

2013 25th International Symposium on Computer Architecture and High Performance Computing

(SBAC-PAD 2013)

**Porto de Galinhas, Pernambuco, Brazil
23-26 October 2013**



**IEEE Catalog Number: CFP13307-POD
ISBN: 978-1-4799-2929-0**

2013 25th International Symposium on Computer Architecture and High Performance Computing SBAC-PAD 2013

Table of Contents

Message from the General Chairs	viii
Message from the Program Chairs	ix
Committees	x
Program Committee and External Reviewers	xii

Session 1: Architecture

Cluster Cache Monitor	1
<i>Guohong Li, Olivier Temam, Zhenyu Liu, Dongsheng Wang, Sanchuan Guo, and Chongmin Li</i>	
Experiences with Disjoint Data Structures in a New Hardware Transactional Memory System	9
<i>Gong Su and Stephen Heisig</i>	
HotStream: Efficient Data Streaming of Complex Patterns to Multiple Accelerating Kernels	17
<i>Sérgio Paiáguas, Frederico Pratas, Pedro Tomás, Nuno Roma, and Ricardo Chaves</i>	
Large Payload Streaming Database Sort and Projection on FPGAs	25
<i>Bharat Sukhwani, Mathew Thoennes, Hong Min, Parijat Dube, Bernard Brezzo, Sameh Asaad, and Donna Dillenberger</i>	

Session 2: Networks

Scalable Many-Field Packet Classification on Multi-core Processors	33
<i>Yun Qu, Shijie Zhou, and Viktor K. Prasanna</i>	
Extending Summation Precision for Network Reduction Operations	41
<i>George Michelogiannakis, Xiaoye S. Li, David H. Bailey, and John Shalf</i>	
Tackling Permanent Faults in the Network-on-Chip Router Pipeline	49
<i>Pavan Poluri and Ahmed Louri</i>	

Fast LH*	57
<i>Juan Chabkinian and Thomas Schwarz</i>	

Session 3: Energy-Efficient Design

Attaining Strictly Increasing and Precise Time Count in Energy-Efficient Computer Systems	65
<i>Diego Leonel Cadette Dutra, Lauro Luis Armondi Whately, and Claudio Luis de Amorim</i>	
Energy Efficient Last Level Caches via Last Read/Write Prediction	73
<i>Marco A.Z. Alves, Carlos Villavieja, Matthias Diener, and Philippe O.A. Navaux</i>	
Dynamic Selective Devectorization for Efficient Power Gating of SIMD Units in a HW/SW Co-Designed Environment	81
<i>Rakesh Kumar, Alejandro Martínez, and Antonio González</i>	
HPC Performance and Energy-Efficiency of Xen, KVM and VMware Hypervisors	89
<i>Sébastien Varrette, Mateusz Guzek, Valentin Plugaru, Xavier Besseron, and Pascal Bouvry</i>	

Session 4: Applications I

Optimizing a 3D-FWT Code in a Heterogeneous Cluster of Multicore CPUs and Manycore GPUs	97
<i>Gregorio Bernabé, Javier Cuenca, and Domingo Giménez</i>	
Preliminary Experiments with XKaapi on Intel Xeon Phi Coprocessor	105
<i>João V.F. Lima, François Broquedis, Thierry Gautier, and Bruno Raffin</i>	
A CPU, GPU, FPGA System for X-Ray Image Processing Using High-Speed Scientific Cameras	113
<i>Alécio Pedro Delazari Binotto, Dionisio Doering, Thorsten Stetzelberger, Patrick McVittie, Sergio Zimmermann, and Carlos Eduardo Pereira</i>	
A Parallel IRAM Algorithm to Compute PageRank for Modeling Epidemic Spread	120
<i>Zifan Liu, Nahid Emad, Soufian Ben Amor, and Michel Lamure</i>	

Session 5: Algorithms and Scheduling

Dealing with Reduction Operations Using Transactional Memory	128
<i>Miguel A. Gonzalez-Mesa, Ricardo Quisiant, Eladio Gutierrez, and Oscar Plata</i>	
Invasive Compute Balancing for Applications with Hybrid Parallelization	136
<i>Martin Schreiber, Christoph Riesinger, Tobias Neckel, and Hans-Joachim Bungartz</i>	
Correct and Efficient Bounded FIFO Queues	144
<i>Nhat Minh Lê, Adrien Guatto, Albert Cohen, and Antoniu Pop</i>	

List Scheduling in Embedded Systems under Memory Constraints	152
<i>Paul-Antoine Arras, Didier Fuin, Emmanuel Jeannot, Arthur Stoutchinin, and Samuel Thibault</i>	

Session 6: GPU Applications

Benchmarking GPUs with a Parallel Lattice-Boltzmann Code	160
<i>Jiri Kraus, Marcello Pivanti, Sebastiano Fabio Schifano, Raffaele Tripiccione, and Marco Zanella</i>	
GPU-NB: A Fast CUDA-Based Implementation of Naïve Bayes	168
<i>Guilherme Andrade, Felipe Viegas, Gabriel Spada Ramos, Jussara Almeida, Leonardo Rocha, Marcos Gonçalves, and Renato Ferreira</i>	
Image Re-ranking Acceleration on GPUs	176
<i>Daniel Carlos Guimarães Pedronette, Ricardo da S. Torres, Edson Borin, and Mauricio Breternitz</i>	
SPLG: A Tuned Signal Processing Library for GPU Architectures	184
<i>Jacobo Lobeiras Blanco, Margarita Amor, and Ramón Doallo</i>	

Session 7: Applications II

Autonomic Malleability in Iterative MPI Applications	192
<i>Alexandre Sena, Felipe Ribeiro, Vinod Rebello, Aline Nascimento, and Cristina Boeres</i>	
On the Performance of Code Block Segmentation for LTE-Advanced: An In-Depth Analysis	200
<i>Karlo Lenzi, Felipe A.P. Figueiredo, José A.B. Filho, and Fabricio L. Figueiredo</i>	
Easily Rendering Token-Ring Algorithms of Distributed and Parallel Applications Fault Tolerant	206
<i>Luciana Arantes and Julien Sopena</i>	
Parallel Algorithms for Using Lagrangian Markers in Immersed Boundary Method with Adaptive Mesh Refinement in FLASH	214
<i>Prateeti Mohapatra, Anshu Dubey, Christopher Daley, Marcos Vanella, and Elias Balaras</i>	
Author Index	221