

Computer Communications and Networks

For other titles published in this series, go to
www.springer.com/series/4198

The **Computer Communications and Networks** series is a range of textbooks, monographs and handbooks. It sets out to provide students, researchers and non-specialists alike with a sure grounding in current knowledge, together with comprehensible access to the latest developments in computer communications and networking.

Emphasis is placed on clear and explanatory styles that support a tutorial approach, so that even the most complex of topics is presented in a lucid and intelligible manner.

Massimo Cafaro • Giovanni Aloisio
Editors

Grids, Clouds and Virtualization

 Springer

Editors

Dr. Massimo Cafaro
Dipartimento di Ingegneria
dell'Innovazione
Università del Salento
Via per Monteroni
73100 Lecce
Italy
massimo.cafaro@unisalento.it

Prof. Giovanni Aloisio
Dipartimento di Ingegneria
dell'Innovazione
Università del Salento
Via per Monteroni
73100 Lecce
Italy
giovanni.aloisio@unisalento.it

Series Editor

Professor A.J. Sammes, BSc, MPhil, PhD,
FBCS, CEng
Centre for Forensic Computing
Cranfield University
DCMT, Shrivenham
Swindon SN6 8LA
UK

ISSN 1617-7975

ISBN 978-0-85729-048-9

e-ISBN 978-0-85729-049-6

DOI 10.1007/978-0-85729-049-6

Springer London Dordrecht Heidelberg New York

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

Library of Congress Control Number: 2010936502

© Springer-Verlag London Limited 2011

Apart from any fair dealing for the purposes of research or private study, or criticism or review, as permitted under the Copyright, Designs and Patents Act 1988, this publication may only be reproduced, stored or transmitted, in any form or by any means, with the prior permission in writing of the publishers, or in the case of reprographic reproduction in accordance with the terms of licenses issued by the Copyright Licensing Agency. Enquiries concerning reproduction outside those terms should be sent to the publishers.

The use of registered names, trademarks, etc., in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant laws and regulations and therefore free for general use.

The publisher makes no representation, express or implied, with regard to the accuracy of the information contained in this book and cannot accept any legal responsibility or liability for any errors or omissions that may be made.

Cover design: VTEX, Vilnius

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Preface

For more than a decade, the development of grid computing was driven by scientific applications. The need to solve large-scale, increasingly complex problems motivated research on grids systems. Many interesting problems have been solved with the help of grids, for instance, the nug30 Quadratic Assignment Problem.

This challenging optimization problem was posed in 1968 and requires, given a set of n facilities, a set of n locations, a distance specified for each pair of locations, and a flow (weight) specified for each pair of facilities (e.g., the amount of supplies transported between the two facilities), assigning all 30 facilities to the 30 different locations with the goal of minimizing the sum of the distances multiplied by the corresponding flows.

Despite its apparent simplicity, the problem is NP-Hard, and the number of possible assignments is extremely large, so that even if you could check a trillion assignments per second, this process would take over 100 times the age of the universe. However, once the algorithms and software necessary to tackle the previously unsolved problem on a computational grid were developed, solving the problem required nearly a week, with a computational endeavor involving more than 1,000 computational resources working simultaneously at eight institutions geographically distributed in different parts of the world.

The FightAIDS@Home project, which is based on the volunteered computing power of the World Community Grid, aims at testing candidate compounds against the variations (or “mutants”) of HIV that can arise and cause drug resistance.

During November 2009, the project identified several fragments as new candidates for a novel binding site on the peripheral surface of HIV protease. These fragments docked well against the “exo site” and in vitro studies (i.e., “wet lab” experiments in test tubes) will assess their potencies. If these wet lab experiments produce promising results, then these fragments could form the foundation for the development of “allosteric inhibitors” of HIV protease (i.e., “flexibility wedges” that can disrupt the conformational changes that HIV protease must undergo in order to function). These allosteric inhibitors could represent a totally new class of anti-AIDS compounds.

These two examples clearly explain why scientists are now routinely supported in their research by grid infrastructures. But what about business and casual users?

Although projects such as BEinGRID have reported some successful business experiments that may profit from execution in grid environments, it appears that there is not a general business case for the grid. However, recent advances in virtualization techniques, coupled with the increased Internet bandwidth now available, led in 2007 to the concept of cloud computing. The emergence of this new paradigm is mainly based on its simplicity and the affordable price for seamless access to both computational and storage resources.

Virtualization enables cloud computing, providing the ability to run legacy applications on older operating systems, creation of a single system image starting from an heterogeneous collection of machines such as those traditionally found in grid environments, and faster job migration within different virtual machines running on the same hardware. For grid and cloud computing, virtualization is the key for provisioning and fair resource allocation. From the security point of view, since virtual machines run isolated in their sandboxes, this provides an additional protection against malicious or faulty codes.

