

Lecture Notes in Computer Science

1151

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

Advisory Board: W. Brauer D. Gries J. Stoer

Özalp Babaoğlu Keith Marzullo (Eds.)

Distributed Algorithms

10th International Workshop, WDAG '96
Bologna, Italy, October 9-11, 1996
Proceedings



Springer

Series Editors

Gerhard Goos, Karlsruhe University, Germany

Juris Hartmanis, Cornell University, NY, USA

Jan van Leeuwen, Utrecht University, The Netherlands

Volume Editors

Özalp Babaoğlu

Department of Computer Science, University of Bologna

Piazza Porta S. Donato, 5, I-40127 Bologna, Italy

E-mail: ozalp@cs.unibo.it

Keith Marzullo

Department of Computer Science and Engineering

University of California, San Diego

9500 Gilman Drive, La Jolla, CA 92093-0114, USA

E-mail: marzullo@cs.ucsd.edu

Cataloging-in-Publication data applied for

Die Deutsche Bibliothek - CIP-Einheitsaufnahme

Distributed algorithms : 10th international workshop ;
proceedings / WDAG '96, Bologna, Italy, October 9 - 11, 1996.
Özalp Babaoğlu ; Keith Marzullo (ed.). - Berlin ; Heidelberg ;
New York ; Barcelona ; Budapest ; Hong Kong ; London ;
Milan ; Paris ; Santa Clara ; Singapore ; Tokyo : Springer, 1996
(Lecture notes in computer science ; Vol. 1151)

ISBN 3-540-61769-8

NE: Babaoğlu, Özalp [Hrsg.]; WDAG <10, 1996, Bologna>; GT

CR Subject Classification (1991): F1. D1.3, C.2.2, C.2.4, D.4.4-5, F2.2

ISSN 0302-9743

ISBN 3-540-61769-8 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 1996

Printed in Germany

Typesetting: Camera-ready by author

SPIN 10513788 06/3142 - 5 4 3 2 1 0 Printed on acid-free paper

Preface

The *Workshop on Distributed Algorithms* is intended to be a forum covering all aspects of distributed algorithms for computation and communication. It is devoted both to the presentation of new research results and to the identification of future research directions.

The year 1996 marks the tenth WDAG conference. It also marks a time at which the field of distributed computing is undergoing a major transition from being the domain of a small research community to that of industry and the public at large. Distributed computing is rapidly becoming the principal computing paradigm for doing business in widely diverse areas, and their requirements present new opportunities and challenges for research in distributed algorithms.

For the tenth WDAG conference, papers were solicited describing original results in all areas of distributed algorithms and their applications, including algorithms for control and communication, distributed searching including resource discovery and retrieval, network protocols, applications and services, fault tolerance and high availability, real-time distributed systems, algorithms for dynamic topology management, mobile computing, distributed intelligent agents, issues in synchrony, asynchrony, scalability and real-time, replicated data management, security in distributed systems, self stabilization, wait-free algorithms, and techniques and paradigms for the design and analysis of distributed systems.

In response to the call for papers, 75 submissions were received from 22 different countries, and 21 papers were accepted for presentation at the workshop. The selection was based on perceived originality and quality. There was considerable interest in papers that examined practical distributed algorithms and their applications. The program committee consisted of:

D. Agrawal (UC Santa Barbara)	D. Malkhi (AT&T Bell Labs, Murray Hill)
J. Anderson (UNC Chapel Hill)	
H. Attiya (Technion, Haifa)	K. Marzullo (UC San Diego, co-chair)
Ö. Babaoğlu (U. of Bologna, co-chair)	A. Panconesi (Freie U., Berlin)
E. N. Elnohazy (CMU, Pittsburgh)	M. Raynal (IRISA, Rennes)
A. Fekete (U. of Sydney)	A. Ricciardi (UT Austin)
P. Jayanti (Dartmouth College, Hanover)	R. Segala (U. of Bologna)
S. Kutten (IBM, Yorktown Heights)	G. Tel (Utrecht U.)

We wish to express our gratitude to the Program Committee members and to the external referees who assisted them during the reviewing process. Caber-Net, the ESPRIT Network of Excellence in Distributed Systems, kindly sponsored WDAG 96, and CIRFID at the University of Bologna graciously hosted the meeting. We would also like to thank the two invited speakers, Domenico Ferrari (Università Cattolica, Piacenza) and Butler Lampson (MIT, Cambridge and Microsoft, Redmond) for joining us in Bologna.

