

LNCS 4346

Luboš Brim  
Boudewijn Haverkort  
Martin Leucker  
Jaco van de Pol (Eds.)

# Formal Methods: Applications and Technology

11th International Workshop, FMICS 2006  
and 5th International Workshop, PDMC 2006  
Bonn, Germany, August 2006, Revised Selected Papers



Springer

TP311-53  
F723  
2006

Luboš Brim Boudewijn Haverkort  
Martin Leucker Jaco van de Pol (Eds.)

# Formal Methods: Applications and Technology

11th International Workshop, FMICS 2006  
and 5th International Workshop, PDMC 2006  
Bonn, Germany, August 26-27, and August 31, 2006  
Revised Selected Papers



Springer



E2007003130

## Volume Editors

Luboš Brim  
Masaryk University  
Botanická 68a, 602 00 Brno, Czech Republic  
E-mail: brim@fi.muni.cz

Boudewijn Haverkort  
University of Twente  
P.O. Box 217, 7500AE Enschede, The Netherlands  
E-mail: brh@cs.utwente.nl

Martin Leucker  
Technische Universität München  
Boltzmannstr. 3, 85748 Garching, Germany  
E-mail: leucker@in.tum.de

Jaco van de Pol  
Centrum voor Wiskunde en Informatica, SEN 2  
P.O. Box 94079, 1090 GB Amsterdam, The Netherlands  
E-mail: Jaco.van.de.Pol@cwi.nl

Library of Congress Control Number: 2007921124

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

LNCS Sublibrary: SL 2 – Programming and Software Engineering

ISSN 0302-9743  
ISBN-10 3-540-70951-7 Springer Berlin Heidelberg New York  
ISBN-13 978-3-540-70951-0 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

© Springer-Verlag Berlin Heidelberg 2007  
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India  
Printed on acid-free paper SPIN: 12021901 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

These are the joint final proceedings of the 11th International Workshop on Formal Methods for Industrial Critical Systems (FMICS 2006) and the fifth International Workshop on Parallel and Distributed Methods in Verification (PDMC 2006). Both workshops were organized as satellite events of CONCUR 2006, the 17th International Conference on Concurrency Theory that was organized in Bonn, August 2006.

The FMICS workshop continued successfully the aim of the FMICS working group – to promote the use of formal methods for industrial applications, by supporting research in this area and its application in industry. The emphasis in these workshops is on the exchange of ideas between researchers and practitioners, in both industry and academia.

This year the Program Committee received a record number of submissions. The 16 accepted regular contributions and 2 accepted tool papers, selected out of a total of 47 submissions, cover formal methodologies for handling large state spaces, model-based testing, formal description and analysis techniques as well as a range of applications and case studies.

The workshop program included two invited talks, by Anna Slobodova from Intel on “Challenges for Formal Verification in an Industrial Setting” and by Edward A. Lee from the University of California at Berkeley on “Making Concurrency Mainstream.” The former full paper can be found in this volume.

Following the tradition of previous workshops, the European Association of Software Science and Technology (EASST) supported a best paper award. This award was granted to Michael Weber and Moritz Hammer for their excellent paper “‘To Store or Not To Store’ Reloaded: Reclaiming Memory on Demand.”

The primary goal of the PDMC workshop series is to present and discuss recent developments in the young area of parallel and distributed methods in verification. Several verification techniques, ranging over model checking, equivalence checking, theorem proving, constraint solving and dependability analysis are addressed by the PDMC community. Verification problems are usually very demanding tasks, especially because the systems that we build and want to verify become increasingly complex.

On the other hand, parallel and distributed computing machinery is widely available. Algorithms and tools must be developed to use this hardware optimally for our verification tasks. Traditionally, we studied algorithms for homogeneous situations, such as parallel shared-memory computers and distributed clusters of PCs. Currently, the emphasis is shifting towards heterogeneous GRIDs. But even modern desktop PCs are quite heterogeneous, consisting of multiple core processors, various memory devices and cache levels, all with their own performance characteristics.

This year's PDMC had nine submissions; six papers were selected for presentation, and four papers were accepted for publication in this volume. In addition, Luboš Brim from Masaryk University, Brno, gave an invited lecture on "Distributed Verification: Exploring the Power of Raw Computing Power." The full paper can also be found in this volume.

We would like to thank all authors for their submissions. We would also like to thank the members of both Program Committees, and the additional referees, for their timely reviewing and lively participation in the subsequent discussion—the quality of the contributions in this volume are also due to their efforts and expertise.

The organizers wish to thank CONCUR for hosting the FMICS and PDMC 2006 workshops and taking care of many administrative aspects, and ERCIM for its financial support of FMICS. Additionally, the organizers would like to thank the EASST (European Association of Software Science and Technology), the Faculty of Informatics, Masaryk University Brno and the Technical University Munich, the CWI (Center of Mathematics and Computer Science, Amsterdam) and the University of Twente for supporting these events.

