# Why and How Using HPC in University Teaching? A Case Study at PoliTo

# Why and How Using HPC in University Teaching? A Case Study at PoliTo

Nicolò Nepote, Elio Piccolo, Claudio G. Demartini and Paolo Montuschi,
*Dipartimento di Automatica e Informatica - Politecnico di Torino*
*C.so Duca Degli Abruzzi 24, 10129 Torino (TO)*
*[nicolo.nepote, elio.piccolo, claudio.demartini, paolo.montuschi]@polito.it*

*After an era of "personal-computers-only", supercomputing facilities and services are coming back to Universities to support research activities and computationally intensive simulations, but with some important differences with respect to the past. Besides the technological issues, while in the seventies-eighties the scene was dominated by mainframes, managed by skilled system managers and most of times used by operators with good computer expertise, today, the widespread and pervasive use of computers has lead to a completely different scenario. The demand for computation resources is emerging from a wide variety of areas and disciplines and mostly by users with basic expertise in computers, who however most of times request to have full control of the computation resources. Within this picture and especially at University level, High Performance Computing (HPC) has emerged as a good tradeoff to meet the different demands, and at the same offering good services at reasonable setup and maintenance costs. Traditionally, HPC has been and is being mostly used in support to applied research, but more recently some questions have emerged: - How much is it reasonable to offer HPC also to some teaching activities? What are the problems, advantages, and drawbacks? Is this the "right" way or should HPC resources be directed to research only? At Politecnico di Torino we tried to respond to these questions and started a test project called HPC-4-teaching. In this paper we present the results achieved by this project on a small set of courses during one full academic year.*

# Why and How Using HPC in University Teaching? A Case Study at PoliTo

Nicolò Nepote, Elio Piccolo, Claudio G. Demartini and Paolo Montuschi,
*Dipartimento di Automatica e Informatica - Politecnico di Torino*
*C.so Duca Degli Abruzzi 24, 10129 Torino (TO)*
*[nicolo.nepote, elio.piccolo, claudio.demartini, paolo.montuschi]@polito.it*

*After an era of "personal-computers-only", supercomputing facilities and services are coming back to Universities to support research activities and computationally intensive simulations, but with some important differences with respect to the past. Besides the technological issues, while in the seventies-eighties the scene was dominated by mainframes, managed by skilled system managers and most of times used by operators with good computer expertise, today, the widespread and pervasive use of computers has lead to a completely different scenario. The demand for computation resources is emerging from a wide variety of areas and disciplines and mostly by users with basic expertise in computers, who however most of times request to have full control of the computation resources. Within this picture and especially at University level, High Performance Computing (HPC) has emerged as a good tradeoff to meet the different demands, and at the same offering good services at reasonable setup and maintenance costs. Traditionally, HPC has been and is being mostly used in support to applied research, but more recently some questions have emerged: - How much is it reasonable to offer HPC also to some teaching activities? What are the problems, advantages, and drawbacks? Is this the "right" way or should HPC resources be directed to research only? At Politecnico di Torino we tried to respond to these questions and started a test project called HPC-4-teaching. In this paper we present the results achieved by this project on a small set of courses during one full academic year.*



## 1. Introduction

Recent years have been witnessing both the birth and widespread usage of High Performance Computing (HPC) Centers in academia [Kindratenko & Trancoso, 2011]. Among the stimulating factors there is the high demand for large computational resources by research groups operating in different disciplines and research areas such as: meteorology, nanotechnology, biotechnology, aerospace engineering, automotive, artificial vision, speech recognition, and many other.

In this scenario, from the viewpoint of the hardware we can observe that the number of possible architectures and implementations of HPC is still large and currently we have several models of parallel systems (MIMD vs. SIMD, CPU/GPU-only vs. hybrid, dedicated systems vs. "HPC in the cloud" open vs. proprietary interconnection, etc.), to choose from, when starting from scratch the design and deployment of a new HPC system, as no one is clearly emerging over the other as the "final-and-only-solution". Similarly, software is under constant and continuous development and today it is already possible to find custom software packages satisfactorily solving some, but unfortunately not all, classes of problems. All of this leads to the conclusion that HPC is a still growing but yet a sufficiently mature technology that can be offered to a wider audience as a tool to solve some types of problems at reasonable costs [Gentzsch 2009].

