Book of Abstracts

HPCSIM2024

Frontiers of High-Performance Computing in Modeling and Simulation

Politecnico di Milano

(3)-4-5 July 2024

Scientific Committee

Luca Formaggia (Politecnico di Milano)
Paolo Cremonesi (Politecnico di Milano)
Carlo de Falco (Politecnico di Milano)
Ivan Fumagalli (Politecnico di Milano)
Ilario Mazzieri (Politecnico di Milano)
Nicola Parolini (Politecnico di Milano)
Carlo Janna (Università di Padova)
Andrea Franceschini (Università di Padova)
Nicolò Spiezia (M3E)

The event HPCSIM2024 has been partially supported by ICSC—Centro Nazionale di Ricerca in High Performance Computing, Big Data, and Quantum Computing funded by European Union - NextGenerationEU and by MUR, DIpartimento di Eccellenza 2023-27.

Sponsors

Platinum Sponsors

Gold Sponsors

Program

Wednesday 3 July

14:00 - 18:00	Daniele Ottaviani (Cineca)	Tutorial on Quantum Computing

Thursday 4 July

8:30 - 9:00	Registration	
9:00 - 9:15	Welcome address	
9:15 - 10:05	Keynote lecture: Carsten Burstedde	Adaptive Mesh Refinement in HPC and Applications in the Geosciences
10:05 - 10:25	Coffee break	
10:25 - 10:50	Pasqua D'Ambra (IAC-CNR, Napoli)	Scalable Linear Solvers for GPU-Accelerated Heterogeneous Supercomputers
10:50 - 11:15	Eric Sonnendrücker (TUM & IPP Garching)	Development of a Geometric Particle-In-Cell (PIC) Code for HPC
11:15 - 11:40	Mattia Tessari (Leonardo S.p.A.)	Open Innovation: Harnessing the Power of HPC Clusters with Open Source Solutions
11:40 - 12:05	Stefano Berrone (Politecnico di Torino)	Numerical Algorithms for Parallel Discrete Fracture Network and Discrete Fracture Matrix Flow Simulations
12:05 - 14:00	Lunch break	
14:00 - 14:50	Keynote lecture: Victorita Dolean Maini	Fast Solution Methods for Wave Propagation Problems: From Classical Domain Decomposition Solvers to Learning
14:50 - 15:10	Coffee Break	
15:10 - 15:35	Michael Bader (TUM, Munich)	Model Flexibility and Performance Portability for Extreme-Scale Earthquake Simulations with SeisSol
15:35 - 16:00	Daniele Gregori (E4)	Boosting Scientific Applications Performance on Novel Computing Architectures
16:00 - 16:25	Antonio Cervone (Università di Bologna)	Development of a Numerical Platform for Code Coupling
16:25 - 16:50	Nicolas Tardieu (EDF)	`code_aster`: a 35-year Journey from Legacy to Massively Parallel Open-Source Finite Element Software

Friday 5 July

9:15 - 10:05	Keynote lecture: Simone Severini	Questions about Quantum Computers
10:05 - 10:25	Coffee break	
10:25 - 10:50	Pasquale Africa (SISSA, Trieste)	Efficient Implementation of Polygonal Discontinuous Galerkin Methods
10:50 - 11:15	Alessia Marruzzo (ENI)	Quantum Computing in the Energy Sector
11:15 - 11:40	Carlo Cavallotti (Politecnico di Milano)	On the Recent Evolution of Theoretical Chemical Kinetics from a Qualitative to a Quantitative Science
11:40 - 12:05	Damiano Casalino (Dassault Systemes)	LBM/VLES Aeroacoustics: New Challenges in the Age of Electric Vertical Take-Off and Landing (eVTOL) Vehicle
12:05 - 14:00	Lunch break	
14:00 - 14:25	Walter Rocchia (IIT)	Continuum Electrostatics Simulations for Biological Applications: an HPC
14:25 - 14:50	Marco Giardino (ASI)	A Cloud Based Bayesian Approach to Calculate the Chemical Composition of the Jupiter Atmosphere
14:50 - 15:15	Nur Fadel (CSCS, Lugano)	Building a Platform for Advanced Scientific Computing at CSCS: Engineering Software for the future of Alps
15:15 - 15:40	Carlo Nardone (NVIDIA)	Accelerate Time-to-Science with the NVIDIA Platform
15:40 - 16:05	Final greetings	
16:05 - 16:35	Farewell cocktail	

