

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, UK

Takeo Kanade

Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler

University of Surrey, Guildford, UK

Jon M. Kleinberg

Cornell University, Ithaca, NY, USA

Alfred Kobsa

University of California, Irvine, CA, USA

Friedemann Mattern

ETH Zurich, Switzerland

John C. Mitchell

Stanford University, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

Oscar Nierstrasz

University of Bern, Switzerland

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

TU Dortmund University, Germany

Madhu Sudan

Microsoft Research, Cambridge, MA, USA

Demetri Terzopoulos

University of California, Los Angeles, CA, USA

Doug Tygar

University of California, Berkeley, CA, USA

Gerhard Weikum

Max Planck Institute for Informatics, Saarbruecken, Germany

Fabio Kon Anne-Marie Kermarrec (Eds.)

Middleware 2011

ACM/IFIP/USENIX

12th International Middleware Conference

Lisbon, Portugal, December 12-16, 2011

Proceedings



Springer

Volume Editors

Fabio Kon

University of São Paulo, Department of Computer Science

Rua do Matão, 1010, 05508-090, São Paulo, SP, Brazil

E-mail: fabio.kon@ime.usp.br

Anne-Marie Kermarrec

INRIA-Rennes

Campus Universitaire de Beaulieu, 35042 Rennes Cedex, France

E-mail: anne-marie.kermarrec@inria.fr

Cover photo: "Torre de Belém (night)", © 2007 Andrew Fechey (andrubby),
made available under an Attribution 2.0 Generic (CC BY 2.0) license

ISSN 0302-9743

e-ISSN 1611-3349

ISBN 978-3-642-25820-6

e-ISBN 978-3-642-25821-3

DOI 10.1007/978-3-642-25821-3

Springer Heidelberg Dordrecht London New York

Library of Congress Control Number: Applied for

CR Subject Classification (1998): C.2, H.4, D.2, H.3, H.5, K.6.5

LNCS Sublibrary: SL 2 – Programming and Software Engineering

© IFIP International Federation for Information Processing 2011

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

The use of general descriptive names, 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 protective laws and regulations and therefore free for general use.

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India

Printed on acid-free paper

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

Preface

Welcome to the 12th edition of the ACM/IFIP/USENIX International Middleware Conference.

The importance of middleware software and systems keeps growing in a world where distribution and heterogeneity are the norm. Middleware abstractions are present everywhere, e.g., from data centers to networks of mobile devices, from multi-core architectures to social networks, bridging the gap between many areas including programming languages, distributed algorithms, networks, and databases.

Among the 125 initial submissions (from 28 countries) to Middleware 2011, 22 research papers and 2 industry papers were selected for inclusion in the technical program and the proceedings of the conference, resulting in an acceptance rate of 19%. All papers were reviewed by at least three reviewers and some of them had four or six reviews. After a discussion period and a rigid selection process, in which some good papers had to be cut out, we selected the 24 papers that appear now in the proceedings. Five of those papers passed through a shepherding process to make sure that minor problems detected by some reviewers were resolved before the papers were published. The papers were judged according to their originality, presentation quality, and relevance to the conference. The accepted papers cover a wide range of topics with a slight bias toward papers related to cloud computing and reliability, which are obviously hot topics at the moment. Many submissions were related to emerging cloud computing, data centers and server farms, as well as scalability and performance for system issues. We thank the Industry Track Chairs Dilma da Silva and Jan de Meer for their help with the selection of industry papers.

We are grateful to Rachid Guerraoui from EPFL for being the keynote speaker of Middleware 2011. He related one of the most successful middleware stories, namely, software transactional memory (STM) systems. An invited paper, included in the proceedings, co-authored by Vincent Gramoli and Rachid Guerraoui, conveys the message that STM systems should not only simplify a programmer's life but also be flexible enough to enable skilled programmers to take the most from this abstraction to ensure STM a promising future.

