Contribution ID: **46** Type: **not specified**

Novel GPU features: Performance and Productivity

Thursday, 11 September 2014 11:00 (45 minutes)

The huge amount of computing power needed for signal processing and off-line simulation makes High-Energy Physics an ideal target for GPUs. Since the first versions of CUDA, considerable progress has been made in demonstrating the benefit of GPUs for these processing pipelines and GPUs are now being deployed for production systems. However, early experiments also showed some of the challenges encountered in HEP specific tasks, including memory footprint, complex control flow, phases of limited concurrency and portability. Many of these concerns have been addressed with recent changes to the GPU hardware and software infrastructure: Unified memory, dynamic parallelism, and priority streams are just some of the features at the developer's disposal to fully take advantage of the available hardware. In addition, recently added boards like the TK1 processor for embedded high performance, low power applications enable now CUDA accelerated applications from the sensor to the offline simulations. In this talk I will present some of the latest additions to the GPU hardware, provide an overview of the recent changes to the software infrastructure and will walk through features added in the latest CUDA version.

Primary author: TBC

Presenter: Dr MESSMER, Peter (NVIDIA)

Session Classification: GPU in Other Applications (1/2)