

Maarten van Steen  
Michi Henning (Eds.)

LNCS 4290

# Middleware 2006

ACM/IFIP/USENIX 7th International Middleware Conference  
Melbourne, Australia, November/December 2006  
Proceedings



Springer



USENIX

Maarten van Steen Michi Henning (Eds.)

# Middleware 2006

ACM/IFIP/USENIX 7th International Middleware Conference  
Melbourne, Australia, November 27–December 1, 2006  
Proceedings



Springer

## Volume Editors

Maarten van Steen  
Vrije Universiteit Amsterdam  
Faculty of Sciences  
Department of Computer Sciences  
De Boelelaan 1081a, 1081 HV Amsterdam, Netherlands  
E-mail: steen@cs.vu.nl

Michi Henning  
ZeroC, Inc.  
4521 PGA Blvd. #335, Palm Beach Gardens, FL 33418, USA  
michi@zeroc.com

Library of Congress Control Number: 2006935877

CR Subject Classification (1998): C.2.4, D.4, C.2, D.1.3, D.3.2, D.2, H.2, H.4

LNCS Sublibrary: SL 2 – Programming and Software Engineering

ISSN	0302-9743
ISBN-10	3-540-49023-X Springer Berlin Heidelberg New York
ISBN-13	978-3-540-49023-4 Springer Berlin Heidelberg New York

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.

Springer is a part of Springer Science+Business Media  
springer.com

© IFIP International Federation for Information Processing, Hofstrasse 3, A-2361 Laxenburg, Austria 2006  
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India  
Printed on acid-free paper SPIN: 11925071 06/3142 5 4 3 2 1 0

*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*

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

*University of Dortmund, Germany*

Madhu Sudan

*Massachusetts Institute of Technology, MA, USA*

Demetri Terzopoulos

*University of California, Los Angeles, CA, USA*

Doug Tygar

*University of California, Berkeley, CA, USA*

Moshe Y. Vardi

*Rice University, Houston, TX, USA*

Gerhard Weikum

*Max-Planck Institute of Computer Science, Saarbruecken, Germany*

# Preface

Middleware is one of those topics in computer science for which it appears difficult to reach consensus on its exact meaning. Broadly speaking, one could say that middleware contains solutions to the distribution of processes, data, and control that are more or less independent from applications, and that allow underlying platforms and hardware to be hidden from applications. In other words, it covers a lot.

However, there does seem to be consensus on the fact that middleware is about distributed systems, and that the solutions incorporated into middleware are applicable to a wide range of applications. Following the trend of past Middleware conferences, this seventh edition has continued to take a broad perspective on what middleware is all about, and there was general agreement among the Program Committee members that we should be open-minded as to what should be considered on topic or not. This open-mindedness is reflected in an interesting collection of papers that cover many fields of middleware, and even touch upon areas that have traditionally belonged more to the systems arena, such as virtualization.

However, not everything changes. As usual, the number of strong submissions was remarkably high, and there were many discussions among committee members as to which papers to accept. (Almost every paper was reviewed by four committee members.) Eventually, we selected 21 out of the 122 submissions, with space limitations forcing us to reject even papers that reflected good and original research.

We would like to thank all authors who submitted papers for Middleware 2006. Also, we both feel that we had a strong committee with members who not only did an excellent job reviewing submissions, but also submitted their reviews on time and acted promptly during the discussion phase, which allowed us to send out notifications to the authors as originally planned. We also gratefully acknowledge the work done by external reviewers, who often provided detailed and high-quality reports.

September 2006

Maarten van Steen and Michi Henning

# Organization

Middleware 2005 was organized under the auspices of IFIP TC6WG6.1 (International Federation for Information Processing, Technical Committee 6 [Communication Systems], Working Group 6.1 [Architecture and Protocols for Computer Networks]).