We were also delighted to grant the 10-Year Best Paper Award to Peter Druschel and Antony Rowstron for their Middleware 2001 paper "Pastry: Scalable, Decentralized Object Location and Routing for Large-Scale Peer-to-Peer Systems." The 10-Year Best Paper Award seeks to reward those papers from the Middleware conference that was held exactly 10 years ago based on their impact on the academy and industry in the past 10 years. This tradition started with the 10th Middleware in 2009 in Urbana-Champaign, USA, and had its second edition in 2010 in Bangalore, India; Middleware 2011 in Lisbon, Portugal, hosted the third edition. This paper, published in 2001, is cited more than 4,000 times

according to Google Scholar and describes probably the most well-known structured peer-to-peer overlay platform. This represents one of the most influential papers in the past decade in distributed systems and we are honored that the Middleware conference was the one to give the authors this award.

We would like to express our deepest thanks to the authors of submitted papers, to the Program Committee members for their work in reviewing the papers and their enthusiasm in the discussions, to Paulo Ferreira and Luís Veiga the General Chairs and their team and, finally, to the members of the Steering Committee for their efforts toward making Middleware one of the major venues in distributed systems.

We hope you enjoy the papers in this volume of *Lecture Notes in Computer Science*.

December 2011

Anne-Marie Kermarrec
Fabio Kon

Organization

Middleware 2011 was organized under the joint sponsorship of the Association for Computing Machinery (ACM), the International Federation for Information Processing (IFIP), and USENIX.

Organizing Committee

General Chairs

Paulo Ferreira INESC-ID Lisboa / IST, Portugal
Luís Veiga INESC-ID Lisboa / IST, Portugal

Program Committee Chairs

Anne-Marie Kermarrec INRIA, France
Fabio Kon IME, University of São Paulo, Brazil

Industrial Track Chairs

Dilma da Silva IBM T.J. Watson Research Center, USA
Jan de Meer SmartSpaceLab, Germany

Workshops Chair

Peter Pietzuch Imperial College, UK

Doctoral Symposium Chair

David Eyers University of Otago, New Zealand

Posters and Demos Chair

Sonia Ben Mokhtar CNRS Lyon, France

Local Arrangements Chair

João Barreto INESC-ID Lisboa / IST, Portugal

Publicity Chair

Carlos Ribeiro INESC-ID Lisboa / IST, Portugal

Financial Chair

João Nuno Silva INESC-ID Lisboa / IST, Portugal

Registration Chair

João Leitão INESC-ID Lisboa / IST, Portugal

Conference Web Chairs

Nuno Carvalho INESC-ID Lisboa / IST, Portugal
Edgar Marques INESC-ID Lisboa / IST, Portugal

Proceedings Chairs

João Garcia INESC-ID Lisboa / IST, Portugal
Luís Veiga INESC-ID Lisboa / IST, Portugal

Steering Committee

Gordon Blair Lancaster University, UK (Chair)
Jean Bacon University of Cambridge, UK
Guruduth Banavar IBM, USA
Roy Campbell University of Illinois at Urbana-Champaign, USA
Brian Cooper Google, USA
Jan De Meer SmartSpaceLab, Germany
Fred Douglass IBM Research, USA
Indranil Gupta University of Illinois at Urbana-Champaign, USA
Valérie Issarny INRIA, France
Hans-Arno Jacobsen University of Toronto, Canada
Wouter Joosen KUL-DistriNet, Belgium
Cecilia Mascolo University of Cambridge, UK
Elie Najm ENST, France

Program Committee

Jean Bacon University of Cambridge, UK
Ken Birman Cornell University, USA
Gordon Blair University of Lancaster, UK
Christian Cachin IBM Research - Zurich, Switzerland
Roy Campbell University of Illinois at Urbana-Champaign, USA
Renato Cerqueira PUC-Rio, Brazil
Ranveer Chandra Microsoft Research Redmond, USA
Lucy Cherkasova HP Labs, USA
Brian F. Cooper Google, USA
Paolo Costa Imperial College, UK
Fábio Costa UFG, Brazil
Geoff Coulson University of Lancaster, UK
Koustuv Dasgupta Xerox, India
Fred Douglass EMC, USA
Frank Eliassen University of Oslo, Norway
Markus Endler PUC-Rio, Brazil
Patrick Eugster Purdue University, USA
David Eyers University of Otago, New Zealand
Paulo Ferreira INESC-ID Lisboa / IST, Portugal
Davide Frey INRIA, France
Vincent Gramoli EPFL, Switzerland
Gang Huang Peking University, China
Valérie Issarny INRIA, France