Keynote speakers

Adaptive Mesh Refinement in HPC and Applications in the Geosciences

Thursday
4 Jul
9:15

Carsten Burstedde

University of Bonn

Whenever a numerical simulation exhibits disparate scales and localized, possibly moving features, the question arises whether the mesh should be updating accordingly. The hope would be that refining the mesh only where needed, coarsening elsewhere, and repartitioning frequently can help optimize the ratio of numerical error to numerical effort as well as the wall time to solution. Moreover, modern simulations venture increasingly into the realms of computational geometry and object and data location. While practical tasks such as intersecting a distributed set of source/sink or observation objects with the mesh are near trivial on uniform meshes, executing them efficiently on arbitrarily partitioned adaptive meshes is far from obvious and only recently being tackled and scaled.
Generally, adaptivity requires a certain investment, namely to enhance the discretization and solver to accommodate the mesh topology, and to execute mesh refinement and search using scalable algorithms. While the first task may seem daunting at first, necessary modifications to the code are mostly local and well isolated. The second task is often too inolved to be addressed by a single simulation scientist, but can be delegated to freely available software libraries. Successfully adopting the conventions imposed by that software is, ideally, rewarded by inheriting its performance.
In this talk, we discuss the algorithm design fundamental to the p4est software for scalable adaptive mesh management, and outline how to benefit from interfacing to it. p4est is being used in many projects, including geoscientific applications, and we close with selected examples.

Fast Solution Methods for Wave Propagation Problems: From Classical Domain Decomposition Solvers to Learning

Thursday
4 Jul
14:00

Victorita Dolean Maini

TU Eindhoven

Wave propagation and scattering problems are of huge importance in many applications in science and engineering - e.g., in seismic and medical imaging and more generally in acoustics and electromagnetics. Large-scale simulations of those applications are one of the hard problems from a computational point of view since requires an interplay between the parsimonious but sufficiently accurate discretisation methods and more sophisticated solution methods. Our aim is to show on one side, how classical domain decomposition methods developed in the latest years coupled with carefully chosen discretisations can help in this endeavour. On the other side, we would like to propose some openings towards the very attractive package of approximation-solution-optimisation offered by the new methods in scientific machine learning.

Friday
5 Jul
9:15

Questions about Quantum Computers

Simone Severini

Amazon Web Services

What is quantum computing? Why is it important? What is not quantum computing? How did we get here? Where are we now? What is the impact today? What are the challenges and risks? What is the outlook? What can industry leaders do now?

Invited speakers

Scalable Linear Solvers for GPU-Accelerated Heterogeneous Supercomputers

Thursday
4 July
10:25

Pasqua D'Ambra

IAC-CNR, Napoli

Solution of large and sparse linear systems is at the core of physics-driven numerical simulations relying on PDEs and often represents a bottleneck in machine learning. Scalability needs and complexity of current supercomputers impose to find good trade-off between optimality of numerical methods and parallelism. I will present new methods and scientific libraries for iterative linear solvers on GPU-accelerated platforms.

Development of a Geometric Particle-In-Cell (PIC) Code for HPC

Thursday
4 July
10:50

Eric Sonnendrücker

TUM Münich and IPP Garching

We have been developing a geometric PIC code for solving the Vlasov-Maxwell system and related models. It the important feature of being able to preserve exactly the main invariant of the system, like energy and the divergence constraints.

Open Innovation: Harnessing the Power of HPC Clusters with Open Source Solutions

Thursday
4 July
11:15

Mattia Tessari

Leonardo S.p.A.

As kinetic simulations of magnetic fusion plasmas in particular are particularly time consuming care needs to be taken so that they can run efficiently on evolving modern computer architectures. For this reason, the new code has been build on top of the AMReX framework, which provides optimized and performance portable data structures, that can run on different architectures by just changing some compiler flags.