## Executive Committee

General Chairs:	Joe Sventek (University of Glasgow, UK) Shanika Karunasekera (University of Melbourne, Australia)
Program Chairs:	Michi Henning (ZeroC, USA) Maarten van Steen (Vrije Universiteit Amsterdam, Netherlands)
Local Arrangements Chairs:	Aaron Harwood and Lars Kulik (University of Melbourne, Australia) Lars Kulik and Lars Kulik (University of Melbourne, Australia)
Workshop Chair:	Antony Rowstron (Microsoft Research Cambridge, UK)
Doctoral Symposium Chair:	Karen Henriksen (National ICT Australia, Australia)
Industry Track Chair:	Rajkumar Buyya (University of Melbourne, Australia)
Publicity Chair:	Egemen Tanin (University of Melbourne, Australia)

## Steering Committee

Gordon Blair  
Elie Najm  
Clem Cole  
Joe Sventek  
Jan de Meer  
Doug Schmidt  
Markus Endler  
Steve Vinoski  
Arno Jacobsen  
Jean-Bernard Stefani  
Gustavo Alonso

## Program Committee

Christiana Amza (University of Toronto, Canada)  
Ozalp Babaoglu (University of Bologna, Italy)  
Mark Baker (Independent Consultant)  
Alberto Bartoli (University of Trieste, Italy)  
Yolande Berbers (Leuven University, Belgium)  
Gordon Blair (Lancaster University, UK)  
Chi-Hung Chi (Tsinghua University, China)  
Michele Colajanni (University of Modena, Italy)  
Geoff Coulson (Lancaster University, UK)  
Fred Douglass (IBM Watson, USA)  
Pascal Felber (University of Neuchâtel, Switzerland)  
Indranil Gupta (University of Illinois, Urbana-Champaign, USA)  
Franz Hauck (Ulm University, Germany)  
Bettina Kemme (McGill University, Canada)  
Anne-Marie Kermarrec (INRIA Rennes, France)  
Fabio Kon (IME/USP, Brazil)  
Ihor Kuz (NICTA, Australia)  
Doug Lea (Oswego State University, USA)  
Mark Little (Arjuna Technologies, UK)  
Ted McFadden (Latent Ventures, USA)  
Philip McKinley (Michigan State University, USA)  
Jishnu Mukerji (HP Labs, USA)  
Bernard Normier (ZeroC, USA)  
Tamer Ozsu (University of Waterloo, Canada)  
Gian Pietro Picco (Politecnico di Milano, Italy)  
Frank Pilhofer (Mercury Computer Systems, USA)  
Misha Rabinovich (Case Western Reserve University, USA)  
Alexander Reinefeld (ZIB, Germany)  
Luis Rodrigues (University of Lisbon, Portugal)  
Emin Gun Sirer (Cornell University, USA)  
Peter Steenkiste (CMU, USA)  
Stefan Tai (IBM Watson, USA)  
Amin Vahdat (UCSD, USA)  
Aad van Moorsel (Newcastle University, UK)  
Steve Vinoski (IONA, USA)  
Craig Wills (Worcester Polytechnic Institute, USA)

## Referees

Jeannie Albrecht	Trevor Armstrong	Ben Beckmann
Gustavo Alonso	Dhouha Ayed	Marin Bertier
Mauro Andreolini	Gilbert Babin	Lindsay Bradford
Artur Andrzejak	Hitesh Ballani	Ryan Braud

Ciaran Bryce  
 Ramon Caceres  
 Raphael Y. Camargo  
 Claudia Canali  
 Antonio Carzaniga  
 Yatin Chawathe  
 Jin Chen  
 Lei Chen  
 Fabio Moreira Costa  
 Paolo Costa  
 Raphael Y. de Camargo  
 Dave Dice  
 Joerg Domaschka  
 Markus Endler  
 Onyeka Ezenwoye  
 Luis Garces  
 Vittorio Ghini  
 Lukasz Golab  
 Alfredo Goldman  
 Charles Gray  
 Gernot Heiser  
 Kirsten Hildrum  
 Mikael Hoegqvist  
 Yyun Huang  
 Danny Hughes  
 Felix Hupfeld

Eric Kasten  
 Oliver Kennedy  
 Chip Killian  
 David Knoester  
 Manolis Koubarakis  
 Riccardo Lancellotti  
 Tobias Langhammer  
 Guoli Li  
 Yi Lin  
 Giorgia Lodi  
 Francesca Mazzoni  
 Madalin Mihailescu  
 Hugo Miranda  
 Alberto Montresor  
 Luca Mottola  
 Erich Nahum  
 David Oppenheimer  
 Sujay Parekh  
 Cesare Pautasso  
 Kathrin Peter  
 Milan Pirca  
 Marius Pirvu  
 Christian Plattner  
 Barry Porter  
 Davy Preuveneers  
 Hans P. Reiser

