

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, UK

Takeo Kanade

Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler

University of Surrey, Guildford, UK

Jon M. Kleinberg

Cornell University, Ithaca, NY, USA

Friedemann Mattern

ETH Zurich, Switzerland

John C. Mitchell

Stanford University, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

Oscar Nierstrasz

University of Bern, Switzerland

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

University of Dortmund, Germany

Madhu Sudan

Massachusetts Institute of Technology, MA, USA

Demetri Terzopoulos

New York University, NY, USA

Doug Tygar

University of California, Berkeley, CA, USA

Moshe Y. Vardi

Rice University, Houston, TX, USA

Gerhard Weikum

Max-Planck Institute of Computer Science, Saarbruecken, Germany

Rachid Guerraoui (Ed.)

Distributed Computing

18th International Conference, DISC 2004
Amsterdam, The Netherlands, October 4-7, 2004
Proceedings

Volume Editor

Rachid Guerraoui
EPFL, I&C – LPD, Bât. IN
1015 Lausanne, Switzerland
E-mail: rachid.guerraoui@epfl.ch

Library of Congress Control Number: 2004112514

CR Subject Classification (1998): C.2.4, C.2.2, F.2.2, D.1.3, F.1.1, D.4.4-5

ISSN 0302-9743

ISBN 3-540-23306-7 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

springeronline.com

© Springer-Verlag Berlin Heidelberg 2004
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Olgun Computergrafik
Printed on acid-free paper SPIN: 11328841 06/3142 5 4 3 2 1 0

Preface

DISC, the International Symposium on Distributed Computing, is an annual conference for the presentation of research on the theory, design, analysis, implementation, and application of distributed systems and network. DISC 2004 was held on October 4–7, 2004, in Amsterdam, The Netherlands.

There were 142 papers submitted to DISC this year. These were read and evaluated by the program committee members, assisted by external reviewers. The quality of submissions was high and we were unable to accept many deserving papers. Thirty one papers were selected at the program committee meeting in Lausanne to be included in these proceedings. The proceedings include an extended abstract of the invited talk by Ueli Maurer. In addition, they include a eulogy for Peter Ruzicka by Shmuel Zaks.

The Best Student Paper Award was split and given to two papers: the paper “Efficient Adaptive Collect Using Randomization”, co-authored by Hagit Attiya, Fabian Kuhn, Mirjam Wattenhofer and Roger Wattenhofer, and the paper “Coupling and Self-stabilization”, co-authored by Laurent Fribourg, Stephane Messika and Claudine Picaronny.

The support of the CWI and EPFL is gratefully acknowledged. The review process and the preparation of this volume were done using CyberChairPRO. I also thank Sebastien Baehni and Sidath Handurukande for their crucial help with these matters.

August 2004

Rachid Guerraoui

Peter Ruzicka 1947–2003



Peter died on Sunday, October 5, 2003, at the age of 56, after a short disease. He was a Professor of Informatics at the Faculty of Mathematics, Physics and Informatics in Comenius University, Bratislava, Slovakia. Those of us who knew him through DISC and other occasions mourn his death and cherish his memory. These words are written in Peter's honor, as a farewell to a true colleague of our community. Peter's death came as a shock to everyone who knew him. His tragic death is an immense loss to his wife Marta and his daughter Kristina. I hope these words bring some comfort to them.

Peter worked in several research areas of theoretical computer science throughout his long career. His earlier works cover topics that include formal languages, unification, graph pebbling and others. Most of his research works since the early 1990s were in the areas of communication networks and distributed computing, and these works connected Peter to our community. These works include studies of complexity issues for various problems and models of interval routing, and for various problems and topics in distributed computing (like oriented and unoriented networks) and in interconnection networks. Peter was a true hard worker, always choosing difficult and challenging problems. His works always use deep and interesting mathematical tools, and are always presented very clearly and precisely. Peter was very persistent in setting up his research goals and following them. His papers contain many original ideas, and cover all relevant aspects comprehensively. Peter did some of these works together with Stefan Dobrev and Rastislav Kralovic, whom he devotedly supervised as his Ph.D. students in this research area, and he did joint works with many other colleagues.

