

2013 IEEE 27th International Symposium on Parallel and Distributed Processing

(IPDPS 2013)

**Cambridge, Massachusetts, USA
20 – 24 May 2013**

Pages 1-660



**IEEE Catalog Number: CFP13023-POD
ISBN: 978-1-4673-6066-1**

2013 IEEE 27th International Symposium on Parallel & Distributed Processing

IPDPS 2013

Table of Contents

Message from the General Co-Chairs.....	xiv
Message from the Program Chair.....	xvi
Message from the Steering Co-Chairs.....	xviii
IPDPS 2013 Organization.....	xix
IPDPS 2013 Technical Program.....	xxv
IPDPS 2013 Reviewers.....	xxvii

Keynote 1

Exascale Computing—A Fact or a Fiction?	3
<i>Shekhar Borkar</i>	

Session 1: Checkpointing

Adaptive Incremental Checkpointing via Delta Compression for Networked Multicore Systems	7
<i>Itthichok Jangjaimon and Nian-Feng Tzeng</i>	
Towards Scalable Checkpoint Restart: A Collective Inline Memory Contents Deduplication Proposal	19
<i>Bogdan Nicolae</i>	
Optimizing Checkpoints Using NVM as Virtual Memory	29
<i>Sudarsun Kannan, Ada Gavrilovska, Karsten Schwan, and Dejan Milojicic</i>	
On Closed Nesting and Checkpointing in Fault-Tolerant Distributed Transactional Memory	41
<i>Aditya Dhoke, Binoy Ravindran, and Bo Zhang</i>	

Session 2: Cloud Computing

Reliable Service Allocation in Clouds	55
<i>Olivier Beaumont, Lionel Eyraud-Dubois, and Hubert Larchevêque</i>	
Scaling and Scheduling to Maximize Application Performance within Budget Constraints in Cloud Workflows	67
<i>Ming Mao and Marty Humphrey</i>	

Optimizing Resource Allocation while Handling SLA Violations in Cloud Computing Platforms	79
<i>Lionel Eyraud-Dubois and Hubert Larchevêque</i>	
V-Cache: Towards Flexible Resource Provisioning for Multi-tier Applications in IaaS Clouds	88
<i>Yanfei Guo, Palden Lama, Jia Rao, and Xiaobo Zhou</i>	

Session 3: Hybrid Systems

High-throughput Analysis of Large Microscopy Image Datasets on CPU-GPU Cluster Platforms	103
<i>George Teodoro, Tony Pan, Tahsin M. Kurc, Jun Kong, Lee A.D. Cooper, Norbert Podhorszki, Scott Klasky, and Joel H. Saltz</i>	
High Performance FFT Based Poisson Solver on a CPU-GPU Heterogeneous Platform	115
<i>Jing Wu and Joseph Jaja</i>	
Design and Implementation of the Linpack Benchmark for Single and Multi-node Systems Based on Intel® Xeon Phi™ Coprocessor	126
<i>Alexander Heinecke, Karthikeyan Vaidyanathan, Mikhail Smeĭyanskiy, Alexander Kobotov, Roman Dubtsov, Greg Henry, Aniruddha G. Shet, George Chrysos, and Pradeep Dubey</i>	
Self-Adaptive OmpSs Tasks in Heterogeneous Environments	138
<i>Judit Planas, Rosa M. Badia, Eduard Ayguadé, and Jesús Labarta</i>	

Session 4: Networks

RAIR: Interference Reduction in Regionalized Networks-on-Chip	153
<i>Lizhong Chen, Kai Hwang, and Timothy M. Pinkston</i>	
An Analytical Performance Model for Partitioning Off-Chip Memory Bandwidth	165
<i>Ruisheng Wang, Lizhong Chen, and Timothy Mark Pinkston</i>	
A Case for Handshake in Nanophotonic Interconnects	177
<i>Lei Wang, Jagadish Jayabalan, Minseon Ahn, Haiyin Gu, Ki Hwan Yum, and Eun Jung Kim</i>	
P-sync: A Photonically Enabled Architecture for Efficient Non-local Data Access	189
<i>David Whelihan, Jeffrey J. Hughes, Scott M. Sawyer, Eric Robinson, Michael Wolf, Sanjeev Mohindra, Julie Mullen, Anna Klein, Michelle Beard, Nadya T. Bliss, Johnnie Chan, Robert Hendry, Keren Bergman, and Luca P. Carloni</i>	