Patrick Reynolds  
 Thomas Rblitz  
 Thomas Roeblitz  
 Bin Rong  
 Leonid Ryzhyk  
 S. Masoud Sadjadi  
 Farshad Samimi  
 Florian Schintke  
 Holger Schmidt  
 Thorsten Schuett  
 Sabina Serbu  
 Nir Shavit  
 Santosh Shrivastava  
 Gokul Soundararajan  
 Kevin Walsh  
 Qiang Wang  
 Zhihua Wen  
 Stuart Wheeler  
 Alex Wun  
 Zhen Xiao  
 Yun Yuan  
 Jingdi Zeng  
 Zhinan Zhou

## Sponsoring Institutions



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



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



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



The Melbourne Convention and Visitors Bureau  
<http://www.mcvb.com.au/>



The University of Melbourne, Australia  
<http://www.unimelb.edu.au/>



P2P Networks and Applications Research Group,  
Univ. Melbourne  
<http://mundula.cs.mu.oz.au:8080/p2p>



Grid Computing and Distributed Systems Laboratory,  
Univ. Melbourne  
<http://www.gridbus.org>



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



National ICT Australia  
<http://www.nicta.com.au/>



ZeroC  
<http://www.zeroc.com/>



ARC Res. Netw. on Intell. Sensors, Sensor Netw. & Info.  
Proc.  
<http://www.ee.unimelb.edu.au/ISSNIP/>

