

# **2009 29th IEEE International Conference on Distributed Computing Systems**

**(ICDCS 2009)**

**Montreal, Quebec, Canada  
22 – 26 June 2009**



**IEEE Catalog Number: CFP09040-PRT**  
**ISBN: 978-1-4244-4481-6**

**Copyright © 2009 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:	CFP09040-PRT
ISBN 13:	978-1-4244-4481-6
ISSN:	1063-6927

**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](mailto:curran@proceedings.com)

# 2009 29th IEEE International Conference on Distributed Computing Systems

## ICDCS 2009

# Table of Contents

### Keynote Address

Rethinking Multicast for Massive-Scale Platforms .....	1
<i>Ken Birman</i>	

### Session 1A: Distributed Systems Infrastructure (I)

A Reinforcement Learning Approach to Online Web Systems	
Auto-configuration .....	2
<i>Xiangping Bu, Jia Rao, and Cheng-Zhong Xu</i>	
Graduated QoS by Decomposing Bursts: Don't Let the Tail Wag Your Server .....	12
<i>Lanyue Lu, Peter Varman, and Kshitij Doshi</i>	
Reducing Disk I/O Performance Sensitivity for Large Numbers of Sequential Streams .....	22
<i>George Panagiotakis, Michail D. Flouris, and Angelos Bilas</i>	

### Session 1B: Privacy

Fragmentation Design for Efficient Query Execution over Sensitive Distributed Databases .....	32
<i>Valentina Ciriani, Sabrina De Capitani di Vimercati, Sara Foresti, Sushil Jajodia, Stefano Paraboschi, and Pierangela Samarati</i>	
Lightweight Secure Search Protocols for Low-cost RFID Systems .....	40
<i>Lars Kulseng, Zhen Yu, Yawen Wei, and Yong Guan</i>	

CAP: A Context-Aware Privacy Protection System for Location-Based Services .....	49
<i>Aniket Pingley, Wei Yu, Nan Zhang, Xinwen Fu, and Wei Zhao</i>	

### **Session 1C: Routing**

The Impact of Communication Models on Routing-Algorithm Convergence .....	58
<i>Aaron D. Jaggard, Vijay Ramachandran, and Rebecca N. Wright</i>	
Selective Protection: A Cost-Efficient Backup Scheme for Link State Routing .....	68
<i>Meijia Hou, Dan Wang, Mingwei Xu, and Jiahai Yang</i>	
Centaur: A Hybrid Approach for Reliable Policy-Based Routing .....	76
<i>Xin Zhang, Adrian Perrig, and Hui Zhang</i>	

### **Session 2A: Wireless Networks**

Joint Optimization of Spectrum Handoff Scheduling and Routing in Multi-hop Multi-radio Cognitive Networks .....	85
<i>Wei Feng, Jiannong Cao, Chisheng Zhang, and Chuda Liu</i>	
Simulation Framework and Performance Analysis of Multimedia Broadcasting Service over Wireless Networks .....	93
<i>Won J. Jeon, Kyungtae Kang, Roy H. Campbell, and Klara Nahrstedt</i>	
Transactional Mobility in Distributed Content-Based Publish/Subscribe Systems .....	101
<i>Songlin Hu, Vinod Muthusamy, Guoli Li, and Hans-Arno Jacobsen</i>	
Multicast Throughput of Hybrid Wireless Networks Under Gaussian Channel Model .....	111
<i>Cheng Wang, Shaojie Tang, Xiang-Yang Li, Changjun Jiang, and Yunhao Liu</i>	

### **Session 2B: Internet Security**

Distributed Key Generation for the Internet .....	119
<i>Aniket Kate and Ian Goldberg</i>	
Clock-like Flow Replacement Schemes for Resilient Flow Monitoring .....	129
<i>Gunwoo Nam, Pushkar Patankar, Seung-Hwan Lim, Bikash Sharma, George Kesidis, and Chita R. Das</i>	
The Taming of the Shrew: Mitigating Low-Rate TCP-Targeted Attack .....	137
<i>Chia-Wei Chang, Seungjoon Lee, Bill Lin, and Jia Wang</i>	
The Design and Evaluation of Accountable Grid Computing System .....	145
<i>Wonjun Lee, Anna Squicciarini, and Elisa Bertino</i>	

### **Session 2C: Performance and System Architecture**

The Case for Spam-Aware High Performance Mail Server Architecture .....	155
<i>Abhinav Pathak, Syed Ali Raza Jafri, and Y. Charlie Hu</i>	

