HPC for Computational Physics at INFN: The SUMA project

https://web2.infn.it/SUMA/

November 16-19, 2015

Introduction

The INFN theoretical community is active in several scientific areas that require significant computational support. These areas stretch over a wide spectrum, requiring in some cases fairly limited computing resources, but in most cases huge computing power is required. Examples in this class are LQCD, coputational fluid-dynamics and astrophysics, dynamical systems and classical and ab-initio simulations of bio-systems.

At the same time, for most groups active in these areas, it is becoming more and more difficult to develop their computational strategies and algorithms in a way that allows to adapt to the increasingly fast changes happening in high performance computing architectures.

Last but not least, several existing INFN projects have produced significant progress on technological developments that may be crucial building blocks for new generation HPC systems.

SUMA plans to support this community, and at the same time aims to explore all suitable ways in which the technological developments made at INFN can be put to good use for the present and future needs of computational physics.

The SUMA project works in close collaboration with academia and computer centers in Italy, such as the Universities of Ferrara, Parma, Pisa and Rome, SISSA (Trieste) and CINECA (Bologna).

SUMA Computing Resources

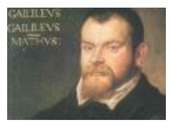

GALILEO
Model: IBM NeXtScale
Nodes: 512
Processors: Intel Haswell 2.4 GHz
Cores: 16 (2x8) per node, 8256 cores in total
RAM: 128 GB/node, 8 GB/core
Network: Infiniband 4x QDR
Accelerators:
- 2 Intel Phi 7120p/node on 384 nodes (768 in total)
- 2 NVIDIA K80/node on 40 nodes (80 in total)

Jointly procured by CINECA and INFN

ZEFIRO
Model : Linux cluser
Nodes: 32
Processors: AMD Opteron 6380 2.5 GHz
Cores: 64 (4x16) per node, 2048 cores in total
RAM: 512 GB / node
Network: Infiniband DDR
Accelerators: no
System installed and managed by INFN-PISA

Relativistic astrophysics

Scientific challenge: Simulation of inspiral and merger phase of binary system involving Neutron Stars and Black Holes and the modelling of the associate gravitational wave signal.
Computational challenge: Time evolution of a set of PDE on cartesian grid using adaptive mesh refinements.

Einstein Toolkit (http://einsteintoolkit.org)
Open source set of tools for simulating and analyzing relativistic astrophysical systems.
Based on Cactus (http://cactuscode.org)
About 500K lines of code (C, C++, Fortran) with openMP and MPI support

Results on ZEFIRO
- Inspected differences MPI vs openMP
- MPI shows batter scaling than OpenMP

Results on GALILEO
- Explored Strong and weak scaling
- Less sensitive to MPI and openMP
- Scaling improves increasing volume

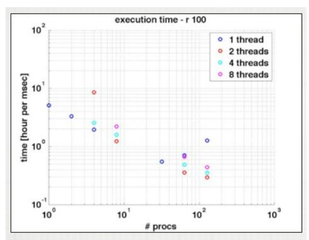

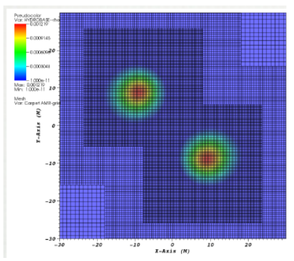

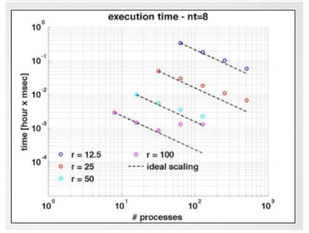

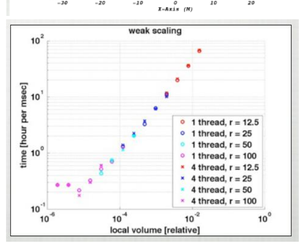

Quantitative Biology

The new opportunities offered by HPC are opening the way to attack new problems in classical and *ab initio* MD simulations for more realistic systems composed by a large number of atoms, thus allowing for a better interpretation of experimental data

Ab initio calculation of X-ray Absorption Spectra

We perform *ab initio* calculation of the X-ray Absorption Spectrum of Cu(II) in water. Spectrum is computed with the help of the ***Xspectra*** code, using the electronic potential computed via DFT by ***QuantumESPRESSO.*** Relevant atomic model structures are produced by ***NAMD***

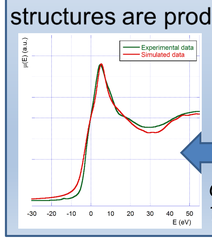

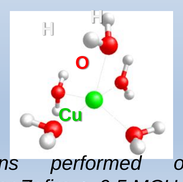

Calculations performed on the Theocluster Zefiro. > 0.5 MCH used.

