Int J Parallel Prog (2011) 39:582–583
DOI 10.1007/s10766-011-0166-5

Guest Editorial: Parallel and Distributed Computing

Can Ozturan · Dan Grigoras

Published online: 25 February 2011

Welcome to this special issue on the 9th edition of International Symposium on Parallel and Distributed Computing (ISPDC), which took place in Istanbul in July 2010. ISPDC is dedicated to the advancement of state-of-the-art in the area of parallel and distributed computing. The symposium program featured 28 full papers accepted out of 71 submitted papers. The program covered a wide spectrum of topics including parallel models and algorithms, interconnection topologies, GPU programming, grid, P2P, multicore, network and distributed systems. The three papers in this special issue were selected among the highest rated accepted papers and cover the topics of algorithm design, compiler optimization, and GPU programming. These three papers were extended and revised after a new round of review process before being published in IJPP.

The first paper titled "Optimizing the reliability of streaming applications under throughput constraints" by Benoit et al. received the best paper award in the symposium. The paper considers the problem of mapping pipelined applications consisting of linear chain of stages onto distributed and parallel platforms in the difficult contexts of when multiple optimization criteria are to be achieved and when the target processors and the communication links are heterogeneous and subject to failures. This problem is NP-hard. A set of efficient heuristics is designed for the problem and the performance of interval and general mapping strategies are compared through extensive simulations. In particular, for the homogeneous processor case, an optimal dynamic

C. Ozturan (✉) · D. Grigoras
Department of Computer Engineering, Boğaziçi University, Istanbul, Turkey
e-mail: ozturaca@boun.edu.tr

D. Grigoras
Computer Science Department, University College Cork, Cork, Ireland
e-mail: grigoras@cs.ucc.ie

programming algorithm for the bi-criteria interval mapping problem is contributed by the authors.

The second paper "Resource-Aware Compiler Prefetching for Fine-Grained Many-Cores" by Caragea et al. is about reducing pressure on memory level parallelism, which becomes more critical on many-core systems. The paper contributes a resource-aware prefetching compiler algorithm and demonstrates it in a GCC-derived compiler, while evaluating its performance using XMT fine-grained many-core architecture. A design-space exploration has also been performed for the XMT architecture by varying the amount of chip resources designated for per-core prefetch storage and off-chip bandwidth.

The emergence of research on Graphics Processing Units (GPU) is evidenced by the third paper titled "Speeding Up Cycle Based Logic Simulation using Graphics Processing Units" by Sen et al. GPUs are special-purpose application accelerators that are routinely used in applications other than graphics, which they were originally developed for. The paper presents two parallel algorithms for logic simulation, which is the most commonly used verification technique for digital designs. The digital designs are first partitioned into independent clusters and then merged and balanced in order to make best use of GPU resources. The algorithms exploit the massively parallel GPU architecture featuring thousands of concurrent threads, fast memory, and memory coalescing for optimizations. Speed-ups up-to 21x on several benchmarks hold promise to reduce the overall digital design cycle.