# Lecture Notes in Computer Science

For information about Vols. 1–4206

please contact your bookseller or Springer

- Vol. 4313: T. Margaria, B. Steffen (Eds.), *Leveraging Applications of Formal Methods*. IX, 197 pages. 2006.
- Vol. 4300: Y.Q. Shi (Ed.), *Transactions on Data Hiding and Multimedia Security I*. IX, 139 pages. 2006.
- Vol. 4292: G. Bebis, R. Boyle, B. Parvin, D. Koracin, P. Remagnino, A. Nefian, G. Meenakshisundaram, V. Pascucci, J. Zara, J. Molineros, H. Theisel, T. Malzbender (Eds.), *Advances in Visual Computing, Part II*. XXXII, 906 pages. 2006.
- Vol. 4291: G. Bebis, R. Boyle, B. Parvin, D. Koracin, P. Remagnino, A. Nefian, G. Meenakshisundaram, V. Pascucci, J. Zara, J. Molineros, H. Theisel, T. Malzbender (Eds.), *Advances in Visual Computing, Part I*. XXXI, 916 pages. 2006.
- Vol. 4290: M. van Steen, M. Henning (Eds.), *Middleware* 2006. XIII, 425 pages. 2006.
- Vol. 4283: Y.Q. Shi, B. Jeon (Eds.), *Digital Watermarking*. XII, 474 pages. 2006.
- Vol. 4281: K. Barkaoui, A. Cavalcanti, A. Cerone (Eds.), *Theoretical Aspects of Computing - ICTAC*. XV, 371 pages. 2006.
- Vol. 4279: N. Kobayashi (Ed.), *Programming Languages and Systems*. XI, 423 pages. 2006.
- Vol. 4278: R. Meersman, Z. Tari, P. Herrero (Eds.), *On the Move to Meaningful Internet Systems 2006: OTM 2006 Workshops, Part II*. XLV, 1004 pages. 2006.
- Vol. 4277: R. Meersman, Z. Tari, P. Herrero (Eds.), *On the Move to Meaningful Internet Systems 2006: OTM 2006 Workshops, Part I*. XLV, 1009 pages. 2006.
- Vol. 4276: R. Meersman, Z. Tari (Eds.), *On the Move to Meaningful Internet Systems 2006: CoopIS, DOA, GADA, and ODBASE, Part II*. XXXII, 752 pages. 2006.
- Vol. 4275: R. Meersman, Z. Tari (Eds.), *On the Move to Meaningful Internet Systems 2006: CoopIS, DOA, GADA, and ODBASE, Part I*. XXXI, 1115 pages. 2006.
- Vol. 4273: I.F. Cruz, S. Decker, D. Allemang, C. Preist, D. Schwabe, P. Mika, M. Uschold, L. Aroyo (Eds.), *The Semantic Web - ISWC 2006*. XXIV, 1001 pages. 2006.
- Vol. 4272: P. Havinga, M. Lijding, N. Meratnia, M. Wegdam (Eds.), *Smart Sensing and Context*. XI, 267 pages. 2006.
- Vol. 4271: F.V. Fomin (Ed.), *Graph-Theoretic Concepts in Computer Science*. XIII, 358 pages. 2006.
- Vol. 4270: H. Zha, Z. Pan, H. Thwaites, A.C. Addison, M. Forte (Eds.), *Interactive Technologies and Sociotechnical Systems*. XVI, 547 pages. 2006.
- Vol. 4269: R. State, S. van der Meer, D. O'Sullivan, T. Pfeifer (Eds.), *Large Scale Management of Distributed Systems*. XIII, 282 pages. 2006.
- Vol. 4268: G. Parr, D. Malone, M. Ó Foghlú (Eds.), *Autonomic Principles of IP Operations and Management*. XIII, 237 pages. 2006.
- Vol. 4267: A. Helmy, B. Jennings, L. Murphy, T. Pfeifer (Eds.), *Autonomic Management of Mobile Multimedia Services*. XIII, 257 pages. 2006.
- Vol. 4266: H. Yoshiura, K. Sakurai, K. Rannenbergh, Y. Murayama, S. Kawamura (Eds.), *Advances in Information and Computer Security*. XIII, 438 pages. 2006.
- Vol. 4265: N. Lavrač, L. Todorovski, K.P. Jantke (Eds.), *Discovery Science*. XIV, 384 pages. 2006. (Sublibrary LNAI).
- Vol. 4264: J.L. Balcázar, P.M. Long, F. Stephan (Eds.), *Algorithmic Learning Theory*. XIII, 393 pages. 2006. (Sublibrary LNAI).
- Vol. 4263: A. Levi, E. Savas, H. Yenigün, S. Balcisoy, Y. Saygin (Eds.), *Computer and Information Sciences - ISCIS 2006*. XXIII, 1084 pages. 2006.
- Vol. 4261: Y. Zhuang, S. Yang, Y. Rui, Q. He (Eds.), *Advances in Multimedia Information Processing - PCM 2006*. XXII, 1040 pages. 2006.
- Vol. 4260: Z. Liu, J. He (Eds.), *Formal Methods and Software Engineering*. XII, 778 pages. 2006.
- Vol. 4259: S. Greco, Y. Hata, S. Hirano, M. Inuiguchi, S. Miyamoto, H.S. Nguyen, R. Słowiński (Eds.), *Rough Sets and Current Trends in Computing*. XXII, 951 pages. 2006. (Sublibrary LNAI).
- Vol. 4257: I. Richardson, P. Runeson, R. Messnarz (Eds.), *Software Process Improvement*. XI, 219 pages. 2006.
- Vol. 4256: L. Feng, G. Wang, C. Zeng, R. Huang (Eds.), *Web Information Systems - WISE 2006 Workshops*. XIV, 320 pages. 2006.
- Vol. 4255: K. Aberer, Z. Peng, E.A. Rundensteiner, Y. Zhang, X. Li (Eds.), *Web Information Systems - WISE 2006*. XIV, 563 pages. 2006.
- Vol. 4254: T. Grust, H. Höpfner, A. Illarramendi, S. Jablonski, M. Mesiti, S. Müller, P.-L. Patranjan, K.-U. Sattler, M. Spiliopoulou (Eds.), *Current Trends in Database Technology - EDBT 2006*. XXXI, 932 pages. 2006.
- Vol. 4253: B. Gabrys, R.J. Howlett, L.C. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part III*. XXXII, 1301 pages. 2006. (Sublibrary LNAI).
- Vol. 4252: B. Gabrys, R.J. Howlett, L.C. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part II*. XXXIII, 1335 pages. 2006. (Sublibrary LNAI).

