CSCS

Swiss National Supercomputing Centre

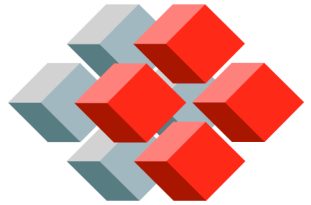

ALPS - The Swiss Grand Challenge Programme on the Cray XT3

CUG 2007, Seattle

Dominik Ulmer, CSCS

CSCS today

- Swiss National Supercomputing Centre, founded in 1991 as part of ETHZ. Since 2004, an autonomous unit working under performance mandate and with global budget
- Appr. 35 staff, providing services in HPC, Grid and visualisation to all academic institutions in Switzerland
- Two major systems
  - Cray XT3, 1650 dual-core CPU (installed 2005)
  - IBM p575 IB cluster, 768 CPU (installed 2006)

Changing the Technology Strategy

- From 1991 to 2003, CSCS' technology strategy was based on
  - Vector computing
  - Proven, mature technology
- Since 2004, CSCS is focussed on
  - MPP computing
  - Early production technology

Introduction of the Cray XT3

- Installed in July/August 2005
  - 1'100 single-core Opterons @ 2.6 GHz
  - 5.6 Tflops peak
- In production since January 2006
- Two major upgrades:
  - August 2006: Additional 50% cabinets (1'664 CPU)
  - April 2007: Dual-core processor exchange (17.2 Tflops)
- Availability > 98%, $MTBI_u$ > 300h, utilisation > 95% (see also talk by Neil Stringfellow on Thursday, 1:30 pm)

Swiss Weather Forecast

- The current 7km forecast suite (1h time-to-solution) was moved from a 16 processor NEC SX-5 to 100 cores of Cray XT3 in January 2007
- The future 2km forecast suite (11 min time-to-solution) will be run on a 5 cabinet Cray XT4 from June 2007 onwards