This rapid growth lets us to expect a rapid and ready-to-come pervasive diffusion of HPC systems [Kindratenko & Trancoso, 2011] similar to the spread of mainframes during the '60s, or microprocessor systems during the '80s, but in a completely different scenario. In fact, while a sort of technological standardization favored the former, this, at least today, is not necessarily holding in modern HPC systems and related software applications. All of this also allows us to expect that in the near future there will be a growing demand for engineers with a strong HPC background, together with very specific, but differentiated, skills.

Clearly, this evolution opens new challenges on the educational side [Glotzer et al. 2009] targeted to the preparation and formation of engineers possibly fitting those profiles [Gibbon et al. 2010].

During the academic year 2011-12 at Politecnico di Torino we have started a test program aimed at offering the students the opportunity to use our HPC systems during normal teaching courses. So far our results have been satisfactory, as we well discuss later in this paper. We are aware that ours is only a first step, and that the architectural and computational skills of our students have to be complemented by notions of mathematics, parallel algorithms, as well as by dedicated knowledge in their own special field of expertise, such as biology, medicine, environmental safety, just to mention a few.

Our aim for the close future is to create an environment where students can work with our HPC in a wide range of laboratory classes, thus improving their own expertise to use the HPC technologies in various fields.

The paper is organized as follows: Section 2 describes the differences between the use of HPC in research and in teaching. Section 3 presents as a case our experience of HPC@POLITO to introduce supercomputing in some courses at the Politecnico di Torino. Section 4 proposes some possible future lines and draws our conclusions.

## 2. HPC supporting teaching activities

It is a matter of fact that in academia HPC systems are mainly used by researchers operating in different disciplines, usually requiring different computing power, involving both small and big data, requiring the execution of both third party and internally developed software, etc. From the viewpoint of users these activities are rather uniform in terms of priorities, and this had led over time to the definition of a standard Service Level Agreement template, commonly adopted by computing centers.

This scenario changes when the same HPC systems need to be used by the students [Joiner et al. 2006], mainly because the priorities, time constraints and latencies are different, as reported in Table 1.

Besides the need to offer the access of the HPC system to a new class of end-users (i.e. students) who in most cases do not have specific expertise in parallel systems, the other big challenge is to offer a user friendly and responsive service where the typical teaching and time priorities have to be sharply met during a laboratory lecture. In the next section we present how we have tackled this challenge at our HPC Center.

| Scenario | Priorities |
| --- | --- |
| Research | • Ensure sufficient computing slots to execute thousands of concurrent jobs by hundreds of users, with no time limit throughout the year;<br>• Management of large amounts of data, typically in the order of TBytes;<br>• Fast data access, especially during job execution;<br>• Job waiting time not longer than a few hours;<br>• Job execution time not longer than a few days;<br>• Data security. |
| Teaching | • Ensure sufficient computing slots to execute hundreds of concurrent jobs by tens of users, within a few hours;<br>• Management of many data packets between 1KB and 250GB;<br>• Fast data transfer before and after job execution;<br>• Fast data access during job execution;<br>• Job waiting time not longer than a few seconds;<br>• Job execution time not longer than three hours. |

**Table 1: Different priorities between HPC in research and teaching**

# 3. Our Case Study

## *3.1 A bit of history*

In order to better focus our experience at Politecnico di Torino, it is worth to have a quick look at the history of the Initiative initially called “Casper project”. A timeline of the entire project history from 2008 to today is presented in Figure 1.

