

2016 IEEE 22nd International Conference on Parallel and Distributed Systems (ICPADS 2016)

**Wuhan, China
13-16 December 2016**

Pages 1-624



**IEEE Catalog Number: CFP16036-POD
ISBN: 978-1-5090-5382-7**

**Copyright © 2016 by the Institute of Electrical and Electronics 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 is a print representation of what appears in the IEEE Digital Library. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP16036-POD
ISBN (Print-On-Demand):	978-1-5090-5382-7
ISBN (Online):	978-1-5090-4457-3
ISSN:	1521-9097

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

2016 IEEE 22nd International Conference on Parallel and Distributed Systems

ICPADS 2016

Table of Contents

Foreword.....	xviii
Organizing Committee.....	xx
Workshop Organizing Committees.....	xxvi
Keynotes.....	xxvii

ICPADS 2016

Session 1: Mobile, Sensor, and Ubiquitous Computing

A Distributed Auction Approach to Crowdsourced Sensing over Smartphones	1
<i>Pingyi Luo, Yanmin Zhu, Jia Peng, and Jiadi Yu</i>	
A Hybrid Approach Based on Collaborative Filtering to Recommending Mobile Apps	8
<i>Xia Wu and Yanmin Zhu</i>	
Bat with Good Eyesight: Using Acoustic Signal and Image to Achieve Accurate Indoor Localization	16
<i>Rui Xi, Yujun Li, Daibo Liu, Siwei Luo, and Mengshu Hou</i>	
Delay Tolerant Routing for Cognitive Radio Vehicular Ad Hoc Networks	24
<i>Jing Wang, Huyin Zhang, and Xing Tang</i>	
Distributed Optimal Source Coding Rate Allocation for Data Aggregation in Wireless Sensor Networks	32
<i>Yang Yang, Songtao Guo, and Yuanyuan Yang</i>	
Distributed Real-Time Pricing Scheme for Local Power Supplier in Smart Community	40
<i>Lan Mu, Nuo Yu, Hejiao Huang, Hongwei Du, and Xiaohua Jia</i>	
Drone-Based Wireless Relay Using Online Tensor Update	48
<i>Yao Xie, Xiao-Yang Liu, Linghe Kong, Fan Wu, Guihai Chen, and Athanasios V. Vasilakos</i>	

EHR: Routing Protocol for Energy Harvesting Wireless Sensor Networks	56
<i>Yifeng Cao, Xiao-Yang Liu, Linghe Kong, Min-You Wu, and Muhammad Khurram Khan</i>	
Enabling Mobile Device Coordination over Distributed Shared Memory	64
<i>Maosen Huang, Hengfeng Wei, and Yu Huang</i>	
Identifying a New Non-Linear CSI Phase Measurement Error with Commodity WiFi Devices	72
<i>Yiwei Zhuo, Hongzi Zhu, and Hua Xue</i>	
NoiseSense: A Crowd Sensing System for Urban Noise Mapping Service	80
<i>Zhaokun Qin and Yanmin Zhu</i>	
OR-Play: An Optimal Relay Placement Scheme for High-Quality Wireless Network Services	88
<i>Ailun Song, Xiaofeng Gao, Fan Wu, Linghe Kong, and Guihai Chen</i>	
Physical Object Model for Smart Environments with Temporal Capability	96
<i>Yuanzheng Wen and Muthucumaru Maheswaran</i>	
ppNav: Peer-to-Peer Indoor Navigation for Smartphones	104
<i>Zuwei Yin, Chenshu Wu, Zheng Yang, Nicholas Lane, and Yunhao Liu</i>	
PUZZLE: Enhancing Throughput by Covering Neighbor's Blanks	112
<i>Jiayue Li</i>	
QoE-Driven Cross-Layer Design for Device-to-Device Video Delivery	120
<i>Lingbao Ye, Yahui Hu, Hongyan Tan, and Liming Wang</i>	
Resource Scheduling Based on Improved FCM Algorithm for Mobile Cloud Computing	128
<i>Wu Hong-Qiang, Li Xiao-Yong, Fang Bin-Xing, and Wang Yi-Ping</i>	
Smart-DJ: Context-Aware Personalization for Music Recommendation on Smartphones	133
<i>Chengkun Jiang and Yuan He</i>	
Smartphone Virtualization	141
<i>Tzi-cker Chiueh, Houcheng Lin, Ares Chao, Anthony, Tan-Gen Wu, Chieh-Min Wang, and Yu-Sung Wu</i>	
TIP: Time-Efficient Identification Protocol for Unknown RFID Tags Using Bloom Filters	151
<i>Yuming Qian, Zongjian He, and Daqiang Zhang</i>	

