

Mario Bravetti
Leïla Kloul
Gianluigi Zavattaro (Eds.)

LNC3 3670

Formal Techniques for Computer Systems and Business Processes

European Performance Engineering Workshop, EPEW 2005
and International Workshop on Web Services
and Formal Methods, WS-FM 2005
Versailles, France, September 2005, Proceedings



Springer

TP3-55
E63
2005

Mario Bravetti Leïla Kloul
Gianluigi Zavattaro (Eds.)

Formal Techniques for Computer Systems and Business Processes

European Performance Engineering Workshop, EPEW 2005
and International Workshop on Web Services
and Formal Methods, WS-FM 2005
Versailles, France, September 1-3, 2005
Proceedings



Springer



E200502122

Volume Editors

Mario Bravetti
Università di Bologna
Corso di laurea in Scienze dell'Informazione
Via Sacchi 3, 47023 Cesena (FC), Italy
E-mail: bravetti@cs.unibo.it

Leïla Kloul
Université de Versailles
Laboratoire PRiSM
45 Avenue des Etats-Unis, 78000 Versailles, France
E-mail: kle@prism.uvsq.fr

Gianluigi Zavattaro
Università di Bologna
Dipartimento di Scienze dell'Informazione
Mura A. Zamboni, 7, 40127 Bologna, Italy
E-mail: zavattar@cs.unibo.it

Library of Congress Control Number: 2005931522

CR Subject Classification (1998): D.2.4, C.2.4, F.3, D.4, C.4, K.4.4, C.2

ISSN 0302-9743
ISBN-10 3-540-28701-9 Springer Berlin Heidelberg New York
ISBN-13 978-3-540-28701-8 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
springeronline.com

© Springer-Verlag Berlin Heidelberg 2005
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 11549970 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

New York University, NY, 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

Lecture Notes in Computer Science

For information about Vols. 1–3564

please contact your bookseller or Springer

Vol. 3687: S. Singh, M. Singh, C. Apte, P. Perner (Eds.), Pattern Recognition and Image Analysis, Part II. XXV, 809 pages. 2005.

Vol. 3686: S. Singh, M. Singh, C. Apte, P. Perner (Eds.), Pattern Recognition and Data Mining, Part I. XXVI, 689 pages. 2005.

Vol. 3674: W. Jonker, M. Petković (Eds.), Secure Data Management. X, 241 pages. 2005.

Vol. 3672: C. Hankin, I. Siveróni (Eds.), Static Analysis. X, 369 pages. 2005.

Vol. 3671: S. Bressan, S. Ceri, E. Hunt, Z.G. Ives, Z. Belahsene, M. Rys, R. Unland (Eds.), Database and XML Technologies. X, 239 pages. 2005.

Vol. 3670: M. Bravetti, L. Kloul, G. Zavattaro (Eds.), Formal Techniques for Computer Systems and Business Processes. XIII, 349 pages. 2005.

Vol. 3664: C. Türker, M. Agosti, H.-J. Schek (Eds.), Peer-to-Peer, Grid, and Service-Oriented in Digital Library Architectures. X, 261 pages. 2005.

Vol. 3663: W. Kropatsch, R. Sablatnig, A. Hanbury (Eds.), Pattern Recognition. XIV, 512 pages. 2005.

Vol. 3662: C. Baral, G. Greco, N. Leone, G. Terracina (Eds.), Logic Programming and Nonmonotonic Reasoning. XIII, 454 pages. 2005. (Subseries LNAI).

Vol. 3660: M. Beigl, S. Intille, J. Rekimoto, H. Tokuda (Eds.), UbiComp 2005: Ubiquitous Computing. XVII, 394 pages. 2005.

Vol. 3659: J.R. Rao, B. Sunar (Eds.), Cryptographic Hardware and Embedded Systems – CHES 2005. XIV, 458 pages. 2005.

Vol. 3654: S. Sajodia, D. Wijesekera (Eds.), Data and Applications Security XIX. X, 353 pages. 2005.

Vol. 3653: M. Abadi, L.d. Alfaro (Eds.), CONCUR 2005 – Concurrency Theory. XIV, 578 pages. 2005.

Vol. 3649: W.M.P. van der Aalst, B. Benatallah, F. Casati, F. Curbera (Eds.), Business Process Management. XII, 472 pages. 2005.

