

Lecture Notes in Computer Science

1499

Edited by G. Goos, J. Hartmanis and J. van Leeuwen

Shay Kutten (Ed.)

Distributed Computing

12th International Symposium, DISC'98
Andros, Greece, September 24-26, 1998
Proceedings



Springer

Series Editors

Gerhard Goos, Karlsruhe University, Germany

Juris Hartmanis, Cornell University, NY, USA

Jan van Leeuwen, Utrecht University, The Netherlands

Volume Editor

Shay Kutten

IBM T.J. Watson Research Center

P.O. Box 704, Yorktown Heights, NY 10598, USA

and

Faculty of Industrial Engineering and Management

The Technion

Haifa 32000, Israel

E-mail: kutten@ie.technion.ac.il

Cataloging-in-Publication data applied for

Die Deutsche Bibliothek - CIP-Einheitsaufnahme

Distributed computing : 12th international symposium ;
proceedings / DISC '98, Andros, Greece, September 24 - 26, 1998.

Shay Kutten (ed.). - Berlin ; Heidelberg ; New York ; Barcelona ;
Budapest ; Hong Kong ; London ; Milan ; Paris ; Singapore ; Tokyo :
Springer, 1998

(Lecture notes in computer science ; Vol. 1499)

ISBN 3-540-65066-0

CR Subject Classification (1991): C.2.4, C.2.2, F.2.2, D.1.3, F.1, D.4.4-5

ISSN 0302-9743

ISBN 3-540-65066-0 Springer-Verlag 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-Verlag. Violations are liable for prosecution under the German Copyright Law.

© Springer-Verlag Berlin Heidelberg 1998

Printed in Germany

Typesetting: Camera-ready by author

SPIN 10638936 06/3142 - 5 4 3 2 1 0 Printed on acid-free paper

Preface

The name of WDAG was changed to DISC (The International Symposium on DIStributed Computing) to reflect the expansion from a workshop to a symposium, as well as the expansion of the area of interest. DISC'98 builds on, and expands, the tradition of WDAG as a truly international symposium on all aspects of distributed computing. It aims to reflect the exciting and rapid developments in this field, and to help lead its continued exploration, by serving as a meeting point for researchers from various sub-fields.

Following 11 WDAGs, DISC'98 is the 12th in the series. This volume of DISC'98 contains the 28 extended abstracts presented at the symposium, held on September 24-26, 1998, in Andros, Greece. Opening the volume are three extended abstracts based on three keynote lectures given at the symposium by Yoram Moses, Nicola Santoro, and Santosh Shrivastava.

The contributed papers were selected from 87 submissions, at a meeting of the program committee held in Athens, following an on-line discussion. The papers were read and evaluated by the program committee members, with the additional very helpful assistance of external reviewers when needed. It is expected that papers based on most of these extended abstracts will appear in scientific journals.

The program committee consisted of: Yair Amir (Johns Hopkins University), Tushar Chandra (IBM T.J. Watson Research Center), Bernadette Charron-Bost (CNRS, Ecole Polytechnique), Alfredo de Santis (University of Salerno), Faith Fich (University of Toronto), Ajei Gopal (IBM, Pervasive Computing), Yuh-Jzer Joung (National University of Taiwan), Shay Kutten – Chair (IBM T.J. Watson Research Center and Technion), Marios Mavronicolas (University of Cyprus), Louise Moser (University of California, Santa Barbara), Noam Nisan (Hebrew University and IDC), Boaz Patt-Shamir (Tel Aviv University), Peter Ruzicka (Comenius University), Nicola Santoro (Carleton University), Marc Shapiro (INRIA), Alex Shvartsman (University of Connecticut), and Paul Spirakis (Patras University and Computer Technology Institute).

The local arrangements (both for the committee meeting, and for the conference) were made by a team headed by Paul Spirakis (local arrangements chair) and Panagiota Fatourou (local arrangements co-chair) from the Computer Technology Institute in Greece. They also took care of advertising the symposium, and the program committee chair would like to thank them for a wonderful job.

We would like to thank all the authors who submitted extended abstracts. There were interesting papers among those we could not accept, and their authors also contributed towards enhancing DISC. We hope to see the continuing support of the next DISC, by everybody. We would also like to thank the invited speakers for joining us in Andros. The program chair would like to thank the program committee members who invested a lot of work and thought in the symposium.

We would like to thank ACM SIGACT for letting us use its automated submission server and its program committee server at sigact.acm.org, and Springer-Verlag for following the tradition and publishing our proceedings. It is a special pleasure to thank

Esther Jennings and Rinat Regev for helping with operating the SIGACT server and Dana Fruchter for her help in putting together the final proceedings.

DISC thanks the following for their generous support: The Greek Ministry of Education, the Greek General Secretary of Research and Technology, the University of Patras, Greek PTT (OTE), Intracom, and Intrasoft.

Continuation of the DISC events is supervised by the DISC Steering Committee, which, for 1998, consists of: Özalp Babağlu (U. Bologna), Bernadette Charron-Bost (E. Polytechnique), Vassos Hadzilacos (U. Toronto), Shay Kutten (IBM and Technion), Marios Mavronicolas (U. Cyprus), Sam Toueg (Cornell) – Chair, and Shmuel Zaks (Technion) - Vice Chair.

