

# **2011 International Conference on Parallel Architectures and Compilation Techniques**

**(PACT 2011)**

**Galveston, Texas, USA  
10 – 14 October 2011**



**IEEE Catalog Number: CFP11073-PRT  
ISBN: 978-1-4577-1794-9**

# 2011 International Conference on Parallel Architectures and Compilation Techniques

## PACT 2011

### Table of Contents

Message from the General Chair.....	xi
Message from the Program Chair.....	xii
Conference Committees.....	xiii
Reviewers.....	xvi

---

#### Session 2a: Scheduling

A Unified Scheduler for Recursive and Task Dataflow Parallelism .....	1
<i>Hans Vandierendonck, George Tzenakis, and Dimitrios S. Nikolopoulos</i>	
No More Backstabbing... A Faithful Scheduling Policy for Multithreaded Programs .....	12
<i>Kishore Kumar Pusukuri, Rajiv Gupta, and Laxmi N. Bhuyan</i>	
Dynamic Fine-Grain Scheduling of Pipeline Parallelism .....	22
<i>Daniel Sanchez, David Lo, Richard M. Yoo, Jeremy Sugerman, and Christos Kozyrakis</i>	

#### Session 2b: Coherence

SPATL: Honey, I Shrunk the Coherence Directory .....	33
<i>Hongzhou Zhao, Arrvindh Shriraman, Sandhya Dwarkadas, and Vijayalakshmi Srinivasan</i>	
POPS: Coherence Protocol Optimization for Both Private and Shared Data .....	45
<i>Hemayet Hossain, Sandhya Dwarkadas, and Michael C. Huang</i>	
An OpenCL Framework for Homogeneous Manycores with No Hardware Cache Coherence .....	56
<i>Jun Lee, Jungwon Kim, Junghyun Kim, Sangmin Seo, and Jaejin Lee</i>	

#### Session 3a: High-Level Programming Frameworks

Compiling Dynamic Data Structures in Python to Enable the Use of Multi-core and Many-core Libraries .....	68
<i>Bin Ren and Gagan Agrawal</i>	
Efficient Parallel Graph Exploration on Multi-Core CPU and GPU .....	78
<i>Sungpack Hong, Tayo Oguntebi, and Kunle Olukotun</i>	

A Heterogeneous Parallel Framework for Domain-Specific Languages .....	89
<i>Kevin J. Brown, Arvind K. Sujeeth, Hyouk Joong Lee, Tiark Rompf, Hassan Chafi, Martin Odersky, and Kunle Olukotun</i>	
<b>Session 3b: Power Efficiency</b>	
PEPSC: A Power-Efficient Processor for Scientific Computing .....	101
<i>Ganesh Dasika, Ankit Sethia, Trevor Mudge, and Scott Mahlke</i>	
Improving Throughput of Power-Constrained GPUs Using Dynamic Voltage/Frequency and Core Scaling .....	111
<i>Jungseob Lee, Vijay Sathisha, Michael Schulte, Katherine Compton, and Nam Sung Kim</i>	
Performance Per Watt Benefits of Dynamic Core Morphing in Asymmetric Multicores .....	121
<i>Rance Rodrigues, Arunachalam Annamalai, Israel Koren, Sandip Kundu, and Omer Khan</i>	
<b>Session 4: Best Paper Session</b>	
Phase-Based Application-Driven Hierarchical Power Management on the Single-chip Cloud Computer .....	131
<i>Nikolas Ioannou, Michael Kauschke, Matthias Gries, and Marcelo Cintra</i>	
Optimizing Data Layouts for Parallel Computation on Multicores .....	143
<i>Yuanrui Zhang, Wei Ding, Jun Liu, and Mahmut Kandemir</i>	
DeNovo: Rethinking the Memory Hierarchy for Disciplined Parallelism .....	155
<i>Byn Choi, Rakesh Komuravelli, Hyojin Sung, Robert Smolinski, Nima Honarmand, Sarita V. Adve, Vikram S. Adve, Nicholas P. Carter, and Ching-Tsun Chou</i>	
<b>PACT 2011 Posters</b>	
A Hierarchical Approach to Maximizing MapReduce Efficiency .....	167
<i>Zhiwei Xiao, Haibo Chen, and Binyu Zang</i>	
Building Retargetable and Efficient Compilers for Multimedia Instruction Sets .....	169
<i>Serge Guelton, Adrien Guinet, and Ronan Keryell</i>	
Compiler Directed Data Locality Optimization for Multicore Architectures .....	171
<i>Wei Ding, Jithendra Srinivas, Mahmut Kandemir, and Mustafa Karakoy</i>	
CriticalFault: Amplifying Soft Error Effect Using Vulnerability-Driven Injection .....	173
<i>Xin Xu and Man-Lap Li</i>	
Understanding the Behavior of Pthread Applications on Non-Uniform Cache Architectures .....	175
<i>Gagandeep S. Sachdev, Kshitij Sudan, Mary W. Hall, and Rajeev Balasubramonian</i>	
Exploiting Mutual Awareness between Prefetchers and On-chip Networks in Multi-cores .....	177
<i>Junghoon Lee, Minjeong Shin, Hanjoon Kim, John Kim, and Jaehyuk Huh</i>	
Decoupled Architectures as a Low-Complexity Alternative to Out-of-order Execution .....	179
<i>Neal C. Crago and Sanjay J. Patel</i>	

