Lecture Notes in Computer Science 11085

More information about this series at http://www.springer.com/series/7407

Zvi Lotker · Boaz Patt-Shamir (Eds.)

Structural Information and Communication Complexity

25th International Colloquium, SIROCCO 2018
Ma'ale HaHamisha, Israel, June 18–21, 2018
Revised Selected Papers

Springer

Editors
Zvi Lotker
Ben-Gurion University of the Negev
Beer-Sheva, Israel

Boaz Patt-Shamir
Tel Aviv University
Tel Aviv, Israel

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-030-01324-0 ISBN 978-3-030-01325-7 (eBook)
https://doi.org/10.1007/978-3-030-01325-7

Library of Congress Control Number: 2018955971

LNCS Sublibrary: SL1 – Theoretical Computer Science and General Issues

© Springer Nature Switzerland AG 2018
This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.
The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.
The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Preface

This volume contains the papers presented at the 25th International Colloquium on Structural Information and Communication Complexity (SIROCCO 2018). This year was particularly special for SIROCCO, as it was the half-jubilee of SIROCCO. The conference and celebration were held during June 18–21, in Ma'ale HaHamisha, Israel.

This year we received 46 submissions in response to the call for papers. Each submission was reviewed by at least three reviewers; we had a total of 21 Program Committee members and 49 external reviewers. The Program Committee decided to accept 23 papers for regular presentations, and eight papers for brief announcements. All these papers are included in this volume.

In addition the conference program included five additional talks: four keynote talks, and one talk by the winner of the SIROCCO Prize for Innovation in Distributed Computing. The invited speakers were Kurt Melhorn, David Peleg, Claire Mathieu, and Seth Pettie. Additionally, there was a talk by Zvi Lotker, the recipient of the 2018 SIROCCO Prize for Innovation in Distributed Computing. Papers representing these talks are also included in this volume.

The Program Committee selected the following paper as the winner of the SIROCCO 2018 Best Student Paper Award: "Mixed Fault Tolerance in Server Assignment: Combining Reinforcement and Backup," by Tal Navon and David Peleg. Selected papers will also appear in a special issue of the *Theoretical Computer Science* journal devoted to SIROCCO 2018.

We would like to thank all of the authors for their high-quality submissions and all of the speakers for their excellent talks. We are grateful to the Program Committee and all external reviewers for their efforts in putting together a great conference program, to the Steering Committee, chaired by Andrzej Pelc, for their help and support, and to everyone who was involved in the local organization for making it possible to have SIROCCO 2018 in lovely Israel. Thanks also to Michael Borokhovich for his work as the proceedings chair.

August 2018

Zvi Lotker
Boaz Patt-Shamir

Organization

Conference Committee

General Chair

Boaz Patt-Shamir	Tel Aviv University, Israel

Web

Chen Avin	Ben Gurion University, Israel

Proceedings

Michael Borokhovich

Jubilee Celebration

Pierre Fraigniaud	CNRS, France

Publicity

Moti Medina	Ben Gurion University, Israel

Steering Committee

Shantanu Das	Aix-Marseille University, France
Andrzej Pelc (Chair)	UQO
Nicola Santoro	Carleton University, Canada
Christian Scheideler	University of Paderborn, Germany
Sébastien Tixeuil	University of Paris 6, France
Jukka Suomela	Aalto University, Finland

Program Committee

Ittai Abraham	Hebrew University, Israel
Amotz Bar-Noy	CUNY, USA
Jurek Czyzowicz	University of Québec Outaouais, Canada
Taisuke Izumi	Nagoya IT, Japan
Paola Flocchini	University of Ottawa, Canada
Luisa Gargano	University of Salerno, Italy
Leszek Gąsieniec	University Liverpool, UK
Danny Hendler	Ben-Gurion University, Israel
Qiang-Sheng Hua	HUST
Adrian Kosowski	Inria, France
Zvi Lotker (Program Co-chair)	Ben-Gurion University, Israel

Euripides Markou	University of Thessaly, Greece
Gopal Pandurangan	University of Houston, USA
Merav Parter	Weizmann Institute of Science, Israel
Boaz Patt-Shamir (Program Co-chair)	Tel Aviv University, Israel
Sriram Pemmaraju	University of Iowa, USA
Harald Räcke	TU Munich, Germany
Sergio Rajsbaum	UNAM
Ivan Rappaport	University of Chile
Adi Rosén	CNRS and University of Paris Diderot, France
Ladislav Stacho	Simon Fraser University, Canada
Lewis Tseng	Boston College, USA

