

2015 44th International Conference on Parallel Processing (ICPP 2015)

**Beijing, China
1-4 September 2015**

Pages 1-489



**IEEE Catalog Number: CFP15127-POD
ISBN: 978-1-4673-7588-7**

**Copyright © 2015 by the Institute of Electrical and Electronic Engineers, Inc
All Rights Reserved**

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved.

******This publication is a representation of what appears in the IEEE Digital Libraries. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP15127-POD
ISBN (Print-On-Demand):	978-1-4673-7588-7
ISSN:	0190-3918

Additional Copies of This Publication Are Available From:

Curran Associates, Inc
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: (845) 758-0400
Fax: (845) 758-2633
E-mail: curran@proceedings.com
Web: www.proceedings.com

CURRAN ASSOCIATES INC.
proceedings
.com

2015 44th International Conference on Parallel Processing

ICPP 2015

Table of Contents

Message from the General Co-Chairs.....	xv
Message from the Program Co-Chairs.....	xvii
Organizing and Program Committees.....	xviii
External Reviewers.....	xxiv

Session 1A: Data Center and Cloud Architectures

Optimization of Resource Allocation and Energy Efficiency in Heterogeneous Cloud Data Centers	1
<i>Amer Qouneh, Ming Liu, and Tao Li</i>	
Exploring Hardware Profile-Guided Green Datacenter Scheduling	11
<i>Weichao Tang, Yu Wang, Haopeng Liu, Tao Zhang, Chao Li, and Xiaoyao Liang</i>	
Designing a Hybrid Scale-Up/Out Hadoop Architecture Based on Performance Measurements for High Application Performance	21
<i>Zhuozhao Li and Haiying Shen</i>	
Profiling and Understanding Virtualization Overhead in Cloud	31
<i>Lihua Chen, Shilkumar Patel, Haiying Shen, and Zhongyi Zhou</i>	

Session 1B: Parallel Programming on Many Cores

Accelerating Spectral Calculation through Hybrid GPU-Based Computing	41
<i>Jian Xiao, Xingyu Xu, Ce Yu, Jiawan Zhang, Shuinai Zhang, Li Ji, and Jizhou Sun</i>	
LBM-IB: A Parallel Library to Solve 3D Fluid-Structure Interaction Problems on Manycore Systems	51
<i>Prateek Nagar, Fengguang Song, Luoding Zhu, and Lan Lin</i>	
Resident Block-Structured Adaptive Mesh Refinement on Thousands of Graphics Processing Units	61
<i>David Beckingsale, Wayne Gaudin, Andrew Herdman, and Stephen Jarvis</i>	

Pattern-Driven Hybrid Multi- and Many-Core Acceleration in the MPAS Shallow-Water Model	71
<i>Peng Zhang, Yulong Ao, Chao Yang, Yiqun Liu, Fangfang Liu, Changmao Wu, and Haitao Zhao</i>	

Session 1C: Networking Algorithms

On Maximizing Reliability of Lifetime Constrained Data Aggregation Tree in Wireless Sensor Networks	81
<i>Mengfan Shan, Guihai Chen, Fan Wu, Xiaobing Wu, Xiaofeng Gao, Pan Wu, and Haipeng Dai</i>	
RepFD - Using Reputation Systems to Detect Failures in Large Dynamic Networks	91
<i>Maxime Véron, Olivier Marin, Sébastien Monnet, and Pierre Sens</i>	
Network Coding for Effective NDN Content Delivery: Models, Experiments, and Applications	101
<i>Kai Lei, Fangxing Zhu, Cheng Peng, and Kuai Xu</i>	
PLP: Protecting Location Privacy Against Correlation-Analysis Attack in Crowdsensing	111
<i>Shanfeng Zhang, Qiang Ma, Tong Zhu, Kebin Liu, Lan Zhang, Wenbo He, and Yunhao Liu</i>	

Session 2A: Architecture Support for Parallelism

Leveraging Error Compensation to Minimize Time Deviation in Parallel Multi-core Simulations	120
<i>Xiaodong Zhu, Junmin Wu, and Tao Li</i>	
DiMP: Architectural Support for Direct Message Passing on Shared Memory Multi-cores	130
<i>Ruben Titos-Gil, Oscar Palomar, Osman Unsal, and Adrian Cristal</i>	
An Energy-Efficient Branch Prediction with Grouped Global History	140
<i>Mingkai Huang, Dan He, Xianhua Liu, Mingxing Tan, and Xu Cheng</i>	
Spatial Locality-Aware Cache Partitioning for Effective Cache Sharing	150
<i>Saurabh Gupta and Huiyang Zhou</i>	