Vol. 3648: J.C. Cunha, P.D. Medeiros (Eds.), Euro-Par 2005 Parallel Processing. XXXVI, 1299 pages. 2005.

Vol. 3645: D.-S. Huang, X.-P. Zhang, G.-B. Huang (Eds.), Advances in Intelligent Computing, Part II. XIII, 1010 pages. 2005.

Vol. 3644: D.-S. Huang, X.-P. Zhang, G.-B. Huang (Eds.), Advances in Intelligent Computing, Part I. XXVII, 1101 pages. 2005.

Vol. 3642: D. Ślezak, J. Yao, J.F. Peters, W. Ziarko, X. Hu (Eds.), Rough Sets, Fuzzy Sets, Data Mining, and Granular Computing, Part II. XXIV, 738 pages. 2005. (Subseries LNAI).

Vol. 3641: D. Ślezak, G. Wang, M.S. Szczuka, I. Düntsch, Y. Yao (Eds.), Rough Sets, Fuzzy Sets, Data Mining, and Granular Computing, Part I. XXIV, 742 pages. 2005. (Subseries LNAI).

Vol. 3639: P. Godefroid (Ed.), Model Checking Software. XI, 289 pages. 2005.

Vol. 3638: A. Butz, B. Fisher, A. Krüger, P. Olivier (Eds.), Smart Graphics. XI, 269 pages. 2005.

Vol. 3637: J. M. Moreno, J. Madrenas, J. Cosp (Eds.), Evolvable Systems: From Biology to Hardware. XI, 227 pages. 2005.

Vol. 3636: M.J. Blesa, C. Blum, A. Roli, M. Sampels (Eds.), Hybrid Metaheuristics. XII, 155 pages. 2005.

Vol. 3634: L. Ong (Ed.), Computer Science Logic. XI, 567 pages. 2005.

Vol. 3633: C. Bauzer Medeiros, M. Egenhofer, E. Bertino (Eds.), Advances in Spatial and Temporal Databases. XIII, 433 pages. 2005.

Vol. 3632: R. Nieuwenhuis (Ed.), Automated Deduction – CADE-20. XIII, 459 pages. 2005. (Subseries LNAI).

Vol. 3629: J.L. Fiadeiro, N. Harman, M. Roggenbach, J. Rutten (Eds.), Algebra and Coalgebra in Computer Science. XI, 457 pages. 2005.

Vol. 3628: T. Gschwind, U. Aßmann, O. Nierstrasz (Eds.), Software Composition. X, 199 pages. 2005.

Vol. 3627: C. Jacob, M.L. Pilat, P.J. Bentley, J. Timmis (Eds.), Artificial Immune Systems. XII, 500 pages. 2005.

Vol. 3626: B. Ganter, G. Stumme, R. Wille (Eds.), Formal Concept Analysis. X, 349 pages. 2005. (Subseries LNAI).

Vol. 3625: S. Kramer, B. Pfahringer (Eds.), Inductive Logic Programming. XIII, 427 pages. 2005. (Subseries LNAI).

Vol. 3624: C. Chekuri, K. Jansen, J.D.P. Rolim, L. Trevisan (Eds.), Approximation, Randomization and Combinatorial Optimization. XI, 495 pages. 2005.

Vol. 3623: M. Liśkiewicz, R. Reischuk (Eds.), Fundamentals of Computation Theory. XV, 576 pages. 2005.

Vol. 3621: V. Shoup (Ed.), Advances in Cryptology – CRYPTO 2005. XI, 568 pages. 2005.

Vol. 3620: H. Muñoz-Avila, F. Ricci (Eds.), Case-Based Reasoning Research and Development. XV, 654 pages. 2005. (Subseries LNAI).

Vol. 3619: X. Lu, W. Zhao (Eds.), Networking and Mobile Computing. XXIV, 1299 pages. 2005.

Vol. 3618: J. Jedrzejowicz, A. Szepietowski (Eds.), Mathematical Foundations of Computer Science 2005. XVI, 814 pages. 2005.

Vol. 3617: F. Roli, S. Vitulano (Eds.), Image Analysis and Processing – ICIAP 2005. XXIV, 1219 pages. 2005.