Radiopeptides vs tumors

MD simulations can be used to study the interaction between peptides and membranes of tumoral cells. With this information we are able to design a vector for radio-nuclides capable of binding to the tumor membrane but at the same time with little affinity to healthy tissue

*MD simulations use the **GROMACS** suite*
Systems with > 42,000 atoms can be simulated on 16 8-cores CPUs at the CINECA Galileo cluster

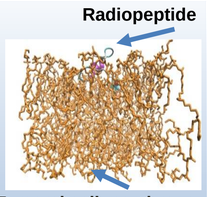

Ab initio X-ray spectra simulations can be profitably exploited in the difficult case of bio-molecules in complex with metals ions. This is relevant to study the formation of protein fibrils that are typically found in the cerebral tissue of people affected by the Alzheimer's and Creutzfeld-Jakob's disease. The process of fibril formation is indeed influenced by the presence of metallic ions [1,2] and can be studied by X-ray spectroscopy .

Machine	GROMACS/NAMD (classical MD)	QuantumEspresso (*ab initio* MD)	XSPECTRA (X-ray spectra calculation)
Galileo	√	√	√
Fermi	√	√	√
Zefiro	√	√	√

Directive based programming of the offload code

Directive based programming is a fundamental component to keep the offloading code **portable, readable and maintainable.** The performance is not at the level of low level programming (e.g. CUDA) but it will improve as compilers implementations are in their early stage.

Directive based compilers available on the SUMA systems (GALILEO) :

- **openMP4** for Xeon PHI : Intel compiler and GCC6 (experimental, for Knight Landing or emulator)
- **openACC2** for NVIDIA K80 : Portland Compiler and GCC6 (experimental)

```
#pragma omp target
#pragma omp parallel for
for (i = 0; i < n; ++i)
        computeIntensiveFunct();
```

```
#pragma acc parallel
#pragma acc loop
for (i = 0; i < n; ++i)
        computeIntensiveFunct();
```

LQCD : twisted mass operator

$$\not{D}_{\rm TM} = \not{D}_{\rm W} 1_f + i\mu\gamma_5\tau_3$$

- kernels for symmetric even-odd preconditioning:
  - Dslash: $\chi = R^{-1}\, \not{D}_{\rm W}\psi$
  - AChiMBDPsi: $\phi = A\chi - b\, \not{D}_{\rm W}\psi$

  with $A = R \equiv 1 + i2\kappa\mu\gamma_5$

  the twist operator.
- …plus daggered versions.
- the successive application of these kernels corresponds to the Schur decomposed operator

$$\tilde{M}_{oo} = R_{oo} - \frac{1}{2}\kappa\, \not{D}_{oe} R_{ee}^{-1}\, \not{D}_{eo}$$

- to implement these kernels we mainly have to work on the code-generator
- the high-level part has to be modified to include the twisted mass parameter μ and to call the corresponding twisted mass low-level kernels

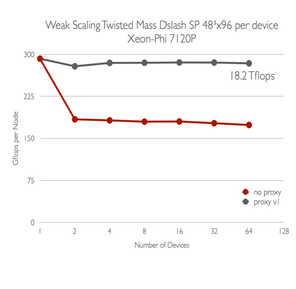

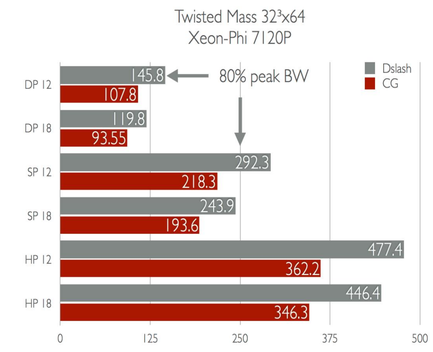

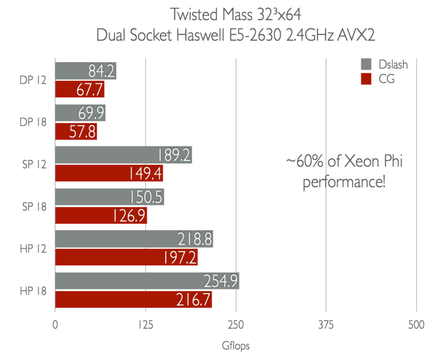

LQCD: QCD in extreme conditions

Part of our research [3,4] is dedicated to the study of strong interactions under extreme conditions, i.e. conditions which have been realized in the early stages of the Universe or which are reproduced in some experiments (e.g., ultrarelativistic heavy ion collisions), and are characterized by extremely high temperatures (exceeding 10^12 Kelvin degrees), densities, or by extremely high magnetic fields (up to 10^16 Tesla). New phenomena are expected in such conditions, such as the deconfinement of quarks and gluons.