- Vol. 4251: B. Gabrys, R.J. Howlett, L.C. Jain (Eds.), Knowledge-Based Intelligent Information and Engineering Systems, Part I. LXVI, 1297 pages. 2006. (Sublibrary LNAI).
- Vol. 4249: L. Goubin, M. Matsui (Eds.), Cryptographic Hardware and Embedded Systems - CHES 2006. XII, 462 pages. 2006.
- Vol. 4248: S. Staab, V. Svátek (Eds.), Managing Knowledge in a World of Networks. XIV, 400 pages. 2006. (Sublibrary LNAI).
- Vol. 4247: T.-D. Wang, X. Li, S.-H. Chen, X. Wang, H. Abbass, H. Iba, G. Chen, X. Yao (Eds.), Simulated Evolution and Learning. XXI, 940 pages. 2006.
- Vol. 4246: M. Hermann, A. Voronkov (Eds.), Logic for Programming, Artificial Intelligence, and Reasoning. XIII, 588 pages. 2006. (Sublibrary LNAI).
- Vol. 4245: A. Kuba, L.G. Nyúl, K. Palágyi (Eds.), Discrete Geometry for Computer Imagery. XIII, 688 pages. 2006.
- Vol. 4244: S. Spaccapietra (Ed.), Journal on Data Semantics VII. XI, 267 pages. 2006.
- Vol. 4243: T. Yakhno, E.J. Neuhold (Eds.), Advances in Information Systems. XIII, 420 pages. 2006.
- Vol. 4242: A. Rashid, M. Aksit (Eds.), Transactions on Aspect-Oriented Software Development II. IX, 289 pages. 2006.
- Vol. 4241: R.R. Beichel, M. Sonka (Eds.), Computer Vision Approaches to Medical Image Analysis. XI, 262 pages. 2006.
- Vol. 4239: H.Y. Youn, M. Kim, H. Morikawa (Eds.), Ubiquitous Computing Systems. XVI, 548 pages. 2006.
- Vol. 4238: Y.-T. Kim, M. Takano (Eds.), Management of Convergence Networks and Services. XVIII, 605 pages. 2006.
- Vol. 4237: H. Leitold, E. Markatos (Eds.), Communications and Multimedia Security. XII, 253 pages. 2006.
- Vol. 4236: L. Breveglieri, I. Koren, D. Naccache, J.-P. Seifert (Eds.), Fault Diagnosis and Tolerance in Cryptography. XIII, 253 pages. 2006.
- Vol. 4234: I. King, J. Wang, L. Chan, D. Wang (Eds.), Neural Information Processing, Part III. XXII, 1227 pages. 2006.
- Vol. 4233: I. King, J. Wang, L. Chan, D. Wang (Eds.), Neural Information Processing, Part II. XXII, 1203 pages. 2006.
- Vol. 4232: I. King, J. Wang, L. Chan, D. Wang (Eds.), Neural Information Processing, Part I. XLVI, 1153 pages. 2006.
- Vol. 4231: J.F. Roddick, R. Benjamins, S. Si-Saïd Cherfi, R. Chiang, C. Claramunt, R. Elmasri, F. Grandi, H. Han, M. Hepp, M. Lytras, V.B. Mišić, G. Poels, I.-Y. Song, J. Trujillo, C. Vangenot (Eds.), Advances in Conceptual Modeling - Theory and Practice. XXII, 456 pages. 2006.
- Vol. 4230: C. Priami, A. Ingólfssdóttir, B. Mishra, H.R. Nielson (Eds.), Transactions on Computational Systems Biology VII. VII, 185 pages. 2006. (Sublibrary LNBI).
- Vol. 4229: E. Najm, J.F. Pradat-Peyre, V.V. Donzeau-Gouge (Eds.), Formal Techniques for Networked and Distributed Systems - FORTE 2006. X, 486 pages. 2006.
- Vol. 4228: D.E. Lightfoot, C.A. Szyperski (Eds.), Modular Programming Languages. X, 415 pages. 2006.
- Vol. 4227: W. Nejdl, K. Tochtermann (Eds.), Innovative Approaches for Learning and Knowledge Sharing. XVII, 721 pages. 2006.
- Vol. 4226: R.T. Mittermeir (Ed.), Informatics Education - The Bridge between Using and Understanding Computers. XVII, 319 pages. 2006.
- Vol. 4225: J.F. Martínez-Trinidad, J.A. Carrasco Ochoa, J. Kittler (Eds.), Progress in Pattern Recognition, Image Analysis and Applications. XIX, 995 pages. 2006.
- Vol. 4224: E. Corchado, H. Yin, V. Botti, C. Fyfe (Eds.), Intelligent Data Engineering and Automated Learning - IDEAL 2006. XXVII, 1447 pages. 2006.
- Vol. 4223: L. Wang, L. Jiao, G. Shi, X. Li, J. Liu (Eds.), Fuzzy Systems and Knowledge Discovery. XXVIII, 1335 pages. 2006. (Sublibrary LNAI).
- Vol. 4222: L. Jiao, L. Wang, X. Gao, J. Liu, F. Wu (Eds.), Advances in Natural Computation, Part II. XLII, 998 pages. 2006.
- Vol. 4221: L. Jiao, L. Wang, X. Gao, J. Liu, F. Wu (Eds.), Advances in Natural Computation, Part I. XLI, 992 pages. 2006.
- Vol. 4219: D. Zamboni, C. Kruegel (Eds.), Recent Advances in Intrusion Detection. XII, 331 pages. 2006.
- Vol. 4218: S. Graf, W. Zhang (Eds.), Automated Technology for Verification and Analysis. XIV, 540 pages. 2006.
- Vol. 4217: P. Cuenca, L. Orozco-Barbosa (Eds.), Personal Wireless Communications. XV, 532 pages. 2006.
- Vol. 4216: M.R. Berthold, R. Glen, I. Fischer (Eds.), Computational Life Sciences II. XIII, 269 pages. 2006. (Sublibrary LNBI).
- Vol. 4215: D.W. Embley, A. Olivé, S. Ram (Eds.), Conceptual Modeling - ER 2006. XVI, 590 pages. 2006.
- Vol. 4213: J. Fürnkranz, T. Scheffer, M. Spiliopoulou (Eds.), Knowledge Discovery in Databases: PKDD 2006. XXII, 660 pages. 2006. (Sublibrary LNAI).
- Vol. 4212: J. Fürnkranz, T. Scheffer, M. Spiliopoulou (Eds.), Machine Learning: ECML 2006. XXIII, 851 pages. 2006. (Sublibrary LNAI).
- Vol. 4211: P. Vogt, Y. Sugita, E. Tuci, C. Nehaniv (Eds.), Symbol Grounding and Beyond. VIII, 237 pages. 2006. (Sublibrary LNAI).
- Vol. 4210: C. Priami (Ed.), Computational Methods in Systems Biology. X, 323 pages. 2006. (Sublibrary LNBI).
- Vol. 4209: F. Crestani, P. Ferragina, M. Sanderson (Eds.), String Processing and Information Retrieval. XIV, 367 pages. 2006.
- Vol. 4208: M. Gerndt, D. Kranzlmüller (Eds.), High Performance Computing and Communications. XXII, 938 pages. 2006.
- Vol. 4207: Z. Ésik (Ed.), Computer Science Logic. XII, 627 pages. 2006.

