

Future Generation Grids

Future Generation Grids

**Proceedings of the Workshop on
Future Generation Grids
November 1-5, 2004, Dagstuhl, Germany**

edited by

Vladimir Getov

*University of Westminster
London, UK*

Domenico Laforenza

*Information Science and Technologies Institute
Pisa, Italy*

Alexander Reinefeld

*Zuse-Institut Berlin and
Humboldt-Universität zu Berlin, Germany*



Springer

Vladimir Getov
University of Westminster
London, UK

Domenico Laforenza
Information Science and
Technologies Institute
Pisa, Italy

Alexander Reinefeld
Zuse-Institut Berlin and
Humboldt-Universität zu Berlin
Germany

Library of Congress Control Number: 2005934732

Future Generation Grids

Proceedings of the Workshop on Future Generation Grids

November 1-5, 2004, Dagstuhl, Germany

edited by Vladimir Getov, Domenico Laforenza and Alexander Reinefeld

ISBN-13: 978-0-387-27935-0

ISBN-10: 0-387-27935-0

e-ISBN-13: 978-0-387-29445-2

e-ISBN-10: 0-387-29445-7

Printed on acid-free paper.

© 2006 Springer Science+Business Media, Inc.

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer Science+Business Media, Inc., 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use in this publication of trade names, trademarks, service marks and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

Printed in the United States of America.

9 8 7 6 5 4 3 2 1

SPIN 11392767, 11570783

springeronline.com

Contents

Foreword	vii
Preface	ix
Contributing Authors	xiii
 Part I Architecture	
From Event-Driven Workflows towards A Posteriori Computing <i>Craig A. Lee, B. Scott Michel, Ewa Deelman, and Jim Blythe</i>	3
On Adaptability in Grid Systems <i>Artur Andrzejak, Alexander Reinefeld, Florian Schintke, and Thorsten Schütt</i>	29
Bringing Knowledge to Middleware – Grid Scheduling Ontology <i>Philipp Wieder and Wolfgang Ziegler</i>	47
Remote Administration and Fault Tolerance in Distributed Computer Infrastructures <i>Volker Lindenstruth, Ralph Panse, Timm Steinbeck, Heinz Tilsner, and Arne Wiebalck</i>	61
 Part II Resource and Data Management	
The Virtual Resource Manager: Local Autonomy versus QoS Guarantees for Grid Applications <i>Lars-Olof Burchard, Hans-Ulrich Heiss, Barry Linnert, Jörg Schneider, Felix Heine, Matthias Hovestadt, Odej Kao, and Axel Keller</i>	83
Resource Management for Future Generation Grids <i>Uwe Schwiegelshohn, Ramin Yahyapour, and Philipp Wieder</i>	99
On Designing and Composing Grid Services for Distributed Data Mining <i>Antonio Congiusta, Domenico Talia, and Paolo Trunfio</i>	113
GDS: An Architecture Proposal for a Grid Data-Sharing Service <i>Gabriel Antoniu, Marin Bertier, Luc Bougé, Eddy Caron, Frédéric Desprez, Mathieu Jan, Sébastien Monnet, and Pierre Sens</i>	133

Part III Intelligent Toolkits

A Search Architecture for Grid Software Components 155

Diego Puppini, Fabrizio Silvestri, Domenico Laforenza, and Salvatore Orlando

Use of A Network-Enabled Server Sys tem for a Sparse Linear Algebra Grid Application 171

Eddy Caron, Frédéric Desprez, Jean-Yves L'Excellent, Christophe Hamerling, Marc Pantel, and Chiara Puglisi-Amestoy

Co-Allocation in Grids: Experiences and Issues 195

Anca Bucur, Dick Epema, and Hashim Mohamed

Part IV Programming and Applications

Structured Implementation of Component-Based Grid Programming Environments 217

Marco Aldinucci, Massimo Coppola, Sonia Campa, Marco Danelutto, Marco Vanneschi, and Corrado Zoccolo

From Grid Middleware to Grid Applications: Bridging the Gap with HOCs 241

Sergei Gorlatch and Jan Dünnweber

HPC Application Execution on Grids 263

Marco Danelutto, Marco Vanneschi, Corrado Zoccolo, Nicola Tonellotto, Salvatore Orlando, Ranieri Baraglia, Tiziano Fagni, Domenico Laforenza, and Alessandro Paccosi

Grid Application Programming Environments 283

Thilo Kielmann, Andre Merzky, Henri Bal, Francoise Baude, Denis Caromel, and Fabrice Huet

Index 307

Foreword

The CoreGRID Network of Excellence (NoE) project began in September 2004. Two months later, in November 2004, the first CoreGRID Integration Workshop was held within the framework of the prestigious international Dagstuhl seminars. CoreGRID aims at strengthening and advancing long-term research, knowledge transfer and integration in the area of Grid and Peer-to-Peer technologies. CoreGRID is a Network of Excellence – a new type of project within the European 6th Framework Programme, to ensure progressive evolution and durable integration of the European Grid research community. To achieve this objective, CoreGRID brings together a critical mass of well-established researchers and doctoral students from forty-two institutions that have constructed an ambitious joint programme of activities.

