Multicore and Manycore Programming

Beniamino Di Martino[1], Fabrizio Petrini[1], Siegfried Benkner[2], Kirk Cameron[2], Dieter Kranzlmüller[2], Jakub Kurzak[2], Davide Pasetto[2], and Jesper Larsson Träff[2]

[1] Topic Chairs
[2] Members

We would like to join the other members of the Program Committee in welcoming you to the Multicore and Manycore Programming Topic of Europar 2010. Europar is one the primary forums where researchers, architects and designers from academia and indutry explore new and emerging technologies in multicore programming and algorithmic development.

This year, we received 43 submissions. Each paper was reviewed by at least three reviewers and we were able to select 17 regular high-quality papers.

The Topic Committee Members handled the paper review process aiming at high-quality and timely review process. Each TPC member was able to handle a high-load, exceeding 20 papers, providing valuable insight and guidance to improve the quality of the scientific contributions.

The accepted papers discuss very interesting issues. In particular, the paper "Parallel Enumeration of Shortest Lattice Vectors" by M. Schneider and O. Dagdelen presents a parallel version of the shortest lattice enumeration algorithm, using multi-core CPU systems. The paper "Exploiting Fine-Grained Parallelism on Cell Processors" by A. Prell, R. Hoffmann and T. Rauber, presents a hierarchically distributed task pool for task parallel programming on Cell processors. The paper "Optimized on-chip-pipelined mergesort on the Cell/B.E." by R. Hulten, C. Kessler and J. Keller works out the technical issues of applying the on-chip pipelining technique to parallel mergesort algorithm for the Cell processor. The paper "A Language-Based Tuning Mechanism for Task and Pipeline Parallelism" by F. Otto, C. A. Schaefer, M. Dempe and W. F. Tichy tackles the issues arising with auto-tuners for parallel applications of requiring several tuning runs to find optimal values for all parameters by introducing a language-based tuning mechanism.

The paper "Near-optimal placement of MPI processes on hierarchical NUMA architecture" by E. Jeannot and G. Mercier describes a novel algorithm called TreeMatch that maps processes to resources in order to reduce the communication cost of the whole application. The paper "Multithreaded Geant4: Semi-Automatic Transformation into Scalable Thread-Parallel Software" by X. Dong, G. Cooperman, J. Apostolakis presents the transformation into scalable thread parallel version of an application case study, Geant4, which is a 750,000 line toolkit first designed in the early 1990s. The paper "Parallel Exact Time Series Motifs Discovery" by A. Narang presents novel parallel algorithms for exact motif discovery on multi-core architectures. The paper "JavaSymphony: A Programming and Execution Environment for Parallel and Distributed Many-core Architectures" by M. Aleem, R. Prodan and T. Fahringer proposes a new Java-based programming

P. D'Ambra, M. Guarracino, and D. Talia (Eds.): Euro-Par 2010, Part II, LNCS 6272, pp. 137–138, 2010.

model for shared memory multi-core parallel computers as an extension to the JavaSymphony distributed programming environment.

The paper "Adaptive Fault Tolerance for Many-Core based Space-Borne Computing" by H. Zima describes an approach for providing software fault tolerance in the context of future deep-space robotic NASA missions, which will require a high degree of autonomy and enhanced on-board computational capabilities, focusing on introspection-based adaptive fault tolerance. The paper "A Parallel GPU Algorithm for Mutual Information based 3D Nonrigid Image Registration" by V. Saxena, J. Rohrer and L. Gong presents parallel design and implementation of 3D non-rigid image registration for the Graphics Processing Units (GPUs). The paper "A Study of a Software Cache Implementation of the OpenMP Memory Model for Multicore and Manycore Architectures" by C. Chen, J. Manzano, G. Gan, G. Gao and V. Sarkar presents an efficient and scalable software cache implementation of OpenMP on multicore and manycore architectures in general, and on the IBM CELL architecture in particular. The paper "Maestro: Data Orchestration for OpenCL Devices" by K. Spafford, J. S Meredith and J. Vetter introduces Maestro, an open source library for automatic data orchestration on OpenCL devices. The paper "Optimized dense matrix multiplication on a manycore architecture" by E. Garcia, I. E. Venetis, R. Khan and G. Gao utilizes dense matrix multiplication as a case of study to present a general methodology to map applications to manycore architectures.

The paper "Multi-GPU and Multi-CPU Parallelization for Interactive Physics Simulations" by E. Hermann, B. Raffin, F. Faure, T. Gautier and J. Allard proposes a parallelization scheme for dynamically balancing work load between multiple CPUs and GPUs. The paper "Long DNA Sequence Comparison on Multicore Architectures" by F. Snchez, F. A. Ramirez and M. Valero analyzes how large scale biology sequences comparison takes advantage of current and future multicore architectures, and investigate which memory organization is more efficient in a multicore environment. The paper "Programming CUDA-based GPUs to simulate two-layer shallow water flows" by M. De la Asuncin, J. Miguel Mantas Ruiz and M. Castro describes an accelerated implementation of a first order well-balanced finite volume scheme for 2D two-layer shallow water systems using GPUs supporting the CUDA programming model and double precision arithmetic. The paper "Scalable Producer-Consumer Pools based on Elimination-Diffraction Trees" by Y. Afek, G. Korland, M. Natanzon and N. Shavit presents new highly distributed pool implementations based on a novel combination combination of the elimination-tree and diracting-tree paradigms. Finally, the paper paper "Productivity and Performance: Improving Consumability of Hardware Transactional Memory through a Real-World Case Study" by H. Wang, G. Yi, Y. Wang and Y. Zou shows how, with well-designed encapsulation, HTM can deliver good consumability for commercial applications.

We would like to take the opportunity of thanking the authors who submitted the contributions, the Euro-Par Chairs Domenico Talia, Pasqua D' Ambra and Mario Guarracino, and the referees with their highly useful comments, whose efforts have made this conference and this topic possible.