Clouds provide access to inexpensive hardware and storage resources through very simple APIs, and are based on a pay-per-use model, so that renting these resources is usually much cheaper than acquiring dedicated new ones. Moreover, people are becoming comfortable with storing their data remotely in a cloud environment. Therefore, clouds are being increasingly used by scientists, small and medium sized enterprises, and casual users.

Grids, clouds, and virtualization are exciting technologies that are going to become prominent in the next few years; we expect a wide proliferation in their use, especially clouds since these distributed computing facilities are already accessible at a reasonable cost to many potential users. We also expect grids and clouds to play an ever increasing role in the field of scientific research. It is therefore necessary a thorough understanding of principles and techniques of these fields, and the main aim of this book is to foster awareness of the essential ideas by exploring current and future developments in the area.

The idea of writing this book dates back to the highly successful Grids, Clouds and Virtualization Workshop that we organized in conjunction with the 4th International Conference on Grid and Pervasive Computing (GPC 2009) held in Geneva, 4–8 May 2009. We were contacted by Mr. Wayne Wheeler of Springer, and, after an insightful discussion, we agreed to serve as the editors for the book. Indeed, it is virtually impossible for a single person to write a book covering all of the important aspects of grids, clouds, and virtualization while maintaining the required depth, consistency, and appeal.

We invited many well-known and internationally recognized experts, asking them to contribute their expertise. The book delves into details of grids, clouds, and virtualization, guiding the reader through a collection of chapters dealing with key topics. The bibliography rather than being exhaustive, covers essential reference material. The aim is to avoid an encyclopedic approach since we believe that an attempt to cover everything will instead fail to convey any useful information to the interested readers, an audience including researchers actively involved in the field, undergraduate and graduate students, system designers and programmers, and IT policy makers.

The book may serve both as an introduction and as a technical reference. Our desire and hope is that it will be useful to many people familiarizing with the subject and will contribute to new advances in the field.

Lecce, Italy

Massimo Cafaro
Giovanni Aloisio

Acknowledgements

Every book requires months of preparations, and this book is no exception. We would like to express our gratitude to the contributors for their participation in this project. Without their technical expertise, patience, and efforts, this book would not have been possible.

We are also indebted with the Springer editorial team for their cooperation efforts that made this book a reality. In particular, we are deeply grateful to Mr. Wayne Wheeler, Senior Editor, for his initial proposal and continuous encouragements. We serve as the editors of this book owing to his incredible skills and energy. Special thanks must also go to Mr. Simon Rees, Senior Editorial Assistant, for his dedication, support, and punctuality.

Contents

1	Grids, Clouds, and Virtualization	1
	Massimo Cafaro and Giovanni Aloisio	
2	Quality of Service for I/O Workloads in Multicore Virtualized Servers	23
	J. Lakshmi and S.K. Nandy	
3	Architectures for Enhancing Grid Infrastructures with Cloud Computing	55
	Eduardo Huedo, Rafael Moreno-Vozmediano, Rubén S. Montero, and Ignacio M. Llorente	
4	Scientific Workflows in the Cloud	71
	Gideon Juve and Ewa Deelman	
5	Auspice: Automatic Service Planning in Cloud/Grid Environments	93
	David Chiu and Gagan Agrawal	
6	Parameter Sweep Job Submission to Clouds	123
	P. Kacsuk, A. Marosi, M. Kozlovsky, S. Ács, and Z. Farkas	
7	Energy Aware Clouds	143
	Anne-Cécile Orgerie, Marcos Dias de Assunção, and Laurent Lefèvre	
8	Jungle Computing: Distributed Supercomputing Beyond Clusters, Grids, and Clouds	167
	Frank J. Seinstra, Jason Maassen, Rob V. van Nieuwpoort, Niels Drost, Timo van Kessel, Ben van Werkhoven, Jacopo Urbani, Cerial Jacobs, Thilo Kielmann, and Henri E. Bal	
9	Application-Level Interoperability Across Grids and Clouds	199
	Shantenu Jha, Andre Luckow, Andre Merzky, Miklos Erdely, and Saurabh Sehgal	

Glossary 231

Index 233

Contributors

S. Ács MTA SZTAKI, P.O. Box 63, 1518 Budapest, Hungary, acs@sztaki.hu

Gagan Agrawal Department of Computer Science and Engineering, The Ohio State University, Columbus, OH 43210, USA, agrawal@cse.ohio-state.edu

Giovanni Aloisio University of Salento, Lecce, Italy, giovanni.aloisio@unisalento.it