December 2006

Luboš Brim  
Boudewijn R. Haverkort  
Martin Leucker  
Jaco van de Pol

# Organization

## FMICS

### Program Chairs

Luboš Brim  
Martin Leucker

Masaryk University Brno, Czech Republic  
Technical University of Munich, Germany

### Program Committee

Rance Cleaveland  
Wan Fokkink

University of Maryland, USA  
Vrije Universiteit Amsterdam and CWI, The Netherlands

Stefania Gnesi  
Susanne Graf  
David Harel  
Klaus Havelund  
Thomas A. Henzinger  
Leszek Holenderski  
Stefan Kowalewski  
Marta Kwiatkowska  
Salvatore La Torre  
Tiziana Margaria  
Radu Mateescu  
Doron Peled  
Ernesto Pimentel  
Andreas Podelski  
Don Sannella  
Joseph Sifakis

ISTI-CNR, Italy  
VERIMAG, France  
Weizmann Institute of Science, Israel  
Kestrel Technology, USA  
EPFL, Switzerland  
Philips Research, The Netherlands  
RWTH Aachen University, Germany  
University of Birmingham, UK  
Università degli Studi di Salerno, Italy  
University of Göttingen, Germany  
INRIA Rhône-Alpes and ENS Lyon, France  
University of Warwick, UK  
University of Malaga, Spain  
Max-Planck-Institut für Informatik, Germany  
University of Edinburgh, UK  
VERIMAG, France

## PDMC

### Program Chairs

Boudewijn Haverkort  
Jaco van de Pol

University of Twente, The Netherlands  
CWI Amsterdam, The Netherlands

### Program Committee

Gerd Behrmann  
Ivana Černá  
Gianfranco Ciardo  
Joerg Denzinger

Aalborg University, Denmark  
Masaryk University Brno, Czech Republic  
University of California at Riverside, USA  
University of Calgary, Canada

## VIII Organization

Hubert Garavel	INRIA Rhône-Alpes, France
Orna Grumberg	Technion, Haifa, Israel
William Knottenbelt	Imperial College, London, UK
Marta Kwiatkowska	University of Birmingham, UK
Martin Leucker	Technical University of Munich, Germany

## Referees (FMICS and PDMC)

C. Artho	I. Černá	M. Kuntz	D. Parker
Y. Atir	F. Ciesinski	F. Lang	G. Parlato
R. Atkey	M. Faella	P. Lopez	G. Salaün
J. Barnat	A. Fantechi	K. MacKenzie	W. Serwe
M. ter Beek	M. Felici	P. Maier	F. Sorrentino
M. van der Bijl	A. J. Fernandez	S. Maoz	J. Tenzer
B. Bollig	M. Fruth	F. Mazzanti	A. Venet
L. Bozzelli	N. Geisweiller	R. Merom	A. Wijs
A. Bucchiarone	A. Goldberg	A. Murano	T. Willemse
D. Calvanese	A. Idani	G. Norman	V. Wolf
M. V. Cengarle	C. Joubert	M. Parente	