Session 2B: Network Routing and Management

DISCS: A DIStributed Collaboration System for Inter-AS Spoofing Defense	160
<i>Bingyang Liu and Jun Bi</i>	
Connecting the Dots: Reconstructing Network Behavior with Individual and Lossy Logs	170
<i>Jiliang Wang, Xiaolong Zheng, Xufei Mao, Zhichao Cao, Daibo Liu, and Yunhao Liu</i>	

MIFO: Multi-path Interdomain Forwarding	180
<i>Ming Zhu, Dan Li, Ying Liu, Dan Pei, K.K. Ramakrishnan, Lili Liu, and Jianping Wu</i>	
Energy-Efficient and Delay-Constrained Broadcast in Time-Varying Energy-Demand Graphs	190
<i>Chenxi Qiu, Haiying Shen, and Lei Yu</i>	

Session 2C: Parallel Algorithm and Performance Optimization

Design and Implementation of a Highly Efficient DGEMM for 64-Bit ARMv8 Multi-core Processors	200
<i>Feng Wang, Hao Jiang, Ke Zuo, Xing Su, Jingling Xue, and Canqun Yang</i>	
In-Place Data Sliding Algorithms for Many-Core Architectures	210
<i>Juan Gómez Luna, Li-Wen Chang, I-Jui Sung, Wen-Mei Hwu, and Nicolás Guil</i>	
Modelling and Developing Co-scheduling Strategies on Multicore Processors	220
<i>Huanzhou Zhu, Ligang He, Bo Gao, Kenli Li, Jianhua Sun, Hao Chen, and Keqin Li</i>	
Optimizing Image Sharpening Algorithm on GPU	230
<i>Mengran Fan, Haipeng Jia, Yunquan Zhang, Xiaojing An, and Ting Cao</i>	

Session 3A: Video Streaming

Energy-Aware Cost-Effective Cooperative Mobile Streaming on Smartphones over Hybrid Wireless Networks	240
<i>Zhihao Qu, Baoliu Ye, Bin Tang, Sanglu Lu, and Song Guo</i>	
Scan without a Glance: Towards Content-Free Crowd-Sourced Mobile Video Retrieval System	250
<i>Cihang Liu, Lan Zhang, Kebin Liu, and Yunhao Liu</i>	
Do Twin Clouds Make Smoothness for Transoceanic Video Telephony?	260
<i>Jian Li, Zhenhua Li, Yao Liu, and Zhi-Li Zhang</i>	

Session 3B: Parallel Applications

Characterizing Multi-media Retrieval Applications	270
<i>Yunping Lu, Xin Wang, Weihua Zhang, Yi Li, and Wenyun Zhao</i>	
Accelerating I/O Performance of Big Data Analytics on HPC Clusters through RDMA-Based Key-Value Store	280
<i>Nusrat Sharmin Islam, Dipti Shankar, Xiaoyi Lu, Md. Wasi-Ur-Rahman, and Dhabaleswar K. (DK) Panda</i>	
SCAN: A Smart Application Platform for Empowering Parallelizations of Big Genomic Data Analysis in Clouds	290
<i>Wei Xing, Wei Jie, and Crispin Miller</i>	

Session 3C: Accelerators

Automatic Performance Tuning of Stencil Computations on GPUs	300
<i>Joseph D. Garvey and Tarek S. Abdelrahman</i>	
Open ACC Programs Examined: A Performance Analysis Approach	310
<i>Robert Dietrich, Guido Juckeland, and Michael Wolfe</i>	
GPGPU Benchmark Suites: How Well Do They Sample the Performance Spectrum?	320
<i>Jee Ho Ryoo, Saddam J. Quirem, Michael Lebeane, Reena Panda, Shuang Song, and Lizy K. John</i>	

Session 4A: Parallelism from Chips to Cities

Fast FCoE: An Efficient and Scale-Up Multi-core Framework for FCoE-Based SAN Storage Systems	330
<i>Yunxiang Wu, Fang Wang, Yu Hua, Dan Feng, Yuchong Hu, Jingning Liu, and Wei Tong</i>	
A Heterogeneity-Aware Region-Level Data Layout for Hybrid Parallel File Systems	340
<i>Shuibing He, Xian-He Sun, Yang Wang, Antonis Kougkas, and Adnan Haider</i>	
Shorter On-Line Warmup for Sampled Simulation of Multi-threaded Applications	350
<i>Chuntao Jiang, Zhibin Yu, Hai Jin, Xiaofei Liao, Lieven Eeckhout, Yonggang Zeng, and Chengzhong Xu</i>	
Reflex-Tree: A Biologically Inspired Parallel Architecture for Future Smart Cities	360
<i>Jason Kane, Bo Tang, Zhen Chen, Jun Yan, Tao Wei, Haibo He, and Qing Yang</i>	

