MHPC Workshop on High Performance Computing

Report of Contributions

https://indico.sissa.it/e/hpcts2016

Contribution ID: **0** Type: **not specified**

HPC in Europe Prace

Friday, February 26, 2016 12:15 PM (30 minutes)

Presenter: BASSINI, Sanzio (CINECA)

Session Classification: HPC in science

Contribution ID: **1** Type: **not specified**

The Quantum ESPRESSO Project

Thursday, February 25, 2016 11:15 AM (15 minutes)

Presenter: GIANNOZZI, Paolo (UNIUD)

Session Classification: HPC projects in FVG, Italy, and Europe

Contribution ID: **2** Type: **not specified**

The MAX Center of excellence

Thursday, February 25, 2016 11:30 AM (15 minutes)

Presenter: MOLINARI, Elisa (CNR NANO)

Session Classification: HPC projects in FVG, Italy, and Europe

Contribution ID: **3** Type: **not specified**

The EXANEST project

Thursday, February 25, 2016 11:45 AM (15 minutes)

Presenter: TAFFONI, Giuliano (OATS - INAF)

Session Classification: HPC projects in FVG, Italy, and Europe

Contribution ID: **4** Type: **not specified**

The INFN vision for the future of HPC and HTC from regional areas to Europe

Thursday, February 25, 2016 12:00 PM (15 minutes)

Presenter: LUCCHESI, Donatella (INFN - University of Padova)

Session Classification: HPC projects in FVG, Italy, and Europe

Contribution ID: **5** Type: **not specified**

Big Data and HPC @ Generali

Thursday, February 25, 2016 2:30 PM (15 minutes)

Generali is one of the most consolidated insurance company in Europe, looking ahead for innovative product development and new markets across the world. In order to better serve business lines as well as to identify customer valuable products, Generali created the Group Chief Data Office function, whose mission is to define and implement strategies and methods to acquire, analyze and govern data. Evaluating and adopting several data analysis modern techniques, the GCDO is addressing the Big Data challenge and will leverage HPC in dedicated researches beyond the traditional insurance modeling analysis.

Presenter: Dr BRANCHESI, Alberto (Generali spa)

Session Classification: HPC in industry: some regional examples

Contribution ID: **6** Type: **not specified**

Best thesis prize, graduation ceremony

Contribution ID: **7** Type: **not specified**

The Personal Computer is dead. Long life the Personal HPC.

Thursday, February 25, 2016 2:00 PM (30 minutes)

The exponential growth of computation is very close to an evolutionary step in the way we use HPC extending and expanding the class of problems they can address. The ongoing digital transformation and software containerization are enabling the use of HPC s in most of the fields of human activities. The new digital hyperconnected world need HPC scientists and not just only Data Scientist.

Presenter: SIAGRI, Roberto (Eurotech spa)

Session Classification: HPC in industry: some regional examples

Contribution ID: **8** Type: **not specified**

Fincantieri

Contribution ID: **9** Type: **not specified**

HPC solutions for Ship Designing @ Fincantieri

Thursday, February 25, 2016 2:45 PM (15 minutes)

Fincantieri is one of the leading ship design company in the world. In 2014 we deployed a HPC cluster that is mainly used to perform Computational Fluid Dynamics and Finite Element Analysis calculations. We equipped our cluster with a Citrix virtualization solution that allows our users to perform their pre-processing activities directly on a few reserved nodes on the cluster. CFD calculations are usually performed to optimize the hull, the appendages and the propellers. FEM calculations are performed to assess the behavior of the Ship internal structure and its response to vibrations.

Presenters: Dr GUSTIN, Gianluca (Fincantieri); Dr CHECHILE, Giuseppe (Fincantieri)

Session Classification: HPC in industry: some regional examples

Contribution ID: **10** Type: **not specified**

MHPC Thesis: A computational ecosystem for near real time processing of satelllte data

Thursday, February 25, 2016 3:00 PM (15 minutes)

The aim of this work is the development of a computational ecosystem for nearly real-time inversion of high spectral resolution infrared data coming from meteorological satellites. The ecosystem has been developed as nearly real-time demonstration project to elaborate the level 2 products derived from MTG-IRS

Primary author: PIANI, Stefano (MHPC - eXact-lab srl)

Presenter: PIANI, Stefano (MHPC - eXact-lab srl)

Session Classification: HPC in industry: some regional examples

Contribution ID: **11** Type: **not specified**

Finite Element Methods at Realistic Complexities

Thursday, February 25, 2016 3:45 PM (40 minutes)

Solving realistic, applied problems with the most modern numerical methods introduces many levels of complexity. In particular, one has to think about not just a single method, but a whole collection of algorithms: a single code may utilize fully adaptive, unstructured meshes; nonlinear, globalized solvers; algebraic multigrid and block preconditioners; and do all this on 1,000 processors or more with realistic material models.