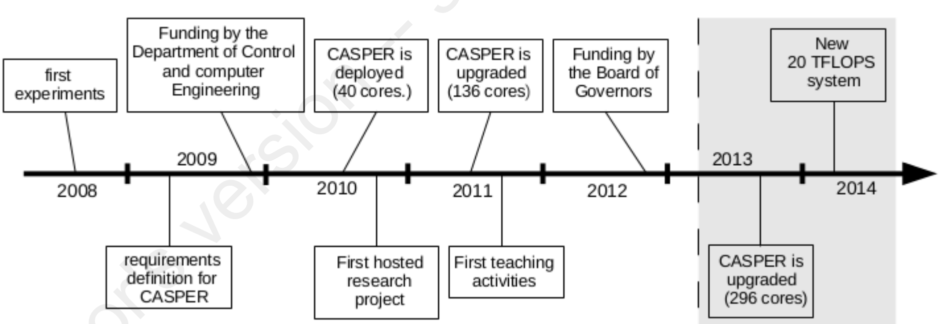


**Figure 1: The history of HPC at Politecnico di Torino**

The possibility of introducing high-performance computing in the teaching activities of the Politecnico di Torino was one of the initial goals and ideas in 2008 as the *DAUIN (Department of Control and Computer Engineering) HPC Initiative* started up, but, because of lack of funding, it was not started in the first phase of deployment of the *Initiative*. In 2011, after two years of tests on a reduced scale and with a limited number of computers, the Initiative scored a first important success when both several research groups provided some funding, and the **first results came out in terms of scientific products** and sustained use of computing resources. At the end of the academic year 2011-12, the *Initiative* counted 25 hosted projects and 12 papers developed thanks to our HPC and published by groups operating in different research areas.

In November 2012 the Board of Governors of the Politecnico di Torino approved a funding of € 250,000 on a 3-years-basis, to support and **transform the DAUIN HPC initiative into a campus level computer center** and renamed it HPC@POLITO.

HPC@POLITO [www.hpc.polito.it] aims at providing computational resources and technical support to both academic research, university teaching and student labs. To pursue these goals, the computing center has set up during the last two years a heterogeneous cluster that can be classified as a PRACE Tier-II facility, called CASPER (Cluster Appliance for Scientific Parallel Execution and Rendering) based on a MIMD distributed shared memory InfiniBand system, with 136 cores, 584 GB of memory and a peak performance of 1.2 TFLOPS.

A detailed technical description of the system in its software and management aspects is available in a previous paper [Della Croce et al. 2011], while the complete up-to-date hardware specifications are found in Table 4.

In our vision **CASPER besides its initial configuration, is a continuously evolving system**, keeping pace with the times and user requirements, developed in collaboration with several research groups, and being maintained and renewed on a regular basis. In our plans, the recent funding of the Board of Governors should bring this cluster to achieve 3 TFLOPS by 2013 and about 20 TFLOPS by 2014.

### *3.2 HPC-4-teaching: participating courses and related bounds*

Based on the very positive results obtained while using HPC@POLITO for research activities, we decided to test the opening of the usage of CASPER to the laboratory-based teaching activities of some selected courses, run by Professors who already had the expertise of the whole system.

The main goal was to **give students the opportunity to use the cluster simulating real use cases**, by crafting ad-hoc laboratory exercises. We planned to provide students with the access to the system by means of industry standard tools, usually offered by major national and international computing centers. For this “extension”, we identified the constraining need to end the computing activities within the time slot(s) of the laboratory lecture, as one of the biggest constraints.

From a practical point of view, this implied that we had to design an environment capable of targeting the end of the computations with the normal laboratory lecture frames, i.e. usually between 90 and 180 minutes. To properly address the time constraints as well as major user potential requirements, we have identified two main types of students, as very likely involved and interested in becoming users of the system:

- **Students from Master Courses** (Laurea Magistrale – L.M.) as, depending on their discipline and expertise, it could be very likely that they will need HPC resources in their future job and/or research activities.
- **PhD students**, with interest and expertise in systems, sharing the perspective to become future system managers, parallel software developers, designers and/or researchers of HPC systems.