Session 2: P2P and Web-Based Computing

A Peer-to-Peer File Sharing System over Named Data Networking	159
<i>Jian Shi, Xuewei Piao, Lihua Li, Yunbo Xun, and Kai Lei</i>	

A Potential Field Based Framework for Publish/Subscribe Service in P2P Cloud	168
<i>Jiangfeng Li, Zhenyu Liao, Zehong Zhou, Jian Xu, Qinpei Zhao, and Chenxi Zhang</i>	
A Weighted Network Model Based on Node Fitness Dynamic Evolution	176
<i>Xiaoheng Deng, You Wu, Deng Li, and Honggang Zhang</i>	
An Efficient Parallel Approach of Parsing and Indexing for Large-Scale XML Datasets	184
<i>Kunfang Song, Hongwei Lu, and Xiao Qin</i>	
DMNS: A Framework to Dynamically Monitor Simulated Network	192
<i>Lun Li, Zhiyu Hao, Yongzheng Zhang, Yaqiong Peng, Xin Deng, and Zhenxi Sun</i>	
DOCO: An Efficient Event Matching Algorithm in Content-Based Publish/Subscribe Systems	200
<i>Jingli Yang, Jing Fan, and Shouda Jiang</i>	
DScheduler: Dynamic Network Scheduling Method for MapReduce in Distributed Controllers	208
<i>Tianxiao Liu, Yi Liu, Ping Song, and Depei Qian</i>	
Garlic Cast: Lightweight and Decentralized Anonymous Content Sharing	216
<i>Chen Qian, Junjie Shi, Zihao Yu, Ye Yu, and Sheng Zhong</i>	
Modeling Traffic of Big Data Platform for Large Scale Datacenter Networks	224
<i>Zhen Xie, Zheng Cao, Zhan Wang, Dawei Zang, En Shao, and Ninghui Sun</i>	
SHSA: A Method of Network Verification with Stateful Header Space Analysis	232
<i>Yufan Yang, Xinli Huang, Shang Cheng, Shiyun Chen, and Peijin Cong</i>	
Study of Intra- and Interjob Interference on Torus Networks	239
<i>Xu Yang, John Jenkins, Misbah Mubarak, Xin Wang, Robert B. Ross, and Zhiling Lan</i>	
Virtual Path Assignment Based on Load Balancing for SDNs	247
<i>Xiaomao Wang, Chuanhe Huang, Kai He, Xiying Fan, and Bin Fu</i>	

Session 3: Cyber-Physical Systems

A Comparison of Road-Network-Constrained Trajectory Compression Methods	256
<i>Yudian Ji, Hao Liu, Xiaoying Liu, Ye Ding, and Wuman Luo</i>	
A Complicated Task Solution Scheme Based on Node Cooperation for Wireless Sensor Networks	264
<i>Jinfang Jiang, Guangjie Han, and Chunsheng Zhu</i>	
A Reliable Depth-Based Routing Protocol with Network Coding for Underwater Sensor Networks	270
<i>Boyu Diao, Yongjun Xu, Qi Wang, Zhao Chen, Chao Li, Zhulin An, and Guangjie Han</i>	

