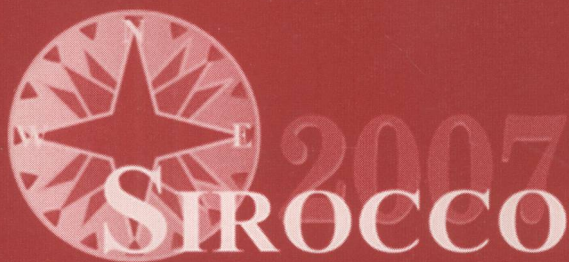


Giuseppe Prencipe
Shmuel Zaks (Eds.)

LNCS 4474

Structural Information and Communication Complexity

14th International Colloquium, SIROCCO 2007
Castiglioncello, Italy, June 2007
Proceedings



TP274-53

S927 Giuseppe Prencipe Shmuel Zaks (Eds.)
2007

Structural Information and Communication Complexity

14th International Colloquium, SIROCCO 2007
Castiglioncello, Italy, June 5-8, 2007
Proceedings



 Springer



E2007003220

Volume Editors

Giuseppe Prencipe
Università degli Studi di Pisa
Dipartimento di Informatica
Largo Bruno Pontecorvo 3, 56127 Pisa, Italy
E-mail: prencipe@di.unipi.it

Shmuel Zaks
Technion
Department of Computer Science
Haifa, Israel
E-mail: zaks@cs.technion.ac.il

Library of Congress Control Number: 2007927754

CR Subject Classification (1998): F.2, C.2, G.2, E.1

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

ISSN 0302-9743
ISBN-10 3-540-72918-6 Springer Berlin Heidelberg New York
ISBN-13 978-3-540-72918-1 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: 12072859 06/3180 5 4 3 2 1 0

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

Preface

The Colloquium on Structural Information and Communication Complexity (SIROCCO) is an annual meeting focused on the relationship between algorithmic aspects of computing and communication. Over its 14 years of existence, SIROCCO has become an acknowledged forum bringing together specialists interested in the fundamental principles underlying the interplay between information, communication, and computing. SIROCCO covers topics such as distributed computing, high-speed networks, interconnection networks, mobile computing, optical computing, parallel computing, sensor networks, wireless networks, and autonomous robots. Its topics of interest include communication complexity, distributed algorithms and data structures, information dissemination, mobile agent computing, models of communication, network topologies, routing protocols, sense of direction, structural properties, and selfish computing.

SIROCCO 2007 was held in Castiglioncello (LI), Italy, June 5–8, 2007. The previous 13 SIROCCO colloquia took place in Ottawa (1994), Olympia (1995), Siena (1996), Ascona (1997), Amalfi (1998), Lacanau-Océan (1999), L'Aquila (2000), Val de Nuria (2001), Andros (2002), Umeå (2003), Smolenice Castle (2004), Mont Saint-Michel (2005), and Chester (2006).

The 66 contributions submitted to SIROCCO 2007 were subject to a thorough refereeing process, and 23 high-quality submissions were selected for publication. We would like to thank the authors of all the submitted papers. The excellent quality of the final program is also due to the dedicated and careful work of the Program Committee members. Our gratitude extends to them. We also thank the numerous sub-referees for their valuable help.

We had four invited speakers: Hans Bodlaender (Utrecht), Luisa Gargano (Salerno), S. Muthukrishnan (Google), and Alessandro Panconesi (Rome). We thank them for accepting our invitation to share their insights on new developments in their areas of interest.

We would like to express our gratitude to the conference Chair Pierre Fraigniaud (CNRS and Paris) for his enthusiasm and invaluable consultations.

Special thanks go to the local Organizing Team of the Dipartimento di Informatica, Università di Pisa, and in particular to Vincenzo Gervasi.

We acknowledge the use of the EasyChair system (for handling the submission of papers, managing the refereeing process, and generating these proceedings).

We thank the Università di Pisa, and its Dipartimento di Informatica, for their generous support. SIROCCO 2007 was co-located with FUN 2007 and with a meeting of the EU COST 293 action (GRAAL - Graphs and Algorithms in Communication Networks). The two sessions of the invited talks were held

jointly with the GRAAL meeting. We thank the Management Committee of GRAAL, and especially their past and present Chairs Xavier Munoz and Arie Koster, for supporting the idea of this joint event and for their generous support.

