

**Sixth IEEE International Symposium on
Cluster Computing and the Grid**

CCGrid06

Spanning the World and Beyond

Proceedings

CCGrid06

Table of Contents

Message from the General and Program Chairs	xiv
Organizing Committee	xvi
Steering Committee	xvii
International Program Committee	xviii
Additional Reviewers	xx
Industrial Sponsors	xxiii

Keynote Speakers

Grid Computing in Drug Discovery	3
<i>Manuel C. Peitsch</i>	
AIST SOA for Building Service Oriented E-infrastructure	4
<i>Satoshi Sekiguchi</i>	
Cyber-Infrastructure in Korea—A Status Report	5
<i>Jysoo Lee</i>	
Building Cyber-Infrastructure for the Biomedical Informatics Research Network	6
<i>Philip M. Papadopoulos</i>	

Track 1

Message Passing Interface (MPI)

MPI-Mitten: Enabling Migration Technology in MPI.....	11
<i>Cong Du and Xian-He Sun</i>	
MPI over uDAPL: Can High Performance and Portability Exist across Architectures?.....	19
<i>Lei Chai, Ranjit Noronha, and Dhabaleswar K. Panda</i>	
Proposal of MPI Operation Level Checkpoint/Rollback and One Implementation.....	27
<i>Yuan Tang, Graham E. Fagg, and Jack J. Dongarra</i>	
Scalable Approaches for Supporting MPI-IO Atomicity.....	35
<i>Peter M. Aarestad, Avery Ching, George K. Thiruvathukal, and Alok N. Choudhary</i>	
Design of High Performance MVAPICH2: MPI2 over InfiniBand.....	43
<i>W. Huang, G. Santhanaraman, H.-W. Jin, Q. Gao, and D. K. Panda</i>	

Peer to Peer Computing (P2P)

Aspect-Oriented Parallel Discrete Optimization on the Cohesion Desktop Grid Platform.....	49
<i>Wolfgang Blochinger, Clemens Dangelmayr, and Sven Schulz</i>	
Density-Based Clustering for Similarity Search in a P2P Network.....	57
<i>Mouna Kacimi and Kokou Yétongnon</i>	
Towards Soft Real-Time Applications on Enterprise Desktop Grids.....	65
<i>Derrick Kondo, Bruno Kindarji, Gilles Fedak, and Franck Cappello</i>	
The Computational and Storage Potential of Volunteer Computing	73
<i>David P. Anderson and Gilles Fedak</i>	
Effective Load Balancing in P2P Systems	81
<i>Zhiyong Xu and Laxmi Bhuyan</i>	
INTCTD: A Peer-to-Peer Approach for Intrusion Detection.....	89
<i>Catalin L. Dumitrescu</i>	
SGII: Combining P2P Data Integration Paradigm and Semantic Web Technology on Top of OGSA-DAI	93
<i>Jingtao Zhou and Mingwei Wang</i>	

Information Services and Resource Discovery (ISRD)

Actor Provenance Capture with Ganglia	99
<i>Ian Wootten, Shrija Rajbhandari, Omer Rana, and J. S. Pahwa</i>	
Architecture Model for Information Service in Large Scale Grid Environments.....	107
<i>Wei Jie, Terence Hung, Stephen J. Turner, and Wentong Cai</i>	
An Extensible Resource Discovery Mechanism for Grid Computing Environments.....	115
<i>Tania Gomes Ramos and Alba Cristina Magalhaes Alves de Melo</i>	
Service Matchmaking with Rough Sets.....	123
<i>Maozhen Li, Bin Yu, Chang Huang, Yong-Hua Song, and Omer F. Rana</i>	
Guarantee of Freshness in Resource Information Cache on WSPE: Web Service Polling Engine.....	131
<i>Fumio Machida, Masahiro Kawato, and Yoshiharu Maeno</i>	
Resource Discovery Using PageRank Technique in Grid Environment.....	135
<i>Noorisyam Hamid, Fazilah Haron, and Chan Huah Yong</i>	

