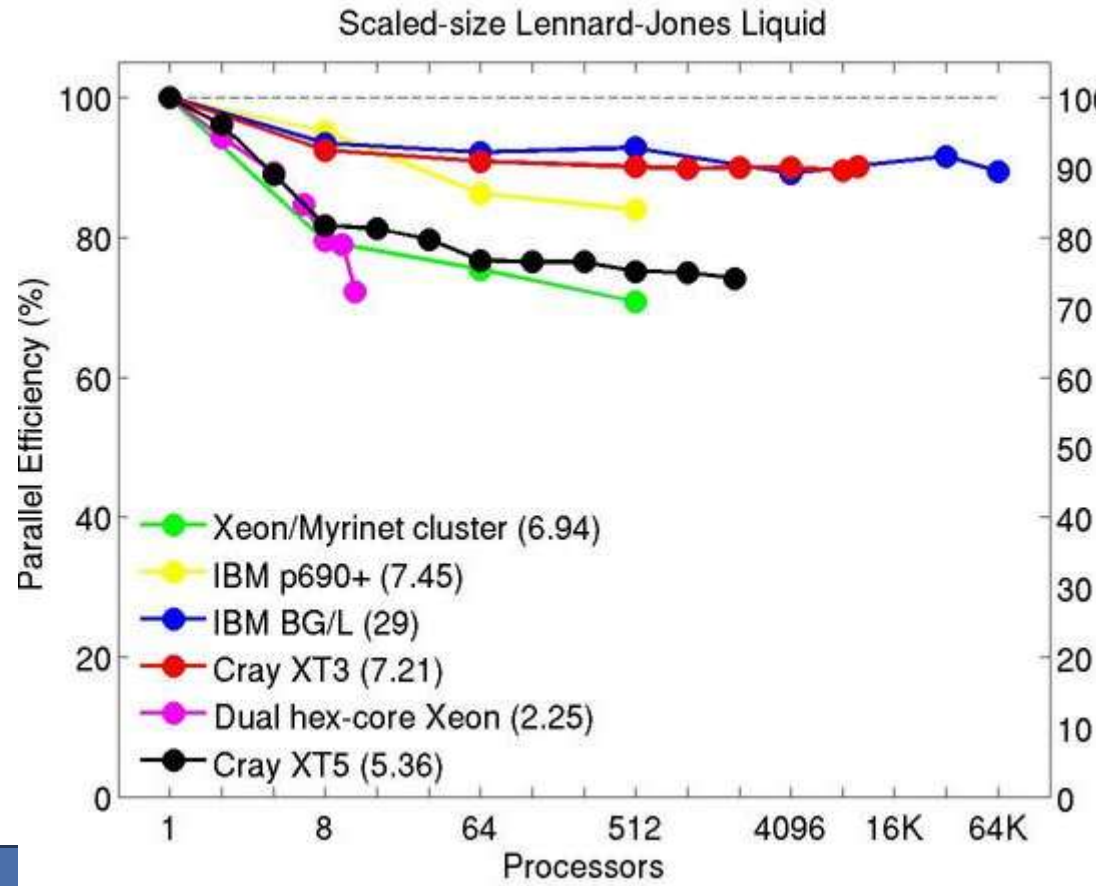


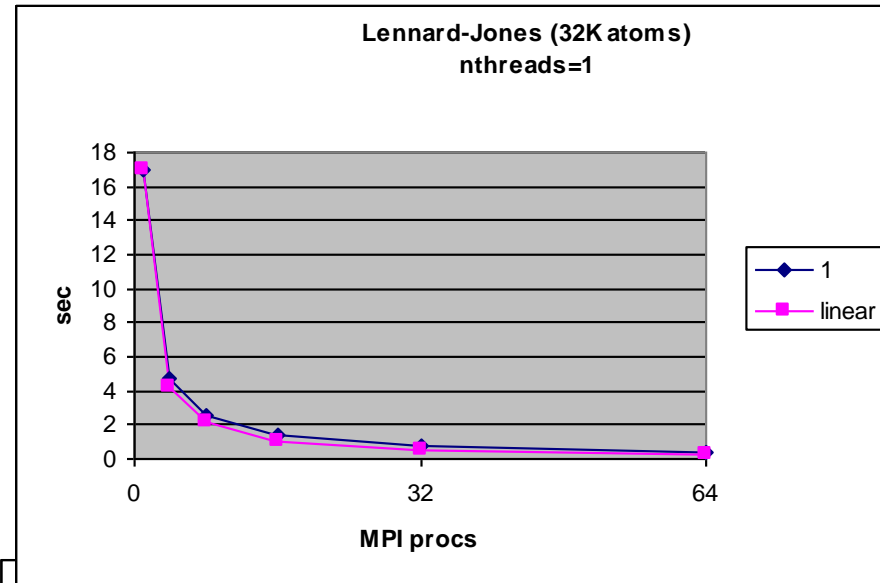
Figure 25 Execution time of NAMD for Test Case B