After an analysis of the portfolio of courses offered at Politecnico di Torino, we then directed our efforts towards the involvement of both courses at the Master level in Engineering and belonging to the PhD curriculum in Computer and Systems

Engineering. In this initial phase of opening HPC to teaching activities, we got a positive response to our invitation to join the Initiative, by the courses reported in Table 2.

| Course | Program | Professor | Activity to be carried out | Students | Lab hours |
|---|---|---|---|---|---|
| Flussi Turbolenti *(Turbulent Flows)* | L.M. in Ingegneria Aerospaziale | D. Tordella | Computational fluid dynamics | 30 | 42 |
| Fluidodinamica *(Fluid Dynamics)* | L.M. in Ingegneria Matematica | D. Tordella | Analytical and Computational Fluid Dynamics | 25 | 20 |
| Realtà Virtuale *(Virtual Reality)* | L.M. in Ingegneria Informatica, Ing. del Cinema; | A. Bottino | Rendering | 64 | 21 |
| Parallel and Distributed Computing | Ph.D. in ingegneria Informatica e dei Sistemi | B. Montrucchio | Parallel programming | 73 | 12 |
| **Total** | | | | **192** | **95** |

**Table 2: Courses involved in our project**

In Table 3 we have reported the most common activities by each of the courses of Table 2, where it can be observed that, as “boundary conditions” are rather different from one course to the other, we have been stimulated to identify and implement different frameworks to meet each course’s requirements.

As shown in Table 3 the courses have very different characteristics. This implies the need to develop different frameworks (i.e. customizations) each one meeting the following (main) constraints and requirements:

- **Real-time execution.** We must ensure that the execution of the jobs submitted by students will begin almost immediately and will finish by the end of the laboratory session.
- **Fast data transfer.** The inbound and outbound transfer times of the whole data necessary to the jobs (the binary code plus optional input, output and intermediate data) has to be minimized and made compatible with the previous item.
- **Out-of-order post-processing.** Optional post-processing (and longer tasks) are also performed on the cluster, but can also be executed in the next lab session or between two sessions.
- **Real-life user experience.** It has to be offered to our students-end-users, a flexible, user-friendly and direct access to the cluster, possibly avoiding the need to create a single account for each student.
- **Tile-rendering system.** It is necessary to deploy a software that is capable of performing tile-rendering operations on any “*.blend*” file. This software must be seamlessly integrated with the *Open GridEngine* by acting as a de facto sub-scheduler.

| Teaching | Tasks | Characteristics |
|---|---|---|
| Flussi turbolenti | Preparing the model and domain of some fluid interaction simulations using in-house Navier-Stokes codes [Iovieno et al. 2001, Tordella & Iovieno 2011], as well as commercial software STAR-CCM+ with MPI enabled. Running the simulations as jobs on the cluster.<br>Extracting images and producing low-res mpegs in post-processing using VisIt. | *Input:* 1 GB<br>*Output:* 250 GB<br>*Lab duration:* 180'<br>*Constraint:* data+time |
| Fluidodinamica | Computation of the temporal life of three-dimensional perturbation waves in sheared flows as the channel flow, the wake and the boundary layer in cross flow (Initial Value Problem for the Orr-Sommerfeld and Squire equations) [Scarsoglio et al. 2010] | *Input:* 5 MB<br>*Output:* 100 GB<br>*Lab duration:* 90'<br>*Constraint:* data+time+SW |
| Realtà Virtuale | Preparing the 3D model of some static high-resolution stereoscopic images using Blender 3D and the Yafaray ray tracer. Distributed rendering on the cluster. Creating and visualizing the resulting image using 3D-Vision active shutter glasses by *nVidia*. | *Input:* 100 MB<br>*Output:* 10 MB<br>*Lab duration:* 90'<br>*Constraint:* time+SW |
| Parallel and Distributed Computing | Writing parallel code in C using the MPI libraries. Running the binaries on the cluster and debugging the code. | *Input:* 10 KB<br>*Output:* 10 KB<br>*Lab duration:* 180'<br>*Constraint:* time |

**Table 3: Tasks to be performed in the lab exercitations of the courses involved**

Each of the points above has required us to pay a special attention while designing how to assemble and configure the hardware and software resources to be, time-by-time, dedicated to each single teaching activity.