Session 5: Graph Algorithms

Optimizations and Analysis of BSP Graph Processing Models on Public Clouds	203
<i>Mark Redekopp, Yogesh Simmhan, and Viktor K. Prasanna</i>	
Parallel Label-Setting Multi-objective Shortest Path Search	215
<i>Peter Sanders and Lawrence Mandow</i>	
Multi-threaded Graph Partitioning	225
<i>Dominique Lasalle and George Karypis</i>	
High-Productivity and High-Performance Analysis of Filtered Semantic Graphs	237
<i>Aydin Buluç, Erika Duriakova, Armando Fox, John R. Gilbert, Shoaib Kamil, Adam Lugowski, Leonid Oliker, and Samuel Williams</i>	

Session 6: Numerical Analysis

Virtual Systolic Array for QR Decomposition	251
<i>Jakub Kurzak, Piotr Luszczek, Mark Gates, Ichitaro Yamazaki, and Jack Dongarra</i>	
Communication-Optimal Parallel Recursive Rectangular Matrix Multiplication	261
<i>James Demmel, David Eiahu, Armando Fox, Shoaib Kamil, Benjamin Lipshitz, Oded Schwartz, and Omer Spillinger</i>	
Improving the Performance of the Symmetric Sparse Matrix-Vector Multiplication in Multicore	273
<i>Theodoros Gkountouvas, Vasileios Karakasis, Kornilios Kourtis, Georgios Goumas, and Nectarios Koziris</i>	
Automated Rapid Prototyping of Regular Grid-Based Numerical Applications Using Generalized Elemental Subroutines	284
<i>Yingchong Situ, Ye Wang, and Zhiyuan Li</i>	

Session 7: Parallel I/O and Server Software

A Transparent Collective I/O Implementation	297
<i>Yongen Yu, Jingjin Wu, Zhiling Lan, Douglas H. Rudd, Nikolay Y. Gnedin, and Andrey Kravtsov</i>	
A Visual Network Analysis Method for Large-Scale Parallel I/O Systems	308
<i>Carmen Sigovan, Chris Muelder, Kwan-Liu Ma, Jason Cope, Kamil Iskra, and Robert Ross</i>	
FlexIO: I/O Middleware for Location-Flexible Scientific Data Analytics	320
<i>Fang Zheng, Hongbo Zou, Greg Eisenhauer, Karsten Schwan, Matthew Wolf, Jai Dayal, Tuan-Anh Nguyen, Jianting Cao, Hasan Abbasi, Scott Klasky, Norbert Podhorszki, and Hongfeng Yu</i>	
Burstiness-aware Server Consolidation via Queuing Theory Approach in a Computing Cloud	332
<i>Zhaoyi Luo and Zhuzhong Qian</i>	

Session 8: Parallel I/O and File Systems

Pattern-Direct and Layout-Aware Replication Scheme for Parallel I/O Systems	345
<i>Yanlong Yin, Jibing Li, Jun He, Xian-He Sun, and Rajeev Thakur</i>	
Disk-Cache and Parallelism Aware I/O Scheduling to Improve Storage System Performance	357
<i>Ramya Prabhakar, Mahmut Kandemir, and Myoungsoo Jung</i>	
Efficient and Scalable Retrieval Techniques for Global File Properties	369
<i>Dong H. Ahn, Michael J. Brim, Bronis R. de Supinski, Todd Gamblin, Gregory L. Lee, Matthew P. Legendre, Barton P. Miller, Adam Moody, and Martin Schulz</i>	
iBridge: Improving Unaligned Parallel File Access with Solid-State Drives	381
<i>Xuechen Zhang, Ke Liu, Kei Davis, and Song Jiang</i>	