- classical molecular dynamics code that models an ensemble of particles in a liquid, solid, or gaseous state. It can model atomic, polymeric, biological, metallic, granular, and coarse-grained systems using a variety of force fields and boundary conditions.
- LAMMPS runs in parallel using MPI and a spatial-decomposition of the simulation domain. Each MPI task can distribute work over several OpenMP threads.
- Benchmarks use two different kinds of input
 - **Protein:** Rhodopsin protein in solvated lipid bilayer
 - **Lennard Jones:** atomic fluid with Lennard-Jones potential
- For both types, input dimension range from 32K to 2M atoms

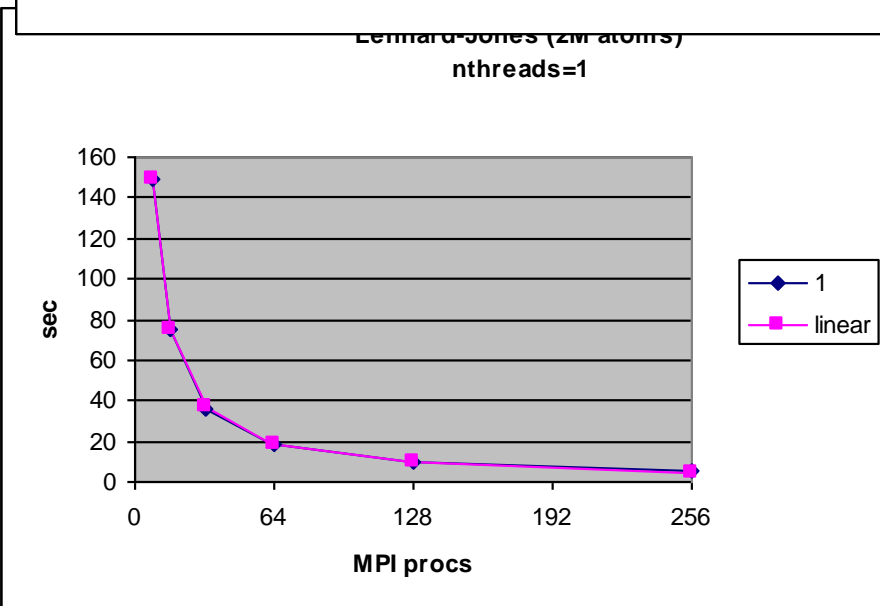




- **LAMMPS** execution time (sec) is reported vs. the increasing number MPI tasks (no OpenMP).

