The Nineteenth International Conference on

Parallel Architectures and Compilation Techniques (PACT 2010)

Vienna, Austria, September 11-15, 2010

POSTER PROGRAM

Welcome to the Poster Session

The poster session takes place on Monday, September 13 from 4:00 to 5:30 pm in the Aula. Posters 22 to 34 are from students taking part in the ACM Student Reseach Competition.

1	On-Chip Network Design Considerations for Compute Accelerators *A. Bakhoda, J. Kim, T. Aamodt*
2	Believe It or Not! Multi-core CPUs can Match GPU Performance for a FLOP-intensive Application! *R. Bordawekar, U. Bondhugula, R. Rao*
3	Less is More: Trading off Work-Efficiency for Scalability in Irregular Programs *M. Hassaan, M. Burtscher, K. Pingali*
4	Moths: Mobile Threads for On-Chip Networks *M. Misler, N. Jerger*
5	Improving Speculative Loop Parallelization via Selective Squash and Speculation Reuse *S. Ananthramu, D. Majeti, S. Aggarwal, M. Chaudhuri*
6	Revisiting Sorting for GPGPU Stream Architectures *D. Merrill, A. Grimshaw*
7	Analyzing Cache Performance Bottlenecks of STM applications and addressing them with compiler's help *S. Mannarswamy, R. Govindarajan*
8	An Intra-Tile Cache Set Balancing Scheme *M. Hammoud, S. Cho, R. Melhem*
9	StatCC: A Statistical Cache Contention Model *D. Eklov, D. Black-Schaffer, E. Hagersten*
10	An Integer Programming Framework for Optimizing Shared Memory Use on GPUs *W. Ma, G. Agrawal*
11	Exploiting Subtrace-Level Parallelism in Clustered Processors *R. Ubal, J. Sahuquillo, S. Petit, P. López, J. Duato*
12	A Case for NUMA-aware Contention Management on Multicore Systems *S. Blagodurov, S. Zhuravlev, A. Fedorova, A. Kamali*
13	DMATiler: Revisiting Loop Tiling for Direct Memory Access *H. Lin, T. Liu, L. Renganarayana*
14	Fidelity and Scaling of the PARSEC Benchmark Inputs *C. Bienia, K. Li*
15	Online Cache Modeling for Commodity Multicore Processors *R. West, P. Zaroo, C. Waldspurger, X. Zhang*
16	NoC-Aware Cache Design for Multithreaded Execution on Tiled Chip Multiprocessors *A. Abousamra, A. Jones, R. Melhem*
17	A Software-SVM-based Transactional Memory for Multicore Accelerator Architectures with Local Memory *J. Lee, S. Seo, J. Lee*
18	NUcache – A Multi-core Cache Organization Based on Next-Use Distance *R. Manikantan, K. Rajan, R. Govindarajan*
19	CoreGenesis: Erasing Core Boundaries for Robust and Configurable Performance *S. Gupta, S. Feng, A. Ansari, S. Mahlke*
20	Automatic Vector Instruction Selection for Dynamic Compilation *R. Barik, J. Zhao, V. Sarkar*
21	Approximating Age-Based Arbitration in On-Chip Networks *M. Lee, J. Kim, D. Abts, M. Marty, J. Lee*
22	Offloading Java to Graphics Processors *P. Calvert*
23	FabScalar: Composing Synthesizable RTL Designs of Arbitrary Cores within a Canonical Superscalar Template *N.K. Choudhary*
24	GLOpenCL: Compiler and Run-Time Support for OpenCL on Hardware- and Software-Managed Cache Multicores *K. Daloukas*
25	Massively Parallel Neural Network Simulation *P. Fox*
26	Autonomic Service Management: Semantic Approach *H. Haitof*
27	Runtime Estimation Methods for Pipelined MPSoCs *H. Javaid*
28	Effective Combination of Transformations via Obstacle Identification *N. Johnson*
29	Graph-based Cluster Assignment for VLIW Architectures *A. Jordan*
30	On stochastic dynamic binary translation in Vx32 *K. Kononenko*
31	A Scalable High-performance Reconfigurable On-chip Network Architecture *M. Modarressi*
32	IIComMP: A Software Framework for Effortless Heterogeneous Architecture Development *R. Reyes*
33	Quantifying the Impact of GPU Specific Optimizations: An Experimental Study on a Weather Forecasting Application *A. Saeed, E. Elwany, E. Tawadros, K. Abdelsalam, P. Yousry, S. Hafez*
34	D^{2}C: Deterministic, Deadlock-free Concurrency *N. Vasudevan*