

# **Thematic CERN School of Computing 2020**

**Sunday, June 7, 2020 - Saturday, June 13, 2020**

**Split, Croatia**

## **Academic programme**

The school will focus on the theme of **Efficient Scientific Software for Heterogeneous Architectures**. The complete programme will offer 25 hours of lectures and hands-on exercises, as well as an additional student presentations session, and a special evening lecture.

## Introduction lecture

### **Scientific and computing challenges in fundamental physics**

*by Ivica Puljak (University of Split)*

Big question and challenges in modern science, with emphasis on fundamental physics

Connecting great theoretical ideas and modern experiments to test them

Future challenges in computing: from traditional increase in data throughput, volume and complexity to emerging concepts of quantum computing, machine learning and artificial intelligence

## Track 1: Technologies and Platforms

4 hours of lectures and 4 hours of hands-on exercises

*by Andrzej Nowak (TIK)*

### **Introduction to efficient computing**

The evolution of computing hardware and what it means in practice

The seven dimensions of performance

Controlling and benchmarking your computer and software

Software that scales with the hardware

Advanced performance tuning in hardware

### **Hardware evolution and heterogeneity**

Accelerators, co-processors, heterogeneity

Memory architectures, hardware caching and NUMA

Compute devices: CPU, GPU, FPGA, ASIC etc.

The role of compilers

### **Data-oriented design**

Hardware vectorization in detail – theory vs. practice

Software design for vectorization and smooth data flow

How can compilers and other tools help?

### **Summary and future technologies overview**

Teaching program summary and wrap-up

Next-generation memory technologies and interconnect

Rack-sized data centres and future computing evolution

## Track 2: Parallel and Optimised Scientific Software

4 hours of lectures and 4 hours of hands-on exercises

*by Sebastien Ponce (CERN)*

### **Writing parallel software**

Amdahl's and Gustafson's laws  
Asynchronous execution  
Finding concurrency, task vs. data parallelism  
Using threading in C++ and Python, comparison with multi-process  
Resource protection and thread safety  
Locks, thread local storage, atomic operations

### **Modern programming languages for HEP**

Why Python and C++ ?  
Recent evolutions: C++ 11/14/17  
Modern features of C++ related to performance  
Templating versus inheritance, pros and cons of virtual inheritance  
Python 3, and switching from Python 2

### **Optimizing existing large codebase**

Measuring performance, tools and key indicators  
Improving memory handling  
The nightmare of thread safety  
Code modernization and low level optimizations  
Data structures for efficient computation in modern C++

### **Practical vectorization**

Measuring vectorization level  
What to expect from vectorization  
Preparing code for vectorization  
Vectorizing techniques in C++: intrinsics, libraries, autovectorization

## **Track 3: Programming for Heterogeneous Architectures**

4 hours of lectures and 4 hours of hands-on exercises  
by *Daniel Campora (Nikhef)*

### **Scientific computing on heterogeneous architectures**

Introduction to heterogeneous architectures and the performance challenge  
From general to specialized: Hardware accelerators and applications  
Type of workloads ideal for different accelerators  
Trade-offs between multi-core and many-core architectures  
Implications of heterogeneous hardware on the design and architecture of scientific software  
Embarrassingly parallel scientific applications in HPC and CERN

### **Programming for GPUs**

From SIMD to SPMD, a programming model transition  
Thread and memory organization  
Basic building blocks of a GPU program  
Debugging and profiling a GPU application

### **Parallel cross-architecture programming**

Data locality, coalesced memory accesses, tiled data processing

Control flow, synchronization, atomics

Other standards: SYCL, HIP, OpenCL

Middleware libraries and cross-architecture compatibility

### **Design patterns and best practices**

GPU streams, pipelined memory transfers

Good practices: single precision, branchless, avoid register spilling, convert the problem

Reusable parallel design patterns with real-life applications

Under the hood: Warps, masked execution, floating point rounding

## **Additional lectures**

### **Student presentations session**

Special evening lecture

### **Future of the Universe and of Humanity**

*by Ivica Puljak (University of Split)*