

# Lecture Notes in Computer Science

647

Edited by G. Goos and J. Hartmanis

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



A. Segall S. Zaks (Eds.)

# Distributed Algorithms

6th International Workshop, WDAG '92

Haifa, Israel, November 2–4, 1992

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  
W-7500 Karlsruhe, FRG

Juris Hartmanis  
Cornell University  
Department of Computer Science  
5149 Upson Hall  
Ithaca, NY 14853, USA

## Volume Editors

Adrian Segall  
Shmuel Zaks  
Department of Computer Science, Israel Institute of Technology  
Technion, Haifa 32000, Israel

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

ISBN 3-540-56188-9 Springer-Verlag Berlin Heidelberg New York  
ISBN 0-387-56188-9 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 1992  
Printed in Germany

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

## Preface

The *Sixth Workshop on Distributed Algorithms* (WDAG 92) took place on November 2-4, 1992 in Haifa, Israel. WDAG is intended to provide a forum for researchers and other parties interested in distributed algorithms and their applications. The aim is to present recent research results, explore directions for future research, and identify common fundamental techniques that serve as building blocks in many distributed algorithms. WDAG 92 follows five successful workshops in Ottawa (1985, proceedings published by Carleton University Press), Amsterdam (1987, see Lecture Notes in Computer Science (LNCS) 312), Nice (1989, LNCS 392), Bari (1990, LNCS 484) and Delphi (1991, LNCS 579).

Papers were solicited describing original results in all areas of distributed algorithms and their applications, including distributed graph algorithms, distributed combinatorial algorithms, design of network protocols, routing and flow control, communication complexity, fault-tolerant distributed algorithms, distributed data structures, distributed database techniques, replica control protocols, distributed optimization algorithms, mechanisms for safety and security in distributed systems, and protocols for real-time distributed systems.

The 24 papers were selected by the Program Committee from 38 submitted papers. The selection was based on perceived originality and quality. The selection process was carried out via email, using a scoring analysis program (courtesy of Robert Schapire, of AT&T). It is expected that the authors will prepare extended versions of the papers appearing in this proceedings, to be submitted for refereed publication in one of the scientific journals. The Program Committee wishes to thank all who submitted papers for consideration.

The Program Committee consisted of

|                                      |   |
|--------------------------------------|---|
| L. Bougé (ENS Lyon)                  | A. Segall ( <i>co-chair</i> , Technion) |
| D. Dolev (Hebrew U. and IBM Almaden) | P. Spirakis (CTI and Patras U.)         |
| S. Kutten (IBM)                      | G. Tel (Utrecht U.)                     |
| M. Merritt (AT&T)                    | S. Toueg (Cornell U.)                   |
| N. Santoro (Carleton U.)             | J. Welch (U. of North Carolina)         |
| A. Schiper (EPF Lausanne)            | S. Zaks ( <i>co-chair</i> , Technion).  |

We wish to express our gratitude to all members of the Program Committee for their cooperation and to the referees who assisted them (see appendix for list).

We wish to thank the Department of Computer Science, Technion for placing at our disposal the best professional assistance it can offer. Judith Tamari did an excellent job in organizing the conference and putting the proceedings together, and Aythan Avior - in handling the computer work and adapting the scoring analysis program to WDAG needs. Our deepest thanks go to them both. We would also like to thank David Cohen of the *S. Neaman Institute* for his experienced help and advice.

The *S. Neaman Institute* selected WDAG 92 as the first workshop to receive sponsorship under the newly-established program to support scientific meetings related to research carried out at the Technion. On behalf of all Workshop participants and the entire Distributed Algorithms scientific community, the Workshop Steering Committee and Co-Chairmen would like to express their deep appreciation to the *S. Neaman Institute* for this support.

Haifa, November 1992

Adrian Segall  
Shmuel Zaks

## **The S. Neaman Institute for Advanced Studies in Science and Technology**

The Samuel Neaman Institute for Advanced Studies in Science and Technology is an independent public-policy research institute, established in 1978 to assist in the search for solutions to national problems in science and technology, education, economy and industry, and social development. As an interdisciplinary think-tank, the institute draws on the faculty and staff of Technion, other institutions and scientists in Israel, and specialists abroad. The Institute serves as a bridge between academia and decision makers through research, workshops and publications.

