Challenges in High Performance Computing

Kimmo.Koski@csc.fi

Berlin

April 7th, 2011

CSC – Tieteen tietotekniikan keskus Oy
CSC – IT Center for Science Ltd.

Topics

- HPC access
- Major challenges in the road to Exascale
- Green computing

The European HPC Ecosystem

- PRACE for high-end (Tier-0) computing
- DEISA for Tier-1 level
- PRACE and DEISA will be merged in 2011

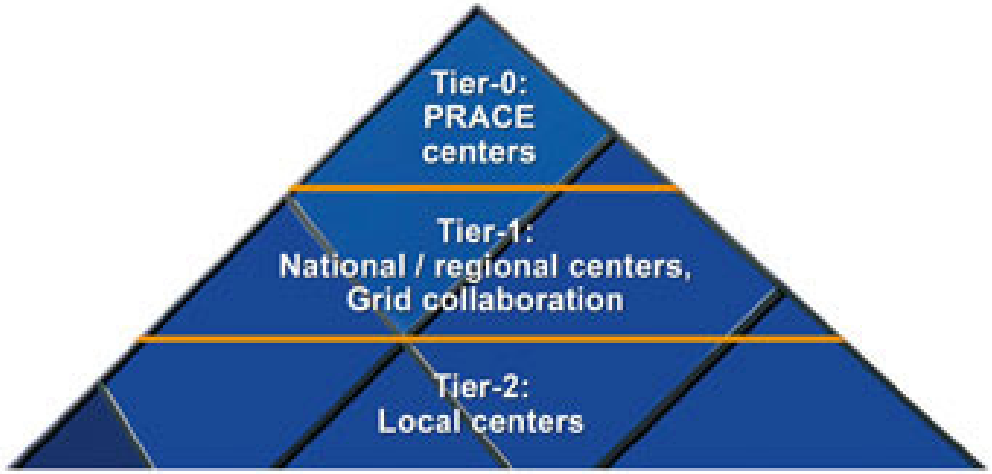

Why?

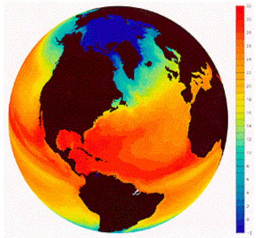
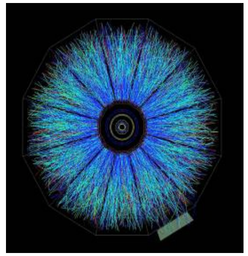
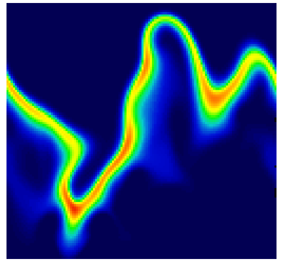

- Weather, Climatology, Earth Science
  - degree of warming, scenarios for our future climate.
  - understand and predict ocean properties and variations
  - weather and flood events
- Astrophysics, Elementary particle physics, Plasma physics
  - systems, structures which span a large range of different length and time scale
  - quantum field theories like QCD, ITER
- Material Science, Chemistry, Nanoscience
  - understanding complex materials, complex chemistry, nanosc
  - the determination of electronic and transport properties
- Life Science
  - system biology, chromatin dynamics, large scale protein dynamics, p
    association and aggregation, supramolecular systems, medicine
- Engineering
  - complex helicopter simulation, biomedical flows,
    gas turbines and internal combustion engines,
    forest fires, green aircraft,
  - virtual power plant

