

# **2012 IEEE 24th International Symposium on Computer Architecture and High Performance Computing**

**(SBAC-PAD 2012)**

**New York, New York, USA  
24-26 October 2012**



**IEEE Catalog Number: CFP12307-PRT  
ISBN: 978-1-4673-4790-7**

# 2012 IEEE 24th International Symposium on Computer Architecture and High Performance Computing **SBAC-PAD 2012**

## Table of Contents

<b>Message from the General Chairs</b> .....	ix
<b>Message from the Program Chairs</b> .....	x
<b>Committees</b> .....	xi
<b>Program Committee</b> .....	xiii
<b>External Reviewers</b> .....	xv
<b>Keynotes</b> .....	xviii

---

### **Session 1: Communication**

The Network Adapter: The Missing Link between MPI Applications and Network Performance .....	1
<i>German Rodriguez, Cyriel Minkenbergh, Ronald P. Luijten, Ramon Beivide, Patrick Geoffray, Jesus Labarta, Mateo Valero, and Steve Poole</i>	
HAT: Heterogeneous Adaptive Throttling for On-Chip Networks .....	9
<i>Kevin Kai-Wei Chang, Rachata Ausavarungnirun, Chris Fallin, and Onur Mutlu</i>	
On the Efficiency of Register File versus Broadcast Interconnect for Collective Communications in Data-Parallel Hardware Accelerators .....	19
<i>Ardavan Pedram, Andreas Gerstlauer, and Robert A. van de Geijn</i>	
Network Endpoints for Clusters of SMPs .....	27
<i>Gabriel Tanase, Gheorghe Almási, Hanhong Xue, and Charles Archer</i>	

### **Session 2: Energy I**

Assessing Energy Efficiency of Fault Tolerance Protocols for HPC Systems .....	35
<i>Esteban Meneses, Osman Sarood, and Laxmikant V. Kalé</i>	
Using Heterogeneous Networks to Improve Energy Efficiency in Direct Coherence Protocols for Many-Core CMPs .....	43
<i>Alberto Ros, Ricardo Fernández-Pascual, and Manuel E. Acacio</i>	

Energy Savings via Dead Sub-Block Prediction .....	51
<i>Marco A.Z. Alves, Khubaib, Eiman Ebrahimi, Veynu T. Narasiman, Carlos Villavieja, Philippe O.A. Navaux, and Yale N. Patt</i>	

Scalable Thread Scheduling in Asymmetric Multicores for Power Efficiency .....	59
<i>Rance Rodrigues, Arunachalam Annamalai, Israel Koren, and Sandip Kundu</i>	

### **Session 3: GPUs / Cloud**

Divergence Analysis with Affine Constraints .....	67
<i>Diogo Sampaio, Rafael Martins, Sylvain Collange, and Fernando Magno Quintão Pereira</i>	

Exploiting Concurrent GPU Operations for Efficient Work Stealing on Multi-GPUs .....	75
<i>João V.F. Lima, Thierry Gautier, Nicolas Maillard, and Vincent Danjean</i>	

Sparse Fast Fourier Transform on GPUs and Multi-core CPUs .....	83
<i>Jiayi Hu, Zhaosen Wang, Qiyuan Qiu, Weijun Xiao, and David J. Lilja</i>	

Cloud Workload Analysis with SWAT .....	92
<i>Mauricio Breternitz, Keith Lowery, Anton Charnoff, Patryk Kaminski, and Leonardo Piga</i>	

### **Session 4: Parallel Algorithms**

Scalable Algorithms for Distributed-Memory Adaptive Mesh Refinement .....	100
<i>Akhil Langer, Jonathan Lifflander, Phil Miller, Kuo-Chuan Pan, Laxmikant V. Kalé, and Paul Ricker</i>	

Compression Speed Enhancements to LZ0 for Multi-core Systems .....	108
<i>Jason Kane and Qing Yang</i>	

Parallelizing Information Set Generation for Game Tree Search Applications .....	116
<i>Mark Richards, Abhishek Gupta, Osman Sarood, and Laxmikant V. Kalé</i>	

A Parallel Implementation of Gomory-Hu’s Cut Tree Algorithm .....	124
<i>Jaime Cohen, Luiz A. Rodrigues, and Elias P. Duarte Jr.</i>	