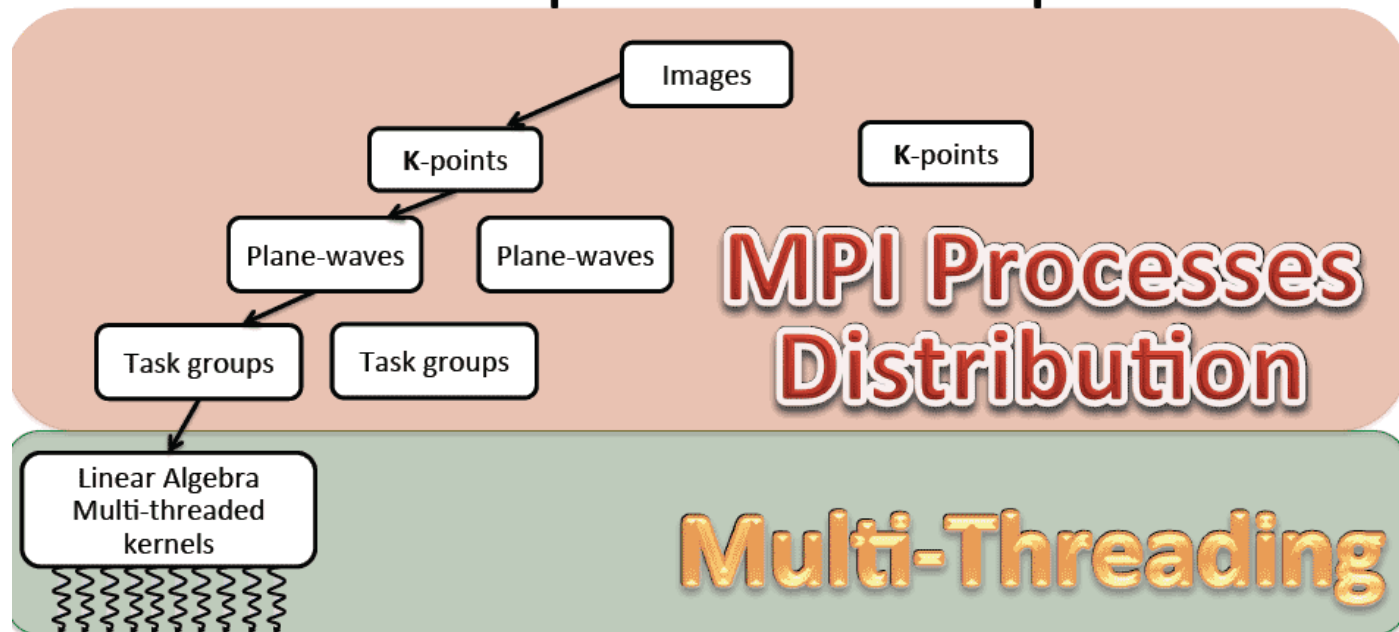


- Scalability is nearly perfect for all the tested cases (small and large, LennardJones and Rhodopsin)



- Quantum ESPRESSO is an integrated suite of computer codes for electronic-structure calculations and materials modeling at the nanoscale. It is based on DFT, plane waves, and pseudopotentials.
- It is written in Fortran/C and is parallelised at different levels (mixed parallelism MPI + OpenMP)

Levels of parallelism in pw.x



- QUANTUM ESPRESSO is written mostly in Fortran90, and parallelised using MPI and OpenMP.
- Test Case B This test case is based on two iterations of a SCF calculation on a functionalised carbon nanotube with a total of 1532 atoms.