-> Talk by Tricia Balle on Thursday, 2:30 pm

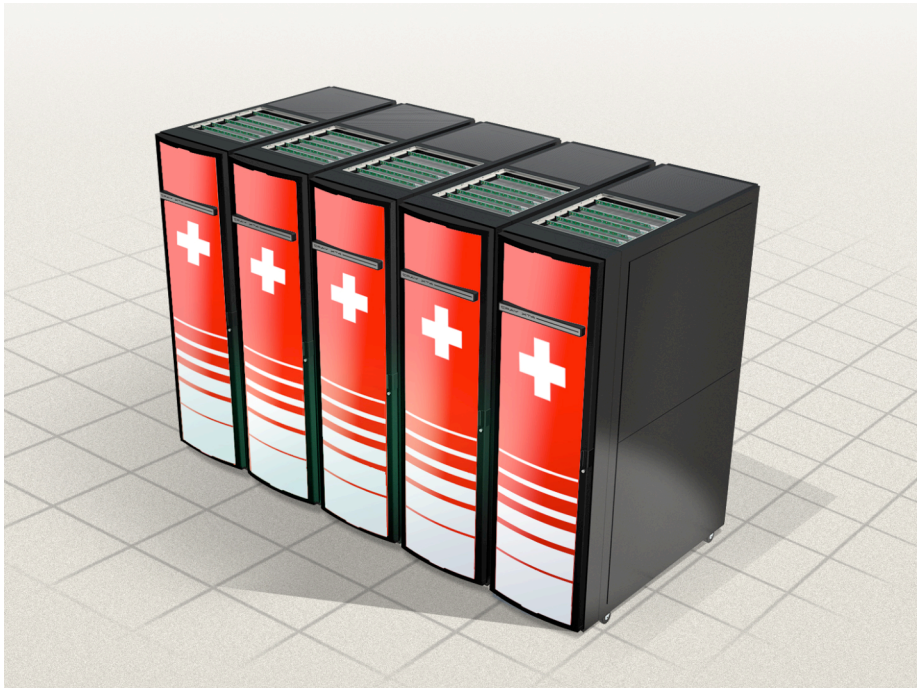

Stabilisation of the Cray XT3

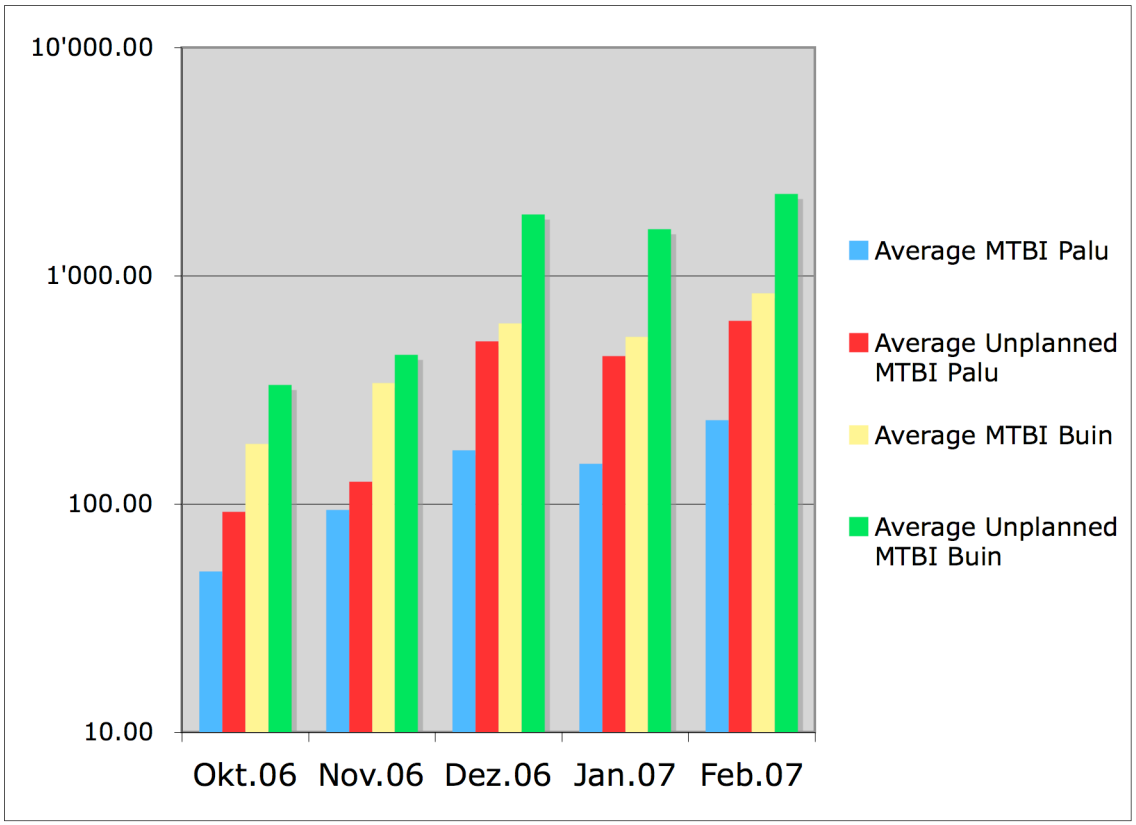

Changing Behaviour

Job size distribution on Cray XT3, Jan 06

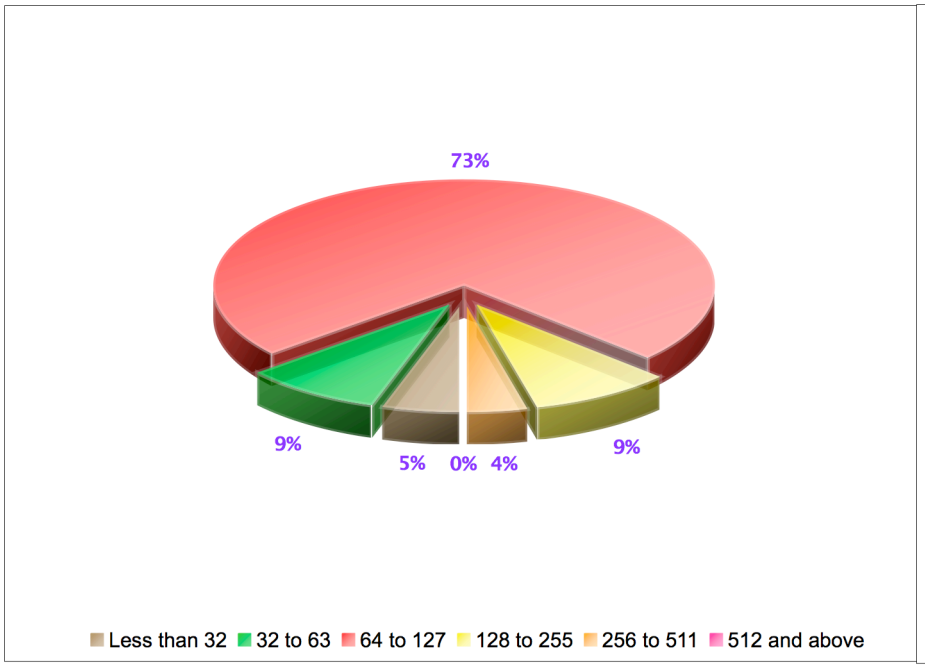

Job size distribution on Cray XT3, Sep 06

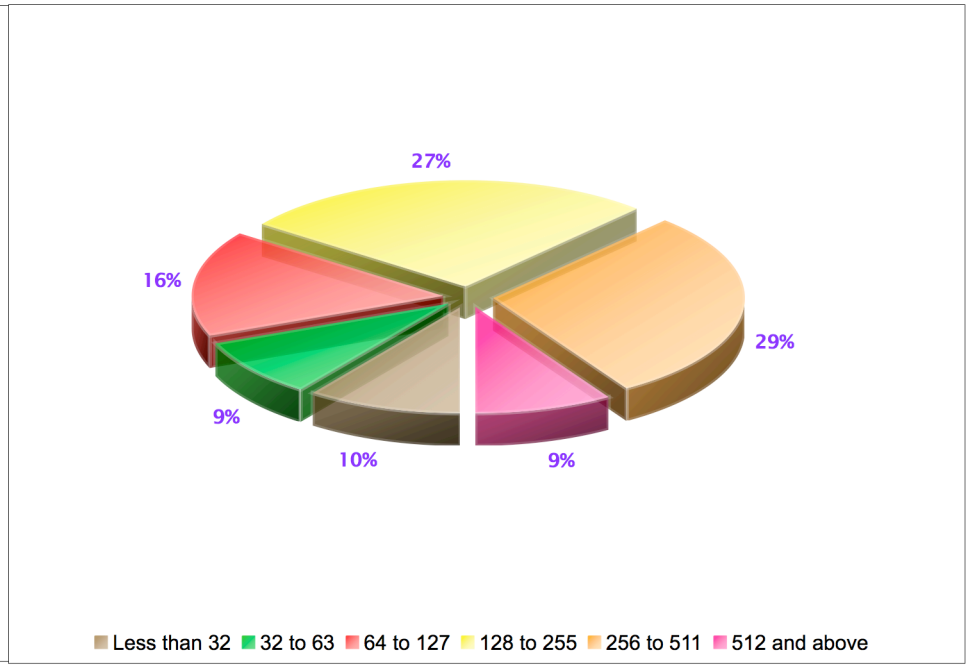

Changing mindset

Job size distribution on IBM SP4, Jan 06

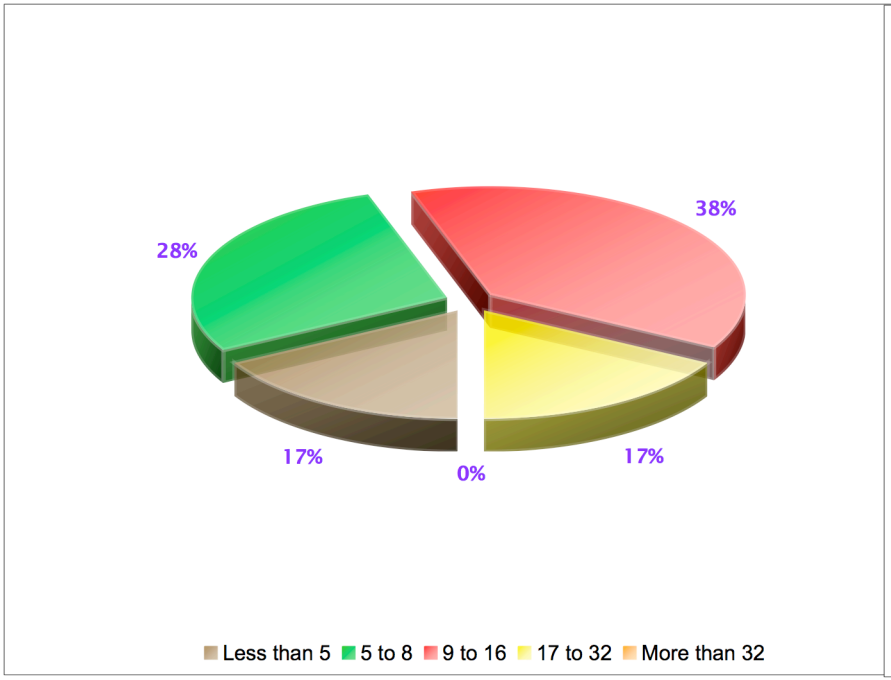

Job size distribution on IBM SP4, Sep 06

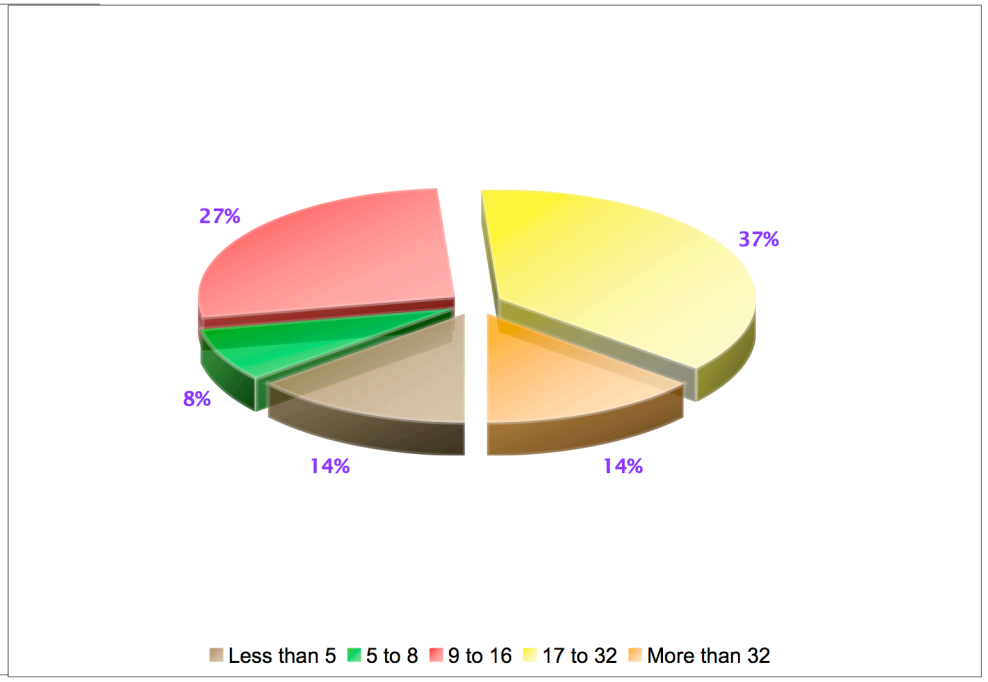

Success of the Cray XT3 introduction

- CSCS made a major step towards deployment and operation of MPP computing facilities
- Users adopted the new technology and a new approach of doing computational research
- Demand for the new machine was immediately very high
  - Up to 25x more than available before scientific review
  - Around 3x to 5x more than available after review
- Do we cut back excellent science too much because of the amount of requests?

Swiss ALPS Programme

- The Swiss ALPS scheme is the new initiative to allow CSCS to engage more closely with scientists and to support very large computational research projects
  - ALPS stands for “Advanced Large Projects in Supercomputing”
  - Targeted at “breakthrough” science
  - Projects are either completely funded or not funded at all
  - Projects must have a collaborative component between CSCS and the research group
  - CSCS provides a key account/project office for the project
- Initially approximately 8.5 million CPU hours of compute time were set aside for this call for proposals
  - These were the allocatable CPU hours over 2 years from the 564 processor expansion of the Cray XT3 in August 2006
- Actual allocation was 16 million CPU hours over 2 years for four projects
  - Projects were only fully funded once the dual-core upgrade was installed in April 2007