### *3.3 Increased computation power and new cluster configuration*

Simulations on the CASPER cluster at HPC@POLITO in its late 2010 configuration shown that the computation power was not sufficient to meet the requirements necessary to be used in a teaching framework like HPC-4-teaching. We then identified the need to perform an overall upgrade of the cluster according to the following main specifications:

- a doubling in the number of cores available on CASPER (at least);
- an increase in the memory/core ratio from 3 up to 4 GB;
- an increased storage capacity from 6 up to 20 TB;
- a high-priority real-time *Execution Queue* and a new designed *Parallel Environment* for the students lab in the scheduler;
- the deployment of an integrated *tile rendering* system.

The requirements were implemented through technical improvements performed on the scheduler and through new acquisitions of hardware:

- four AMD Opteron *Magny-Cours* computational nodes were added to the cluster, providing a total addition of 96 cores and 512 GB of RAM;

- a new 20 TB midrange NAS with an aggregated link of 4 Gbps was added to the nodes;
- two new symmetrical *Parallel Environment* and two related fast *execution queues* spanning through respectively 40 and 96 cores, were created in the scheduler;
- a web portal was created to facilitate the monitoring of the waiting and running jobs by the students

To address the rendering issues, we have considered what open source software was offering. Unfortunately, as we have not found any tool meeting our specific needs, we have solved the problem by provisionally sketching a dispatcher that:

- divides the so called “cameras” inside a *.blend* file into small tiles;
- has them dispatched by the *Sun GridEngine* scheduler between the compute nodes of the cluster;
- controls the rendering operated by *Blender 3D;*
- finally collects each tile and rebuilds the whole rendered image using *ImageMagick*.

These upgrades have led CASPER to reach the current configuration (mid 2011) summarized in Table 4.

| Item | value |
|---|---|
| Architecture | Linux Infiniband-DDR MIMD Distributed Shared-Memory Heterogeneous Cluster |
| Interconnection | Infiniband DDR 20 Gb/s (some nodes) |
| Network | Gigabit Ethernet 1+1 Gb/s |
| CPU family | AMD Magny-Cours and Intel Nehalem |
| Sustained performance (Top500 index) | 1.2 TFLOPS (R*max*=1.238e+03 Gflops) |
| Peak performance | 1.4 TFLOPS (R*peak*=1.474e+03 Gflops) |
| Green500 Index | 422.31 MFLOPS/W |
| Power consumption | 2.6 kW |
| Cores | 136 |
| Compute nodes | 9 |
| Total memory | 584 GB DDR3 |
| Storage | 20 TB |
| OS layer | ROCKS Clusters |
| Scheduler | Open Grid Engine (also known as Sun GridEngine) |

**Table 4: Current hardware configuration of the CASPER cluster**

At this stage, i.e. before the funding assigned by the Board of Governors, all purchases have been entirely funded through fellowship agreements with research groups using the cluster for research purposes. It is worth to mention that the software to supervise the system and to implement the necessary features was entirely developed by internal members of the HPC@POLITO staff.

### *3.4 Laboratory design, user support and final results*

After identifying the necessary constraints and the minimum hardware requirements and to having the cluster ready for HPC-4-teaching, we have directed our effort to setting up the tools to support a cluster efficient and simple usage by professors and students.

In this initial phase, we have addressed this complex issue by:

- Offering technical advice and support to professors while designing their lab lectures as well as when running them
- Customizing the cluster according to the specific frameworks and needs by each course and professor
- Preparing some how-to simple guides for students explaining the basic use of the cluster
- Offering a minimal technical support to students during the lab lectures
- Offering an email-based offline support service to professors and students

In this way we have helped professors and students to meet the following goals:

- To enhance the overall quality of lab lectures
- To improve the positive perception of the course by students, especially related to the use of the computing facilities as tools in support to education, learning and research
- To maximize the efficiency and effectiveness of lab lectures

All lab lectures by the courses in Table 2 were hosted at the *Laboratorio Didattico di Informatica Avanzata "Labinf"* [www.labinf.polito.it] of the *Department of Control and Computer Engineering.* LABINF counts 70 workstations and a fully gigabit network and can host up to 95 people. The typical organization of a lab lecture is reported in Figure 2.