Although excellence is a goal to which CoreGRID is committed, durable integration is our main concern. It means that CoreGRID has to carry out activities to improve the effectiveness of European research in Grid by coordinating and adapting the participants' activities in Grid research, to share resources such as Grid testbeds, to encourage exchange of research staff and students, and to ensure close collaboration and wide dissemination of its results to the international community. Organising CoreGRID Integration Workshops is one of the activities that aims at identifying and promoting durable collaboration between partners involved in the network. It is thus expected that this series of Integration Workshops will provide opportunities for CoreGRID and other researchers to confront their ideas and approaches to solving challenging problems in Grid research, as well as to present the results of their joint research activities. The first Integration Workshop has already demonstrated that close collaborative activities are producing publishable joint result achieved by at least two different CoreGRID partners. At the time this proceedings is being compiled, several indicators show that integration has increased and I encourage you to visit our website¹ to get access to the latest results produced by the network.

¹<http://www.coregrid.net>

Before you start reading this book, I would like to extend my gratitude to the organizers of this first CoreGRID Integration Workshop who did a wonderful job by editing these high quality proceedings. I wish you enjoyable reading of this second volume of the CoreGRID project series of publications.

Thierry Priol, CoreGRID Scientific Co-ordinator

Preface

Since their invention two decades ago, the Internet and the Web have had a significant impact on our life. By allowing us to discover and access information on a global scale, they have enabled the rapid growth of an entirely new industry and brought new meaning to the term “surfing”. However, simply being able to offer and access information on the Web is ultimately unsatisfactory – we want processing and, increasingly, we want collaborative processing within distributed teams. This need has led to the creation of the Grid, an infrastructure that enables us to share capabilities, integrate services and resources within and across enterprises, and allows active collaborations across distributed, multi-organizational environments.

Powered by on-demand access to computer resources, seamless access to data, and dynamic composition of distributed services, the Grid promises to enable fundamentally new ways of interacting with our information technology infrastructure, doing business, and practicing science. It represents perhaps the final step in the great disappearing act that will take computing out of our homes and machine rooms and into the fabric of society, where it will stand alongside telephone exchanges, power generators, and the other invisible technologies that drive the modern world.

Future applications will not only use individual computer systems, but a large set of networked resources. This scenario of computational and data grids is attracting a lot of attention from application scientists, as well as from computer scientists. In addition to the inherent complexity of current high-end systems, the sharing of resources and the transparency of the actual available resources introduce not only new research challenges, but also a completely new vision and novel approaches to designing, building, and using future generation Grid systems.

The Dagstuhl Seminar 04451 on Future Generation Grids (FGG) was held in the International Conference and Research Centre (IBFI), Schloss Dagstuhl¹

¹<http://www.dagstuhl.de>

from 1st to 5th November 2004. The focus of the seminar was on open problems and future challenges in the design of next generation Grid systems.

The seminar brought together 45 scientists and researchers in the area of Grid technologies in an attempt to draw a clearer picture of future generation Grids and to identify some of the most challenging problems on the way to achieving the “invisible” Grid ideas in our society. The participants came from France (12), Germany (10), Italy (8), the United Kingdom (5), the Netherlands (3), Belgium (1), Cyprus (1), the Czech Republic (1), Poland (1), Spain (1), Switzerland (1), and the U.S.A. (1).

This was the first workshop of a series of scientific events planned by the EU Network of Excellence project CoreGRID, the “European Research Network on Foundations, Software Infrastructures and Applications for large scale distributed, GRID and Peer-to-Peer Technologies”. The CoreGRID Network of Excellence, which started in September 2004, aims at strengthening and advancing scientific and technological excellence in the area of Grid and Peer-to-Peer systems.

Additional impetus for the organization of the FGG workshop came from another EU project, the “ERA Pilot on a Coordinated Europe-Wide Initiative in Grid Research” (GridCoord). Its main objective is to strengthen Europe’s position on Grid research and to overcome the fragmentation and dispersion across the EU research programmes. The workshop also gave an overview of the various Grid initiatives and projects and thereby provided a good snapshot of Grid related activities in Europe. Furthermore, the seminar was inspired by the results published in two recent reports by an EU expert group on Next Generation Grids².