Semantic Web and Web Services (SWWS)

Compilation of XSLT into Dataflow Graphs for Web Service Composition.....	141
<i>Peter M. Kelly, Paul D. Coddington, and Andrew L. Wendelborn</i>	
Statistical Properties of Task Running Times in a Global-Scale Grid Environment.....	150
<i>Menno Dobber, Rob van der Mei, and Ger Koole</i>	
Dynamic Workflow Management Using Performance Data.....	154
<i>Lican Huang, David W. Walker, Omer F. Rana, and Yan Huang</i>	
An Ontology-Based Conceptual Mapping Framework for Translating FBPML to the Web Services Ontology.....	158
<i>Gayathri Nadarajan and Yun-Heh Chen-Burger</i>	
WS-Messenger: A Web Services Based Messaging System for Service-Oriented Grid Computing.....	166
<i>Yi Huang, Aleksander Slominski, Chathura Herath, and Dennis Gannon</i>	
A WSRF-Based Resource Management System of Manufacturing Grid.....	174
<i>Yu'an He, Tao Yu, Lilan Liu, Bin Shen, and Haiyang Sun</i>	
A Core Grid Ontology for the Semantic Grid.....	178
<i>Wei Xing, Marios D. Dikaiakos, and Rizos Sakellariou</i>	

Applications

Integrating Gridcomputing and Metamodeling	185
<i>Dirk Gorissen, Wouter Hendrickx, Karel Crombecq, and Tom Dhaene</i>	

OGSA-DAI and Bioinformatics Grids: Challenges, Experience and Strategies.....	193
<i>S. Y. Crompton, B. M. Matthews, W. A. Gray, A. C. Jones, R. J. White, and J. S. Pahwa</i>	
Biodiversity World: A Problem-Solving Environment for Analysing Biodiversity Patterns.....	201
<i>J. S. Pahwa, P. Brewer, T. Sutton, C. Yesson, M. Burgess, X. Xu, A. C. Jones, R. J. White, W. A. Gray, N. J. Fiddian, F. A. Bisby, A. Culham, N. Caithness, M. Scoble, P. Williams, and S. Bhagwat</i>	
Evaluating Performance and Scalability of Advanced Accelerator Simulations.....	209
<i>Jungmin Lee, Zhiling Lan, J. Amundson, and P. Spentzouris</i>	
The Signal Data Explorer: A High Performance Grid Based Signal Search Tool for Use in Distributed Diagnostic Applications.....	217
<i>Martyn Fletcher, Tom Jackson, Mark Jessop, Bojian Liang, and Jim Austin</i>	
Distribution of a World Space for Real-Time 3D Applications	225
<i>David A. Heitbrink and Sam K. Makki</i>	
PACS-CS: A Large-Scale Bandwidth-Aware PC Cluster for Scientific Computations.....	233
<i>Taisuke Boku, Mitsuhsisa Sato, Akira Ukawa, Daisuke Takahashi, Shinji Sumimoto, Kouichi Kumon, Takashi Moriyama, and Masaaki Shimizu</i>	
Deploying Scientific Applications to the PRAGMA Grid Testbed: Strategies and Lessons.....	241
<i>David Abramson, Amanda Lynch, Hiroshi Takemiya, Yusuke Tanimura, Susumu Date, Haruki Nakamura, Karpjoo Jeong, Suntae Hwang, Ji Zhu, Zhong-hua Lu, Celine Amoreira, Kim Baldridge, Hurng-Chun Lee, Chi-Wei Wang, Horng-Liang Shih, Tomas Molina, Wilfred W. Li, and Peter W. Arzberger</i>	
ResGrid: A Grid-Aware Toolkit for Reservoir Uncertainty Analysis	249
<i>Zhou Lei, Dayong Huang, Archit Kulshrestha, Santiago Pena, Gabrielle Allen, Xin Li, Christopher White, Richard Duff, John R. Smith, and Subhash Kalla</i>	
Interactive Grid-Based Free-Form Shape Modeling.....	253
<i>Anthony Chong, Konstantin Levinski, and Alexei Sourin</i>	
Supporting the Collaborative Construction of Learning Objects Using the Grid	257
<i>Marta Rosatelli, Hermes Senger, Fabrício Silva, Silvio Stanzani, and Cesar Nunes</i>	
A Parallel Genetic Algorithm for the Optimal Design of Multi-body Model Vehicle Suspensions.....	261
<i>Jingjun Zhang, Guangyuan Liu, Ruizhen Gao, and Kanghua Lou</i>	
Network Bandwidth Predictor (NBP): A System for Online Network Performance Forecasting	265
<i>Alaknantha Eswaradass, Xian-He Sun, and Ming Wu</i>	

