

Luca de Alfaro
Stephen Gilmore (Eds.)

LNCS 2165

Process Algebra and Probabilistic Methods

Performance Modelling and Verification

**Joint International Workshop, PAPM-PROBMIV 2001
Aachen, Germany, September 2001
Proceedings**



Springer

Luca de Alfaro Stephen Gilmore (Eds.)

Process Algebra and Probabilistic Methods

Performance Modelling and Verification

Joint International Workshop, PAPM-PROBMIV 2001
Aachen, Germany, September 12-14, 2001
Proceedings



Springer

Series Editors

Gerhard Goos, Karlsruhe University, Germany
Juris Hartmanis, Cornell University, NY, USA
Jan van Leeuwen, Utrecht University, The Netherlands

Volume Editors

Luca de Alfaro
University of California at Berkeley, Dept. of Electr. Eng. and Computer Science
479 Cory Hall, Berkeley, CA 94720-1770, USA
E-mail: dealfaro@eecs.berkeley.edu

Stephen Gilmore
The University of Edinburgh, Laboratory for Foundations of Computer Science
Edinburgh EH9 3JZ, UK
E-mail: stg@dcs.ed.ac.uk

Cataloging-in-Publication Data applied for

Die Deutsche Bibliothek - CIP-Einheitsaufnahme

Process algebra and probabilistic methods : performance modeling and verification ; joint international workshop ; proceedings / PAPM PROBMIV 2001, Aachen, Germany, September 12 - 14, 2001. Luca de Alfaro ; Stephen Gilmore (ed.). - Berlin ; Heidelberg ; New York ; Barcelona ; Hong Kong ; London ; Milan ; Paris ; Tokyo : Springer, 2001
(Lecture notes in computer science ; Vol. 2165)
ISBN 3-540-42556-X

CR Subject Classification (1998): F.3.1, F.3, D.2.4, D.3.1, C.4

ISSN 0302-9743

ISBN 3-540-42556-X Springer-Verlag 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-Verlag. Violations are liable for prosecution under the German Copyright Law.

Springer-Verlag Berlin Heidelberg New York
a member of BertelsmannSpringer Science+Business Media GmbH

<http://www.springer.de>

© Springer-Verlag Berlin Heidelberg 2001
Printed in Germany

Typesetting: Camera-ready by author, data conversion by PTP-Berlin, Stefan Sossna
Printed on acid-free paper SPIN: 10840347 06/3142 5 4 3 2 1 0

Preface

This volume contains the proceedings of the first joint PAPM-PROBMIV Workshop, held at the Rheinisch-Westfälische Technische Hochschule (RWTH) Aachen, Germany, 12–14 September 2001.

The PAPM-PROBMIV workshop results from the combination of two workshops: PAPM (Process Algebras and Performance Modeling) and PROBMIV (Probabilistic Methods in Verification). The aim of the joint workshop is to bring together the researchers working across the whole spectrum of techniques for the modeling, specification, analysis, and verification of probabilistic systems. Probability is widely used in the design and analysis of software and hardware systems, as a means to derive efficient algorithms (e.g. randomization), as a model for unreliable or unpredictable behavior (as in the study of fault-tolerant systems and computer networks), and as a tool to study performance and dependability properties. The topics of the workshop include specification, models and semantics of probabilistic systems, analysis and verification techniques, probabilistic methods for the verification of non-probabilistic systems, and tools and case studies.

The first PAPM workshop was held in Edinburgh in 1993; the following ones were held in Regensburg (1994), Edinburgh (1995), Torino (1996), Enschede (1997), Nice (1998), Zaragoza (1999), and Geneva (2000). The first PROBMIV workshop was held in Indianapolis, Indiana (1998); the next one took place in Eindhoven (1999). In 2000, PROBMIV was replaced by a Dagstuhl seminar on Probabilistic Methods in Verification.

The PAPM-PROBMIV workshop is held in conjunction with two other workshops: 11th GI/ITG Conference on Measuring, Modeling, and Evaluation of Computer and Communications Systems (MMB), and the 9th International Workshop on Petri Nets and Performance Models (PNPM). Together, these three workshops form the *2001 Aachen Multiconference on Measurement, Modeling, and Evaluation of Computer-Communication Systems*. We hope that this setting fosters the exchange of ideas with neighboring research fields and allows for a comparison of different viewpoints towards similar problems.