In an attempt to provide an overview of the status of the various national Grid initiatives – a topic deemed important especially for the GridCoord project – the following Grid initiatives were presented as part of the discussion sessions:

- DAS-2 (The Netherlands)
- D-Grid (Germany)
- e-Science (UK)
- Grid.it (Italy)
- SGIGrid (Poland)
- ACI GRID’s Grid’5000 project (France)

While the general goal of establishing a national Grid for the benefit of science and research in the respective countries is similar, each of these initiatives

²<http://www.cordis.lu/ist/grids/index.htm>

puts an emphasis on slightly different aspects. Most apparent are perhaps the “virtual laboratories” approach in the Netherlands, the more experimental character of the French Grid 5000 project as part of the ACI GRID initiative, and the strong trend towards the deployment of productive application scenarios in the UK e-Science initiative. However, it is difficult to summarize the subtle differences in the initiatives in this brief preface and therefore, a more detailed analysis must be left for the future.

The discussion session on next generation Grid technologies focused largely on the importance of making Grid systems “autonomic” in the sense that future Grid components should be able to autonomously cope with failures without affecting the other “healthy” components. Even more emphasis was put on the discussion of the newly established Web Services Resources Framework (WSRF) versus the previous Open Grid Service Infrastructure (OGSI), Web Services, and Service Oriented Architectures (SOA) in general.

In this volume, we present a selection of articles based on the topics and results presented at the workshop in Dagstuhl. They are a snapshot of some recent research activities bringing together scientists and researchers in the Grid area. The contents of the proceedings are organised in four parts: Architecture, Resource and Data Management, Intelligent Toolkits, and Programming and Applications.

To conclude, we would like to thank all the participants for their contributions in making the workshop a resounding success; all the staff at Dagstuhl for their professional support in the organization; and, last but not least, all the authors that contributed articles for publication in this volume.

Our thanks also go to the European Commission for sponsoring this volume of selected articles from the workshop via the CoreGRID NoE project, grant number 004265.

Vladimir Getov, Domenico Laforenza, Alexander Reinefeld

Contributing Authors

Marco Aldinucci Institute of Information Science and Technologies, CNR, Via Moruzzi 1, 56100 Pisa, Italy (aldinuc@di.unipi.it)

Artur Andrzejak Zuse Intitute Berlin, Takustr. 7, 14195 Berlin-Dahlem, Germany (andrzejak@zib.de)

Gabriel Antoniu IRISA/INRIA, Campus universitaire de Beaulieu, 35042 Rennes cedex, France (Gabriel.Antoniou@irisa.fr)

Henri Bal Department of Computer Science, Vrije Universiteit, De Boelelaan 1081A, 1081 HV Amsterdam, The Netherlands (bal@cs.vu.nl)

Ranieri Baraglia Institute of Information Science and Technologies, CNR, Via Moruzzi 1, 56100 Pisa, Italy (ranieri.baraglia@isti.cnr.it)

Francoise Baude INRIA, 2004 Rte des Lucioles, BP 93, 06902 Sophia Antipolis cedex, France (Francoise.Baude@sophia.inria.fr)

Marin Bertier IRISA/INRIA, Campus universitaire de Beaulieu, 35042 Rennes cedex, France (marin.bertier@irisa.fr)

Jim Blythe USC/Information Sciences Institute, 4676 Admiralty Way, Suite 1001, Marina del Rey, CA 90292, USA (blythe@isi.edu)

Luc Bougé IRISA/INRIA, Campus universitaire de Beaulieu, 35042 Rennes cedex, France (Luc.Bouge@bretagne.ens-cachan.fr)

Anca Bucur Philips Research, Prof. Holstlaan 4, 5656 AA Eindhoven, The Netherlands (anca.bucur@philips.com)

Sonia Campa Department of Computer Science, University of Pisa, Largo Pontecorvo 3, 56127 Pisa, Italy (compa@di.unipi.it)

Lars-Olof Burchard Communications and Operating Systems Group, Technische Universität Berlin, Fakultät IV, Einsteinufer 17, 10587 Berlin, Germany (baron@cs.tu-berlin.de)

Eddy Caron LIP Laboratory, ENS Lyon, 46 allée d'Italie, 69364 Lyon cedex 7, France (Eddy.Caron@ens-lyon.fr)

Denis Caromel INRIA, 2004 Rte des Lucioles, BP 93, 06902 Sophia Antipolis cedex, France (Denis.Caromel@sophia.inria.fr)