- Vol. 3615: B. Ludäscher, L. Raschid (Eds.), Data Integration in the Life Sciences. XII, 344 pages. 2005. (Subseries LNBI).
- Vol. 3614: L. Wang, Y. Jin (Eds.), Fuzzy Systems and Knowledge Discovery, Part II. XLI, 1314 pages. 2005. (Subseries LNAI).
- Vol. 3613: L. Wang, Y. Jin (Eds.), Fuzzy Systems and Knowledge Discovery, Part I. XLI, 1334 pages. 2005. (Subseries LNAI).
- Vol. 3612: L. Wang, K. Chen, Y. S. Ong (Eds.), Advances in Natural Computation, Part III. LXI, 1326 pages. 2005.
- Vol. 3611: L. Wang, K. Chen, Y. S. Ong (Eds.), Advances in Natural Computation, Part II. LXI, 1292 pages. 2005.
- Vol. 3610: L. Wang, K. Chen, Y. S. Ong (Eds.), Advances in Natural Computation, Part I. LXI, 1302 pages. 2005.
- Vol. 3608: F. Dehne, A. López-Ortiz, J.-R. Sack (Eds.), Algorithms and Data Structures. XIV, 446 pages. 2005.
- Vol. 3607: J.-D. Zucker, L. Saitta (Eds.), Abstraction, Reformulation and Approximation. XII, 376 pages. 2005. (Subseries LNAI).
- Vol. 3606: V. Malyshev (Ed.), Parallel Computing Technologies. XII, 470 pages. 2005.
- Vol. 3604: R. Martin, H. Bez, M. Sabin (Eds.), Mathematics of Surfaces XI. IX, 473 pages. 2005.
- Vol. 3603: J. Hurd, T. Melham (Eds.), Theorem Proving in Higher Order Logics. IX, 409 pages. 2005.
- Vol. 3602: R. Eigenmann, Z. Li, S.P. Midkiff (Eds.), Languages and Compilers for High Performance Computing. IX, 486 pages. 2005.
- Vol. 3599: U. Aßmann, M. Aksit, A. Rensink (Eds.), Model Driven Architecture. X, 235 pages. 2005.
- Vol. 3598: H. Murakami, H. Nakashima, H. Tokuda, M. Yasumura, Ubiquitous Computing Systems. XIII, 275 pages. 2005.
- Vol. 3597: S. Shimojo, S. Ichii, T.W. Ling, K.-H. Song (Eds.), Web and Communication Technologies and Internet-Related Social Issues - HSI 2005. XIX, 368 pages. 2005.
- Vol. 3596: F. Dau, M.-L. Mugnier, G. Stumme (Eds.), Conceptual Structures: Common Semantics for Sharing Knowledge. XI, 467 pages. 2005. (Subseries LNAI).
- Vol. 3595: L. Wang (Ed.), Computing and Combinatorics. XVI, 995 pages. 2005.
- Vol. 3594: J.C. Setubal, S. Verjovski-Almeida (Eds.), Advances in Bioinformatics and Computational Biology. XIV, 258 pages. 2005. (Subseries LNBI).
- Vol. 3593: V. Mafik, R. W. Brennan, M. Pěchouček (Eds.), Holonic and Multi-Agent Systems for Manufacturing. XI, 269 pages. 2005. (Subseries LNAI).
- Vol. 3592: S. Katsikas, J. Lopez, G. Pernul (Eds.), Trust, Privacy and Security in Digital Business. XII, 332 pages. 2005.
- Vol. 3591: M.A. Wimmer, R. Traummüller, Å. Grönlund, K.V. Andersen (Eds.), Electronic Government. XIII, 317 pages. 2005.
- Vol. 3590: K. Bauknecht, B. Pröll, H. Werthner (Eds.), E-Commerce and Web Technologies. XIV, 380 pages. 2005.
- Vol. 3589: A.M. Tjoa, J. Trujillo (Eds.), Data Warehousing and Knowledge Discovery. XVI, 538 pages. 2005.
