


James H. Anderson
Giuseppe Prencipe
Roger Wattenhofer (Eds.)

LNCS 3974

Principles of Distributed Systems

9th International Conference, OPODIS 2005
Pisa, Italy, December 2005
Revised Selected Papers



Springer

TP302.1-53

061

2005

James H. Anderson Giuseppe Prencipe
Roger Wattenhofer (Eds.)

Principles of Distributed Systems

9th International Conference, OPODIS 2005
Pisa, Italy, December 12-14, 2005
Revised Selected Papers



Springer



E2007001415

Volume Editors

James H. Anderson
University of North Carolina at Chapel Hill
Department of Computer Science
Chapel Hill, USA
E-mail: anderson@cs.unc.edu

Giuseppe Prencipe
Università di Pisa
Dipartimento di Informatica
Largo Bruno Pontecorvo 3, 56100 Pisa, Italy
E-mail: prencipe@di.unipi.it

Roger Wattenhofer
ETH Zurich
Computer Engineering and Networks Laboratory
8092 Zurich, Switzerland
E-mail: wattenhofer@tik.ee.ethz.ch

Library of Congress Control Number: 2006938989

CR Subject Classification (1998): C.2.4, D.1.3, D.2.7, D.2.12, D.4.7, C.3

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

ISSN 0302-9743
ISBN-10 3-540-36321-1 Springer Berlin Heidelberg New York
ISBN-13 978-3-540-36321-7 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: 11795490 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

The 9th International Conference on Principles of Distributed Systems (OPODIS 2005) was held during December 12–14, 2005 in Pisa, Italy. It continued a tradition of successful conferences with friendly and pleasant atmospheres. Previous OPODIS conferences were held in Chantilly (1997), Amiens (1998), Hanoi (1999), Paris (2000), Mexico (2001), Reims (2002), La Martinique (2003), and Grenoble (2004).

The OPODIS conference constitutes an open forum for the exchange of state-of-the-art knowledge on distributed computing and systems among researchers from around the world. Following the tradition of the previous events, the 2005 program was composed of high-quality contributed papers by experts of international caliber in this scientific area. Papers were sought soliciting original research contributions to the theory, specifications, design and implementation of distributed systems, including:

- communication and synchronization protocols
- distributed algorithms, multiprocessor algorithms
- distributed collaborative environments
- embedded systems
- fault-tolerance, reliability, availability
- grid and cluster computing
- location- and context-aware systems
- mobile computing and networks
- peer-to-peer systems, overlay networks
- performance analysis of distributed algorithms and systems
- real-time systems
- security issues in distributed computing and systems
- sensor networks
- specification verification and testing of distributed systems

This year, a particular focus was placed on real-time systems and wireless networks.

In response to the call for papers for OPODIS 2005, 109 papers in the above areas were submitted from 30 countries from around the world. Each paper was reviewed by at least three reviewers, and judged according to scientific and presentation quality, originality and relevance to the conference topics. The Program Committee selected 30 papers for presentation at the conference. In addition to the submitted technical papers, the program included two exciting invited talks, by David Peleg (Weizmann Institute of Science, Israel), and by Giorgio Buttazzo (Scuola Superiore S. Anna of Pisa). We are grateful that these two distinguished experts accepted our invitation to share with us their views on various aspects of the field.

It is impossible to organize a successful program without the help of many individuals. We would like to express our appreciation to the authors of the submitted papers, the Program Committee members and the external referees. We would also like to thank the OPODIS Steering Committee members, in particular the chairman Philippos Tsigas, who supervised and supported the continuation of this event. We owe special thanks to the Organizing Committee chair, Giuseppe Prencipe (Università di Pisa, Italy), the publicity chair, Thomas Mosciro (ETH Zurich, Switzerland), and Andreas Wetzel (ETH Zurich, Switzerland) for his assistance with the electronic submission and reviewing system.

December 2005

James H. Anderson and Roger Wattenhofer

Organization

OPODIS 2005 was organized by the Dipartimento di Informatica, Università di Pisa, Italy.

