The PRACE Project and using the high performance computing for large Monte Carlo simulations *Krassimir Georgiev, IICT-BAS and NCSA, Bulgaria*









Overview (1 of 8)

Building a world-class pan-European High Performance Computing (HPC) Service is a highly ambitious undertaking that involves governments, funding agencies, centers capable to host and manage the supercomputers, and the scientific and industrial user communities with leading edge applications.



Overview (2 of 8)

In contrast to Research Infrastructures that focus on a single scientific instrument an HPC Infrastructure has two unique characteristics: supercomputers serve all scientific disciplines and tier-0 supercomputers have a three year depreciation cycle as tier-0 implies leading edge services.



Overview (3 of 8)

In the roadmap published in 2006, the European Strategy Forum on Research Infrastructures (ESFRI) identified HPC as a strategic priority for Europe – an area where for the last decade Europe has taken a back seat to the United States and Japan.

Since Europe has no significant HPC hardware industry, the process of nurturing European research is based on a coherent integration of infrastructures.



Overview (4 of 8)

PRACE is creating a persistent pan-European High
Performance Computing RI and related services.
Four nations (France, Germany, Italy and Spain) have
agreed to provide 400 million Euro to implement
supercomputers with a combined computing power in
the multi Petaflop/s range over the next 5 years.



Overview (5 of 8)

This funding is complemented by up to 70 million Euros from the European Commission which is supporting the preparation and implementation of this infrastructure. These leadership class systems will help European scientists and engineers to remain internationally competitive.



Overview (6 of 8)

PRACE will maintain a pan-European HPC service consisting of up to six top of the line leadership systems (Tier-0) well integrated into the European HPC ecosystem. Each system will provide computing power of several Petaflop/s (one quadrillion operations per second) in midterm^{*}. On the longer term (2019) Exaflop/s (one quintillion) computing power will be targeted by PRACE.



Overview (7 of 8)

Users are supported by experts in porting, scaling, and optimizing applications to novel, highly parallel computer architectures. An in-depth training program accompanies the PRACE offering teaching scientists and students how to best exploit the unprecedented capabilities of the systems.



Overview (8 of 8)

A scientific steering committee provide advice to PRACE and operate alongside a bespoke peer review process through which access to the Tier-0 resources will be granted based on scientific excellence.

Supercomputing Drives Science through Simulation





Member States

PRACE consists of 20 members, representing Austria, Bulgaria, Cyprus, Czech Republic, Finland, France, Germany, Greece, <u>Hungary</u>, Ireland, Italy, The Netherlands, Norway, Poland, Portugal, Serbia, Spain, Sweden, Switzerland, Turkey and the UK. Partnership For Advanced Computing IN Europe

Member States in more details (1 of 3)

- <u>Austria</u> (JKU University Linz, Institute for Computer Architecture)
- <u>Bulgaria</u> (NCSA National Center for Supercomputing Applications)
- <u>Cyprus</u> (CaSToRC Computation-based Science and Technology Research Center)
- <u>Czech Republic</u> (VŠB Technical University of Ostrava)
- Finland (CSC IT Center for Science)
- <u>France</u> (GENCI Grand Equipement national pour le Calcul Intensif)
- <u>Germany</u> (GCS GAUSS Centre for Supercomputing)
- Greece (GRNET Greek Research and Technology Network)

Member States in more details (2 of 3)

- <u>Hungary</u> (NIIFI National Information Infrastructure Development Institute)
- Ireland (ICHEC Irish Centre for High-End Computing)
- <u>Italy</u> (CINECA Consorzio Interuniversitario)
- <u>The Netherlands</u> (NCF Netherlands Computing Facilities Foundation)
- <u>Norway</u> (SIGMA UNINETT Sigma AS The Norwegian Metacenter for Computational Science)
- <u>Poland</u> (PSNC Poznan Supercomputing and Networking Center)
- <u>Portugal</u> (FCTUC Faculdade Ciencias e Tecnologia da Universidade de Coimbra)
- <u>Serbia</u> (IPB Institute of Physics, Belgrade)

Partnership for Advanced Computing IN Europe

Member States in more details (3 of 3)

- <u>Spain (BSC Barcelona Supercomputing Center Centro</u> <u>Nacional de Supercomputación</u>)
- <u>Sweden</u> (SNIC Swedish National Infrastructure for Computing)
- <u>Switzerland</u> (ETH Zurich Swiss Federal Institute of Technology Zurich, CSCS – Swiss National Supercomputing Centre)
- <u>Turkey</u> (UYBHM National Center for High Performance Computing),
- <u>UK</u> (EPSRC Engineering and Physical Sciences Research Council)

www.prace-project.eu

Partnership For Advanced Computing IN Europe

Three pillars of the PRACE Mission

HPC for Science:

Implement the European Strategy Forum on Research Infrastructures (ESFRI) vision of an European HPC-service at the top of an HPC provisioning pyramid for the scientific community of Europe.