# Lecture Notes in Computer Science

For information about Vols. 1–4300

please contact your bookseller or Springer

- Vol. 4429: C. Ullrich, J.H. Siekmann, R. Lu (Eds.), *Cognitive Systems*. X, 162 pages. 2007. (Sublibrary LNAI).
- Vol. 4405: L. Padgham, F. Zambonelli (Eds.), *Agent-Oriented Software Engineering VII*. XII, 225 pages. 2007.
- Vol. 4403: S. Obayashi, K. Deb, C. Poloni, T. Hiroyasu, T. Murata (Eds.), *Evolutionary Multi-Criterion Optimization*. XIX, 954 pages. 2007.
- Vol. 4397: C. Stephanidis, M. Pieper (Eds.), *Universal Access in Ambient Intelligence Environments*. XV, 467 pages. 2007.
- Vol. 4396: J. García-Vidal, L. Cerdà-Alabern (Eds.), *Wireless Systems and Mobility in Next Generation Internet*. IX, 271 pages. 2007.
- Vol. 4394: A. Gelbukh (Ed.), *Computational Linguistics and Intelligent Text Processing*. XVI, 648 pages. 2007.
- Vol. 4393: W. Thomas, P. Weil (Eds.), *STACS 2007*. XVIII, 708 pages. 2007.
- Vol. 4392: S.P. Vadhan (Ed.), *Theory of Cryptography*. XI, 595 pages. 2007.
- Vol. 4390: S.O. Kuznetsov, S. Schmidt (Eds.), *Formal Concept Analysis*. X, 329 pages. 2007. (Sublibrary LNAI).
- Vol. 4385: K. Coninx, K. Luyten, K.A. Schneider (Eds.), *Task Models and Diagrams for Users Interface Design*. XI, 355 pages. 2007.
- Vol. 4384: T. Washio, K. Satoh, H. Takeda, A. Inokuchi (Eds.), *New Frontiers in Artificial Intelligence*. IX, 401 pages. 2007. (Sublibrary LNAI).
- Vol. 4383: E. Bin, A. Ziv, S. Ur (Eds.), *Hardware and Software, Verification and Testing*. XII, 235 pages. 2007.
- Vol. 4381: J. Akiyama, W.Y.C. Chen, M. Kano, X. Li, Q. Yu (Eds.), *Discrete Geometry, Combinatorics and Graph Theory*. XI, 289 pages. 2007.
- Vol. 4380: S. Spaccapietra, P. Atzeni, F. Fages, M.-S. Hacid, M. Kifer, J. Mylopoulos, B. Pernici, P. Shvaiko, J. Trujillo, I. Zaihrayeu (Eds.), *Journal on Data Semantics VIII*. XV, 219 pages. 2007.
- Vol. 4378: I. Virbitskaite, A. Voronkov (Eds.), *Perspectives of Systems Informatics*. XIV, 496 pages. 2007.
- Vol. 4377: M. Abe (Ed.), *Topics in Cryptology – CT-RSA 2007*. XI, 403 pages. 2006.
- Vol. 4376: E. Frachtenberg, U. Schwiegelshohn (Eds.), *Job Scheduling Strategies for Parallel Processing*. VII, 257 pages. 2007.
- Vol. 4373: K. Langendoen, T. Voigt (Eds.), *Wireless Sensor Networks*. XIII, 358 pages. 2007.
- Vol. 4372: M. Kaufmann, D. Wagner (Eds.), *Graph Drawing*. XIV, 454 pages. 2007.
- Vol. 4371: K. Inoue, K. Satoh, F. Toni (Eds.), *Computational Logic in Multi-Agent Systems*. X, 315 pages. 2007. (Sublibrary LNAI).
- Vol. 4370: P.P. Lévy, B. Le Grand, F. Poulet, M. Soto, L. Darago, L. Toubiana, J.-F. Vibert (Eds.), *Pixelization Paradigm*. XV, 279 pages. 2007.
- Vol. 4369: M. Umeda, A. Wolf, O. Bartenstein, U. Geske, D. Seipel, O. Takata (Eds.), *Declarative Programming for Knowledge Management*. X, 229 pages. 2006. (Sublibrary LNAI).
- Vol. 4368: T. Erlebach, C. Kaklamanis (Eds.), *Approximation and Online Algorithms*. X, 345 pages. 2007.
- Vol. 4367: K. De Bosschere, D. Kaeli, P. Stenström, D. Whalley, T. Ungerer (Eds.), *High Performance Embedded Architectures and Compilers*. XI, 307 pages. 2007.
- Vol. 4366: K. Tuyls, R. Westra, Y. Saeys, A. Nowé (Eds.), *Knowledge Discovery and Emergent Complexity in Bioinformatics*. IX, 183 pages. 2007. (Sublibrary LNBI).
- Vol. 4364: T. Kühne (Ed.), *Models in Software Engineering*. XI, 332 pages. 2007.
- Vol. 4362: J. van Leeuwen, G.F. Italiano, W. van der Hoek, C. Meinel, H. Sack, F. Pláčil (Eds.), *SOFSEM 2007: Theory and Practice of Computer Science*. XXI, 937 pages. 2007.
- Vol. 4361: H.J. Hoogeboom, G. Păun, G. Rozenberg, A. Salomaa (Eds.), *Membrane Computing*. IX, 555 pages. 2006.
- Vol. 4360: W. Dubitzky, A. Schuster, P.M.A. Sloot, M. Schroeder, M. Romberg (Eds.), *Distributed, High-Performance and Grid Computing in Computational Biology*. X, 192 pages. 2007. (Sublibrary LNBI).
- Vol. 4358: R. Vidal, A. Heyden, Y. Ma (Eds.), *Dynamical Vision*. IX, 329 pages. 2007.
- Vol. 4357: L. Buttyán, V. Gligor, D. Westhoff (Eds.), *Security and Privacy in Ad-Hoc and Sensor Networks*. X, 193 pages. 2006.
- Vol. 4355: J. Julliand, O. Kouchnarenko (Eds.), *B 2007: Formal Specification and Development in B*. XIII, 293 pages. 2006.
- Vol. 4354: M. Hanus (Ed.), *Practical Aspects of Declarative Languages*. X, 335 pages. 2006.
- Vol. 4353: T. Schwentick, D. Suciu (Eds.), *Database Theory – ICDT 2007*. XI, 419 pages. 2006.
- Vol. 4352: T.-J. Cham, J. Cai, C. Dorai, D. Rajan, T.-S. Chua, L.-T. Chia (Eds.), *Advances in Multimedia Modeling*, Part II. XVIII, 743 pages. 2006.
- Vol. 4351: T.-J. Cham, J. Cai, C. Dorai, D. Rajan, T.-S. Chua, L.-T. Chia (Eds.), *Advances in Multimedia Modeling*, Part I. XIX, 797 pages. 2006.

