

**2012 12th IEEE/ACM  
International Symposium on  
Cluster, Cloud and Grid  
Computing**

**(CCGrid 2012)**

**Ottawa, Ontario, Canada  
13 – 16 May 2012**



**IEEE Catalog Number: CFP12276-PRT  
ISBN: 978-1-4673-1395-7**

# 2012 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing

## CCGrid 2012

### Table of Contents

<b>Message from the General Chair</b> .....	xviii
<b>Message from the Program Co-Chairs</b> .....	xx
<b>Chairs and Committees</b> .....	xxii
<b>Program Committee Members</b> .....	xxiv
<b>Steering Committee</b> .....	xxvii
<b>External Reviewers</b> .....	xxviii
<b>Keynote Addresses</b> .....	xxxi

---

#### **Session 1A: Programming Models and File Systems**

Design and Implementation of Portable and Efficient Non-blocking Collective Communication .....	1
<i>Akihiro Nomura, Yutaka Ishikawa, Naoya Maruyama, and Satoshi Matsuoka</i>	
Productive Parallel Linear Algebra Programming with Unstructured Topology Adaption .....	9
<i>Peter Gottschling and Torsten Hoefler</i>	
User Transparent Data and Task Parallel Multimedia Computing with Pyxis-DT .....	17
<i>Timo van Kessel, Niels Drost, Jason Maassen, Henri E. Bal, and Frank J. Seinstra</i>	
CEFLS: A Cost-Effective File Lookup Service in a Distributed Metadata File System .....	25
<i>Xiuqiao Li, Bin Dong, Limin Xiao, Li Ruan, and Dongmei Liu</i>	
Fine-Grained Access Control in the Chirp Distributed File System .....	33
<i>Patrick Donnelly and Douglas Thain</i>	

## Session 1B: Map Reduce and Workflows

Towards Trusted Services: Result Verification Schemes for MapReduce .....	41
<i>Chu Huang, Sencun Zhu, and Dinghao Wu</i>	
MARLA: MapReduce for Heterogeneous Clusters .....	49
<i>Zacharia Fadika, Elif Dede, Jessica Hartog, and Madhusudhan Govindaraju</i>	
A Map-Reduce Based Framework for Heterogeneous Processing Element Cluster Environments .....	57
<i>Yu Shyang Tan, Bu-Sung Lee, Bingsheng He, and Roy H. Campbell</i>	
Workflow Scheduling to Minimize Data Movement Using Multi-constraint Graph Partitioning .....	65
<i>Masahiro Tanaka and Osamu Tatebe</i>	
Decentralized Orchestration of Data-centric Workflows Using the Object Modeling System .....	73
<i>Bahman Javadi, Martin Tomko, and Richard O. Sinnott</i>	

## Session 1C: QoS and Architecture

Efficient Disk I/O Scheduling with QoS Guarantee for Xen-based Hosting Platforms .....	81
<i>Xiao Ling, Hai Jin, Shadi Ibrahim, Wenzhi Cao, and Song Wu</i>	
MORPHOSYS: Efficient Colocation of QoS-Constrained Workloads in the Cloud .....	90
<i>Vatche Ishakian and Azer Bestavros</i>	
Reward Scheduling for QoS in Cloud Applications .....	98
<i>Ahmed Elnably, Kai Du, and Peter Varman</i>	
Evaluating Dynamics and Bottlenecks of Memory Collaboration in Cluster Systems .....	107
<i>Ahmad Samih, Ren Wang, Christian Maciocco, Tsung-Yuan Charlie Tai, Ronghui Duan, Jiengang Duan, and Yan Solihin</i>	
Energy- and Cost-Efficiency Analysis of ARM-Based Clusters .....	115
<i>Zhonghong Ou, Bo Pang, Yang Deng, Jukka K. Nurminen, Antti Ylä-Jääski, and Pan Hui</i>	

