

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

Andrzej Pelc (Ed.)

Distributed Computing

21st International Symposium, DISC 2007
Lemesos, Cyprus, September 24-26, 2007
Proceedings



Springer

Volume Editor

Andrzej Pelc
Département d'informatique
Université du Québec en Outaouais
Gatineau, Québec J8X 3X7, Canada
E-mail: pelc@uqo.ca

Library of Congress Control Number: 2007935053

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

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

ISSN 0302-9743
ISBN-10 3-540-75141-6 Springer Berlin Heidelberg New York
ISBN-13 978-3-540-75141-0 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: 12162741 06/3180 5 4 3 2 1 0

Preface

DISC, the International Symposium on Distributed Computing, is an annual forum for presentation of research on all aspects of distributed computing, including the theory, design, implementation and applications of distributed algorithms, systems and networks. The 21st edition of DISC was held during September 24–26, 2007, in Lemesos, Cyprus.

This volume of proceedings begins with abstracts of three invited talks. The keynote speakers of DISC 2007 were: Burkhard Monien from the University of Paderborn, Germany, David Peleg from The Weizmann Institute of Science, Israel, and Michel Raynal from IRISA, Université de Rennes, France.

There were 100 ten-page-long extended abstracts submitted to DISC this year and this volume contains 32 contributions selected by the Program Committee among these 100 submissions. Every submitted paper was read and evaluated by Program Committee members assisted by external reviewers. The final decisions regarding acceptance or rejection of each paper were made during the electronic Program Committee meeting held in June/July 2007. Revised and expanded versions of a few best selected papers will be considered for publication in a special issue of the journal *Distributed Computing*.

The Best Student Paper Award of DISC 2007 was awarded to David Eisenstat for the paper “Fast Robust Approximate Majority” coauthored with Dana Angluin and James Aspnes.

This volume of proceedings also contains nine two-page-long brief announcements (BA). These BAs present ongoing work or recent results whose full description is not yet ready; it is expected that full papers containing those results will soon appear in other conferences or journals. The main purpose of the BA track is to announce ongoing projects to the distributed computing community and to obtain feedback for the authors. Each BA was also read and evaluated by the Program Committee.

This volume concludes with a section devoted to the 20th anniversary of the DISC conferences that took place during DISC 2006, held September 18–20, 2006, in Stockholm, Sweden.

DISC 2007 was organized in cooperation with the University of Cyprus. The main sponsor of DISC 2007 was CYTA - Cyprus Telecommunications Authority. The support of the Cyprus Tourism Organisation, Microsoft (Cyprus) and COST Action 295 DYNAMO is also gratefully acknowledged.

The 2007 Edsger W. Dijkstra Prize in Distributed Computing

The 2007 Edsger W. Dijkstra Prize in Distributed Computing was presented at DISC 2007 for the paper “Consensus in the Presence of Partial Synchrony” by Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer, which appeared in the *Journal of the ACM* (Vol. 35, No. 2, April, 1988. pages 288–323). A preliminary version appeared in PODC 1984.

This paper introduces a number of practically motivated partial synchrony models that lie between the completely synchronous and the completely asynchronous models, and in which consensus is solvable. It gives practitioners the right tool for building fault-tolerant systems, and contributes to the understanding that safety can be maintained at all times, despite the impossibility of consensus, and progress is facilitated during periods of stability. These are the pillars on which every fault-tolerant system has been built for two decades. This includes academic projects such as Petal, Frangipani, and Boxwood, as well as real-life data centers, such as the Google file system.

In distributed systems, balancing the pragmatics of building software that works against the need for rigor is particularly difficult because of impossibility results such as the FLP theorem. The publication by Dwork, Lynch, and Stockmeyer was in many respects the first to suggest a path through this thicket, and has been enormously influential. It presents consensus algorithms for a number of partial synchrony models with different timing requirements and failure assumptions: crash, authenticated Byzantine, and Byzantine failures. It also proves tight lower bounds on the resilience of such algorithms.

The eventual synchrony approach introduced in this paper is used to model algorithms that provide safety at all times, even in completely asynchronous runs, and guarantee liveness once the system stabilizes. This has since been established as the leading approach for circumventing the FLP impossibility result and solving asynchronous consensus, atomic broadcast, and state-machine replication.