HPC for Industry:

Guaranteeing an European independent access to HPCcompetence for science and industry

HPC by Vendors:

Helping European (hard- & software) vendors to foster their technology HPC-competence



European computer service model



A European model of a sustainable high performance ecosystem consists of a small number of supercomputer centres offering computing service at the highest performance level; national and regional centres with supercomputers offering a the performance to run most of the advanced computing; and the local computing centres in universities, research labs or in other organizations strengthening software development and researchers' competence in computational science.

The PRACE Tier – 0 systems (1)



The Gauss Center for Supercomputing IBM Blue Gene/P system JUGENE

(Juelicher Blue Gene/P) at Forschungzentrum Juelich (FZJ), Jülich, Germany is the first PRACE system and has been available since July 1, 2010 for successful PRACE resource application.

JUGENE has a peak performance of 1 Petaflop. It is composed of 294,912 processing cores with 4 cores forming a node with 2 GB of memory for a total of 147 TB

The PRACE Tier – 0 systems (2)

CURIE, the French Tier-0 system, funded by GENCI, installed at CEA, Bruyères-Le-Châtel, France, designed and built by BULL, is the second PRACE system.

CURIE is based on x86 architecture CPUs with a mix of thin and fat nodes interconnected through a QDR Infiniband interconnect. It has a total of 92,160 processing cores with 4 GB/core and 360 TBytes in total.

The peak performance of the fat nodes partition is 105 Teraflops and the total peak performance (thin nodes and fat nodes) is 1.6 Petaflops.

http://www-hpc.cea.fr/en/complexe/tgcc-curie.htm

The PRACE Tier – 0 systems (3.1)



The Gauss Center for Supercomputing CRAY XE6 system HERMIT at Hoechstleistungsrechenzentrum Stuttgart (HLRS), Stuttgart, Germany is the third PRACE system and will be available in November 2011 for successful PRACE resource applications. As of now a test- and development system is available.



The PRACE Tier – 0 systems (3.2)

HERMIT has a peak performance of 1 Petaflop and is designed for sustained application performance and highly scalable applications. It is composed of 3552 dual socket nodes equipped with AMD Interlagos Processors leading to overall 113664 processing cores. Nodes are equipped with 32GB or 64GB main memory. The pre- and postprocessing infrastructure aims to support users with complex workflows and with advanced access methods including remote graphics rendering and simulation steering.

http://www.hlrs.de/systems/platforms/cray-xe6-hermit



The PRACE Tier-1 systems (1)

Seventeen different **Tier-1** systems are also available for European researchers through PRACE calls.

The Tier-1 resources were previously provided by DEISA (Distributed European Infrastructure for Supercomputing Applications) in so called DECI (DEISA Extreme Computing Initiative) calls, now part of PRACE.



The PRACE Tier-1 systems (2)

- Cray XT4/5/6 and Cray XE6 three large Cray XE and XT systems at EPCC (UK), KTH (Sweden) and CSC (Finland)
- IBM Blue Gene/P three BG/P systems at IDRIS (France), RZG (Germany) and NCSA (Bulgaria).
- IBM Power 6 three IBM Power 6 systems at RZG (Germany), SARA (The Netherlands) and CINECA (Italy).

The PRACE Tier-1 systems (3)

Eight clusters at:

- FZJ (Germany, Bull Nehalem cluster)
- LRZ (Germany, Xeon cluster)
- **HLRS** (Germany, NEC Nehalem cluster + GP/GPU cluster)
- **CINES** (France, SGI ICE 8200)
- **BSC** (Spain, IBM PowerPC)
- **CINECA** (Italy, Westmere plus GP/GPU cluster)
- **PSNC** (**Poland**, Bullx plus GP/GPU cluster and HP cluster)
- **ICHEC** (Ireland, SGI ICE 8200).

