KEYNOTERS – HPCS 2012

Tuesday Keynote: Europe Back in the HPC Race: Building a European Ecosystem to Recover and Maintain the Capacity of Designing and Building Large Computers

Jean Gonnord
Commissariat à l'Energie Atomique (CEA), France

Wednesday Keynote I: The Blue Brain Project: Challenges and Opportunities for HPC

Felix Schürmann
Blue Brain Project, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

Wednesday Keynote II: Modelling, Simulation and Analytics in the Exascale Era

James C. Sexton
Computational Science Center, IBM T. J. Watson Research Center, Yorktown Heights, NY, USA

Thursday Keynote: Productivity and Performance: The OmpSs Programming Model and Environment

Jesus Labarta
Barcelona Supercomputing Center & Technical University of Catalonia (UPC), Spain

---

Tuesday Keynote: Europe Back in the HPC Race: Building a European Ecosystem to Recover and Maintain the Capacity of Designing and Building Large Computers

Jean Gonnord
Commissariat à l'Energie Atomique (CEA), France

ABSTRACT

High Power Computing is unanimously recognized today as strategic for research, industry, society and defence. After a long absence, Europe has set up in 2008, as part of PCRD7, the PRACE project to give European research world class computing access. CEA, which has always been a major actor in this field, has a more global vision covering from the multiple usage of HPC to the development of the necessary technology. To support its strategy, CEA established a scientific computing complex, one of the largest in the world with two pétaflop/s machines, set up a research platform on machine architectures, parallelism, and software environment shared with industrial and academic partners, and shows an ambitious roadmap to the exaflop/s. The last result of CEA strategy is the BULL TERA100 machine, the first pétaflop/s computer ever designed and built in Europe. The TERA100 machine, followed by the Curie machine of the PRACE project, the HELIOS machine of the Fusion program, and a lot of industrial success are demonstrating the capability of Europe to come back on this challenging market. With this background the European Commission has called in 2012, as part of its HORIZON 2020 program, for an ambitious project:

- Provide a world-class European HPC infrastructure, benefitting a broad range of academic and industry users, and especially SMEs, including a workforce well trained in HPC;
- Ensure independent access to HPC technologies, systems and services for the EU.

The talk will cover what has been done in Europe in this period and what is in preparation to answer the European Commission roadmap.

Wednesday Keynote I: The Blue Brain Project: Challenges and Opportunities for HPC

Felix Schürmann

Blue Brain Project, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

ABSTRACT

Many areas of science and engineering have adopted simulation-based research as a novel tool for discovery and insight. The sustained performance growth in supercomputer performance allows ever more detailed models, which makes supercomputing nowadays also a viable tool for biology. However, the heterogeneity of biological systems challenges many aspects of supercomputing: intricate workflows are required for model generation, mathematical formulations are volatile, and memory requirements are demanding. At the same time, the weak scaling properties of many biological systems are enormous and therefore are a good match for today's massive parallelism, whereas the multiple time scales inherent to biological systems requires out-of-the-box thinking. On the example of the Swiss Blue Brain Project computational workflows and high performance computing challenges are discussed for data-driven simulation-based research in neuroscience. The talk will conclude on opportunities for HPC that can arise from brain simulations.

Wednesday Keynote II: Modelling, Simulation and Analytics in the Exascale Era

James C. Sexton

Computational Science Center, IBM T. J. Watson Research Center, Yorktown Heights, NY, USA

ABSTRACT

As the world HPC community launches into development of Exascale systems over the next 10 years, a critical challenge will be to develop new classes and capabilities in software and applications to support the integrated modeling, simulation and analytics which those systems enable. An essential impact that is both enabled and required will be to develop integrated services and solutions which deploy modeling, simulation and analytics for research, industrial and commercial benefit. This presentation describes some of the key opportunities and challenges to be addressed and discusses some of the activities in which IBM Research is engaging to support aggressive innovation in and impact of the use of High Performance Computing in the Exascale Era.

Thursday Keynote: Productivity and Performance: The OmpSs Programming Model and Environment

Jesus Labarta

Barcelona Supercomputing Center & Technical University of Catalonia (UPC), Spain

ABSTRACT

The talk will address the issues of programmability and performance in high performance computing and in particular how to achieve a good balance between them. The OmpSs programming model and its associated development and analysis environment constitutes a pragmatic approach opening a huge set of possibilities. OmpSs aims at supporting higher levels of asynchrony and overlap than most widely used models while providing an easy incremental path for porting existing codes. The model addresses the limitations that current hybrid or accelerated approaches have when addressing large-scale parallel execution but it also aims at general purpose multicore programming. OmpSs is an extension of OpenMP where directionality clauses provide with information on the access pattern which is used by the runtime to dynamically detect and exploit parallelism as well as to automatically manage data and perform locality optimizations. The objective is decoupling the way programs are written, focusing on the specification of algorithms and computations to be performed, from how they are executed, dealing with actual resources used and order of computations. This separation of responsibilities offers an homogeneous vie at heterogeneity where a single source code can be mapped to the available target machine by an intelligent runtime. A powerful prediction and analysis framework is needed to help the programmer understand the potential and exploit the huge capabilities that the model offers.

The talk will describe the OmpSs programming model, the issues faced when introducing high levels of asynchrony, the development environment and examples of large-scale scientific and non-numeric applications where the model is being used.