## Session 2A: GPU

Transparent Accelerator Migration in a Virtualized GPU Environment .....	124
<i>Shucaï Xiao, Pavan Balaji, James Dinan, Qian Zhu, Rajeev Thakur, Susan Coghlan, Heshan Lin, Gaojin Wen, Jue Hong, and Wu-chun Feng</i>	
GPU Performance Enhancement via Communication Cost Reduction: Case Studies of Radix Sort and WSN Relay Node Placement Problem .....	132
<i>Che-Rung Lee, Shih-Hsiang Lo, Nan-Hsi Chen, Yeh-Ching Chung, and I-Hsin Chung</i>	
Scheduling Concurrent Applications on a Cluster of CPU-GPU Nodes .....	140
<i>Vignesh T. Ravi, Michela Becchi, Wei Jiang, Gagan Agrawal, and Srimat Chakradhar</i>	
A Highly Parallel Multi-class Pattern Classification on GPU .....	148
<i>Mahdi Nabyouni and Delasa Aghamirzaie</i>	

## Session 2B: Cloud Services I

Client Classification Policies for SLA Enforcement in Shared Cloud Datacenters .....	156
<i>Mario Macías and Jordi Guitart</i>	
Policy-Based Automation of SLA Establishment for Cloud Computing Services .....	164
<i>Mohan Baruwal Chhetri, Quoc Bao Vo, and Ryszard Kowalczyk</i>	
SLA-based Optimization of Power and Migration Cost in Cloud Computing .....	172
<i>Hadi Goudarzi, Mohammad Ghasemazar, and Massoud Pedram</i>	
Service Level Agreement for Distributed Mutual Exclusion in Cloud Computing .....	180
<i>Jonathan Lejeune, Luciana Arantes, Julien Sopena, and Pierre Sens</i>	

## Session 2C: I/O and File Systems

On Urgency of I/O Operations .....	188
<i>Mahmut Kandemir, Taylan Yemliha, Ramya Prabhakar, and Myoungsoo Jung</i>	
Boosting Application-Specific Parallel I/O Optimization Using IOSIG .....	196
<i>Yanlong Yin, Surendra Byna, Huaiming Song, Xian-He Sun, and Rajeev Thakur</i>	
SERA-IO: Integrating Energy Consciousness into Parallel I/O Middleware .....	204
<i>Rong Ge, Xizhou Feng, and Xian-He Sun</i>	

Supporting User-Defined Subsetting and Aggregation over Parallel NetCDF Datasets .....	212
<i>Yu Su and Gagan Agrawal</i>	

### **Session 3A: Communication and Networks**

Scalable Multi-purpose Network Representation for Large Scale Distributed System Simulation .....	220
<i>Laurent Bobelin, Arnaud Legrand, David A. González Márquez, Pierre Navarro, Martin Quinson, Frédéric Suter, and Christophe Thiéry</i>	
Topology Agnostic Dynamic Quick Reconfiguration for Large-Scale Interconnection Networks .....	228
<i>Frank Olaf Sem-Jacobsen and Olav Lysne</i>	
Scalable Memcached Design for InfiniBand Clusters Using Hybrid Transports .....	236
<i>Jithin Jose, Hari Subramoni, Krishna Kandalla, Md. Wasi-ur-Rahman, Hao Wang, Sundeep Narravula, and Dhabaleswar K. Panda</i>	
Lowering Inter-datacenter Bandwidth Costs via Bulk Data Scheduling .....	244
<i>Thyaga Nandagopal and Krishna P.N. Puttaswamy</i>	
A Scalable Parallel Debugging Library with Pluggable Communication Protocols .....	252
<i>Chao Jin, David Abramson, Minh Ngoc Dinh, Andrew Gontarek, Robert Moench, and Luiz DeRose</i>	