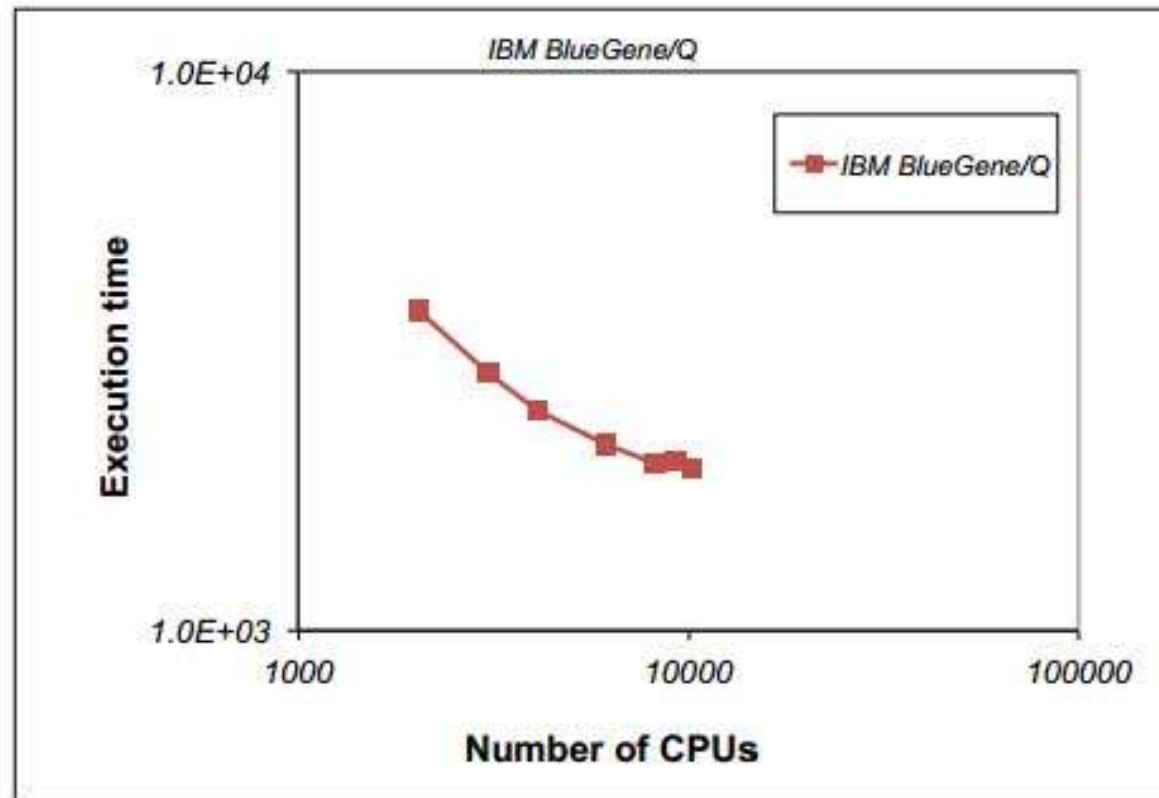


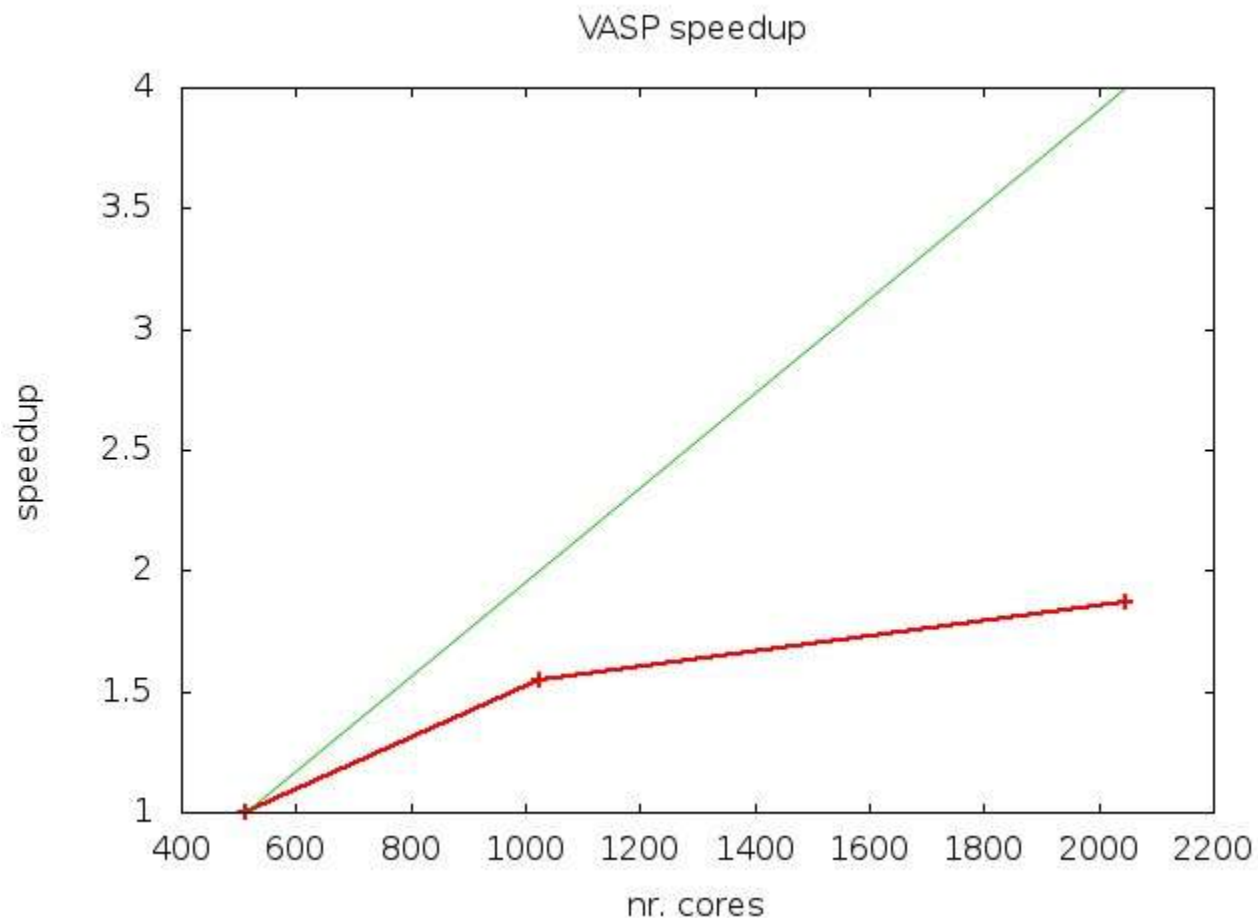
Figure 41 Execution time of Quantum Espresso for Test Case B