Of the 23 regular papers, 12 were accepted for presentation at the workshop and are included in the present volume. The workshop is preceded by three tutorials, given by Joost-Pieter Katoen (University of Twente) on *Probabilistic verification of Markov chains*, by Marina Ribaldo (University of Torino) on *An introduction to stochastic process algebras*, and by Roberto Segala (University of Bologna) on *Nondeterminism in probabilistic verification*. The workshop includes three invited presentations, by Shankar Sastry (University of California, Berkeley), Markus Siegle (Friedrich-Alexander Universität Erlangen-Nürnberg), and Frits Vaandrager (University of Nijmegen).

We thank all the members of the program committee, and their sub-referees, for selecting the papers to be presented. Special thanks are due to Boudewijn

Haverkort (University of Aachen), the general chair of the multi-conference and local organization, and Peter Kemper (University of Dortmund), the tool session chair. Our thanks go to the following organizations for their generous sponsorship of the Aachen multiconference: German Research Association (DFG), IBM Deutschland, Siemens AG München (Information and Communication Networks), T-Nova Deutsche Telekom Innovationsgesellschaft mbH, and TENOVIS. Our thanks also go to all the authors for meeting the tight deadlines which we set without compromising on the rigor or clarity of their papers.

July 2001

Luca de Alfaro
Stephen Gilmore

Scientific Organization

PAPM-PROBMIV 2001 Program Committee

Luca de Alfaro (Co-chair, University of California, Berkeley)
Stephen Gilmore (Co-chair, University of Edinburgh)
Gianfranco Ciardo (College of William and Mary)
Rance Cleaveland (SUNY Stony Brook)
Roberto Gorrieri (University of Bologna)
Boudewijn Haverkort (RWTH Aachen)
Ulrich Herzog (University of Erlangen)
Michael Huth (Kansas State University)
Radha Jagadeesan (Loyola University)
Joost-Pieter Katoen (University of Twente)
Leila Kloul (University of Versailles)
Kim Larsen (University of Aalborg)
Corrado Priami (University of Verona)
Marina Ribaudo (University of Turin)
Roberto Segala (University of Bologna)
Moshe Vardi (Rice University)

General Chair

Boudewijn Haverkort (University of Aachen)

Tool Session Chair

Peter Kemper (University of Dortmund)

PROBMIV Steering Committee

Marta Kwiatkowska (Chair, University of Birmingham)
Luca de Alfaro (University of California, Berkeley)
Rajeev Alur (University of Pennsylvania)
Christel Baier (University of Mannheim)
Michael Huth (Kansas State University)
Joost-Pieter Katoen (University of Twente)
Prakash Panangaden (McGill University)
Roberto Segala (University of Bologna)

PAPM Steering Committee

Ed Brinksmas (University of Twente)
Roberto Gorrieri (University of Bologna)
Ulrich Herzog (University of Erlangen)
Jane Hillston (University of Edinburgh)

Referees

Alessandro Aldini	Purush Iyer
Luca de Alfaro	Paul Jackson
Suzana Andova	Radha Jagadeesan
Marco Bernardo	Joost-Pieter Katoen
Henrik Bohnenkamp	Peter Kemper
Jeremy Bradley	Uli Klehmet
Mario Bravetti	Leïla Kloul
Peter Buchholz	Rom Langerak
Gianfranco Ciardo	Kim Larsen
Graham Clark	Joachim Meyer-Kayser
Rance Cleaveland	Gethin Norman
Lucia Cloth	Prakash Panangaden
Pedro d'Argenio	Dave Parker
Jean-Michel Fourneau	Corrado Priami
Stephen Gilmore	Marina Ribaudo
Roberto Gorrieri	Roberto Segala
Marco Gribaudo	Riccardo Sisto
Vineet Gupta	Gene Stark
Boudewijn Haverkort	Nigel Thomas
Holger Hermanns	Joanna Tomasik
Ulrich Herzog	Moshe Vardi
Andras Hovarth	Katinka Wolter
Michael Huth	Wlodek Zuberek

Organization

Sponsors

German Research Association (DFG)