### **Session 3B: Faults, Failures and Reliability**

Task Scheduling Algorithm for Multicore Processor System for Minimizing Recovery Time in Case of Single Node Fault .....	260
<i>Shohei Gotoda, Minoru Ito, and Naoki Shibata</i>	
Speculative Memory State Transfer for Active-Active Fault Tolerance .....	268
<i>Maohua Lu and Tzi-cker Chiueh</i>	
Checkpointing Orchestration: Toward a Scalable HPC Fault-Tolerant Environment .....	276
<i>Hui Jin, Tao Ke, Yong Chen, and Xian-He Sun</i>	
An Autonomous Reliability-Aware Negotiation Strategy for Cloud Computing Environments .....	284
<i>Amir Vahid Dastjerdi and Rajkumar Buyya</i>	
Separating Performance Anomalies from Workload-Explained Failures in Streaming Servers .....	292
<i>Carlos Augusto Cunha and Luis Moura e Silva</i>	

### **Session 3C: Workflows**

A Multi-objective Approach for Workflow Scheduling in Heterogeneous Environments .....	300
<i>Hamid Mohammadi Fard, Radu Prodan, Juan Jose Durillo Barrionuevo, and Thomas Fahringer</i>	
Automating Data-Throttling Analysis for Data-Intensive Workflows .....	310
<i>Ricardo J. Rodríguez, Rafael Tolosana-Calasanz, and Omer F. Rana</i>	
Self-Healing of Operational Workflow Incidents on Distributed Computing Infrastructures .....	318
<i>Rafael Ferreira da Silva, Tristan Glatard, and Frédéric Desprez</i>	
A Workflow-Aware Storage System: An Opportunity Study .....	326
<i>Emalayan Vairavanathan, Samer Al-Kiswany, Lauro Beltrão Costa, Zhao Zhang, Daniel S. Katz, Michael Wilde, and Matei Ripeanu</i>	
Using Model Checking to Analyze the System Behavior of the LHC Production Grid .....	335
<i>Daniela Remenska, Tim A.C. Willemse, Kees Verstoep, Wan Fokkink, Jeff Templon, and Henri Bal</i>	

### **Session 3D: Scheduling and Monitoring**

WSCOM: Online Task Scheduling with Data Transfers .....	344
<i>Jean-Noël Quintin and Frédéric Wagner</i>	
Improving Grid Resource Usage: Metrics for Measuring Fragmentation .....	352
<i>Luis Tomás, Blanca Caminero, and Carmen Carrión</i>	
Malleable Model Coupling with Prediction .....	360
<i>Daihee Kim, J. Walter Larson, and Kenneth Chiu</i>	
Load Balancing Query Processing in Metric-Space Similarity Search .....	368
<i>Veronica Gil-Costa and Mauricio Marin</i>	
Distributed Monitoring with Collaborative Prediction .....	376
<i>Dawei Feng, Cécile Germain-Renaud, and Tristan Glatard</i>	

### **Session 4A: Programming Models**

Delta Send-Recv for Dynamic Pipelining in MPI Programs .....	384
<i>Bin Bao, Chen Ding, Yaoqing Gao, and Roch Archambault</i>	
Global Futures: A Multithreaded Execution Model for Global Arrays-based Applications .....	393
<i>Daniel Chavarría-Miranda, Sriram Krishnamoorthy, and Abhinav Vishnu</i>	