# Table of Contents

## Performance I

Caching Dynamic Web Content: Designing and Analysing an Aspect-Oriented Solution . . . . .	1
<i>Sara Bouchenak, Alan Cox, Steven Dropsho, Sumit Mittal, Willy Zwaenepoel</i>	
Non-intrusive Performance Management for Computer Services . . . . .	22
<i>Magnus Karlsson, Christos Karamanolis</i>	

## Composition

True and Transparent Distributed Composition of Aspect-Components . . . . .	42
<i>Bert Lagaisse, Wouter Joosen</i>	
Policy-Driven Middleware for Self-adaptation of Web Services Compositions . . . . .	62
<i>Abdelkarim Erradi, Piyush Maheshwari, Vladimir Tosic</i>	

## Management I

Living with Nondeterminism in Replicated Middleware Applications . . . . .	81
<i>Joseph Slember, Priya Narasimhan</i>	
Trading Off Resources Between Overlapping Overlays . . . . .	101
<i>Brian F. Cooper</i>	

## Publish/Subscribe Technology

Efficient Probabilistic Subsumption Checking for Content-Based Publish/Subscribe Systems . . . . .	121
<i>Aris M. Ouksel, Oana Jurca, Ivana Podnar, Karl Aberer</i>	
Dynamic Load Balancing in Distributed Content-Based Publish/Subscribe . . . . .	141
<i>Alex King Yeung Cheung, Hans-Arno Jacobsen</i>	
Decentralized Message Ordering for Publish/Subscribe Systems . . . . .	162
<i>Cristian Lumezanu, Neil Spring, Bobby Bhattacharjee</i>	

## Databases

DBFarm: A Scalable Cluster for Multiple Databases .....	180
<i>Christian Plattner, Gustavo Alonso, M. Tamer Özsu</i>	
Queryll: Java Database Queries Through Bytecode Rewriting .....	201
<i>Ming-Yee Iu, Willy Zwaenepoel</i>	

## Mobile and Ubiquitous Computing

Contory: A Middleware for the Provisioning of Context Information on Smart Phones .....	219
<i>Oriana Riva</i>	
Efficient Semantic Service Discovery in Pervasive Computing Environments .....	240
<i>Sonia Ben Mokhtar, Anupam Kaul, Nikolaos Georgantas, Valérie Issarny</i>	