Peter was a regular participant in many conferences in this area, including DISC and SIROCCO; he was on the program committees of DISC 1998 and 2000 and SIROCCO 1999. He participated and was very actively involved in other conferences, including ICALP, SOFSEM and MFCS. These activities included co-chairing MFCS 1994, chairing MFCS 1997 and chairing SOFSEM 2001, and being on the program committees of MFCS 2000 and SOFSEM 1997. In 1998, when I was the chair of DISC, I approached Peter, suggesting that he organize DISC in Bratislava, thus bringing DISC, for its first time, to an East-European country. Peter accepted this challenge enthusiastically, and with no hesitation whatsoever. Together with his local team, he then devoted the fol-

lowing year to the organization of DISC 1999 in Bratislava. I was in touch with Peter throughout that year and noticed his conscientious efforts to ensure that the conference would be successful. Indeed, this conference proved to be a great success, and Peter and his team were consequently highly appreciated by the Steering Committee and the whole audience.

I knew that Peter was sick, and I contacted him by email before I left home to attend DISC 2003 in Sorrento, at the end of September. When I returned home, a message waited for me, in which his wife Marta informed me about his death. It turned out that on the same days that we all met in Sorrento, Peter was fighting for his life. He died on Sunday, October 5, 2003, when we all made our way back home from the conference. I learnt from his wife that to the very end Peter proved very strong and was full of hope.

Personally, I found Peter an excellent colleague. I spent a few weeks with Peter, during our common visits to each other's place. The time I spent with him was of the highest value to me. He was always open to discussions of new ideas, very skillful as a researcher, very friendly, and very patient and gentle. I will always cherish Peter and Marta's warm hospitality during my visits to Bratislava.

We will all miss a very precious member of our community. I will miss a dear colleague and friend.

Prof. Shmuel Zaks
Department of Computer Science
Technion, Israel

Organization

DISC, the International Symposium on Distributed Computing, is an annual forum for research presentations on all facets of distributed computing. The symposium was called the International Workshop on Distributed Algorithms (WDAG) from 1985 to 1997. DISC 2004 was organized in cooperation with the European Association for Theoretical Computer Science (EATCS).



Conference Chairs	Jaap-Henk Hoepman, University of Nijmegen Paul Vitányi, CWI and University of Amsterdam
Program Chair	Rachid Guerraoui, EPFL
Organizing Committee	Wilmy van Ojik, CWI

Program Committee

Mustaque Ahamad	Georgia Tech
Lorenzo Alvisi	UT Austin
James Anderson	University of North Carolina at Chapel Hill
Paul Attie	Northeastern University
Ozalp Babaoglu	University of Bologna
Carole Delporte	University of Paris VII
Shlomi Dolev	Ben-Gurion University
Pierre Fraigniaud	CNRS
Felix Gaertner	University of Aachen
Maurice Herlihy	Brown University
Nancy Lynch	MIT
Michael Merritt	AT&T
Achour Mostefaoui	IRISA
Mike Reiter	CMU
Robert van Renesse	Cornell University
Luis Rodrigues	University of Lisbon
Paul Spirakis	University of Patras
Philippas Tsigas	Chalmers University
Paul Vitányi	CWI and University of Amsterdam
Roman Vitenberg	UCSB

