Topic 15: High Performance and Scientific Applications

Thomas Ludwig, Costas Bekas, Alice Koniges, and Kengo Nakajima

Topic Committee

Many fields of science and engineering are characterized by an increasing demand for computational resources. Coupled with important algorithmic advances, high performance computing allows for the gain of new results and insights by utilizing powerful computers and big storage systems. Workflows in science and engineering produce huge amounts of data through numerical simulations and derive new knowledge by a subsequent analysis of this data. Progress in these fields depends on the availability of HPC environments, from medium sized teraflops systems up to leading petaflops systems.

The High Performance and Scientific Applications Topic highlights recent progress in the use of high performance parallel computing and puts an emphasis on success stories, advances of the state-of-the-art and lessons learned in the development and deployment of novel scientific and engineering applications.

Today's complex research issues have to be mapped onto complex compute and storage environments: We find powerful computers being composed of hundreds and thousands of nodes which themselves are shared memory parallel computers with many processor cores and sometimes additional accelerator hardware. Storage hardware consists at least of high volume disk systems but could also comprise tape libraries. With the advent of the Exascale era a way to increased resource usage might enforce hardware-software co-design.

The papers accepted for this workshop characterize typical steps on the way to exploit maximum system performance. Four papers focus on an optimal adaptation of the software system onto the available hardware system. They deal with the layout of data structures in memory and the mapping of software components onto hardware in multi-core environments. One contribution analyses how data volumes could be reduced when a certain loss of quality is acceptable. Finally, we include an interesting report on how quality of data can be assessed and how it influences subsequent post-processing. We invite you to read this collection of papers and get inspired by the results of our colleagues.

Kunaseth et al. present in "Memory-Access Optimization of Parallel Molecular Dynamics Simulation via Dynamic Data Reordering", a novel data-reordering scheme aimed to optimize runtime memory access in the context of large scale molecular dynamics simulations.

Reiter et al. present in "On Analyzing Quality of Data Influences on Performance of Finite Elements driven Computational Simulations" a thorough analysis of mechanisms with which data quality can dramatically influence the results as well as the performance of scientific simulations.

C. Kaklamanis et al. (Eds.): Euro-Par 2012, LNCS 7484, pp. 779–780, 2012.

Malakar et al. demonstrate in “Performance Evaluation and Optimization of Nested High Resolution Weather Simulations” a significant reduction in run time of complex climate simulations by exploiting a careful combination of compiler optimizations coupled with overlapping computation and communication.

Fietz et al. present in “Optimized Hybrid Parallel Lattice Boltzmann Fluid Flow Simulations on Complex Geometries” an optimized hybrid parallelization strategy, that is capable of solving large-scale fluid flow problems on complex computational domains.

Aktulga et al. quantitatively show in “Topology-aware Mappings for Large-Scale Eigenvalue Problems” that topology-aware mapping of processes to physical processors can have a significant impact on the efficiency of high-performance computing applications, with a particular view to modern large-scale multi-core architectures.

Finally, Iverson et al. developed in “Fast and Effective Lossy Compression Algorithms for Scientific Datasets” effective and efficient algorithms for compressing scientific simulation data computed on structured and unstructured grids.