- Vol. 3588: K.V. Andersen, J. Debenham, R. Wagner (Eds.), Database and Expert Systems Applications. XX, 955 pages. 2005.
- Vol. 3587: P. Perner, A. Imiya (Eds.), Machine Learning and Data Mining in Pattern Recognition. XVII, 695 pages. 2005. (Subseries LNAI).
- Vol. 3586: A.P. Black (Ed.), ECOOP 2005 - Object-Oriented Programming. XVII, 631 pages. 2005.
- Vol. 3584: X. Li, S. Wang, Z.Y. Dong (Eds.), Advanced Data Mining and Applications. XIX, 835 pages. 2005. (Subseries LNAI).
- Vol. 3583: R.W. H. Lau, Q. Li, R. Cheung, W. Liu (Eds.), Advances in Web-Based Learning - ICWL 2005. XIV, 420 pages. 2005.
- Vol. 3582: J. Fitzgerald, I.J. Hayes, A. Tarlecki (Eds.), FM 2005: Formal Methods. XIV, 558 pages. 2005.
- Vol. 3581: S. Miksch, J. Hunter, E. Keravnou (Eds.), Artificial Intelligence in Medicine. XVII, 547 pages. 2005. (Subseries LNAI).
- Vol. 3580: L. Caires, G.F. Italiano, L. Monteiro, C. Palamidessi, M. Yung (Eds.), Automata, Languages and Programming. XXV, 1477 pages. 2005.
- Vol. 3579: D. Lowe, M. Gaedke (Eds.), Web Engineering. XXII, 633 pages. 2005.
- Vol. 3578: M. Gallagher, J. Hogan, F. Maire (Eds.), Intelligent Data Engineering and Automated Learning - IDEAL 2005. XVI, 599 pages. 2005.
- Vol. 3577: R. Falcone, S. Barber, J. Sabater-Mir, M.P. Singh (Eds.), Trusting Agents for Trusting Electronic Societies. VIII, 235 pages. 2005. (Subseries LNAI).
- Vol. 3576: K. Etessami, S.K. Rajamani (Eds.), Computer Aided Verification. XV, 564 pages. 2005.
- Vol. 3575: S. Wermter, G. Palm, M. Elshaw (Eds.), Biomimetic Neural Learning for Intelligent Robots. IX, 383 pages. 2005. (Subseries LNAI).
- Vol. 3574: C. Boyd, J.M. González Nieto (Eds.), Information Security and Privacy. XIII, 586 pages. 2005.
- Vol. 3573: S. Etalle (Ed.), Logic Based Program Synthesis and Transformation. VIII, 279 pages. 2005.
- Vol. 3572: C. De Felice, A. Restivo (Eds.), Developments in Language Theory. XI, 409 pages. 2005.
- Vol. 3571: L. Godo (Ed.), Symbolic and Quantitative Approaches to Reasoning with Uncertainty. XVI, 1028 pages. 2005. (Subseries LNAI).
- Vol. 3570: A. S. Patrick, M. Yung (Eds.), Financial Cryptography and Data Security. XII, 376 pages. 2005.
- Vol. 3569: F. Bacchus, T. Walsh (Eds.), Theory and Applications of Satisfiability Testing. XII, 492 pages. 2005.
- Vol. 3568: W.-K. Leow, M.S. Lew, T.-S. Chua, W.-Y. Ma, L. Chaisorn, E.M. Bakker (Eds.), Image and Video Retrieval. XVII, 672 pages. 2005.
- Vol. 3567: M. Jackson, D. Nelson, S. Stirk (Eds.), Database: Enterprise, Skills and Innovation. XII, 185 pages. 2005.
- Vol. 3566: J.-P. Banâtre, P. Fradet, J.-L. Giavitto, O. Michel (Eds.), Unconventional Programming Paradigms. XI, 367 pages. 2005.
- Vol. 3565: G.E. Christensen, M. Sonka (Eds.), Information Processing in Medical Imaging. XXI, 777 pages. 2005.