Hans-Arno Jacobsen	University of Toronto, Canada
Wouter Joosen	KUL-DistriNet, Belgium
Flavio Junqueira	Yahoo Research, Barcelona
Anne-Marie Kermarrec	INRIA, France
Steve Ko	SUNY at Buffalo, USA
Fabio Kon	IME, University of São Paulo, Brazil
Cecilia Mascolo	University of Cambridge, UK
Jan De Meer	SmartSpaceLab, Germany
Erwan Le Merrer	Technicolor Research, France
Dejan Milojicic	HP Labs, USA
Klara Nahrstedt	University of Illinois at Urbana-Champaign, USA
Gian Pietro Picco	University of Trento, Italy
Peter Pietzuch	Imperial College, UK
Oriana Riva	Microsoft Research, USA
Étienne Rivière	University of Neuchâtel, Switzerland
Antony Rowstron	Microsoft Research Cambridge, UK
Rick Schantz	BBN Technologies, USA
Douglas Schmidt	SEI-CMU, USA
Dilma da Silva	IBM T.J. Watson Research Center, USA
Maarten van Steen	VU University Amsterdam, The Netherlands
Francois Taiani	Lancaster University, UK
Peter Triantafyllou	University of Patras, Greece
Nalini Venkatasubramanian	University of California at Irvine, USA
Roman Vitenberg	University of Oslo, Norway
Marko Vukolic	Eurecom, France
Ben Zhao	University of California Santa Barbara, USA

Referees

Gustavo L.B. Baptista	Nikos Ntarmos	Gilles Straub
Christos Efstratiou	Partha Pal	Guido Urdaneta
Nebil Ben Mabrouk	Animesh Pathak	José Valerio
Francesco Gadaleta	Aaron Paulos	Luís Veiga
Matthew Gillen	Hangwei Qian	Abhishek Verma
Stefan Guna	Kiran K. Rachuri	José Viterbo Filho
Leila Jalali	Olaf Resch	Jack Xing
Kyungbaek Kim	Kurt Rohloff	Maysam Yabandeh
Bert Lagaisse	Mo Sadoghi	Chunyang Ye
Ilias Leontiadis	Gerhard Schimpf	Xiujian Yi
Walter Mascarenhas	Vinay Setty	Young Yoon
Giuliano Mega	Reza Sherafat	Kaiwen Zhang
Hein Meling	Francisco Silva e Silva	
Kianoosh Mokhtarian	Alessandro Sorniotti	

Sponsoring Institutions



International Federation for Information Processing
<http://www.ifip.org>



Association for Computing Machinery
<http://www.acm.org>



Advanced Computing Systems Association
<http://www.usenix.org>



Instituto Superior Técnico
<http://www.ist.utl.pt>

INSTITUTO
SUPERIOR
TÉCNICO



INESC-ID Lisboa
<http://www.inesc-id.pt>

Corporate Sponsors



BBN Technologies
<http://www.bbn.com>



Hewlett-Packard Company
www.hp.com



Innovation Makers
<http://www.innovmakers.com>

Table of Contents

Invited Paper

Democratizing Transactional Programming	1
<i>Vincent Gramoli and Rachid Guerraoui</i>	

Social Networks

Scaling Microblogging Services with Divergent Traffic Demands	20
<i>Tianyin Xu, Yang Chen, Lei Jiao, Ben Y. Zhao, Pan Hui, and Xiaoming Fu</i>	

Contrail: Enabling Decentralized Social Networks on Smartphones	41
<i>Patrick Stuedi, Iqbal Mohamed, Mahesh Balakrishnan, Z. Morley Mao, Venugopalan Ramasubramanian, Doug Terry, and Ted Wobber</i>	

Confidant: Protecting OSN Data without Locking It Up	61
<i>Dongtao Liu, Amre Shakimov, Ramón Cáceres, Alexander Varshavsky, and Landon P. Cox</i>	

