

LNCS 4337

S. Arun-Kumar
Naveen Garg (Eds.)

FSTTCS 2006: Foundations of Software Technology and Theoretical Computer Science

26th International Conference
Kolkata, India, December 2006
Proceedings



Springer

TP 31-53
F 771
2006

S. Arun-Kumar Naveen Garg (Eds.)

FSTTCS 2006: Foundations of Software Technology and Theoretical Computer Science

26th International Conference
Kolkata, India, December 13-15, 2006
Proceedings



Springer



E2007000050

Volume Editors

S. Arun-Kumar

Department of Computer Science and Engineering

Indian Institute of Technology Delhi

New Delhi 110016, India

E-mail: sak@cse.iitd.ac.in

Naveen Garg

Department of Computer Science and Engineering

Indian Institute of Technology Delhi

New Delhi 110016, India

E-mail: naveen@cse.iitd.ac.in

Library of Congress Control Number: 2006937157

CR Subject Classification (1998): F.3, D.3, F.4, F.2, F.1, G.2

LNCS Sublibrary: SL 1 – Theoretical Computer Science and General Issues

ISSN 0302-9743

ISBN-10 3-540-49994-6 Springer Berlin Heidelberg New York

ISBN-13 978-3-540-49994-7 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 2006

Printed in Germany

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

Printed on acid-free paper SPIN: 11944836 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

Welcome to the proceedings of the 26th annual conference on the Foundations of Software Technology and Theoretical Computer Science (FSTTCS). FSTTCS is organized by the Indian Association for Research in Computer Science (IARCS) which it helped create by bringing together academic computer scientists from various parts of India. Over the years several changes have taken place in the conference, beginning with the first Springer LNCS edition in 1984. The first three proceedings were printed and published in India by the Tata Institute of Fundamental Research. Since then other changes, such as an international Programme Committee and pre-conference and post-conference workshops, have helped nurture and enhance the status of this conference among computer scientists interested in foundational research.

Along with these changes there are quite a few invariant properties this conference enjoys. It has always been held in India, and always in the second or third week of December, which is a good time to travel to and within the country. It is also the most convenient time to meet Indian researchers, most of whom would be temporarily free of teaching and administrative commitments.

This year for the first time in its history, FSTTCS was held in the historic city of Kolkata (formerly Calcutta) at the Indian Statistical Institute (ISI). FSTTCS is one of the events to commemorate the Platinum jubilee year of ISI and we are happy to be part of this celebration. Subhas C. Nandy of ISI Kolkata took full responsibility as chairman of the Organizing Committee.

The conference attracted 155 submissions with authors from 29 countries. We thank the authors for their interest in this conference. The reputation of a conference is effectively determined by its Programme Committee, the refereeing process and the Invited talks. This year, as in the past, we were able to get highly respected researchers to serve on our Programme Committee. The refereeing process too has been of a very high quality. The Programme Committee deliberated on each of the submitted papers and the accompanying referee reports and finally decided to select 34 submissions as being worthy of publication.

The invited speakers this year were Gordon Plotkin of the University of Edinburgh, UK, Emo Welzl of ETH Zurich, Switzerland, Gérard Boudol of INRIA, Sophia Antipolis, France, David Shmoys of Cornell University, USA and Eugene Asarin of Université Paris 7, France.

In keeping with what has now become tradition, two workshops on recent advances in areas of current interest to the community were also organized. The themes this year were *Timed Systems* (organized by Deepak D'Souza and Supratik Chakraborty) and *Approximation Algorithms* (organized by Anupam Gupta and Amit Kumar).

We would like to thank the authors who submitted papers and the Programme Committee for helping us maintain the standard of this conference. We especially thank the authors of the selected papers which form the bulk of this volume. We

would like to thank Springer for agreeing to publish these proceedings in their prestigious *Lecture Notes in Computer Science* series, which has in no small way contributed to the status of this conference in academic circles.

We would also like to thank ISI Kolkata and Subhas Nandy's team for hosting FSTTCS, IIT Delhi and MPI-Informatik for providing support and of course, IARCS, the parent organization that FSTTCS gave birth to!

December 2006

Naveen Garg
S. Arun-Kumar
Programme Chairs
FSTTCS 2006

Organization

FSTTCS 2006 and the associated workshops were held at the Indian Statistical Institute, Kolkata.