## Security

A Middleware System for Protecting Against Application Level Denial of Service Attacks .....	260
<i>Mudhakar Srivatsa, Arun Iyengar, Jian Yin, Ling Liu</i>	
Generalized Access Control of Synchronous Communication .....	281
<i>Constantin Serban, Naftaly Minsky</i>	

## Datamining Techniques

FMware: Middleware for Efficient Filtering and Matching of XML Messages with Local Data .....	301
<i>K. Selçuk Candan, Mehmet E. Dönderler, Yan Qi, Jaikannan Ramamoorthy, Jong W. Kim</i>	
Synergy: Sharing-Aware Component Composition for Distributed Stream Processing Systems .....	322
<i>Thomas Repantis, Xiaohui Gu, Vana Kalogeraki</i>	

## Performance II

Enforcing Performance Isolation Across Virtual Machines in Xen .....	342
<i>Diwaker Gupta, Ludmila Cherkasova, Rob Gardner, Amin Vahdat</i>	

Low-Overhead Message Tracking for Distributed Messaging .....	363
<i>Seung Jun, Mark Astley</i>	

## Management II

Utility-Driven Proactive Management of Availability in Enterprise-Scale Information Flows .....	382
<i>Zhongtang Cai, Vibhore Kumar, Brian F. Cooper, Greg Eisenhower, Karsten Schwan, Robert E. Strom</i>	

Model Driven Provisioning: Bridging the Gap Between Declarative Object Models and Procedural Provisioning Tools .....	404
<i>Kaoutar El Maghraoui, Alok Meghranjani, Tamar Eilam, Michael Kalantar, Alexander V. Konstantinou</i>	

<b>Author Index .....</b>	<b>425</b>
---------------------------	------------

# Caching Dynamic Web Content: Designing and Analysing an Aspect-Oriented Solution

Sara Bouchenak<sup>1</sup>, Alan Cox<sup>2</sup>, Steven Dropsho<sup>3</sup>,  
Sumit Mittal<sup>4,\*</sup>, and Willy Zwaenepoel<sup>3</sup>

<sup>1</sup> INRIA, 655, av. de l'Europe, Montbonnot, 38334 St. Ismier Cedex, France  
Sara.Bouchenak@inria.fr

<sup>2</sup> Rice University, Department of Computer Science, Houston, TX USA  
alc@cs.rice.edu

<sup>3</sup> EPFL, Department of Computer Science, CH-1015 Lausanne, Switzerland  
{Steven.Dropsho, Willy.Zwaenepoel}@epfl.ch

<sup>4</sup> IBM India Research Lab, Block-1, IIT, Hauz Khas, New Delhi, India  
sumittal@in.ibm.com

**Abstract.** Caching dynamic web content is an effective approach to reduce Internet latency and server load. An ideal caching solution is one that can be added transparently by the developers and provides complete consistency of the cached documents, while minimizing false cache invalidations. In this paper, we design and implement AutoWebCache, a middleware system for adding caching of dynamic content transparently to J2EE server-side applications having a backend database. For this purpose, we first present the principles involved in caching dynamic web content, including our logic to ensure consistency of the cached entries. Thereafter, we demonstrate the use of aspect-oriented (AOP) techniques to implement our system, showing how AOP provides modularity and transparency to the entire process. Further, we evaluate the effectiveness of AutoWebCache in reducing response times of applications, thereby improving throughput. We also analyze the transparency of our system for a general application suite, considering issues such as dynamic web pages aggregating data from multiple sources, presence of insufficiently structured interfaces for exchanging information and the use of application semantics while caching. We use two standard J2EE web benchmark applications, RUBiS and TPC-W, to conduct our experiments and discuss the results obtained.

**Keywords:** Caching, aspect-oriented programming, J2EE applications, dynamic content.

## 1 Introduction

Dynamically generated web content represents a large portion of web requests. The rate at which dynamic documents are delivered is often orders of magnitudes slower than static documents [9,11]. Therefore, caching dynamic web content is

---

\* Work done while being at Rice University, Houston and EPFL University, Lausanne.

an appealing approach to reduce Internet latency and server load. Web sites for dynamic content are usually based on a multi-tier J2EE architecture using several middleware systems [27]: an HTTP server as a web front-end and provider of static content, an application server to execute the business logic and generate the dynamic web content, and a database to store the persistent data required by the application. Dynamic content generation places a significant burden on the servers, often leading to performance bottlenecks. Caching dynamic web content can directly address these bottlenecks.