Codes at this level of complexity can no longer be written from scratch. However, over the past decade, many high quality libraries have been developed that make writing advanced computational software simpler. In this talk, I will briefly introduce the deal.II finite element library (http://www.dealii.org) whose development I lead and show how it has enabled us to develop the ASPECT code (http://aspect.dealii.org) for simulation of thermal convection. The project also builds on a variety of other libraries (e.g., p4est, Threading Building Blocks, and Trilinos) that provide parallelism at various levels. I will discuss some of the results obtained with this code and comment on the lessons learned from developing this massively parallel code for the solution of a complex problem.

Presenter: BANGERTH, Wolfgang (Texas A&M)

Session Classification: HPC in Mathematics

Contribution ID: **12** Type: **not specified**

MHPC Thesis: Hybrid Parallelisation Strategies for Boundary Element Methods

Thursday, February 25, 2016 4:25 PM (20 minutes)

Whenever a mathematical problem admits a boundary integral representation, it can be straightforwardly discretised by Boundary Element Methods (BEM). In this work, we present an efficient hybrid parallel solver for FSI problems based on collocation BEM.

The major bottlenecks for a serial implementations of BEM is the computational cost and memory requirements needed to respectively assemble and store the BEM full matrices. Both memory storage and assembling CPU times scale with the square of the number of degrees of freedom. We present two different strategies to parallelise BEM implementations. The first uses an MPI strategy, in which we distribute both assemblage workload and storage requirement among different processors, maintaining the classical BEM structure (and algorithm complexity). This approach leads to optimal strong and weak scalability for the matrix assemble cycles and vector matrix multiplication, although the overall algorithm remains of order O(N^2). In the second strategy, we employ a Fast Multipole Method (FMM) to reduce the computational cost and memory allocation of the BEM problem resolution to O(N), and we use a hybrid MPI and multi-threaded parallelization strategy.

This implementation combines direct BEM close range interactions with FMM long range couplings, and represents the state of the art in parallel BEM solvers. The BEM-FMM algortihm calls for a hybrid solution, since the algorithm requires inherently a lot of communication among different processors. We address the main parallelisation techniques to be used in a hybrid parallel BEM-FMM implementation, for which we used the Intel Threaded Building Block paradigm to handle multicore platform, and MPI for the communication between different processors. We present strong and weak scalability results together with an optimality result concerning the way to proper set the hierarchical FMM space subdivision.

Primary author: GIULIANI, Nicola

Presenter: Mr GIULIANI, Nicola (MHPC - SISSA)

Session Classification: HPC in Mathematics

Contribution ID: **13** Type: **not specified**

High performance computing for computational electrocardiology. Part II: scalable solvers.

Thursday, February 25, 2016 5:15 PM (30 minutes)

The complex interaction between the cardiac bioelectrical and mechanical phenomena is modeled by a system of non-linear partial differential equations (PDEs), known as cardiac electromechanical coupling (EMC) model. Due to the extremely different spatial and temporal scales of the physical phenomena occurring during a single heartbeat, the discretization of the EMC model with finite elements in space and finite differences in time yields the solution of thousands of large scale linear systems, with $O(10^6 - 10^8)$ degrees of freedom each. The effective solution of such linear systems requires the use of hundreds/thousands processors and, consequently, of highly scalable preconditioners. In this presentation, we will first introduce two classes of Domain Decomposition preconditioners, the Multilevel Additive Schwarz (MAS) and the Balancing Domain Decomposition by Constraints (BDDC) preconditioners, in the simple setting a scalar elliptic PDE. Then, we will extend such preconditioners to the solution of the reaction-diffusion PDEs and of the non-linear elasticity system constituting the EMC model. Finally, the results of three-dimensional parallel simulations will demonstrate the effectiveness of the resulting algorithms.

Presenter: SCACCHI, Simone (UNIMI)

Session Classification: HPC in Mathematics

Contribution ID: **14** Type: **not specified**

High performance computing for computational electrocardiology. Part I: motivation and mathematical models.