Thursday
4 July
11:40

Numerical Algorithms for Parallel Discrete Fracture Network and Discrete Fracture Matrix Flow Simulations

Stefano Berrone

Politecnico di Torino

In the framework of underground exploitation procedures, a large number of numerical simulations at the scale of a geological basin are typically necessary. To properly account for the strong and highly variable directionality of underground flows due to the presence of fractures, an explicit representation of the rock fractures crossing the basin is required, as fractures provide preferential flow pathways. This is achieved, for example, by utilizing Discrete Fracture Network (DFN) models.
A new optimization approach for flow problems in large-scale DFNs has recently been presented. The approach is based on a reformulation of the problem as the minimization of a functional measuring the hydraulic head mismatch and the flux imbalance at fracture intersections, constrained by the flow equations on the fractures. This optimization approach can be extended to the coupling of DFN with the surrounding rock matrix. The minimization is performed via the conjugate gradient method and is suitably developed to be efficiently parallelized using several approaches and/or different parallel architectures. Scalability results will be presented for a distributed memory system.
The same problem has been tackled using a traditional conforming Galerkin method based on polygonal and polyhedral elements. In this framework, several partitioning approaches have been tested, and a parallel adaptive mesh generator is under development.

Thursday
4 July
15:10

Model Flexibility and Performance Portability for Extreme-Scale Earthquake Simulations with SeisSol

Michael Bader

TUM Münich

The earthquake simulation software SeisSol (www.seissol.org) implements high-order discontinuous Galerkin schemes for various seismic wave propagation models and has been optimised for CPU and GPU architectures. We report how performance portability and model flexibility are addressed and show performance results for various supercomputers and scenarios, such as seismo-acoustic nuisance patterns from induced earthquakes, coupled tsunami-earthquake events or dynamic rupture on complex fault geometries.

Thursday
4 July
15:35

Boosting Scientific Applications Performance on Novel Computing Architectures

Daniele Gregori

E4

E4 Computer Engineering is an Italian company, leader in the HPC sector, always focused on new technologies and able to provide emerging computing platforms. We are members of several European research projects in HPC and Quantum Computing, to which we provide, in addition to our expertise in design, access to computing clusters (some prototypes), present in our datacenter. In this presentation we will show some comparative tests performed on different architectures such as: x86, ARM, RISC-V. The tests performed are both synthetic and application-oriented, particularly in the areas of research in which we are involved as members of the centers of excellence: SPACE, MAX and EoCoE. The aim is to provide an overview of the state of the art of the different architectures and try to intercept possible trends in their evolution.

Development of a Numerical Platform for Code Coupling

Thursday
4 July
16:00

Antonio Cervone

Università di Bologna

The talk will present relevant implementation details in developing a framework to couple open-source numerical codes and tackle multi-scale and multi-physics simulations. The coupling interface is based on the MED library for robustness, performance, and ease of extensibility. The platform is tested with an in-house finite element code coupled with the finite volume OpenFOAM package.

---

code_aster: a 35-year Journey from Legacy to Massively Parallel Open-Source Finite Element Software

Thursday
4 July
16:25

Nicolas Tardieu

EDF

code_aster is a legacy code developed since 1989 by Electricité de France (EDF), released under the GNU GPL licence in 2001. Over the last years, we have tackled the challenge of running it in a massively parallel framework. We are now able to run multi-physics, multi-finite elements including multipoint constraints in a fully parallel workflow, from mesh reading to result writing.

Efficient Implementation of Polygonal Discontinuous Galerkin Methods

Friday
5 July
10:25

Pasquale Africa

SISSA, Trieste

This talk aims to investigate the parallel implementation of polygonal Discontinuous Galerkin methods. While traditional Finite Element Methods (FEMs) struggle with complex geometries due to limitations in grid generation, polygonal elements offer flexibility, reducing degrees of freedom and improving efficiency. We introduce polyDEAL, which extends the open-source FEM library deal.II to support these elements, ensuring efficient simulations and parallel scalability.

Quantum Computing in the Energy Sector

Friday
5 July
10:50

Alessia Marruzzo

ENI

The interest in the potentialities of quantum computers firstly started in the field of quantum mechanics: the exponential growth of the state space would allow efficient simulations of quantum systems. For certain problems, quantum algorithms have theoretically proven to provide exponential speed up over their classical counterpart. Starting with these results, we have been trying to understand which problems across our entire value chain could be approached with a quantum algorithm.