- Vol. 4349: B. Cook, A. Podelski (Eds.), Verification, Model Checking, and Abstract Interpretation. XI, 395 pages. 2007.
- Vol. 4348: S.T. Taft, R.A. Duff, R.L. Brukardt, E. Ploedereder, P. Leroy (Eds.), Ada 2005 Reference Manual. XXII, 765 pages. 2006.
- Vol. 4347: J. Lopez (Ed.), Critical Information Infrastructures Security. X, 286 pages. 2006.
- Vol. 4346: L. Brim, B. Haverkort, M. Leucker, J. van de Pol (Eds.), Formal Methods: Applications and Technology. X, 363 pages. 2007.
- Vol. 4345: N. Maglaveras, I. Chouvarda, V. Koutkias, R. Brause (Eds.), Biological and Medical Data Analysis. XIII, 496 pages. 2006. (Sublibrary LNBI).
- Vol. 4344: V. Gruhn, F. Oquendo (Eds.), Software Architecture. X, 245 pages. 2006.
- Vol. 4342: H. de Swart, E. Orłowska, G. Schmidt, M. Roubens (Eds.), Theory and Applications of Relational Structures as Knowledge Instruments II. X, 373 pages. 2006. (Sublibrary LNAI).
- Vol. 4341: P.Q. Nguyen (Ed.), Progress in Cryptology - VIETCRYPT 2006. XI, 385 pages. 2006.
- Vol. 4340: R. Prodan, T. Fahringer, Grid Computing. XXIII, 317 pages. 2007.
- Vol. 4339: E. Ayguadé, G. Baumgartner, J. Ramanujam, P. Sadayappan (Eds.), Languages and Compilers for Parallel Computing. XI, 476 pages. 2006.
- Vol. 4338: P. Kalra, S. Peleg (Eds.), Computer Vision, Graphics and Image Processing. XV, 965 pages. 2006.
- Vol. 4337: S. Arun-Kumar, N. Garg (Eds.), FSTTCS 2006: Foundations of Software Technology and Theoretical Computer Science. XIII, 430 pages. 2006.
- Vol. 4335: S.A. Brueckner, S. Hassas, M. Jelasity, D. Yamins (Eds.), Engineering Self-Organising Systems. XII, 212 pages. 2007. (Sublibrary LNAI).
- Vol. 4334: B. Beckert, R. Hähnle, P.H. Schmitt (Eds.), Verification of Object-Oriented Software. XXIX, 658 pages. 2007. (Sublibrary LNAI).
- Vol. 4333: U. Reimer, D. Karagiannis (Eds.), Practical Aspects of Knowledge Management. XII, 338 pages. 2006. (Sublibrary LNAI).
- Vol. 4332: A. Bagchi, V. Atluri (Eds.), Information Systems Security. XV, 382 pages. 2006.
- Vol. 4331: G. Min, B. Di Martino, L.T. Yang, M. Guo, G. Ruenger (Eds.), Frontiers of High Performance Computing and Networking - ISPA 2006 Workshops. XXXVII, 1141 pages. 2006.
- Vol. 4330: M. Guo, L.T. Yang, B. Di Martino, H.P. Zima, J. Dongarra, F. Tang (Eds.), Parallel and Distributed Processing and Applications. XVIII, 953 pages. 2006.
- Vol. 4329: R. Barua, T. Lange (Eds.), Progress in Cryptology - INDOCRYPT 2006. X, 454 pages. 2006.
- Vol. 4328: D. Penkler, M. Reitenspiess, F. Tam (Eds.), Service Availability. X, 289 pages. 2006.
- Vol. 4327: M. Baldoni, U. Endriss (Eds.), Declarative Agent Languages and Technologies IV. VIII, 257 pages. 2006. (Sublibrary LNAI).
- Vol. 4326: S. Göbel, R. Malkewitz, I. Iurgel (Eds.), Technologies for Interactive Digital Storytelling and Entertainment. X, 384 pages. 2006.
- Vol. 4325: J. Cao, I. Stojmenovic, X. Jia, S.K. Das (Eds.), Mobile Ad-hoc and Sensor Networks. XIX, 887 pages. 2006.
- Vol. 4323: G. Doherty, A. Blandford (Eds.), Interactive Systems. XI, 269 pages. 2007.
- Vol. 4320: R. Gotzhein, R. Reed (Eds.), System Analysis and Modeling: Language Profiles. X, 229 pages. 2006.
- Vol. 4319: L.-W. Chang, W.-N. Lie (Eds.), Advances in Image and Video Technology. XXVI, 1347 pages. 2006.
- Vol. 4318: H. Lipmaa, M. Yung, D. Lin (Eds.), Information Security and Cryptology. XI, 305 pages. 2006.
- Vol. 4317: S.K. Madria, K.T. Claypool, R. Kannan, P. Uppuluri, M.M. Gore (Eds.), Distributed Computing and Internet Technology. XIX, 466 pages. 2006.
- Vol. 4316: M.M. Dalkilic, S. Kim, J. Yang (Eds.), Data Mining and Bioinformatics. VIII, 197 pages. 2006. (Sublibrary LNBI).
- Vol. 4314: C. Freksa, M. Kohlhase, K. Schill (Eds.), KI 2006: Advances in Artificial Intelligence. XII, 458 pages. 2007. (Sublibrary LNAI).
- Vol. 4313: T. Margaria, B. Steffen (Eds.), Leveraging Applications of Formal Methods. IX, 197 pages. 2006.
- Vol. 4312: S. Sugimoto, J. Hunter, A. Rauber, A. Morishima (Eds.), Digital Libraries: Achievements, Challenges and Opportunities. XVIII, 571 pages. 2006.
- Vol. 4311: K. Cho, P. Jacquet (Eds.), Technologies for Advanced Heterogeneous Networks II. XI, 253 pages. 2006.
- Vol. 4310: T. Boyanov, S. Dimova, K. Georgiev, G. Nikolov (Eds.), Numerical Methods and Applications. XIII, 715 pages. 2007.
- Vol. 4309: P. Inverardi, M. Jazayeri (Eds.), Software Engineering Education in the Modern Age. VIII, 207 pages. 2006.
- Vol. 4308: S. Chaudhuri, S.R. Das, H.S. Paul, S. Tirthapura (Eds.), Distributed Computing and Networking. XIX, 608 pages. 2006.
- Vol. 4307: P. Ning, S. Qing, N. Li (Eds.), Information and Communications Security. XIV, 558 pages. 2006.
- Vol. 4306: Y. Avrithis, Y. Kompatsiaris, S. Staab, N.E. O'Connor (Eds.), Semantic Multimedia. XII, 241 pages. 2006.
- Vol. 4305: A.A. Shvartsman (Ed.), Principles of Distributed Systems. XIII, 441 pages. 2006.
- Vol. 4304: A. Sattar, B.-H. Kang (Eds.), AI 2006: Advances in Artificial Intelligence. XXVII, 1303 pages. 2006. (Sublibrary LNAI).
- Vol. 4303: A. Hoffmann, B.-H. Kang, D. Richards, S. Tsumoto (Eds.), Advances in Knowledge Acquisition and Management. XI, 259 pages. 2006. (Sublibrary LNAI).
- Vol. 4302: J. Domingo-Ferrer, L. Franconi (Eds.), Privacy in Statistical Databases. XI, 383 pages. 2006.
- Vol. 4301: D. Pointcheval, Y. Mu, K. Chen (Eds.), Cryptology and Network Security. XIII, 381 pages. 2006.