For these studies, we are currently performing numerical simulations on lattices as big as 48^3 x 96 and lattice spacings below 0.1 fm. This is possible also due to various supercomputing resources available within the SUMA project.
Future prospects will rely on the development of new computing infrastructures. Our present efforts are directed towards the use of multiGPU architectures: we have already done some progress along this line in the recent past [5] and we are currently exploring new programming platforms (OpenACC) and direct communications among GPUs

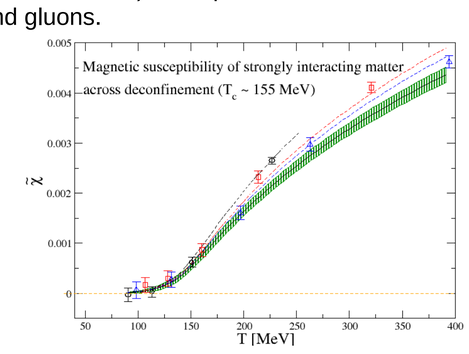

Computational fluid dynamics

Lattice Boltzmann methods are widely used in computational fluid dynamics, to describe flows in two and three dimensions:

- D2Q37 Lattice Boltzmann method
- 7600 DP operations /site
- Good scaling over tens of GPUs

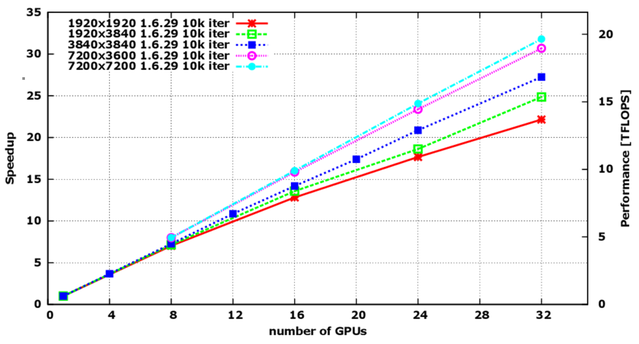

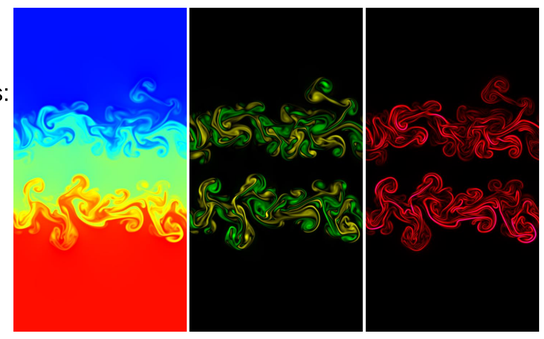

Simulation of the Rayleigh-Taylor instability. The pictures show temperature-map (left), vorticity (center) and temperature-gradient (right).

The plot on the left shows the speedup and the aggregate performance on a cluster of NVIDIA K80 GPUs.

References

[1] F. Stellato et al., Copper-Zinc cross modulation inprion protein binding. European Biophysics Journal 43 (2015) 631-642.
[2] P. Giannozzi, et al., Zn induced structural aggregation patterns of β-amyloid peptides by first-principle simulations and XAS measurements. Metallomics (2012) 4, 156-165
[3] C. Bonati et al., Magnetic Susceptibility of Strongly Interacting Matter across the Deconfinement Transition, Phys. Rev. Lett. 111 (2013) 182001, [arXiv:1307.8063 [hep-lat]].
[4] C. Bonati et al., Curvature of the chiral pseudocritical line in QCD, Phys. Rev. D 90, (2014) 114025, [arXiv:1410.5758 [hep-lat]].
[5] C. Bonati et al., QCD simulations with staggered fermions on GPUs, Comput. Phys. Commun. 183, (2012) 853, [arXiv:1106.5673 [hep-lat]].
[6] G. Crimi et al., Early Experience on Porting and Running a Lattice Boltzmann Code on the Xeon-Phi Co-Processor, Proceedings of the International Conference on Computational Science, ICCS 2013, Procedia Computer Science 18 (2013) 551-560.
[7] F. Mantovani et al., Exploiting parallelism in many-core architectures: a test case based on Lattice Boltzmann Models, Proceedings of the Conference on Computational Physics, Kobe, Japan (in press).
[8] A. Bertazzo et al., Implementation and Optimization of a Thermal Lattice Boltzmann Algorithm on a multi-GPU cluster, Proceedings of Innovative Parallel Computing (INPAR) 2012.
[9] F. Mantovani et al., Performance issues on many-core processors: A D2Q37 Lattice Boltzmann scheme as a test-case, Computers and Fluids 88 (2013) 743-752, 10.1016/j.compfuid.2013.05.014
[10] L. Biferale et al., Optimization of Multi - Phase Compressible Lattice Boltzmann Codes on Massively Parallel Multi - Core Systems, International Conference on Computational Science (ICCS), Procedia Science 4 (2011) 994-1003
[11] L. Biferale et al., An Optimized D2Q37 Lattice Boltzmann Code on GP-GPUs Computers and Fluids 80 (2013) 55-62, 10.1016/j.compfuid.2012.06.003

.