Antonio Congiusta DEIS, University of Calabria, Via P. Bucci 41c, 87036 Rende, Italy (acongiusta@deis.unical.it)

Massimo Coppola Institute of Information Science and Technologies, CNR, Via Moruzzi 1, 56100 Pisa, Italy (coppola@di.unipi.it)

Marco Danelutto Department of Computer Science, University of Pisa, Largo Pontecorvo 3, 56127 Pisa, Italy (marcod@di.unipi.it)

Ewa Deelman USC/Information Sciences Institute, 4676 Admiralty Way, Suite 1001, Marina del Rey, CA 90292, USA (deelman@isi.edu)

Frédéric Desprez LIP Laboratory, ENS Lyon, 46 allée d'Italie, 69364 Lyon cedex 7, France (Frederic.Desprez@ens-lyon.fr)

Jan Dünnweber Institute for Informatics, University of Münster, Einsteinstrasse 62, 48149 Münster, Germany (duennweb@math.uni-muenster.de)

Dick Epema Dept. of Software Technology, Delft University of Technology, Mekelweg 4, 2628 CD Delft, The Netherlands (d.h.j.epema@ewi.tudelft.nl)

Tiziano Fagni Institute of Information Science and Technologies, CNR, Via Moruzzi 1, 56100 Pisa, Italy (tiziano.fagni@isti.cnr.it)

Sergei Gorlatch Institute for Informatics, University of Münster, Einsteinstrasse 62, 48149 Münster, Germany (gorlatch@math.uni-muenster.de)

Christophe Hamerling CERFACS Laboratory, 42 Avenue Gaspard Coriolis, 31057 Toulouse cedex 1, France (Christophe.Hamerling@cerfacs.fr)

Felix Heine Paderborn Center for Parallel Computing (PC²), Universität Paderborn, Fürstenallee 11, 33102 Paderborn, Germany (fh@upb.de)

Hans-Ulrich Heiss Communications and Operating Systems Group, Technische Universität Berlin, Fakultät IV, Einsteinufer 17, 10587 Berlin, Germany (heiss@cs.tu-berlin.de)

Matthias Hovestadt Paderborn Center for Parallel Computing (PC²), Universität Paderborn, Fürstenallee 11, 33102 Paderborn, Germany (mahol@upb.de)

Fabrice Huet INRIA, 2004 Rte des Lucioles, BP 93, 06902 Sophia Antipolis cedex, France (Fabrice.Huet@sophia.inria.fr)

Mathieu Jan IRISA/INRIA, Campus universitaire de Beaulieu, 35042 Rennes cedex, France (Mathieu.Jan@irisa.fr)

Odej Kao Paderborn Center for Parallel Computing (PC²), Universität Paderborn, Fürstenallee 11, 33102 Paderborn, Germany (okao@upb.de)

Axel Keller Paderborn Center for Parallel Computing (PC²), Universität Paderborn, Fürstenallee 11, 33102 Paderborn, Germany (kel@upb.de)

Thilo Kielmann Department of Computer Science, Vrije Universiteit, De Boelelaan 1081A, 1081 HV Amsterdam, The Netherlands (kielmann@cs.vu.nl)

Domenico Laforenza Institute of Information Science and Technologies, CNR, Via Moruzzi 1, 56100 Pisa, Italy (domenico.laforenza@isti.cnr.it)

Craig A. Lee Computer Systems Research Department, The Aerospace Corporation, P.O. Box 92957, El Segundo, CA 90009, USA (lee@aero.org)

Jean-Yves L'Excellent LIP Laboratory, ENS Lyon, 46 allée d'Italie, 69364 Lyon cedex 7, France (Jean-Yves.L.Excellent@ens-lyon.fr)

Volker Lindenstruth Kirchhoff Institute for Physics, Im Neuenheimer Feld 227, 69120 Heidelberg, Germany (ti@kip.uni-heidelberg.de)

Barry Linnert Communications and Operating Systems Group, Technische Universität Berlin, Fakultät IV, Einsteinufer 17, 10587 Berlin, Germany (linnert@cs.tu-berlin.de)

Andre Merzky Department of Computer Science, Vrije Universiteit, De Boelelaan 1081A, 1081 HV Amsterdam, The Netherlands (merzky@cs.vu.nl)

B. Scott Michel Computer Systems Research Department, The Aerospace Corporation, P.O. Box 92957, El Segundo, CA 90009, USA (scottm@aero.org)