Referees

Maha Abdallah
 Ittai Abraham
 Uri Abraham
 Adnan Agbaria
 Divyakant Agrawal
 Marcos Aguilera
 Amitanand S. Aiyer
 Emmanuelle Anceaume
 Krzysztof Apt
 Luciana Arantes
 Filipe Araujo
 James Aspnes
 Hagit Attiya
 Yossi Azar
 R. Badrinath
 Sebastien Baehni
 Mahesh Balakrishnan
 Roberto Baldoni
 Amos Beimel
 Zinaida Benenson
 Mark Bickford
 Costas Busch
 Gregory Chockler
 Allen Clement
 Sylvie Delaët
 Yefim Dinitz
 Frederick Ducatelle
 Partha Dutta
 Rui Fan
 Panagiota Fatourou
 Hugues Fauconnier
 Dimitris Fotakis
 Roy Friedman
 Charles Fry
 Juan A. Garay
 Anders Gidenstam
 Garth Goodson
 Eric Goubault
 Maria Gradinariu
 Abishek Gutta
 Phuong Ha
 Martin Haenggi
 David Hales
 Sidath Handurukande

Mor Harchol-Balter
 Yinnon Haviv
 Jean-Michel Hélary
 Ted Herman
 Phuong Hoai Ha
 Michel Hurfin
 Amos Israeli
 Michael A. Jaeger
 Mark Jelasity
 Colette Johnen
 Alon Kama
 Ronen I. Kat
 Idit Keidar
 Anne-Marie Kermarrec
 Alex Kesselman
 Roger Khazan
 Boris Koldehofe
 Lei Kong
 Spyros Kontogiannis
 Petr Kouznetsov
 Danny Krizanc
 Limor Lahiani
 Kofi Laing
 Subramanian
 Lakshmanan
 Gerard Le Lann
 Ron Levy
 Feng Li
 Xiaozhou (Steve) Li
 Zvi Lotker
 Dahlia Malkhi
 Deepak Manohar
 Jean-Philippe Martin
 Marios Mavronicolas
 Maged Michael
 Martin Mink
 Hugo Miranda
 Mark Moir
 Alberto Montresor
 Gero Mühl
 Jeff Napper
 Priya Narasimhan
 Tina Nolte
 Michael Okun

Rui Oliveira
 Florin Oprea
 Vicky Papadopoulou
 Marina Papatriantafilou
 Simon Patarin
 Lucia Draque Penso
 Sara Tucci Piergiovanni
 Lexi Pimenidis
 Stefan Pleisch
 Bastian Pochon
 Michel Raynal
 Robert van Renessee
 Matthieu Roy
 Eric Ruppert
 Asad Samar
 Christian Scheideler
 Elad Schiller
 Pierre Sens
 Noam Singer
 Jesper Spring
 Maarten van Steen
 Frank Stomp
 Gideon Stupp
 Håkan Sundell
 Mikkel Thorup
 Sebastien Tixeuil
 Peter Triantafillou
 Nir Tzachar
 Sankarapandian
 Vijayaraghavan
 Jean-Marc Vincent
 Antonino Virgillito
 Ramesh Viswanath
 Hagen Völzer
 Jaksa Vuckovic
 Marko Vukolic
 Jay J. Wylie
 Reuven Yagel
 Praveen Yalagandula
 Christos Zaroliagis
 Zhiyuan Zhan
 Yi Zhang
 Lantian Zheng

Table of Contents

The Synchronous Condition-Based Consensus Hierarchy	1
<i>Achour Mostefaoui, Sergio Rajsbaum, and Michel Raynal</i>	
Synchronous Condition-Based Consensus Adapting to Input-Vector Legality	16
<i>Taisuke Izumi and Toshimitsu Masuzawa</i>	
Group-Solvability	30
<i>Eli Gafni</i>	
The Notion of Veto Number and the Respective Power of $\diamond\mathcal{P}$ and $\diamond\mathcal{S}$ to Solve One-Shot Agreement Problems	41
<i>Roy Friedman, Achour Mostefaoui, and Michel Raynal</i>	
The Black-White Bakery Algorithm and Related Bounded-Space, Adaptive, Local-Spinning and FIFO Algorithms	56
<i>Gadi Taubenfeld</i>	
Local-Spin Group Mutual Exclusion Algorithms	71
<i>Robert Danek and Vassos Hadzilacos</i>	
On Quorum Systems for Group Resources with Bounded Capacity	86
<i>Yuh-Jzer Joung</i>	
Bounded Version Vectors	102
<i>José Bacelar Almeida, Paulo Sérgio Almeida, and Carlos Baquero</i>	
An Optimistic Approach to Lock-Free FIFO Queues	117
<i>Edya Ladan-Mozes and Nir Shavit</i>	
A Single-Enqueuer Wait-Free Queue Implementation	132
<i>Matei David</i>	
Practical Lock-Free and Wait-Free LL/SC/VL Implementations Using 64-Bit CAS	144
<i>Maged M. Michael</i>	
Efficient Adaptive Collect Using Randomization	159
<i>Hagit Attiya, Fabian Kuhn, Mirjam Wattenhofer, and Roger Wattenhofer</i>	
Nonblocking Concurrent Data Structures with Condition Synchronization .	174
<i>William N. Scherer III and Michael L. Scott</i>	

