ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

CSCS
Swiss National Supercomputing Centre

High-Performance and High-Productivity Computing (HP2C) Platform

Thomas C. Schulthess

Overarching goal of HP2C

Prepare computational sciences to make effective use of next generation supercomputers

Specific goal

Emerge with several high-impact scientific applications that scale and run efficiently on leadership computing platforms in 2012/13 timeframe

Build on, and multiply the early science applications experience on Jaguar at ORNL

DCA++: simulate models of high-temperature superconductivity
first sustained petaflop/s in production runs (Gordon Bell Prize 2008)

WL-LSMS: simulate termodynamics properties in magnetic nanoparticles sustained petaflop/s in production runs (Gordon Bell Prize 2009)

Elements of the DCA++ project

- Start with a real and challenging scientific problem
  - Simulations to understand high-temperature superconductivity
- Take the best methods known in the field
  - Quantum cluster theory and quantum Monte Carlo (DCA/QMC)
- Understand their limitations on today's emerging hardware
  - QMC kernels are memory bandwidth limited
- Algorithmic reengineering to better map method to hardware
  - Introduced delayed updates ("Ed" updates) into QMC algorithm, method is no longer limited by memory bandwidth
- Take aggressive view on rewriting codes – don't hesitate to rewrite major portions or even starting from scratch
  - DCA++ was a total re-write, more general, and extensible
- Pay close attention to new hardware developments
  - GPGPU work started in 2007 and motivated mixed precision solvers

Learning from the Oak Ridge experience: Covering all aspect of the simulation system

Distributing the tasks in Switzerland:

Example based on ORNL's early science teams that run on the first petaflop/s systems

Simulations

Models,

Methods,

& Implementation

Map to Hardware

System operation

System design

Users

Domain projects at Swiss univers.

Sci. comp. research CSCS & U. of Lugano

CSCS's (HPC Centers) traditional role

Systems research D-INFK & vendors

Physics (chemistry, ...)

Application software

Comp. mathematics

Computer Science

Computer Center

Hardware vendor

High-Performance and High-Productivity Computing (HP2C)