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

Massimo Cafaro Dipartimento di Ingegneria dell'Innovazione, Università del Salento, Via per Monteroni, 73100 Lecce, Italy, massimo.cafaro@unisalento.it

David Chiu School of Engineering and Computer Science, Washington State University, Vancouver, WA 98686, USA, david.chiu@wsu.edu

Marcos Dias de Assunção INRIA, LIP Laboratory (UMR CNRS, INRIA, ENS, UCB), University of Lyon, 46 allée d'Italie, 69364 Lyon Cedex 07, France, marcos.dias.de.assuncao@ens-lyon.fr

Ewa Deelman University of Southern California, Marina del Rey, CA, USA, deelman@isi.edu

Niels Drost Department of Computer Science, Vrije Universiteit, De Boelelaan 1081A, 1081 HV Amsterdam, The Netherlands, niels@cs.vu.nl

Miklos Erdely University of Pannonia, Veszprem, Hungary, erdelyim@gmail.com

Z. Farkas MTA SZTAKI, P.O. Box 63, 1518 Budapest, Hungary, zfarkas@sztaki.hu

Eduardo Huedo Universidad Complutense de Madrid, 28040 Madrid, Spain, ehuedo@fdi.ucm.es

Ceriel Jacobs Department of Computer Science, Vrije Universiteit, De Boelelaan 1081A, 1081 HV Amsterdam, The Netherlands, ceriel@cs.vu.nl

Shantenu Jha Louisiana State University, Baton Rouge, 70803, USA,
sjha@cct.lsu.edu

Gideon Juve University of Southern California, Marina del Rey, CA, USA,
juve@usc.edu

P. Kacsuk MTA SZTAKI, P.O. Box 63, 1518 Budapest, Hungary,
kacsuk@sztaki.hu

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

M. Kozlovsky MTA SZTAKI, P.O. Box 63, 1518 Budapest, Hungary,
m.kozlovsky@sztaki.hu

J. Lakshmi SERC, Indian Institute of Science, Bangalore 560012, India, jlakshmi@serc.iisc.ernet.in

Laurent Lefèvre INRIA, LIP Laboratory (UMR CNRS, INRIA, ENS, UCB), University of Lyon, 46 allée d'Italie, 69364 Lyon Cedex 07, France,
laurent.lefevre@inria.fr

Ignacio M. Llorente Universidad Complutense de Madrid, 28040 Madrid, Spain,
llorente@dacya.ucm.es

Andre Luckow Louisiana State University, Baton Rouge, 70803, USA,
aluckow@cct.lsu.edu

Jason Maassen Department of Computer Science, Vrije Universiteit, De Boelelaan 1081A, 1081 HV Amsterdam, The Netherlands, jason@cs.vu.nl

A. Marosi MTA SZTAKI, P.O. Box 63, 1518 Budapest, Hungary, atisu@sztaki.hu

Andre Merzky Louisiana State University, Baton Rouge, 70803, USA, andre@merzky.net

Rubén S. Montero Universidad Complutense de Madrid, 28040 Madrid, Spain,
rubensm@dacya.ucm.es

Rafael Moreno-Vozmediano Universidad Complutense de Madrid, 28040 Madrid, Spain, rmoreno@dacya.ucm.es

S.K. Nandy SERC, Indian Institute of Science, Bangalore 560012, India,
nandy@serc.iisc.ernet.in

Anne-Cécile Orgerie ENS Lyon, LIP Laboratory (UMR CNRS, INRIA, ENS, UCB), University of Lyon, 46 allée d'Italie, 69364 Lyon Cedex 07, France,
annececile.orgerie@ens-lyon.fr

A.J. Sammes, Centre for Forensic Computing, Cranfield University, DCMT, Shrivensham, Swindon SN6 8LA, UK

Saurabh Sehgal Louisiana State University, Baton Rouge, 70803, USA

Frank J. Seinstra Department of Computer Science, Vrije Universiteit, De Boelelaan 1081A, 1081 HV Amsterdam, The Netherlands, fjseins@cs.vu.nl

Jacopo Urbani Department of Computer Science, Vrije Universiteit, De Boelelaan 1081A, 1081 HV Amsterdam, The Netherlands, jacopo@cs.vu.nl

Timo van Kessel Department of Computer Science, Vrije Universiteit, De Boelelaan 1081A, 1081 HV Amsterdam, The Netherlands, timo@cs.vu.nl

Rob V. van Nieuwpoort Department of Computer Science, Vrije Universiteit, De Boelelaan 1081A, 1081 HV Amsterdam, The Netherlands, rob@cs.vu.nl

Ben van Werkhoven Department of Computer Science, Vrije Universiteit, De Boelelaan 1081A, 1081 HV Amsterdam, The Netherlands, ben@cs.vu.nl