Programme Committee

Amit Kumar (IIT Delhi, India)	Manindra Agrawal (IIT Kanpur, India)
Anil Seth (IIT Kanpur, India)	Marco Pistore (DIT Trento, Italy)
Anuj Dawar (Cambridge University, UK)	Marina Papatriantafilou (Chalmers UT, Sweden)
Anupam Gupta (CMU, USA)	Naveen Garg (IIT Delhi, India), Co-chair
Ashish Tiwari, (SRI, USA)	Neal Young (UC Riverside, USA)
Astrid Kiehn (IIT Delhi, India)	Nobuko Yoshida (Imperial College, UK)
Dale Miller (INRIA-Futurs, France)	Prakash Panangaden (McGill University, Canada)
Deepak D'Souza (IISc, India)	Radha Jagadeesan (DePaul University, USA)
Edgar Ramos (UIUC, USA)	Rohit Khandekar (University Waterloo, Canada)
Giuseppe Italiano (Università di Roma, Italy)	S. Arun-Kumar (IIT Delhi, India), Co-chair
Helmut Veith (TU München, Germany)	Sriram K. Rajamani (Microsoft Research, USA)
Javier Esparza (University of Stuttgart, Germany)	Subhas C. Nandy (ISI Kolkata, India)
Joost-Pieter Katoen (RWTH, Germany)	Supratik Chakraborty (IIT Bombay, India)
Kavitha Telikepalli (IISc, India)	Susanne Albers (University of Freiburg, Germany)
Madhavan Mukund (CMI, India)	Yuval Rabani (Technion, Israel)

Referees

Abhik Roychoudhury	Aditya Nori	Andrey Rybalchenko
Abraham Flaxman	Amitabha Bagchi	Andrzej Murawski
Achim Blumensath	Andrea Pacifici	Angelika Mader

Antoine Meyer	Kohei Honda	Rahul Jain
Arijit Bishnu	Konstantin Korovin	Ramesh Hariharan
Ashish Goel	Krishna S.	Ran Canetti
Ashwin Nayak	Krishnendu Chatterjee	Rana Barua
Avik Chaudhuri	Krishnendu	Richard Mayr
B. Meenakshi	Mukhopadhyaya	Robert Schweller
Barbara Koenig	Kunal Talwar	Rohit Parikh
Benedikt Bollig	Kunsoo Park	Ron van der Meyden
Bharat Adsul	Leonard Schulman	Ronald de Wolf
Bhargab Bhattacharya	Lukasz Kowalik	Russ Harmer
Bimal Roy	Luke Ong	Ryan Williams
Binay K. Bhattacharya	Madhu Gopinathan	S. P. Suresh
Boris Koldehofe	Madhu Sudan	S. Ramesh
C. A. Murthy	Madhusudan P.	Sanjiva Prasad
Carsten Kern	Marco Carbone	Sariel Har-Peled
Catuscia Palamidessi	Marielle Stoelinga	Satyanarayana Valluri
Christian Schallhart	Marko Samer	Saugata Basu
Daniel Hirschhoff	Markus Holzer	Sergey Bereg
Daniel Willems	Martin Berger	Soeren Laue
Daniele Varacca	Martin Neuhusser	Srinath Sridhar
David N. Jansen	Massimiliano Caramia	Sriram S.
Dino Distefano	Matthias Nickles	Stefan Funke
Doina Precup	Maurizio Naldi	Stefan Rieger
Elad Schiller	Meena Mahajan	Stephen Kobourov
Erika Abraham	Michael Kohler	Sudeb P. Pal
Fedor Fomin	Nabil Mustafa	Sudhakar G.
Franck van Breugel	Nicolas Markey	Suman Roy
G. Sivakumar	Nir Piterman	Sumit Ganguly
G. Paun	Norm Ferns	Sumit Jha
Guido Proietti	Oded Regev	Surender Baswana
Herbert Wiklicky	Oleg Sokolsky	Susmita Sur-Kolay
Hoeteck Wee	Olivier Carton	T. S. Hsu
Holger Schlingloff	Om P. Damani	Thomas Noll
Igor Walukiewicz	P. S. Thiagarajan	Thomas Thierauf
Irit Katriel	Palash Sarkar	Tingting Han
J. Radhakrishnan	Pallab Dasgupta	Umesh Bellur
Jayalal Sarma M. N.	Paolo Zuliani	Véronique Bruyère
Jochen Konemann	Paritosh Pandya	V. Arvind
Johannes Kinder	Pavithra Prabhakar	Valentin Goranko
K. Narayan Kumar	Peter Bro Miltersen	Vinayaka Pandit
Kamal Lodaya	Prakash Chandrasekaran	Vinodchandran Variyam
Kapil Vaswani	Pranab Sen	Walter Vogler
Kasturi R. Varadarajan	Prasad Naldurg	Willem Conradie
Katarzyna Paluch	Prasad Sistla	Yaoyun Shi
Khaled Elbassioni	Pushpak Bhattacharya	Yichen Xie
Kishan Gupta	R. Ramanujam	