June 2007

Giuseppe Prencipe
Shmuel Zaks

Conference Organization

Steering Committee

Paola Flocchini	University of Ottawa, Canada
Pierre Fraigniaud	CNRS, France (Chair)
Leszek Gasieniec	University of Liverpool, UK
Lefteris Krousis	University of Patras, Greece
Rastislav Kráľovič	Comenius University, Slovakia
Evangelos Kranakis	Carleton University, Canada
Danny Krizanc	Wesleyan University, USA
Bernard Mans	Macquarie University, Australia
Andrzej Pelc	Université du Québec en Outaouais, Canada
David Peleg	Weizmann Institute, Israel
Giuseppe Prencipe	Università di Pisa, Italy
Michel Raynal	IRISA, France
Nicola Santoro	Carleton University, Canada
Paul Spirakis	University of Patras, Greece
Shmuel Zaks	Technion, Israel

Program Committee

Christoph Ambühl	University of Liverpool, UK
Ioannis Caragiannis	University of Patras, Greece
Lenka Carr-Motyckova	Palacky University, Czech Republic
Bogdan Chlebus	University of Colorado, USA
Thomas Erlebach	University of Leicester, UK
Rastislav Kráľovič	Comenius University, Slovakia
Evangelos Kranakis	Carleton University, Canada
Jan van Leeuwen	Utrecht University, Netherlands
Euripides Markou	McMaster University, Canada
Toshimitsu Masuzawa	Osaka University, Japan
Yves Métivier	Université Bordeaux 1, France
Giuseppe Prencipe	Università di Pisa, Italy (Co-chair)
Guido Proietti	Università de L'Aquila, Italy
Andrzej Proskurowski	University of Oregon, USA
Geppino Pucci	Università di Padova, Italy
Tomasz Radzik	King's College London, UK
Sergio Rajsbaum	UNAM, Mexico
Michel Raynal	IRISA, France
Tami Tamir	The Interdisciplinary Center, Israel
Savio Tse	Bilkent University, Turkey
Peter Widmayer	ETHZ, Switzerland
Shmuel Zaks	Technion, Israel (Co-chair)

Local Organization

Vincenzo Gervasi, Dipartimento di Informatica, Università di Pisa

Sponsoring Institutions

Università di Pisa

EU COST 293 action (GRAAL)

External Reviewers

Ittai Abraham

Hagit Attiya

Alberto Bertoldo

Vittorio Bilò

Davide Bilò

Maria Blesa

Anat Bremler-Barr

Marco Bressan

Alfred Bruckstein

Keren Censor

Milind Dawande

Xavier Defago

Stefan Dobrev

Frederick Ducatelle

Toby Ehrenkrantz

Michael Elkin

Angelo Fanelli

Arthur Farley

Laura Feeney

Michele Flammini

Paola Flocchini

Luca Forlizzi

Leszek Gasieniec

Ornan (Ori) Gerstel

Vincenzo Gervasi

Sukumar Ghosh

Emmanuel Godard

Luciano Gualà

Erez Hadad

Panos Hilaris

Martin Hoefer

Jiung-yao Huang

David Ilcinkas

Michiko Inoue

Tomas Johansson

Hirotsugu Kakugawa

Panagiotis Kanellopoulos

Branislav Katreniak

Ralf Klasing

Lukasz Kowalik

Dariusz Kowalski

Richard Královič

Danny Krizanc

Qin Lv

Francesca Martelli

Russell Martin

Valia Mitsou

Gianpiero Monaco

Mohamed Mosbah

Luca Moscardelli

Alfredo Navarra

Francesco Nidito

Fukuhito Ooshita

Aris Pagourtzis

Evi Papaioannou

Dana Pardubska

Susanna Pelagatti

Andrzej Pelc

David Peleg

Enoch Peserico

Andrea Pietracaprina

Katerina Potika

Danny Raz

Nicola Santoro

Fabiano Sarracco

Christian Schindelhauer

Alexander Shvartsman
Francesco Silvestri
Shakhar Smorodinsky
Laura Spinsanti
Paul Spirakis
Rob van Stee
Alessandro Tommasi

Fabio Vandin
Imrich Vrt'ó
Koichi Wada
Prudence Wong
Yu Wu
Jun-Ming Xu
Yukiko Yamauchi