Thursday, February 25, 2016 4:45 PM (30 minutes)

Life sciences could benefit immensely from the massive growth of HPC processing power occurred in the last ten years. Indeed, complex biological systems are described by sophisticated mathematical models, whose solution requires highly scalable solvers. In particular, for what concerns cardiac electrophysiology, the simulation of the electrical excitation of the heart muscle, and the subsequent contraction-relaxation process, represents a challenging computational task. In the present talk, we will describe the main mathematical model of the cardiac electrical and mechanical interactions, the so-called cardiac electro mechanical coupling model. This model consists of a system of non-linear partial differential equations (PDEs), constituted by four sub-models: the quasi-static anisotropic finite elasticity equations describing the macroscopic deformation of the cardiac tissue; the active tension model, i.e. a system of ordinary differential equations (ODEs) describing the intracellular calcium dynamics and the consequent generation of the cellular force; the anisotropic Bidomain model, i.e. a system of degenerate parabolic reaction-diffusion PDEs describing the electrical current flow through the tissue; the membrane model, i.e. a stiff system of ODEs describing the bioelectrical activity of the membrane of cardiac cells. We will finally present the results of three-dimensional simulations of the full cardiac excitation-contraction process.

Presenter: COLLI FRANZONE, Piero (UNIPV)

Session Classification: HPC in Mathematics

Contribution ID: **15** Type: **not specified**

Here and now: the intersection of computational science and computer science

Friday, February 26, 2016 8:45 AM (40 minutes)

Quantum-mechanical simulations have become dominant and widely used tools for scientific discovery and technological advancement; since they are performed without any experimental input or parameter they can streamline, accelerate, or replace actual physical experiments. This is a far-reaching paradigm shift, substituting the cost- and time-scaling of brick-and-mortar facilities, equipment, and personnel with those, very different, of computing engines.

Nevertheless, computational science remains anchored to a renaissance model of individual artisans gathered in a workshop, under the guidance of an established practitioner. Great benefits could follow from rethinking such model, while adopting concepts and tools from computer science for the automation, management, preservation, analytics, and dissemination of these computational efforts.

I will offer my perspective on the current state-of-the-art in the field, its power and limitations, and the role and opportunities of high-throughput computing (HTC, rather than HPC), of open source codes and workflows, and of big data available on demand.

Presenter: MARZARI, Nicola (EPFL)

Session Classification: HPC in science

Contribution ID: **16** Type: **not specified**

High Performance Computing and Materials Science: How atomistic simulations can pave the way for clean and sustainable energy

Friday, February 26, 2016 9:25 AM (20 minutes)

The availability of cheap and abundant energy was one of the main drivers of the industrial revolution. Until today, energy remains an essential ingredient for many aspects of human activity. Is is recognized that a major challenge of our times is the transition towards sustainable energy conversion, moving away from carbon-based fossil fuels. Developing more efficient and cheaper ways to convert wind or solar radiation into electricity or to store electric energy are important steps in this transition.

Computer simulations at the atomic scale can lead to a detailed understanding of the fundamental steps during energy conversion. In this presentation, I will illustrate a few cases where such a "computational microscope" can be used by materials scientists to develop better solar cells or to more efficiently use solar light to split water into hydrogen and oxygen. In these cases, high performance computing allows for a screening of potential materials, before they have even been synthesized in a laboratory.

Presenter: GEBAUER, Ralph (ICTP)

Session Classification: HPC in science

Contribution ID: **17** Type: **not specified**

MHPC Thesis: Analysis of Hybrid Parallelization strategies: Simulation of Anderson Localization

Friday, February 26, 2016 3:35 PM (20 minutes)

This thesis presents two experiences of hybrid programming applied to condensed matter and high energy physics. The two projects differ in various aspects, but both of them aim to analyse the benefits of using accelerated hardware to speedup the calculations in current science-research scenarios.

The first project enables massively parallelism in a simulation of the Anderson localisation phenomenon in a disordered quantum system. The code represents a Hamiltonian in momentum space, then it executes a diagonalization of the corresponding matrix using linear algebra libraries, and finally it analyses the energy-levels spacing statistics averaged over several realisations of the disorder.

The implementation combines different parallelization approaches in an hybrid scheme. The averaging over the ensemble of disorder realisations exploits massively parallelism with a master-slave config- uration based on both multi-threading and message passing interface (MPI). This framework is designed and implemented to easily interface similar application commonly adopted in scientific research, for ex- ample in Monte Carlo simulations. The diagonalization uses multi-core and GPU hardware interfacing with MAGMA, PLASMA or MKL libraries. The access to the libraries is modular to guarantee portability, maintainability and the extension in a near future.

The second project is the development of a Kalman Filter, including the porting on GPU architectures and autovectorization for online LHCb triggers. The developed codes provide information about the viability and advantages for the application of GPU technologies in the first triggering step for Large Hadron Collider beauty experiment (LHCb).