¥490.88元

Preface

This volume contains the proceedings of two international workshops EPEW and WS-FM held at the Université de Versailles Saint-Quentin-en-Yvelines, Versailles, France, 1–3 September 2005.

EPEW (European Performance Engineering Workshop) and WS-FM (International Workshop on Web Services and Formal Methods) were colocated to gather the researchers working across the spectrum of techniques for modelling, specification, analysis and verification of the behavior of computer systems and business processes.

This proceedings contains a selection of 20 research contributions, out of 59 submissions, which went through a rigorous review process by international reviewers. We therefore owe special thanks to all members of both program committees of EPEW and WS-FM and their sub-referees for the excellent work they did in the short time they had.

Additionally, this proceedings includes four invited papers, by Gianfranco Ciardo (University of California at Riverside), Peter G. Harrison (Imperial College London), Cosimo Laneve (University of Bologna) and Wil van der Aalst (Eindhoven University of Technology). These contributions brought an additional dimension to the technical and the scientific merit of these workshops.

Finally, our thanks go to the University of Versailles Saint-Quentin-en-Yvelines, its Laboratoire PRiSM and the CNRS for hosting the workshops and providing technical and financial support.

September 2005

Mario Bravetti
Leïla Kloul
Gianluigi Zavattaro

EPEW

The European Performance Engineering Workshop aims to bring together researchers working on all aspects of performance modelling and analysis of computer and telecommunication systems. Of the 32 regular papers submitted to EPEW 2005, after a rigorous review process, only 10 were accepted for presentation. We are very pleased with the quality of these papers and hope that you will find them interesting. The topics of this workshop are various and include queueing theory, bounding techniques, stochastic model checking, communication schemes analysis for high-speed LAN, QoS analysis in wireless ad hoc networks and optical networks analysis.

Workshop Chair

Leïla Kloul

Université de Versailles, France

Program Committee

Jeremy Bradley	Imperial College London (UK)
Mario Bravetti	University of Bologna (Italy)
Tadeusz Czachórski	IITiS PAN, Gliwice (Poland)
Jean-Michel Fourneau	Université de Versailles (France)
Stephen Gilmore	University of Edinburgh (UK)
Holger Hermanns	University of the Saarland (Germany)
Alain Jean-Marie	Université de Montpellier and INRIA (France)
Helen Karatza	Aristotle University of Thessaloniki (Greece)
Leïla Kloul	Université de Versailles (France)
Kim G. Larsen	University of Aalborg (Denmark)
Fernando López Pelayo	Universidad Castilla-La Mancha (Spain)
Raymond Marie	Université de Rennes and IRISA (France)
Andrew Miner	Iowa State University (USA)
Manuel Núñez	Universidad Complutense de Madrid (Spain)
Brigitte Plateau	ID-IMAG, Grenoble (France)
Ramon Puigjaner	Universidad Illes Balears (Spain)
Marina Ribaudó	University of Genoa (Italy)
Mark Squillante	IBM T.J. Watson Research Center, NY (USA)
Nigel Thomas	University of Newcastle (UK)
Fabrice Valois	CITI, INSA Lyon (France)
Katinka Wolter	Humboldt-Universität zu Berlin (Germany)
Wlodek M. Zuberek	Memorial University (Canada)

Additional Referees

D. Barth	P. Harrison	N. Mezzetti
D. Cazorla	S. Johr	N. Pekergin
G. Chiola	J.P. Katoen	R. Pulungan
F. Cuartero	A. Langville	I. Rodríguez
J.L. Ferrer-Gomila	D. Lime	P.P. Sancho
S. Galmés	N. López	I. Sbeity
C. Guerrero Tomé	G. Massone	W.J. Stewart
U. Harder	M.G. Merayo	

Local Arrangements Committee

A. Baffert	L. Kloul	A. Mokhtari
C. Ducoin	C. Lequere	

Sponsoring Institutions

CNRS
Laboratoire PRiSM
Université de Versailles Saint-Quentin-en-Yvelines

WS-FM

The International Workshop on Web Services and Formal Methods aims to bring together researchers working on Web Services and formal methods in order to activate a fruitful collaboration in this direction of research. This, potentially, could also have a great impact on the current standardization phase of Web Services technologies. This second edition of the workshop (WS-FM 2005) featured 10 papers selected among 27 submissions after a rigorous review process. The main topics of the conference include: protocols and standards for WS (SOAP, WSDL, UDDI, etc.); languages and description methodologies for choreography/orchestration/workflow (BPML, XLANG and BizTalk, WSFL, WS-BPEL, etc.); coordination techniques for WS (transactions, agreement, coordination services, etc.); semantics-based dynamic WS discovery services (based on Semantic Web/Ontology techniques or other semantic theories); security, performance evaluation and quality of service of WS; semi-structured data and XML related technologies; and comparisons with different related technologies/approaches.

