Introduction

Mitsuhisa Sato, Denis Barthou, Pedro C. Diniz, and P. Saddayapan

Topic chairs

This topic deals with architecture design and 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; 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, program 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 scalability issues in performance/power in contemporary high-performance architectures.

The paper "Filtering directory lookups in CMPs with write-through caches" by Ana Bosque, Victor Viñals, Pablo Ibañez and Jose Maria Llaberia proposes an architectural enhancement to reduce the number of directory lookups for a directory-based coherence protocol, in CMP shared caches. The authors describe a hardware filter reducing the associativity of the lookups, eliminating these lookups in some cases and reducing power consumption.

The paper "FELI: HW/SW support for On-Chip Distributed Shared Memory in Multicores" by Carlos Villavieja, Yoav Etsion, Alex Ramirez and Nacho Navarro proposes a set of operating system mechanisms to automatically manage memory on a CMP with on-chip scratchpad memories. The authors describe how the virtual memory paging mechanism is leveraged to achieve this and reduce power consumption.

The paper "Token3D: Reducing Temperature in 3D die-stacked CMPs through Cycle-level Power Control Mechanisms" by Juan M. Cebrián, Juan L. Aragón and Stefanos Kaxiras describes a token-based power management algorithm for multi-core architectures organized as stacks of chips. The paper describes how to take into account temperature and layer information when balancing power, giving higher priority to cool cores over hot ones.

The paper "Unified Locality-sensitive Signatures for Transactional Memory" by Ricardo Quislant, Eladio D Gutierrez, Oscar Plata and Emilio Zapata describes a new design for the hardware support of transactional memory. The authors proposes to combine the use of locality-sensitive signatures with unified hash functions for read and write sets.

Last but not least, two papers address prefetching issues with new hardware prefetching schemes. The paper "Bandwidth Constrained Coordinated HW/SW Prefetching For Multicores" by Sai Prashanth Muralidhara, Mahmut

E. Jeannot, R. Namyst, and J. Roman (Eds.): Euro-Par 2011, LNCS 6852, Part I, pp. 267–268, 2011.

Taylan Kandemir and Yuanrui Zhang presents a hierarchical management and bandwidth-constrained prefetching algorithm for multi-cores. The authors describe a prefetching scheme and the metrics used to adjust dynamically the aggressiveness of the prefetch, handling bandwidth contention and performance. The paper "Using runtime activity to dynamically filter out inefficient data prefetches" by Gamoudi, Nathalie Drach and Karine Heydemann describes a hardware prefetching strategy exploiting runtime activity information and a history-based algorithm to filter out inefficient data prefetches. The paper describe a method to correlate runtime activities with prefetching effects in order to increase prefetching efficiency.

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.