Accomplishing Information Consistency under OSPF in General Networks	278
<i>Jianhui Lv, Xingwei Wang, Min Huang, Fuliang Li, Keqin Li, and Hui Cheng</i>	
An Optimized RM Algorithm by Task Affinity on Multi-Core Processor	286
<i>Ying Li, Jianwei Niu, Jiong Zhang, Mohammed Atiquzzaman, and Xiang Long</i>	
An Optimized RPL Protocol for Wireless Sensor Networks	294
<i>Zhenfei Wang, Liying Zhang, Zhiyun Zheng, and Junfeng Wang</i>	
Efficient Privacy-Preserving kNN Query in Mobile Sensing Systems	300
<i>Shiming He, Ying Li, Kun Xie, Xin Su, and Weini Zeng</i>	
FPGA-Based Parallel Implementation of SURF Algorithm	308
<i>Wenjie Chen, Shuaishuai Ding, Zhilei Chai, Daojing He, Weihua Zhang, Guanhua Zhang, Qiwei Peng, and Wang Luo</i>	
Hierarchically Social-Aware Incentivized Caching for D2D Communications	316
<i>Wenting Zhi, Konglin Zhu, Yuan Zhang, and Lin Zhang</i>	
HMF: Heatmap and WiFi Fingerprint-Based Indoor Localization with Building Layout Consideration	324
<i>Xiting Liu, Banghui Lu, Jianwei Niu, Lei Shu, and Yuanfang Chen</i>	
mtCloudSim: A Flow-Level Network Simulator for Multi-Tenant Cloud	332
<i>Junjie Xie and Yuhui Deng</i>	
Optimized Virtual Network Functions Migration for NFV	340
<i>Jing Xia, Zhiping Cai, and Ming Xu</i>	
PROAR: A Weak Consistency Model for Ceph	347
<i>Jiayuan Zhang, Yongwei Wu, and Yeh-Ching Chung</i>	
TaxiCast: Efficient Broadcasting of Multimedia Advertisements in Vehicular Ad-Hoc Networks	354
<i>Peng Liu, Jia Xu, Biao Xu, and Tingting Fu</i>	
Very Low-Resolution Face Recognition via Semi-Coupled Locality-Constrained Representation	362
<i>Tao Lu, Wei Yang, Yanduo Zhang, Xiaolin Li, and Zixiang Xiong</i>	
Watch Traffic in the Sky: A Method for Path Selection in Packet Transmission between V2V from Macro Perspective	368
<i>Wen Cui, Xiaoqing Gong, Chen Liu, Dan Xu, Zhuang Yin, Xiaojiang Chen, and Dingyi Fang</i>	

Session 4: Security and Trustworthy Computing

A Provably Secure Blind Signature Based on Coding Theory	376
<i>Siyuan Chen, Peng Zeng, and Kim-Kwang Raymond Choo</i>	
A Users Collaborative Scheme for Location and Query Privacy	383
<i>Zhu Haitao, Zhang Lei, Feng Weimiao, and Ma Chunguang</i>	

An Approach of Anti-Eavesdropping Linear Network Coding in Wireless Network	391
<i>Zuoting Ning, Dafang Zhang, and Kun Xie</i>	
Automatic Security Bug Classification: A Compile-Time Approach	398
<i>Tie Du, Long Zheng, Shaopeng Chen, and Hai Jin</i>	
Baseline Is Fragile: On the Effectiveness of Stack Pivot Defense	406
<i>Fei Yan, Fan Huang, Lei Zhao, Huirong Peng, and Qian Wang</i>	
CC-Paxos: Integrating Consistency and Reliability in Wide-Area Storage Systems	414
<i>Yili Gong, Chuang Hu, Wentao Ma, and Wenjie Wang</i>	
Multi-User Location Correlation Protection with Differential Privacy	422
<i>Lu Ou, Zheng Qin, Yonghe Liu, Hui Yin, Yupeng Hu, and Hao Chen</i>	
Opportunistic Probe: An Efficient Adaptive Detection Model for Collaborative Intrusion Detection	430
<i>Dali Zhu, Na Pang, Gang Li, and Wenjing Rong</i>	
Outsourcing Large-Scale Systems of Linear Matrix Equations in Cloud Computing	438
<i>Jian Zhang, Yang Yang, and Zhibo Wang</i>	
RuleCache: A Mobility Pattern Based Multi-Level Cache Approach for Location Privacy Protection	448
<i>Qiuwei Yang and Pan Kong</i>	
Towards a Framework to Facilitate the Mobile Advertising Ecosystem	456
<i>Gong Chen, Shouling Ji, and John A. Copeland</i>	

