TUTORIAL: Scientific Computing with GPUs

Organizers and Presenters:

Dominik Göddeke http://www.ians.uni-stuttgart.de/cmcs/
Robert Strzodka https://asc.ziti.uni-heidelberg.de/en

Abstract:

GPU Computing has evolved from an obscure niche to a commonplace technique in the arsenal of simulation scientists. However, there is an increasing gap between hardware capabilities and their application: Hardware and software capabilities continue to evolve at a higher pace than their adoption in scientific codes. The goal of this tutorial is to bridge this gap, by providing a hands-on overview on recent advances, focusing on simulation workloads rather than topics advertised elsewhere. Each session incorporates a detailed practical.

Target audience:

The tutorial is centered on CUDA by NVIDIA. We aim at an audience already familiar with C/C++ and CUDA. Participants are expected to bring their own laptops in order to connect to the provided GPU servers with ssh.

Schedule:

Session 1: **12:00 – 13:00**
basic CUDA review, hardware capabilities of current GPU architectures Kepler, Pascal and Volta

Lunch: **13:00 – 13:30**

Session 2: **13:30 – 15:30**
streams, host-device async. parallelism, buffered IO, memory optimizations, ...

Break: **15:30 – 16:00**

Session 3: **16:00 – 18:00**
unified memory, multi-GPU and GPU-MPI,
Summary and Roundup