IBM Deutschland

Siemens AG München, Information and Communication Networks

T-Nova Deutsche Telekom Innovationsgesellschaft mbH

TENOVIS

Local Arrangements Committee

Alexander Bell

Henrik Bohnenkamp

Rachid El Abdouni Kayari

Lucia Cloth

Boudewijn Haverkort

Ramin Sadre

Finance Chair

Volker Schanz (ITG, Frankfurt/Main)

Lecture Notes in Computer Science

For information about Vols. 1–2084
please contact your bookseller or Springer-Verlag

- Vol. 1982: S. Näher, D. Wagner (Eds.), *Algorithm Engineering. Proceedings, 2000. VIII*, 243 pages. 2001.
- Vol. 1996: P.L. Lanzi, W. Stolzmann, S.W. Wilson (Eds.), *Advances in Learning Classifier Systems. Proceedings, 2000. VIII*, 273 pages. 2001. (Subseries LNAI).
- Vol. 2085: J. Mira, A. Prieto (Eds.), *Bio-Inspired Applications of Connectionism. Proceedings, 2001. Part II. XXVII*, 848 pages. 2001.
- Vol. 2086: M. Luck, V. Mařík, O. Štěpánková, R. Trappl (Eds.), *Multi-Agent Systems and Applications. Proceedings, 2001. X*, 437 pages. 2001. (Subseries LNAI).
- Vol. 2087: G. Kern-Isberner, *Conditionals in Non-monotonic Reasoning and Belief Revision. X*, 190 pages. 2001. (Subseries LNAI).
- Vol. 2088: S. Yu, A. Păun (Eds.), *Implementation and Application of Automata. Proceedings, 2000. XI*, 343 pages. 2001.
- Vol. 2089: A. Amir, G.M. Landau (Eds.), *Combinatorial Pattern Matching. Proceedings, 2001. VIII*, 273 pages. 2001.
- Vol. 2090: E. Brinksma, H. Hermanns, J.-P. Katoen (Eds.), *Lectures on Formal Methods and Performance Analysis. Proceedings, 2000. VII*, 431 pages. 2001.
- Vol. 2091: J. Bigun, F. Smeraldi (Eds.), *Audio- and Video-Based Biometric Person Authentication. Proceedings, 2001. XIII*, 374 pages. 2001.
- Vol. 2092: L. Wolf, D. Hutchison, R. Steinmetz (Eds.), *Quality of Service – IWQoS 2001. Proceedings, 2001. XII*, 435 pages. 2001.
- Vol. 2093: P. Lorenz (Ed.), *Networking – ICN 2001. Proceedings, 2001. Part I. XXV*, 843 pages. 2001.
- Vol. 2094: P. Lorenz (Ed.), *Networking – ICN 2001. Proceedings, 2001. Part II. XXV*, 899 pages. 2001.
- Vol. 2095: B. Schiele, G. Sagerer (Eds.), *Computer Vision Systems. Proceedings, 2001. X*, 313 pages. 2001.
- Vol. 2096: J. Kittler, F. Roli (Eds.), *Multiple Classifier Systems. Proceedings, 2001. XII*, 456 pages. 2001.
- Vol. 2097: B. Read (Ed.), *Advances in Databases. Proceedings, 2001. X*, 219 pages. 2001.
- Vol. 2098: J. Akiyama, M. Kano, M. Urabe (Eds.), *Discrete and Computational Geometry. Proceedings, 2000. XI*, 381 pages. 2001.
- Vol. 2099: P. de Groote, G. Morrill, C. Retoré (Eds.), *Logical Aspects of Computational Linguistics. Proceedings, 2001. VIII*, 311 pages. 2001. (Subseries LNAI).
- Vol. 2100: R. Küsters, *Non-Standard Inferences in Description Logics. X*, 250 pages. 2001. (Subseries LNAI).
- Vol. 2101: S. Quaglini, P. Barahona, S. Andreassen (Eds.), *Artificial Intelligence in Medicine. Proceedings, 2001. XIV*, 469 pages. 2001. (Subseries LNAI).
- Vol. 2102: G. Berry, H. Comon, A. Finkel (Eds.), *Computer-Aided Verification. Proceedings, 2001. XIII*, 520 pages. 2001.
- Vol. 2103: M. Hannebauer, J. Wendler, E. Pagello (Eds.), *Balancing Reactivity and Social Deliberation in Multi-Agent Systems. VIII*, 237 pages. 2001. (Subseries LNAI).
- Vol. 2104: R. Eigenmann, M.J. Voss (Eds.), *OpenMP Shared Memory Parallel Programming. Proceedings, 2001. X*, 185 pages. 2001.
- Vol. 2105: W. Kim, T.-W. Ling, Y.-J. Lee, S.-S. Park (Eds.), *The Human Society and the Internet. Proceedings, 2001. XVI*, 470 pages. 2001.
- Vol. 2106: M. Kerckhove (Ed.), *Scale-Space and Morphology in Computer Vision. Proceedings, 2001. XI*, 435 pages. 2001.
- Vol. 2107: F.T. Chong, C. Kozyrakis, M. Oskin (Eds.), *Intelligent Memory Systems. Proceedings, 2000. VIII*, 193 pages. 2001.
- Vol. 2108: J. Wang (Ed.), *Computing and Combinatorics. Proceedings, 2001. XIII*, 602 pages. 2001.
- Vol. 2109: M. Bauer, P.J. Gymtrasiewicz, J. Vassileva (Eds.), *User Modelind 2001. Proceedings, 2001. XIII*, 318 pages. 2001. (Subseries LNAI).
- Vol. 2110: B. Hertzberger, A. Hoekstra, R. Williams (Eds.), *High-Performance Computing and Networking. Proceedings, 2001. XVII*, 733 pages. 2001.
- Vol. 2111: D. Helmbold, B. Williamson (Eds.), *Computational Learning Theory. Proceedings, 2001. IX*, 631 pages. 2001. (Subseries LNAI).
- Vol. 2116: V. Akman, P. Bouquet, R. Thomason, R.A. Young (Eds.), *Modeling and Using Context. Proceedings, 2001. XII*, 472 pages. 2001. (Subseries LNAI).
- Vol. 2117: M. Beynon, C.L. Nehaniv, K. Dautenhahn (Eds.), *Cognitive Technology: Instruments of Mind. Proceedings, 2001. XV*, 522 pages. 2001. (Subseries LNAI).
- Vol. 2118: X.S. Wang, G. Yu, H. Lu (Eds.), *Advances in Web-Age Information Management. Proceedings, 2001. XV*, 418 pages. 2001.
- Vol. 2119: V. Varadharajan, Y. Mu (Eds.), *Information Security and Privacy. Proceedings, 2001. XI*, 522 pages. 2001.
- Vol. 2120: H.S. Delugach, G. Stumme (Eds.), *Conceptual Structures: Broadening the Base. Proceedings, 2001. X*, 377 pages. 2001. (Subseries LNAI).
- Vol. 2121: C.S. Jensen, M. Schneider, B. Seeger, V.J. Tsotras (Eds.), *Advances in Spatial and Temporal Databases. Proceedings, 2001. XI*, 543 pages. 2001.
- Vol. 2123: P. Perner (Ed.), *Machine Learning and Data Mining in Pattern Recognition. Proceedings, 2001. XI*, 363 pages. 2001. (Subseries LNAI).