In particular, the distributed systems engineering community has been increasingly drawn towards systems architectures that reflect the basic split between safety and liveness cited above. Dwork, Lynch, and Stockmeyer thus planted the seed for a profound rethinking of the ways that we should build, and reason about, this class of systems. Following this direction are many foundational solutions. First, these include state-machine replication methods such as Lamport’s seminal Paxos algorithm and many group communication methods. Another important branch of research that directly follows this work is given by Chandra and Toueg’s unreliable failure detector abstraction, which is realized in the eventual synchrony model of this paper. As Chandra and Toueg write: “we argue that partial synchrony assumptions can be encapsulated in the

unreliability of failure detectors. For example, in the models of partial synchrony considered in Dwork et al. it is easy to implement a failure detector that satisfies the properties of $\diamond W$." Finally, the insight by Dwork, Lynch, and Stockmeyer also led to various timed-based models of partial synchrony, such as Cristian and Fetzer's Timed-Asynchronous model and others.

The award committee would like to acknowledge the sincere efforts by the nominators of this work, as well as all other (worthy!) nominations which came short of winning.

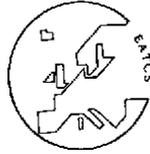
The Committee wishes to pay a special tribute via this award to Larry Stockmeyer, who passed away on July 31, 2004. Larry's impact on the field through this paper and many others will always be remembered.

The Committee of the 2007 Edsger W. Dijkstra Prize
in Distributed Computing:

Hagit Attiya
Dahlia Malkhi
Keith Marzullo
Marios Mavronicolas
Andrzej Pelc
Roger Wattenhofer (Chair)

Organization

DISC, the International Symposium on Distributed Computing, is an annual forum for presentation of research on all aspects of distributed computing. It is organized in cooperation with the European Association for Theoretical Computer Science (EATCS). The symposium was established in 1985 as a biannual International Workshop on Distributed Algorithms on Graphs (WDAG). The scope was soon extended to cover all aspects of distributed algorithms as WDAG came to stand for International Workshop on Distributed AlGorithms, and in 1989 it became an annual symposium. To reflect the expansion of its area of interest, the name was changed to DISC (International Symposium on DIStributed Computing) in 1998. The name change also reflects the opening of the symposium to all aspects of distributed computing.



Program Committee Chair

Andrzej Pelc, Université du Québec en Outaouais, Canada

Organizing Committee Chair

Chryssis Georgiou, University of Cyprus, Cyprus

Steering Committee Chair

Alexander Shvartsman, University of Connecticut, USA

Organizing Committee

Chryssis Georgiou	University of Cyprus, Cyprus (Chair)
Marios Mavronicolas	University of Cyprus, Cyprus
Nicolas Nicolaidis	Congresswise, Cyprus (Financial Officer)
Anna Philippou	University of Cyprus, Cyprus

Steering Committee

Hagit Attiya	Technion, Israel
Shlomi Dolev	Ben Gurion University, Israel
Pierre Fraigniaud	CNRS and Université Paris 7, France
Rachid Guerraoui	EPFL, Switzerland, (Vice Chair)

Andrzej Pelc	Université du Québec en Outaouais, Canada
Sergio Rajsbaum	Universidad Nacional Autónoma de México, México
Alexander Shvartsman	University of Connecticut, USA (Chair)

Program Committee

James Aspnes	Yale University, USA
Reuven Cohen	Technion, Israel
Sajal Das	University of Texas at Arlington, USA
Paola Flocchini	University of Ottawa, Canada
Eli Gafni	UCLA, USA
Leszek Gąsieniec	University of Liverpool, UK
Cyril Gavoille	University of Bordeaux, France
Chryssis Georgiou	University of Cyprus, Cyprus
Amir Herzberg	Bar Ilan University, Israel
Alex Kesselman	Intel, Israel
Rastislav Kráľovič	Comenius University, Slovakia
Zvi Lotker	Ben Gurion University, Israel
Marios Mavronicolas	University of Cyprus, Cyprus
Michael Merritt	AT&T Research, USA
Thomas Moscibroda	Microsoft Research at Redmond, USA
Achour Mostefaoui	IRISA, France
Andrzej Pelc	Université du Québec en Outaouais, Canada (Chair)
Michael Reiter	Carnegie Mellon University, USA
Eric Ruppert	York University, Canada
Arun Somani	Iowa State University, USA
Paul Spirakis	Computer Technology Institute, Greece
Sam Toueg	University of Toronto, Canada
Jennifer Welch	Texas A&M University, USA
Udi Wieder	Microsoft Research at Silicon Valley, USA
Masafumi Yamashita	Kyushu University, Japan