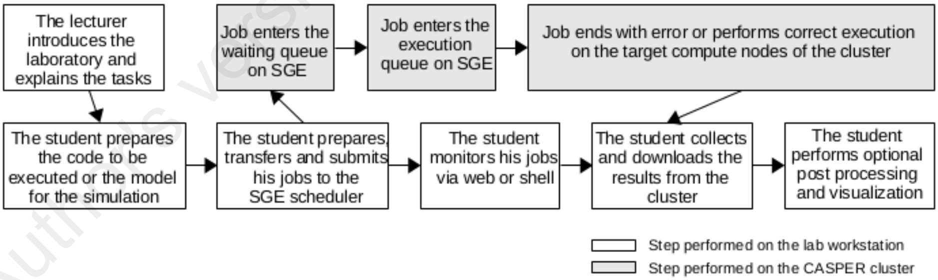


**Figure 2: The steps of a typical lab lecture**

The courses of Table 2 took place from September 2011 to June 2012. Table 5 outlines the results that have been achieved by the students of each course. We provide also a link to a set of full-scale images and short mpegs showing their works. Some smaller examples are shown in Figure 3.

| Teaching | Results |
|---|---|
| Flussi Turbolenti | High-resolution multi-scale images of velocity, pressure and vorticity flow fields |
| Fluidodinamica | Three-dimensional travelling waves perturbation in channel and wake flows |
| Realtà Virtuale | High-resolution stereoscopic images of interiors, with photo textures. |
| Parallel and Distributed Computing | Parallel solution of well known linear algebra problems using MPI |

**Table 5: Final results obtained by students at the end of the courses involved**

Some relevant results were presented by the students of the first semester courses and their advisors during the first workshop on high performance computing organized by HPC@POLITO in December 2011.

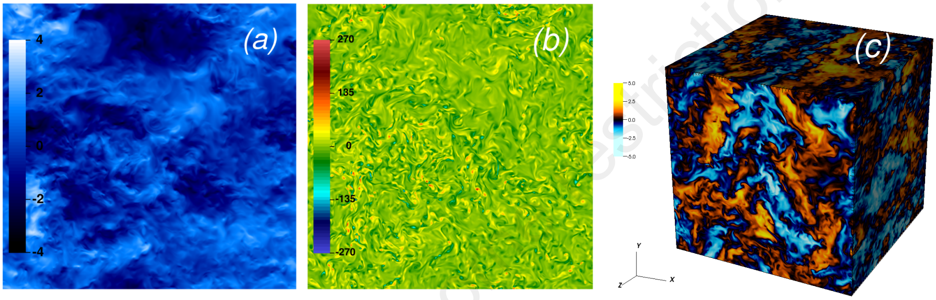


**Figure 3: Example of student work. Visualization of a high resolution turbulent field: a) velocity, b) vorticity, c) zoomed 3D velocity. Reynolds number = 50000.**

## 4. Future Works and Conclusions

Given the progress achieved in the academic year 2011-2012, the same courses are currently underway in the academic year 2012-13 with the addition of a fifth lesson related to the field of Computer Graphics. However, it is our intention for the future to improve the support of HPC-based teaching by introducing some other goals:

- Assess the satisfaction of teachers and students through questionnaires;
- Give the opportunity to other courses to take advantage of this opportunity by further expanding the CASPER cluster using the recent funding by the Board of Governors;
- Develop software tools that could ease and automate the use of the cluster by the students to shorten the laboratory start-up phase.

In this paper we have presented the results of our project HPC-4-teaching, run at Politecnico di Torino on a small number of courses over one full academic year. **The results have been encouraging and we plan to extend the project to slightly increase the number of courses**. We strongly believe that HPC should be carefully evaluated also in the light of teaching and not only research activities. One of the main reasons, besides the technical issues, is that to use HPC resources brings at least a double benefit to students involved. First, within their discipline of interest they get results to computationally intensive problems, thus improving their personal technical knowledge. Second, they can learn how to interact, even at a User level, with a complex computation system, even without having dedicated and specific expertise in computers, thus improving their overall professional capabilities in the light of a future job.