Workshop Co-chairs

Mario Bravetti
Gianluigi Zavattaro

University of Bologna (Italy)
University of Bologna (Italy)

Program Committee

Marco Aiello
Jean-Pierre Banatre
Boualem Benatallah
Karthik Bhargavan
Manfred Broy
Roberto Bruni
Michael Butler
Fabio Casati
Rocco De Nicola
Schahram Dustdar
Gianluigi Ferrari
Jose Luiz Fiadeiro
Peter Furniss
Stephanie Gnesi
Reiko Heckel
Nickolas Kavantzaz

University of Trento (Italy)
University of Rennes 1 and INRIA (France)
University of New South Wales (Australia)
Microsoft Research, Cambridge (UK)
Technische Universität München (Germany)
University of Pisa (Italy)
University of Southampton (UK)
HP Labs (USA)
University of Florence (Italy)
Wien University of Technology (Austria)
University of Pisa (Italy)
University of Leicester (UK)
Choreology Ltd. (UK)
CNR Pisa (Italy)
University of Leicester (UK)
Oracle Co. (USA)

Leïla Kloul	Université de Versailles (France)
Mark Little	Arjuna Technologies Ltd. (UK)
Natalia López	University Complutense of Madrid (Spain)
Roberto Lucchi	University of Bologna (Italy)
Jeff Magee	Imperial College London (UK)
Fabio Martinelli	CNR Pisa (Italy)
Shin Nakajima	National Institute of Informatics and JST (Japan)
Manuel Nunez	University Complutense of Madrid (Spain)
Fernando Pelayo	University of Castilla-La Mancha, Albacete (Spain)
Marco Pistore	University of Trento (Italy)
Wolfgang Reisig	Humboldt University, Berlin (Germany)
Vladimiro Sassone	University of Sussex (UK)
Frank Van Breugel	York University, Toronto (Canada)
Friedrich Vogt	Technical University of Hamburg-Harburg (Germany)

Additional Referees

M. ter Beek	A. Lazovik	S. Ripon
N. Diernhofer	P. Massuthe	B.A. Schmit
G. Frankova	M. Mazzara	E. Tuosto
C. Guidi	H. Melgratti	G. Wimmel
W. Knottenbelt	M.G. Merayo	
J.C. Laclaustra	T. Priol	

Organizing Committee

M. Bravetti	R. Lucchi
L. Kloul	G. Zavattaro

Sponsoring Institutions

CNRS
Laboratoire PRiSM
Université de Versailles Saint-Quentin-en-Yvelines

Table of Contents

Invited Speakers

Performance Engineering and Stochastic Modelling <i>Peter Harrison</i>	1
Implicit Representations and Algorithms for the Logic and Stochastic Analysis of Discrete-State Systems <i>Gianfranco Ciardo</i>	15
PiDuce: A Process Calculus with Native XML Datatypes <i>Allen L. Brown Jr., Cosimo Laneve, L. Gregory Meredith</i>	18
Life After BPEL? <i>W.M.P. van der Aalst, M. Dumas, A.H.M. ter Hofstede, N. Russell, H.M.W. Verbeek, P. Wohed</i>	35

EPEW

On Moments of Discrete Phase-Type Distributions <i>Tuğrul Dayar</i>	51
Zero-Automatic Queues <i>Thu-Hà Dao-Thi, Jean Mairesse</i>	64
A Unified Approach to the Moments Based Distribution Estimation – Unbounded Support <i>Árpád Tari, Miklós Telek, Peter Buchholz</i>	79
Bounds for Point and Steady-State Availability: An Algorithmic Approach Based on Lumpability and Stochastic Ordering <i>Ana Bušić, Jean-Michel Fourneau</i>	94
Stochastic Model Checking with Stochastic Comparison <i>Nihal Pekergin, Sana Younès</i>	109
Delay Analysis of the Go-Back-N ARQ Protocol over a Time-Varying Channel <i>Koen De Turck, Sabine Wittevrongel</i>	124