Dynamic Memory ABP Work-Stealing	188
<i>Danny Hendler, Yossi Lev, and Nir Shavit</i>	
Coupling and Self-stabilization	201
<i>Laurent Fribourg, Stéphane Messika, and Claudine Picaronny</i>	
Optimal Randomized Self-stabilizing Mutual Exclusion on Synchronous Rings	216
<i>Philippe Duchon, Nicolas Hanusse, and Sébastien Tixeuil</i>	
Virtual Mobile Nodes for Mobile <i>Ad Hoc</i> Networks	230
<i>Shlomi Dolev, Seth Gilbert, Nancy A. Lynch, Elad Schiller, Alex A. Shvartsman, and Jennifer L. Welch</i>	
Contention-Free MAC Protocols for Wireless Sensor Networks	245
<i>Costas Busch, Malik Magdon-Ismaïl, Fikret Sivrikaya, and Bülent Yener</i>	
Relationships Between Broadcast and Shared Memory in Reliable Anonymous Distributed Systems	260
<i>James Aspnes, Faith Fich, and Eric Ruppert</i>	
A Local Algorithm for Ad Hoc Majority Voting via Charge Fusion	275
<i>Yitzhak Birk, Liran Liss, Assaf Schuster, and Ran Wolff</i>	
Message-Optimal and Latency-Optimal Termination Detection Algorithms for Arbitrary Topologies	290
<i>Neeraj Mittal, Subbarayan Venkatesan, and Sathya Peri</i>	
Routing with Improved Communication-Space Trade-Off	305
<i>Ittai Abraham, Cyril Gavoille, and Dahlia Malkhi</i>	
Active and Concurrent Topology Maintenance	320
<i>Xiaozhou Li, Jayadev Misra, and C. Greg Plaxton</i>	
Distributed Weighted Matching	335
<i>Mirjam Wattenhofer and Roger Wattenhofer</i>	
Exploiting Content Localities for Efficient Search in P2P Systems	349
<i>Lei Guo, Song Jiang, Li Xiao, and Xiaodong Zhang</i>	
Compact Routing Schemes for Bounded Tree-Length Graphs and for k -Chordal Graphs	365
<i>Yon Dourisboure</i>	
Towards a Theory of Consistency Primitives	379
<i>Ueli Maurer</i>	
Fault-Tolerant Storage in a Dynamic Environment	390
<i>Uri Nadav and Moni Naor</i>	

Non-skipping Timestamps for Byzantine Data Storage Systems	405
<i>Rida A. Bazzi and Yin Ding</i>	
Efficient Verification for Provably Secure Storage and Secret Sharing in Systems Where Half the Servers Are Faulty	420
<i>Rida A. Bazzi and Goran Konjevod</i>	
Optimal Dispersal of Certificate Chains	435
<i>Eunjin Jung, Ehab S. Elmallah, and Mohamed G. Gouda</i>	
On Byzantine Agreement over $(2, 3)$ -Uniform Hypergraphs	450
<i>D.V.S. Ravikant, V. Muthuramakrishnan, V. Srikanth, K. Srinathan, and C. Pandu Rangan</i>	
Author Index	465