We have shown that **HPC for teaching implies a number of additional issues with respect to HPC to research only**, and we have also presented how our solutions to address most of them. We are very confident, also by direct experience that the final results are worth having played the game. Our secret dream is that, based on the good results obtained for research activities, HPC can and hopefully will take off and travel high in the skies of teaching, thanks also to the small additional funding which is necessary to extend a traditional HPC Center from research to teaching activities too.

## Acknowledgments

We are grateful to Proff. Bottino, Montrucchio and Tordella for their valuable and continuous collaboration in this first phase of the HPC-4-teaching project. We wish also to thank the Rector, the Board of Governors of Politecnico di Torino, as well as the research groups so far involved and the Chair of the Department of Control and Computer Engineering, for their support to the HPC@POLITO initiative. We also thank the students for their valuable comments/feedback and their help to identify and solve the critical “customization” issues.

## References

**[Della Croce et al. 2011]** F. Della Croce, N. Nepote, E. Picoolo, “A Terascale Cost-Effective Open Solution for Academic Computing: Early Experience of the DAUIN HPC Initiative”, Proc. of the 49th AICA national conference on smart tech and smart innovation, Turin, Italy, 2011

**[Gentzsch 2009]** W. Gentzsch “Sustainable HPC infrastructures - supercomputers, grids, and clouds: the DEISA experience”, Proc. of the 18th ACM international symposium on High performance distributed computing, New York, NY, USA, 2009

**[Gibbon et al. 2010]** A. F. Gibbon, D. A. Joiner, H. Neeman, C. Peck and S. Thompson, “Teaching High Performance Computing to Undergraduate Faculty and Undergraduate Students”. Proc. of the 2010 TeraGrid Conference, August 2010, Pittsburgh, PA, USA, 2010.

**[Glotzer et al. 2009]** S. C. Glotzer, B. Panoff, and S. Lathrop, “Challenges and opportunities in preparing students for petascale computational science and engineering”. In: Computing in Science and Engineering, vol. 11, n. 5, pp. 22-27, 2009.

**[Iovieno et al. 2001]** M. Iovieno; C. Cavazzoni; Tordella D. (2001) “A new tecnique for a parallel dealiased pseudospectral Navier-Stokes code, Comp. Phys. Comm“ In: COMPUTER PHYSICS COMMUNICATIONS, vol. 141, pp. 365-374.

**[Joiner et al. 2006]** D. A. Joiner, P. Gray, T. Murphy, and C. Peck, “Teaching parallel computing to science faculty: Best practices and common pitfalls”. Proc. of the eleventh ACM SIGPLAN Symposium on principles and practice of parallel programming, New York, NY, USA, 2006.

**[Kindratenko & Trancoso, 2011]** V. Kindratenko, P. Trancoso, “Trends in High-Performance Computing”. In: IEEE Computing in Science & Engineering, vol. 13, n. 3, pp. 92-95, Sept. 2011.

**[Scarsoglio et al. 2010]** Scarsoglio S; Tordella D.; Criminale W.O (2010) “The role of long waves in the stability of the plane wake“ In: PHYSICAL REVIEW E, STATISTICAL, NONLINEAR, AND SOFT MATTER PHYSICS, vol. 81 n. 3, 036326-1-036326-9.

**[Tordella & Iovieno 2011]** Tordella D., Iovieno M. (2011) “Small scale anisotropy in turbulent shearless mixing“ In: PHYSICAL REVIEW LETTERS, vol. 107 n. 19, 194501/1-194501/5.

**[www.hpc.polito.it]** “HPC@POLITO”, Academic Computing Center at Politecnico di Torino

**[www.labinf.polito.it]** “Laboratorio Didattico di Informatica Avanzata” at Politecnico di Torino