### **Session 5: Energy II**

Beyond CPU Frequency Scaling for a Fine-grained Energy Control of HPC Systems .....	132
<i>Ghislain Landry Tsafack Chetsa, Laurent Lefevre, Jean-Marc Pierson, Patricia Stolf, and Georges Da Costa</i>	

BTL: A Framework for Measuring and Modeling Energy in Memory Hierarchies .....	139
<i>Ioannis Manousakis and Dimitrios S. Nikolopoulos</i>	

Energy-Performance Tradeoffs in Software Transactional Memory .....	147
<i>Alexandro Baldassin, João P.L. de Carvalho, Leonardo A.G. Garcia, and Rodolfo Azevedo</i>	

Runtime Procedure for Energy Savings in Applications with Point-to-Point Communications .....	155
<i>Vaibhav Sundriyal, Masha Sosonkina, and Alexander Gaenko</i>	

## Session 6: Architecture-Specific Implementations of Parallel Algorithms

Scalable Triadic Analysis of Large-Scale Graphs: Multi-core vs. Multi-processor vs. Multi-threaded Shared Memory Architectures .....	163
<i>George Chin Jr., Andres Marquez, Sutanay Choudhury, and John Feo</i>	
Efficient Sorting on the Tiler Manycore Architecture .....	171
<i>Alessandro Morari, Antonino Tumeo, Oreste Villa, Simone Secchi, and Mateo Valero</i>	
Level-3 BLAS on the TI C6678 Multi-core DSP .....	179
<i>Murtaza Ali, Eric Stotzer, Francisco D. Igual, and Robert A. van de Geijn</i>	
Parallel Exact Inference on Multicore Using MapReduce .....	187
<i>Nam Ma, Yinglong Xia, and Viktor K. Prasanna</i>	

## Session 7: Operating Systems and Run-Time Systems

An OS-Hypervisor Infrastructure for Automated OS Crash Diagnosis and Recovery in a Virtualized Environment .....	195
<i>Joefon Jann, R. Sarma Burugula, Ching-Farn E. Wu, and Kaoutar El Maghraoui</i>	
VPC: Scalable, Low Downtime Checkpointing for Virtual Clusters .....	203
<i>Peng Lu, Binoy Ravindran, and Changsoo Kim</i>	
FusedOS: Fusing LWK Performance with FWK Functionality in a Heterogeneous Environment .....	211
<i>Yoonho Park, Eric Van Hensbergen, Marius Hillenbrand, Todd Inglett, Bryan Rosenburg, Kyung Dong Ryu, and Robert W. Wisniewski</i>	
Transactional Forwarding: Supporting Highly-Concurrent STM in Asynchronous Distributed Systems .....	219
<i>Mohamed M. Saad and Binoy Ravindran</i>	

## Session 8: Global Data/Shared Data

Exploiting Phase-Change Memory in Cooperative Caches .....	227
<i>Luiz Ramos and Ricardo Bianchini</i>	
Global Data Re-allocation via Communication Aggregation in Chapel .....	235
<i>Alberto Sanz, Rafael Asenjo, Juan López, Rafael Larrosa, Angeles Navarro, Vassily Litvinov, Sung-Eun Choi, and Bradford L. Chamberlain</i>	
Integrating Dataflow Abstractions into the Shared Memory Model .....	243
<i>Vladimir Gajinov, Srdjan Stipic, Osman S. Unsal, Tim Harris, Eduard Ayguadé, and Adrián Cristal</i>	
CSHARP: Coherence and SHaring Aware Cache Replacement Policies for Parallel Applications .....	252
<i>Biswabandan Panda and Shankar Balachandran</i>	

## **Session 9: Instruction-Level Issues**

Low Overhead Instruction-Cache Modeling Using Instruction Reuse Profiles .....	260
<i>Muneeb Khan, Andreas Sembrant, and Erik Hagersten</i>	
Data and Instruction Uniformity in Minimal Multi-threading .....	270
<i>Teo Milanez, Sylvain Collange, Fernando Magno Quintão Pereira, Wagner Meira Jr., and Renato A. Ferreira</i>	
ACCGen: An Automatic ArchC Compiler Generator .....	278
<i>Rafael Auler, Paulo Cesar Centoducatte, and Edson Borin</i>	
Efficiently Handling Memory Accesses to Improve QoS in Multicore Systems under Real-Time Constraints .....	286
<i>José Luis March, Salvador Petit, Julio Sahuquillo, Houcine Hassan, and José Duato</i>	
<b>Author Index</b> .....	294