Evaluation of BPEL to Scientific Workflows.....	269
<i>Asif Akram, David Meredith, and Rob Allan</i>	

Programming Models and Environments (PME)

How the JSDL Can Exploit the Parallelism?.....	275
<i>Ivan Rodero, Francesc Guim, Julita Corbalán, and Jesús Labarta</i>	
Transparent Adaptive Library-Based Checkpointing for Master-Worker Style Parallelism.....	283
<i>Gene Cooperman, Jason Ansel, and Xiaoqin Ma</i>	
SDCS: Simplified Data Communications in Parallel/Distributed Applications	292
<i>Yong Mao, Yunhong Gu, Jia Chen, and Robert L. Grossman</i>	
Integrating Computing Resources on Multiple Grid-Enabled Job Scheduling Systems through a Grid RPC System	296
<i>Yoshihiro Nakajima, Mitsuhsia Sato, Yoshiaki Aida, Taisuke Boku, and Franck Cappello</i>	
JaSkel: A Java Skeleton-Based Framework for Structured Cluster and Grid Computing.....	301
<i>J. F. Ferreira, J. L. Sobral, and A. J. Proença</i>	
GRASG—A Framework for “Gridifying” and Running Applications on Service-Oriented Grids	305
<i>Quoc-Thuan Ho, Terence Hung, Wei Jie, Hoong-Maeng Chan, Emilda Sindhu, Ganesan Subramaniam, Tianyi Zang, and Xiaorong Li</i>	

Track 2

Performance Evaluation (PE)

GRENCHEMARK: A Framework for Analyzing, Testing, and Comparing Grids	313
<i>Alexandru Iosup and Dick Epema</i>	
CPU Load Predictions on the Computational Grid.....	321
<i>Yuanyuan Zhang, Wei Sun, and Yasushi Inoguchi</i>	
Statistical Data Reduction for Efficient Application Performance Monitoring.....	327
<i>Lingyun Yang, Jennifer M. Schopf, Catalin L. Dumitrescu, and Ian Foster</i>	
Improving a Local Learning Technique for Queue Wait Time Predictions	335
<i>Hui Li, Juan Chen, Ying Tao, David Groep, and Lex Wolters</i>	
A Semi-Empirical Model for Maximal LINPACK Performance Predictions	343
<i>Chau-Yi Chou, Hsi-Ya Chang, Shuen-Tai Wang, and Chang-Hsing Wu</i>	

Security

gSET: Trust Management and Secure Accounting for Business in the Grid.....	349
<i>Thomas Weishäupl, Christoph Witzany, and Erich Schikuta</i>	
A Shibboleth-Protected Privilege Management Infrastructure for E-science Education	357
<i>J. Watt, O. Ajayi, J. Jiang, J. Koetsier, and R. O. Sinnott</i>	
Evaluating Provenance-Based Trust for Scientific Workflows	365
<i>Shrija Rajbhandari, Ian Wootten, Ali Shaikh Ali, and Omer F. Rana</i>	
Towards a Secure, Tamper-Proof Grid Platform.....	373
<i>Andrew Cooper and Andrew Martin</i>	
Towards Reliable Trust Establishment in Grid: A Pre-evaluating Set Based Reputation Evaluation Approach.....	381
<i>Xiangli Qu, Xuejun Yang, and Jingwei Zhong</i>	
Trustworthy Auctions for Grid-Style Economies	386
<i>Kris Bubendorfer, Ian Welch, and Blayne Chard</i>	
Policy-Based Access Control Framework for Grid Computing	391
<i>Jin Wu, Chokchai Box Leangsuksun, Vishal Rampure, and Hong Ong</i>	
End-to-End Trustworthy Data Access in Data-Oriented Scientific Computing	395
<i>Sangmi Lee Pallickara, Beth Plale, Liang Fang, and Dennis Gannon</i>	