Productivity and Performance of Global-View Programming with XcalableMP PGAS Language .....	402
<i>Masahiro Nakao, Jinpil Lee, Taisuke Boku, and Mitsuhsisa Sato</i>	
Distributed S-Net: Cluster and Grid Computing without the Hassle .....	410
<i>Clemens Grellck, Jukka Julku, and Frank Penczek</i>	
<b>Session 4B: Map Reduce</b>	
Investigation of Data Locality in MapReduce .....	419
<i>Zhenhua Guo, Geoffrey Fox, and Mo Zhou</i>	
TomusBlobs: Towards Communication-Efficient Storage for MapReduce Applications in Azure .....	427
<i>Radu Tudoran, Alexandru Costan, Gabriel Antoniu, and Hakan Soncu</i>	
Maestro: Replica-Aware Map Scheduling for MapReduce .....	435
<i>Shadi Ibrahim, Hai Jin, Lu Lu, Bingsheng He, Gabriel Antoniu, and Song Wu</i>	
SciMATE: A Novel MapReduce-Like Framework for Multiple Scientific Data Formats .....	443
<i>Yi Wang, Wei Jiang, and Gagan Agrawal</i>	
<b>Session 4C: Cloud Computing I</b>	
Pricing Cloud Compute Commodities: A Novel Financial Economic Model .....	451
<i>Bhanu Sharma, Ruppa K. Thulasiram, Parimala Thulasiraman, Saurabh K. Garg, and Rajkumar Buyya</i>	
A Time-Series Pattern Based Noise Generation Strategy for Privacy Protection in Cloud Computing .....	458
<i>Gaofeng Zhang, Yun Yang, Xiao Liu, and Jinjun Chen</i>	
COCA: Computation Offload to Clouds Using AOP .....	466
<i>Hsing-Yu Chen, Yue-Hsun Lin, and Chen-Mou Cheng</i>	
ParaLite: Supporting Collective Queries in Database System to Parallelize User-Defined Executable .....	474
<i>Ting Chen and Kenjiro Taura</i>	
<b>Session 5A: Virtualization</b>	
Snooze: A Scalable and Autonomic Virtual Machine Management Framework for Private Clouds .....	482
<i>Eugen Feller, Louis Rilling, and Christine Morin</i>	

RO-BURST: A Robust Virtualization Cost Model for Workload Consolidation over Clouds .....	490
<i>Jianzong Wang, Rui Hua, Yifeng Zhu, Jiguang Wan, Changsheng Xie, and Yanjun Chen</i>	
A Stable Network-Aware VM Placement for Cloud Systems .....	498
<i>Ofer Biran, Antonio Corradi, Mario Fanelli, Luca Foschini, Alexander Nus, Danny Raz, and Ezra Silvera</i>	
Surreptitious Deployment and Execution of Kernel Agents in Windows Guests .....	507
<i>Tzi-cker Chiueh, Matthew Conover, and Bruce Montague</i>	
<b>Session 5B: Cloud Services II</b>	
Cloud Service Negotiation: Concession vs. Tradeoff Approaches .....	515
<i>Xianrong Zheng, Patrick Martin, and Kathryn Brohman</i>	
Interactive Use of Cloud Services: Amazon SQS and S3 .....	523
<i>Hobin Yoon, Ada Gavrilovska, Karsten Schwan, and Jim Donahue</i>	
Dynamic Replication in Service-Oriented Systems .....	531
<i>Mathias Björkqvist, Lydia Y. Chen, and Walter Binder</i>	
Selling T-shirts and Time Shares in the Cloud .....	539
<i>Daniel Gmach, Jerry Rolia, and Ludmila Cherkasova</i>	
<b>Session 5C: Data on the Cloud</b>	
Scalable Join Queries in Cloud Data Stores .....	547
<i>Zhou Wei, Guillaume Pierre, and Chi-Hung Chi</i>	
Privacy Preserving Access Control with Authentication for Securing Data in Clouds .....	556
<i>Sushmita Ruj, Milos Stojmenovic, and Amiya Nayak</i>	
A Cost-Effective Mechanism for Cloud Data Reliability Management Based on Proactive Replica Checking .....	564
<i>Wenhao Li, Yun Yang, Jinjun Chen, and Dong Yuan</i>	
A Model and Decision Procedure for Data Storage in Cloud Computing .....	572
<i>Arkaitz Ruiz-Alvarez and Marty Humphrey</i>	
<b>Session 6A: Multicore Architectures</b>	
A Bandwidth-Optimized Multi-core Architecture for Irregular Applications .....	580
<i>Simone Secchi, Antonino Tumeo, and Oreste Villa</i>	
Parallel Real-Time OLAP on Multi-core Processors .....	588
<i>Frank Dehne and Hamdireza Zaboli</i>	