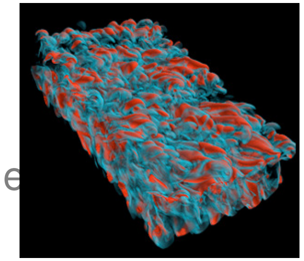
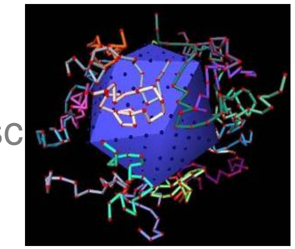
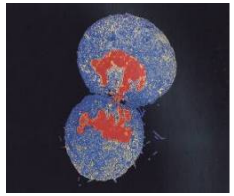
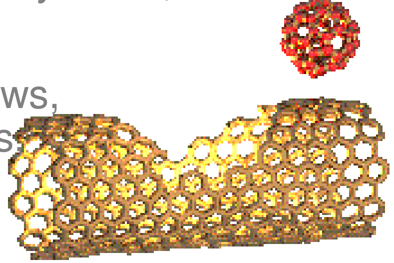
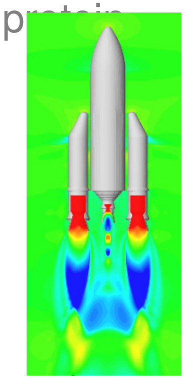
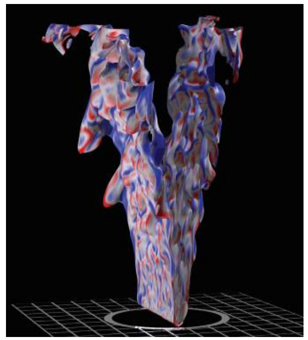

PRACE
The Partnership for Advance Computing in Europe is *the* European HPC Research Infrastructure

- PRACE enables world-class science through large scale simulations
- PRACE provides HPC services on leading edge capability systems on a diverse set of architectures
- PRACE operates Tier-0 systems as a single legal entity including user and application support
- PRACE offers its resources through a single pan-European peer review process
- PRACE is providing services since August 2010
- The first Tier-0 system is the fastest Supercomputer in Europe

Ways to access

- PRACE Tier-0 can provide high amount of cycles from the most efficient computers in Europe
  - Challenge: resources not guaranteed, peer review process
- PRACE Tier-1 can provide cycles in a reasonably efficient supercomputer
  - Challenge: better hit rate, but resources still not guaranteed
- Dedicated computer centers
  - Guaranteed resources but expensive to maintain
- Contracts with various stakeholders with HPC resources
  - Access needs to be negotiated
  - Usage of national computing resources possible, but how to handle international user communities such as ESFRI projects?

Balanced environment, better results

- “Cutting few days from computing helps only little, if it still takes six months to understand the results”
- Various components of HPC Ecosystem
  - Computing in different performance levels
  - Data management
  - Network
  - Application development, algorithm design and scalability
  - Visualization
  - Training and education

Exascale

- “We are talking about exascale even if we have not seen much petascale in real codes, yet”
- Major issues
  - Application development: requires code scalable to millions of cores
  - Choice of hardware: CPU or GPU and need for re-programming
  - Energy consumption: running exascale with reasonable power requires dramatic innovation
- Requirement of global collaboration becomes crucial
  - No sense to establish exascale systems in every location
  - Demand for work (for example in applications) is huge
- Very few HPC vendors in Europe
  - What is the role for Europe in exascale development?
- Global collaboration through IESP program

Scientific Computing World, Feb-March 2011

CSC

Bill Dally,
chief scientist at Nvidia

'There are three main challenges to reaching exascale – as defined and sustained performance on real and meaningful problems, and not just as appears on a benchmark. Probably the most difficult is power – because we don't just want to build an exascale machine, we want to construct one that can run in a machine room with a total power budget not much higher than today. We have tentatively adopted the goal of exascale at 20MW and to do that requires almost 100:1 improvement in performance per Watt, which is a very tough goal to meet. One of the main things we're working on is finding ways of improving the energy efficiency of computing so that we can close that gap and reach exascale computing in a power envelope that will be acceptable.