Organizing Institutes



Organizing Committee

Vincenzo Gervasi	Università di Pisa, Italy
Thomas Moscibroda	ETH Zurich, Switzerland (Publicity)
Giuseppe Prencipe	Università di Pisa, Italy (Chair)
Andreas Wetzel	ETH Zurich, Switzerland (Submissions)

Steering Committee

Alain Bui	Université de Reims Champagne-Ardenne, France
Marc Bui	Université de Paris 8/EPHE, France
Hacène Fouchal	Université Antilles-Guyane, France
Roberto Gomez	CEM-ITESM, Mexico
Nicola Santoro	Carleton University, Canada
Philippas Tsigas	University of Chalmers, Sweden (Chair)

Program Committee

Tarek Abdelzaher	University of Virginia, USA
Luis Almeida	Universidade de Aveiro, Portugal
James Anderson	University of North Carolina at Chapel Hill, USA (Co-chair)
Anish Arora	Ohio State University, USA
Ted Baker	Florida State University, USA
Joffroy Beauquier	Université de Paris Sud, France

Riccardo Bettati	Texas A&M University, USA
Costas Busch	Rensselaer Polytechnic Institute, USA
Christian Cachin	IBM Zurich Research, Switzerland
Danny Dolev	Hebrew University, Israel
Shlomi Dolev	Ben-Gurion University, Israel
Panagiota Fatourou	University of Ioannina, Greece
Paola Flocchini	Ottawa University, Canada
Gerhard Fohler	Malardalen University, Sweden
Hacene Fouchal	Université Antilles-Guyane
Shelby Funk	University of Georgia, USA
Mohamed Gouda	The University of Texas at Austin, USA
Maurice Herlihy	Brown University, USA
Teruo Higashino	Osaka University, Japan
Danny Krizanc	Wesleyan University, USA
Xiang-Yang Li	Illinois Institute of Technology, USA
Giuseppe Lipari	Scuola Superiore S. Anna Pisa, Italy
Giuseppe Prencipe	Università di Pisa, Italy
Mike Reiter	Carnegie Mellon University, USA
Ulrich Schmid	TU Vienna, Austria
Alex A. Shvartsman	University of Connecticut, USA
Srikanta Tirthapura	Iowa State University, USA
Roger Wattenhofer	ETH Zurich, Switzerland (Co-chair)
Masafumi Yamashita	Kyushu University, Japan

Referees

Ittai Abraham	Shantanu Das	Philippe Hunel
Sudhir Aggarwal	Jean-Domin. Decotignie	Martin Hutle
José Alberto	Gianna M. Del Corso	Damir Isovich
Emilio Ancillotti	Carole Delporte	Gian Paolo Jesi
Elliot Anshelevich	Yefim Dinitz	Colette Johnen
Thibaut Bernard	Stefan Dobrev	Eunjin Jung
Pascal Berthommé	Tullio Facchinetti	Koushik Kar
Marin Bertier	Hugues Fauconnier	Rastislav Kralovic
Marko Bertogna	Joaquim Ferreira	Pierre Lemarinier
Martin Biely	Hen Fitoussi	Antonio Mancina
Enrico Bini	Paola Flocchini	Ernesto Martins
Maurizio Bonuccelli	Pedro Fonseca	Alessia Milani
Olga Brukman	Vincenzo Gervasi	Florent Nolot
Bernadette Charron-Bost	Seth Gilbert	Linda Pagli
Soohyun Cho	Kartik Gopalan	Paulo Pedreiras
Young-ri Choi	Roberto Grossi	Evaggelia Pitoura
Michele Cirinei	Rachid Guerraoui	Matteo Repeti
Massimo Coppola	Yinnon Haviv	Laurent Rosaz
Fernanda Coutinho	Thomas Héroult	Eric Ruppert

Sébastien Tixeul
Nir Tzachar
Akira Uchiyama
Takaaki Umedu
Paolo Valente

Nicolas Vidot
Shengquan Wang
Bettina Weiss
Josef Widder
Reuven Yagel

Hirozumi Yamaguchi
Shen Zhenhui

Sponsoring Institution



Università di Pisa

Lecture Notes in Computer Science

For information about Vols. 1–4263

please contact your bookseller or Springer

Vol. 4377: M. Abe (Ed.), Topics in Cryptology – CT-RSA 2007. XI, 403 pages. 2006.

