

Press release: "PRACE evaluates Technologies for Future Multi-Petaflop/s Systems"

For release: 30/03/2009

## PRACE evaluates Technologies for Future Multi-Petaflop/s Systems

*PRACE, the Partnership for Advanced Computing in Europe, has selected a range of promising system and component prototypes for multi-Petaflop/s class systems to be deployed beyond 2010. Prototypes will be installed at eight partner sites starting in 2009.*

In 2008 PRACE selected a set of promising architectures for Petaflop/s class systems to be deployed in 2009/2010. These production system prototypes are being installed at six partner sites in the end of 2008 and the beginning of 2009. To evaluate promising technologies for future Petaflop/s systems beyond 2010 PRACE will deploy a second set of prototypes. This selection was guided by the translation of user and application requirements to hardware and software components and prototypes.

The new prototypes will be installed at the following PRACE partner sites:

**CINECA** (HPC consortium of 36 universities, Italy) will address metadata performance of I/O subsystem solutions for petascale machines on an HP XC4 prototype and study NFS and pNFS over RDMA. The usage of SSD based technology to accelerate metadata performance will be tested with Lustre, NFS and pNFS and compared with traditional disks technology based solutions.

**EPSRC-EPCC** (Edinburgh Parallel Computing Centre, UK) within the FPGA High Performance Computing Alliance (FHPCA) will evaluate porting efforts and the ratio of performance versus power consumption of PRACE benchmarks on their "Maxwell" FPGA prototype supercomputer.

**ETHZ-CSCS** (Swiss National Supercomputing Centre, Switzerland) will study new parallel programming paradigms like the Partitioned Global Address Space (PGAS) programming approaches (Co-Array Fortran, UPC) and the upcoming DARPA High Productivity Computer System (HPCS) language (like Cray's Chapel) on a 3328 cores Cray XT3 system.

**FZJ** (Forschungszentrum Jülich, Germany) will provide a power efficient special-purpose architecture called eQPACE for lattice Quantum Chromodynamics (QCD). This 25.6 TFlop/s peak performance prototype is based on IBM PowerXCell 8i processors and a custom 3d-torus interconnect implemented within FPGAs supporting presently only nearest-neighbor communication. One of the main goals will be to extend the concept to general all-to-all communication.

**GENCI-CEA** (Commissariat à l'Energie Atomique, France) offers a hybrid system composed by nVIDIA Tesla S1070 coupled with BULL Novascale R425 systems. The purpose of this prototype will

[www.prace-project.eu](http://www.prace-project.eu)

### Contact

Herbert Huber, LRZ/GAUSS  
e-mail: herbert.huber@lrz.de



be to evaluate different programming environments like CUDA, HMPP (from CAPS Entreprise), OpenCL and the GPU aware version of Alinea's DDT debugger.

**BAdW-LRZ** (Leibniz Supercomputing Centre, Germany) will assess the productivity of the new data stream parallel programming language RapidMind on x86 multicore systems and multiple accelerators (nVIDIA and AMD/ATI GPUs, IBM Cell and Intel Larrabee).

**BAdW-LRZ and GENCI-CINES** (Centre Informatique National de l'Enseignement Supérieur, France) will install a joint prototype in Garching and Montpellier based on SGI thin nodes (ICE system) and fat nodes (UltraViolet) coupled with ClearSpeed (e710 boards) and Intel Larrabee GPUs. The two partners are planning to evaluate novel hardware (Intel Nehalem-EP/EX processors, NUMalink5 and 4X QDR Infiniband networks, ClearSpeed accelerators and Intel Larrabee manycore GPUs) as well as software components (Lustre filesystem) on synthetic benchmarks and real applications.

**NCF-SARA** (Stichting Academisch Rekencentrum Amsterdam, Netherlands) will install a compact hybrid system composed by 12 standard Intel compute nodes coupled with 12 ClearSpeed CATS 700 units primarily dedicated to the evaluation of large-scale applications (astrophysics, iterative solvers, geophysics and image analysis). Comparisons with GPU based versions of some applications are planned as well as collaboration with the joint BAdW-LRZ / GENCI-CINES prototype.

**Notes for editors:** In computing Flop/s is an acronym meaning Floating point Operations per Second. The Flop/s is a measure of a computer's performance. One PetaFlop/s means computer performance of  $10^{15}$  Flop/s.

NFS means Network File System, pNFS means parallel Network File System, RDMA means Remote Data Memory Access, SSD means Solid State Drive, FPGA means Field Programmable Gate Array, DARPA means Defense Advanced Research Projects Agency, HMPP means Hybrid Multicore Parallel Programming, DDT means Distributed Debugging Tool, GPU means Graphics Processing Unit, CUDA means Compute Unified Device Architecture, OpenCL means Open Computing Language and QDR means Quad Data Rate.

**More information:** Contact, Herbert Huber from LRZ/GAUSS, the PRACE work package leader on Future Technologies. E-mail: [herbert.huber@lrz.de](mailto:herbert.huber@lrz.de)

**About PRACE:** The Partnership for Advanced Computing in Europe (PRACE) prepares the creation of a persistent pan-European HPC service, consisting of several tier-0 centres providing European researchers with access to capability computers and forming the top level of the European HPC ecosystem. PRACE is a project funded in part by the EU's 7th Framework Programme. The PRACE project receives funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° RI-211528.

[www.prace-project.eu](http://www.prace-project.eu)

#### Contact

Herbert Huber, LRZ/GAUSS  
e-mail: [herbert.huber@lrz.de](mailto:herbert.huber@lrz.de)