Implementing caching as a middleware solution is particularly attractive. Of course, an ideal solution is one that can be added transparently by the developers, possibly even as an after-thought. Some examples of transparently adding caching to an application are given in [17,6,4], but these ignore consistency of the cached entries. Other solutions provide consistency, but ignore transparency, requiring manual insertion [10]. There are some projects that provide both consistency and transparency, such as those caching SQL query result sets [8] at the back-end. The interesting property of data from result sets of SQL queries is that it is from a single interface and hence, of one type (homogeneous). An open question is whether similar techniques can be successful for more complex content such as web pages that aggregate data from multiple sources (*i.e.*, heterogeneous).

In this paper, we present the design and implementation of *AutoWebCache*, a middleware solution for caching dynamically generated content in J2EE applications. A goal is to move the caching as far forward in the multi-tier architecture to not only reduce the database activity in the back-end but also the business logic activity, which is becoming ever more complex and costly at the middle tier. Unlike caching data such as JDBC SQL results at a single well-specified interface, caching fully formed web pages requires interfacing to both the front-end (e.g., Tomcat servlet engine) *and* the back-end (e.g., JDBC interface). Caching at this level requires information from both interfaces to maintain consistency of the cached documents. To keep the caching transparent, we cast caching as an aspect of the application and use an aspect oriented programming (AOP) framework to capture the information flowing through various interfaces. We give details of the *AutoWebCache* cache system based on AOP principles and the AspectJ [2] weaving rules that add the caching logic transparently to the application.

We evaluate the performance of our middleware solution with the help of two J2EE benchmarks - RUBiS and TPC-W. RUBiS implements the core functionality of an auction-site: selling, browsing and bidding [1], while TPC-W simulates an online-bookstore [30]. We demonstrate the gains in response times using *AutoWebCache* for each. We also analyze the transparency of *AutoWebCache* for a general application suite. We argue that for the general case, issues can arise when caching dynamic content at the front-end due to 1) dynamic web pages aggregating data from multiple sources, 2) some sources not having sufficiently structured interfaces for exchanging information and 3) the need to consider semantics of the application while caching. Although our benchmark applications are servlets-based and use SQL queries to incorporate dynamism, we believe that the results and arguments presented in this paper hold true for a general architecture as well.

The contributions of this paper can be summarized as follows:

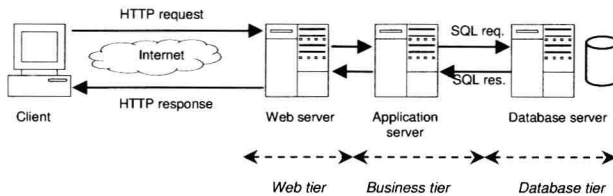
1. Design, implementation and evaluation of AutoWebCache, a middleware solution that caches dynamic web pages at the front-end while maintaining consistency with the back-end database(s).
2. Demonstrating that dynamic web caching can be considered a crosscutting aspect and, therefore, AOP methods should be considered as a flexible and easy-to-use tool to develop the middleware support.

The remainder of this paper is organized as follows. Section 2 gives some background on dynamic web applications and aspect-oriented programming. Section 3 outlines the principles involved in designing a dynamic web cache and gives an overview of our AutoWebCache system. Section 4 describes the implementation of AutoWebCache using aspect-oriented techniques, and analyzes its transparency with respect to an application. Sections 5 and 6 present our evaluation environment and the results of our evaluation, respectively. Section 7 provides a discussion of our experiences. Section 8 discusses some related work and finally, Section 9 draws our conclusions.

## 2 Background

### 2.1 J2EE Web Applications

Java 2 Platform, Enterprise Edition (J2EE) defines a model for developing distributed applications, *e.g.*, web applications, in a multi-tiered architecture [27]. Such applications usually start with requests from web clients that flow through an HTTP server front-end and provider of static content, then to an application server to execute the business logic and generate web pages on-the-fly, and finally to a database that stores resources and data (see Figure 1).



**Fig. 1.** Architecture of Dynamic Web Applications

Upon an HTTP client request, either the request targets a static web document that the web server can return directly; or the request refers to a dynamic document, in which case the web server forwards that request to the application server. The application server runs one or more software components (*e.g.*, Servlets, EJB) that query a database through a JDBC driver (Java DataBase Connection driver) [28] and retrieve data to generate a web document on-the-fly.