- Vol. 2124: W. Skarbek (Ed.), Computer Analysis of Images and Patterns. Proceedings, 2001. XV, 743 pages. 2001.
- Vol. 2125: F. Dehne, J.-R. Sack, R. Tamassia (Eds.), Algorithms and Data Structures. Proceedings, 2001. XII, 484 pages. 2001.
- Vol. 2126: P. Cousot (Ed.), Static Analysis. Proceedings, 2001. XI, 439 pages. 2001.
- Vol. 2127: V. Malyshev (Ed.), Parallel Computing Technologies. Proceedings, 2001. XII, 516 pages. 2001.
- Vol. 2129: M. Goemans, K. Jansen, J.D.P. Rolim, L. Trevisan (Eds.), Approximation, Randomization, and Combinatorial Optimization. Proceedings, 2001. IX, 297 pages. 2001.
- Vol. 2130: G. Dorffner, H. Bischof, K. Hornik (Eds.), Artificial Neural Networks – ICANN 2001. Proceedings, 2001. XXII, 1259 pages. 2001.
- Vol. 2132: S.-T. Yuan, M. Yokoo (Eds.), Intelligent Agents. Specification. Modeling, and Application. Proceedings, 2001. X, 237 pages. 2001. (Subseries LNAI).
- Vol. 2136: J. Sgall, A. Pultr, P. Kolman (Eds.), Mathematical Foundations of Computer Science 2001. Proceedings, 2001. XII, 716 pages. 2001.
- Vol. 2138: R. Freivalds (Ed.), Fundamentals of Computation Theory. Proceedings, 2001. XIII, 542 pages. 2001.
- Vol. 2139: J. Kilian (Ed.), Advances in Cryptology – CRYPTO 2001. Proceedings, 2001. XI, 599 pages. 2001.
- Vol. 2141: G.S. Brodal, D. Frigioni, A. Marchetti-Spaccamela (Eds.), Algorithm Engineering. Proceedings, 2001. X, 199 pages. 2001.
- Vol. 2142: L. Fribourg (Ed.), Computer Science Logic. Proceedings, 2001. XII, 615 pages. 2001.
- Vol. 2143: S. Benferhat, P. Besnard (Eds.), Symbolic and Quantitative Approaches to Reasoning with Uncertainty. Proceedings, 2001. XIV, 818 pages. 2001. (Subseries LNAI).
- Vol. 2146: J.H. Silverman (Eds.), Cryptography and Lattices. Proceedings, 2001. VII, 219 pages. 2001.
- Vol. 2147: G. Brebner, R. Woods (Eds.), Field-Programmable Logic and Applications. Proceedings, 2001. XV, 665 pages. 2001.
- Vol. 2149: O. Gascuel, B.M.E. Moret (Eds.), Algorithms in Bioinformatics. Proceedings, 2001. X, 307 pages. 2001.
- Vol. 2150: R. Sakellariou, J. Keane, J. Gurd, L. Freeman (Eds.), Euro-Par 2001 Parallel Processing. Proceedings, 2001. XXX, 943 pages. 2001.
- Vol. 2151: A. Caplinskas, J. Eder (Eds.), Advances in Databases and Information Systems. Proceedings, 2001. XIII, 381 pages. 2001.
- Vol. 2152: R.J. Boulton, P.B. Jackson (Eds.), Theorem Proving in Higher Order Logics. Proceedings, 2001. X, 395 pages. 2001.
- Vol. 2153: A.L. Buchsbaum, J. Snoeyink (Eds.), Algorithm Engineering and Experimentation. Proceedings, 2001. VIII, 231 pages. 2001.
- Vol. 2154: K.G. Larsen, M. Nielsen (Eds.), CONCUR 2001 – Concurrency Theory. Proceedings, 2001. XI, 583 pages. 2001.