'Another main challenge is reliability as an exascale machine is going to have many more components than is seen today. As components become smaller, the same size fluctuations can cause larger disturbances. We need to develop techniques that make our machines very robust in the presence of transient errors, and build highly-reliable

PARALLEL MACHINES ARE ALREADY HARD TO PROGRAM AND IF YOU SCALE EVERYTHING UP AND KEEP IT AS BUSINESS AS USUAL IT'S GOING TO BE INTRACTABLE

machines out of somewhat unreliable components. The third major challenge is one of programmability. Parallel machines are already hard to program and if you scale everything up and keep it as business as usual it's going to be intractable.

'It's clear to me that the solution is a heterogeneous supercomputer where the nodes incorporate CPUs to do the serial portions of the code and GPUs to do the throughput sensitive portions. I think that's something parts of the community, and certainly the Chinese, have embraced. Realistically as GPUs are so much more energy efficient measured in Watts per Flop than their counterparts, this is the only way we can build machines of this scale with the energy efficiency we need. There are certainly those who are dragging their feet on this transition and an important message to get across is that to get to exascale we have to take the leap and embrace heterogeneous computing – it's the way of the future.'

Green ICT

- Location close to experiments or further where the energy consumption and delivery is optimal?
  - Network bandwidth
- Green aspects
  - Minimizing energy utilization (PUE, advanced technology, water cooling, cool climate, …)
  - Type of electricity and impact to energy ecosystem
  - Carbon footprint

Example: Energy efficiency (PUE cat.3)