Friday
5 July
11:15

On the Recent Evolution of Theoretical Chemical Kinetics from a Qualitative to a Quantitative Science

Carlo Cavallotti

Politecnico di Milano

Driven by synergic advancements in high performance computing and theory, the capability to estimate rate constants for gas phase reactions from first principles without the need of any empirical correction has evolved considerably in the recent years. In this presentation the breakthroughs leading to this success will be reviewed and the current theoretical challenges discussed.
Friday
5 July
11:40

Friday
5 July
14:00

Continuum Electrostatics Simulations for Biological Applications: an HPC Perspective

Walter Rocchia

IIT

Electrostatics plays a fundamental role in many fields of Physics, Chemistry and Biology and it is instrumental in regulating a huge number of mechanisms at diFerent scales. In computational biophysics, for instance, estimating the electrostatic component of free energy of some process can provide key information to explain the occurrence of recognition and binding processes. On the other side, in the exa-scale computing era, the demand for precise description and simulation of increasingly complex systems presents formidable challenges. Despite the burgeoning computational power of High Performance Computing (HPC) architectures, existing methodologies often fall short in harnessing their full potential.
I will present two applications that try to combine advanced models with HPC implementation, leading to promising results on highly challenging biological systems. The first application is NanoShaper, a software for building and analysing the molecular surface of atomistic systems at the nanoscale, while the second is NextGenPB, a Poisson-Boltzmann equation solver where analytical calculations are leveraged to increase the accuracy/computational cost trade oF. NextGenPB is based on the finite element method on hierarchically refinedCartesian grids. By suitably adopting analytical results and exploiting local (de)refinement of the grid, we manage to increase accuracy and eFiciency in computing electrostatic potential and energy, crucial for characterizing diverse protonation state populations in proteins.
Specific applications and preliminary results are presented, discussing future developments and persistent challenges.

Friday
5 July
14:25

A Cloud Based Bayesian Approach to Calculate the Chemical Composition of the Jupiter Atmosphere

Marco Giardino

ASI

A recurring need for science research is the need to re-engineer native algorithms, preserving the scientist contribution while adopting HPC techniques. We present a proof of concept about a retrieval algorithm developed in Fortran language by INAF, engineered by ASI to exploit the HPC features of Google Cloud Platform, overtaking the initial limitations mainly due to hardware platform and non-portability.

Building a Platform for Advanced Scientific Computing at CSCS: Engineering Software for the future of Alps

Friday
5 July
14:50

Nur Fadel

CSCS, Lugano

We are facing a revolution that affects several areas of the HPC (High Performance Computing) world in science, ranging from the hardware to the way a supercomputer works to the way scientists use it. These innovations bring with them new challenges, new risks, but above all great opportunities for similar machines in the coming years. I will briefly introduce Alps, the new technological solution that CSCS will publicly present in September. Alps will be CSCS's new flagship machine, destined to revolutionise the world of scientific HPC computing thanks to a series of new technologies. I will present the new hardware and the new user environment, designed to be more flexible than in the past. In addition, I will discuss new challenges, such as energy consumption, and ongoing projects to enable researchers to make the most of the new machine while minimising the effort required to fully utilise it.

Accelerate Time-to-Science with the NVIDIA Platform

Friday
5 July
15:15

Carlo Nardone

Nvidia

Accelerated computing is nowadays, de-facto, accepted as the path forward to deploy large-scale energy-efficient scientific and technical computing (including Exascale). Today there are multiple valuable choices in programming model to accelerate any codebase to leverage GPU. Current GPU system design often adopt dense node, with multi-socket and multi-GPU connected to each other via PCIe links or custom high-speed buses (NVLink). Trends in simulation technology and artificial intelligence are evolving at incredible pace, we developed the GH200 Grace Hopper Superchip product to address the current and future computational challenges. The purpose of this talk is to briefly introduce the NVIDIA platform, hardware and software, showcasing few examples of engineering applications that have been successfully accelerated using NVIDIA GPU and set the stage for the future in computing which involve classic HPC simulations coupled or argument by AI methods.

Author Index