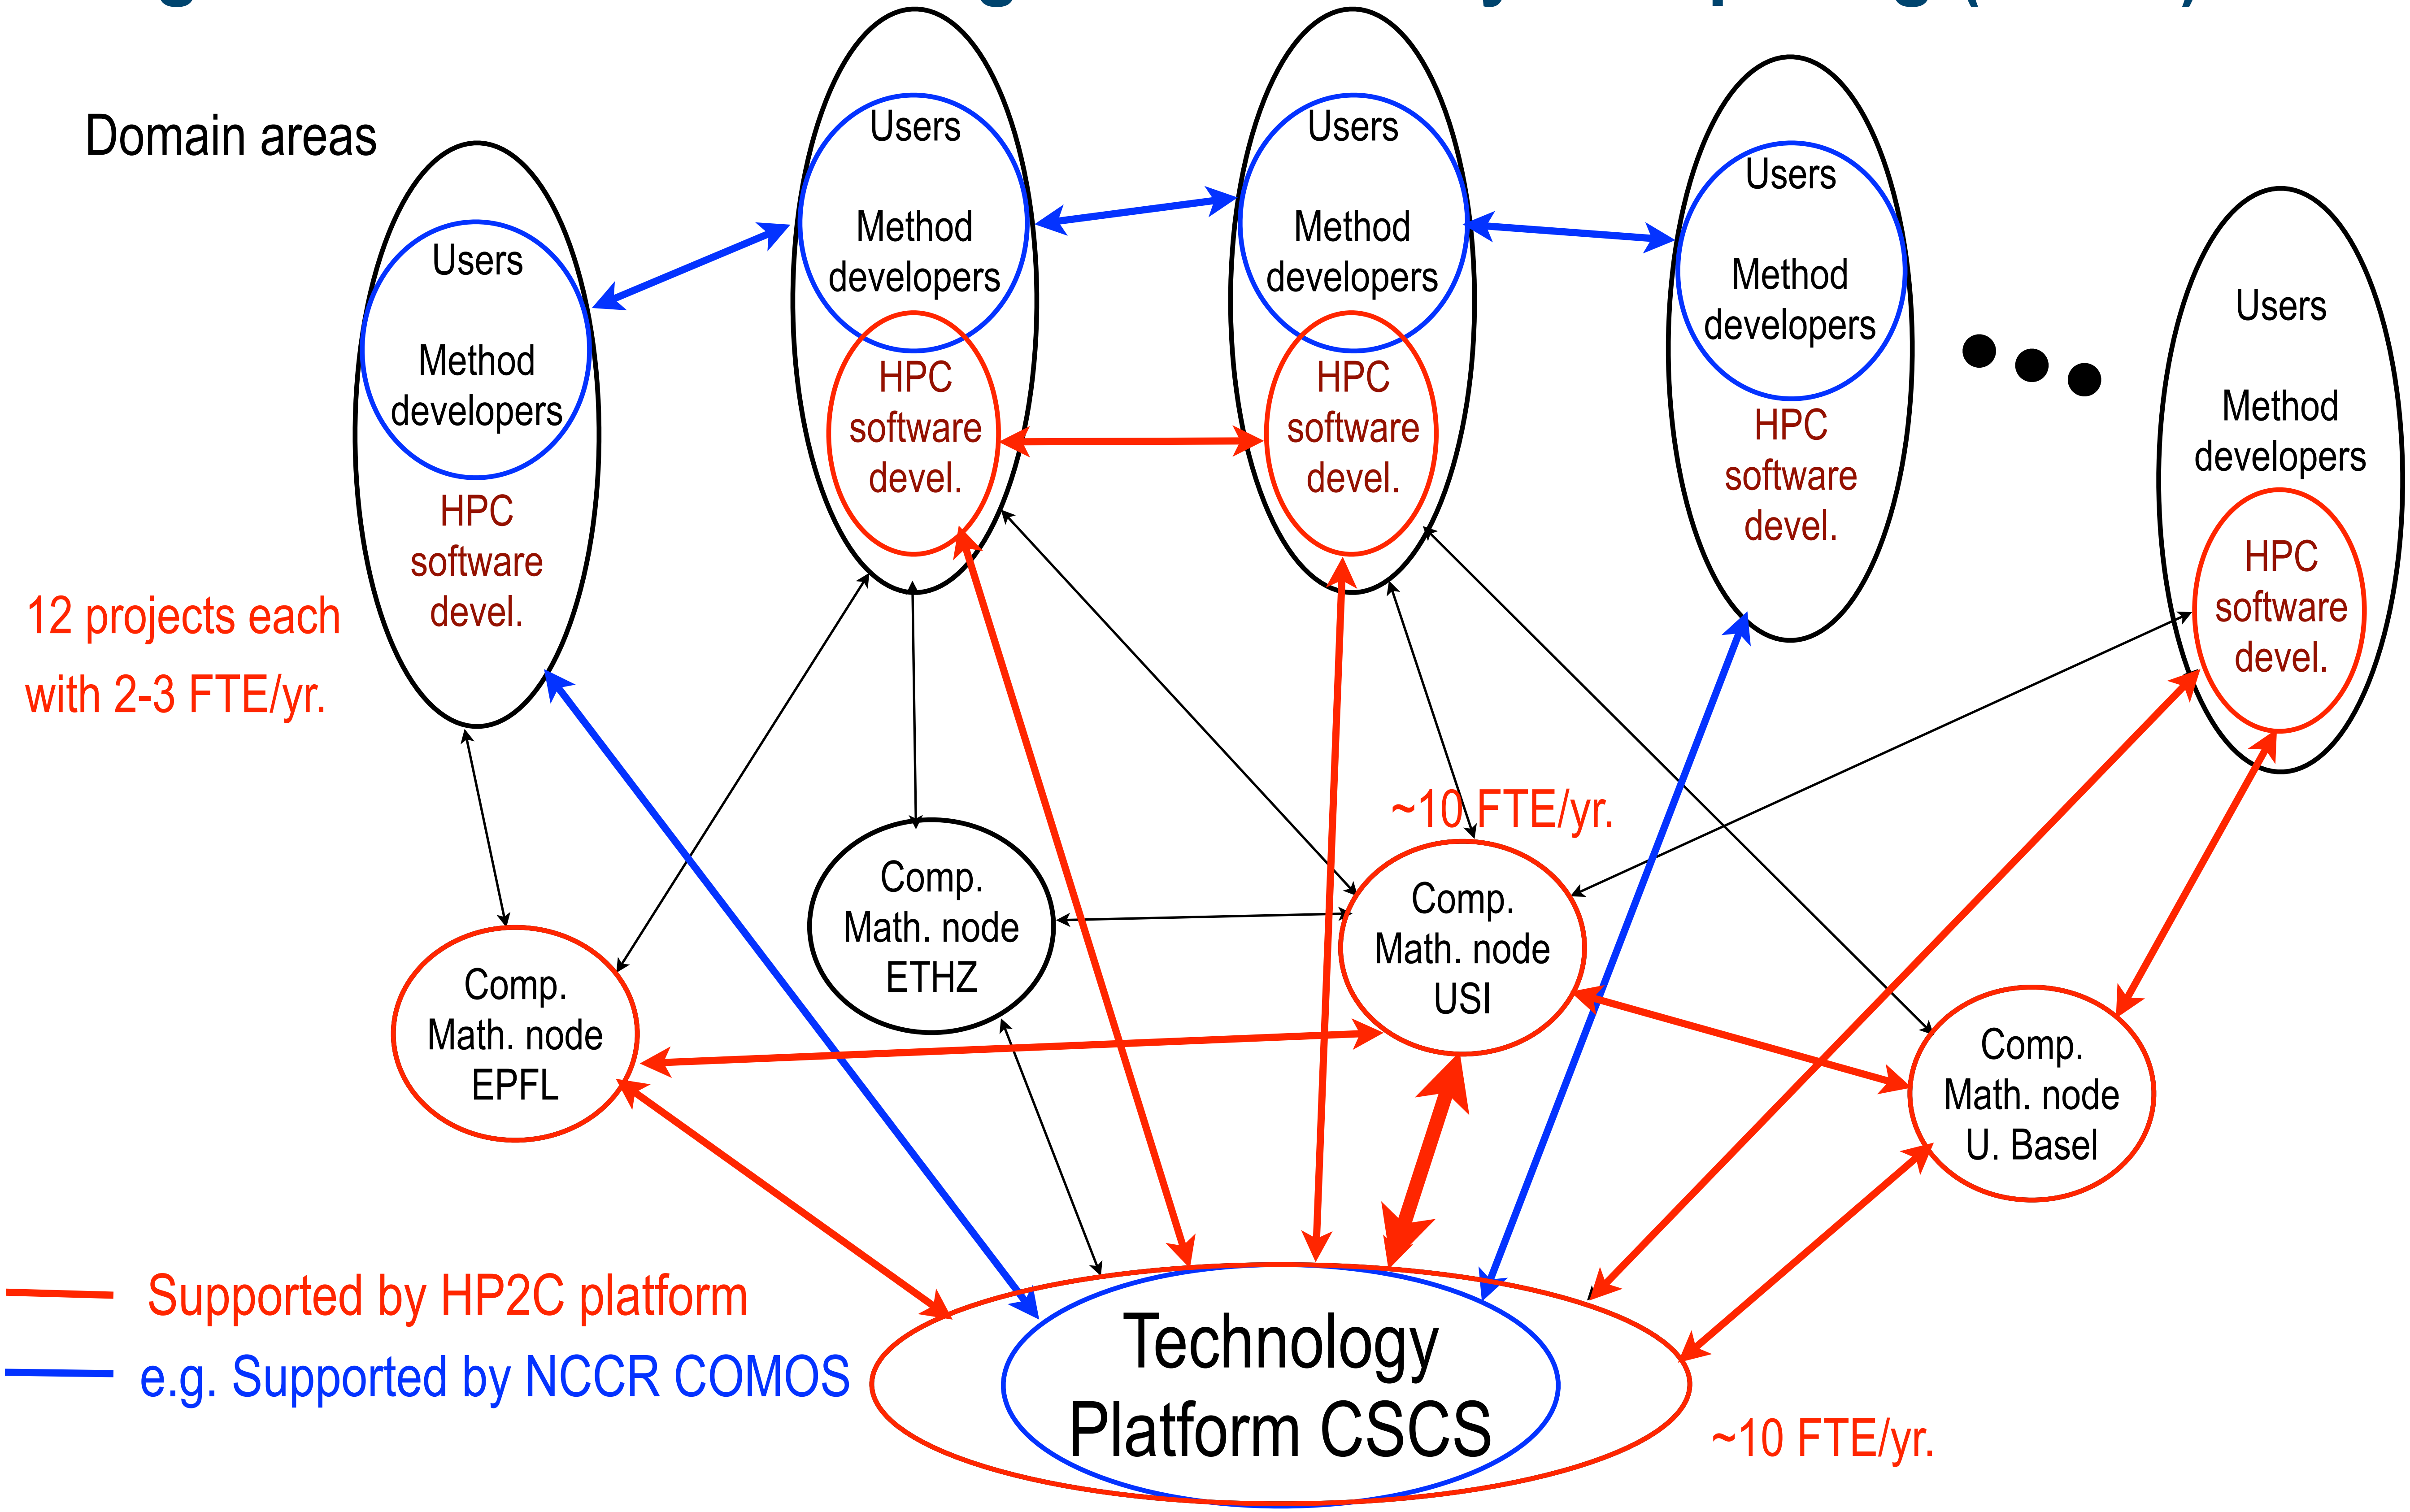

HP2C Projects have to face “brutal facts of HPC”

- Massive concurrency: applications will have to put up with millions (billions) of threads
- Less and (relatively) slower memory per thread: memory consideration should be integral part of complexity analysis
- Only slow improvements in inter-processor and inter thread communications - remember that speed of light is constant!
- Stagnant I/O subsystems: you don’t want to limit progress in simulation capabilities with rate of progress in long-term storage technologies
- Resilience and fault tolerance: resilience towards failure of individual components; (energy) cost to error detection and correction is non-negligible

Expected research priorities of HP2C projects

- Significant problems that require orders of magnitude more computer power than what is available today
- Significant re-engineering of algorithms and refactoring of codes - scientific progress cannot be limited by legacy software
- Consider emerging parallel programming models - multiple levels of parallelism, PGAS, DARPA HPCS languages, hybrid multi-core
