Proceedings of the

15th IEEE International Symposium on High Performance Distributed Computing

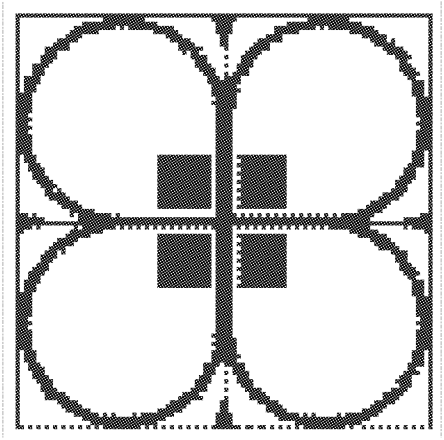

HPDC-15
Paris, France
June 19-23 2006

The papers in this book comprise the digest of the meeting mentioned on the cover and title page. They reflect the authors' opinions and are published as presented and without change, in the interest of timely dissemination. Their inclusion in this publication does not necessarily constitute endorsement by the editors, the Institute of Electrical and Electronics Engineers, Inc.

IEEE Catalog Number:	06TH8878
ISBN:	1-4244-0307-3
ISSN:	1082-8907

TABLE OF CONTENTS

Welcome xiii
Message from the Program Chairs xv
Organization xvii
Reviewers xix
Program xxi
Keynote Abstracts 1

Full Papers

Peer-to-Peer Systems and Overlay Networks

Peer to Peer Size Estimation in Large and Dynamic Networks: A Comparative Study 7
Erwan Le Merrer ..... France Telecom R&D, France
Anne-Marie Kermarrec ..... INRIA/IRISA, France
Laurent Massoulié ..... Microsoft Research, UK

IQ-Paths: Predictably High Performance Data Streams across Dynamic Network Overlays 18
Zhongtang Cai, Vibhore Kumar and Karsten Schwan .....
..... Georgia Institute of Technology, USA

WOW: Self-Organizing Wide Area Overlay Networks of Virtual Workstations 30
Arijit Ganguly, Abhishek Agrawal, P. Oscar Boykin and
Renato Figueiredo ..... University of Florida, USA

Applications

A Case Study Using Automatic Performance Tuning for Large-Scale Scientific Programs 45
I-Hsin Chung ..... IBM Thomas J. Watson Research Center, USA
Jeffrey K. Hollingsworth ..... University of Maryland, USA

Path Grammar Guided Trace Compression and Trace Approximation 57
Xiaofeng Gao ..... University of California at San Diego, USA
Allan Snavely ..... San Diego Supercomputing Center, USA
Larry Carter ..... University of California at San Diego, USA

Filecules in High-Energy Physics: Characteristics and Impact on Resource Management 69
Adriana Iamnitchi and Shyamala Doraimani .....
..... University of South Florida, USA
Gabriele Garzoglio ..... Fermi National Laboratory, USA

Fault Tolerance and Reliability

Fault Tolerance of Tornado Codes for Archival Storage.....83
Matthew Woitaszek and Henry M. Tufo .....
.....University of Colorado at Boulder, USA

Resource Availability Prediction in Fine-Grained Cycle Sharing Systems.....93
Xiaojuan Ren, Seyong Lee, Rudolf Eigenmann and Saurabh Bagchi.....
..... Purdue University, USA

Replicating Nondeterministic Services on Grid Environments.....105
Xianan Zhang and Flavio Junqueira .....
..... University of California at San Diego, USA
Matti Hiltunen.....AT&T Labs, USA
Keith Marzullo..... University of California at San Diego, USA
Richard D. Schlichting.....AT&T Labs, USA

Resource Management

Service Contracts and Aggregate Utility Functions 119
Alvin AuYoung..... University of California at San Diego, USA
Laura Grit.....Duke University, USA
Janet Wiener and John Wilkes.....Hewlett Packard, USA

Market-Based Resource Allocation using Price Prediction in a High Performance Computing Grid for Scientific Applications 132
Thomas Sandholm .....Royal Institute of Technology, Sweden
Kevin Lai ..... Hewlett-Packard Laboratories, USA
Jorge Andrade Ortíz and Jacob Odeberg .....
.....Royal Institute of Technology, Sweden

Optimal Bandwidth Sharing in Grid Environments 144
Loris Marchal, Pascale Vicat-Blanc Primet, Yves Robert and
Jingdi Zeng.....École Normale Supérieure de Lyon, France

A Tool for Prioritizing DAGMan Jobs and Its Evaluation.....156
Grzegorz Malewicz.....Google, USA
Ian Foster ..... University of Chicago/Argonne National Labs, USA
Arnold L. Rosenberg..... University of Massachusetts, USA
Michael Wilde.....Argonne National Labs, USA

Software Environments

Motor: A Virtual Machine for High Performance Computing 171
Wojtek Goscinski and David Abramson... Monash University, Australia