Session 4B: Cloud Algorithms

PDTL: Parallel and Distributed Triangle Listing for Massive Graphs	370
<i>Ilias Giechaskiel, George Panagopoulos, and Eiko Yoneki</i>	
Good Work Deserves Good Pay: A Quality-Based Surplus Sharing Method for Participatory Sensing	380
<i>Shuo Yang, Fan Wu, Shaojie Tang, Xiaofeng Gao, Bo Yang, and Guihai Chen</i>	
Optimal Node Selection for Data Regeneration in Heterogeneous Distributed Storage Systems	390
<i>Qingyuan Gong, Jiaqi Wang, Dongsheng Wei, Jin Wang, and Xin Wang</i>	
Joint Media Streaming Optimization of Energy and Rebuffering Time in Cellular Networks	400
<i>Zeqi Lai, Yong Cui, Yayun Bao, Jiangchuan Liu, Yingchao Zhao, and Xiao Ma</i>	

Session 4C: Reliability

What Is Wrong with the Transmission? A Comprehensive Study on Message Passing Related Bugs	410
<i>Mingxing Zhang, Yongwei Wu, Kang Chen, and Weimin Zheng</i>	
Bit Flipping Errors in High Performance Linpack at Exascale and Beyond	420
<i>Erlin Yao and Guangming Tan</i>	
Is Your Graph Algorithm Eligible for Nondeterministic Execution?	430
<i>Zhiyuan Shao, Lin Hou, Yan Ai, Yu Zhang, and Hai Jin</i>	
Assessing the Impact of Partial Verifications against Silent Data Corruptions	440
<i>Aurélien Cavelan, Saurabh K. Raina, Yves Robert, and Hongyang Sun</i>	

Session 5A: Reliable Parallel Storage Architectures

Code 5-6: An Efficient MDS Array Coding Scheme to Accelerate Online RAID Level Migration	450
<i>Chentao Wu, Xubin He, Jie Li, and Minyi Guo</i>	
PPM: A Partitioned and Parallel Matrix Algorithm to Accelerate Encoding/Decoding Process of Asymmetric Parity Erasure Codes	460
<i>Shiyi Li, Qiang Cao, Shenggang Wan, Wenhui Zhang, Changsheng Xie, Xubin He, and Pradeep Subedi</i>	
CoRec: A Cooperative Reconstruction Pattern for Multiple Failures in Erasure-Coded Storage Clusters	470
<i>Jianzhong Huang, Erwei Dai, Changsheng Xie, and Xiao Qin</i>	
EC-FRM: An Erasure Coding Framework to Speed Up Reads for Erasure Coded Cloud Storage Systems	480
<i>Yingxun Fu, Jiwu Shu, and Zhirong Shen</i>	

Session 5B: Cloud Computing

PCS: Predictive Component-Level Scheduling for Reducing Tail Latency in Cloud Online Services	490
<i>Rui Han, Junwei Wang, Siguang Huang, Chenrong Shao, Shulin Zhan, Jianfeng Zhan, and Jose Luis Vazquez-Poletti</i>	
Cloud Fog: Towards High Quality of Experience in Cloud Gaming	500
<i>Yuhua Lin and Haiying Shen</i>	
SLA-Based Resource Scheduling for Big Data Analytics as a Service in Cloud Computing Environments	510
<i>Yali Zhao, Rodrigo N. Calheiros, Graeme Gange, Kotagiri Ramamohanarao, and Rajkumar Buyya</i>	

Executing Large Scale Scientific Workflow Ensembles in Public Clouds	520
<i>Qingye Jiang, Young Choon Lee, and Albert Y. Zomaya</i>	

Session 5C: Resource Allocation and Contention Management

A Penalty Aware Memory Allocation Scheme for Key-Value Cache	530
<i>Jianqiang Ou, Marc Patton, Michael Devon Moore, Yuehai Xu, and Song Jiang</i>	
Reducing Synchronization Cost in Distributed Multi-resource Allocation Problem	540
<i>Jonathan Lejeune, Luciana Arantes, Julien Sopena, and Pierre Sens</i>	
Green-CM: Energy Efficient Contention Management for Transactional Memory	550
<i>Shady Issa, Paolo Romano, and Mats Brorsson</i>	
Matchmaking Applications and Partitioning Strategies for Efficient Execution on Heterogeneous Platforms	560
<i>Jie Shen, Ana Lucia Varbanescu, Xavier Martorell, and Henk Sips</i>	