Sponsoring Institutions



CYTA - Cyprus Telecommunications Authority



University of Cyprus



Cyprus Tourism Organisation



COST Action 295 DYNAMO

Microsoft®

Microsoft

DISC 2007 Webmasters

Chryssis Georgiou, University of Cyprus, Cyprus

Richard Kráľovič, Comenius University, Slovakia

Referees

N. Agmon	M. Elkin	S. L. Horn
M. K. Aguilera	X. Euthimiou	Z. Hu
L. Alvisi	G. Even	R. Kat
H. Attiya	P. Fatourou	L. Katzir
Y. Bartal	H. Fauconnier	D. Kaynar
S. Baswana	A. Fernandez	I. Keidar
A. Beimel	F. Freiling	A. Kinalis
P. Bille	S. Funke	R. Klasing
E. Bortnikov	S. Ganguly	S. Kontogiannis
D. Benjamin Carbajal	P. W. Goldberg	A. Korman
J. Chalopin	O. Goussevskaia	D. R. Kowalski
B. Charron-Bost	R. Guerraoui	R. Kráľovič
I. Chatzigiannakis	T. Harris	D. Krizanc
G. Chockler	Y. Haviv	F. Kuhn
S. Das	D. Hendler	P. Kuznetov
C. Delporte	M. Herlihy	C. Lavault
S. Dolev	E. Hillel	P. Leone
A. Dvir	J.-H. Hoepman	V. Liagkou

M. Liu
A. Lopez-Ortiz
N. Lynch
R. Majundar
D. Malkhi
G. De Marco
E. Markou
T. Mchenry
L. Michael
M. Moir
P. M. Musial
G. Mylonas

N. C. Nicolaou
A. Nisgav
F. Oprea
P. Panagopoulou
D. Pardubska
B. Patt-Shamir
T. Plachetka
G. Prencipe
R. De Prisco
G. Proietti
C. Raptopoulos
E. M. Schiller

G. Shegalov
H. Shpungin
S. Smorodinski
G. Stupp
C. Travers
U. Vaccaro
R. Wattenhofer
T. Wong
Q. Xin
I. Yoffe
S. Zaks
J. Zhang

Table of Contents

Invited Talks

Routing and Scheduling with Incomplete Information	1
<i>Burkhard Monien and Karsten Tiemann</i>	
Time-Efficient Broadcasting in Radio Networks	3
<i>David Peleg</i>	
A Subjective Visit to Selected Topics in Distributed Computing	5
<i>Michel Raynal</i>	

Regular Papers

Bounded Wait-Free Implementation of Optimally Resilient Byzantine Storage Without (Unproven) Cryptographic Assumptions	7
<i>Amitanand S. Aiyer, Lorenzo Alvisi, and Rida A. Bazzi</i>	
A Simple Population Protocol for Fast Robust Approximate Majority	20
<i>Dana Angluin, James Aspnes, and David Eisenstat</i>	
A Denial-of-Service Resistant DHT	33
<i>Baruch Awerbuch and Christian Scheideler</i>	
Mobility Versus the Cost of Geocasting in Mobile Ad-Hoc Networks	48
<i>Roberto Baldoni, Kleoni Ioannidou, and Alessia Milani</i>	
Self-stabilizing Counting in Mobile Sensor Networks with a Base Station	63
<i>Joffroy Beauquier, Julien Clement, Stephane Messika, Laurent Rosaz, and Brigitte Rozoy</i>	
Scalable Load-Distance Balancing	77
<i>Edward Bortnikov, Israel Cidon, and Idit Keidar</i>	
Time Optimal Asynchronous Self-stabilizing Spanning Tree	92
<i>Janna Burman and Shay Kutten</i>	
Rendezvous of Mobile Agents in Unknown Graphs with Faulty Links . . .	108
<i>J�r�mie Chalopin, Shantanu Das, and Nicola Santoro</i>	
Weakening Failure Detectors for k -Set Agreement Via the Partition Approach	123
<i>Wei Chen, Jialin Zhang, Yu Chen, and Xuezheng Liu</i>	

Amnesic Distributed Storage	139
<i>Gregory Chockler, Rachid Guerraoui, and Idit Keidar</i>	
Distributed Approximations for Packing in Unit-Disk Graphs	152
<i>Andrzej Czygrinow and Michał Hańcówkiak</i>	
From Crash-Stop to Permanent Omission: Automatic Transformation and Weakest Failure Detectors	165
<i>Carole Delporte-Gallet, Hugues Fauconnier, Felix C. Freiling, Lucia Draque Penso, and Andreas Tielmann</i>	
Deterministic Distributed Construction of Linear Stretch Spanners in Polylogarithmic Time	179
<i>Bilel Derbel, Cyril Gavoille, and David Peleg</i>	
On Self-stabilizing Synchronous Actions Despite Byzantine Attacks	193
<i>Danny Dolev and Ezra N. Hoch</i>	
Gossiping in a Multi-channel Radio Network: An Oblivious Approach to Coping with Malicious Interference (Extended Abstract)	208
<i>Shlomi Dolev, Seth Gilbert, Rachid Guerraoui, and Calvin Newport</i>	
The Space Complexity of Unbounded Timestamps	223
<i>Faith Ellen, Panagiota Fatourou, and Eric Ruppert</i>	
Approximating Wardrop Equilibria with Finitely Many Agents	238
<i>Simon Fischer, Lars Olbrich, and Berthold Vöcking</i>	
Energy and Time Efficient Broadcasting in Known Topology Radio Networks	253
<i>Leszek Gąsieniec, Erez Kantor, Dariusz R. Kowalski, David Peleg, and Chang Su</i>	
A Distributed Algorithm for Finding All Best Swap Edges of a Minimum Diameter Spanning Tree	268
<i>Beat Gfeller, Nicola Santoro, and Peter Widmayer</i>	
On the Message Complexity of Indulgent Consensus	283
<i>Seth Gilbert, Rachid Guerraoui, and Dariusz R. Kowalski</i>	
Gathering Autonomous Mobile Robots with Dynamic Compasses: An Optimal Result	298
<i>Taisuke Izumi, Yoshiaki Katayama, Nobuhiro Inuzuka, and Koichi Wada</i>	
Compact Separator Decompositions in Dynamic Trees and Applications to Labeling Schemes	313
<i>Amos Korman and David Peleg</i>	