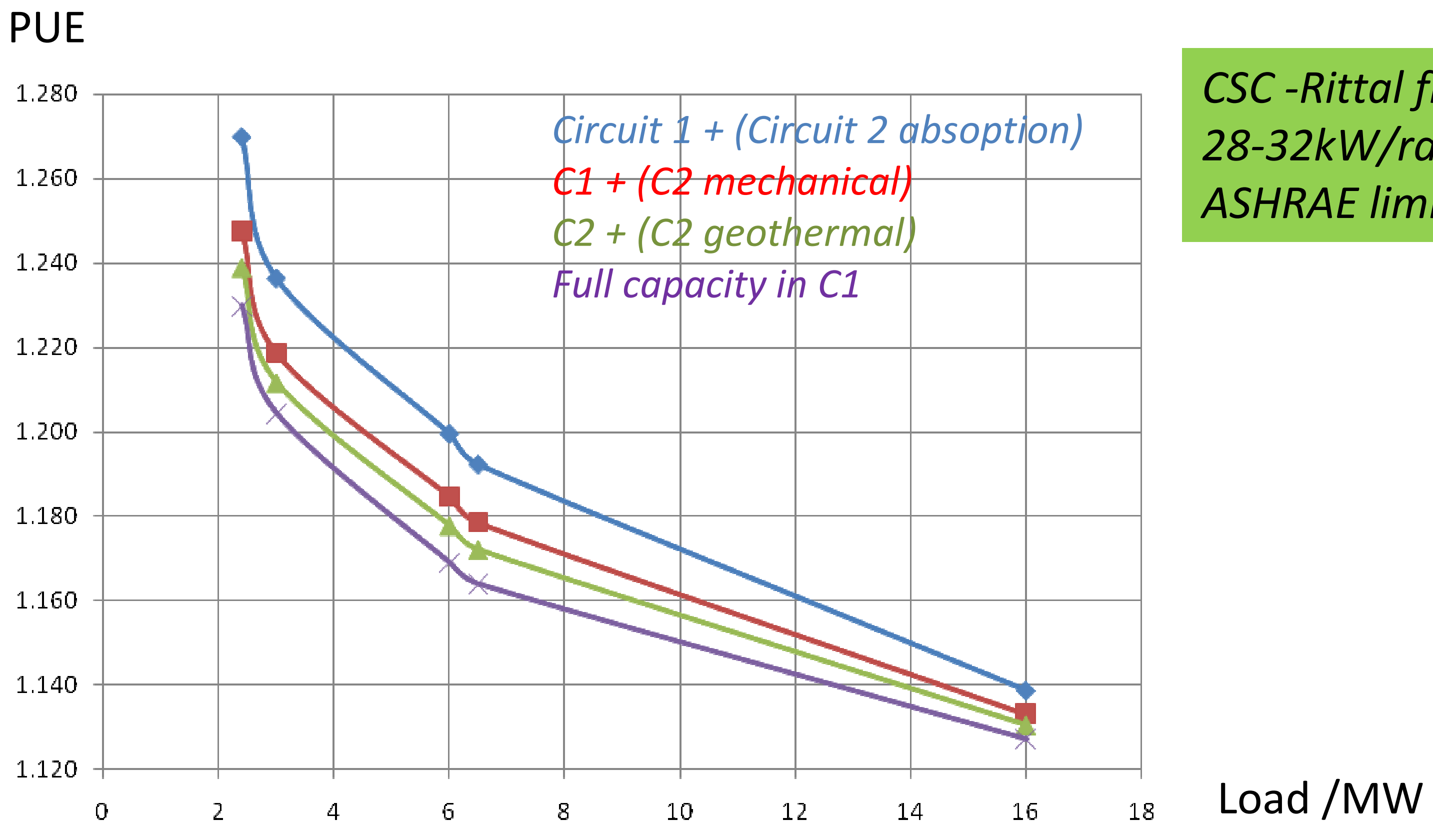

CSC -Rittal field tests: 28-32kW/rack within ASHRAE limits

*C1 load based 1400-16000kW on Rittal model with 21-23^{O}C high limit for water intake from river**
C2 load 1000-2000kW Supercomputer with 15.5^{O}C high limit for water intake (up to 13^{O}C from river)

Case CERN

- Part of computing needs close to accelerator
  - Dedicated systems for capturing data
  - Real-time processing requirements do not allow data transfer through network
- Solution for off-line computation planned
  - Processing events require huge amount of computing capacity
  - Cost efficiency and green energy play a major role
  - Datacenter remote from CERN considered

Example: CSC datacenter in 2012

- *Modern and reliable infrastructure (national power grid, roads, airline connections, data networks)*
- *World class education system and competences on ICT & energy*
- *Steady economical and political conditions*
- *Cheap energy (www.energy.eu) and strongly increasing CO_2 –free capacity*
- *Cool climate and water resources*
- *No major earthquakes (4.1 biggest ever)*
- *No major storms or other dangerous natural phenomena*
- *European Gateway to Russia*