Session 5: Big Data and Massive Storage Systems

A Novel Method of Keyword Query for RDF Data Based on Bipartite Graph	466
<i>Zhiyun Zheng, Yang Ding, Zhentao Wang, and Zhenfei Wang</i>	
A Unified Access Manner for Storage-Class Memory	474
<i>Yuchuan Tian and Fang Wang</i>	
Application-Aware and Software-Defined SSD Scheme for Tencent Large-Scale Storage System	482
<i>Jianquan Zhang, Dan Feng, Jianlin Gao, Wei Tong, Jingning Liu, Yu Hua, Yang Gao, Caihua Fang, Wen Xia, Feiling Fu, and Yaqing Li</i>	
Catenae: Low Latency Transactions across Multiple Data Centers	491
<i>Ying Liu, Qinjin Wang, and Vladimir Vlassov</i>	
Checking the Inconsistent Data in Concurrent Systems by Petri Nets with Data Operations	501
<i>Dongming Xiang, Guanjun Liu, Chungang Yan, and Changjun Jiang</i>	

CloudBB: Scalable I/O Accelerator for Shared Cloud Storage	509
<i>Tianqi Xu, Kento Sato, and Satoshi Matsuoka</i>	
DEC: An Efficient Deduplication-Enhanced Compression Approach	519
<i>Zijin Han, Wen Xia, Yuchong Hu, Dan Feng, Yucheng Zhang, Yukun Zhou, Min Fu, and Liang Gu</i>	
Exploiting the Data Redundancy Locality to Improve the Performance of Deduplication-Based Storage Systems	527
<i>Suzhen Wu, Xiao Chen, and Bo Mao</i>	
High Throughput Log-Based Replication for Many Small In-Memory Objects	535
<i>Kevin Beineke, Stefan Nothaas, and Michael Schoettner</i>	
Improving MLC Flash Performance with Workload-Aware Differentiated ECC	545
<i>Qianbin Xia and Weijun Xiao</i>	
Improving Write Performance of LSMT-Based Key-Value Store	553
<i>Weitao Zhang, Yinlong Xu, Yongkun Li, and Dinglong Li</i>	
Lifting Wavelet Compression Based Data Aggregation in Big Data Wireless Sensor Networks	561
<i>Ledan Cheng, Songtao Guo, Ying Wang, and Yuanyuan Yang</i>	
masFS: File System Based on Memory and SSD in Compute Nodes for High Performance Computers	569
<i>Xin Liu, Ying Lu, Yutong Lu, Chunjia Wu, and Jieting Wu</i>	
NVMCFS: Complex File System for Hybrid NVM	577
<i>Cai Tao, Niu Dejiao, He Yao, and Zhu Yeqing</i>	
Optimizing Data Placement of MapReduce on Ceph-Based Framework under Load-Balancing Constraint	585
<i>Edwin H.-M. Sha, Yutong Liang, Weiwen Jiang, Xianzhang Chen, and Qingfeng Zhuge</i>	
ScalaRDF: A Distributed, Elastic and Scalable In-Memory RDF Triple Store	593
<i>Chunming Hu, Xixu Wang, Renyu Yang, and Tianyu Wo</i>	
SLA-DO: A SLA-Based Data Distribution Strategy on Multiple Cloud Storage Systems	602
<i>Chang Guo, Ying Li, and Zhonghai Wu</i>	
The Implementation of Supporting Uniform Data Distribution with Software-Defined Storage Service on Heterogeneous Cloud Storage	610
<i>Wei-Hsun Cheng, Chun-I Chiang, Chao-Tung Yang, Shuo-Tsung Chen, and Jung-Chun Liu</i>	
Utilizing SSD to Alleviate Chunk Fragmentation in De-Duplicated Backup Systems	616
<i>Longxin Lin, Kun Xiao, and Wenjie Liu</i>	