Session 9: Potpourri Algorithms 1

Locally Self-Adjusting Tree Networks	395
<i>Chen Avin, Bernhard Haeupler, Zvi Lotker, Christian Scheideler, and Stefan Schmid</i>	
A Network Configuration Algorithm Based on Optimization of Kirchhoff Index	407
<i>Adam Hackett, Deepak Ajwani, Shoukat Ali, Steve Kirkland, and John P. Morrison</i>	

Malleable Sorting	418
<i>Patrick Flick, Peter Sanders, and Jochen Speck</i>	

Adapting Particle Filter Algorithms to Many-Core Architectures	427
<i>Mehdi Chitchian, Alexander S. van Amesfoort, Andrea Simonetto, Tamás Keviczky, and Henk J. Sips</i>	

Session 10: GPU Scheduling

Guided Region-Based GPU Scheduling: Utilizing Multi-thread Parallelism to Hide Memory Latency	441
---	-----

Jianmin Chen, Xi Tao, Zhen Yang, Jih-Kwon Peir, Xiaoyuan Li, and Shih-Lien Lu

Optimizing and Auto-Tuning Iterative Stencil Loops for GPUs with the In-Plane Method	452
--	-----

Wai Teng Tang, Wen Jun Tan, Ratna Krishnamoorthy, Yi Wen Wong, Shyh-Hao Kuo, Rick Siow Mong Goh, Stephen John Turner, and Weng-Fai Wong

Data-Driven Versus Topology-driven Irregular Computations on GPUs	463
---	-----

Rupesh Nasre, Martin Burtscher, and Keshav Pingali

HQL: A Scalable Synchronization Mechanism for GPUs	475
--	-----

Ayse Yilmazer and David Kaeli

Session 11: Fault Tolerance and Contention Resolution

Pluggable Watchdog: Transparent Failure Detection for MPI Programs	489
--	-----

Keun Soo Yim, Zbigniew Kalbarczyk, and Ravishankar K. Iyer

Improving the Computing Efficiency of HPC Systems Using a Combination of Proactive and Preventive Checkpointing	501
---	-----

Mohamed Slim Bouguerra, Ana Gainaru, Leonardo Bautista Gomez, Franck Cappello, Satoshi Matsuoka, and Naoya Maruyama

CASTED: Core-Adaptive Software Transient Error Detection for Tightly Coupled Cores	513
--	-----

Konstantina Mitropoulou, Vasileios Porpodas, and Marcelo Cintra

Contention Resolution in a Non-synchronized Multiple Access Channel	525
---	-----

Gianluca De Marco and Dariusz R. Kowalski

Session 12: Communication and Routing 1

Generalized Hierarchical All-to-All Exchange Patterns	537
---	-----

Bogdan Prisacari, German Rodriguez, and Cyriel Minkenbergh

Minimizing Communication in All-Pairs Shortest Paths	548
--	-----

Edgar Solomonik, Aydin Buluç, and James Demmel

Programmable and Scalable Reductions on Clusters	560
--	-----

Jan Ciesko, Javier Bueno, Nikola Puzovic, Alex Ramirez, Rosa M. Badia, and Jesús Labarta

JVM-Bypass for Efficient Hadoop Shuffling	569
---	-----

Yandong Wang, Cong Xu, Xiaobing Li, and Weikuan Yu

Symposium Tutorial

Resource Management in VMware Powered Cloud: Concepts and Techniques	581
<i>Pradeep Padala</i>	

Keynote 2

Communication-Avoiding Algorithms for Linear Algebra and Beyond	585
<i>James Demmel</i>	

Session 13: Data Centers

Oversubscription Bounded Multicast Scheduling in Fat-Tree Data Center Networks	589
<i>Zhiyang Guo, Jun Duan, and Yuanyuan Yang</i>	
Replicate and Bundle (RnB) — A Mechanism for Relieving Bottlenecks in Data Centers	601
<i>Shachar Raindel and Yitzhak Birk</i>	
Profit Aware Load Balancing for Distributed Cloud Data Centers	611
<i>Shuo Liu, Shaolei Ren, Gang Quan, Ming Zhao, and Shangping Ren</i>	
Joint Host-Network Optimization for Energy-Efficient Data Center Networking	623
<i>Hao Jin, Tosmate Cheochnngarn, Dmita Levy, Alex Smith, Deng Pan, Jason Liu, and Niki Pissinou</i>	