Additional Reviewers

James Aspnes
Przemysław Uznański
Yossi Azar
Lelia Blin
Lucas Boczkowski
Benedikt Bollig
Costas Busch
Soumyottam Chatterjee
Huda Chuangpishit
Gennaro Cordasco
Shantanu Das
Giuseppe Di Luna
Gabriele Di Stefano
Nguyen Dinh Pham
Michal Dory
Arun Ganesh
Vijay Garg
Konstantinos Georgiou
Sukumar Ghosh
Robert Gmyr
Emmanuel Godard
Maurice Herlihy
Riko Jacob
Eleni Kanellou
Ryan Killick
David Kirkpatrick
Christian Konrad
Kishori Konwar
Rastislav Kralovic
Arnaud Labourel
Nikos Leonardos
Reut Levi
Pedro Montealegre
William Moses
Yoram Moses
Alfredo Navarra
Shreyas Pai
Paolo Penna
Joseph Peters
Franck Petit
Adele Rescigno
Talal Riaz
Eric Ruppert
Dimitris Sakavalas
Stefan Schmid
Paulo Sérgio Almeida
Gokarna Sharma
Giovanni Viglietta
Mengchuan Zou

Sponsoring Institutions

We gratefully acknowledge the generous support that was provided to the SIROCCO 2018 Conference by the Israel Science Foundation, the Israel Ministry of Science Technology and Space, and Springer.

Invited Talks (Abstracts)

The Distributed Lovász Local Lemma Problem

Seth Pettie

University of Michigan
seth@pettie.net

Abstract. The Lovász Local Lemma (LLL) is a well known tool to prove the *existence* of a combinatorial object, by showing that a randomly chosen object satisfies some property with positive (but small) probability. The LLL has been applied in numerous areas, e.g., to compute graph colorings, packet-routing schedules, and satisfying assignments to CNF-SAT formulae. *Algorithmic* versions of the LLL can compute such objects efficiently, in polynomial time.

In this talk I will define the *Distributed LLL* problem and survey its role in algorithm design and complexity theory in the LOCAL model. Among the take-away messages from this talk are the following:

- The LLL is instrumental for designing fast algorithms for edge-coloring, defective coloring, frugal coloring, and other problems.
- There is an exponential gap between randomized and deterministic complexity in the LOCAL model, and the Distributed LLL is the foremost problem realizing this gap.
- The randomized Distributed LLL is *complete* for sublogarithmic randomized time. In particular, any sublogarithmic time algorithm for a locally checkable labeling problem can be automatically sped up to match the time of the Distributed LLL.
- The *deterministic* complexity of the Distributed LLL is inextricably linked to computing network decompositions deterministically. On the one hand, network decompositions are the basis of the fastest Distributed LLL algorithms. Conversely, a deterministic $\text{polylog}(n)$ LLL algorithm implies a deterministic $(\text{polylog}(n), \text{polylog}(n))$-network decomposition algorithm. (The Distributed LLL is PSLOCAL-hard.)

Keywords: LOCAL model · Probabilistic method · Graph coloring
Lovász local lemma

Supported by NSF grants CCF-1514383 and CCF-1637546.

On Fair Division for Indivisible Goods

Kurt Mehlhorn

Max Planck Institute for Informatics, Saarland Informatics Campus (SIC),
66123 Saarbrücken, Germany

We consider the task of dividing indivisible goods among a set of n agents in a fair manner. More precisely, we consider the following scenario. We have m distinct goods. Goods are available in several copies or items; there are k_j items of good j. The agents have decreasing utilities for the different items of a good, i.e., for all i and j

$$u_{i,j,1} \geq u_{i,j,2} \geq \ldots \geq u_{i,j,k_j}.$$

An allocation assigns the items to the agents. For an allocation x, x_i denotes the multi-set of items assigned to agent i, and $m(j, x_i)$ denotes the multiplicity of j in x_i. The total utility of bundle x_i under valuation u_i is given by

$$u_i(x_i) := \sum_j \sum_{1 \leq \ell \leq m(j,x_i)} u_{i,j,\ell}.$$

Each agent has a utility cap c_i. The utility of bundle x_i for agent i is defined as

$$\bar{u}_i(x_i) = \min(c_i, u_i(x_i)).$$

Our notion of fairness is *Nash social welfare* (NSW) [Nas50], i.e., the goal is to maximize the geometric mean

$$\mathrm{NSW}(x) = \left(\prod_{1 \leq i \leq n} \bar{u}_i(x_i) \right)^{1/n}$$

of the capped utilities. All utilities and caps are assumed to be integers.