The optimisation introduced on both codes for CPU and GPU delivered a relevant speedup on the Kalman Filter. The two GPU versions in CUD and OpenCL have similar performances and are adequate to be considered in the upgrade and in the corresponding implementations of the Gaudi framework.

In both projects we implement optimisation techniques in the CPU code. This report presents exten- sive benchmark analyses of the correctness and of the performances for both projects.

Presenter: AGUILAR MENA, Jimmy (MHPC - ICTP)

Session Classification: HPC in science

Contribution ID: **18** Type: **not specified**

Simulating Cosmic Structure Formation

Friday, February 26, 2016 10:45 AM (30 minutes)

Numerical simulations on supercomputers play an ever more important role in astrophysics. They have become the tool of choice to predict the non-linear outcome of the initial conditions left behind by the Big Bang, providing crucial tests of cosmological theories. However, the problem of galaxy and star formation confronts us with a staggering multi-physics complexity and an enormous dynamic range that severely challenges existing numerical methods. In my talk, I review current strategies to address these problems, focusing on recent developments in the field such as hierarchical time integration schemes, improved particle- and mesh-based hydrodynamical solvers. I will also discuss a selection of current results and highlight some challenges for the future.

Presenter: SPRINGEL, Volker (The Heidelberg Institute for Theoretical Studies)

Session Classification: HPC in science

Contribution ID: **19** Type: **not specified**

Numerical simulations of galaxies and galaxy clusters at the Trieste Observatory

Friday, February 26, 2016 11:15 AM (20 minutes)

Presenter: MURANTE, Giuseppe (OATS-INAF)

Session Classification: HPC in science

Contribution ID: **20** Type: **not specified**

HPC for Earth Sciences: training opportunities and research challenges

Friday, February 26, 2016 11:35 AM (20 minutes)

Presenter: SALON, Stefano (OGS)

Session Classification: HPC in science

Contribution ID: **21** Type: **not specified**

MHPC Thesis: Shyfem Parallelization: An innovative task approach for coastal environment FEM software

Friday, February 26, 2016 11:55 AM (20 minutes)

SHYFEM is a finite element hydrodynamic code written by Georg Umgiesser in the 80s to model Venice lagoon for his master thesis; its development has been continued by CNR-ISMAR group. It is one of the few opensource codes for coastal areas that use a finite element approach. SHYFEM is a very important resource because it is focused on coastal areas and can be coupled it with other software in order to increase the simulation accuracy in such areas. Coastal areas are strategic because many human activities are here concentrated. This means that a software that produces an accurate representation in coastal areas may also advantage socio-economical activities. SHYFEM has been already and successfully applied to several coastal and lagoon environments; for example, it is used to produce tidal forecasts in the Venice lagoon and other lagoons in Mediterranean sea. It is also used in the Danube Delta and to estimate its effects on the Black Sea, and in Malta to produce coastal forecasts. The main goal of this work is to obtain a new version of SHYFEM that may be faster, parallel, capable to use efficiently modern hardware, and easily coupled with other software.

Presenter: PASCOLO, Eric (MHPC - OGS)

Session Classification: HPC in science

Contribution ID: **22** Type: **not specified**

The path toward High Performance Computing in High Energy Physics

Friday, February 26, 2016 2:15 PM (40 minutes)

Presenter: CARMINATI, Federico (CERN)

Session Classification: HPC in science

Contribution ID: **23** Type: **not specified**

Studies of Flavor and Hadron Physics using Lattice QCD simulations with modern HPC hardware

Friday, February 26, 2016 2:55 PM (20 minutes)

Presenter: SIMULA, Silvano (INFN - Roma3)

Session Classification: HPC in science

Contribution ID: **24** Type: **not specified**

High Performance Computing in the ALICE experiment

Friday, February 26, 2016 3:15 PM (20 minutes)

Presenter: MASERA, Massimo (Università di Torino)

Session Classification: HPC in science

Contribution ID: **25** Type: **not specified**

Improving Performance of Basis-set-free Hartree-Fock Calculations Through Grid-based Massively Parallel Techniques

Friday, February 26, 2016 10:00 AM (15 minutes)

Multicenter numerical integration scheme for polyatomic molecules has been implemented as an initial step to develop a complete basis-set-free Hartree-Fock (HF) software. The validation of the integration scheme includes the integration of the total density and the calculation of Coulomb potentials for several diatomic molecules. A finite difference method is used to solve Poisson's equation for the Coulomb potential on numerical orbitals expanded on the interlocking multicenter quadrature grid. The implementation which rely on OpenMP and CUDA shows a speedup up to 30x.