- Vol. 2157: C. Rouveirol, M. Sebag (Eds.), Inductive Logic Programming. Proceedings, 2001. X, 261 pages. 2001. (Subseries LNAI).
- Vol. 2158: D. Shepherd, J. Finney, L. Mathy, N. Race (Eds.), Interactive Distributed Multimedia Systems. Proceedings, 2001. XIII, 258 pages. 2001.
- Vol. 2159: J. Kelemen, P. Sosfk (Eds.), Advances in Artificial Life. Proceedings, 2001. XIX, 724 pages. 2001. (Subseries LNAI).
- Vol. 2161: F. Meyer auf der Heide (Ed.), Algorithms – ESA 2001. Proceedings, 2001. XII, 538 pages. 2001.
- Vol. 2162: Ç. K. Koç, D. Naccache, C. Paar (Eds.), Cryptographic Hardware and Embedded Systems – CHES 2001. Proceedings, 2001. XIV, 411 pages. 2001.
- Vol. 2164: S. Pierre, R. Glitho (Eds.), Mobile Agents for Telecommunication Applications. Proceedings, 2001. XI, 292 pages. 2001.
- Vol. 2165: L. de Alfaro, S. Gilmore (Eds.), Process Algebra and Probabilistic Methods. Proceedings, 2001. XII, 217 pages. 2001.
- Vol. 2166: V. Matoušek, P. Mautner, R. Mouček, K. Taušer (Eds.), Text, Speech and Dialogue. Proceedings, 2001. XIII, 452 pages. 2001. (Subseries LNAI).
- Vol. 2170: S. Palazzo (Ed.), Evolutionary Trends of the Internet. Proceedings, 2001. XIII, 722 pages. 2001.
- Vol. 2172: C. Batini, F. Giunchiglia, P. Giorgini, M. Mecella (Eds.), Cooperative Information Systems. Proceedings, 2001. XI, 450 pages. 2001.
- Vol. 2176: K.-D. Althoff, R.L. Feldmann, W. Müller (Eds.), Advances in Learning Software Organizations. Proceedings, 2001. XI, 241 pages. 2001.
- Vol. 2177: G. Butler, S. Jarzabek (Eds.), Generative and Component-Based Software Engineering. Proceedings, 2001. X, 203 pages. 2001.
- Vol. 2181: C. Y. Westort (Eds.), Digital Earth Moving. Proceedings, 2001. XII, 117 pages. 2001.
- Vol. 2184: M. Tucci (Ed.), Multimedia Databases and Image Communication. Proceedings, 2001. X, 225 pages. 2001.
- Vol. 2186: J. Bosch (Ed.), Generative and Component-Based Software Engineering. Proceedings, 2001. VIII, 177 pages. 2001.
- Vol. 2188: F. Bomarius, S. Komi-Sirviö (Eds.), Product Focused Software Process Improvement. Proceedings, 2001. XI, 382 pages. 2001.
- Vol. 2189: F. Hoffmann, D.J. Hand, N. Adams, D. Fisher, G. Guimaraes (Eds.), Advances in Intelligent Data Analysis. Proceedings, 2001. XII, 384 pages. 2001.
- Vol. 2190: A. de Antonio, R. Aylett, D. Ballin (Eds.), Intelligent Virtual Agents. Proceedings, 2001. VIII, 245 pages. 2001. (Subseries LNAI).
- Vol. 2191: B. Radig, S. Florczyk (Eds.), Pattern Recognition. Proceedings, 2001. XVI, 452 pages. 2001.
- Vol. 2193: F. Casati, D. Georgakopoulos, M.-C. Shan (Eds.), Technologies for E-Services. Proceedings, 2001. X, 213 pages. 2001.