Data Grid and Management (DGM)

Designing Efficient Cooperative Caching Schemes for Multi-tier Data-Centers over RDMA-Enabled Networks	401
<i>S. Narravula, H.-W. Jin, K. Vaidyanathan, and D. K Panda</i>	
Nexus: A Novel Weighted-Graph-Based Prefetching Algorithm for Metadata Servers in Petabyte-Scale Storage Systems	409
<i>Peng Gu, Yifeng Zhu, Hong Jiang, and Jun Wang</i>	
Optimal Replica Placement Strategy for Hierarchical Data Grid Systems.....	417
<i>Pangfeng Liu and Jan-Jan Wu</i>	
Efficient Multi-source Data Transfer in Data Grids	421
<i>Chien-Min Wang, Chun-Chen Hsu, Hsi-Min Chen, and Jan-Jan Wu</i>	
Relaxed Data Consistency with CONStanza	425
<i>Andrea Domenici, Flavia Donno, Gianni Pucciani, and Heinz Stockinger</i>	
Enabling Transparent Data Sharing in Component Models	430
<i>Gabriel Antoniu, Hinde Lilia Bouziane, Landry Breuil, Mathieu Jan, and Christian Pérez</i>	

Replica Placement Design with Static Optimality and Dynamic Maintainability	434
<i>Rashedur M. Rahman, Ken Barker, and Reda Alhajj</i>	
Efficient Many-to-One Communication for a Distributed RAID	438
<i>Alessandro Di Marco and Giuseppe Ciaccio</i>	
ReCon: A Fast and Reliable Replica Retrieval Service for the Data Grid.....	446
<i>XiaoLi Zhou, Eunsung Kim, Jai Wug Kim, and Heon Y. Yeom</i>	
GridFTP-APT: Automatic Parallelism Tuning Mechanism for Data Transfer Protocol GridFTP	454
<i>Takeshi Ito, Hiroyuki Ohsaki, and Makoto Imase</i>	
Integrating Logical and Physical File Models in the MPI-IO Implementation for “Clusterfile”	462
<i>Florin Isailă, David Singh, Jesús Carretero, Félix Garcia, Gábor Szeder, and Thomas Moschny</i>	
 Middleware and Infrastructure (MI)	
DIMVisual: Data Integration Model for Visualization of Parallel Programs Behavior	473
<i>Lucas Mello Schnorr, Philippe Olivier Alexandre Navaux, and Benhur de Oliveira Stein</i>	
Heimdhal: A History-Based Policy Engine for Grids.....	481
<i>Pedro Gama, Carlos Ribeiro, and Paulo Ferreira</i>	
PhoenixG: A Unified Management Framework for Industrial Information Grid.....	489
<i>Jianfeng Zhan, Gengpu Liu, Lei Wang, Bibo Tu, Yi Jin, Yang Li, Yan Hao, Xuehai Hong, Dan Meng, and Ninghui Sun</i>	
Mirroring Resources or Mapping Requests: Implementing WS-RF for Grid Workflows.....	497
<i>Thomas Heinis, Cesare Pautasso, and Gustavo Alonso</i>	
Assessing Data Virtualization for Irregularly Replicated Large Datasets	505
<i>Bruno Diniz, Diêgo L. Nogueira, André Cardoso, Renato A. Ferreira, Dorgival Guedes, and Wagner Meira Jr.</i>	
Virtual Clusters for Grid Communities	513
<i>I. Foster, T. Freeman, K. Keahey, D. Scheftner, B. Sotomayor, and X. Zhang</i>	
Design and Evaluation of Nemesis, a Scalable, Low-Latency, Message-Passing Communication Subsystem	521
<i>Darius Buntinas, Guillaume Mercier, and William Gropp</i>	
Exploit Failure Prediction for Adaptive Fault-Tolerance in Cluster Computing.....	531
<i>Yawei Li and Zhiling Lan</i>	
Calder Query Grid Service: Insights and Experimental Evaluation	539
<i>Nithya N. Vijayakumar, Ying Liu, and Beth Plale</i>	