Parameterized Micro-benchmarking: An Auto-tuning Approach for Complex Applications .....	181
<i>Wenjing Ma, Sriram Krishnamoorthy, and Gagan Agrawal</i>	
Prediction Based DRAM Row-Buffer Management in the Many-Core Era .....	183
<i>Manu Awasthi, David W. Nellans, Rajeev Balasubramonian, and Al Davis</i>	
Program Interferometry .....	185
<i>Zhe Wang and Daniel A. Jiménez</i>	
Regulating Locality vs. Parallelism Tradeoffs in Multiple Memory Controller Environments .....	187
<i>Syed Minhaj Hassan, Dhruv Choudhary, Michelle Rasquinha, and Sudhakar Yalamanchili</i>	
Row-Buffer Reorganization: Simultaneously Improving Performance and Reducing Energy in DRAMs .....	189
<i>Nagendra Gulur, R. Manikantan, R. Govindarajan, and Mahesh Mehendale</i>	
Scalable Proximity-Aware Cache Replication in Chip Multiprocessors .....	191
<i>Chongmin Li, Haixia Wang, Yibo Xue, Dongsheng Wang, and Jian Li</i>	
Scalable and Efficient Bounds Checking for Large-Scale CMP Environments .....	193
<i>Baik Song An, Ki Hwan Yum, and Eun Jung Kim</i>	
An Alternative Memory Access Scheduling in Manycore Accelerators .....	195
<i>Yonggon Kim, Hyunseok Lee, and John Kim</i>	
Beforehand Migration on D-NUCA Caches .....	197
<i>Javier Lira, Timothy M. Jones, Carlos Molina, and Antonio González</i>	
SymptomTM: Symptom-Based Error Detection and Recovery Using Hardware Transactional Memory .....	199
<i>Gulay Yalcin, Osman S. Unsal, Adrian Cristal, Ibrahim Hur, and Mateo Valero</i>	
rPRAM: Exploring Redundancy Techniques to Improve Lifetime of PCM-based Main Memory .....	201
<i>Jie Chen, Zachary Winter, Guru Venkataramani, and H. Howie Huang</i>	
Pi-TM: Pessimistic Invalidation for Scalable Lazy Hardware Transactional Memory .....	203
<i>Anurag Negi, Per Stenstrom, Rubén Titos-Gil, Manuel E. Acacio, and José M. Garcia</i>	
MCFQ: Leveraging Memory-level Parallelism and Application's Cache Friendliness for Efficient Management of Quasi-partitioned Last-level Caches .....	205
<i>Dimitris Kaseridis, Muhammad Faisal Iqbal, Jeffrey Stuecheli, and Lizy Kurian John</i>	
MRAC: A Memristor-based Reconfigurable Framework for Adaptive Cache Replacement .....	207
<i>Ping Zhou, Bo Zhao, Youtao Zhang, Jun Yang, and Yiran Chen</i>	

## **PACT 2011 SRC Posters**

Sampling Temporal Touch Hint (STTH) Inclusive Cache Management Policy .....	209
<i>Yingying Tian and Daniel A. Jiménez</i>	
Exploiting Rank Idle Time for Scheduling Last-Level Cache Writeback .....	210
<i>Zhe Wang and Daniel A. Jiménez</i>	
TIDeFlow: A Parallel Execution Model for High Performance Computing Programs .....	211
<i>Daniel Orozco</i>	
Decoupled Cache Segmentation: Mutable Policy with Automated Bypass .....	212
<i>Samira Khan and Daniel A. Jiménez</i>	
A Software-Managed Coherent Memory Architecture for Manycores .....	213
<i>Jungho Park, Choonki Jang, and Jaejin Lee</i>	
Improving Last-Level Cache Performance by Exploiting the Concept of MRU-Tour .....	214
<i>Alejandro Valero, Julio Sahuquillo, Salvador Petit, Pedro López, and José Duato</i>	
A Compiler-assisted Runtime-prefetching Scheme for Heterogenous Platforms .....	215
<i>Baojiang Shou, Xionghui Hou, and Li Chen</i>	
Improving Run-Time Scheduling for General-Purpose Parallel Code .....	216
<i>Alexandros Tzannes, Rajeev Barua, and Uzi Vishkin</i>	
Collaborative Caching for Unknown Cache Sizes .....	217
<i>Xiaoming Gu</i>	
Programming Strategies for GPUs and their Power Consumption .....	218
<i>Sayan Ghosh and Barbara Chapman</i>	
An Architecture to Enable Lifetime Full Chip Testability in Chip Multiprocessors .....	219
<i>Rance Rodrigues, Israel Koren, and Sandip Kundu</i>	
Probabilistic Models Towards Optimal Speculation of DFA Applications .....	220
<i>Zhijia Zhao and Bo Wu</i>	

