

ETH

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

CSCS

Swiss National Supercomputing Centre



Panel on Future Directions in GPU HW

Sadaf Alam, Swiss National Supercomputing Centre (CSCS)

SOS14 Conference

March 8-11, 2010

Savannah, GA

Questions

1. Status: Is your organization currently pursuing heterogeneous computing? Small scale, large scale? Why or why not?
2. Apps: Which specific applications at your site (could) exploit heterogeneous manycore, if any?
3. Software: What programming models and paradigms do you use, need for these architectures to be successful?
4. Arch: What specific architectural features do/could your application exploit?

Acknowledgements: feedback from CSCS colleagues especially Ugo Varetto; and discussions during the recent Hybrid Multi-core Consortium (HMC) and PRACE WP 8 meetings

Answer #1: Yes

About

- Context of this Initiative
- Developments in HPC
- Organization of HP2C
- Updates

Welcome to hp2c

Welcome to HP2C, the *Swiss Platform for High-Performance and High-Productivity Computing*.

The HP2C platform aims at developing applications to run at scale and make efficient use of the next generation of supercomputers. Presently, in 2009, this will be the generation of computing technologies available in 2013 timeframe.

The platform consists of domain science projects that are lead by research groups at Swiss universities and Institutes of the ETH Domain, and supported by a core group of scientific computing experts in the Lugano area.

HP2C is jointly operated by the *Swiss National Supercomputer Center (CSCS)* and the *Institute for Computational Sciences of the University of Lugano (USI)*. Project teams engage in high-risk and high-impact application development for HPC systems at scale.

Updates

First Projects Accepted
The first 8 projects to be supported by hp2c have been accepted. A decision about 4 additional...

Proposal submission closed
The proposal submission for HP2C closed on September 30, 2009.

The Swiss HPC Service Provider Community

For additional information about HPC activities in Switzerland you may want to have a look to www.hpc-ch.org, the web pages of the Swiss HPC Service Provider Community.



Answer # 2: HP2C Apps

- **Advanced gyrokinetic numerical simulations of turbulence in fusion plasmas;** Prof. Laurent Villard, EPF Lausanne
- **CP2K - New Frontiers in ab initio Molecular Dynamics;** Prof. Dr. Juerg Hutter, Uni Zürich
- **Computational Cosmology on the Petascale;** Prof. Dr. George Lake, Uni Zürich
- **Selectome, looking for Darwinian evolution in the tree of life;** Prof. Dr. Marc Robinson-Rechavi, Uni Lausanne
- **HPC for Cardiovascular System Simulations;** Prof. Alfio Quarteroni, EPFL
- **MAQUIS - Modern Algorithms for Quantum Interacting Systems;** Prof. Thierry Giamarchi, University of Geneva
- **PETAQUAKE - Large-Scale Parallel Nonlinear Optimization for High Resolution 3D-Seismic Imaging;** Dr. Olaf Schenk, Uni Basel
- **Productive 3D Models of Stellar Explosions;** Dr. Matthias Liebendörfer, Uni Basel
- And 3 or 4 more ...

Answer # 3 & 4: Programming Models & Arch Evolution for Scalable Systems

③ Programming models

– Today

- MPI + CUDA; MPI + pthreads + CUDA; MPI + PGI directives; MPI + HMPP directives; MPI + BSC superscalar directives; etc.

(the hybrid models come under no guarantee and warranty. On different variants of operating systems, versions of CUDA, drivers, MPI libraries, etc. program behavior could change without any warnings. Users and code developers are encouraged to exercise caution when moving data to and from the device and overlapping communication on the network and GPU devices. Any expectation of debugging and performance tools is deemed premature..)

– For *scalable* heterogeneous systems

- Encapsulation on distinct memory hierarchies across systems; Evolutionary path for existing applications; Performance consistency across variants of heterogeneous nodes; IO models integration; ...
- Eventually some tools enabling discoveries of programming errors and omissions, and some tuning and optimization

④ Expectations for scalable architectures

- Upgradable components (GPU chip and memory)
- Chipsets for HPC / server technologies
- Resolution of class wars between different types of cores and their access to system resources
- Support for graceful failures and recovery
- Of course some more memory and bandwidth won't hurt!