Sponsoring Institutions

Indian Association for Research in Computer Science

Indian Institute of Technology, Delhi

Indian Statistical Institute, Kolkata

Lecture Notes in Computer Science

For information about Vols. 1–4231

please contact your bookseller or Springer

Vol. 4337: S. Arun-Kumar, N. Garg (Eds.), *FSTTCS 2006: Foundations of Software technology and Theoretical Computer Science*. XIII, 430 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*. XXVII, 1141 pages. 2006.

Vol. 4329: R. Barua, T. Lange (Eds.), *Progress in Cryptology – INDOCRYPT 2006*. XXXXX, 454000 pages. 2006.

Vol. 4318: H. Lipmaa, M. Yung, D. Lin (Eds.), *Information Security and Cryptology*. XI, 305 pages. 2006.

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. 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. 4304: A. Sattar, B.-H. Kang (Eds.), *AI 2006: Advances in Artificial Intelligence*. XXVII, 1303 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.

Vol. 4300: Y.Q. Shi (Ed.), *Transactions on Data Hiding and Multimedia Security I*. IX, 139 pages. 2006.

Vol. 4296: M.S. Rhee, B. Lee (Eds.), *Information Security and Cryptology – ICISC*. XIII, 358 pages. 2006.

Vol. 4295: J.D. Carswell, T. Tezuka (Eds.), *Web and Wireless Geographical Information Systems*. XI, 269 pages. 2006.

Vol. 4293: A. Gelbukh, C.A. Reyes-Garcia (Eds.), *MICA 2006: Advances in Artificial Intelligence*. XXVIII, 1232 pages. 2006. (Sublibrary LNAI).

Vol. 4292: G. Bebis, R. Boyle, B. Parvin, D. Koracin, P. Remagnino, A. Nefian, G. Meenakshisundaram, V. Pascucci, J. Zara, J. Molineros, H. Theisel, T. Malzbender (Eds.), *Advances in Visual Computing, Part II*. XXXII, 906 pages. 2006.

Vol. 4291: G. Bebis, R. Boyle, B. Parvin, D. Koracin, P. Remagnino, A. Nefian, G. Meenakshisundaram, V. Pascucci, J. Zara, J. Molineros, H. Theisel, T. Malzbender (Eds.), *Advances in Visual Computing, Part I*. XXXI, 916 pages. 2006.

Vol. 4290: M. van Steen, M. Henning (Eds.), *Middleware 2006*. XIII, 425 pages. 2006.

Vol. 4289: M. Ackermann, B. Berendt, M. Grobelnik, A. Hotho, D. Mladenich, G. Semeraro, M. Spiliopoulou, G. Stumme, V. Svatek, M. van Someren (Eds.), *Semantics, Web and Mining*. X, 197 pages. 2006. (Sublibrary LNAI).

Vol. 4288: T. Asano (Ed.), *Algorithms and Computation*. XXXX, 753333 pages. 2006.

Vol. 4285: Y. Matsumoto, R. Sproat, K.-F. Wong, M. Zhang (Eds.), *Computer Processing of Oriental Languages*. XXXX, 55000 pages. 2006. (Sublibrary LNAI).

Vol. 4284: X. Lai, K. Chen (Eds.), *Advances in Cryptology – ASIACRYPT 2006*. XIV, 468 pages. 2006.

Vol. 4283: Y.Q. Shi, B. Jeon (Eds.), *Digital Watermarking*. XII, 474 pages. 2006.

Vol. 4282: Z. Pan, A. Cheok, M. Haller, R.W.H. Lau, H. Saito, R. Liang (Eds.), *Advances in Artificial Reality and Tele-Existence*. XXIII, 1347 pages. 2006.

Vol. 4281: K. Barkaoui, A. Cavalcanti, A. Cerone (Eds.), *Theoretical Aspects of Computing – ICTAC 2006*. XV, 371 pages. 2006.