Lecture Notes in Computer Science

For information about Vols. 1–4421

please contact your bookseller or Springer

- Vol. 4538: F. Escolano, M. Vento (Eds.), *Graph-Based Representations in Pattern Recognition*. XII, 416 pages. 2007.
- Vol. 4534: I. Tomkos, F. Neri, J. Solé Pareta, X. Masip Bruin, S. Sánchez Lopez (Eds.), *Optical Network Design and Modeling*. XI, 460 pages. 2007.
- Vol. 4531: J. Indulska, K. Raymond (Eds.), *Distributed Applications and Interoperable Systems*. XI, 337 pages. 2007.
- Vol. 4530: D.H. Akehurst, R. Vogel, R.F. Paige (Eds.), *Model Driven Architecture- Foundations and Applications*. X, 219 pages. 2007.
- Vol. 4526: M. Malek, M. Reitenspieß, A. van Moorsel (Eds.), *Service Availability*. X, 155 pages. 2007.
- Vol. 4525: C. Demetrescu (Ed.), *Experimental and Efficient Algorithms*. XIII, 448 pages. 2007.
- Vol. 4523: Y.-H. Lee, H.-N. Kim, J. Kim, Y. Park, L.T. Yang, S.W. Kim (Eds.), *Embedded Software and Systems*. XIX, 829 pages. 2007.
- Vol. 4521: J. Katz, M. Yung (Eds.), *Applied Cryptography and Network Security*. XIII, 498 pages. 2007.
- Vol. 4519: E. Franconi, M. Kifer, W. May (Eds.), *The Semantic Web: Research and Applications*. XVIII, 830 pages. 2007.
- Vol. 4517: F. Boavida, E. Monteiro, S. Mascolo, Y. Koucheryavy (Eds.), *Wired/Wireless Internet Communications*. XIV, 382 pages. 2007.
- Vol. 4515: M. Naor (Ed.), *Advances in Cryptology - EUROCRYPT 2007*. XIII, 591 pages. 2007.
- Vol. 4514: S.N. Artemov, A. Nerode (Eds.), *Logical Foundations of Computer Science*. XI, 513 pages. 2007.
- Vol. 4510: P. Van Hentenryck, L. Wolsey (Eds.), *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*. X, 391 pages. 2007.
- Vol. 4509: Z. Kobti, D. Wu (Eds.), *Advances in Artificial Intelligence*. XII, 552 pages. 2007. (Sublibrary LNAI).
- Vol. 4508: M.-Y. Kao, X.-Y. Li (Eds.), *Algorithmic Applications in Management*. VIII, 428 pages. 2007.
- Vol. 4506: D. Zeng, I. Gotham, K. Komatsu, C. Lynch, M. Thurmond, D. Madigan, B. Lober, J. Kvach, H. Chen (Eds.), *Intelligence and Security Informatics: Bio-surveillance*. XI, 234 pages. 2007.
- Vol. 4505: G. Dong, X. Lin, W. Wang, Y. Yang, J.X. Yu (Eds.), *Advances in Data and Web Management*. XXII, 896 pages. 2007.
- Vol. 4504: J. Huang, R. Kowalczyk, Z. Maamar, D. Martin, I. Müller, S. Stoutenburg, K.P. Sycara (Eds.), *Service-Oriented Computing: Agents, Semantics, and Engineering*. X, 175 pages. 2007.
- Vol. 4501: J. Marques-Silva, K.A. Sakallah (Eds.), *Theory and Applications of Satisfiability Testing – SAT 2007*. XI, 384 pages. 2007.
- Vol. 4500: N. Streitz, A. Kameas, I. Mavrommati (Eds.), *The Disappearing Computer*. XVIII, 304 pages. 2007.
- Vol. 4496: N.T. Nguyen, A. Grzech, R.J. Howlett, L.C. Jain (Eds.), *Agent and Multi-Agent Systems: Technologies and Applications*. XXI, 1046 pages. 2007. (Sublibrary LNAI).
- Vol. 4493: D. Liu, S. Fei, Z. Hou, H. Zhang, C. Sun (Eds.), *Advances in Neural Networks – ISNN 2007, Part III*. XXVI, 1215 pages. 2007.
- Vol. 4492: D. Liu, S. Fei, Z. Hou, H. Zhang, C. Sun (Eds.), *Advances in Neural Networks – ISNN 2007, Part II*. XXVII, 1321 pages. 2007.
- Vol. 4491: D. Liu, S. Fei, Z.-G. Hou, H. Zhang, C. Sun (Eds.), *Advances in Neural Networks – ISNN 2007, Part I*. LIV, 1365 pages. 2007.
- Vol. 4490: Y. Shi, G.D. van Albada, J. Dongarra, P.M.A. Sloot (Eds.), *Computational Science – ICCS 2007, Part IV*. XXXVII, 1211 pages. 2007.
- Vol. 4489: Y. Shi, G.D. van Albada, J. Dongarra, P.M.A. Sloot (Eds.), *Computational Science – ICCS 2007, Part III*. XXXVII, 1257 pages. 2007.
- Vol. 4488: Y. Shi, G.D. van Albada, J. Dongarra, P.M.A. Sloot (Eds.), *Computational Science – ICCS 2007, Part II*. XXXV, 1251 pages. 2007.
- Vol. 4487: Y. Shi, G.D. van Albada, J. Dongarra, P.M.A. Sloot (Eds.), *Computational Science – ICCS 2007, Part I*. LXXXI, 1275 pages. 2007.
- Vol. 4486: M. Bernardo, J. Hillston (Eds.), *Formal Methods for Performance Evaluation*. VII, 469 pages. 2007.
- Vol. 4485: F. Sgallari, A. Murli, N. Paragios (Eds.), *Scale Space and Variational Methods in Computer Vision*. XV, 931 pages. 2007.
- Vol. 4484: J.-Y. Cai, S.B. Cooper, H. Zhu (Eds.), *Theory and Applications of Models of Computation*. XIII, 772 pages. 2007.
- Vol. 4483: C. Baral, G. Brewka, J. Schlipf (Eds.), *Logic Programming and Nonmonotonic Reasoning*. IX, 327 pages. 2007. (Sublibrary LNAI).
- Vol. 4482: A. An, J. Stefanowski, S. Ramanna, C.J. Butz, W. Pedrycz, G. Wang (Eds.), *Rough Sets, Fuzzy Sets, Data Mining and Granular Computing*. XIV, 585 pages. 2007. (Sublibrary LNAI).