An Approach to Sharing Legacy TV/Arcade Games for Real-Time Collaboration .....	165
<i>Sili Zhao, Du Li, Hansu Gu, Bin Shao, and Ning Gu</i>	
Minimizing Latency in Fault-Tolerant Distributed Stream Processing Systems .....	173
<i>Andrey Brito, Christof Fetzer, and Pascal Felber</i>	
Automatic Performance Tuning for the Virtualized Cluster System .....	183
<i>Chuliang Weng, Minglu Li, Zhigang Wang, and Xinda Lu</i>	
<b>Session 3A: Distributed Data Management (I)</b>	
m-LIGHT: Indexing Multi-Dimensional Data over DHTs .....	191
<i>Yuzhe Tang, Jianliang Xu, Shuigeng Zhou, and Wang-chien Lee</i>	
Two-Tier Air Indexing for On-Demand XML Data Broadcast .....	199
<i>Weiwei Sun, Ping Yu, Yongrui Qing, Zhuoyao Zhang, and Baihua Zheng</i>	
Distributed Processing of Spatial Alarms: A Safe Region-Based Approach .....	207
<i>Bhuvan Bamba, Ling Liu, Arun Iyengar, and Philip S. Yu</i>	
<b>Session 3B: Vehicular Ad-hoc Networks</b>	
Maintaining Probabilistic Consistency for Frequently Offline Devices in Mobile Ad Hoc Networks .....	215
<i>Wenzhong Li, Edward Chan, Daoxu Chen, and Sanglu Lu</i>	
Roadcast: A Popularity Aware Content Sharing Scheme in VANETs .....	223
<i>Yang Zhang, Jing Zhao, and Guohong Cao</i>	
TBD: Trajectory-Based Data Forwarding for Light-Traffic Vehicular Networks .....	231
<i>Jaehoon Jeong, Shuo Guo, Yu Gu, Tian He, and David Du</i>	
<b>Session 3C: Distributed Systems Infrastructure (II)</b>	
PADD: Power Aware Domain Distribution .....	239
<i>Min Yeol Lim, Freeman Rawson, Tyler Bletsch, and Vincent W. Freeh</i>	
REMO: Resource-Aware Application State Monitoring for Large-Scale Distributed Systems .....	248
<i>Shicong Meng, Srinivas R. Kashyap, Chitra Venkatramani, and Ling Liu</i>	
On the Utility of Inference Mechanisms .....	256
<i>Ethan Blanton, Sonia Fahmy, and Greg N. Frederickson</i>	
<b>Keynote Address</b>	
Networking the Cloud .....	264
<i>Albert Greenberg</i>	

## Session 4A: Sensor Networks (I)

On the Connected k-Coverage Problem in Heterogeneous Sensor Nets: The Curse of Randomness and Heterogeneity .....	265
<i>Habib M. Ammari and John Giudici</i>	
Sampling Based (epsilon, delta)-Approximate Aggregation Algorithm in Sensor Networks .....	273
<i>Siyao Cheng and Jianzhong Li</i>	
Available Bandwidth in Multirate and Multihop Wireless Sensor Networks .....	281
<i>Feng Chen, Hongqiang Zhai, and Yuguang Fang</i>	
FLASH: Fine-Grained Localization in Wireless Sensor Networks Using Acoustic Sound Transmissions and High Precision Clock Synchronization .....	289
<i>Evangelos Mangas and Angelos Bilas</i>	

## Session 4B: Distributed Data Management (II)

Q-Tree: A Multi-Attribute Based Range Query Solution for Tele-immersive Framework .....	299
<i>Md Ahsan Arefin, Md Yusuf Sarwar Uddin, Indranil Gupta, and Klara Nahrstedt</i>	
Optimizing File Retrieval in Delay-Tolerant Content Distribution Community .....	308
<i>Ying Huang, Yan Gao, Klara Nahrstedt, and Wenbo He</i>	
On Optimal Concurrency Control for Optimistic Replication .....	317
<i>Weihan Wang and Cristiana Amza</i>	
File Versioning for Block-Level Continuous Data Protection .....	327
<i>Maohua Lu and Tzicker Chiueh</i>	

## Session 4C: Distributed Algorithms and Protocols (I)

Model Checking Transactional Memory with Spin .....	335
<i>John O'Leary, Bratin Saha, and Mark R. Tuttle</i>	
A Distributed Termination Detection Algorithm for Dynamic Asynchronous Systems .....	343
<i>Paul Johnson and Neeraj Mittal</i>	
Towards Optimal Resource Utilization in Heterogeneous P2P Streaming .....	352
<i>Dongyu Liu, Fei Li, and Songqing Chen</i>	

## Session 5A: Dependable Distributed Systems

NAT-resilient Gossip Peer Sampling .....	360
<i>Anne-Marie Kermarrec, Alessio Pace, Vivien Quéma, and Valerio Schiavoni</i>	
Fault-Tolerant Consensus in Unknown and Anonymous Networks .....	368
<i>Carole Delporte-Gallet, Hugues Fauconnier, and Andreas Tielmann</i>	
A Practical Study of Regenerating Codes for Peer-to-Peer Backup Systems .....	376
<i>Alessandro Duminuco and Ernst Biersack</i>	

CARP: Handling Silent Data Errors and Site Failures in an Integrated Program and Storage Replication Mechanism .....	385
<i>Lanyue Lu, Prasenjit Sarkar, Dinesh Subhraveti, Soumitra Sarkar, Mark Seaman, Reshu Jain, and Ahmed Bashir</i>	