- Revisit workflows, in particular to minimize I/O

What we are today

- HP2C platform funded in June 2009 (currently through 2012)
- Call for project proposals issued on July 2009
- 17 proposals received by September 30, 2009
- Oct./Nov. 2010: external peer review and ranking of projects
  - High expectations on scientific impact and relevance of supercomputing
- Funding decision by steering committee on Dec. 18, 2009:
  - 8 projects selected for immediate start
  - 2 teams asked to resubmit revised proposal
  - 2 proposals need more review
  - 5 proposals were rejected
- Negotiations with computer industry over development / procurement of prototype (should be complete by spring 2010)

Tier 1 projects of the HP2C platform

- Advanced Gyrokinetic Numerical Simulations of Turbulence in **Fusion Plasmas** – Laurent Villard, EPF Lausanne
- New Frontiers in **ab initio Molecular Dynamics** – Jürg Hutter (CP2K), Univ. of Zurich
- Computational **Cosmology** on the Petascale – Geoge Lake, Univ. of Zurich
- Selectome, looking for **Darwinian evolution** in the tree of life – Marc Robinson-Rechavi, Univ. of Lausanne
- HPC for **Cardiovascular Systems Simulations** – Alfio Quarteroni, EPF Lausanne
- Modern Algorithms for **Quantum Interacting Systems** – Thierry Giamarchi, Univ. of Geneva
- Large-Scale Parallel Nonlinear Optimization for High Resolution **3D-Seismic Imaging** – Olaf Schenk, Univ. of Basel
- Productive 3D Models of **Stellar Explosions** – Matthias Liebendörfer, Univ. of Basel

Priorities of the core group

- Apply PGAS model to real applications – focus on
  - Fortran: co-array Fortran (CAF)
  - C: Unified Parallel C (UPC)
  - C++: create a solution based on extending C++/STL
- Hybrid-Multicore – focus still under development
  - Programming models / environments
  - Using CUDA
  - Exploring CUDA-Fortran as well as various compiler options
- Generic implementations of sparse matrix-vector and block-sparse matrix-matrix as well as stencils
- Computational mathematic priorities at University of Lugano are still under development
- **Important:** actively engage in collaborations with efforts in USA and elsewhere in Europe

QUESTIONS / COMMENTS?