- Vol. 4481: J. Yao, P. Lingras, W.-Z. Wu, M. Szczuka, N.J. Cercone, D. Ślęzak (Eds.), *Rough Sets and Knowledge Technology*. XIV, 576 pages. 2007. (Sublibrary LNAI).
- Vol. 4480: A. LaMarca, M. Langheinrich, K.N. Truong (Eds.), *Pervasive Computing*. XIII, 369 pages. 2007.
- Vol. 4479: I.F. Akyildiz, R. Sivakumar, E. Ekici, J.C.d. Oliveira, J. McNair (Eds.), *NETWORKING 2007. Ad Hoc and Sensor Networks, Wireless Networks, Next Generation Internet*. XXVII, 1252 pages. 2007.
- Vol. 4478: J. Martí, J.M. Benedí, A.M. Mendonça, J. Serrat (Eds.), *Pattern Recognition and Image Analysis, Part II*. XXVII, 657 pages. 2007.
- Vol. 4477: J. Martí, J.M. Benedí, A.M. Mendonça, J. Serrat (Eds.), *Pattern Recognition and Image Analysis, Part I*. XXVII, 625 pages. 2007.
- Vol. 4476: V. Gorodetsky, C. Zhang, V.A. Skormin, L. Cao (Eds.), *Autonomous Intelligent Systems: Multi-Agents and Data Mining*. XIII, 323 pages. 2007. (Sublibrary LNAI).
- Vol. 4475: P. Crescenzi, G. Prencipe, G. Pucci (Eds.), *Fun with Algorithms*. X, 273 pages. 2007.
- Vol. 4474: G. Prencipe, S. Zaks (Eds.), *Structural Information and Communication Complexity*. XI, 342 pages. 2007.
- Vol. 4472: M. Haindl, J. Kittler, F. Roli (Eds.), *Multiple Classifier Systems*. XI, 524 pages. 2007.
- Vol. 4471: P. Cesar, K. Chorianopoulos, J.F. Jensen (Eds.), *Interactive TV: a Shared Experience*. XIII, 236 pages. 2007.
- Vol. 4470: Q. Wang, D. Pfahl, D.M. Raffo (Eds.), *Software Process Dynamics and Agility*. XI, 346 pages. 2007.
- Vol. 4468: M.M. Bonsangue, E.B. Johnsen (Eds.), *Formal Methods for Open Object-Based Distributed Systems*. X, 317 pages. 2007.
- Vol. 4467: A.L. Murphy, J. Vitek (Eds.), *Coordination Models and Languages*. X, 235 pages. 2007.
- Vol. 4465: T. Chahed, B. Tuffin (Eds.), *Network Control and Optimization*. XIII, 305 pages. 2007.
- Vol. 4464: E. Dawson, D.S. Wong (Eds.), *Information Security Practice and Experience*. XIII, 361 pages. 2007.
- Vol. 4463: I. Măndoiu, A. Zelikovsky (Eds.), *Bioinformatics Research and Applications*. XV, 653 pages. 2007. (Sublibrary LNBI).
- Vol. 4462: D. Sauveron, K. Markantonakis, A. Bilas, J.-J. Quisquater (Eds.), *Information Security Theory and Practices*. XII, 255 pages. 2007.
- Vol. 4459: C. Cérin, K.-C. Li (Eds.), *Advances in Grid and Pervasive Computing*. XVI, 759 pages. 2007.
- Vol. 4453: T. Speed, H. Huang (Eds.), *Research in Computational Molecular Biology*. XVI, 550 pages. 2007. (Sublibrary LNBI).
- Vol. 4452: M. Fasli, O. Shehory (Eds.), *Agent-Mediated Electronic Commerce*. VIII, 249 pages. 2007. (Sublibrary LNAI).
- Vol. 4451: T.S. Huang, A. Nijholt, M. Pantic, A. Pentland (Eds.), *Artificial Intelligence for Human Computing*. XVI, 359 pages. 2007. (Sublibrary LNAI).
