The Abdus Salam International Centre for Theoretical Physics

SISSA

2019-2020

MHPC

Master in High Performance Computing

TRAINING THE EXCELLENCE

The Master in High Performance Computing (MHPC) is hosted and organized by **SISSA** (International School for Advanced Studies) and **ICTP** (Abdus Salam International Center for Theoretical Physics) in Trieste.

MHPC is an **innovative specialization program** devoted to training students in the booming field of HPC.

Students that complete the Master have a solid background in **advanced and parallel computing approaches**, **algorithms**, and **machine learning techniques**.

The program combines lectures with hands-on tutorials, that are held by **internationally renowned scientists**.

MHPC
Master in High Performance Computing

2019-2020

THE PROGRAM

During the **12-months-program**, students follow courses by international **lecturers from academia and from the industrial world**, who teach them the latest applications of HPC technologies and innovations.

FIRST PART
INTRODUCTION: HPC & PROGRAMMING
6 MONTHS

SECOND PART
HPC ALGORITHMS FOR SCIENCE&TECH
3 MONTHS

THIRD PART
THESIS PROJECT
6-9 MONTHS

Courses should include **practical and exercise-based sessions** and the majority of the first part ones spawns several months.
Some courses of the second part could be optional, but a minimun of CFUs is needed to complete the program.

Master in High Performance Computing

2019-2020

from September to June

COURSES

Part I - from September to February

1.1 Scientific Programming Environment*
1.2 Foundation of HPC*
1.3 Advanced Programming*
1.4 Numerical Analysis*
1.5 Parallel Programming*
1.6 Scientific Data Management*
1.7 Advanced linear algebra libraries*
1.8 High-Performance Computing Technologies*
1.9 Best Programming Practices*

Part II - from February to June

2.1 Data Structures and Searching and Sorting Algorithms*
2.2 Unsupervised Machine Learning
2.3 The Finite Element Method Using deal.II
2.4 Spatial Locality Algorithms
2.5 Reduced Order Modelling
2.6 Electronic structure: from blackboard to source code
2.7 Approximation and Interpolation
2.8 Fast Fourier Transforms
2.9 Supervised Machine Learning
2.10 Monte Carlo method
2.11 Reinforcement Learning
2.12 Advanced optimization techniques
2.13 Deep Learning
2.14 Deep Generative Models with Tensorflow 2
2.15 Molecular Dynamics

* mandatory

Master in High Performance Computing

2019-2020

from September to February

FIRST PART

Scientific Programming Environment

Stefano Cozzini - Nicolas Salles - Alberto Sartori

This course will introduce Unix-like operating systems, show how to setup the scientific programming environment in such operating systems. It will present the modern software tools required to provide such an environment and discuss important points like documentation and testing.

Foundation of HPC

Stefano Cozzini - Luca Tornatore

Introduction to key topics in computer architecture needed in HPC environment including a detailed overview on parallel architectures.

Advanced Programming

Alberto Sartori

Provide advanced knowledge of both theoretical and practical programming in C++ and Python, with regard to the principles of object-oriented programming and best practices of software development (advanced use of version control systems, continuous integration, unit testing).

MHPC
Master in High Performance Computing

2019-2020

from September to February

FIRST PART

Numerical Analysis

Luca Heltai - Gianluigi Rozza

Introduction to numerical analysis, with focus on linear algebra, polynomial approximation, numerical integration and numerical solution of ODEs.

Parallel Programming

Ivan Girotto

Introduction to key topics in parallel computing. Main parallel programming paradigms: message passing (MPI) and multi-threading (OpenMP).

Scientific Data Management

Stefano Cozzini

The module introduces modern techniques to deal with the large amount of data in scientific and technical computing.

from September to February

FIRST PART

Advanced linear algebra libraries

Piotr Luszczek

Large supercomputing installations and scientific clusters will be discussed with the emphasis on their architectural features that are essential for good performance and scalability. Although majority of the presented code will be explained, familiarity with programming is welcome and will be helpful in following along.

High-Performance Computing Technologies

Richard Berger - Fernando Posada

This module introduces state-of-the-art technologies and innovation in High Performance Computing. Main components of computing infrastructure are analyzed and discussed. Students will install and configure a HPC Linux Cluster and will also be exposed to the use of Cloud and Grid Infrastructures.

Best Programming Practices

Axel Kohlmeyer

A module where students are introduced to best practices in scientific computing from different perspective: software development with modern software engineering techniques, optimal exploitation of different HPC platforms, usage and maintenance of large scientific software packages..