¥562.00元

# Table of Contents

## Invited Contributions

Challenges for Formal Verification in Industrial Setting .....	1
<i>Anna Slobodová</i>	
Distributed Verification: Exploring the Power of Raw Computing Power .....	23
<i>Luboš Brim</i>	

## FMICS

An Easy-to-Use, Efficient Tool-Chain to Analyze the Availability of Telecommunication Equipment .....	35
<i>Kai Lampka, Markus Siegle, and Max Walter</i>	
“To Store or Not To Store” Reloaded: Reclaiming Memory on Demand .....	51
<i>Moritz Hammer and Michael Weber</i>	
Discovering Symmetries .....	67
<i>Hassen Saïdi</i>	
On Combining Partial Order Reduction with Fairness Assumptions .....	84
<i>Luboš Brim, Ivana Černá, Pavel Moravec, and Jiří Šimša</i>	
Test Coverage for Loose Timing Annotations .....	100
<i>C. Helmstetter, F. Maraninchi, and L. Maillet-Contoz</i>	
Model-Based Testing of a WAP Gateway: An Industrial Case-Study ...	116
<i>Anders Hessel and Paul Pettersson</i>	
Heuristics for <b>ioco</b> -Based Test-Based Modelling .....	132
<i>Tim A.C. Willemse</i>	
Verifying VHDL Designs with Multiple Clocks in SMV .....	148
<i>A. Smrčka, V. Řehák, T. Vojnar, D. Šafránek, P. Matoušek, and Z. Řehák</i>	
Verified Design of an Automated Parking Garage .....	165
<i>Aad Mathijssen and A. Johannes Pretorius</i>	
Evaluating Quality of Service for Service Level Agreements .....	181
<i>Allan Clark and Stephen Gilmore</i>	