- Vol. 4450: T. Okamoto, X. Wang (Eds.), *Public Key Cryptography – PKC 2007*. XIII, 491 pages. 2007.
- Vol. 4448: M. Giacobini et al. (Ed.), *Applications of Evolutionary Computing*. XXIII, 755 pages. 2007.
- Vol. 4447: E. Marchiori, J.H. Moore, J.C. Rajapakse (Eds.), *Evolutionary Computation, Machine Learning and Data Mining in Bioinformatics*. XI, 302 pages. 2007.
- Vol. 4446: C. Cotta, J. van Hemert (Eds.), *Evolutionary Computation in Combinatorial Optimization*. XII, 241 pages. 2007.
- Vol. 4445: M. Ebner, M. O'Neill, A. Ekárt, L. Vanneschi, A.I. Esparcia-Alcázar (Eds.), *Genetic Programming*. XI, 382 pages. 2007.
- Vol. 4444: T. Reps, M. Sagiv, J. Bauer (Eds.), *Program Analysis and Compilation, Theory and Practice*. X, 361 pages. 2007.
- Vol. 4443: R. Kotagiri, P.R. Krishna, M. Mohania, E. Nantajeewarawat (Eds.), *Advances in Databases: Concepts, Systems and Applications*. XXI, 1126 pages. 2007.
- Vol. 4440: B. Liblit, *Cooperative Bug Isolation*. XV, 101 pages. 2007.
- Vol. 4439: W. Abramowicz (Ed.), *Business Information Systems*. XV, 654 pages. 2007.
- Vol. 4438: L. Maicher, A. Sigel, L.M. Garshol (Eds.), *Leveraging the Semantics of Topic Maps*. X, 257 pages. 2007. (Sublibrary LNAI).
- Vol. 4433: E. Şahin, W.M. Spears, A.F.T. Winfield (Eds.), *Swarm Robotics*. XII, 221 pages. 2007.
- Vol. 4432: B. Beliczynski, A. Dzielinski, M. Iwanowski, B. Ribeiro (Eds.), *Adaptive and Natural Computing Algorithms, Part II*. XXVI, 761 pages. 2007.
- Vol. 4431: B. Beliczynski, A. Dzielinski, M. Iwanowski, B. Ribeiro (Eds.), *Adaptive and Natural Computing Algorithms, Part I*. XXV, 851 pages. 2007.
- Vol. 4430: C.C. Yang, D. Zeng, M. Chau, K. Chang, Q. Yang, X. Cheng, J. Wang, F.-Y. Wang, H. Chen (Eds.), *Intelligence and Security Informatics*. XII, 330 pages. 2007.
- Vol. 4429: R. Lu, J.H. Siekmann, C. Ullrich (Eds.), *Cognitive Systems*. X, 161 pages. 2007. (Sublibrary LNAI).
- Vol. 4427: S. Uhlig, K. Papagiannaki, O. Bonaventure (Eds.), *Passive and Active Network Measurement*. XI, 274 pages. 2007.
- Vol. 4426: Z.-H. Zhou, H. Li, Q. Yang (Eds.), *Advances in Knowledge Discovery and Data Mining*. XXV, 1161 pages. 2007. (Sublibrary LNAI).
- Vol. 4425: G. Amati, C. Carpineto, G. Romano (Eds.), *Advances in Information Retrieval*. XIX, 759 pages. 2007.
- Vol. 4424: O. Grumberg, M. Huth (Eds.), *Tools and Algorithms for the Construction and Analysis of Systems*. XX, 738 pages. 2007.
- Vol. 4423: H. Seidl (Ed.), *Foundations of Software Science and Computational Structures*. XVI, 379 pages. 2007.
- Vol. 4422: M.B. Dwyer, A. Lopes (Eds.), *Fundamental Approaches to Software Engineering*. XV, 440 pages. 2007.