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. 4361: H.J. Hoogeboom, G. Păun, G. Rozenberg, A. Salomaa (Eds.), Membrane Computing. IX, 555 pages. 2006.

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. 4353: T. Schwentick, D. Suciu (Eds.), Database Theory – ICDT 2007. XI, 419 pages. 2006.

Vol. 4348: S.T. Taft, R.A. Duff, R.L. Brukardt, E. Ploederer, 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. 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. Orlowska, G. Schmidt, M. Roubens (Eds.), Theory and Application 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. 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. 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. 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. 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. 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.
- Vol. 4300: Y.Q. Shi (Ed.), Transactions on Data Hiding and Multimedia Security I. IX, 139 pages. 2006.
- Vol. 4299: S. Renals, S. Bengio, J.G. Fiscus (Eds.), Machine Learning for Multimodal Interaction. XII, 470 pages. 2006.
- Vol. 4297: Y. Robert, M. Parashar, R. Badrinath, V.K. Prasanna (Eds.), High Performance Computing - HiPC 2006. XXIV, 642 pages. 2006.
- Vol. 4296: M.S. Rhee, B. Lee (Eds.), Information Security and Cryptology - ICISC 2006. XIII, 358 pages. 2006.
- Vol. 4295: J.D. Carswell, T. Tezuka (Eds.), Web and Wireless Geographical Information Systems. XI, 269 pages. 2006.
- Vol. 4294: A. Dan, W. Lamersdorf (Eds.), Service-Oriented Computing - ICSOC 2006. XIX, 653 pages. 2006.
- Vol. 4293: A. Gelbukh, C.A. Reyes-Garcia (Eds.), MICAI 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. Mladenici, 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. XX, 766 pages. 2006.
- Vol. 4287: C. Mao, T. Yokomori (Eds.), DNA Computing. XII, 440 pages. 2006.
- Vol. 4286: P. Spirakis, M. Mavronicolas, S. Kontogiannis (Eds.), Internet and Network Economics. XI, 401 pages. 2006.
- Vol. 4285: Y. Matsumoto, R. Sproat, K.-F. Wong, M. Zhang (Eds.), Computer Processing of Oriental Languages. XVII, 544 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. XXIV, 805 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. Weggdam (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).

¥513.00元

Table of Contents

Invited Talk 1

Distributed Algorithms for Systems of Autonomous Mobile Robots <i>David Peleg</i>	1
--	---

Invited Talk 2

Real-Time Issues in Mobile Wireless Networks <i>Giorgio Buttazzo</i>	2
---	---

Session 1: Nonblocking Synchronization

A Lazy Concurrent List-Based Set Algorithm <i>Steve Heller, Maurice Herlihy, Victor Luchangco, Mark Moir, William N. Scherer III, Nir Shavit</i>	3
Efficiently Implementing a Large Number of LL/SC Objects <i>Prasad Jayanti, Srdjan Petrovic</i>	17
Can Memory Be Used Adaptively by Uniform Algorithms? <i>Burkhard Englert, Darin Goldstein</i>	32
Randomized Wait-Free Consensus Using an Atomicity Assumption <i>Ling Cheung</i>	47

Session 2: Fault-Tolerant Broadcast and Consensus

Optimal Randomized Fair Exchange with Secret Shared Coins <i>Felix Freiling, Maurice Herlihy, Lucia Draque Penso</i>	61
Two Abstractions for Implementing Atomic Objects in Dynamic Systems <i>Roy Friedman, Michel Raynal, Corentin Travers</i>	73
Parsimonious Asynchronous Byzantine-Fault-Tolerant Atomic Broadcast <i>HariGovind V. Ramasamy, Christian Cachin</i>	88

Session 3: Self-stabilizing Systems

Self-stabilizing Population Protocols
Dana Angluin, James Aspnes, Michael J. Fischer, Hong Jiang 103

A Self-stabilizing Link-Coloring Protocol Resilient to Unbounded Byzantine Faults in Arbitrary Networks
Toshimitsu Masuzawa, Sébastien Tixeuil 118

Timed Virtual Stationary Automata for Mobile Networks
Shlomi Dolev, Seth Gilbert, Limor Lahiani, Nancy Lynch, Tina Nolte 130

