



Tutorial on Node-Level Performance Engineering at University of Basel, Switzerland March 13-14, 2017

The advent of multi- and many-core chips has led to a further opening of the gap between peak and application performance for many scientific codes. This trend is accelerating as we move from petascale to exascale. Paradoxically, bad node-level performance helps to "efficiently" scale to massive parallelism, but at the price of increased overall time to solution. If the user cares about time to solution on any scale, optimal performance on the node level is often the key factor. We teach the architectural features of current processor chips, multiprocessor nodes, and accelerators, as far as they are relevant for the practitioner. Peculiarities like SIMD vectorization, shared vs. separate caches, bandwidth bottlenecks, and ccNUMA characteristics are introduced, and the influence of system topology and affinity on the performance of typical parallel programming constructs is demonstrated. Performance engineering and performance models are powerful tools that help the user understand the bottlenecks at hand and to assess the impact of possible code optimizations. A cornerstone of these concepts is the Roofline model, which is described in detail. Hands-on exercises will provide ample opportunity for attendees to put all concepts to the test.

Lecturers: [Dr. Georg Hager](#) and [Prof. Gerhard Wellein](#), FAU Erlangen, Germany

Location: Kollegienhaus, Petersplatz, 4003 Basel, Regenzimmer 111

Contact: Mrs. Yvonne Wegmüller (yvonne.wegmueller@unibas.ch)

Registration deadline: March 10, 2017 (by email)

Fees: MA/PhD students CHF 50.-, Academic staff CHF 150.-, Industry staff CHF 250.-.

Requirements: Laptop with an SSH client installed.

Agenda

Monday, March 13, 2017

08:00 – 09:00	Registration
09:00 – 10:15	Computer architecture and simple multicore tools
10:15 – 10:45	Coffee break
10:45 – 12:00	Microbenchmarking for architectural exploration (and more)
12:15 – 13:30	Lunch (provided)
13:30 – 14:30	Introduction to the Roofline performance model
14:30 – 17:00	Hands-on exercise (with on-the-fly coffee break)

Tuesday, March 14, 2017

09:00 – 09:30	Hardware performance counters
09:30 – 10:15	Roofline case study: dense matrix-vector multiplication
10:15 – 10:45	Coffee break
10:45 – 12:00	Roofline case study: Jacobi smoother
12:15 – 13:30	Lunch (provided)
13:30 – 14:30	Optimal use of parallel resources: SIMD & ccNUMA
14:30 – 17:00	Hands-on exercise (with on-the-fly coffee break)

Further tutorial details:

<http://goo.gl/0Qv45o>