Vol. 4280: A.K. Datta, M. Gradinariu (Eds.), *Stabilization, Safety, and Security of Distributed Systems*. XVII, 590 pages. 2006.

Vol. 4279: N. Kobayashi (Ed.), *Programming Languages and Systems*. XI, 423 pages. 2006.

Vol. 4278: R. Meersman, Z. Tari, P. Herrero (Eds.), *On the Move to Meaningful Internet Systems 2006: OTM 2006 Workshops, Part II*. XLV, 1004 pages. 2006.

Vol. 4277: R. Meersman, Z. Tari, P. Herrero (Eds.), *On the Move to Meaningful Internet Systems 2006: OTM 2006 Workshops, Part I*. XLV, 1009 pages. 2006.

Vol. 4276: R. Meersman, Z. Tari (Eds.), *On the Move to Meaningful Internet Systems 2006: CoopIS, DOA, GADA, and ODBASE, Part II*. XXXII, 752 pages. 2006.

Vol. 4275: R. Meersman, Z. Tari (Eds.), *On the Move to Meaningful Internet Systems 2006: CoopIS, DOA, GADA, and ODBASE, Part I*. XXXI, 1115 pages. 2006.

Vol. 4274: Q. Huo, B. Ma, E.-S. Chng, H. Li (Eds.), *Chinese Spoken Language Processing*. XXXX, 8000 pages. 2006. (Sublibrary LNAI).

Vol. 4273: I. Cruz, S. Decker, D. Allemang, C. Preist, D. Schwabe, P. Mika, M. Uschold, L. Aroyo (Eds.), *The Semantic Web – ISWC 2006*. XXIV, 1001 pages. 2006.