Presenter: Mr POSADA, Fernando (MHPC)

Session Classification: HPC in science

Contribution ID: **26** Type: **not specified**

Introduction by MHPC

Contribution ID: **27** Type: **not specified**

Welcome address by ICTP Director

Thursday, February 25, 2016 9:05 AM (5 minutes)

Session Classification: Welcome and general introduction

Contribution ID: **28** Type: **not specified**

Welcome address by SISSA Director

Thursday, February 25, 2016 9:00 AM (5 minutes)

Session Classification: Welcome and general introduction

Contribution ID: **29** Type: **not specified**

Welcome address by Governmental Institutions

Thursday, February 25, 2016 9:10 AM (15 minutes)

Session Classification: Welcome and general introduction

Contribution ID: **30** Type: **not specified**

Perspectives of HPC in FVG

Primary author: SCANDOLO, Sandro (ICTP)

Presenter: SCANDOLO, Sandro (ICTP)

Contribution ID: **31** Type: **not specified**

Perspectives of HPC in FVG

Presenter: SCANDOLO, Sandro (ICTP)

Contribution ID: **32** Type: **not specified**

MHPC: first achievements and future perspectives

Thursday, February 25, 2016 9:25 AM (20 minutes)

Presenter: COZZINI, Stefano (CNR/IOM - MHPC)

Session Classification: MHPC and HPC trends

Contribution ID: **33** Type: **not specified**

Architecture-aware Algorithms and Software for Peta and Exascale Computing

Thursday, February 25, 2016 9:45 AM (45 minutes)

In this talk we examine how high performance computing has changed over the last 10-year and look toward the future in terms of trends. These changes have had and will continue to have a major impact on our software. Some of the software and algorithm challenges have already been encountered, such as management of communication and memory hierarchies through a combination of compile–time and run–time techniques, but the increased scale of computation, depth of memory hierarchies, range of latencies, and increased run–time environment variability will make these problems much harder.

We will look at five areas of research that will have an importance impact in the development of software and algorithms.

We will focus on following themes:

- Redesign of software to fit multicore and hybrid architectures
- Automatically tuned application software
- Exploiting mixed precision for performance
- The importance of fault tolerance
- Communication avoiding algorithms

Presenter: DONGARRA, Jack (University of Tennessee)

Session Classification: MHPC and HPC trends

Contribution ID: **36** Type: **not specified**

QE, main strategies of parallelization and levels of parallelisms

Wednesday, February 24, 2016 2:15 PM (45 minutes)

Presenter: AFFINITO, Fabio

Session Classification: Tutorial: HPC aspects of Quantum-Espresso package

Contribution ID: **37** Type: **not specified**

QE, methodologies to develop/maintain/testing toward code modernization and code sustainability

Wednesday, February 24, 2016 3:00 PM (45 minutes)

Presenter: SPIGA, Filippo (QE Foundation - University of Cambridge)

Session Classification: Tutorial: HPC aspects of Quantum-Espresso package

Contribution ID: **38** Type: **not specified**

QE and many-core architectures

Wednesday, February 24, 2016 4:15 PM (45 minutes)

Presenter: AFFINITO, Fabio

Session Classification: Tutorial: HPC aspects of Quantum-Espresso package

Contribution ID: **39** Type: **not specified**

QE and heterogeneous architectures

Wednesday, February 24, 2016 5:00 PM (45 minutes)

Presenter: SPIGA, Filippo (QE Foundation - University of Cambridge)

Session Classification: Tutorial: HPC aspects of Quantum-Espresso package

Contribution ID: **40** Type: **not specified**

MHPC thesis: High-performance implementation of the Density Peak clustering algorithm

Friday, February 26, 2016 9:45 AM (15 minutes)

We developed a parallel implementation of the "Density Peak"clustering algorithm, exploiting C++11, OpenMP and the FLANN library for k-nearest-neighbour search. The modified algorithm is approximately 50 times faster than the original version on datasets with half a million points, and scales almost linearly with the dataset size. Thanks to improvements on the density estimation and assignation procedure, the algorithm is also unsupervised and non-parametric.

Presenter: BORELLI, Marco (MHPC - SISSA)

Session Classification: HPC in science

Contribution ID: **41** Type: **not specified**

Perspectives of HPC in FVG

Thursday, February 25, 2016 11:00 AM (15 minutes)

Presenter: SCANDOLO, Sandro (ICTP)

Session Classification: HPC projects in FVG, Italy, and Europe