### **Session 5B: P2P Systems**

A Commutative Replicated Data Type for Cooperative Editing .....	395
<i>Nuno Preguica, Joan Manuel Marquès, Marc Shapiro, and Mihai Letia</i>	
Logoot: A Scalable Optimistic Replication Algorithm for Collaborative Editing on P2P Networks .....	404
<i>Stéphane Weiss, Pascal Urso, and Pascal Molli</i>	
ISP—Friend or Foe? Making P2P Live Streaming ISP-Aware .....	413
<i>Fabio Picconi and Laurent Massoulié</i>	
A Case Study of Traffic Locality in Internet P2P Live Streaming Systems .....	423
<i>Yao Liu, Lei Guo, Fei Li, and Songqing Chen</i>	

### **Session 5C: Sensor Networks (II)**

EUL: An Efficient and Universal Localization Method for Wireless Sensor Network .....	433
<i>Wei Xi, Jizhong Zhao, Xue Liu, Xiang-Yang Li, and Yong Qi</i>	
CLIQUE: Role-Free Clustering with Q-Learning for Wireless Sensor Networks .....	441
<i>Anna Förster and Amy L. Murphy</i>	
QVS: Quality-Aware Voice Streaming for Wireless Sensor Networks .....	450
<i>Liqun Li, Guoliang Xin, Limin Sun, and Yan Liu</i>	
TDMA-ASAP: Sensor Network TDMA Scheduling with Adaptive Slot-Stealing and Parallelism .....	458
<i>Sameh Gabriel, Daniel Mossé, and Robert Cleric</i>	

### **Session 6A: Distributed Algorithms and Protocols (II)**

A Note on Distributed Stable Matching .....	466
<i>Alex Kipnis and Boaz Patt-Shamir</i>	
Communication Efficiency in Self-Stabilizing Silent Protocols .....	474
<i>Stéphane Devismes, Toshimitsu Masuzawa, and Sébastien Tixeuil</i>	
On the Impossibility of Maximal Scheduling for Strong Fairness with Interleaving .....	482
<i>Matthew Lang and Paolo A.G. Sivilotti</i>	

## Session 6B: Network Coding and Protocols

Pushing the Envelope: Extreme Network Coding on the GPU .....	490
<i>Hassan Shojania and Baochun Li</i>	
Stochastic Multicast with Network Coding .....	500
<i>Ajay Gopinathan and Zongpeng Li</i>	
High-Speed Flow Nature Identification .....	510
<i>Amir R. Khakpour and Alex X. Liu</i>	

## Session 6C: Distributed System Resource Efficiency

Locality-Preserving Clustering and Discovery of Wide-Area Grid Resources .....	518
<i>Haiying Shen and Kai Hwang</i>	
MOPS: Providing Content-Based Service in Disruption-Tolerant Networks .....	526
<i>Feng Li and Jie Wu</i>	
FlashLite: A User-Level Library to Enhance Durability of SSD for P2P File Sharing .....	534
<i>Hyojun Kim and Umakishore Ramachandran</i>	

## Keynote Address

Better Architectures and New Security Applications for Network Monitoring .....	542
<i>Michael Reiter</i>	

## Session 7A: Programmer-Visible Optimization in Distributed Systems

Explicit Batching for Distributed Objects .....	543
<i>Eli Tilevich, William R. Cook, and Yang Jiao</i>	
Minimizing the Hidden Cost of RDMA .....	553
<i>Philip Werner Frey and Gustavo Alonso</i>	
Autonomous Resource Selection for Decentralized Utility Computing .....	561
<i>Paolo Costa, Jeff Napper, Guillaume Pierre, and Maarten van Steen</i>	
Information Value-Driven Near Real-Time Decision Support Systems .....	571
<i>Ying Yan, Wen-Syan Li, and Jian Xu</i>	

## Session 7B: Wireless Security and Intrusion Detection

Protecting Neighbor Discovery Against Node Compromises in Sensor Networks .....	579
<i>Donggang Liu</i>	
The Digital Marauder's Map: A New Threat to Location Privacy .....	589
<i>Xinwen Fu, Nan Zhang, Aniket Pingley, Wei Yu, Jie Wang, and Wei Zhao</i>	
Characterization and Solution to a Stateful IDS Evasion .....	597
<i>Issam Aib, Tung Tran, and Raouf Boutaba</i>	
Down the Block and Around the Corner—The Impact of Radio Propagation on Inter-vehicle Wireless Communication .....	605
<i>John S. Otto, Fabián E. Bustamante, and Randall A. Berry</i>	



## **Session 7C: Reliable Distributed Systems**

Deterministic Replay for Transparent Recovery in Component-Oriented Middleware .....	615
<i>Rob Strom, Chitra Dorai, Thomas Huining Feng, and Wei Zheng</i>	
Modeling Probabilistic Measurement Correlations for Problem Determination in Large-Scale Distributed Systems .....	623
<i>Jing Gao, Guofei Jiang, Haifeng Chen, and Jiawei Han</i>	
Collaboration-Oriented Data Recovery for Mobile Disk Arrays .....	631
<i>Tao Xie and Abhinav Sharma</i>	
Implementing a Register in a Dynamic Distributed System .....	639
<i>Roberto Baldoni, Silvia Bonomi, Anne-Marie Kermarrec, and Michel Raynal</i>	

## **Author Index**