Table of Contents

Invited Paper

Advances in Model Representations	1
<i>M. Siegle</i>	

Contributed Papers

Faster and Symbolic CTMC Model Checking	23
<i>J.-P. Katoen, M. Kwiatkowska, G. Norman, and D. Parker</i>	
Reachability Analysis of Probabilistic Systems by Successive Refinements	39
<i>P.R. D'Argenio, B. Jeannet, H.E. Jensen, and K.G. Larsen</i>	
Beyond Memoryless Distributions: Model Checking Semi-Markov Chains	57
<i>G.G. Infante López, H. Hermanns, and J.-P. Katoen</i>	
Coin Lemmas with Random Variables	71
<i>K. Folegati and R. Segala</i>	
MoDeST – A Modelling and Description Language for Stochastic Timed Systems	87
<i>P.R. D'Argenio, H. Hermanns, J.-P. Katoen, and R. Klaren</i>	
Randomization Helps in LTL Model Checking	105
<i>L. Brim, I. Černá, and M. Nečesal</i>	
An Efficient Kronecker Representation for PEPA Models	120
<i>J. Hillston and L. Kloul</i>	
Reward Based Congruences: Can We Aggregate More?	136
<i>M. Bernardo and M. Bravetti</i>	
Using Max-Plus Algebra for the Evaluation of Stochastic Process Algebra Prefixes	152
<i>L. Cloth, H. Bohnenkamp, and B. Haverkort</i>	
Expressing Processes with Different Action Durations through Probabilities	168
<i>M. Bravetti and A. Aldini</i>	
Quantifying the Dynamic Behavior of Process Algebras	184
<i>P. Buchholz and P. Kemper</i>	

Implementing a Stochastic Process Algebra within
the Möbius Modeling Framework..... 200
 G. Clark and W.H. Sanders

Author Index 217

Advances in Model Representations

Markus Siegle

Lehrstuhl für Informatik 7, University of Erlangen-Nürnberg, Germany
siegle@informatik.uni-erlangen.de

Abstract. We review high-level specification formalisms for Markovian performability models, thereby emphasising the role of structuring concepts as realised par excellence by stochastic process algebras. Symbolic representations based on decision diagrams are presented, and it is shown that they quite ideally support compositional model construction and analysis.

