High Performance Architectures and Compilers

Pedro C. Diniz[1], Marco Danelutto[1], Denis Barthou[2], Marc Gonzales[2], and Michael Hübner[2]

[1] Topic Chairs
[2] Members

This topic deals with architecture design and compilation for high performance systems. The areas of interest range from microprocessors to large-scale parallel machines; from general-purpose platforms to specialized hardware (e.g., graphic coprocessors, low-power embedded systems); and from hardware design to compiler technology. On the compilation side, topics of interest include programmer productivity issues, concurrent and/or sequential language aspects, program analysis, transformation, automatic discovery and/or management of parallelism at all levels, and the interaction between the compiler and the rest of the system. On the architecture side, the scope spans system architectures, processor micro-architecture, memory hierarchy, and multi-threading, and the impact of emerging trends. All the papers submitted to this track highlight the growing significance of Chip Multi-Processors (CMP) and Simultaneous Multi-Threaded (SMT) processors in contemporary high-performance architectures.

Fourteen papers were submitted to this topic, of which five were accepted for presentation at the conference in two different sessions.

The paper "Power-Efficient Spilling Techniques for chip Multiprocessors", by Enric Herrero, José González and Ramon Canal presents two spilling techniques for improving the efficiency the *N-chance* forwarding mechanism used to improve unused cache space in chip multiprocessors.

The paper "Scalable Object-Aware Hardware Transactional Memory (SO-HTM)" by Behram Khan, Mikel Lujan, Ian Watson and Matthew J. Horsnell reflects a growing trend of hardware support for transactions and advanced programming abstraction for concurrency specification and management. The paper describes the design and evaluation of an object-aware hardware transactional memory system that relies on multicast to allow multiple transactional commits to occur concurrently.

The paper "Extending the Cell SPE with Energy Efficient Branch Prediction" by Martijn Briejer, Cor Meenderinck and Ben Juurlink examines the key issue of the impact of branch prediction in terms of energy used and execution time for a multi-core system, in this particular instance the Cell Synergistic Processing Elements (SPE). The authors describe the application of three dynamic branch prediction schemes with a combination of branch prediction hints and warning instructions and the use of the Branch Target Buffer (BTB) hardware mechanism.

Last but not the least, we have included two companion papers that address the issue of shared caching access latency in multi-core architectures. The first paper "Efficient Address Mapping of Shared Cache for On-Chip Many-Core

P. D'Ambra, M. Guarracino, and D. Talia (Eds.): Euro-Par 2010, Part I, LNCS 6271, pp. 254–255, 2010.
© Springer-Verlag Berlin Heidelberg 2010

Architecture" by Fenglong Song, Zhiyong Liu, Dongrui Fan, Junchao Zhang and Lei Yu, describes the use of a variant of the traditional XOR-based mapping schemes for accessing multi-banked caches and evaluates it using well-known parallel application benchmarks. The second paper "Thread Owned Block Cache: Managing Latency in Many-Core Architecture" by Fenglong Song, Zhiyong Liu, Dongrui Fan, Lei Yu and Shibin Tang, describes and evaluates two techniques for the reduction of conflicts between blocks at the last level of shared caches. Whereas one technique relies on partitioning of the shared cache space such that blocks of different threads are segregated, the second technique relies on hardware to promote caching of thread owned data blocks.

We would like to take this opportunity to thank the authors who submitted a contribution, as well as the Euro-Par Organizing Committee, and the referees with their highly useful comments, whose efforts have made this conference and this topic possible.