- Vol. 4272: P. Havinga, M. Lijding, N. Meratnia, M. Wegdam (Eds.), Smart Sensing and Context. XI, 267 pages. 2006.
- Vol. 4271: F.V. Fomin (Ed.), Graph-Theoretic Concepts in Computer Science. XIII, 358 pages. 2006.
- Vol. 4270: H. Zha, Z. Pan, H. Thwaites, A.C. Addison, M. Forte (Eds.), Interactive Technologies and Sociotechnical Systems. XVI, 547 pages. 2006.
- Vol. 4269: R. State, S. van der Meer, D. O'Sullivan, T. Pfeifer (Eds.), Large Scale Management of Distributed Systems. XIII, 282 pages. 2006.
- Vol. 4268: G. Parr, D. Malone, M. Ó Foghlú (Eds.), Autonomic Principles of IP Operations and Management. XIII, 237 pages. 2006.
- Vol. 4267: A. Helmy, B. Jennings, L. Murphy, T. Pfeifer (Eds.), Autonomic Management of Mobile Multimedia Services. XIII, 257 pages. 2006.
- Vol. 4266: H. Yoshiura, K. Sakurai, K. Rannenber, Y. Murayama, S. Kawamura (Eds.), Advances in Information and Computer Security. XIII, 438 pages. 2006.
- Vol. 4265: L. Todorovski, N. Lavrač, K.P. Jantke (Eds.), Discovery Science. XIV, 384 pages. 2006. (Sublibrary LNAI).
- Vol. 4264: J.L. Balcázar, P.M. Long, F. Stephan (Eds.), Algorithmic Learning Theory. XIII, 393 pages. 2006. (Sublibrary LNAI).
- Vol. 4263: A. Levi, E. Savas, H. Yenigün, S. Balcisoy, Y. Saygin (Eds.), Computer and Information Sciences – ISCIS 2006. XXIII, 1084 pages. 2006.
- Vol. 4262: K. Havelund, M. N'ou'nez, G. Ro'csu, B. Wolff (Eds.), Formal Approaches to Software Testing and Runtime Verification. XXXXXX, 24555555 pages. 2006.
- Vol. 4261: Y. Zhuang, S. Yang, Y. Rui, Q. He (Eds.), Advances in Multimedia Information Processing - PCM 2006. XXII, 1040 pages. 2006.
- Vol. 4260: Z. Liu, J. He (Eds.), Formal Methods and Software Engineering. XII, 778 pages. 2006.
- Vol. 4259: S. Greco, Y. Hata, S. Hirano, M. Inuiguchi, S. Miyamoto, H.S. Nguyen, R. Stowinski (Eds.), Rough Sets and Current Trends in Computing. XXII, 951 pages. 2006. (Sublibrary LNAI).
- Vol. 4257: I. Richardson, P. Runeson, R. Messnarz (Eds.), Software Process Improvement. XI, 219 pages. 2006.
- Vol. 4256: L. Feng, G. Wang, C. Zeng, R. Huang (Eds.), Web Information Systems – WISE 2006 Workshops. XIV, 320 pages. 2006.
- Vol. 4255: K. Aberer, Z. Peng, E.A. Rundensteiner, Y. Zhang, X. Li (Eds.), Web Information Systems – WISE 2006. XIV, 563 pages. 2006.
- Vol. 4254: T. Grust, H. Höpfner, A. Illarramendi, S. Jablonski, M. Mesiti, S. Müller, P.-L. Patranjan, K.-U. Sattler, M. Spiliopoulou, J. Wijsen (Eds.), Current Trends in Database Technology – EDBT 2006. XXXI, 932 pages. 2006.
- Vol. 4253: B. Gabrys, R.J. Howlett, L.C. Jain (Eds.), Knowledge-Based Intelligent Information and Engineering Systems, Part III. XXXII, 1301 pages. 2006. (Sublibrary LNAI).
- Vol. 4252: B. Gabrys, R.J. Howlett, L.C. Jain (Eds.), Knowledge-Based Intelligent Information and Engineering Systems, Part II. XXXIII, 1335 pages. 2006. (Sublibrary LNAI).
- Vol. 4251: B. Gabrys, R.J. Howlett, L.C. Jain (Eds.), Knowledge-Based Intelligent Information and Engineering Systems, Part I. LXVI, 1297 pages. 2006. (Sublibrary LNAI).
- Vol. 4250: H.J. van den Herik, S.-C. Hsu, T.-s. Hsu, H.H.L.M. Donkers (Eds.), Advances in Computer Games. XIV, 273 pages. 2006.
- Vol. 4249: L. Goubin, M. Matsui (Eds.), Cryptographic Hardware and Embedded Systems - CHES 2006. XII, 462 pages. 2006.
- Vol. 4248: S. Staab, V. Svátek (Eds.), Managing Knowledge in a World of Networks. XIV, 400 pages. 2006. (Sublibrary LNAI).
- Vol. 4247: T.-D. Wang, X. Li, S.-H. Chen, X. Wang, H. Abbass, H. Iba, G. Chen, X. Yao (Eds.), Simulated Evolution and Learning. XXI, 940 pages. 2006.
- Vol. 4246: M. Hermann, A. Voronkov (Eds.), Logic for Programming, Artificial Intelligence, and Reasoning. XIII, 588 pages. 2006. (Sublibrary LNAI).
- Vol. 4245: A. Kuba, L.G. Nyúl, K. Palágyi (Eds.), Discrete Geometry for Computer Imagery. XIII, 688 pages. 2006.
- Vol. 4244: S. Spaccapietra (Ed.), Journal on Data Semantics VII. XI, 267 pages. 2006.
- Vol. 4243: T. Yakhno, E.J. Neuhold (Eds.), Advances in Information Systems. XIII, 420 pages. 2006.
- Vol. 4242: A. Rashid, M. Aksit (Eds.), Transactions on Aspect-Oriented Software Development II. IX, 289 pages. 2006.
- Vol. 4241: R.R. Beichel, M. Sonka (Eds.), Computer Vision Approaches to Medical Image Analysis. XI, 262 pages. 2006.
- Vol. 4239: H.Y. Youn, M. Kim, H. Morikawa (Eds.), Ubiquitous Computing Systems. XVI, 548 pages. 2006.
- Vol. 4238: Y.-T. Kim, M. Takano (Eds.), Management of Convergence Networks and Services. XVIII, 605 pages. 2006.
- Vol. 4237: H. Leitold, E. Markatos (Eds.), Communications and Multimedia Security. XII, 253 pages. 2006.
- Vol. 4236: L. Breveglieri, I. Koren, D. Naccache, J.-P. Seifert (Eds.), Fault Diagnosis and Tolerance in Cryptography. XIII, 253 pages. 2006.
- Vol. 4234: I. King, J. Wang, L. Chan, D. Wang (Eds.), Neural Information Processing, Part III. XXII, 1227 pages. 2006.
- Vol. 4233: I. King, J. Wang, L. Chan, D. Wang (Eds.), Neural Information Processing, Part II. XXII, 1203 pages. 2006.
- Vol. 4232: I. King, J. Wang, L. Chan, D. Wang (Eds.), Neural Information Processing, Part I. XLVI, 1153 pages. 2006.