¥562.00元

Table of Contents

Session 1. Invited Talks

Fast Distributed Algorithms Via Primal-Dual	1
<i>Alessandro Panconesi</i>	
Time Optimal Gathering in Sensor Networks	7
<i>Luisa Gargano</i>	
Treewidth: Structure and Algorithms	11
<i>Hans L. Bodlaender</i>	

Session 2. Autonomous Systems: Graph Exploration

Fast Periodic Graph Exploration with Constant Memory	26
<i>Leszek Gąsieniec, Ralf Klasing, Russell Martin, Alfredo Navarra, and Xiaohui Zhang</i>	
Why Robots Need Maps	41
<i>Miroslaw Dynia, Jakub Łopuszański, and Christian Schindelhauer</i>	
Graph Searching with Advice	51
<i>Nicolas Nisse and David Soguet</i>	

Session 3. Distributed Algorithms: Fault Tolerance

From Renaming to Set Agreement	66
<i>Achour Mostefaoui, Michel Raynal, and Corentin Travers</i>	
A Self-stabilizing Algorithm for the Median Problem in Partial Rectangular Grids and Their Relatives	81
<i>Victor Chepoi, Tristan Fevat, Emmanuel Godard, and Yann Vaxès</i>	
A New Self-stabilizing Maximal Matching Algorithm	96
<i>Fredrik Manne, Morten Mjelde, Laurence Pilard, and Sébastien Tixeuil</i>	

Session 4. Distributed Algorithms and Data Structures

Labeling Schemes with Queries	109
<i>Amos Korman and Shay Kutten</i>	

A Simple Optimistic Skiplist Algorithm	124
<i>Maurice Herlihy, Yossi Lev, Victor Luchangco, and Nir Shavit</i>	
Data Aggregation in Sensor Networks: Balancing Communication and Delay Costs	139
<i>Peter Korteweg, Alberto Marchetti-Spaccamela, Leen Stougie, and Andrea Vitaletti</i>	

Session 5. Autonomous Systems: Location Problems

Optimal Moves for Gossiping Among Mobile Agents	151
<i>Tomoko Suzuki, Taisuke Izumi, Fukuhito Ooshita, Hirotsugu Kakugawa, and Toshimitsu Masuzawa</i>	
Swing Words to Make Circle Formation Quiescent	166
<i>Yoann Dieudonné and Franck Petit</i>	
Distributed Algorithms for Partitioning a Swarm of Autonomous Mobile Robots	180
<i>Asaf Efrima and David Peleg</i>	

Session 6. Wireless Networks

Local Edge Colouring of Yao-Like Subgraphs of Unit Disk Graphs	195
<i>Jurek Czyzowicz, Stefan Dobrev, Evangelos Kranakis, Jaroslav Opatrny, and Jorge Urrutia</i>	
Proxy Assignments for Filling Gaps in Wireless Ad-Hoc Lattice Computers	208
<i>Tiziana Calamoneri, Emanuele G. Fusco, Anil M. Shende, and Sunil M. Shende</i>	
Location Oblivious Distributed Unit Disk Graph Coloring	222
<i>Mathieu Couture, Michel Barbeau, Prosenjit Bose, Paz Carmi, and Evangelos Kranakis</i>	