Bologna, October 1996

Özalp Babaoğlu
Keith Marzullo

List of Referees

Yehuda Afek	Jeremiah W. James	Sampath Rangarajan
Mustaque Ahamad	Kevin Jeffay	Lihu Rappoport
Lorenzo Alvisi	Theodore Johnson	K. V. Ravikanth
Alessandro Amoroso	Rajeev Joshi	Michael Reiter
Anish Arora	Idit Keidar	Arend Rensink
James Aspnes	Sanjay Khanna	Eric J. Rothfus
Alberto Bartoli	Sandeep Kulkarni	Andre Schiper
Rozoy Brigitte	Akhil Kumar	Marco Schneider
Harry Buhrman	Wei-Cheng Lai	Marc Shapiro
Nadia Busi	Raimondas Lencevicius	Nir Shavit
Tushar Chandra	Jon C. Little	Alex Shvartsman
Craig Chase	Victor Luchangco	Gurdip Singh
Manhoi Choy	Bernard Mans	C. Skawratananond
Renzo Davoli	James Roger Mitchell	Shaune R. Stark
Luca de Alfaro	Alberto Montresor	John Hamilton Slye
Roberto de Prisco	Yoram Moses	Scott Stoller
Shlomi Dolev	Achour Mostefaoui	David Stotts
Paulo Ferreira	Venkat Murty	Jeremy Sussman
Riccardo Focardi	Nancy Norris	Yu-Pang Tang
Roy Friedman	Graziano Obertelli	Gadi Taubenfeld
Juan Garay	Michael Ogg	Oliver E. Theel
Vijay K. Garg	Gerard Padiou	Dan Touitou
Brian Grayson	Marina Papatrantaflou	John Tromp
Rachid Guerraoui	David Peleg	Philippas Tsigas
Sibsankar Halder	Evelyn Pierce	Hong Va Leong
Jean-Michel Helary	Ian Piumarta	Robbert VanRenesse
Maurice Herlihy	David Plainfosse	George Varghese
Ted Herman	Noel Plouzeau	Avishai Wool
Jaap-Henk Hoepman	Philippe Queinnec	Daniel Wu
Michel Hurfin	Tal Rabin	Shmuel Zaks
Valerie Issarny	Sergio Rajsbaum	Asaph Zernach
Alon Itai	Haritha Ramavarjula	

Table of Contents

Invited Papers

How to Build a Highly Available System Using Consensus <i>B. W. Lampson</i>	1
Distributed Admission Control Algorithms for Real-Time Communication <i>D. Ferrari</i>	18

Regular Papers

Randomization and Failure Detection: A Hybrid Approach to Solve Consensus <i>M. K. Aguilera, S. Toueg</i>	29
Levels of Authentication in Distributed Agreement <i>M. Borcherdig</i>	40
Efficient and Robust Sharing of Memory in Message-Passing Systems <i>H. Attiya</i>	56
Plausible Clocks: Constant Size Logical Clocks for Distributed Systems <i>F. J. Torres-Rojas, M. Ahamad</i>	71
Abstracting Communication to Reason about Distributed Algorithms <i>M. Charpentier, M. Filali, P. Mauran, G. Padiou, P. Quéinnec</i>	89
Simulating Reliable Links with Unreliable Links in the Presence of Process Crashes <i>A. Basu, B. Charron-Bost, S. Toueg</i>	105
A Cyclic Distributed Garbage Collector for Network Objects <i>H. Rodrigues, R. Jones</i>	123
Incremental, Distributed Orphan Detection and Actor Garbage Collection Using Graph Partitioning and Euler Cycles <i>P. Dickman</i>	141
A Framework for the Analysis of Non-Deterministic Clock Synchronization Algorithms <i>P. Fonseca, Z. Mammeni</i>	159
Optimal Time Broadcasting in Faulty Star Networks <i>A. Mei, F. Bao, Y. Hamada, Y. Igarashi</i>	175

A Lower Bound for Linear Interval Routing <i>T. Eilam, S. Moran, S. Zaks</i>	191
Topological Routing Schemes <i>G. Gambosi, P. Vocca</i>	206
Maintaining a Dynamic Set of Processors in a Distributed System <i>S. Fujita, M. Yamashita</i>	220
Collective Consistency <i>C. Dwork, C.-T. Ho, R. Strong</i>	234
Planar Quorums <i>R. A. Bazzi</i>	251
“Gamma-Accurate” Failure Detectors <i>R. Guerraoui, A. Schiper</i>	269
Fast Long-Lived Renaming Improved and Simplified <i>M. Moir, J. A. Garay</i>	287
A Timestamp Based Transformation of Self-Stabilizing Programs for Distributed Computing Environments <i>M. Mizuno, H. Kakugawa</i>	304
The Combinatorial Structure of Wait-Free Solvable Tasks <i>H. Attiya, S. Rajsbaum</i>	322
On the Robustness of h_m^r <i>S. Moran, L. Rappoport</i>	344
Understanding the Set Consensus Partial Order Using the Borowsky-Gafni Simulation <i>S. Chaudhuri, P. Reiners</i>	362
Author Index	381