Wamalloc: An Efficient Wear-Aware Allocator for Non-Volatile Memory	625
<i>Jiashun Zhu, Sumin Li, and Linpeng Huang</i>	
Zero-Chunk: An Efficient Cache Algorithm to Accelerate the I/O Processing of Data Deduplication	635
<i>Hongyuan Gao, Chentao Wu, Jie Li, and Minyi Guo</i>	
Session 6: Cloud Computing and Services	
A Field-Based Model for Representing Dynamic and Evolving Features of Cloud Services	643
<i>Fu Hou and Xinjun Mao</i>	
A Game-Theoretic Analysis of Pricing Strategies for Competing Cloud Platforms	653
<i>Bing Shi, Yalong Huang, Jinwen Wang, and Shengwu Xiong</i>	
A Task Scheduling Method for Energy-Efficient Cloud Video Surveillance System Using a Time-Clustering-Based Genetic Algorithm	661
<i>Dongping Fu, Yonghua Xiong, Chengda Lu, Min Wu, and Keyuan Jiang</i>	
An Energy-Aware Ant Colony Algorithm for Network-Aware Virtual Machine Placement in Cloud Computing	669
<i>Chuangen Gao, Hua Wang, Linbo Zhai, Yanqing Gao, and Shanwen Yi</i>	
An Online Auction for Deadline-Aware Dynamic Cloud Resource Provisioning	677
<i>Kai He, Chuanhe Huang, Zongpeng Li, Aiwu Shi, and Jiaoli Shi</i>	
Authenticated Spatio-Textual Similarity Joins in Untrusted Cloud Environments	685
<i>Han Yan, Xiang Cheng, Sen Su, Qiying Zhang, and Jianliang Xu</i>	
Benchmarking Sentiment Analysis Approaches on the Cloud	695
<i>Richard O. Sinnott and Shucheng Cui</i>	
Data Dissemination Protocols Based on Opportunistic Sharing for Data Offloading in Mobile Social Networks	705
<i>Na Jiang, Longjiang Guo, Jinbao Li, Meirui Ren, Sisi Cheng, and Xiaodan Guo</i>	
Efficient Snapshot KNN Join Processing for Large Data Using MapReduce	713
<i>Yupeng Hu, Chong Yang, Cun Ji, Yang Xu, and Xueqing Li</i>	
Game Theoretic Energy Allocation for Renewable Powered In-Situ Server Systems	721
<i>Junlong Zhou, Jianfei Chen, Kun Cao, Tongquan Wei, and Mingsong Chen</i>	
GDSW: A General Framework for Distributed Sliding Window over Data Streams	729
<i>Huan Chen, Yijie Wang, Yuan Wang, and Xingkong Ma</i>	

Implementation of an Energy Saving Cloud Infrastructure with Virtual Machine Power Usage Monitoring and Live Migration on OpenStack	737
<i>Tsung-Yueh Wan, Chun-I Chiang, Chao-Tung Yang, Shuo-Tsung Chen, and Jung-Chun Liu</i>	
Improving the Performance of Data Sharing in Dynamic Peer-to-Peer Mobile Cloud	743
<i>Wenhua Xiao, Weidong Bao, Xiaomin Zhu, Wen Zhou, and Peizhong Luy</i>	
On Autonomous Service Migrations in the Cloud for Mobile Accesses	753
<i>Yang Wang, Shuibing He, Fuji Ren, Lujia Wang, and Chengzhong Xu</i>	
Secure Conjunctive Multi-Keyword Search for Multiple Data Owners in Cloud Computing	761
<i>Hui Yin, Zheng Qin, Jixin Zhang, Wenjie Li, Lu Ou, Yupeng Hu, and Keqin Li</i>	
VNF Placement in Hybrid NFV Environment: Modeling and Genetic Algorithms	769
<i>Jiuyue Cao, Yan Zhang, Wei An, Xin Chen, Yanni Han, and Jiyan Sun</i>	

Session 7: Multicore Computing and Parallel/Distributed Architecture

A Fine-Grained Parallel Intra Prediction for HEVC Based on GPU	778
<i>Wenbin Jiang, Ye Chi, Hai Jin, Xiaofei Liao, Yangsong Zhang, and Geyan Ye</i>	
A Fine-Grained Parallel Power Flow Method for Large Scale Grid Based on Lightweight GPU Threads	785
<i>Han Jiang, Deyang Chen, Yalou Li, and Ran Zheng</i>	
Accelerating Spark RDD Operations with Local and Remote GPU Devices	791
<i>Yasuhiro Ohno, Shin Morishima, and Hiroki Matsutani</i>	
Accelerating the Simulation of Thermal Convection in the Earth's Outer Core on Tianhe-2	800
<i>Changmao Wu, Fangfang Liu, Chao Yang, Ligang Li, Haitao Zhao, Yutong Lu, Leisheng Li, and Yunfei Du</i>	
An Energy-Efficient Scheduler for Throughput Guaranteed Jobs on Asymmetric Multi-Core Platforms	810
<i>Ching-Chi Lin, Hsiang-Hsin Li, Jan-Jan Wu, and Pangfeng Liu</i>	
Balanced Parity Update Algorithm with Queueing Length Awareness for RAID Arrays	818
<i>Youxu Chen, Yinlong Xu, Yongkun Li, and Jun Xu</i>	
Compile-Time Automatic Synchronization Insertion and Redundant Synchronization Elimination for GPU Kernels	826
<i>Lifeng Liu, Meilin Liu, Chong-Jun Wang, and Jun Wang</i>	