Minimizing the Data Transfer Time Using Multicore End-System Aware Flow Bifurcation .....	595
<i>Vishal Ahuja, Dipak Ghosal, and Matthew Farrens</i>	
Cache Conscious Task Regrouping on Multicore Processors .....	603
<i>Xiaoya Xiang, Bin Bao, Chen Ding, and Kai Shen</i>	
<b>Session 6B: Cloud Computing II</b>	
An Analysis of Provisioning and Allocation Policies for Infrastructure-as-a-Service Clouds .....	612
<i>David Villegas, Athanasios Antoniou, Seyed Masoud Sadjadi, and Alexandru Iosup</i>	
The Impact of User Rationality in Federated Clouds .....	620
<i>Marian Mihailescu and Yong Meng Teo</i>	
Automated Tagging for the Retrieval of Software Resources in Grid and Cloud Infrastructures .....	628
<i>Ioannis Katakis, George Pallis, Marios D. Dikaiakos, and Onisiforos Onoufriou</i>	
Time and Cost Sensitive Data-Intensive Computing on Hybrid Clouds .....	636
<i>Tekin Bicer, David Chiu, and Gagan Agrawal</i>	
<b>Session 6C: Applications</b>	
Lightweight Resource Scaling for Cloud Applications .....	644
<i>Rui Han, Li Guo, Moustafa M. Ghanem, and Yike Guo</i>	
Performance Modeling and Comparative Analysis of the MILC Lattice QCD Application su3_rmd .....	652
<i>Greg Bauer, Steven Gottlieb, and Torsten Hoefler</i>	
An Efficient Parallel Implementation for Three-Dimensional Incompressible Pipe Flow Based on SIMPLE .....	660
<i>Ji-Lin Zhang, Li-Ting Zhu, Jian Wan, Jie Mao, Xiang-Hua Xu, Cong-Feng Jiang, and Peng Di</i>	
Parallel Simulation of Peer-to-Peer Systems .....	668
<i>Martin Quinson, Cristian Rosa, and Christophe Thiéry</i>	
<b>Poster Papers</b>	
Automated Construction of Performance Models for High Performance Distributed Applications .....	676
<i>Ahmad Mizan and Greg Franks</i>	

A Cloud Infrastructure for Optimization of a Massive Parallel Sequencing Workflow .....	678
<i>Olivier Terzo, Lorenzo Mossucca, Andrea Acquaviva, Francesco Abate, and Rosalba Provenzano</i>	
Pseudo Random Number Generation for Parallelized Jobs on Clusters .....	680
<i>Mike Mikailov, Sithu D. Sudarsan, and Fu-Jyh Luo</i>	
Creating Your Own Private Cloud: Ezilla Toolkit—For Coordinated Storage, Computing, and Networking Services .....	682
<i>Yi-Lun Pan, Chang-Hsing Wu, Hsi-En Yu, Hui-Shan Chen, and Weicheng Huang</i>	
On Effective Quality of Service Negotiation .....	684
<i>Khalid Mansour, Ryszard Kowalczyk, and Mohan Baruwal Chhetri</i>	
IDSaaS: Intrusion Detection System as a Service in Public Clouds .....	686
<i>Turki Alharkan and Patrick Martin</i>	
Distributed Shared Memory and Compiler-Induced Scalable Locality for Scalable Cluster Performance .....	688
<i>Mohamed Abdalkader, Ian Burnette, Tim Douglas, and David G. Wonnacott</i>	
A Comparative Study of Cloud Computing Middleware .....	690
<i>Chaker El Amrani, Kaoutar Bahri Filali, Kaoutar Ben Ahmed, Amadou Tidiane Diallo, Stéphane Telolahy, and Tarek El-Ghazawi</i>	
Automatic Adaptive Page-Size Control for Remote Memory Paging .....	694
<i>Hiroko Midorikawa and Joe Uchiyama</i>	
Analyzing Effect of Network Processor’s Cache Dependent Parameter on MPI Broadcast Performance .....	697
<i>Kedar Kulkarni and Geetanjali Gadre</i>	
Advanced MAC in HPC Systems: Performance Improvement .....	699
<i>D. Gros, M. Blanc, J. Briffaut, and C. Toinard</i>	
Cluster as a Service for Self-Deployable Cloud Applications .....	703
<i>Shigetoshi Yokoyama and Nobukazu Yoshioka</i>	
Accelerating 2-opt and 3-opt Local Search Using GPU in the Travelling Salesman Problem .....	705
<i>Kamil Rocki and Reiji Suda</i>	
Distributed Shared Memory Programming in the Cloud .....	707
<i>Ahmad Anbar, Vikram K. Narayana, and Tarek El-Ghazawi</i>	
A Fault Tolerance Framework for High Performance Computing in Cloud .....	709
<i>Ifeanyi P. Egwutuoha, Shiping Chen, David Levy, and Bran Selic</i>	
Business Process Engine Simulator .....	711
<i>Suraj Pandey, Surya Nepal, and Shiping Chen</i>	