1 Introduction

Stochastic models have a long tradition in the areas of performance and dependability evaluation. Since their specification at the level of the Markov chain is tedious and error-prone, several high-level model specification formalisms have been developed, such as queueing networks, stochastic Petri nets and networks of stochastic automata, which allow humans to describe the intended behaviour at a convenient level of abstraction. Although under Markovian assumptions the analysis of the underlying stochastic process does not pose any conceptual problems, the size of the underlying state space often renders models intractable in practice. Structuring concepts have shown to be of great value in order to alleviate this well-known state space explosion problem.

A process algebra is a mathematically founded specification formalism which provides compositional features, such as parallel composition of components, abstraction from internal actions, and the replacing of components by behaviourally equivalent ones. Therefore, stochastic extensions of process algebras are among the methods of choice for constructing complex, hierarchically structured stochastic models.

Recently, decision diagrams, which were originally developed as memory-efficient representations of Boolean functions in the area of hardware verification, have been extended in order to capture the numerical information which is contained in stochastic models. They have already been successfully used as the underlying data structure in prototype tools for performance analysis and verification of probabilistic systems. In this paper, it is shown that symbolic representations based on decision diagrams are particularly attractive if applied in a compositional context, as provided, for example, by a process algebraic specification formalism. In many cases, decision diagrams allow extremely compact representations of huge state spaces, and it has been demonstrated that all steps of model construction, manipulation and analysis (be it model checking, numerical analysis, or a combination of the two) can be carried out on the decision

diagram based representations. Thus, we argue that decision diagrams fit in well with structured modelling formalisms and open new ways towards increasing the range of manageable performance evaluation and verification problems.

This paper does not intend to present new research results, but to survey the history of structured model representations, with special emphasis on process algebras and symbolic encodings. We provide many pointers to further reading, without attempting to be exhaustive.

The paper is organised as follows: In Sec. 2, we survey the evolution from monolithic to modular model specification formalisms. Sec. 3 reviews the concept of stochastic process algebras. Sec. 4 introduces the symbolic representation of Markovian models with the help of multi-terminal binary decision diagrams and describes compositional model construction, manipulation and analysis on the basis of this data structure. The paper concludes with Sec. 5.

2 From Unstructured to Structured Models

2.1 Monolithic Model Representations

Continuous Time Markov Chains (CTMC) are the basic formalism for specifying performance and dependability models¹. A CTMC consists of a (finite, for our purpose) set of states and a finite set of transitions between states. The transitions are labelled by positive reals, called the transition rates, which determine transition probabilities and state sojourn times (the latter being exponentially distributed). Time-dependent state probabilities can be derived by solving a system of ordinary differential equations, and steady-state probabilities are calculated by solving a linear system of equations (see, for instance [99]). In order to save memory space, CTMCs are commonly represented as sparse matrices, where essentially only the non-zero entries are stored.

The direct specification of a CTMC at the level of individual states and state-to-state transitions is tedious and error-prone, and therefore only feasible for very small models. This motivated researchers to develop high-level specification formalisms for defining Markovian models at a level of abstraction which is more convenient for the human modeller. The most popular of these formalisms are queueing networks and stochastic Petri nets.

Queueing networks (QN), developed mainly in the 1960ies and 1970ies for modelling time-sharing and polling systems, describe customers moving between stations where they receive service after possibly waiting for a service unit to become available. The aim of analysis is typically the mean or distribution of the number of customers at a station, the customer throughput at a station, or the waiting time. The success of queueing networks stems mainly from the fact that for the class of product form networks [5] very efficient analysis algorithms, such as Buzen's algorithm [24] or mean-value analysis [86], are known, and that software tools for the specification and analysis of QN models were available at

¹ In this paper, we do not consider the line of research on non-Markovian models such as described, for example, in [47].

an early stage [89,100]. Although QN have been extended in various directions, e.g. in order to model the forking and synchronisation of jobs (fork-join QNs, [3,68,80]), the formalism of QNs is not suitable for the modelling of arbitrary systems, but specialised to the application area of shared resource systems.

