

# Topic 4: High-Performance Architectures and Compilers (Introduction)

Denis Barthou, Wolfgang Karl, Ramón Doallo, Evelyn Duesterwald,  
and Sami Yehia

Topic Committee

The topic "High Performance Architectures and Compilers" deals with architecture design and compilation for high performance systems. The areas of interest range from microprocessors to large-scale parallel machines (including multi-core, possibly heterogeneous, architectures); 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, 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.

The papers submitted to this topic were thoroughly reviewed and discussed. For each of the papers we obtained four reviews. We would like to thank all reviewers who helped in this process. Finally, four papers were accepted which are summarized below.

The paper "Adaptive Granularity Control in Task Parallel Programs using Multiversioning" by Peter Thoman, Herbert Jordan, and Thomas Fahringer introduces a method to adapt dynamically the granularity of fine-grained parallel programs. The method has two stages: first a set of versions of the input program are generated by a compiler, basically this compilation relies on the unrolling techniques to generate versions of the code with different granularity, then, the optimal granularity is adapted dynamically at run time using a simple algorithm based in heuristics. The evaluation mostly uses benchmarks from BOTS and compares the approach with Cilk, Intel ICC and GCC with OpenMP. The experimental results show that the method is effective to increase the efficiency of recursive parallel programs.

The paper "Adaptive Snoop Granularity in Hardware Transactional Memory" by Ehsan Atoofian presents an approach to reduce the coherency traffic in a Hardware transactional memory (HTM) system. The idea relies on remembering regions of conflicts in a snoop granularity table and filtering out subsequent snoops if it is known that the region is not conflicting with others. The evaluation shows that this approach is effective and also reduces the energy consumption of the bus.

The paper "Towards Efficient Dynamic LLC Home Bank Mapping with NoC-Level Support" by Mario Lodde, José Flich, and Manuel E. Acacio proposes a new approach for optimizing the use of shared last level caches in tiled CMPs. The principle is to dynamically determine a home bank, taking into account the topology of the CMP and the occupation of each LLC. A migration mechanism is included in order to better locate shared blocks. In the context of tiled CMPs, this addresses an important issue concerning scalability and performance.

Finally, the paper "Online Dynamic Dependence Analysis for Speculative Polyhedral Parallelization" by Alexandra Jimborean, Philippe Clauss, Juan Manuel Martinez, and Aravind Sukumaran-Rajam presents a runtime dependence analysis for speculative parallelization, based on VMAD, a framework for program analysis and instrumentation. The dependence analysis used is a combination of range and GCD test, using dependence distance vector as abstraction.