The problem has a long history. For divisible goods, maximizing Nash Social Welfare (NSW) for any set of valuation functions can be expressed via an Eisenberg-Gale program [EG59]. For *additive valuations* ($c_i = \infty$ for each agent i and $k_j = 1$ for each good j) this program is equivalent to a Fisher market with identical budgets and maximizing NSW is achieved via the well-known fairness notion of competitive equilibrium with equal incomes (CEEI) [Mou03].

For indivisible goods, the problem is NP-complete [NNRR14] and APX-hard [Lee17]. Several constant-factor approximation algorithms are known for the case of additive valuations. They use different approaches.

The first one was pioneered by Cole and Gkatzelis [CG15] and uses spending-restricted Fisher markets. Each agent comes with one unit of money to the market. Spending is restricted in the sense that no seller wants to earn more than one unit of money. If the price p of a good is higher than one in equilibrium, only a fraction $1/p$ of the good is sold. Cole and Gkatzelis showed how to compute a spending restricted equilibrium in polynomial time and how to round its allocation to an integral allocation with good NSW. In the original paper they obtained an approximation ratio of $2e^{1/e} \approx 2.889$. Subsequent work [CDG+17] improved the ratio to 2.

The second approach is via stable polynomials. Anari et al. [AGSS17] obtained an approximation factor of e.

The third approach is via integral allocations that are Pareto-optimal and envy-free up to one good introduced by Barman et al. [BMV17]. Let x_i be the set of goods that are allocated to agent i. An allocation is envy-free up to one good if for any two agents i and k, there is a good j such that $u_i(x_k - j) \leq u_i(x_i)$, i.e., after removal of one good from k's bundle its value for i is no larger than the value of i's bundle for i. Caragiannis et al. [CKM+16] have shown that an allocation maximizing NSW is Pareto-optimal and envy-free up to one good. Barman et al. [BMV17] studied allocations that are Pareto-optimal and almost envy-free up to one good (ε-EF1), i.e., $u_i(x_k - g) \leq (1+\varepsilon)u_i(x_i)$, where ε is an approximation parameter. They showed that a Pareto-optimal and ε-EF1 allocation approximates NSW up to a factor $e^{1/e} + \varepsilon \approx 1.445 + \varepsilon$. They also showed how to compute such an allocation in polynomial time.

There are also constant-factor approximation algorithms beyond additive utilities.

Garg et al. [GHM18] studied budget-additive utilities ($k_j = 1$ for all goods j and arbitrary c_i). They showed how to generalize the Fisher market approach and obtained an $2e^{1/2e} \approx 2.404$-approximation.

Anari et al. [AMGV18] investigated multi-item concave utilities ($c_i = \infty$ for all i and k_j arbitrary). They generalized the Fisher market and the stable polynomial approach and obtained approximation factors of 2 and e^2, respectively.

In [CCG+18] is shown that the envy-free allocation approach can handle both generalizations combined and yields an approximation ratio of $e^{1/e} + \varepsilon \approx 1.445 + \varepsilon$. The approach via envy-freeness does not only yield better approximation ratios, it is also easier to state and to analyse.

References

[AGSS17] Anari, N., Gharan, S.O., Saberi, A., Singh, M.: Nash social welfare, matrix permanent, and stable polynomials. In: ITCS, pp. 36:1–36:12 (2017)

[AMGV18] Anari, N., Mai, T., Gharan, S.O., Vazirani, V.V.: Nash social welfare for indivisible items under separable, piecewise-linear concave utilities. In: SODA, pp. 2274–2290 (2018)

[BMV17] Barman, S., Murthy, S.K.K., Vaish, R.: Finding fair and efficient allocations. *CoRR*, abs/1707.04731 (2017). To appear in EC 2018

[CCG+18] Cheung, Y.K., Chaudhuri, B., Garg, J., Garg, N., Hoefer, M., Mehlhorn, K.: On Fair Division of Indivisible Items. *CoRR*, abs/1805.06232 (2018)