~~¥803.00元~~ ¥479.00元

Table of Contents

Invited Papers

Shared-Variable Concurrency: A Proposal	1
<i>G�rard Boudol</i>	
Hennessy-Plotkin-Brookes Revisited	4
<i>Gordon Plotkin</i>	
Approximation Algorithms for 2-Stage Stochastic Optimization Problems	5
<i>Chaitanya Swamy, David B. Shmoys</i>	
The Number of Crossing Free Configurations on Finite Point Sets in the Plane	20
<i>Emo Welzl</i>	

Contributed Papers – Track A

Normal and Feature Approximations from Noisy Point Clouds	21
<i>Tamal K. Dey, Jian Sun</i>	
Coresets for Discrete Integration and Clustering	33
<i>Sariel Har-Peled</i>	
Self-assembling Classes of Shapes with a Minimum Number of Tiles, and in Optimal Time	45
<i>Florent Becker, Ivan Rapaport, �ric R�mila</i>	
One-Input-Face MPCVP Is Hard for L, But in LogDCFL	57
<i>Tanmoy Chakraborty, Samir Datta</i>	
Hardness of Approximation Results for the Problem of Finding the Stopping Distance in Tanner Graphs	69
<i>K. Murali Krishnan, L. Sunil Chandran</i>	
Multi-stack Boundary Labeling Problems	81
<i>Michael A. Bekos, Michael Kaufmann, Katerina Potika, Antonios Symvonis</i>	
Computing a Center-Transversal Line	93
<i>Pankaj K. Agarwal, Sergio Cabello, J. Antoni Sellar�s, Micha Sharir</i>	

On Obtaining Pseudorandomness from Error-Correcting Codes	105
<i>Shankar Kalyanaraman, Christopher Umans</i>	
Fast Edge Colorings with Fixed Number of Colors to Minimize Imbalance	117
<i>Gruia Calinescu, Michael J. Pelsmajer</i>	
Zero Error List-Decoding Capacity of the $q/(q - 1)$ Channel	129
<i>Sourav Chakraborty, Jaikumar Radhakrishnan, Nandakumar Raghunathan, Prashant Sasatte</i>	
Fast Exponential Algorithms for Maximum r -Regular Induced Subgraph Problems	139
<i>Sushmita Gupta, Venkatesh Raman, Saket Saurabh</i>	
Solving Connected Dominating Set Faster Than 2^n	152
<i>Fedor V. Fomin, Fabrizio Grandoni, Dieter Kratsch</i>	
Linear-Time Algorithms for Two Subtree-Comparison Problems on Phylogenetic Trees with Different Species	164
<i>Sun-Yuan Hsieh</i>	
Computationally Sound Symbolic Secrecy in the Presence of Hash Functions	176
<i>Véronique Cortier, Steve Kremer, Ralf Küsters, Bogdan Warinschi</i>	
Some Results on Average-Case Hardness Within the Polynomial Hierarchy	188
<i>A. Pavan, Rahul Santhanam, N.V. Vinodchandran</i>	
Unbiased Rounding of Rational Matrices	200
<i>Benjamin Doerr, Christian Klein</i>	

Contributed Papers – Track B

Rational Behaviour and Strategy Construction in Infinite Multiplayer Games	212
<i>Michael Ummels</i>	
The Anatomy of Innocence Revisited	224
<i>Russ Harmer, Olivier Laurent</i>	
Testing Probabilistic Equivalence Through Reinforcement Learning	236
<i>Josée Desharnais, François Laviolette, Sami Zhioua</i>	
On Decidability of LTL Model Checking for Process Rewrite Systems	248
<i>Laura Bozzelli, Mojmir Křetínský, Vojtěch Řehák, Jan Strejček</i>	
Monitoring of Real-Time Properties	260
<i>Andreas Bauer, Martin Leucker, Christian Schallhart</i>	

