



PAUL SCHERRER INSTITUT

Invitation

LMU-Seminar

Title: Dynamic Kernel Scheduler (DKS) - a thin software layer between host application and hardware accelerators

Speaker: Uldis Locans (Paul Scherrer Institute)

Time: Tuesday, March 3, 2015

Place: WBGB/021

Abstract:

Emerging processor architectures such as GPUs and Intel MICs provide a huge performance potential for high performance computing. However developing software using these hardware accelerators introduces additional challenges for the developer such as exposing additional parallelism, dealing with different hardware designs (Nvidia GPUs, AMD GPUs and APUs, Intel MICs) and using multiple development frameworks in order to use devices from different vendors (CUDA, OpenCL, OpenACC, OpenMP).

Dynamic Kernel Scheduler (DKS) is being developed in order to provide a software layer between host application and different hardware accelerators. DKS handles the communication between the host and the hardware accelerators, schedules task execution, and provides a library of built-in algorithms. Algorithms available in the DKS library will be written in CUDA, OpenCL and OpenMP. Depending on the available hardware, the DKS can select the appropriate implementation of the algorithm.

First DKS prototype was developed using CUDA and tested on Nvidia Tesla K20 GPU. DKS was integrated in OPAL (Object-oriented Parallel Accelerator Library) to speed up FFT Poisson solver and Monte Carlo simulations for proton therapy degrader. DKS was also used together with Minuit2 for parameter fitting, where χ^2 and max-log-likelihood functions were offloaded to hardware accelerator. The concepts of DKS and the results of first tests will be presented. Results show that substantial speedups can be achieved on hardware accelerators using DKS.