NVIDIA Application Lab Kick-off Workshop

Status: 11.09.2012

Wednesday, 19.9.2012

9:00	10:30	Peter Messmer (NVIDIA)	Kepler and CUDA 5 – Overview and first experiences
10:30	11:00	*Coffee*	
11:00	11:45	Markus Axer/ Marcel Huysegoms (FZJ, INM-1)	How to build a 3D human brain model from thousands of histological sections
11:45	12:30	Frank Winter (U Edinburgh)	Migration of a Data-Parallel, Expression Template based Vector Library to GPU-enabled Multi-Node Systems
12:30	14:00	*Lunch*	
14:00	14:45	John Romein (ASTRON)	Signal Processing on GPUs for Radio Telescopes
14:45	15:30	Richard Leys (U Heidelberg)	EXTOLL – A feature rich interconnect for HPC
15:30	16:00	*Coffee*	
16:00	17:30	Jiri Kraus (NVIDIA)	Cuda-aware MPI (incl. Hands-on)

19:30	…	*Dinner (Kaiserhof)*

Thursday, 20.9.2012

9:00	9:45	Rainer Spurzem (NAOC/CAS Beijing and ARI/ZAH U Heidelberg)	GPU Supercomputing in China and Germany and Application to super-massive black holes and gravitational waves from galactic nuclei
9:45	10:30	Fabio Schifano (U Ferrara)	Implementation and optimization of a thermal Lattice Boltzmann Algorithm on a multi-GPU cluster
10:30	11:00	*Coffee*	
11:00	11:45	Jan Meinke (FZJ, JSC)	Modelling Molecules with Python and GPUs
11:45	13:00	Peter Philippen (FZJ, JSC)	Score-P - Scalable performance analysis on accelerated architectures (incl. Hands-on)

NVIDIA Application Lab Kick-off Workshop
Abstracts

Rainer Spurzem: GPU Supercomputing in China and Germany and Application to super-massive black holes and gravitational waves from galactic nuclei

To use graphical processing units (GPU) for general purpose computations is becoming more and more ubiquitous on all scales - from user's desktops to the most powerful supercomputers. China has been in the past two years pioneering this field, now operating some of the most powerful GPU accelerated supercomputers in the world. An overview of our research with these clusters in China and Germany is presented, especially also the new system installed at FZ Jülich few months ago in collaboration between University of Heidelberg, Munich, FZ Jülich and the German Science Foundation (DFG). We discuss particle- and mesh-based algorithms for astrophysics using hundreds to thousands of GPUs for one single application run in a parallel message passing environment, some with detailed timing models. Future perspectives for GPU and FPGA accelerated computing will be shown. GPU and similar multi-core accelerator hardware could be a stepping stone on the path to reach Exascale supercomputing. As an example our astrophysical application - computer models of star clusters containing super-massive black holes - is explained. We use large high-order direct N-body simulations, parallelised with MPI across a large number of nodes and across many GPU thread processors on each node. Our simulations proceed from a merger of two galaxies with central black holes to the complete relativistic merger of the black holes, including Post-Newtonian corrections to gravitational forces. The expected gravitational radiation from our model sources will be briefly touched. We discuss the relevance of this for current and future instruments to measure gravitational waves.