Improving MapReduce Performance in Heterogeneous Network Environments and Resource Utilization .....	714
<i>Zhenhua Guo and Geoffrey Fox</i>	
Perspectives of UnaCloud: An Opportunistic Cloud Computing Solution for Facilitating Research .....	717
<i>Juan D. Osorio, Harold Castro, and Francisco Brasileiro</i>	
Provisioning-Based Resource Management for Effective Workflow Scheduling on Utility Grids .....	719
<i>Vahid Khajevand, Hossein Pedram, and Mostafa Zandieh</i>	
Mining Concept Drifting Network Traffic in Cloud Computing Environments .....	721
<i>Sai Kiran Mukkavilli and Sachin Shetty</i>	

## **Doctoral Symposium**

### **Cloud Scheduling**

VM Economics for Java Cloud Computing: An Adaptive and Resource-Aware Java Runtime with Quality-of-Execution .....	723
<i>José Simão and Luís Veiga</i>	
Combinatorial Auction-Based Mechanisms for VM Provisioning and Allocation in Clouds .....	729
<i>Sharrukh Zaman and Daniel Grosu</i>	
Optimistic Scheduling with Geographically Replicated Services in the Cloud Environment (COLOR) .....	735
<i>Wenbo Zhu and Murray Woodside</i>	

### **Clusters and Data Centers**

Kernel-Assisted MPI Collective Communication among Many-core Clusters .....	741
<i>Teng Ma</i>	
Enabling Application Resilience with and without the MPI Standard .....	746
<i>Wesley Bland</i>	
Leveraging Heterogeneity for Energy Minimization in Data Centers .....	752
<i>Marina Zapater, José L. Ayala, and José M. Moya</i>	

### **Workflows and MapReduce**

Executing Data-Intensive Workloads in a Cloud .....	758
<i>Rizwan Mian and Patrick Martin</i>	

Integration of Workflow Partitioning and Resource Provisioning .....	764
<i>Weiwei Chen and Ewa Deelman</i>	
Hierarchical MapReduce Programming Model and Scheduling Algorithms .....	769
<i>Yuan Luo and Beth Plale</i>	
Resource Management for Elastic Cloud Workflows .....	775
<i>Li Yu and Douglas Thain</i>	

## Workshops

### **International Workshop on Cloud for Business, Industry and Enterprises (C4BIE 2012)**

Enhanced Energy-Efficient Scheduling for Parallel Applications in Cloud .....	781
<i>Qingjia Huang, Sen Su, Jian Li, Peng Xu, Kai Shuang, and Xiao Huang</i>	
The Cloud: Requirements for a Better Service .....	787
<i>Eileen Marie Hanna, Nader Mohamed, and Jameela Al-Jaroodi</i>	
Reducing Operational Costs through Consolidation with Resource Prediction in the Cloud .....	793
<i>Jian Li, Kai Shuang, Sen Su, Qingjia Huang, Peng Xu, Xiang Cheng, and Jie Wang</i>	
Cellular STEM: Promoting Interest in Science, Technology, Engineering, and Math Education Using Cellular Messaging, Cloud Computing, and Web-Based Social Networks .....	799
<i>Brandon C. Judd and Corey A. Graves</i>	
Design and Implementation of a Secure Healthcare Social Cloud System .....	805
<i>Ryan Wooten, Roger Klink, Frank Sinek, Yan Bai, and Meeta Sharma</i>	