- Vienna Ab initio Simulation Package: a computer program for atomic scale materials modelling (electronic structure calculations and quantum-mechanical molecular dynamics, from first principles)
- Based on DFT or HF and hybrid functionals, Green's functions methods and many-body perturbation theory
- Uses plane wave basis sets, norm-conserving or ultrasoft pseudopotentials (projector-augmented-wave method)
- Fortran code, parallelisation via MPI library



- Test performed within a Prace project. It shows a quite good speedup up to 1 K processors.
- Scalability depends on the method and input data.
- Memory is critical: for large cases it is better to use ranks-per-nodes= 8 or lower
- Modifying input parameters allow for important gains in performance (50-80%).



Resources

- Hardware
- Software
 - Application Software
 - Chemistry
 - Physics
 - Life Science
 - Engineering
 - Astronomy
 - Visualisation
 - Maths Libraries
 - Data Libraries
 - All Software
 - Production Environment
- Reports on systems usage

[Home](#) > [Resources](#) > [Software](#)

Application Software for Science

Content:

- [Chemistry](#)
- [Physics](#)
- [Life Science](#)
- [Engineering](#)
- [Astronomy](#)
- [Visualisation](#)
- [Maths Libraries](#)
- [Data Libraries](#)
- [All Software](#)

Application Software for Science & Technology

Cineca offers a variety of third-party applications and community codes that are installed on its HPC systems.

Most of the third-party software are installed using the software modules mechanism.

The packages available and detailed descriptions of them can be viewed for discipline by selecting the menu on the left.


If you want to see the full catalog, please get [all Software in alphabetical order](#).

The information in this list may not reflect all software products available at Cineca. If you do not see an application you are interested in, or if you have questions about software that is currently available, please contact the [Help Desk](#)

Information and announcements

For users

- My portal
- Getting started
- Get in touch**
- Help desk
- Documentation



Home > For users

Get in touch

How to get Center's announcements (HPC-news)

We manage a mailing list (HPC-news) for posting announcements, scheduled downs, software updates, any problems and so on, about our HPC computing resources. It is advisable for HPC users to be included in that list!

You can subscribe (or unsubscribe) to HPC-news by sending an email from the address you want to subscribe. You can consult the archive browsing the archive web site.

To subscribe to HPC-new:

send a mail to listserv@list.cineca.it
 in the body --> "subscribe hpc-news"
 in the subject --> any string...

To unsubscribe to HPC-new:

send a mail to listserv@list.cineca.it
 in the body --> "unsubscribe hpc-news"
 in the subject --> any string...

To consult the archive:

for the recent archive go to → [Center News](#)
 for the full archive go to → <http://list.cineca.it/archives/hpc-news.html>

Help desk



Service Closed

Center news

- 23/11/2012
[Power outage](#)
- 23/11/2012
[Reboot of FERMI login node fen01](#)
- 20/11/2012
[FERMI shut down for management operation](#)
- 14/11/2012
[Call for proposals for DECI-10 call for Tier-1 resources](#)

» more Center News