

Lecture Notes in Computer Science

725

Edited by G. Goos and J. Hartmanis

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



André Schiper (Ed.)

Distributed Algorithms

7th International Workshop, WDAG '93
Lausanne, Switzerland, September 27-29, 1993
Proceedings

Springer-Verlag

Berlin Heidelberg New York
London Paris Tokyo
Hong Kong Barcelona
Budapest

Series Editors

Gerhard Goos
Universität Karlsruhe
Postfach 69 80
Vincenz-Priessnitz-Straße 1
D-76131 Karlsruhe, Germany

Juris Hartmanis
Cornell University
Department of Computer Science
4130 Upson Hall
Ithaca, NY 14853, USA

Volume Editor

André Schiper
Ecole Polytechnique Fédérale de Lausanne
Département d'Informatique
Laboratoire de Systèmes d'Exploitation
CH-1015 Lausanne, Switzerland

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

ISBN 3-540-57271-6 Springer-Verlag Berlin Heidelberg New York
ISBN 0-387-57271-6 Springer-Verlag New York Berlin Heidelberg

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 1993
Printed in Germany

Typesetting: Camera-ready by author
Printing and binding: Druckhaus Beltz, Hemsbach/Bergstr.
45/3140-543210 - Printed on acid-free paper

Preface

The SEVENTH INTERNATIONAL WORKSHOP ON DISTRIBUTED ALGORITHMS (WDAG 93) was held September 27-29 in Lausanne, Switzerland (more precisely in *Les Diablerets*, a village located near Lausanne). The workshop followed six successive workshops held in Ottawa (1985, proceedings published by Carleton University Press), Amsterdam (1987, proceedings published by Springer Verlag, LNCS 312), Nice (1989, LNCS 392), Bari (1990, LNCS 484), Delphi (1991, LNCS 579) and Haifa (1992, LNCS 647). The WDAG provides an international forum for the presentation of new research results and the identification of future research directions in the area of distributed algorithms.

Submissions were solicited in all areas of distributed algorithms and their applications, including distributed algorithms for control and communication, fault-tolerant distributed algorithms, network protocols, algorithms for managing replicated data, protocols for real-time distributed systems, issues of asynchrony, synchrony and real-time, mechanisms for security in distributed systems, techniques for the design and analysis of distributed algorithms, distributed database techniques, distributed combinatorial and optimization algorithms, and distributed graph algorithms.

A total of 72 papers were received within the submission deadline (33 submissions from Europe, 29 from North America, 6 from the Middle East, 3 from the Far East, and 1 from Australia). The Program Committee wishes to thank all the authors who submitted papers for consideration. Out of the 72 submissions the Program Committee was able to select the 22 papers appearing in these proceedings (6 from Europe, 13 from North America, and 3 from the Middle East). The selection was based on originality and quality. Relevance of the papers to the field of distributed computing was also considered carefully.

The Program Committee was composed of:

D. Dolev (Hebrew U. and IBM Almaden)	M. Herlihy (DEC CRL)
G. Le Lann (INRIA, Paris)	K. Marzullo (Cornell U. and UC San Diego)
F. Mattern (U. of Saarland)	M. Raynal (IRISA, Rennes)
M. Merritt (AT&T)	P. Spirakis (CTI and Patras U.)
A. Schiper (<i>chair</i> , EPF Lausanne and Cornell U.)	J. van Leeuwen (U. of Utrecht)
J. Welch (Texas A&M U.)	S. Zaks (Technion, Haifa)

I wish to thank all the members of the Program Committee and all the referees who assisted them for their careful reviewing carried out within a very short period of time. My thanks go also to Alain Sandoz for his excellent organization of the Workshop, and for his smooth handling of the submitted and accepted papers.

For the first time this year, a tutorial was organized during the Workshop. The title of the tutorial was *Specifications and Algorithms for Fault-Tolerant*

Broadcasts – A Modular Approach. It was presented by Sam Toueg from Cornell University. The idea in combining a tutorial with research papers was to attract young researchers to the Workshop, and hopefully to the field of distributed algorithms. Financial support for the tutorial was provided by the 3ème Cycle Romand d'Informatique, for which we are grateful.