July 1998

Shay Kutten
DISC'98 Program Committee Chair

External Reviewers

The Program Committee wishes to thank the following persons, who acted as referees for DISC'98:

Y. Afek	M. Greenwald	J. Pershing
D. Agrawal	R. Guerraoui	F. Petit
M. Aguilera	M. Herlihy	G. Pierre
A. Allist	L. Higham	K. Potter-Kihlstrom
J. H. Anderson	G. Hunt	S. Rajsbaum
A. Baggio	J. James	A. Roncato
G. Banavar	E. Jennings	E. M. Royer
C. Bastarrica	M. Kalantar	E. Santos
A. Basu	V. Kalogeraki	A. Schiper
Y. Ben-Asher	D. Kaminsky	H. Shachnai
K. Berket	S. Katz	D. Shaw
X. Blondel	I. Keidar	G. Singh
P. Boldi	C. Khourey	J. Stanton
R. Borgstrom	R. R. Koch	R. Strom
S. Chari	E. Kushilevitz	D. Sturman
O. Cheiner	F. Le Fessant	R. B. Tan
M. S. Chen	C. L. Lei	G. Taubenfeld
P. Dasgupta	G. Le Lann	G. Tel
A. K. Datta	S. Leonardi	L. A. Tewksbury
S. Dandamudi	G. Lindstrom	S. Toueg
O. Dedieu	T. J. Liu	D. Touitou
C. Delporte-Gallet	D. Malkhi	Y. K. Tsay
S. Demurjian	B. Mans	G. Tsudik
S. Dolev	C. Mascolo	M. Tuttle
N. Dorta	F. Mattern	R. VanRenesse
T. Eilam	P. M. Melliar-Smith	V. Villain
B. Englert	M. Moir	D. W. Wang
O. Etzion	B. Mukherjee	F. Wang
H. Fauconnier	P. Narasimhan	M. J. Ward
P. Flocchini	L. Narayanan	S. Weiss
B. Folliot	G. Neiger	A. Wool
C. Fournet	R. Netzer	D. Wu
P. Fraigniaud	S. Onn	L. C. Wu
J. A. Garay	J. Parameswaran	S. Zaks
C. Gavoille	D. Peleg	

Table of Contents

Keynote Lectures

Sense of Direction in Distributed Computing P. Flocchini, B. Mans, N. Santoro	1
Top-Down Considerations on Distributed Computing R. van der Meyden, Y. Moses	16
Inter-Task Co-ordination in Long-Lived Distributed Applications S. K. Shrivastava	20

Contributed Papers

Seamlessly Selecting the Best Copy from Internet-Wide Replicated Web Servers Y. Amir, A. Peterson, D. Shaw	22
Wait-Free Synchronization in Quantum-Based Multiprogrammed Systems J. H. Anderson, R. Jain, D. Ott	34
Computing in Totally Anonymous Asynchronous Shared Memory Systems H. Attiya, A. Gorbach, S. Moran	49
Transient Fault Detectors J. Beauquier, S. Delaët, S. Dolev, S. Tixeuil	62
Directed Virtual Path Layouts in ATM Networks J.-C. Bermond, N. Marlin, D. Peleg, S. Perennes	75
A Decision-Theoretic Approach to Reliable Message Delivery F. Chu, J. Halpern	89
Propagation and Leader Election in a Multihop Broadcast Environment I. Cidon, O. Mokryn	104
The Arrow Distributed Directory Protocol M. J. Demmer, M. P. Herlihy	119
Efficient Byzantine Agreement Secure Against General Adversaries Matthias Fitzi, Ueli Maurer	134
Long-lived, Fast, Waitfree Renaming with Optimal Name Space and High Throughput W. Eberly, L. Higham, J. Warpechowska-Gruca	149

The Compactness of Interval Routing for Almost All Graphs <i>C. Gavoille, D. Peleg</i>	161
A Wait-Free Classification of Loop Agreement Tasks <i>M. Herlihy, S. Rajsbaum</i>	175
A Stabilizing Repair Timer <i>T. Herman</i>	186
Java: Memory Consistency and Process Coordination <i>L. Higham, J. Kawash</i>	201
A Complete and Constant Time Wait-Free Implementation of CAS from LL/SC and Vice Versa <i>P. Jayanti</i>	216
Failure Detection and Consensus in the Crash-Recovery Model <i>M. K. Aguilera, W. Chen, S. Toueg</i>	231
A More Committed Quorum-Based Three Phase Commit Protocol <i>T. Kempster, C. Stirling, P. Thanisch</i>	246
Multicast Group Communication as a Base for a Load-Balancing Replicated Data Service <i>R. Khazan, A. Fekete, N. Lynch</i>	258
Efficient Deadlock-Free Multi-dimensional Interval Routing in Interconnection Networks <i>R. Kráľovič, B. Rován, P. Ružička and D. Štefankovič</i>	273
A New Protocol for Efficient Cooperative Transversal Web Caching <i>J.-M. Menaud, V. Issarny, M. Banâtre</i>	288
Fairness of Shared Objects <i>M. Merritt, G. Taubenfeld</i>	303
Optimistic Atomic Broadcast <i>F. Pedone, A. Schiper</i>	318
Approximate Agreement with Mixed Mode Faults: Algorithm and Lower Bound <i>R. Plunkett, A. Fekete</i>	333
Using Remote Access Histories for Thread Scheduling in Distributed Shared Memory Systems <i>A. Schuster, L. Shalev</i>	347
The Bancomat Problem: An Example of Resource Allocation in a Partitionable Asynchronous System <i>J. Sussman, K. Marzullo</i>	363

Lifetime Based Consistency Protocols for Distributed Objects <i>F. J. Torres-Rojas, M. Ahamad, M. Raynal</i>	378
Deriving a Scalable Algorithm for Mutual Exclusion <i>Y.-K. Tsay</i>	393
OFC: A Distributed Fossil-Collection Algorithm for Time-Warp <i>C. H. Young, N. B. Abu-Ghazaleh, P. A. Wilsey</i>	408
Author Index	419