



Introduction to Node-Level Performance Engineering

Dr Georg Hager
(Erlangen Regional Computing Centre - RRZE)

Performance engineering does not identify the search of hotspots and bottlenecks of computer software solely; instead, it covers the process of understanding the deeper relations between software and hardware, to exploit the underlying architecture improving the efficiency of computational codes.

At the end of this all-day course, participants will have a clearer picture of the basic architecture of a computational node, and how to measure code performances on it. Moreover, a way to predict how to gain benefits from code optimisation will be introduced.

Schedule:

Slot 1 [10:00-11:30]

Introduction

- Scalability vs. performance
- 1000 x 0 = 0: serial performance does matter

Basic node architecture

- Pipelines, SMT, SIMD, cores, sockets
- ccNUMA structure
- Data transfers and caches

Slot 2 [11:30-13:00]

A short introduction to OpenMP

Topology and affinity

- Thread binding for OpenMP

Microbenchmarking: learning from simple loops

- Motivation: vector triad
- Divide throughput
- OpenMP overhead

Slot 3 [14:30-16:00]

Working with hardware performance counters

- General guidelines
- Example: likwid-perfctr
- Case study: Detecting and fixing load imbalance

The Roofline model, part 1

- Basic assumptions
- Architectural comparisons
- Limits of applicability

Slot 4 [16:00-17:30]

The Roofline model, part 2

- Basic example: dense MVM
- Advanced: stencil algorithms
- Outlook

Wrap-up: Q&A

Prerequisites:

Basic knowledge of C or Fortran

Language:

The course will be in English

Info:

When May 31, 2019
Where DISMA - Politecnico di Torino
Aula Buzano [third floor]

Contacts:

stefano.berrone@polito.it
federico.tesser@polito.it
excellence.disma@polito.it