Lausanne, September 1993

André Schiper

List of Referees

Anceaume, E.	Mattern, F.
Attiya, H.	Merritt, M.
Bakker, E.M.	Minet, P.
Breitbach, T.	Moran, S.
Buehler, P.	Papatriantafillou, M.
Charron-Bost, B.	Plouzeau, N.
Cidon, I.	Priebe, V.
Clegg, M.	Rachman, O.
Dolev, D.	Raynal, M.
Dolev, S.	Ricciardi, A.
Fix, L.	Richter, J.
Fuenfrocken, S.	Sandoz, A.
Guerraoui, R.	Schiper, A.
Haldar, S.	Schwarz, R.
Helary, J.-M.	Spirakis, P.
Herlihy, M.	Tampakas, B.
Herman, T.	Tan, R.B.
Israeli, A.	Tel, G.
Le Lann, G.	Tsigas, Ph.
Lueling, R.	van Leeuwen, J.
Malki, D.	Welch, J.
Marzullo, K.	Zaks, S.
Matsliach, G.	Zamsky, A.

Table of Contents

Wait-free synchronization

Efficient wait-free implementation of a concurrent priority queue <i>A. Israeli, L. Rappoport</i>	1
Binary snapshots <i>J.-H. Hoepman, J. Tromp</i>	18
Linear-time snapshot protocols for unbalanced systems <i>A. Israeli, A. Shaham, A. Shirazi</i>	26
Towards a necessary and sufficient condition for wait-free synchronization <i>J.H. Anderson, M. Moir</i>	39

Shared memory model

Efficient algorithms for checking the atomicity of a run of read and write operations <i>L.M. Kirousis, A.G. Veneris</i>	54
Benign failure models for shared memory <i>Y. Afek, M. Merritt, G. Taubenfeld</i>	69
Generalized agreement between concurrent fail-stop processes <i>J.E. Burns, R.I. Cruz, M.C. Loui</i>	84
Controlling memory access concurrency in efficient fault-tolerant parallel algorithms <i>P.C. Kanellakis, D. Michailidis, A.A. Shvartsman</i>	99

Miscellaneous

Asynchronous epoch management in replicated databases <i>M. Rabinovich, E.D. Lazowska</i>	115
Crash resilient communication in dynamic networks <i>S. Dolev, J.L. Welch</i>	129
Distributed job scheduling using snapshots <i>M. Choy, A.K. Singh</i>	145

Fault tolerance

Optimal time self stabilization in dynamic systems <i>S. Dolev</i>	160
Tolerating transient and permanent failures <i>E. Anagnostou, V. Hadzilacos</i>	174
Quick atomic broadcast <i>P. Berman, A.A. Bharali</i>	189
Time bounds for decision problems in the presence of timing uncertainty and failures <i>H. Attiya, T. Djerassi-Shintel</i>	204

Networks and rings

Boolean routing <i>M. Flammini, G. Gambosi, S. Salomone</i>	219
Notes on sorting and counting networks <i>N. Hardavellas, D. Karakos, M. Mavronicolas</i>	234
A simple, efficient algorithm for maximum finding on rings <i>L. Higham, T. Przytycka</i>	249
Wang tilings and distributed orientation on anonymous torus networks <i>V.R. Syrotiuk, C.J. Colbourn, J. Pahl</i>	264

Miscellaneous

Fairness of N-party synchronization and its implementation in a distributed environment <i>C. Wu, G. v. Bochmann, M. Yao</i>	279
Programming distributed reactive systems: a strong and weak synchronous coupling <i>F. Boniol, M. Adelantado</i>	294
Using message semantics to reduce rollback in the time warp mechanism <i>H.V. Leong, D. Agrawal, J.R. Agre</i>	309
List of Authors	325