Session 7. Communication Networks: Fault Tolerance

Edge Fault-Diameter of Cartesian Product of Graphs	234
<i>Iztok Banič and Janez Žerovnik</i>	
Rapid Almost-Complete Broadcasting in Faulty Networks	246
<i>Rastislav Kráľovič and Richard Kráľovič</i>	
Design of Minimal Fault Tolerant On-Board Networks: Practical Constructions	261
<i>Jean-Claude Bermond, Frédéric Giroire, and Stéphane Pérennes</i>	

Session 8. Autonomous Systems: Fault Tolerance

Dynamic Compass Models and Gathering Algorithms for Autonomous Mobile Robots	274
<i>Yoshiaki Katayama, Yuichi Tomida, Hiroyuki Imazu, Nobuhiro Inuzuka, and Koichi Wada</i>	
Fault-Tolerant Simulation of Message-Passing Algorithms by Mobile Agents	289
<i>Shantanu Das, Paola Flocchini, Nicola Santoro, and Masafumi Yamashita</i>	

Session 9. Communication Networks: Parallel Computing and Selfish Routing

Optimal Conclusive Sets for Comparator Networks	304
<i>Guy Even, Tamir Levi, and Ami Litman</i>	
Selfish Routing with Oblivious Users	318
<i>George Karakostas, Taeyon Kim, Anastasios Viglas, and Hao Xia</i>	
Upper Bounds and Algorithms for Parallel Knock-Out Numbers	328
<i>Hajo Broersma, Matthew Johnson, and Daniël Paulusma</i>	
Author Index	341

Fast Distributed Algorithms Via Primal-Dual (Extended Abstract)

Alessandro Panconesi

Informatica, Sapienza University
via Salaria 113, 00198 Roma, Italy
ale@di.uniroma1.it

*When a trick works once, it is a trick.
If it works twice, it is a technique.
If it works three times, it is a method.*

Juris Hartmanis

1 Introduction

We would like to discuss what seems to be a general methodology to develop fast distributed algorithms for optimization problems on graphs, based on the primal-dual schema. The kind of problems we have in mind are of the following type. We have a synchronous, message-passing network that is to compute a global function of its own topology. Examples of such functions are maximal independent sets, vertex and edge colorings, small dominating sets, vertex covers and so on. Crucially, nodes only know their neighbours and have very little or no global information. In what follows, the only global information allowed will be n , the number of nodes in the network (or an upper bound on it). In this setting the running time of a protocol is given by the number of communication rounds needed to compute the output. By the end of the algorithm each node or edge will have decided its final status: its own color, whether or not to be part of the dominating set etc. In many situations of interest the cost of communication is orders of magnitude larger than local computation cost, and the model provides a rough, but quite useful, quantitative framework to develop and analyze interesting algorithms.

The combinatorial objects that we are interested in computing are useful both on theoretical and practical grounds. Small dominating sets for instance are the method of choice to set up the routing infrastructure of ad-hoc networks (the so-called backbone). Edge colorings have been repeatedly used to parallelize data transfers in distributed architectures, and so on. Maximal independent sets on the other hand appear to be a basic building block of many distributed algorithms.

Note the basic challenge here: If a protocol runs for t many rounds, each node will be able to collect information only from nodes at distance t . If t is much smaller than the diameter of the network, what we are trying to do is to compute a global function of the entire network, based on local information alone.

Distributed algorithms for graph problems is a very wide and active area. In what follows we do not try to be complete or encyclopaedic. Rather, we focus on the issues

that are immediately relevant to our topic. This means that we will not do a proper job of acknowledging the large and relevant body of brilliant literature that exists, and we offer our apologies in advance for this lack of completeness. The papers we cite and the references therein provide a good entry point for the research areas we will be discussing.

2 Distributed Algorithms Via the Primal-Dual Schema

The primal-dual schema is a very powerful methodology to develop efficient algorithms for combinatorial optimization problems. In recent years it has been applied with good success to NP-hard problems for which it yields many sophisticated approximation algorithms with performance guarantee. The main thesis of this talk is that in general a primal-dual algorithm exhibits certain locality properties that make it amenable to a fast, distributed implementation. This point of view is cogently developed in the PhD dissertation [8] which is the basis of the discussion to follow. We shall outline the method by discussing the example of vertex cover. The algorithm we describe is the first example of this general methodology [6]. The algorithm was developed in the dark days when people were interested in PRAM algorithms, but the solution is in fact fully distributed.