[CDG+17] Cole, R., Devanur, N.R., Gkatzelis, V., Jain, K., Mai, T., Vazirani, V.V., Yazdanbod, S.: Convex program duality, fisher markets, and Nash social welfare. In: EC, pp. 459–460 (2017)

[CG15] Cole, R., Gkatzelis, V.: Approximating the Nash social welfare with indivisible items. In: STOC, pp. 371–380 (2015)

[CKM+16] Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A.D., Shah, N., Wang, J.: The unreasonable fairness of maximum Nash welfare. In: EC, pp. 305–322 (2016)

[EG59] Eisenberg, E., Gale, D.: Consensus of subjective probabilities: the pari-mutuel method. Ann. Math. Statist. **30**, 165–168 (1959)

[GHM18] Garg, J., Hoefer, M., Mehlhorn, K.: Approximating the nash social welfare with budget-additive valuations. In: SODA 2018, pp. 2326–2340 (2018)

[Lee17] Lee, E.: APX-hardness of maximizing Nash social welfare with indivisible items. Inf. Process. Lett. **122**, 17–20 (2017)

[Mou03] Moulin, H.: Fair Division and Collective Welfare. MIT Press (2003)

[Nas50] Nash, J.: The bargaining problem. Econometrica **18**, 155–162 (1950)

[NNRR14] Nguyen, N.-T., Nguyen, T.T., Roos, M., Rothe, J.: Computational complexity and approximability of social welfare optimization in multiagent resource allocation. Autonom. Agents Multi-Agent Syst. **28**(2), 256–289 (2014)

College Admissions in Practice

Claire Mathieu

CNRS, Paris, France

Abstract. The Gale-Shapley algorithm is the standard method in practice for stable marriage in large matching markets, but must be adapted to the constraints of each situation. We study the design of college admissions in a setting with the following features and constraints:

- Lack of trust in the platform: students worry that the rankings of students by schools will factor in their own ranking of schools
- Simplicity: the general public must be able to understand the method
- Transparency: the final result must not be given as a black box but come with an “explanation” that helps rebuild trust
- Quotas: schools have a legal obligation to respect certain quotas of student types. The types and quotas vary from school to school
- Housing: schools provide need-based housing to some of their students. Some students can only afford to attend if housing is provided. The offers must thus take into account both the students’ academic ranking and their ranking according to need.

I will present some preliminary work to address such issues, with an application to the French higher education admissions problem.

This is ongoing joint work with Hugo Gimbert.

Taking Turing to the Theater (Abstract of Award Lecture)

Zvi Lotker

Ben Gurion University in Israel
`zvilo@bgu.ac.il`

Abstract. Computer science has grown out of the seed of imitation. From von Neumann's machine to the famous Turing test, which sparked the field of AI, algorithms have always tried to imitate humans and nature. Examples of such "imitation algorithms" are simulated annealing which imitates thermodynamics, genetic algorithms which imitate biology, or deep learning which imitates human learning.

In this talk, I describe an algorithm which imitates human psychology. Specifically, I discuss *M* algorithms, which serve as a simple example of psychology-based imitation algorithms. The *M* algorithm is one of the simplest natural language processing (NLP) algorithms.

Respecting the long tradition of imitation algorithms, the *M* algorithm is simple yet powerful. Like other imitation algorithms, the *M* algorithm is able to efficiently solve difficult problems. The *M* algorithm pinpoints critical events in films, theater productions, and other scripts, revealing the rhythm of the texts.

At first glance, when trying to design an algorithm which pinpoints critical events of a text, it seems necessary for the algorithm to understand the complete text. Additionally, it would be expected that all layers of the narrative, background information, etc., would also be necessary. In short, it would be expected that the algorithm would imitate the human process of comprehending a text.

Surprisingly, the *M* algorithm utilizes the structure of the complete text itself without understanding even a *single* word, sentence, or character in order to discover critical events. The content of the narrative is not necessary for the algorithm to work. Other than an awareness of the illusion of time, borrowed from psychology, the *M* algorithm circumvents the human process of reading.

In the link below, we can see the computerized summary of several movies and relevant data. The *M* algorithm extracted the critical points on all those movies. As you can see these synopsis provides an "executive" summary of the movies. https://zvilotker.myportfolio.com/psychological-alg.

This talk is based on my upcoming book (in process).

Contents

Invited Talks and Brief Announcments

Full Papers