Session 14: Energy Modeling and Scheduling

Energy-Efficient Scheduling for Best-Effort Interactive Services to Achieve High Response Quality	637
<i>Zhihui Du, Hongyang Sun, Yuxiong He, Yu He, David A. Bader, and Huazhe Zhang</i>	
Perfect Strong Scaling Using No Additional Energy	649
<i>James Demmel, Andrew Gearhart, Benjamin Lipshitz, and Oded Schwartz</i>	
A Roofline Model of Energy	661
<i>Jee Whan Choi, Daniel Bedard, Robert Fowler, and Richard Vuduc</i>	
A Simplified and Accurate Model of Power-Performance Efficiency on Emergent GPU Architectures	673
<i>Shuaiwen Song, Chunyi Su, Barry Rountree, and Kirk W. Cameron</i>	

Session 15: Communication and Routing 2

Acceleration of an Asynchronous Message Driven Programming Paradigm on IBM Blue Gene/Q	689
<i>Sameer Kumar, Yanhua Sun, and Laximant V. Kalé</i>	
Communication-Based Mapping Using Shared Pages	700
<i>Matthias Diener, Eduardo H.M. Cruz, and Philippe O.A. Navaux</i>	
Integrating Asynchronous Task Parallelism with MPI	712
<i>Sanjay Chatterjee, Sagnak Tasirlar, Zoran Budimlic, Vincent Cavé, Milind Chabbi, Max Grossman, Vivek Sarkar, and Yonghong Yan</i>	

DTN-FLOW: Inter-Landmark Data Flow for High-Throughput Routing in DTNs	726
<i>Kang Chen and Haiying Shen</i>	

Session 16: Peer to Peer Systems

WHATSUP: A Decentralized Instant News Recommender	741
<i>Antoine Boutet, Davide Frey, Rachid Guerraoui, Arnaud Jégou, and Anne-Marie Kermarrec</i>	
Crowdsourcing under Real-Time Constraints	753
<i>Ioannis Boutsis and Vana Kalogeraki</i>	
Replication-Based Load Balancing in Distributed Content-Based Publish/Subscribe	765
<i>Weixiong Rao, Chao Chen, Pan Hui, and Sasu Tarkoma</i>	
ZHT: A Light-Weight Reliable Persistent Dynamic Scalable Zero-Hop Distributed Hash Table	775
<i>Tonglin Li, Xiaobing Zhou, Kevin Brandstatter, Dongfang Zhao, Ke Wang, Anupam Rajendran, Zhao Zhang, and Ioan Raicu</i>	

Session 17: Programming Frameworks

A Theoretical Framework for Algorithm-Architecture Co-design	791
<i>Kenneth Czechowski and Richard Vuduc</i>	
Wait-free Hyperobjects for Task-Parallel Programming Systems	803
<i>Martin Wimmer</i>	
Cyclops Tensor Framework: Reducing Communication and Eliminating Load Imbalance in Massively Parallel Contractions	813
<i>Edgar Solomonik, Devin Matthews, Jeff Hammond, and James Demmel</i>	
Scaling Techniques for Massive Scale-Free Graphs in Distributed (External) Memory	825
<i>Roger Pearce, Maya Gokhale, and Nancy M. Amato</i>	

Session 18: Scheduling 1

Scheduling Tree-Shaped Task Graphs to Minimize Memory and Makespan	839
<i>Loris Marchal, Oliver Sinnen, and Frédéric Vivien</i>	
On Graphs, GPUs, and Blind Dating: A Workload to Processor Matchmaking Quest	851
<i>Abdullah Gharaibeh, Lauro Beltrão Costa, Elizeu Santos-Neto, and Matei Ripeanu</i>	
Non Linear Divisible Loads: There is No Free Lunch	863
<i>Olivier Beaumont, Hubert Larchevêque, and Loris Marchal</i>	
SIPMaP: A Tool for Modeling Irregular Parallel Computations in the Super Instruction Architecture	874
<i>Nakul Jindal, Victor Lotrich, Erik Deumens, and Beverly A. Sanders</i>	