Stochastic Petri nets (SPN) were developed in the 1980ies for modelling complex synchronisation schemes which cannot easily be expressed by queueing models [79]. The modelling primitives of Petri nets (places, transitions, markings) are very basic and do not carry any application-specific semantics. For that reason, Petri nets are universally applicable and very flexible, which is reflected by the fact that they have been successfully applied to many different areas of application. In the class of generalised SPNs (GSPN) [1,2], transitions are either timed or immediate. Timed transitions are associated with an exponentially distributed firing time, while immediate transitions fire as soon as they are enabled. During the analysis of a GSPN, the reachability graph is generated and the so-called vanishing markings, which are due to the firing of immediate transitions, are eliminated. The result is a CTMC whose analysis yields (steady-state or transient) state probabilities, i.e. the probabilities of the individual net markings, from which high-level measures can be computed.

Some software tools for performance modelling, e.g. USENUM [90], MARCA [98], MOSEL [8] and DNAmaca [70], implement their own specialised model description languages, which can also be considered as high-level specification formalisms for CTMCs.

With the help of the high-level model specification formalisms considered so far it is possible to specify larger CTMCs than at the state-to-state level, but these formalisms do not support the concepts of modularity, hierarchy or composition of submodels. As a result, the models are monolithic and may be difficult to understand and debug. Moreover, state space generation and numerical analysis of very large monolithic CTMCs is often not feasible in practice due to memory and CPU time limitations, which is referred to as the notorious state space explosion problem.

A large state space may become tractable if it is decomposed into smaller parts [95,33]. Instead of analysing one large system, the decomposition approach relies on analysing several small subsystems, analysing an aggregated overall system, and afterwards combining the subsystems' solutions accordingly. In general, this approach works well for nearly completely decomposable (NCD) systems whose state space can be partitioned into disjoint subsets of states, such that there is a lot of interaction between states belonging to the same subset, but little interaction between states belonging to different subsets. For the class of reversible Markov chains, the decomposition/aggregation approach yields exact results [32]. We mention that the approach may also be applied iteratively [34,25]. The major question is, of course, how to best partition a given state space, and in general this information should be derived from a modular high-level model specification. Approximate decomposition-based analysis methods for stochastic process algebra models (see Sec. 3) are discussed in [73], where time scale decomposition is based on the concept of NCD Markov chains [65,72],

and response time approximation relies on a structural decomposition for the special class of decision-free processes [74]. Another such approach, based on the exploitation of the structure of a special class of process algebraic models, is described in [7]. Approximate decomposition-based analysis for nearly-independent GSPN structures is considered in [27,30].

2.2 Modular Model Representations

Queueing models, stochastic Petri nets and the tool-specific modelling languages mentioned above do not offer the possibility of composing an overall model from components which can be specified in isolation. Such a composition, however, is a highly desirable feature when modelling complex systems, since it enables human users to focus on manageable parts from which a whole system can be constructed. For instance, modern performance analysis advocates a separation of the load model and the machine model, an idea developed already in [69,62,63], and similar ideas are also applied in stochastic rendezvous networks [101] and layered queueing networks [87]. As another, specific example, suppose one wished to model a communication system where two partners communicate over some communication medium. The model should reflect this structure, i.e. it should consist of three interacting submodels, one for each partner, and one for the medium. The user should be able to specify these three submodels more or less independently of each other and then simply specify the way in which they interact.

In the basic GSPN formalism, a model consists of a single net which covers the whole system to be studied. Therefore, GSPN models of complex systems tend to become very large and confused and suffer from the state space explosion problem. Stochastic activity networks [88,35] constitute an approach to the structuring of GSPNs through the sharing of places between different subnets. In the presence of symmetric submodels, they tackle the state space explosion problem by directly generating a reduced reachability graph in which all mutually symmetric markings are collapsed into one. Symmetries also play a predominant role for the analysis of stochastic well-formed coloured Petri nets [26,43], where a reduced reachability graph is constructed directly from the net description, without the need to construct the full reachability graph first. Another line of research is concerned with building SPNs in a structured way, basically by synchronising subnets via common transitions, which is an instance of the Kronecker approach described below.

Stochastic automata networks (SAN)², developed in the 1980ies and 1990ies [82,83,84,85], consist of several stochastic automata, basically CTMCs whose transitions are labelled with event names, which run in parallel and may perform certain synchronising events together. Thus, the SAN formalism is truly structured, since it allows the user to specify an overall model as a collection of interacting submodels. The major attraction of SANs is their memory-efficient

² The acronym SAN is also used for stochastic activity networks (see above), but in this paper it stands for stochastic automata networks.