Session 6A: Network Applications

Social VoD: A Social Feature-Based P2P System	570
<i>Wei Chang and Jie Wu</i>	
Crowdsourcing Sensing Workloads of Heterogeneous Tasks: A Distributed Fairness-Aware Approach	580
<i>Wei Sun, Yanmin Zhu, Lionel M. Ni, and Bo Li</i>	
Parallel (Probable) Lock-Free Hash Sieve: A Practical Sieving Algorithm for the SVP	590
<i>Artur Mariano, Christian Bischof, and Thijs Laarhoven</i>	

Session 6B: Hardware Exploitation for Computing

Efficient Use of Hardware Transactional Memory for Parallel Mesh Generation	600
<i>Tetsu Kobayashi, Shigeyuki Sato, and Hideya Iwasaki</i>	
Zebra: An East-West Control Framework for SDN Controllers	610
<i>Haisheng Yu, Keqiu Li, Heng Qi, Wenxin Li, and Xiaoyi Tao</i>	
Software-Based Lightweight Multithreading to Overlap Memory-Access Latencies of Commodity Processors	619
<i>Cihang Jiang, Youhui Zhang, and Weimin Zheng</i>	

Session 6C: Energy Efficiency

Using Per-Loop CPU Clock Modulation for Energy Efficiency in OpenMP Applications	629
<i>Wei Wang, Allan Porterfield, John Cavazos, and Sridutt Bhalachandra</i>	
Evaluating Latency-Sensitive Applications: Performance Degradation in Datacenters with Restricted Power Budget	639
<i>Song Wu, Chuxiong Yan, Haibao Chen, Hai Jin, Wei Guo, Zhen Wang, and Deqing Zou</i>	
REED: A Reliable Energy-Efficient RAID	649
<i>Shu Yin, Xuewu Li, Kenli Li, Jianzhong Huang, Xiaojun Ruan, Xiaomin Zhu, Wei Cao, and Xiao Qin</i>	

Session 7A: Data Center Networks

TAPS: Software Defined Task-Level Deadline-Aware Preemptive Flow Scheduling in Data Centers	659
<i>Lili Liu, Dan Li, and Jianping Wu</i>	
Sheriff: A Regional Pre-alert Management Scheme in Data Center Networks	669
<i>Xiaofeng Gao, Wen Xu, Fan Wu, and Guihai Chen</i>	
Dual-centric Data Center Network Architectures	679
<i>Dawei Li, Jie Wu, Zhiyong Liu, and Fa Zhang</i>	
Slowing Little Quickens More: Improving DCTCP for Massive Concurrent Flows	689
<i>Mao Miao, Peng Cheng, Fengyuan Ren, and Ran Shu</i>	

Session 7B: Large Scale Systems

Study on Partitioning Real-World Directed Graphs of Skewed Degree Distribution	699
<i>Jie Yan, Guangming Tan, and Ninghui Sun</i>	
GPSA: A Graph Processing System with Actors	709
<i>Jianhua Sun, Dongwei Zhou, Hao Chen, Cheng Chang, Zhiwen Chen, Wentao Li, and Ligang He</i>	
LCIndex: A Local and Clustering Index on Distributed Ordered Tables for Flexible Multi-dimensional Range Queries	719
<i>Chen Feng, Xi Yang, Fan Liang, Xian-He Sun, and Zhiwei Xu</i>	
MAMS: A Highly Reliable Policy for Metadata Service	729
<i>Jiang Zhou, Yong Chen, Weiping Wang, and Dan Meng</i>	

Session 7C: Caches, Shared Memory, and Architectures

Cache Coherence Protocol and Memory Performance of the Intel Haswell-EP Architecture	739
<i>Daniel Molka, Daniel Hackenberg, Robert Schöne, and Wolfgang E. Nagel</i>	
Optimal Cache Partition-Sharing	749
<i>Jacob Brock, Chencheng Ye, Chen Ding, Yechen Li, Xiaolin Wang, and Yingwei Luo</i>	
Characterizing Loop-Level Communication Patterns in Shared Memory	759
<i>Arya Mazaheri, Ali Jannesari, Abdolreza Mirzaei, and Felix Wolf</i>	
Enhancing Garbage Collection Synchronization Using Explicit Bit Barriers	769
<i>Jochen Hollmann, Ruben Titos-Gil, and Per Stenström</i>	