### **Workshop on Cloud Computing Optimization (CCOPT 2012)**

Optimizing Completion Time and Resource Provisioning of Pig Programs .....	811
<i>Zhuoyao Zhang, Ludmila Cherkasova, Abhishek Verma, and Boon Thau Loo</i>	
Pricing and Resource Allocation in a Cloud Computing Market .....	817
<i>Linna Du</i>	
Towards the Automated Engineering of Dependable Adaptive Services .....	823
<i>Thar Baker, Martin Randles, and A. Taleb-Bendiab</i>	
Integrity Verification of Multiple Data Copies over Untrusted Cloud Servers .....	829
<i>Ayad F. Barsoum and M. Anwar Hasan</i>	
Optimal Reconfiguration of the Cloud Network for Maximum Energy Savings .....	835
<i>Burak Kantarci and Hussein T. Mouftah</i>	

Optimal Location of Data Centers and Software Components in Cloud Computing Network Design .....	841
<i>Federico Larumbe and Brunilde Sansò</i>	
<b>Workshop on Modeling and Simulation on Grid and Cloud Computing (MSGC 2012)</b>	
Reducing Complexity in Management of eScience Computations .....	845
<i>Per-Olov Östberg, Andreas Hellander, Brian Drawert, Erik Elmroth, Sverker Holmgren, and Linda Petzold</i>	
Conservative Distributed Discrete Event Simulation on Amazon EC2 .....	853
<i>Kurt Vanmechelen, Silas De Munck, and Jan Broeckhove</i>	
Integrating HLA and Service-Oriented Architecture in a Simulation Framework .....	861
<i>Monica Dragoicea, Laurentiu Bucur, Wei-Tek Tsai, and Hessam Sarjoughian</i>	
Simulation in the Cloud Using Handheld Devices .....	867
<i>Emilio Mancini, Gabriel Wainer, Khaldoon Al-Zoubi, and Olivier Dalle</i>	
Goal-Directed Grid-Enabled Computing for Legacy Simulations .....	873
<i>Ernest H. Page, Laurie Litwin, Matthew T. McMahon, Brian Wickham, Mike Shadid, and Elizabeth Chang</i>	
U.S. Army Modeling and Simulation Executable Architecture Deployment Cloud Virtualization Strategy .....	880
<i>Shaun Murphy, Scott Gallant, Chris Gaughan, and Manny Diego</i>	
Simulation Processes in the Cloud for Emergency Planning .....	886
<i>Judicaël Ribault and Gabriel Wainer</i>	
Development of a Metamodel for Medical Database Management on a Grid Network: Application to Health Watch and Epidemiology for Cancer and Perinatal Health .....	892
<i>Sébastien Capière, Paul De Vlieger, David Sarramia, David R.C. Hill, and Lydia Maigne</i>	

## **First International Workshop on Data--Intensive Process Management in Large-Scale Sensor Systems (DPMSS 2012)**

Towards Ontology-based Data Quality Inference in Large-Scale Sensor Networks .....	898
<i>Sam Esswein, Sebastien Goasguen, Chris Post, Jason Hallstrom, David White, and Gene Eidson</i>	
End-to-End QoS on Shared Clouds for Highly Dynamic, Large-Scale Sensing Data Streams .....	904
<i>Rafael Tolosana-Calasanz, José Ángel Bañares, Congduc Pham, and Omer Rana</i>	
Human Postures Recognition Based on D-S Evidence Theory and Multi-sensor Data Fusion .....	912
<i>Wenfeng Li, Junrong Bao, Xiuwen Fu, Giancarlo Fortino, and Stefano Galzarano</i>	
On Managing Very Large Sensor-Network Data Using Bigtable .....	918
<i>Byunggu Yu, Alfredo Cuzzocrea, Dong Jeong, and Sergey Maydebura</i>	
Data Outsourcing Simplified: Generating Data Connectors from Confidentiality and Access Policies .....	923
<i>Konrad Jünemann, Jens Köhler, and Hannes Hartenstein</i>	
<b>Author Index</b> .....	<b>931</b>