Simulation-Based Performance Analysis of a Medical Image-Processing Architecture .....	195
<i>P.J.L. Cuijpers and A.V. Fyukov</i>	
BLASTing Linux Code .....	211
<i>Jan Tobias Mühlberg and Gerald Lüttgen</i>	
A Finite State Modeling of AFDX Frame Management Using Spin .....	227
<i>Indranil Saha and Suman Roy</i>	
UML 2.0 State Machines: Complete Formal Semantics Via Core State Machines .....	244
<i>Harald Fecher and Jens Schönborn</i>	
Automated Incremental Synthesis of Timed Automata .....	261
<i>Borzoo Bonakdarpour and Sandeep S. Kulkarni</i>	
SAT-Based Verification of LTL Formulas .....	277
<i>Wenhui Zhang</i>	
jmlc: A Tool for Executing JML Specifications Via Constraint Programming .....	293
<i>Ben Krause and Tim Wahls</i>	
Goanna—A Static Model Checker .....	297
<i>Ansgar Fehnker, Ralf Huuck, Patrick Jayet, Michel Lussenburg, and Felix Rauch</i>	
<b>PDMC</b>	
Parallel SAT Solving in Bounded Model Checking .....	301
<i>Erika Ábrahám, Tobias Schubert, Bernd Becker, Martin Fränzle, and Christian Herde</i>	
Parallel Algorithms for Finding SCCs in Implicitly Given Graphs .....	316
<i>Jiří Barnat and Pavel Moravec</i>	
Can Saturation Be Parallelised? – On the Parallelisation of a Symbolic State-Space Generator .....	331
<i>Jonathan Ezekiel, Gerald Lüttgen, and Radu Siminiceanu</i>	
Distributed Colored Petri Net Model-Checking with CYCLADES .....	347
<i>Christophe Pajault and Jean-François Pradat-Peyre</i>	
<b>Author Index</b> .....	363

# Challenges for Formal Verification in Industrial Setting

Anna Slobodová

Intel

`anna.slobodova@intel.com`

**Abstract.** Commercial competition is forcing computer companies to get better products to market more rapidly, and therefore the time for validation is shrinking relative to the complexity of microprocessor designs. Improving time-to-market performance cannot be solved by just growing the size of design and validation teams. Design process automation is increasing, and the adoption of more rigorous methods, including formal verification, is unavoidable because for achieving the quality demanded by the marketplace.

Intel is one of the strongest promoters of the use of formal methods across all phases of the design development. Intel's design teams use high-level modeling of protocols and algorithms, formal verification of floating-point libraries, design exploration systems based on formal methods, full proofs and property verification of RTL specifications, and equivalence checking to verify that transistor-level schematics correspond to their RTL specifications. Even with the best effort to adopt the progress in formal methods quickly, there is a large gap between an idea published at a conference and a development of a tool that can be used on industrial-sized designs. These tools and methods need to scale well, be stable during a multi-year design effort, and be able to support efficient debugging. The use of formal methods on a live design must allow for ongoing changes in the specification and the design. The methodology must be flexible enough to permit new design features, such as scan and power-down logic, soft error detection, etc. In this paper, I will share my experience with the formal verification of the floating-point unit on an Itanium(R) microprocessor design and point out how it may influence future microprocessor-design projects.

## 1 Introduction

Floating-point (FP) arithmetic is, with respect to functional validation, one of the critical parts of modern microprocessor designs. Even though the algorithms for FP arithmetic are well known, optimization for high performance, reliability, testability and low power, may introduce bugs into a design. The huge input data space that needs to be explored to ensure correctness of floating-point designs is beyond the limits of traditional simulation techniques (hereafter referred to as *simulation*). Fortunately, formal methods are well suited for this area and they can enhance a verification effort substantially. Formal semantics of floating-point

operations can be expressed in a succinct way and the IEEE Floating-Point standard [IEEE] serves as a guide for many instruction-set architectures. In this paper, floating-point algorithms are not our concern. Instead, I focus on the correctness of their register-transfer level (RTL) implementations.

For almost a decade, papers reporting compliance proofs for circuit models with respect to the IEEE standard and particular instruction set architecture have been published. While the early research was focused on answering a principal question of the feasibility of formal proofs for computer arithmetic (e.g., [AS95, CB98, OZ+99]), recent work emanates from commercial industry. Formal tools and methods have reached the maturity necessary for their deployment in real design projects.

There are substantial differences between the methodologies used at different companies, depending on their target and available tools and resources. Intel and AMD were among first companies that applied formal methods, at first to verification of floating-point algorithms and then to RTL design. Methods reported from AMD design team in [Rus98, RF00, FK+02] are solely based on theorem proving using ACL2 system <sup>1</sup>. Although lot of automation has been added to building ACL2 models from RTL descriptions, and an ACL2 library of Floating-Point Arithmetic has been created to avoid repetition of implementation independent proofs, the methodology still requires high-level expertise in theorem proving and a perfect understanding of the design.