Efficient Distributed Data Structures for Future Many-Core Architectures	835
<i>Panagiota Fatourou, Nikolaos D. Kallimanis, Eleni Kanellou, Odysseas Makridakis, and Christi Symeonidou</i>	
Enabling Tissue-Scale Cardiac Simulations Using Heterogeneous Computing on Tianhe-2	843
<i>Johannes Langguth, Qiang Lan, Namit Gaur, Xing Cai, Mei Wen, and Chun-Yuan Zhang</i>	
Exploiting Longer SIMD Lanes in Dynamic Binary Translation	853
<i>Ding-Yong Hong, Sheng-Yu Fu, Yu-Ping Liu, Jan-Jan Wu, and Wei-Chung Hsu</i>	
Increasing Lifetime and Security of Phase-Change Memory with Endurance Variation	861
<i>Wen Zhou, Dan Feng, Yu Hua, Jingning Liu, Fangting Huang, and Pengfei Zuo</i>	
Joint Hybrid Frequent Value Cache and Multi-Coding for Data Bus Energy Saving	869
<i>Mingquan Zhang, Zhihua Gan, Jizan Zhang, and Zhimin Gu</i>	
Large Page Address Mapping in Massive Parallel Processor Systems	877
<i>Yichun Sun, Xiaodong Yi, and Hengzhu Liu</i>	
Lightweight Dependency Checking for Parallelizing Loops with Non-Deterministic Dependency on GPU	884
<i>Hongyuan Liu, King Tin Lam, Huanxin Lin, Cho-Li Wang, and Junchao Ma</i>	
Machine Learning Approach for the Predicting Performance of SpMV on GPU	894
<i>Akrem Benatia, Weixing Ji, Yizhuo Wang, and Feng Shi</i>	
Parallelizing Back Propagation Neural Network on Speculative Multicores	902
<i>Yaobin Wang, Hong An, Zhiqin Liu, Tao Liu, and Dongmei Zhao</i>	
Selectively GPU Cache Bypassing for Un-Coalesced Loads	908
<i>Chen Zhao, Fei Wang, Zhen Lin, Huiyang Zhou, and Nanning Zheng</i>	
Speeding Up Virtualized Transaction Logging with vTrans	916
<i>Yong Tang, Lingxiao Chen, Jiwei Liu, and Dingding Li</i>	
Study of Neocortex Simulations with GENESIS on High Performance Computing Resources	924
<i>Sean McDaniel, David L. Boothe, Joshua C. Crone, Song J. Park, Dale R. Shires, Alfred B. Yu, and Michela Taufer</i>	
System-Level Scalable Checkpoint-Restart for Petascale Computing	932
<i>Jiajun Cao, Kapil Arya, Rohan Garg, Shawn Matott, Dhabaleswar K. Panda, Hari Subramoni, Jérôme Vienne, and Gene Cooperman</i>	
Using Supercomputer to Speed up Neural Network Training	942
<i>Yue Yu, Jinrong Jiang, and Xuebin Chi</i>	