Performance Tuning of Failure Detectors in Wireless Ad-hoc Networks: Modelling and Experiments <i>Corine Marchand, Jean-Marc Vincent</i>	139
Hypergraph Partitioning for Faster Parallel PageRank Computation <i>Jeremy T. Bradley, Douglas V. de Jager, William J. Knottenbelt, Aleksandar Trifunović</i>	155
Prediction of Communication Latency over Complex Network Behaviors on SMP Clusters <i>Maxime Martinasso, Jean-François Méhaut</i>	172
A Diffusion Approximation Model of an Electronic-Optical Node <i>Tadeusz Czachórski, Ferhan Pekergin</i>	187
WS-FM	
Choreographing Security and Performance Analysis for Web Services <i>Stephen Gilmore, Valentin Haenel, Leïla Kloul, Monika Maidl</i>	200
Application of Formal Methods to the Analysis of Web Services Security <i>Llanos Tobarra, Diego Cazorla, Fernando Cuartero, Gregorio Díaz</i> ...	215
Automatic Translation of WS-CDL Choreographies to Timed Automata <i>Gregorio Díaz, Juan-José Pardo, María-Emilia Cambronero, Valentín Valero, Fernando Cuartero</i>	230
Executable Semantics for Compensating CSP <i>Michael Butler, Shamim Ripon</i>	243
Verifying the Conformance of Web Services to Global Interaction Protocols: A First Step <i>M. Baldoni, C. Baroglio, A. Martelli, V. Patti, C. Schifanella</i>	257
From Theory to Practice in Transactional Composition of Web Services <i>Roberto Bruni, Gianluigi Ferrari, Hernán Melgratti, Ugo Montanari, Daniele Stollo, Emilio Tuosto</i>	272
Timing Issues in Web Services Composition <i>Manuel Mazzara</i>	287
A Compositional Operational Semantics for OWL-S <i>Barry Norton, Simon Foster, Andrew Hughes</i>	303

A Parametric Communication Model for the Verification of BPEL4WS
Compositions

Raman Kazhamiakin, Marco Pistore 318

Reasoning About Interaction Patterns in Choreography

Roberto Gorrieri, Claudio Guidi, Roberto Lucchi 333

Author Index 349

Performance Engineering and Stochastic Modelling

Peter Harrison

Imperial College London, South Kensington Campus, London SW7 2AZ, UK
pgh@doc.ic.ac.uk

Abstract. Performance engineering has become a central plank in the design of complex, time-critical systems. It is supported by stochastic modelling, a brief history of which is given, going back to Erlang as long ago as 1909. This in turn developed according to successive new generations of communication and computer architectures and other operational systems. Its evolution through queues and networks is reviewed, culminating in the unification of many specification and solution techniques in a common formalism, stochastic process algebra. Recent results are given on the automatic computation of separable solutions for the equilibrium state probabilities in systems specified in such a formalism. A performance engineering support environment is proposed to integrate these methods with others such as response time analysis and fluid models, which are better suited to large scale aggregation of similar components in a continuous space.

1 Introduction

Performance engineering is about the development of efficient computer and communication systems by providing crucial performance analysis during the design phase. For several decades, computer and communication system design was left to ‘experienced engineers’ who used their knowledge and experience to successfully construct and ‘optimise’ both hardware and software architectures. These systems grew from simple stand-alone computers, through mainframes with virtual storage management and multi-access, to client-server systems organised as networks of clusters. More recently still we have distributed internet applications and the Grid. Increasingly from the late nineteen eighties/early nineties, the traditional ad hoc approach has been found inadequate in view of the complexity and sheer size of modern systems. This is hardly surprising in view of other engineering disciplines where the construction of an artefact is invariably preceded by the construction of a *model*, usually a mathematical model, upon which design decisions and parameterisations are based. In fact, the moves to ‘structured programming’, beginning in the 1970s, specialised programming languages and ‘formal methods’ defined the basis of software engineering which mirrors the conventional engineering paradigm.

The tools of performance engineering are *quantitative models* that predict metrics which reflect a view of performance. This view depends on who’s interested. For the user it is some form of quality of service (QoS), often response

time, reliability or availability. For the owner of a business that sells Grid resources, it is revenue, based more on the number of customers processed per day. Whatever the metric, the whole system needs to be modelled at an appropriate level of detail. There are several modelling techniques, each playing a vital role in different circumstances. An evolving, dynamic system can be considered as passing through a sequence of states that change upon the occurrence of events. Thus, discrete event simulation is the most natural and general modelling formalism since it is possible to mimic in an abstract model the dynamic behaviour of almost any system. The main issues in simulation are reliable interpretation of the simulation outputs, which are essentially sample paths of an underlying stochastic process, and run-time, which may be prohibitively long in detailed simulations and increase rapidly as tighter confidence bands are imposed.