A recent paper from IBM by Jacobi *et al.* [JW+05] presents a verification method based on symbolic simulation of a RTL model and its comparison to a high-level model written in VHDL. Although highly automated, the approach is not as rigorous as the one described by AMD, and lacks the scope of the methodology developed at Intel [AJ+00][AJ+00, KA00, KN02, KK03, Sch03]. It skips the verification of the more difficult part of the design - multiplier, by removing it from the cone of influence, hence proving merely correctness of the adder and rounder. In contrast, in our approach, no abstraction or design modification took place. We return to the comparison of their work to our results in Section 7.

At Intel, an important work in the area of the verification of floating-point algorithms, and in particular, floating-point libraries for Itanium<sup>(R)</sup> has been done by John Harrison (see [Har05] for an overview, and [Har00a, Har00b, Har03] for details). However, in hardware verification, while many papers have been published on verification of Pentium<sup>(R)</sup> design, the first report on the formal verification of floating-point arithmetic for the Itanium<sup>(R)</sup> microprocessor family was reported on Designing Correct Circuits (DCC) Workshop <sup>2</sup>, in March, 2004, in Barcelona [SN04]. The main result was the first successful formal verification of floating-point fused multiply-add instruction which is in repertoire of the IA-64 Instruction Set Architecture (ISA). Proofs have been constructed for a live project aimed at a next generation of Itanium<sup>(R)</sup> microprocessor. We continued our work presented at DCC by extending the scope and verifying correctness

<sup>1</sup> <http://www.cs.utexas.edu/moore/acl2/>

<sup>2</sup> <http://www.math.chalmers.se/~ms/DCC04/>

of the rest of floating-point instructions (about 40) issued to execution pipe. All instructions have been verified with respect to eight precisions and four IEEE rounding modes, including dynamic rounding specified in floating-point status register. The verification was based on symbolic trajectory evaluation and arithmetic libraries previously proven within the same system [KN02]. The proof includes correctness of the result, update of floating-point status register, and correctness of more than a dozen interrupt signals. Behavior of the floating-point circuitry for invalid instructions and/or instructions with false qualifying predicates have been considered as well. All proofs have been regularly rerun as a regression suite to ensure the consistency of any changes in the design. Formal sequential equivalence checking was used to finish the validation of low-level design proving its correctness with respect to the RTL. However this last phase of verification is out of scope of this paper.

In the process of constructing our proofs we found many bugs and issues that required RTL changes. Our work also helped to clarify incomplete and ambiguous parts of our micro-architectural specification, and it contributed to some hardware optimizations. The proofs are automated and portable to other Itanium<sup>(R)</sup> micro-processor designs.

The goal of this paper is to describe the scope and results of our work, and to provide some insight into challenges of using formal verification in an industrial environment, where a fine balance between rigorous verification methods and traditional simulation-based methods is crucial for success of the validation. Although the approach we choose is a combination of known techniques already documented in context of the verification of floating-point adders and multipliers, we believe that it has many aspects that might be interesting to researchers in academia as well as validation engineers.

The paper is organized in following way: Next section describes tools and methodology developed for formal verification of floating-point arithmetic at Intel Corporation and specifics of our approach. The core of our work is described in Section 3, where we dive into details of the verification of the most interesting operation - fused multiply-add, and report what has been covered by our proofs. Since debugging of failing proofs is one of the concerns in the use of formal methods, we touch this question in Section 4. Section 5 focuses on benefits of our effort for the design project. We describe our experience with proof management in Section 6. Concluding section contains summary of our work and detailed comparison to related published work.

## 2 Our Approach to Formal Verification of FP Arithmetic

Intel's approach to the validation of floating-point arithmetic includes a huge database of corner test-cases and pseudo-random generators for simulation, as well as Intel's FORTE formal verification tool that combines theorem-proving with model-checking capabilities<sup>3</sup>. The methodology described below does not

<sup>3</sup> A publicly available version of the tool that can be used for non-commercial purposes can be downloaded from <http://www.intel.com/software/products/opensource/>

rely on FORTE specifics and can be reproduced using any tool with capability of symbolic trajectory evaluation (STE) and some means of composing results obtained by STE. We believe that formal proofs coupled with traditional pseudo-random and focused simulation is a good way to achieve thorough functional validation. In our project, the formal and simulation based validation teams mutually benefited from their collaboration. However, this is out of the scope of this paper and we will focus on formal verification only.