A Proof System for the Linear Time μ -Calculus	273
<i>Christian Dax, Martin Hofmann, Martin Lange</i>	
Tree Automata Make Ordinal Theory Easy	285
<i>Thierry Cachat</i>	
Context-Sensitive Dependency Pairs	297
<i>Beatriz Alarcón, Raúl Gutiérrez, Salvador Lucas</i>	
On Reduction Criteria for Probabilistic Reward Models	309
<i>Marcus Größer, Gethin Norman, Christel Baier, Frank Ciesinski, Marta Kwiatkowska, David Parker</i>	
Distributed Synthesis for Well-Connected Architectures	321
<i>Paul Gastin, Nathalie Sznajder, Marc Zeitoun</i>	
The Meaning of Ordered SOS	333
<i>MohammadReza Mousavi, Iain Phillips, Michel A. Reniers, Irek Ulidowski</i>	
Almost Optimal Strategies in One Clock Priced Timed Games	345
<i>Patricia Bouyer, Kim G. Larsen, Nicolas Markey, Jacob Ilum Rasmussen</i>	
Expressivity Properties of Boolean BI Through Relational Models	357
<i>Didier Galmiche, Dominique Larchey-Wendling</i>	
On Continuous Timed Automata with Input-Determined Guards	369
<i>Fabrice Chevalier, Deepak D'Souza, Pavithra Prabhakar</i>	
Safely Freezing LTL	381
<i>Ranko Lazić</i>	
Branching Pushdown Tree Automata	393
<i>Rajeev Alur, Swarat Chaudhuri</i>	
Validity Checking for Finite Automata over Linear Arithmetic Constraints	405
<i>Gary Wassermann, Zhendong Su</i>	
Game Semantics for Higher-Order Concurrency	417
<i>J. Laird</i>	
Author Index	429

Shared-Variable Concurrency: A Proposal

(Abstract)

G rard Boudol

INRIA Sophia Antipolis

In this talk I discuss the semantics of shared-variable concurrency, aka multi-threading. There are two well-known ways of managing concurrent threads: one either uses a *preemptive* or a *cooperative* scheduling discipline. In the former, a program, or more precisely its executable version, can be interrupted at any time during its execution by an external device, the scheduler, and the resources needed for execution are then given to another concurrent component for a while. This is perfect for executing concurrent *processes*, which do not share memory. In this case, the programmer does not have to care about the relative performance of the various processes in the system: this is the task of the scheduler. Unfortunately, it is very difficult to program multi-threaded applications with this model. The main difficulty is with *data races*, that is conflicting concurrent accesses to the memory. Although it is very easy to provide a formal “interleaving” semantics for preemptive multi-threading, this semantics usually does not coincide with what is actually implemented. In particular, the grain of atomicity is generally not preserved by the implementation, and a program may be time-sliced at some points of its execution which make no sense at the user level, and the consequence is that there is no clear semantics for the race conditions (see [14] for instance). It is therefore necessary to complement preemptive multi-threading with elaborate synchronization techniques, that require a real expertise from the programmer to be used (see [3]), and to design methods to analyze concurrent programs in order to make them “thread safe”, avoiding or detecting race conditions [1,6,11]. We shall not follow the preemptive approach in our proposal for shared-memory concurrency semantics.

In the *cooperative* programming model, a thread decides, by means of specific instructions (like *yield* for instance), when to leave its turn to another concurrent thread, and the scheduling is therefore distributed among the components. This model, in which there is no data race, has been advocated as a better model than the preemptive one for programming some modern, massively concurrent applications. However, this model also has its drawbacks, the main one being that if the active thread runs into an error, or raises an uncaught exception, or diverges, then the model is broken, in the sense that no other component will have a chance to execute. In particular, in cooperative programming, we have to avoid divergence in some way. Still, we need to be able to program non-terminating applications. Any server for instance should conceptually have an infinite life duration, and should not be programmed to stop after a while. Such a

server should not enter into an infinite loop however, it should rather be infinitely often waiting for a new request. In other words, in cooperative programming, programs *must be cooperative*, or *fair*, that is, they should be guaranteed to either terminate or suspend themselves infinitely often. Since we would also like to be able to reuse sequential code in a multi-threaded context, a challenge is: can we design a concurrency semantics that would be basically cooperative, in order to avoid data races, but where any thread – and in particular purely sequential code – is guaranteed to be fair?