The Institute pursues a policy of inquiry and analysis designed to identify significant public policy problems, to determine possible courses of action to deal with the problems, and to evaluate the consequences of the identified courses of action.

As an independent not-for-profit research organization, the Institute does not advocate any specific policy or embrace any particular social philosophy. As befits a democratic society, the choices among policy alternatives are the prerogative and responsibility of the elected representatives of the citizenry. The Samuel Neaman Institute endeavors to contribute to a climate of informed choice.

The Institute undertakes sponsored advanced research, formulates invitational workshops, implements continuing education activities on topics of significance for the development of the State of Israel, and maintains a publications program for the dissemination of research and workshop findings. Specific topics for research may be initiated by the Institute, researchers, government agencies, foundations, industry or other concerned institutions. Each research program undertaken by the Institute is designed to be a significant scholarly study worthy of publication and public attention.

# Table of Contents

|   |     |
|---|-----|
| Sparsen: a paradigm for running distributed algorithms<br><i>Y. Afek, M. Ricklin</i> .....  | 1   |
| Closed schedulers: constructions and applications to consensus protocols<br><i>R. Lubitch, S. Moran</i> .....   | 11  |
| Efficient atomic snapshots using lattice agreement<br><i>H. Attiya, M. Herlihy, O. Rachman</i> .....  | 35  |
| Choice coordination with multiple alternatives<br><i>D.S. Greenberg, G. Taubenfeld, D.-W. Wang</i> .....  | 54  |
| Some results on the impossibility, universality, and decidability<br>of consensus<br><i>P. Jayanti, S. Toueg</i> .....  | 69  |
| Wait-free test-and-set<br><i>Y. Afek, E. Gafni, J. Tromp, P.M.B. Vitanyi</i> .....  | 85  |
| A concurrent time-stamp scheme which is linear in time and space<br><i>A. Israeli, M. Pinhasov</i> .....  | 95  |
| Tentative and definite distributed computations: an optimistic<br>approach to network synchronization<br><i>J. Garofalakis, S. Rajsbaum, P. Spirakis, B. Tampakas</i> ..... | 110 |
| Semisynchrony and real time<br><i>S. Ponzio, R. Strong</i> .....  | 120 |
| Optimal time Byzantine agreement for $t < n/8$ with linear messages<br><i>A. Zamsky, A. Israeli, S.S. Pinter</i> .....  | 136 |
| A continuum of failure models for distributed computing<br><i>J.A. Garay, K.J. Perry</i> .....  | 153 |
| Simulating crash failures with many faulty processors<br><i>R. Bazzi, G. Neiger</i> .....   | 166 |

|   |     |
|---|-----|
| An efficient topology update protocol for dynamic networks<br><i>B. Awerbuch, Y. Mansour</i> .....              | 185 |
| Memory adaptive self-stabilizing protocols<br><i>E. Anagnostou, R. El-Yaniv, V. Hadzilacos</i> .....            | 203 |
| Optimal early stopping in distributed consensus<br><i>P. Berman, J.A. Garay, K.J. Perry</i> .....               | 221 |
| Traffic-light scheduling on the grid<br><i>G. Kortsarz, D. Peleg</i> .....                                      | 238 |
| Distributed computing on anonymous hypercubes with faulty<br>components<br><i>E. Kranakis, N. Santoro</i> ..... | 253 |
| Message terminate algorithms for anonymous rings of unknown size<br><i>I. Cidon, Y. Shavitt</i> .....           | 264 |
| Distributed resource allocation algorithms<br><i>J. Bar-Ilan, D. Peleg</i> .....                                | 277 |
| Membership algorithms for multicast communication groups<br><i>Y. Amir, D. Dolev, S. Kramer, D. Malki</i> ..... | 292 |
| The granularity of waiting<br><i>J.H. Anderson, J.-H. Yang, M.G. Gouda</i> .....                                | 313 |
| The cost of order in asynchronous systems<br><i>A. Ricciardi, K. Birman, P. Stephenson</i> .....                | 329 |
| Efficient, strongly consistent implementations of shared memory<br><i>M. Mavronicolas, D. Roth</i> .....        | 346 |
| Optimal primary-backup protocols<br><i>N. Budhiraja, K. Marzullo, F.B. Schneider, S. Toueg</i> .....            | 362 |
| List of Referees .....  | 379 |
| List of Authors .....   | 380 |