## 2.1 FORTE System and STE

The history of formal verification of floating-point arithmetic at Intel has been motivated by two controversial trends: promising results in academia that were followed by proof of concept at Intel Research Lab [OZ+99]; and bugs that escaped to the micro-processor products [Coe96, Fis97]. Today, formal proofs developed for Pentium<sup>(R)</sup> designs [AJ+00, KA00, KN02, KK03, Sch03] are re-used and even put into hands of validation engineers that are not experts on formal methods. These proofs have been done using FORTE – a system built on top of VOSS. In this section, we give a rather informal description of the technology inside the FORTE system, just enough to understand the paper; details can be found in the referred publications. FORTE includes a light-weight theorem prover and a symbolic trajectory evaluation (STE) engine [STE]. The theorem prover is based on a higher-order logic. The interface language for FORTE is FL – a strongly-typed functional language in the ML family [Pau96]. One good property of FL as a specification language is its executability. While creating specifications, we often ran sanity checks. For instance, the translation from the memory format to register-file format was written as specified by the Software Developers Manual [ISA], and then checked whether consequent inverse translations yield consistent values. FL includes Binary Decision Diagrams (BDDs)[Bry86] as first-class objects and STE as a built-in function. For more information we refer the interested reader to the online documentation for the FORTE system and [KA00]. Here we describe the basic mechanisms of STE and the framework in which we work.

STE is a weak form of model-checking where a formal (gate-level) model is subjected to a symbolic simulation. The idea of a symbolic simulator is similar to that of standard simulator but it differs in that symbolic values (besides explicit binary values) are assigned to each signal and these values propagated through the design model. Results of such simulations are formulas for specified signals at specified times.

STE is an enhancement of symbolic simulation where Boolean logic has been extended to a lattice [STE] with  $X$  as a bottom (no information) and  $T$  as a top element (overconstrained).  $X$  is automatically assigned (by the STE simulator) to signals to which no value has been specified.  $X$  can be thought of as an unknown value. Its semantics and use are discussed later. Symbolic values are bound to signals at specified times to form signal trajectories. Trajectories that prune possible computations by restricting the values of some signals at specific



times are called *antecedents*; they can be interpreted as assumptions. Trajectories that specify expected responses of the circuit are called *consequents*. A specification is written in a form of Boolean expressions that constrain symbolic values in antecedents and consequents. Trajectory evaluation correctness statement  $\models_{ckt} [ant \implies cons]$  means: all circuit computations that satisfy antecedent *ant* also satisfy consequent *cons*. If any of consequent is violated, a STE run (proof) fails and a counterexample can be extracted from this failure. In fact, the failed proof provides all possible counterexamples and the user may select one for debugging purposes. If all consequents hold at every point of the simulation, success is reported by the tool.

## 2.2 Pre- and Post-condition Framework

Because of capacity limitations inherit in the STE engine, we may be forced to break our model into smaller pieces. In this case, we make sure that those pieces perfectly fit together. Informally, this means that the border signals of the decomposition match exactly and that nothing is left out of the design. Also the times at which we extract the values of the signals must be consistent. In terms of STE, consequents that include border signals serve as antecedents in the following step of the proof. In this way, we can use facts proved in one part as assumptions for later proofs.

The idea of proof (de)composition described above comes from the *pre-and-post-condition theory* used for verification of sequential programs. It was first applied to STE by Kaivola and Aagaard [KA00]. It allows one to prove the statements of the form  $\{P\}S\{Q\}$ , where  $P$  and  $Q$  are logical properties and  $S$  is a program. In our case, the program is replaced by a circuit and trajectories that bind values inputs and outputs of the circuit at specific times.  $\{P(x)\}(pretr_x, ckt, posttr_y)\{Q(x, y)\}$  represents the statement: if  $pretr_x$  binds the Boolean vector  $x$  to signals (usually inputs) of the circuit  $ckt$  and  $posttr_y$  binds the Boolean vector  $y$  to signals of the circuit (usually outputs), then the property  $P(x)$  guarantees property  $Q(x, y)$ .

$\{P(x)\}(pretr_x, ckt, posttr_y)\{Q(x, y)\}$  is a shorthand for the following formula:

$$\begin{aligned} \forall x(P(x) \Rightarrow (\exists y(\models_{ckt} [pretr_x \implies posttr_y])) \wedge \\ (\forall y((\models_{ckt} [pretr_x \implies posttr_y]) \Rightarrow Q(x, y)))) \end{aligned} \quad (1)$$

In our methodology,  $P$  is a conjunction of an *initial condition* that describes the restriction of inputs to the circuit, and an *auxiliary pre-condition* that is used to further restrict the simulation. For consistency, we use the same initial conditions throughout all proofs for every instruction analyzed, except when we weaken an initial condition to *true*. An example of an initial input condition is a statement that the specified input signals have value of a specific opcode. Auxiliary pre-conditions are usually used to simplify a particular STE run by restricting symbolic values (meaning that the inputs or internal nodes are restricted). An example of an auxiliary pre-condition is a restriction specifying