The PRACE benchmark suite (1 of 3)

- QCD
- NAMD
- CPMD
- Code Saturne **dynamics**
- GADGET
- TORB/EUTERPE
- WRF

- particle physics
- Quantum Espresso computational chemistry
 - computational chemistry
 - computational chemistry
 - computational fluid
 - astronomy and cosmology
 - plasma physics
 - atmospheric modelling

The PRACE benchmark suite (2 of 3)

- **GROMACS**
- NS3D
- AVBP
- HELIUM
- TRIPOLI-4
- **PEPC**
- **GPAW**

- computational chemistry
- computational fluid dynamics
- computational fluid dynamics
- computational physics
- computational engineering
- plasma physics
- computational chemistry

The PRACE benchmark suite (3 of 3)

- ELMER
- SPECFEM3D
- NEMO
- **CP2K**
- OCTOPUS
- BSIT
- ALYA

- computational engineering
- computational geophysics
- ocean modelling
- computational chemistry
- computational chemistry
- computational geophysics
- computational mechanics

The PRACE IBM Blue Gene/P computer NCSA, Sofia, Bulgaria

Consists of two racks, 2048 Power PC 450 based compute nodes, 8192 processor cores and a total of 4TB random access memory. Each processor core has a double precision, dual pipe floating point core accelerator.

Sixteen I/O nodes are connected via fibre optics to a 10 Gb/s Ethernet switch. The smallest partition size, available currently, is 128 compute nodes (512 processor cores).

The theoretical performance of the computer is 27.85 Tflops while the maximum LINPACK performance achieved is 23.42 Tlops (\approx 84%).

http://www.bgsc.acad.bg/



Application software and libraries on IBM Blue Gene/P computer, NCSA, Sofia (1 of 3)

Molecular dynamics

GROMACS 4.0.5, NAMD 2.7b2, LAMMPS 05012010, DL POLY, MMFF94

Quantum and computational chemistry

CPMD 3.13.2, CP2K, NWChem, Quantum Espresso, GAMESS, DALTON, Qbox

In silico drug design

DOCK 6.4, ROSETTA 3.1, ROSETTA 2.3



Application software and libraries on IBM Blue Gene/P computer, NCSA, Sofia (2 of 3)

<u>Seismic Analysis</u> SPECFEM3D

<u>Fluid dynamics, thermodynamics and hydrodynamics</u> Code Saturne 1.3.2, Code Aster, Code Salome, Salome-Meca platform PLANETICS

<u>Multiphysical models</u> ELMER



Application software and libraries on IBM Blue Gene/P computer, NCSA, Sofia (3 of 3)

<u>Libraries</u>

PETCs, ParMETIS, LAPACK, ScalaBLAST, mpiBLAST, MUMS, SuperLU, FFTW, HYPRE, Trilinos, ParFE, GotoBLAS



Calls for Proposals (1 of 3)

PRACE resources are available through three forms of access: *Programme Access*, *Project Access* and *Preparatory access*

<u>Programme Access</u> is available to major Europen projects and infrastructures that can benefit from PRACE resources and for which Project access is not appropriate



Calls for Proposals (2 of 3)

<u>Project Access</u> is intended for individual researchers and research groups including multi-national research groups and has a one year duration. Calls for Proposals for Project Access are issued twice yearly (May and November).

<u>Preparatory Access</u> is intended for resource use required to prepare proposals for Project Access. Applications for Preparatory Access are accepted at any time. Partnership for Advanced Computing IN Europe

Calls for Proposals (3 of 3)

Programme and **Project Access** are subject to a PRACE peer review process. This process includes a <u>technical</u> <u>evaluation</u> followed by a <u>scientific peer review</u>. The Peer Review process, managed by PRACE staff, is overseen by the PRACE Scientific Steering Committee composed of leading European scientists appointed by the PRACE Council.

Preparatory Access requests are only subject to a technical evaluation.



Open Calls

- PRACE Preparatory Access Continuously open for proposals. Next cutoff dates: 5 May 2011; 1 July 2011; 1 September 2011; 1 December 2011; 1 March 2012; 1 June 2012
- It is strongly advised that applicants read the information on PRACE Resource Access before applying for PRACE resources.



PRACE Events

The PRACE Autumn School 2011 will take place, from the 25th to the 27th of October 2011, in France at the Très Grand Centre de Calcul (TGCC) operated by CEA in Bruyères-le-Châtel (near Paris) and where the French CURIE Tier0 system is installed.

This summer, CURIE will be upgraded with a 1.5 Pflop/s partition of x86 thin nodes and a 200 Tflop/s hybrid partition based on the latest NVIDIA M2090 GPUs, leading to a global peak performance of more than 1.8 Pflop/s.

This hybrid fraction, opened to European scientists through preparatory access calls, will be the first hybrid system available into the PRACE RI.

Co-organized by CEA and GENCI, the PRACE Autumn School will focus on advanced hybrid programming, benefiting from the brand new CURIE hybrid fraction for online practicing. For this three-day event, advanced courses on OpenCL and HMPP programming as well as profiling and debugging tools will be delivered.