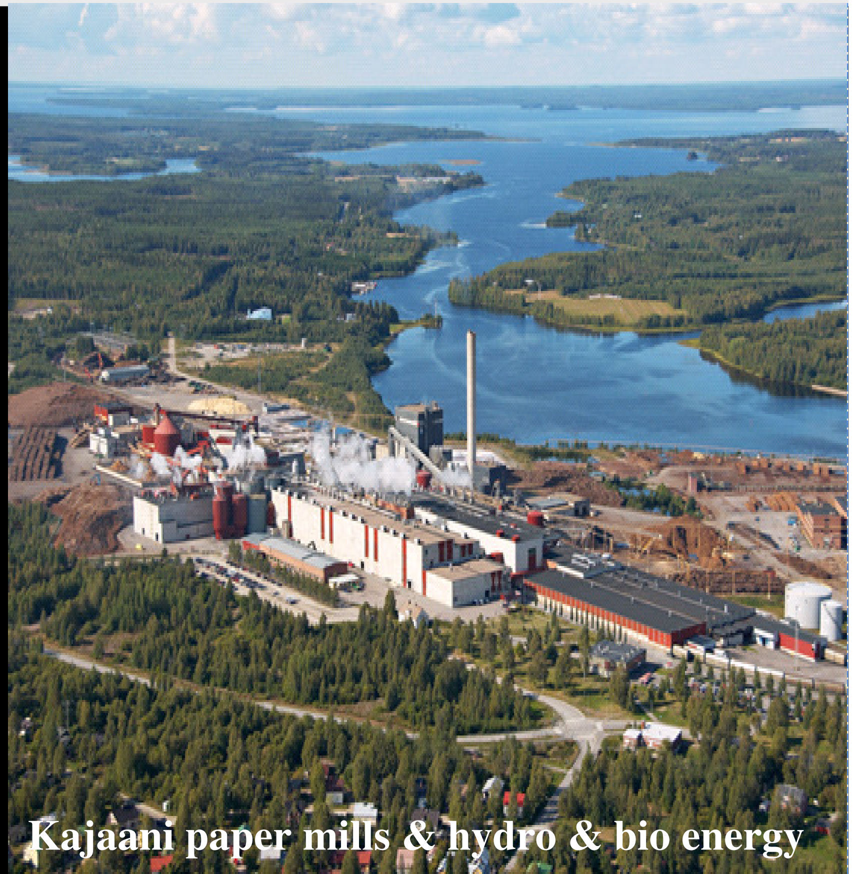

CSC


75t
51114


IDEARK

RAMIRI

Realising and Managing International Research Infrastructures

Learning Programme 2011 & 2012

Amsterdam 14-15 June

Trieste 24-25 October

SEVENTH FRAMEWORK PROGRAMME

www.ramiri.eu

Topics covered by the RAMIRI Programme

The programme for each group will consist of six learning modules spread across the two workshops. Each participant must therefore commit to attend both workshops in their learning cycle. Areas discussed will include:

- Understanding and developing European and national policies in relation to research infrastructures and access policies for users
- Planning, project management and cost control for research infrastructures along their life cycle (design, construction, operation, decommissioning)
- Financing research infrastructures, including structural funds and debt-funding through the Risk-Sharing Finance Facility
- Selecting the best legal foundation for a new research infrastructure, and developing statutes and governance principles
- Formulating policy in relation to intellectual property and data management
- Recruiting, training and retaining in Europe the best staff for an international research infrastructure

INTERNATIONAL
SUPERCOMPUTING CONFERENCE

ISC'11 – The HPC Event

Hamburg, Germany | June 19 – 23, 2011

Tutorials & Workshops, Sunday, June 19

- 3 full-day and 8 half-day Tutorials
- HPC in Asia and Mid Market HPC Workshop

Conference, Monday June 20 – Thursday, June 23

Partial List:

- 37th TOP500 List
- Debate on GPU Computing
- Heterogeneous Systems & Their Challenges to HPC Systems
- Addressing the (HPC) Data Challenges: Today & Tomorrow
- Parallel Filesystems
- Sustained Performance on Petascale Systems
- Trans-petaflop/s Initiatives
- Storage & Memory Access – The Real Challenge for Future HPC Systems
- How to Teach Parallel Programming for Millions of Cores
- State of the Art in Visualization
- Panels: Expectation for Exascale Computing & Net Zero Carbon
- HPC Applications
- Various Birds-of-a-Feather (BoF) Sessions, Research Papers & Posters
- Start-Up & Exhibitor Forums, Hot Seats
- …and much more

The 2011 Keynotes

Henry Markram, EPFL
Simulating the Brain – The Next Decisive Years

Philippe Vanier, Bull
The Road to Exascale Computing – A European Challenge?

Thomas Sterling, LSU
HPC Achievement & Impact 2011

Dean Klein, Micron Technology
Future Trends in Memory Systems: Showstopper or Performance Potential for HPC?

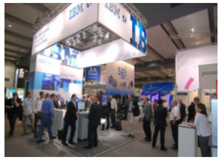

Exhibition

Monday, June 20 – Wednesday, June 22

Network while having a great time

For more information, please visit: **www.isc11.org**