Dealing with a higher-order imperative programming model à la ML for sequential code (some other choices are obviously possible), our proposal to solve the above mentioned challenge is to introduce a touch of preemptive scheduling into the cooperative model. The idea is very simple: it is to consider that every *recursive call* to a function should be regarded as a *suspensive* operation, that yields the scheduler for a while. We think that this is a natural and intuitive idea, since recursion appears to be the source of non-termination. Moreover, this does not introduce any data race. However, we then have to face a technical problem, which is that recursion, or more accurately non-termination, may occur in indirect ways in a ML-like, or, for that matter, a SCHEME-like language. Indeed, it is well-known that, on the one hand, recursion can be encoded in the pure λ -calculus, and, on the other hand, that recursion can also be encoded using circular reference (this is indeed the way it is implemented), as shown by Landin long ago [8]. The well-known method to recover from the first difficulty is to use a *type system*. In this talk I show that we can use a type and *effect* system [9] for dealing with the second difficulty.

To preclude circular references, we stratify the memory into *regions*, in such a way that functional values stored in a given region may only have a latent effect, such as reading a reference, in strictly “lower” regions. This allows us to define, by induction on the stratification, a *realizability interpretation* of the types and effects, for which the type and effect system is sound. From this we conclude that typable expressions can only diverge if they perform an infinite number of recursive calls, that is, if they suspend infinitely often, in our mixed cooperative/preemptive model. This termination argument is quite classical and general (see [2,5]): it was first used by Tait in [15], under the name of “convertibility”, and then by Girard (see [5]) with his “candidats de réductibilité”, and by Tait again [16] who called it the “realizability” technique, since it relies on an interpretation of types closely related to Kleene’s original recursive realizability interpretation [7]. The realizability interpretation is also a special case of a “logical relation” [10,13]. As far as I can see, this technique has not been previously used for higher-order imperative languages: the work that is technically the closest to ours, and which was our main source of inspiration, is the one by Pitts and Stark [12], but their logical relation (intended to provide a means to prove observational equivalence of programs, not to prove termination) is restricted to a language where the memory can only contain values of basic types. In the talk I mention another application of typing termination, namely in typing “termination leaks” in information flow control (see [4]).

References

1. M. ABADI, C. FLANAGAN, S.N. FREUND, *Types for safe locking: static race detection for Java*, ACM TOPLAS Vol. 28 No. 2 (2006) 207-255.
2. H. BARENDREGT, *Lambda Calculi with Types*, in Handbook of Logic in Computer Science, Vol. 2 (S. Abramsky, Dov M. Gabbay & T.S.E. Maibaum, Eds.), Oxford University Press (1992) 117-309.
3. A.D. BIRREL, *An introduction to programming with threads*, SRC Report 35, DEC (1989).
4. G. BOUDOL, *On typing information flow*, Intern. Coll. on Theoretical Aspects of Computing, Lecture Notes in Comput. Sci. 3722 (2005) 366-380.
5. J.-Y. GIRARD, Y. LAFONT, P. TAYLOR, *Proofs and Types*, Cambridge Tracts in Theoretical Computer Science 7, Cambridge University Press (1989).
6. D. GROSSMAN, *Type-safe multithreading in Cyclone*, TLDI'03 (2003) 13-25.
7. S.C. KLEENE, *On the interpretation of intuitionistic number theory*, J. of Symbolic Logic, Vol. 10 (1945) 109-124.
8. P.J. LANDIN, *The mechanical evaluation of expressions*, Computer Journal Vol. 6 (1964) 308-320.
9. J.M. LUCASSEN, D.K. GIFFORD, *Polymorphic effect systems*, POPL'88 (1988) 47-57.
10. J.C. MITCHELL, *Foundations for Programming Languages*, MIT Press (1996).
11. P. O'HEARN, *Resources, concurrency and local reasoning*, CONCUR'04, Lecture Notes in Comput. Sci. 3170 (2004) 49-67.
12. A. PITTS, I. STARK, *Operational reasoning for functions with local state*, in Higher-Order Operational Techniques in Semantics (A. Gordon and A. Pitts, Eds), Publications of the Newton Institute, Cambridge Univ. Press (1998) 227-273.
13. G. PLOTKIN, *Lambda-definability and logical relations*, Memo SAI-RM-4, University of Edinburgh (1973).
14. J.C. REYNOLDS, *Toward a grainless semantics for shared-variable concurrency*, FST-TCS'04, Lecture Notes in Comput. Sci. 3328 (2004) 35-48.
15. W. TAIT, *Intensional interpretations of functionals of finite type I*, J. of Symbolic Logic 32 (1967) 198-212.
16. W. TAIT, *A realizability interpretation of the theory of species*, Logic Colloquium, Lecture Notes in Mathematics 453 (1975) 240-251.