As it is well-known, in the vertex cover problem we are given an undirected network and we are to compute a so-called cover, i.e. a set of vertices such that, for every edge at least one of the two endpoints lies in the cover. Among all covers, we are interested in computing one of the smallest possible size. When vertices have positive integer costs we seek a cover of the smallest possible aggregate cost. This problem is NP-hard even with unit costs. What we will do is the following: (a) develop a sequential 2-approximated algorithm for it, and then (b) show how to parallelize it efficiently by means of a distributed implementation. Note that there are two separate issues here: to make the process distributed *and* fast. To do so, one must be able to exploit the inherent parallelism.

We begin by formulating the problem as an integer program (IP):

$$\min \sum_{v \in V} c(v) \cdot x_v \quad (\text{IP})$$

$$\text{s.t. } x_v + x_u \geq 1 \quad \forall e = (u, v) \in E \quad (1)$$

$$x_v \in \{0, 1\} \quad \forall v \in V \quad (2)$$

The cover is defined to be the set of all vertices v such that the corresponding binary indicator variable $x_v = 1$. The set of constraints (1) ensure that it is indeed a cover, for at least one endpoint for every edge must be included in it.

We now let (LP) be the standard LP relaxation obtained from (IP) by replacing the constraints (2) by

$$x_v \geq 0 \quad \forall v \in V \quad (3)$$

In the linear-programming dual of (LP) we associate a variable α_e with constraint (1) for every $e \in E$. The linear programming dual (D) of (LP) is then

$$\max \sum_{e \in E} \alpha_e \quad (\text{D})$$

$$\text{s.t.} \quad \sum_{e=(u,v) \in E} \alpha_e \leq c(v) \quad \forall v \in V \quad (4)$$

$$\alpha_e \geq 0 \quad \forall e \in E \quad (5)$$

We will build a cover working with the dual variables. Note that we have a constraint of type (4) for every node v , denoted as $(4)_v$. Consider the following continuous process. We let all the variables α_e grow at uniform speed. Sooner or later a constraint of type (4) will be satisfied with equality. If $(4)_v$ is the constraint we say that it becomes *tight*. When $(4)_v$ becomes tight we add v to the cover that we are computing. We do this by setting $x_v = 1$ (initially all primal variables are set to 0). Then we freeze the values α_e of the edges incident to v . The α -values of frozen edges stop growing, so that the constraint considered remains tight. The process continues with the remaining edges, until all edges are frozen.

At this point we have a set of vertices C containing all vertices v such that $x_v = 1$. We want to show that (a) it is a cover and that (b) its size is at most twice the optimum cost. To see why it is a cover, suppose not. But then there is an edge $e = uv$ which is not covered, i.e. $x_u = x_v = 0$. This means that the constraints $(4)_u$ and $(4)_v$ corresponding to u and v are not tight and α_e can continue to grow, a contradiction.

The cost of C is upper-bounded by twice the cost of the dual solution:

$$\sum_{v \in V} c(v) x_v = \sum_{v \in C} c(v) \leq \sum_{v \in C} \sum_{e=(u,v) \in E} \alpha_e \leq 2 \sum_{e \in E} \alpha_e.$$

You can think of this chain of inequalities in the following way. At the end of the algorithm we have a value α_e for every edge in the network. By doubling each α_e we have enough “money” to pay for the cost of vertices we put in C . This is true locally: for every $v \in C$, since v is tight, the sum of the α_e ’s, where e is incident to v , is equal to its cost. Since we doubled every α_e , each edge e has enough cash to pay the cost of both vertices it is incident to. Thus, twice the sum of the α_e ’s covers the cost of the solution we computed.

That the solution computed is 2-approximate follows by weak duality. Weak duality states that any dual feasible solution is no more than any primal solution. That is, denoting with z^* the optimal value of the primal solution we have

$$\sum_{e \in E} \alpha_e \leq z^*.$$

The primal is a relaxation of the original integer program and thus $z^* \leq \text{opt}$. The claim follows.

The continuous process above can be easily turned into a discrete one that runs in polynomial-time. The problem is how to make it both distributed and fast. It is apparent