Session 8: Parallel/Distributed Algorithms and Applications

A C-SVM Based Anomaly Detection Method for Multi-Dimensional Sequence over Data Stream	948
<i>Han Bao and Yijie Wang</i>	
A Key Frame Selection Algorithm Based on Sliding Window and Image Features	956
<i>Jigang Cao, Linchen Yu, Mianlong Chen, and Xinyue Cui</i>	
Accelerating Deep Learning with Shrinkage and Recall	963
<i>Shuai Zheng, Abhinav Vishnu, and Chris Ding</i>	
An Energy-Efficient Implementation of LU Factorization on Heterogeneous Systems	971
<i>Canqun Yang, Cheng Chen, Tao Tang, Xuhao Chen, Jianbin Fang, and Jingling Xue</i>	
Analysis, Modeling, and Simulation of Hadoop YARN MapReduce	980
<i>Thomas C. Bressoud and Qiuyi (Jessica) Tang</i>	
Application-Level Determinism in Distributed Systems	989
<i>Christopher Boelmann, Lorenz Schwittmann, Marian Waltereit, Matthäus Wander, and Torben Weis</i>	
Asynchronous Progress Design for a MPI-Based PGAS One-Sided Communication System	999
<i>Huan Zhou and José Gracia</i>	
CHIME: A Checkpoint-Based Approach to Improving the Performance of Shared Clusters	1007
<i>Yiyang Shao, Xiaomin Zhu, Weidong Bao, Wen Zhou, and Wenhua Xiao</i>	
CPU Frequency Tuning to Improve Energy Efficiency of MapReduce Systems	1015
<i>Nidhi Tiwari, Umesh Bellur, Santonu Sarkar, and Maria Indrawan</i>	
Energy Proportionality in Heterogeneous Data Center Supporting Applications with Variable Load	1023
<i>Violaine Villebonnet, Georges Da Costa, Laurent Lefevre, Jean-Marc Pierson, and Patricia Stolf</i>	
Evaluation of Flash-Based Out-of-Core Stencil Computation Algorithms for SSD-Equipped Clusters	1031
<i>Hiroko Midorikawa and Hideyuki Tan</i>	
ICE: A General and Validated Energy Complexity Model for Multithreaded Algorithms	1041
<i>Vi Ngoc-Nha Tran and Phuong Hoai Ha</i>	

Parallel Gene Upstream Comparison via Multi-Level Hash Tables on GPU	1049
<i>Andrew Todd, Huan Truong, Justin Deters, John Long, Gavin Conant, and Michela Becchi</i>	
PFrauDector: A Parallelized Graph Mining Approach for Efficient Fraudulent Phone Call Detection	1059
<i>Josh Jia-Ching Ying, Ji Zhang, Che-Wei Huang, Kuan-Ta Chen, and Vincent S. Tseng</i>	
PIE: A Pipeline Energy-Efficient Accelerator for Inference Process in Deep Neural Networks	1067
<i>Yangyang Zhao, Qi Yu, Xuda Zhou, Xuehai Zhou, Xi Li, and Chao Wang</i>	
Reverse Engineering of Dynamic Parallel Program Behavior from Execution Traces	1075
<i>Shin-Chieh Tsai, Chiu-Ping Chang, and Chung-Ta King</i>	
Scalable Single-Source SimRank Computation for Large Graphs	1083
<i>Xingkun Gao, Nianyuan Bao, Jie Liu, Jie Tang, and Gangshan Wu</i>	
suCAQR: A Simplified Communication-Avoiding QR Factorization Solver Using the TBLAS Framework	1092
<i>Weijian Zheng, Fengguang Song, Lan Lin, and Zizhong Chen</i>	
Tapas: An Implicitly Parallel Programming Framework for Hierarchical N-Body Algorithms	1100
<i>Keisuke Fukuda, Motohiko Matsuda, Naoya Maruyama, Rio Yokota, Kenjiro Taura, and Satoshi Matsuoka</i>	
Timed Dataflow: Reducing Communication Overhead for Distributed Machine Learning Systems	1110
<i>Peng Sun, Yonggang Wen, Ta Nguyen Binh Duong, and Shengen Yan</i>	
Towards Scalable Subgraph Pattern Matching over Big Graphs on MapReduce	1118
<i>Bo Suo, Zhanhuai Li, Qun Chen, and Wei Pan</i>	
Towards Seamless Resynchronization for Active-Active Database Clustering	1127
<i>Xiang Gao and Tzi-cker Chiueh</i>	
Understanding Software Platforms for In-Memory Scientific Data Analysis: A Case Study of the Spark System	1135
<i>Xuechen Zhang, Ujjwal Khanal, Xinghui Zhao, and Stephen Ficklin</i>	

ICPADSW 2016

A Categorical Approach in Handling Event-Ordering in Distributed Systems	1145
<i>Dimitrios Sisiaridis, Veronika Kuchta, and Olivier Markowitch</i>	

