

Thematic CERN School of Computing 2022

Sunday, 1 May 2022 - Saturday, 7 May 2022

Cargèse, Corsica, France

Academic programme

The school will focus on the theme of **Scientific Software for Heterogeneous Architectures**. The complete programme will offer 28 hours of lectures and hands-on exercises, and a student presentations session.

Introduction lecture

Preparing for the HL-LHC computational challenge

by *Daniilo Piparo (CERN)*

HEP data processing and analysis workflows

Upgrades of the LHC accelerator and experiments

Evolution of hardware and computing infrastructure

Impact on HEP data processing software

Track 1: Technologies and Platforms

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

by *Andrzej Nowak*

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

Future computing evolution

Track 2: Parallel and Optimised Scientific Software

4 hours of lectures and 5 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 6 hours of hands-on exercises
by Dorothea vom Bruch (CPPM/CNRS)
and Daniel Campora (University of Maastricht)

Scientific computing on heterogeneous architectures (*D. vom Bruch*)

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 (*D. vom Bruch*)

From SIMD to SPMD, a programming model transition

Thread and memory organization
Basic building blocks of a GPU program
Control flow, synchronization, atomics

Performant programming for GPUs (*D. Campora*)

Data locality, coalesced memory accesses, tiled data processing
GPU streams, pipelined memory transfers
Under the hood: branchless, warps, masked execution
Debugging and profiling a GPU application

Design patterns and best practices (*D. Campora*)

Good practices: single precision, floating point rounding, avoid register spilling, prefer single source
Other standards: SYCL, HIP, OpenCL
Middleware libraries and cross-architecture compatibility
Reusable parallel design patterns with real-life applications

Additional lectures

Student lightning talks session