Runtime Support for Memory Adaptation in Scientific Applications via Local Disk and Remote Memory 183
Chuan Yue ..... College of William and Mary, USA
Richard T. Mills ..... Oak Ridge National Laboratory, USA
Andreas Stathopoulos and Dimitrios Nikolopoulos .....
..... College of William and Mary, USA

Building a Generic SOAP Framework over Binary XML 195
Wei Lu ..... Indiana University, USA
Kenneth Chiu ..... State University of New York at Binghamton, USA
Dennis Gannon ..... Indiana University, USA

I/O

Improving I/O Performance of Clustered Storage Systems by Adaptive Request Distribution 207
Changxun Wu and Randal Burns ..... Johns Hopkins University, USA

Improving the Performance of Remote I/O Using Asynchronous Primitives 218
Nawab Ali and Mario Lauria ..... Ohio State University, USA

Exploring I/O Strategies for Parallel Sequence-Search Tools with S3aSim 229
Avery Ching ..... Northwestern University, USA
Wu-chun Feng ..... Virginia Tech, USA
Heshan Lin and Xiaosong Ma..... North Carolina State University, USA
Alok Choudhary ..... Northwestern University, USA

Task Scheduling and File Replication for Data-Intensive Jobs with Batch-shared I/O 241
Gaurav Khanna, Nagavijayalakshmi Vydyanathan, Umit Catalyurek,
Tahsin Kurc, Sriram Krishnamoorthy, P. Sadayappan and Joel Saltz.....
..... Ohio State University, USA

Scheduling

On the Harmfulness of Redundant Batch Requests 255
Henri Casanova ..... University of Hawaii at Manoa, USA

How to Avoid Herd: A Novel Stochastic Algorithm in Grid Scheduling 267
Qinghua Zheng ..... Chinese Academy of Sciences, China
Haijun Yang ..... Beihang University, China
Yuzhong Sun ..... Chinese Academy of Sciences, China

ALPS: An Application-Level Proportional-Share Scheduler 279
Travis Newhouse and Joseph Pasquale .....
..... University of California at San Diego, USA

Scheduling Mixed Workloads in Multi-grids: The Grid Execution Hierarchy 291
Mark Silberstein, Dan Geiger and Assaf Schuster ..... Technion, Israel
Miron Livny ..... University of Wisconsin at Madison, USA

Short Papers

Hot Topics

Data Mining-based Fault Prediction and Detection on the Grid 305
Rubing Duan, Radu Prodan and Thomas Fahringer .....
..... University of Innsbruck, Austria

Troubleshooting Distributed Systems via Data Mining 309
David A. Cieslak, Douglas Thain and Nitesh V. Chawla .....
..... University of Notre Dame, USA

Policy Driven Virtual Machine Monitor for Protected Grids 313
Fabrizio Baiardi and Laura Ricci ..... Università di Pisa, Italy
Paolo Mori and Anna Vaccarelli .....
..... Consiglio Nazionale delle Ricerche, Italy

Autonomic Adaptation of Virtual Distributed Environments in a Multi-Domain Infrastructure 317
Dongyan Xu, Paul Ruth, Junghwan Rhee, Rick Kennell and
Sebastien Goasguen ..... Purdue University, USA

Materializing Highly Available Grids 321
Mark Silberstein, Gabriel Kliot, Artyom Sharov and Assaf Schuster .....
..... Technion, Israel
Miron Livny ..... University of Wisconsin at Madison, USA

Toward Self Organizing Grids 324
Nael Abu-Ghazaleh and Michael J. Lewis .....
..... State University of New York at Binghamton, USA

Posters

Performance and Practicability of Dynamic Adaptation for Parallel Computing 331
Jérémy Buisson ..... IRISA/INSA de Rennes, France
Françoise André ..... Université de Rennes, France
Jean-Louis Pazat ..... IRISA/INSA de Rennes, France

Efficient Services Composition for Grid-Enabled Data-Intensive Applications 333
Tristan Glatard and Johan Montagnat ..... CNRS, France
Xavier Pennec ..... INRIA, France

Ensuring Numerical Quality in Grid Computing 335
Andreas Frommer and Matthias Hüsken .....
..... Bergische Universit¨at Wuppertal, Germany

Integration of Legacy Grid Systems with Emerging Grid Standards 337
A. Grimshaw, W. Kang, D. Merrill and M. Morgan .....
..... University of Virginia, USA

How Should You Structure Your Hierarchical Scheduler? 339
Pushpinder Kaur Chouhan, Holly Dail, Eddy Caron and
Frédéric Vivien ..... École Normale Supérieure de Lyon, France

Robust Resource Allocation for Large-scale Distributed Shared Resource Environments 341
Yang-Suk Kee, Ken Yocum and Andrew A. Chien .....
..... University of California at San Diego, USA
Henri Casanova ..... University of Hawaii at Manoa, USA