A Molecular Computation Model to Compute Inversion over Finite Field GF(2 ⁿ)	1151
<i>Yongnan Li, Limin Xiao, and Huawei Tian</i>	
A Secure and Reliable Hybrid Model for Cloud-of-Clouds Storage Systems	1157
<i>Dagang Li and Ying Zhou</i>	
Coordinated Broadcast-Based Request-Reply and Group Management for Tightly-Coupled Wireless Systems	1163
<i>Manos Koutsoubelias and Spyros Lalis</i>	
Dual LWE-Based Fully Homomorphic Encryption with Errorless Key Switching	1169
<i>Zengpeng Li, Chunguang Ma, Gang Du, and Ouyang Weiping</i>	
IRMD: Malware Variant Detection Using Opcode Image Recognition	1175
<i>Jixin Zhang, Zheng Qin, Hui Yin, Lu Ou, and Yupeng Hu</i>	
MrBayes 3.2.6 on Tianhe-1A: A High Performance and Distributed Implementation of Phylogenetic Analysis	1181
<i>Cheng Ling, Arong Luo, and Jingyang Gao</i>	
Resources Renting with Reserved and On-Demand Instances for Cloud Workflow Applications	1187
<i>Long Chen and Xiaoping Li</i>	
RESS: A Reliable Energy-Efficient Storage System	1193
<i>Shu Yin, Zhaoyu Xiao, Kenli Li, Jianzhong Huang, Xiaojun Ruan, Xiaomin Zhu, and Xiao Qin</i>	

AIR 2016

Accelerating Learning to Rank via SVM with OpenCL and OpenMP on Heterogeneous Platforms	1199
<i>Huming Zhu, Zheng Luo, Yanfei Wu, Pei Li, Peng Zhang, Shuiping Gou, and L.C. Jiao</i>	
Context-Aware Video Object Proposals	1203
<i>Wenjing Geng and Gangshan Wu</i>	
Convolutional Neural Network Simplification Based on Feature Maps Selection	1207
<i>Ting Rui, Junhua Zou, You Zhou, Jianchao Fei, and Chengsong Yang</i>	
Image Enhancement Based on Bi-Histogram Equalization with Non-Parametric Modified Technology	1211
<i>Zhijun Yao, Quan Zhou, Zhongyuan Lai, Zhiming Ren, and Liming Liu</i>	
Multiple Cartesian K-Medoids for a Fine Quantization	1216
<i>Lihua Tian, Shanmin Pang, and Chen Li</i>	
Online Object Tracking Based on Convex Hull Representation	1221
<i>Chunjuan Bo and Dong Wang</i>	

Research on Semi-Supervised Learning for Hyperspectral Remote Sensing Imaging Classification Base on Confidence Entropy	1225
<i>Chunyang Wang, Zhifang Xu, Shuangting Wang, Hebing Zhang, and Zhichao Chen</i>	
Sparse Autoencoder Based Deep Neural Network for Voxelwise Detection of Cerebral Microbleed	1229
<i>Yu-Dong Zhang, Xiao-Xia Hou, Yi-Ding Lv, Hong Chen, Yin Zhang, and Shui-Hua Wang</i>	
Steganalysis via Deep Residual Network	1233
<i>Songtao Wu, Sheng-Hua Zhong, and Yan Liu</i>	
Super Resolving of the Depth Map for 3D Reconstruction of Underwater Terrain Using Kinect	1237
<i>Yu Nakagawa, Keita Kihara, Ryunosuke Tadoh, Seiichi Serikawa, Huimin Lu, Yudong Zhang, and Yujie Li</i>	

SOWN 2016

A Secure Data Deduplication Scheme Based on Differential Privacy	1241
<i>Jun Ren, Zhiqiang Yao, Jinbo Xiong, Yuanyuan Zhang, and Ayong Ye</i>	
Changing Pseudonym Based on Virtual Users in Vehicular Networks	1247
<i>Liu Hong-Yue, Lyu Xiao-Yang, Yang Wei-Dong, and Li Xing-Hua</i>	
Design and Implementation of Mixed Extended-Precision Package in MATLAB	1251
<i>Peibing Du, Hao Jiang, and Lizhi Cheng</i>	

Author Index