Asynchronous and Fully Self-stabilizing Time-Adaptive Majority Consensus
Janna Burman, Ted Herman, Shay Kutten, Boaz Patt-Shamir 146

Session 4: Peer-to-Peer Systems and Collaborative Environments

Stable Predicate Detection in Dynamic Systems
Donald Darling, Jean Mayo, Xinli Wang 161

MTcast: Robust and Efficient P2P-Based Video Delivery for Heterogeneous Users
Tao Sun, Morihiko Tamai, Keiichi Yasumoto, Naoki Shibata, Minoru Ito, Masaaki Mori 176

Towards a Theory of Self-organization
Emmanuelle Anceaume, Xavier Défago, Maria Gradinariu, Matthieu Roy 191

Node Discovery in Networks
Kishori M. Konwar, Dariusz Kowalski, Alex A. Shvartsman 206

Session 5: Sensor Networks and Mobile Computing

Optimal Clock Synchronization Under Energy Constraints in Wireless Ad-Hoc Networks
Hagit Attiya, David Hay, Jennifer L. Welch 221

Half-Space Proximal: A New Local Test for Extracting a Bounded Dilation Spanner of a Unit Disk Graph
Edgar Chavez, Stefan Dobrev, Evangelos Kranakis, Jaroslav Opatrny, Ladislav Stacho, Héctor Tejada, Jorge Urrutia 235

A State-Based Model of Sensor Protocols <i>Mohamed G. Gouda, Young-ri Choi</i>	246
---	-----

Session 6: Security and Verification

Approximation Bounds for Black Hole Search Problems <i>Ralf Klasing, Euripides Markou, Tomasz Radzik, Fabiano Sarracco</i>	261
Revising UNITY Programs: Possibilities and Limitations <i>Ali Ebnenasir, Sandeep S. Kulkarni, Borzoo Bonakdarpour</i>	275

Session 7: Real-Time Systems

The Partitioned, Static-Priority Scheduling of Sporadic Real-Time Tasks with Constrained Deadlines on Multiprocessor Platforms <i>Nathan Fisher, Sanjoy Baruah</i>	291
New Schedulability Tests for Real-Time Task Sets Scheduled by Deadline Monotonic on Multiprocessors <i>Marko Bertogna, Michele Cirinei, Giuseppe Lipari</i>	306
Static-Priority Scheduling of Sporadic Messages on a Wireless Channel <i>Björn Andersson, Eduardo Tovar</i>	322
Implementing Reliable Distributed Real-Time Systems with the Θ -Model <i>Jean-François Hermant, Josef Widder</i>	334

Session 8: Peer-to-Peer Systems

Reconfigurable Distributed Storage for Dynamic Networks <i>Gregory Chockler, Seth Gilbert, Vincent Gramoli, Peter M. Musial, Alex A. Shvartsman</i>	351
Skip B-Trees <i>Ittai Abraham, James Aspnes, Jian Yuan</i>	366
Bounding Communication Cost in Dynamic Load Balancing of Distributed Hash Tables <i>Marcin Bienkowski, Mirosław Korzeniowski</i>	381

Distributed Algorithms for Systems of Autonomous Mobile Robots

David Peleg*

Department of Computer Science, The Weizmann Institute of Science, Rehovot, Israel
david.peleg@weizmann.ac.il

Over the last five decades, mobile robots have been the focus of extensive research and development activities, with numerous applications for industrial tasks, military operations, search and rescue missions and space exploration, as well as some home applications.

Systems consisting of a group of autonomously operating mobile robots (sometimes referred to as *robot swarms*) have attracted considerable interest throughout the past twenty years, due to their potential for providing flexible, low-cost solutions in hazardous situations (e.g., military operations, toxic environments or fire fighting). The idea is to deal with such applications using swarms consisting of many small and simple robots, with very limited capabilities (e.g., low energy sources, limited communication means and weak processors). The use of tiny, functionally simple and cheap robots may make it acceptable to lose some of the robots, so long as the team manages to achieve its collective goals.