Hence, where possible, analytical models are preferred, based on a mathematical analysis of the stochastic process that precisely describes the dynamics of the system being modelled. The trouble now is that further approximating assumptions must be made to achieve mathematical and numerical tractability. As a rule, the more tractable the model needs to be, the more approximations are necessary. Hence, when using approximate analytical models, validation must be done, often initially mutual validation between the analytical and simulation models. It is typical for an analytical model to be validated against simulation in simple scenarios, resulting in confidence in both models. Then either model can be used to make predictions in more complex cases. As complexity increases, one of the models will ultimately cease to be numerically tractable, but after validation, confidence is high in the surviving model. Indeed, it may be that it is best to validate a simulation in simple cases against an *exact* analytical model and then use that simulation to validate an approximate analytical model, which is the one intended for practical use.

Although analytical models produce precise predictions with respect to their assumptions and approximations, it is not always easy to interpret their output due to subtle properties like interdependence between, and even existence of, model components. Beware of such subtleties! For example, consider the following gamble. Two boxes are placed in front of you with a piece of gold in each. You can keep the gold from either box (your choice) and you have been told there is exactly double the amount of gold in one box than in the other, but you don't know which is which. You are allowed to open one box, look in and then decide which box's contents you want to keep. What should you do? If you find an amount x of gold in the box you open, you know there is either $2x$ or $x/2$ in the other, with equal probability by the symmetry of your random choice of box. So the average amount in the other box is

$$\frac{2x + \frac{x}{2}}{2} = 1.25x$$

Clearly, then, the best strategy is to choose a box at random, open it, and take the gold from the other box! What's wrong with that?

This paper is organised as follows. In section 2, a brief history is given of stochastic modelling, with a somewhat personal view. This begins with the early

work of A.K. Erlang and progresses through queues and queueing networks of various kinds up to the present day, not forgetting the underpinnings of A.A. Markov. In section 3, relatively recent advances in the field of stochastic process algebra (SPA) are reviewed and it is shown how these combine many previous, apparently diverse, Markov models into a unified framework. The use of SPA is then argued as a basis for future performance engineering support environments, ideally to be integrated into software engineering support environments. The paper closes in section 6.

2 A Brief History

The origins of stochastic modelling, as applied to communication systems and operations research, is widely credited to A.K. Erlang, who was the first person to consider the problem of telephone networks quantitatively. By studying a village telephone exchange, he worked out a formula, now known as Erlang's formula, to calculate the fraction of callers attempting to call someone outside the village that must wait because all of the lines are in use. Although Erlang's model is a simple one, the mathematics underlying today's complex communication networks is still based on it in many ways. Erlang continued to investigate these problems and published perhaps his best known work "Solution of some Problems in the Theory of Probabilities of Significance in Automatic Telephone Exchanges" in 1917. Meanwhile, the theory underlying many of today's models, including those considered here, was being developed by A.A. Markov, the well known Russian mathematician, who began publishing in the area shortly after the turn of the 20th century; an interesting review of Markov's life and contributions is [1].

More recently, the focus of attention became Markovian networks of queues, with the famous product-form solutions of Jackson (open and closed networks) in 1963 and Gordon & Newell (closed networks) in 1967 [23,15]. These were generalised to multiple classes of customers, with class-dependent service and routing, and queueing disciplines other than first come first served by Baskett et. al. in 1974 and by Kelly at about the same time. Kelly used a different approach based on so-called *quasi-reversibility*, which is defined in terms of the reversed process of a stationary continuous time Markov chain (CTMC) and actually leads to further generalisations, for example to symmetric queues [24]. Interestingly, the reversed process is key to the unified methodology discussed later in this paper.

Various minor extensions were produced in the 1980s, more sophisticated individual queues were analysed and the focus shifted somewhat to approximate methods that required much less restrictive assumptions and so applied to more realistic scenarios. However, at the end of the 1980s, Gelenbe introduced the concept of the *negative customer*, which removes a customer (said to 'kill' one) from a queue rather than adds to it. This model was motivated by inhibitor signals in biological neural networks – these appear discrete in nature rather than continuously varying potentials. Surprisingly, the resulting Markovian network