Storage and Performance Management

Live Deduplication Storage of Virtual Machine Images in an Open-Source Cloud	81
<i>Chun-Ho Ng, Mingcao Ma, Tsz-Yeung Wong, Patrick P.C. Lee, and John C.S. Lui</i>	

Scalable Load Balancing in Cluster Storage Systems	101
<i>Gae-won You, Seung-won Hwang, and Navendu Jain</i>	

Predico: A System for What-If Analysis in Complex Data Center Applications	123
<i>Rahul Singh, Prashant Shenoy, Maitreya Natu, Vaishali Sadaphal, and Harrick Vin</i>	

Green Computing and Resource Management

GreenWare: Greening Cloud-Scale Data Centers to Maximize the Use of Renewable Energy	143
<i>Yanwei Zhang, Yefu Wang, and Xiaorui Wang</i>	

Resource Provisioning Framework for MapReduce Jobs with Performance Goals 165
Abhishek Verma, Ludmila Cherkasova, and Roy H. Campbell

Resource-Aware Adaptive Scheduling for MapReduce Clusters 187
Jordà Polo, Claris Castillo, David Carrera, Yolanda Becerra, Ian Whalley, Malgorzata Steinder, Jordi Torres, and Eduard Ayguadé

Notification and Streaming

A Content-Based Publish/Subscribe Matching Algorithm for 2D Spatial Objects 208
Athanasios Konstantinidis, Antonio Carzaniga, and Alexander L. Wolf

FAIDECS: Fair Decentralized Event Correlation 228
Gregory Aaron Wilkin, K.R. Jayaram, Patrick Eugster, and Ankur Khetrapal

AmbiStream: A Middleware for Multimedia Streaming on Heterogeneous Mobile Devices 249
Emil Andriescu, Roberto Speicys Cardoso, and Valérie Issarny

Virtualizing Stream Processing 269
Michael Duller, Jan S. Rellermeyer, Gustavo Alonso, and Nesime Tatbul

Replication and Caching

Leader Election for Replicated Services Using Application Scores 289
Diogo Becker, Flavio Junqueira, and Marco Serafini

PolyCert: Polymorphic Self-optimizing Replication for In-Memory Transactional Grids 309
Maria Couceiro, Paolo Romano, and Luis Rodrigues

A Trigger-Based Middleware Cache for ORMs 329
Priya Gupta, Nickolai Zeldovich, and Samuel Madden

Security and Interoperability

Deploy, Adjust and Readjust: Supporting Dynamic Reconfiguration of Policy Enforcement 350
Gabriela Gheorghe, Bruno Crispo, Roberto Carbone, Lieven Desmet, and Wouter Joosen

A Middleware Layer for Flexible and Cost-Efficient Multi-tenant Applications	370
<i>Stefan Walraven, Eddy Truyen, and Wouter Joosen</i>	
Bridging the Interoperability Gap: Overcoming Combined Application and Middleware Heterogeneity	390
<i>Yérom-David Bromberg, Paul Grace, Laurent Réveillère, and Gordon S. Blair</i>	
Run-Time (Re)configuration and Inspection	
The Role of Ontologies in Emergent Middleware: Supporting Interoperability in Complex Distributed Systems	410
<i>Gordon S. Blair, Amel Bennaceur, Nikolaos Georgantas, Paul Grace, Valérie Issarny, Vatsala Nundloll, and Massimo Paolucci</i>	
Co-managing Software and Hardware Modules through the Juggle Middleware	431
<i>Jan S. Rellermeyer and Ramon Küpfer</i>	
A Generic Solution for Agile Run-Time Inspection Middleware	451
<i>Wouter De Borger, Bert Lagaisse, and Wouter Joosen</i>	
Industry	
A Comparison of Secure Multi-Tenancy Architectures for Filesystem Storage Clouds	471
<i>Anil Kurmus, Moitrayee Gupta, Roman Pletka, Christian Cachin, and Robert Haas</i>	
SAFEWEB: A Middleware for Securing Ruby-Based Web Applications	491
<i>Petr Hosek, Matteo Migliavacca, Ioannis Papagiannis, David M. Eyers, David Evans, Brian Shand, Jean Bacon, and Peter Pietzuch</i>	
Author Index	513