Hashim Mohamed Dept. of Software Technology, Delft University of Technology, Mekelweg 4, 2628 CD Delft, The Netherlands (h.h.mohamed@ewi.tudelft.nl)

Sébastien Monnet IRISA, Campus universitaire de Beaulieu, 35042 Rennes cedex, France (Sebastien.Monnet@irisa.fr)

Salvatore Orlando Dipartimento di Informatica, Università Ca'Foscari di Venezia, Via Torino 155, 30172 Mestre, Italy (orlando@unive.it)

Alessandro Paccosi Institute of Information Science and Technologies, CNR, Via Moruzzi 1, 56100 Pisa, Italy (alessandro.paccosi@isti.cnr.it)

Ralph Panse Kirchhoff Institute for Physics, Im Neuenheimer Feld 227, 69120 Heidelberg, Germany (ti@kip.uni-heidelberg.de)

Marc Pantel IRT Laboratory, ENSEEIHT, 2 rue Camichel, BP 7122, 31071 Toulouse cedex 7, France (Marc.Pantel@enseeiht.fr)

Chiara Puglisi-Amestoy IRIT Laboratory, ENSEEIHT, 2 rue Camichel, BP 7122, 31071 Toulouse cedex 7, France (Chiara.Puglisi@enseeiht.fr)

Diego Puppin Institute of Information Science and Technologies, CNR, Via Moruzzi 1, 56100 Pisa, Italy (diego.puppin@isti.cnr.it)

Alexander Reinefeld Zuse Institute Berlin, Takustr. 7, 14195 Berlin-Dahlem; and Humboldt-Universität zu Berlin, Rudower Chaussee 25 (WBC), 12489 Berlin-Adlershof, Germany (ar@zib.de)

Pierre Sens INRIA/LIP6-MSI, 8 rue du Capitaine Scott, 75015 Paris, France (Pierre.Sens@lip6.fr)

Florian Schintke Zuse Institute Berlin, Takustr. 7, 14195 Berlin-Dahlem, Germany (schintke@zib.de)

Jörg Schneider Communications and Operating Systems Group, Technische Universität Berlin, Fakultät IV, Einsteinufer 17, 10587 Berlin, Germany (komm@cs.tu-berlin.de)

Thorsten Schütt Zuse Institute Berlin, Takustr. 7, 14195 Berlin-Dahlem, Germany (schuett@zib.de)

Uwe Schwiegelshohn Computer Engineering Institute, University of Dortmund, Otto-Hahn-Str. 8, 44221 Dortmund, Germany (uwe.schwiegelshohn@udo.edu)

Fabrizio Silvestri Institute of Information Science and Technologies, CNR, Via Moruzzi 1, 56100 Pisa, Italy (fabrizio.silvestri@isti.cnr.it)

Timm Steinbeck Kirchhoff Institute for Physics, Im Neuenheimer Feld 227, 69120 Heidelberg, Germany (ti@kip.uni-heidelberg.de)

Domenico Talia DEIS, University of Calabria, Via P. Bucci 41c, 87036 Rende, Italy (talia@deis.unical.it)

Heinz Tilsner Kirchhoff Institute for Physics, Im Neuenheimer Feld 227, 69120 Heidelberg, Germany (ti@kip.uni-heidelberg.de)

Nicola Tonellotto Institute of Information Science and Technologies, CNR, Via Moruzzi 1, 56100 Pisa, Italy (nicola.tonellotto@isti.cnr.it)

Paolo Trunfio DEIS, University of Calabria, Via P. Bucci 41c, 87036 Rende, Italy (trunfio@deis.unical.it)

Marco Vanneschi Department of Computer Science, University of Pisa, Largo Pontecorvo 3, 56127 Pisa, Italy (vannesch@di.unipi.it)

Arne Wiebalck Kirchhoff Institute for Physics, Im Neuenheimer Feld 227, 69120 Heidelberg, Germany (ti@kip.uni-heidelberg.de)

Philipp Wieder Central Institute for Applied Mathematics, Research Centre Jülich, 52425 Jülich, Germany (ph.wieder@fz-juelich.de)

Ramin Yahyapour Computer Engineering Institute, University of Dortmund, Otto-Hahn-Str. 8, 44221 Dortmund, Germany (ramin.yahyapour@udo.edu)

Wolfgang Ziegler Fraunhofer Institute SCAI, Department for Web-based Applications, Schloss Birlinghoven, 53754 Sankt Augustin, Germany (wolfgang.ziegler@scai.fraunhofer.de)

Corrado Zoccolo Department of Computer Science, University of Pisa, Largo Pontecorvo 3, 56127 Pisa, Italy (zoccolo@di.unipi.it)