On the Communication Surplus Incurred by Faulty Processors	328
<i>Dariusz R. Kowalski and Michał Strojnowski</i>	
Output Stability Versus Time Till Output (Extended Abstract)	343
<i>Shay Kutten and Toshimitsu Masuzawa</i>	
A Distributed Maximal Scheduler for Strong Fairness	358
<i>Matthew Lang and Paolo A.G. Sivilotti</i>	
Cost-Aware Caching Algorithms for Distributed Storage Servers	373
<i>Shuang Liang, Ke Chen, Song Jiang, and Xiaodong Zhang</i>	
Push-to-Pull Peer-to-Peer Live Streaming	388
<i>Thomas Locher, Remo Meier, Stefan Schmid, and Roger Wattenhofer</i>	
Probabilistic Opaque Quorum Systems	403
<i>Michael G. Merideth and Michael K. Reiter</i>	
Detecting Temporal Logic Predicates on Distributed Computations	420
<i>Vinit A. Ogale and Vijay K. Garg</i>	
Optimal On-Line Colorings for Minimizing the Number of ADMs in Optical Networks (Extended Abstract)	435
<i>Mordechai Shalom, Prudence W.H. Wong, and Shmuel Zaks</i>	
Efficient Transformations of Obstruction-Free Algorithms into Non-blocking Algorithms	450
<i>Gadi Taubenfeld</i>	
Automatic Classification of Eventual Failure Detectors	465
<i>Piotr Zieliński</i>	

Brief Announcements

When $3f + 1$ Is Not Enough: Tradeoffs for Decentralized Asynchronous Byzantine Consensus	480
<i>Alysson Neves Bessani, Miguel Correia, Henrique Moniz, Nuno Ferreira Neves, and Paulo Verissimo</i>	
On the Complexity of Distributed Greedy Coloring	482
<i>Cyril Gavoille, Ralf Klasing, Adrian Kosowski, and Alfredo Navarra</i>	
Fault-Tolerant Implementations of the Atomic-State Communication Model in Weaker Networks	485
<i>Colette Johnen and Lisa Higham</i>	
Transaction Safe Nonblocking Data Structures	488
<i>Virendra J. Marathe, Michael F. Spear, and Michael L. Scott</i>	

Long Live Continuous Consensus	490
<i>Tal Mizrahi and Yoram Moses</i>	
Fully Distributed Algorithms for Convex Optimization Problems	492
<i>Damon Mosk-Aoyama, Tim Roughgarden, and Devavrat Shah</i>	
On the Power of Impersonation Attacks	494
<i>Michael Okun</i>	
Perfectly Reliable and Secure Communication in Directed Networks Tolerating Mixed Adversary	496
<i>Arpita Patra, Ashish Choudhary, Kannan Srinathan, and Chandrasekharan Pandu Rangan</i>	
A Formal Analysis of the Deferred Update Technique	499
<i>Rodrigo Schmidt and Fernando Pedone</i>	
 DISC 20th Anniversary	
DISC at Its 20th Anniversary (Stockholm, 2006)	501
<i>Michel Raynal, Sam Toueg, and Shmuel Zaks</i>	
Time, Clocks, and the Ordering of My Ideas About Distributed Systems (DISC 20th Anniversary: Invited Talk)	504
<i>Leslie Lamport</i>	
My Early Days in Distributed Computing Theory: 1979-1982 (DISC 20th Anniversary: Invited Talk)	505
<i>Nancy Lynch</i>	
Provably Unbreakable Hyper-Encryption Using Distributed Systems (DISC 20th Anniversary: Invited Talk)	506
<i>Michael O. Rabin</i>	
 Author Index	 509