## **Session 6a: Transactional Memory**

STM2: A Parallel STM for High Performance Simultaneous Multithreading Systems .....	221
<i>Gokcen Kestor, Roberto Gioiosa, Tim Harris, Osman S. Unsal, Adrian Cristal, Ibrahim Hur, and Mateo Valero</i>	
Making STMs Cache Friendly with Compiler Transformations .....	232
<i>Sandya Mannarswamy and Ramaswamy Govindarajan</i>	

## **Session 8a: Locality**

Enhancing Data Locality for Dynamic Simulations through Asynchronous Data Transformations and Adaptive Control .....	243
<i>Bo Wu, Eddy Z. Zhang, and Xipeng Shen</i>	

SFMalloc: A Lock-Free and Mostly Synchronization-Free Dynamic Memory Allocator for Manycores .....	253
<i>Sangmin Seo, Junghyun Kim, and Jaejin Lee</i>	
Coherent Profiles: Enabling Efficient Reuse Distance Analysis of Multicore Scaling for Loop-based Parallel Programs .....	264
<i>Meng-Ju Wu and Donald Yeung</i>	
<b>Session 8b: Customized Processors</b>	
StVEC: A Vector Instruction Extension for High Performance Stencil Computation .....	276
<i>Naser Sedaghati, Renji Thomas, Louis-Noël Pouchet, Radu Teodorescu, and P. Sadayappan</i>	
OpenMDSP: Extending OpenMP to Program Multi-Core DSP .....	288
<i>Jiangzhou He, Wenguang Chen, Guangri Chen, Weimin Zheng, Zhizhong Tang, and Handong Ye</i>	
ARIADNE: Agnostic Reconfiguration in a Disconnected Network Environment .....	298
<i>Konstantinos Aisopos, Andrew DeOrio, Li-Shiuan Peh, and Valeria Bertacco</i>	
<b>Session 9a: SPMD Analysis</b>	
Correctly Treating Synchronizations in Compiling Fine-Grained SPMD-Threaded Programs for CPU .....	310
<i>Ziyu Guo, Eddy Zheng Zhang, and Xipeng Shen</i>	
Divergence Analysis and Optimizations .....	320
<i>Bruno Coutinho, Diogo Sampaio, Fernando Magno Quintão Pereira, and Wagner Meira Jr.</i>	
Large Scale Verification of MPI Programs Using Lamport Clocks with Lazy Update .....	330
<i>Anh Vo, Ganesh Gopalakrishnan, Robert M. Kirby, Bronis R. de Supinski, Martin Schulz, and Greg Bronevetsky</i>	
<b>Session 9b: Memory Hierarchies</b>	
DiDi: Mitigating the Performance Impact of TLB Shootdowns Using a Shared TLB Directory .....	340
<i>Carlos Villavieja, Vasileios Karakostas, Lluís Vilanova, Yoav Etsion, Alex Ramirez, Avi Mendelson, Nacho Navarro, Adrian Cristal, and Osman S. Unsal</i>	
Linear-time Modeling of Program Working Set in Shared Cache .....	350
<i>Xiaoya Xiang, Bin Bao, Chen Ding, and Yaoqing Gao</i>	
Using a Reconfigurable L1 Data Cache for Efficient Version Management in Hardware Transactional Memory .....	361
<i>Adrià Armejach, Azam Seyedi, Rubén Titos-Gil, Ibrahim Hur, Adrián Cristal, Osman S. Unsal, and Mateo Valero</i>	

## **Session 10a: Compiler Optimizations**

An Evaluation of Vectorizing Compilers .....	372
<i>Saeed Maleki, Yaoqing Gao, María J. Garzarán, Tommy Wong, and David A. Padua</i>	
Modeling and Performance Evaluation of TSO-Preserving Binary Optimization .....	383
<i>Cheng Wang and Youfeng Wu</i>	
Exploiting Task Order Information for Optimizing Sequentially Consistent Java Programs .....	393
<i>Christoph M. Angerer and Thomas R. Gross</i>	

## **Session 10b: Potpourri**

Memory Architecture for Integrating Emerging Memory Technologies .....	403
<i>Kun Fang, Long Chen, Zhao Zhang, and Zhichun Zhu</i>	
Speculative Parallelization in Decoupled Look-ahead .....	413
<i>Alok Garg, Raj Parihar, and Michael C. Huang</i>	
Optimizing Regular Expression Matching with SR-NFA on Multi-Core Systems .....	424
<i>Yi-Hua E. Yang and Viktor K. Prasanna</i>	
<b>Author Index</b> .....	434