Session 8A: Sensor, DTN, and RFID

Fine-Grained Loss Tomography in Dynamic Sensor Networks	779
<i>Chenhong Cao, Yi Gao, Wei Dong, and Jiajun Bu</i>	
Joint Wireless Charging and Sensor Activity Management in Wireless Rechargeable Sensor Networks	789
<i>Yuan Gao, Cong Wang, and Yuanyuan Yang</i>	
A Buffer Management Strategy on Spray and Wait Routing Protocol in DTNs	799
<i>En Wang, Yongjian Yang, Jie Wu, and Wenbin Liu</i>	
Towards Constant-Time Cardinality Estimation for Large-Scale RFID Systems	809
<i>Binbin Li, Yuan He, and Wenyuan Liu</i>	

Session 8B: Scientific Applications and Scaling

SZTS: A Novel Big Data Transportation System Benchmark Suite	819
<i>Wen Xiong, Zhibin Yu, Lieven Eeckhout, Zhengdong Bei, Fan Zhang, and Chengzhong Xu</i>	
GEM: A Framework for Developing Shared-Memory Parallel Genomic Applications on Memory Constrained Architectures	829
<i>Mucahid Kutlu and Gagan Agrawal</i>	
A Responsive Knapsack-Based Algorithm for Resource Provisioning and Scheduling of Scientific Workflows in Clouds	839
<i>Maria A. Rodriguez and Rajkumar Buyya</i>	
A Testing Engine for High-Performance and Cost-Effective Workflow Execution in the Cloud	849
<i>V.K. Pallipuram, T. Estrada, and M. Tauber</i>	

Session 8C: Programming and Performance Modeling

GLAF: A Visual Programming and Auto-tuning Framework for Parallel Computing	859
<i>Konstantinos Krommydas, Ruchira Sasanka, and Wu-Chun Feng</i>	
DPX10: An Efficient X10 Framework for Dynamic Programming Applications	869
<i>Chen Wang, Ce Yu, Jizhou Sun, and Xiangfei Meng</i>	
LPM: Concurrency-Driven Layered Performance Matching	879
<i>Yu-Hang Liu and Xian-He Sun</i>	
Region-Based May-Happen-in-Parallel Analysis for C Programs	889
<i>Peng Di, Yulei Sui, Ding Ye, and Jingling Xue</i>	

Session 9A: Smartphones and Crowdsensing

Towards Redundancy-Aware Data Utility Maximization in Crowdsourced Sensing with Smartphones	899
<i>Juan Li, Yanmin Zhu, Jiadi Yu, Qian Zhang, and Lionel M. Ni</i>	
Privacy Preserving Market Schemes for Mobile Sensing	909
<i>Yuan Zhang, Yunlong Mao, He Zhang, and Sheng Zhong</i>	
LMDD: Light-Weight Magnetic-Based Door Detection with Your Smartphone	919
<i>Yiyang Zhao, Chen Qian, Liangyi Gong, Zhenhua Li, and Yunhao Liu</i>	

Session 9B: Cloud Applications

Elastic and Efficient Virtual Network Provisioning for Cloud-Based Multi-tier Applications	929
<i>Meng Shen, Ke Xu, Fan Li, Kun Yang, Liehuang Zhu, and Lei Guan</i>	
Optimizing MapReduce Based on Locality of K-V Pairs and Overlap between Shuffle and Local Reduce	939
<i>Jianjiang Li, Jie Wu, Xiaolei Yang, and Shiqi Zhong</i>	
PIC: Enable Large-Scale Privacy Preserving Content-Based Image Search on Cloud	949
<i>Lan Zhang, Taeho Jung, Puchun Feng, Kebin Liu, Xiang-Yang Li, and Yunhao Liu</i>	

Session 9C: GPU Code Generation and Performance Tuning

Automatic OpenCL Code Generation for Multi-device Heterogeneous Architectures	959
<i>Pei Li, Elisabeth Brunet, François Trahay, Christian Parrot, Gaël Thomas, and Raymond Namyst</i>	

Generating Efficient Tensor Contractions for GPUs	969
<i>Thomas Nelson, Axel Rivera, Prasanna Balaprakash, Mary Hall, Paul D. Hovland, Elizabeth Jessup, and Boyana Norris</i>	
Nested Parallelism on GPU: Exploring Parallelization Templates for Irregular Loops and Recursive Computations	979
<i>Da Li, Hancheng Wu, and Michela Becchi</i>	

Author Index