The main research efforts invested so far in mobile robots focused on the main engineering aspects of providing physical functionalities. Nevertheless, it seems clear that the design of very large robots swarms makes it essential to reconsider also control and coordination issues. For instance, managing the movements of a robot swarm involves new and interesting algorithmic problems due to the need to coordinate the movements of the individual robots and avoid collisions and over-crowding. Coordination tasks studied so far in the literature include gathering a robot swarm to a single point, pattern formation, flocking (or following a leader), partitioning, spreading and searching.

Most existing experimental settings of robot swarms involve small swarms (of, say, up to a dozen robots), which allow centralized control. However, future robot swarms, consisting of tens of thousands of robots, can no longer be controlled centrally in an efficient manner, and it seems that certain tasks may need to be managed by distributed protocols. Indeed, there have been a number of recent studies on distributed coordination and control protocols for robot swarms. From the point of view of the community of distributed algorithms and systems, this presents an interesting new distributed model that differs in a number of key aspects from the traditional models, and raises some intriguing research problems.

The talk will review this exciting research area, present some of the main problems and issues raised by it, and discuss directions for future study.

* Supported by the Israel Science Foundation (grant No. 693/04).

Real-Time Issues in Mobile Wireless Networks

Giorgio Buttazzo

Scuola Superiore Sant'Anna, Pisa

The use of cooperating mobile robots is requested in an increasing number of application domains, including civil protection, surveillance, environmental monitoring, under-water exploration, and space missions. In most of these applications, the robot units are required to acquire sensory information, localize themselves in the environment, plan trajectories, avoid obstacles, and cooperate with the other robots to reach a common goal.

To achieve such objectives, the development of a team of cooperating robots poses several interesting problems from a research point of view, such as the real-time execution of acquisition and control processes, the efficient management of computational resources, the software control of energy consumption, the real-time communication protocols on wireless networks, and the development of distributed agreement algorithms for reaching a consensus in collective decisions. Moreover, small mobile robots are often controlled by micro-controllers having low computational power and limited resources, hence satisfying timing constraints requires the use of efficient operating systems and algorithms that can guarantee a predictable behavior both in normal and overload conditions.

This talk will present some of the most challenging problems to be solved in order to support the development of mobile wireless networks of cooperating robots.

A Lazy Concurrent List-Based Set Algorithm

Steve Heller¹, Maurice Herlihy², Victor Luchangco¹, Mark Moir¹,
William N. Scherer III³, and Nir Shavit¹

¹ Sun Microsystems Laboratories

² Brown University

³ University of Rochester

Abstract. List-based implementations of sets are a fundamental building block of many concurrent algorithms. A skiplist based on the lock-free list-based set algorithm of Michael will be included in the JavaTM Concurrency Package of *JDK 1.6.0*. However, Michael's lock-free algorithm has several drawbacks, most notably that it requires all list traversal operations, including membership tests, to perform cleanup operations of logically removed nodes, and that it uses the equivalent of an atomically markable reference, a pointer that can be atomically “marked,” which is expensive in some languages and unavailable in others.

We present a novel “lazy” list-based implementation of a concurrent set object. It is based on an optimistic locking scheme for inserts and removes, eliminating the need to use the equivalent of an atomically markable reference. It also has a novel wait-free membership test operation (as opposed to Michael's lock-free one) that does not need to perform cleanup operations and is more efficient than that of all previous algorithms.

Empirical testing shows that the new lazy-list algorithm consistently outperforms all known algorithms, including Michael's lock-free algorithm, throughout the concurrency range. At high load, with 90% membership tests, the lazy algorithm is more than twice as fast as Michael's. This is encouraging given that typical search structure usage patterns include around 90% membership tests. By replacing the lock-free membership test of Michael's algorithm with our new wait-free one, we achieve an algorithm that slightly outperforms our new lazy-list (though it may not be as efficient in other contexts as it uses Java's RTTI mechanism to create pointers that can be atomically marked).

1 Introduction

Lists are a fundamental building block for concurrent data structures, both in their own right, and as the basis for many types of search and dictionary data types [12]. We consider three kinds of list operations: inserting a list entry, removing a list entry, and testing whether an entry is in the list.

This paper introduces the *lazy list*, a simple new concurrent *list-based set* algorithm with a number of novel concurrency-related properties. To explain the novel aspects of lazy lists, we start with an overview of different ways to synchronize lists. *Coarse-grained* locking, which uses a single lock to protect the entire