Symposium Panel

Big Data in 10 Years	887
<i>Raghu Ramakrishnan</i>	

Keynote 3

HPC Cloud Bad; HPC in the Cloud Good	891
<i>Josh Simons</i>	

Plenary Session: Best Papers

Implementing a Blocked Aasen’s Algorithm with a Dynamic Scheduler on Multicore Architectures	895
<i>Grey Ballard, Dulceneia Becker, James Demmel, Jack Dongarra, Alex Druinsky, Inon Peled, Oded Schwartz, Sivan Toledo, and Ichitaro Yamazaki</i>	
DLOOP: A Flash Translation Layer Exploiting Plane-Level Parallelism	908
<i>Abdul R. Abdurrab, Tao Xie, and Wei Wang</i>	
Exploring Traditional and Emerging Parallel Programming Models Using a Proxy Application	919
<i>Ian Karlin, Abhinav Bhatnagar, Jeff Keasler, Bradford L. Chamberlain, Jonathan Cohen, Zachary Devito, Riyaz Haque, Dan Laney, Edward Luke, Felix Wang, David Richards, Martin Schulz, and Charles H. Still</i>	
Extending the Generality of Molecular Dynamics Simulations on a Special-Purpose Machine	933
<i>Daniele P. Scarpazza, Douglas J. Jerardi, Adam K. Lerer, Kenneth M. Mackenzie, Albert C. Pan, Joseph A. Bank, Edmond Chow, Ron O. Dror, J.P. Grossman, Daniel Killebrew, Mark A. Moraes, Cristian Predescu, John K. Salmon, and David E. Shaw</i>	

Session 19: Scheduling 2

Algorithms for the Thermal Scheduling Problem	949
<i>Koyel Mukherjee, Samir Khuller, and Amol Deshpande</i>	
Lock-Free and Wait-Free Slot Scheduling Algorithms	961
<i>Pooja Aggarwal and Smruti R. Sarangi</i>	
Distributed Algorithms for Scheduling on Line and Tree Networks with Non-uniform Bandwidths	973
<i>Venkatesan T. Chakaravarthy, Anamitra R. Choudhury, Sambuddha Roy, and Yogish Sabharwal</i>	
Analysis of Randomized Work Stealing with False Sharing	985
<i>Richard Cole and Vijaya Ramachandran</i>	

Session 20: GPU Software

Extending OpenSHMEM for GPU Computing	1001
<i>S. Potluri, D. Bureddy, H. Wang, H. Subramoni, and D.K. Panda</i>	
Deploying Graph Algorithms on GPUs: An Adaptive Solution	1013
<i>Da Li and Michela Becchi</i>	
GPU-based Runtime Verification	1025
<i>Shay Berkovich, Borzoo Bonakdarpour, and Sebastian Fischmeister</i>	
Kernel Specialization for Improved Adaptability and Performance on Graphics Processing Units (GPUs)	1037
<i>Nicholas Moore, Miriam Leeser, and Laurie Smith King</i>	

Session 21: Scientific Computing

The Bounded Data Reuse Problem in Scientific Workflows	1051
<i>Mohsen Zohrevandi and Rida A. Bazzi</i>	
Performance Analysis of the Lattice Boltzmann Model Beyond Navier-Stokes	1063
<i>Amanda Peters Randles, Vivek Kale, Jeff Hammond, William Gropp, and Efthimos Kaxiras</i>	
A Communication-Optimal N-Body Algorithm for Direct Interactions	1075
<i>Michael Driscoll, Evangelos Georganas, Penporn Koanantakool, Edgar Solomonik, and Katherine Yelick</i>	
Exploring SIMD for Molecular Dynamics, Using Intel® Xeon® Processors and Intel® Xeon Phi™ Coprocessors	1085
<i>Simon J. Pennycook, Chris J. Hughes, M. Smelyanskiy, and S.A. Jarvis</i>	