Andrew Jackson: “Convection and Magnetic Field Generation in Earth and Other Planets” (1/2)

- 3,500,000 CPU hours for
  - Modelling the dynamics where the Coriolis forces are large enough compared to the viscous forces that they resemble the situation in the Earth’s core
  - Simulation of thermal convection in gas giants and how this leads to the observed surface cloud features and zonal wind pattern
  - Carrying out the first simulations to account for the oblate geometries of the gas giants as opposed to simulating these planets as purely spherical objects.
  - Performing calculations to gain insights into magnetic polarity reversals.

Andrew Jackson: “Convection and Magnetic Field Generation in Earth and Other Planets” (2/2)

- Home-grown codes scaling up to 1’000 processors or more
- Very large storage requirements
  - 200 Terabytes of offline storage
  - From 9 to 41 Terabytes of online storage, to be accessed for pre- and post-processing as well as visualisation
  - Initialisation of a centre-wide parallel file system project at CSCS
- Optimisation of I/O strategies

Michele Parrinello: “Modelling Protein-Protein Interactions at the Atomic Level” (1/2)

- 4,000,000 CPU hours for simulating molecular dynamics of proteins for
  - Amyloid fibrils, which aggregate in the brain of Alzheimer patients
  - Cyclin-dependent kinases, which are related to cell replication and therefore, to uncontrolled multiplication of cancer cells, as well as also being implicated in Alzheimer’s disease.
  - HIV protease and the interaction with local elementary structures, in order to gain a better insight into methods of developing protease inhibitors - a key for developing drugs against HIV

Michele Parrinello: “Modelling Protein-Protein Interactions at the Atomic Level” (2/2)

- Simulations carried out with two community codes: NAMD and ORAC
- Development of of a robust, asynchronous parallelisation scheme for multiple instances of NAMD
- improvements to the scalability of the ORAC code

Christoph Schär: “Climate Change and the Hydrological Cycle from Global to European/Alpine Scales” (1/2)

- 3,750,000 CPU hours for coupling climate and weather simulations of Europe from 1950 to 2050 for
  - Detailed inclusion of aerosol processes in the climate simulations
  - The use of very-high resolution climate change experiments (down to 2km)
  - The combination of global and regional/local modelling frameworks

Christoph Schär: “Climate Change and the Hydrological Cycle from Global to European/Alpine Scales” (2/2)

- Interleaving of a climate model, using the code ECHAM-HAM ported by Mark Cheeseman from CSCS to the Cray XT3 platform, and a weather model, based on CLM (derivate of the German weather forecasting code LM, also used for Swiss forecasting)
- 185 TB offline storage
- 5 TB online storage
- Focus on I/O optimisation strategies (use of NetCDF?) and improved scaling

Viola Vogel: “Towards Simulating a Cell Adhesion Site at Ångström Resolution” (1/2)

- 4,800,000 CPU hours for investigating
  - the mechanisms by which force is transmitted in force-bearing proteins
  - how force might change the structure and function of proteins at cell surfaces

Viola Vogel: “Towards Simulating a Cell Adhesion Site at Ångström Resolution” (2/2)

- Molecular dynamics simulations based on NAMD
- Combination of computational and laboratory experiments
  - Investigations of the mechanisms involved in activating these adhesive molecules is greatly facilitated by scientific visualisation
  - CSCS becomes involved in further developments around VMD

Current status

- Comparison study of single and dual core performance
  - Very good speed-up for molecular dynamics codes (1.98 for NAMD, others > 1.8)
  - Investigation of the new XT4 Meteo system for the climatology project and the geomagnetics project
- Centre-wide parallel file system project started

Conclusion and outlook

- The installation of the Cray XT3 enabled CSCS to support large-scale numerical experiments
- Working more closely with scientists from these larger projects will give CSCS advantages for the future
  - Challenges faced with these projects will serve as input for future procurement decisions and for infrastructure development
  - Problem sizes ran from these projects may be the future problem sizes tackled by many users from our other schemes
- Code development will benefit a wider community
  - Direct benefits from community code developments
  - Indirect benefits from extra skills gained by CSCS staff

CSCS

Swiss National Supercomputing Centre