Nova: An Approach to On-Demand Virtual Execution Environments for Grids.....	544
<i>Srikanth Sundarrajan, Hariprasad Nellitheertha, Subhabrata Bhattacharya, and Neel Arurkar</i>	
A Robust Reputation System for the Grid	548
<i>Jianshu Weng, Chunyan Miao, and Angela Goh</i>	
Extending the Entry Consistency Model to Enable Efficient Visualization for Code-Coupling Grid Applications	552
<i>Gabriel Antoniu, Loïc Cudennec, and Sébastien Monnet</i>	

Track 3

Resource Management and Scheduling (RMS)

A Co-ordinate Based Resource Allocation Strategy for Grid Environments.....	561
<i>Benjamin Boon Tat Khoo, Bharadwaj Veeravalli, Terence Hung, and Simon See Chong Wee</i>	
Scalable Grid Application Scheduling via Decoupled Resource Selection and Scheduling	568
<i>Yang Zhang, Anirban Mandal, Henri Casanova, Andrew A. Chien, Yang-Suk Kee, Ken Kennedy, and Charles Koelbel</i>	
Learning-Based Negotiation Strategies for Grid Scheduling.....	576
<i>Jiadao Li and Ramin Yahyapour</i>	
A Feedback Mechanism for Network Scheduling in LambdaGrids	584
<i>Pallab Datta, Sushant Sharma, and Wu-Chun Feng</i>	
Dynamic Co-scheduling of Distributed Computation and Replication	592
<i>Huadong Liu, Micah Beck, and Jian Huang</i>	
QoS Oriented Resource Reservation in Shared Environments.....	601
<i>Ming Wu, Xian-He Sun, and Yong Chen</i>	
Analysis of Query Matching Criteria and Resource Monitoring Models for Grid Application Scheduling	609
<i>Ronak Desai, Sameer Tilak, Bhavin Gandhi, Michael J. Lewis, and Nael B. Abu-Ghazaleh</i>	
Adaptive Policing for Token-Exchange Based Management of Shared Computing Resources.....	617
<i>Percival Xavier, Wentong Cai, and Bu-Sung Lee</i>	
Management Middleware for Enterprise Grids	625
<i>Sven Graupner, Nigel Cook, Derek Coleman, and Tilo Nitzsche</i>	
A Portable Grid Infrastructure for Resource-Aware Applications	633
<i>Michal Wegiel, Grzegorz Czajkowski, Laurent Daynes, and Krzysztof Palacz</i>	

A Repository Adapter for Resource Management Information.....	641
<i>T. M. Ong, L. T. Chia, and B. S. Lee</i>	
A Failure-Aware Scheduling Strategy in Large-Scale Cluster System.....	645
<i>Liping Wu, Dan Meng, Jianfeng Zhan, Lei Wang, and Bibo Tu</i>	
Gang Scheduling and Adaptive Resource Allocation to Mitigate Advance Reservation Impact.....	649
<i>Angela C. Sodan, Chintan Doshi, Lawrence Barsanti, and Darren Taylor</i>	
Imprecise Computation Grid Application Model for Flexible Market-Based Resource Allocation	654
<i>Kyong Hoon Kim, Rajkumar Buyya, and Jong Kim</i>	
Author Index	659