Session 22: Wireless and Sensor Systems

Multi-vehicle Coordination for Wireless Energy Replenishment in Sensor Networks	1101
<i>Cong Wang, Ji Li, Fan Ye, and Yuanyuan Yang</i>	
On Feasibility of Fingerprinting Wireless Sensor Nodes Using Physical Properties	1112
<i>Xiaowei Mei, Donggang Liu, Kun Sun, and Dingbang Xu</i>	
Distributed Algorithms for Joint Routing and Frame Aggregation in 802.11n Wireless Mesh Networks	1122
<i>Dawei Gong and Yuanyuan Yang</i>	
Distributed Low-Latency Out-of-Order Event Processing for High Data Rate Sensor Streams	1133
<i>Christopher Mutschler and Michael Philippsen</i>	

Session 23: Potpourri Algorithms 2

Agreement via Symmetry Breaking: On the Structure of Weak Subconsensus Tasks	1147
<i>Armando Castañeda, Sergio Rajsbaum, and Michel Raynal</i>	
A Multi-partitioning Approach to Building Fast and Accurate Counting Bloom Filters	1159
<i>Kun Huang, Jie Zhang, Dafang Zhang, Gaogang Xie, Kave Salamatian, Alex X. Liu, and Wei Li</i>	
Composing Relaxed Transactions	1171
<i>Vincent Gramoli, Rachid Guerraoui, and Mihai Letia</i>	
Throughput Enhancement through Selective Time Sharing and Dynamic Grouping	1183
<i>Junliang Chen, Bing Bing Zhou, Chen Wang, Peng Lu, Penghao Wang, and Albert Y. Zomaya</i>	

Session 24: Potpourri Applications

Novel Parallelization Schemes for Large-Scale Likelihood-based Phylogenetic Inference	1195
<i>Alexandros Stamatakis and Andre J. Aberer</i>	
Integrating Online Compression to Accelerate Large-Scale Data Analytics Applications	1205
<i>Tekin Bicer, Jian Yin, David Chiu, Gagan Agrawal, and Karen Schuchardt</i>	

Massively Parallel Model of Extended Memory Use in Evolutionary Game Dynamics	1217
<i>Amanda Peters Randles, David G. Rand, Christopher Lee, Greg Morrisett, Jayanta Sircar, Martin A. Nowak, and Hanspeter Pfister</i>	

Early Experience on the Blue Gene/Q Supercomputing System	1229
<i>Vitali Morozov, Kalyan Kumaran, Venkatram Vishwanath, Jiayuan Meng, and Michael E. Papka</i>	

Session 25: Potpourri Systems

Adaptive Cache Bypassing for Inclusive Last Level Caches	1243
<i>Saurabh Gupta, Hongliang Gao, and Huiyang Zhou</i>	

Hardware-Accelerated Regular Expression Matching with Overlap Handling on IBM PowerENTM Processor	1254
<i>Kubilay Atasu, Florian Doerfler, Jan van Lunteren, and Christoph Hagleitner</i>	

TM-dietlibc: A TM-aware Real-World System Library	1266
<i>Vesna Smiljkovic, Martin Nowack, Nebojša Miletic, Timothy Harris, Osman Ünsal, Adrián Cristal, and Mateo Valero</i>	

Cura: A Cost-Optimized Model for MapReduce in a Cloud	1275
<i>Balaji Palanisamy, Aameek Singh, Ling Liu, and Bryan Langston</i>	

Session 26: Programming Frameworks

A Scalable Heterogeneous Parallelization Framework for Iterative Local Searches	1289
<i>Martin Burtscher and Hassan Rabeti</i>	

XKaaapi: A Runtime System for Data-Flow Task Programming on Heterogeneous Architectures	1299
<i>Thierry Gautier, João V.F. Lima, Nicolas Maillard, and Bruno Raffin</i>	

A Study of the Behavior of Synchronization Methods in Commonly Used Languages and Systems	1309
<i>Daniel Cederman, Bapi Chatterjee, Nhan Nguyen, Yiannis Nikolakopoulos, Marina Papatriantafidou, and Philippos Tsigas</i>	

Managing Asynchronous Operations in Coarray Fortran 2.0	1321
<i>Chaoran Yang, Karthik Murthy, and John Mellor-Crummey</i>	

Author Index