MHPC
Master in High Performance Computing

2019-2020

from February to June

SECOND PART

Data Structures and Searching and Sorting Algorithms

Axel Kohlmeyer

Introduction to fundamental data structures and their impact on performance and memory consumption. Study parallelization issues.

Unsupervised Machine Learning

Alex Rodriguez

During this module, we will follow the whole procedure of exploratory data analysis with this technique, starting from the raw data and finishing in the validation of the results. This path will lead us to introduce some concepts of other unsupervised machine learning techniques like dimensional reduction. The duration of the module is for three days, with both theoretical and practical lessons.

The Finite Element Method Using deal.II

Luca Heltai - special guest

Hands-on module that guides the students to solve a simple poisson problem.

from February to June

SECOND PART

2019-2020

Spatial Locality Algorithms

Riccardo Valdarnini

Theory and applications of algorithms for spatial locality.

Reduced Order Methods for Computational Mechanics

Gianluigi Rozza

In this course we present reduced basis (RB) approximation and associated a posteriori error estimation for rapid and reliable solution of parametrized partial differential equations (PDEs).

Electronic Structure: from Blackboard to Source Code

Stefano de Gironcoli

Material science and condensed matter theorists extensively employ in their research ab initio atomistic simulations as implemented in a number of widely available software codes. Most often these tools are used as 'black boxes' with little or only partial knowledge of the practical implementation of the general theoretical ideas they are based on.This is particularly severe drawback when it hinders the development of new analysis tools or computational experiments due to the lack of insight on the internal structure of the employed research software.

Master in High Performance Computing

2019-2020

from February to June

SECOND PART

Approximation and Interpolation

Nicola Seriani

Introduction to several techniques for efficient approximation of numerical functions to varying degrees of accuracy.

Fast Fourier Transforms

Ralph Gebauer - Ivan Girotto

Introduction to the Discrete Fourier Transform (DFT) and its application to real problems. From the Discrete to the "Fast" version (FFT). Analysis of a most common algorithm for the solution of a multi-dimensional FFT on parallel systems.

Supervised Machine Learning

Valerio Consorti

Data are becoming the new gold mine in modern companies. The ability to retrieve important information from very large data-sets is more and more requested on the market. This course is focused on teaching how to handle a complete complex data analytics process, by leveraging on supervised machine-learning techniques.

MHPC
Master in High Performance Computing

2019-2020

from February to June

SECOND PART

Monte Carlo Method

Sandro Sorella

Theory and applications of the Monte Carlo methods. Hands-on with examples, analysis of simulations and parallelization.

Reinforcement Learning

Antonio Celani

An introduction to the basic concepts and algorithms that stand at the foundations of Reinforcement Learning.

Advanced Optimization Techniques

Christopher Dahnken

The course presents advanced topics in optimization techniques needed in HPC environment. In particular it will focus on the use of application accelerators in high-performance and scientific computing and issues that surround it.

from February to June

SECOND PART

Deep Learning

Alessio Ansuini - Cristiano De Nobili

In the first part of the course, we will introduce the basics of deep learning. Starting from the simple artificial neuron (perceptron) we will introduce artificial neural networks and their most common architectures, such as fully connected and convolutional models. Then, we will see how these networks learn to accomplish specific tasks from data. Backpropagation and optimization of the loss function will be explained together with the standard pipeline to train, validate, and test these models. We will also introduce some basic concepts of unsupervised deep learning. In the second part, we will introduce one of the main areas where deep learning has demonstrated to be successful: Natural Language Processing (NLP). In particular, we will show the power of Contextual Word Embedding and use state-of-the-art algorithms such as BERT-like models to solve a couple of complex tasks. We will touch upon the theoretical part of BERT-like architectures but the lectures will be practical. We will use PyTorch as the main deep learning framework, while SpaCy (by Explosion AI) and Transformers (by Hugging Face) as NLP libraries.

Deep Generative Models with TensorFlow 2

Matt Archer - Piero Coronica

Introduction to the Deep Learning library TensorFlow 2, its APIs and architecture. The hands-on sessions focus on the fascinating field of generative modeling.

Master in High Performance Computing

2019-2020

from February to June

SECOND PART

Other courses from the past editions:

- Big Data Processing with MapReduce
- Lattice Boltzmann
- Molecular Dynamics
- Object Oriented Programming
- HPC Applications in Science and Technology
- Lookup Tables, Cell and Neighbor Lists
- Domain Decomposition Methods

CONTACT US

SISSA
Via Bonomea, 265
34136 Trieste
ITALY

ICTP
Strada Costiera, 11
34151 Trieste
ITALY

www.mhpc.it
info@mhpc.it