Adaptive I/O Scheduling for Distributed Multi-applications Environments 343
Adrien Lebre, Yves Denneulin and Guillaume Huard.....
..... Laboratoire Informatique et Distribution, France
Przemyslaw Sowa ..... Czestochowa University pf Technology, Poland

Dynamic Optimization of Communications over High Speed Networks 345
Elisabeth Brunet, Olivier Aumage and Raymond Namyst .....
..... Université Bordeaux 1, France

Cooperative Caching in the pCFS Parallel Cluster File System 347
Paulo A. Lopes and Pedro D. Medeiros.....
..... Universidade Nova de Lisboa, Portugal

PetaCache: A Memory-Based Data-Server System 349
Chuck Boeheim, Stephen J. Gowdy, Andy Hanushevsky, David Leith, Randy Melen, Richard Mount, Teela Pulliam and Bill Weeks ..... Stanford Linear Accelerator Center, USA

A High Throughput Approach to Combinatorial Search on Grids 351
Yan Liu, Alberto Maria Segre and Shaowen Wang ..... University of Iowa, USA

Is Unmetered, Scalable Computation Worth the Price? 353
Huadong Liu and Micah Beck ..... University of Tennessee, USA

Geometrical Interpretation for Data partitioning on a Grid Architecture 355
Dominique Bernardi ..... Université P. & M. Curie-Paris 6, France
Christophe Cérin ..... Université de Paris Nord, France
Hazem Fkaier and Mohamed Jemni ..... Ecole Supérieure des Sciences et Techniques de Tunis, Tunisia
Michel Koskas ..... Université de Picardie Jules Verne, France

XtremLab: A System for Characterizing Internet Desktop Grids 357
Paul Malecot, Derrick Kondo and Gilles Fedak ..... Université Paris Sud, France

Effective Prediction of Job Processing Times in a Large-Scale Grid Environment 359
Menno Dobber, Rob van der Mei and Ger Koole ..... Vrije Universiteit, The Netherlands

When Jobs Play Nice: The Case for Symbiotic Space-Sharing 361
Jonathan Weinberg and Allan Snavely ..... University of California at San Diego, USA

PARM: Physics Aware Runtime Manager for Large-scale Scientific and Engineering Applications 363
Yeliang Zhang, Salim Hariri, Jianwei Xiang and Jim Yeh ..... University of Arizona, USA

Using File Grouping to Improve the Disk Performance 365
Tsozen Yeh, Joseph Arul, Jia-Shian Wu, I-Fan Chen and Kuo-Hsin Tan ..... Fu Jen Catholic University, Taiwan

Improving Resource Matching Through Estimation of Actual Job Requirements 367
Elad Yom-Tov and Yariv Aridor ..... IBM Haifa Research Laboratory, Israel

Bob++: Framework for Solving Optimization Problems with Branch-and-Bound Methods 369
A. Djerrah and B. Le Cun ..... Universit de Versailles-Saint Quentin en Yvelines, France
V-D. Cung ..... ENSGI-INPG, France,
C. Roucairol ..... Universit de Versailles-Saint Quentin en Yvelines, France

An Implementation of the Message Passing Interface over an Adaptive Peer-to-Peer Network 371
Lei Ni and Aaron Harwood ..... University of Melbourne, Australia

Multidimensional Replica Selection in the Data Grid 373
Sridhar Ramakrishnan and Philip J. Rhodes ..... University of Mississippi, USA

Asynchronous programming with Tarragon 375
Pietro Cicotti and Scott B. Baden ..... University of California at San Diego, USA

Toward a New Direction on Data Management in Grids 377
Aurélien Ortiz, Jacques Jorda and Abdelaziz M'zoughi ..... Institut de Recherche en Informatique de Toulouse, France

PROOF - The Parallel ROOT Facility 379
B. Bellenot, R. Brun, G. Ganis, J. Iwaszkiewicz, G. Kickinger,
A.J. Peters and F. Rademakers ..... CERN, Switzerland
M. Ballintijn, C. Loizides and C. Reed ..... MIT, USA
P. Canal ..... FNAL, USA
D. Feichtinger ..... PSI, Switzerland

Dynamic Programming Based Approach for Bi-criteria Workflow Scheduling on the Grid 381
Marek Wieczorek, Radu Prodan and Thomas Fahringer ..... University of Innsbruck, Austria

RENATER Dark Fibre Project Architecture 383
Philippe d'Anfray and Franck Simon ..... GIP RENATER, France

The SIMGRID Project Simulation and Deployment of Distributed Applications 385
Arnaud Legrand ..... Laboratoire Informatique et Distribution, France
Martin Quinson ..... Université Henri Poincaré, France
Henri Casanova and Kayo Fujiwara ..... University of Hawaii at Manoa, USA

Author Index 387