THEORY OF COMPUTING, Volume 19 (5), 2023, pp. 1–4
www.theoryofcomputing.org

SPECIAL ISSUE: APPROX-RANDOM 2020

Guest Editors' Foreword

Jarosłav Byrka Raghu Meka

November 14, 2023

This collection contains the expanded and fully refereed versions of selected papers presented at the 24th International Conference on Randomization and Computation (RANDOM 2020), and the 23rd International Conference on Approximation Algorithms for Combinatorial Optimization Problems (APPROX 2020). Due to COVID-related travel restrictions, the conferences were jointly organized in a virtual setup on August 17–19, 2020.

The proceedings of the meetings was published in the Dagstuhl LIPIcs series. APPROX published 34 contributed papers and RANDOM published 30 contributed papers, selected by the program committees. The total number of submissions was 67 for each conference.

In consultation with the RANDOM PC, PC chair and guest editor Raghu Meka invited three papers to this Special Issue and the authors of the following two papers accepted the invitation:

- "Iterated Decomposition of Biased Permutations Via New Bounds on the Spectral Gap of Markov chains" by Sarah Miracle, Amanda Streib and Noah Streib
- "Reaching a Consensus on Random Networks: The Power of Few" by Linh Tran and Van Vu.

In consultation with the APPROX PC, PC chair and guest editor Jarosław Byrka invited the following three papers to this Special Issue.

- "Pinning Down the Strong Wilber 1 Bound for Binary Search Trees" by Parinya Chalermsook, Julia Chuzhoy and Thatchaphol Saranurak
- "Maximizing the Correlation: Extending Grothendieck's Inequality to Large Domains" by Dor Katzelnick and Roy Schwartz

ACM Classification: G.3, F.2

AMS Classification: 68Q25

Key words and phrases: foreword, special issue, RANDOM, APPROX, RANDOM 2019, APPROX 2020

 DOI: 10.4086/toc.2023.v019a005

- "Parametrized Metrical Task Systems" by Sébastien Bubeck and Yuval Rabani.

We would like to thank the authors for their contributions and the anonymous referees for their hard work that helped improve this issue. We are especially indebted to the referees for their detailed and high quality reviews.

It was a pleasure to edit this Special Issue for Theory of Computing.

APPROX 2020 Program Committee

Nikhil Bansal, CWI & TU Eindhoven
Jarosław Byrka, U. Wrocław, PC chair
Andreas Emil Feldmann, Charles U., Prague
Naveen Garg, IIT Delhi
Anupam Gupta, Carnegie Mellon
Pasin Manurangsi, Google Research
Evangelos Markakis, AUEB, Athens
Nicole Megow, U. Bremen
Marcin Mucha, U. Warsaw
Harald Räcke, TU Munich
Laura Sanità, TU Eindhoven & U. Waterloo
Chaitanya Swamy, U. Waterloo
Jakub Tarnawski, Microsoft Research
Anke van Zuylen, William and Mary
David Williamson, Cornell

RANDOM 2020 Program Committee

Nima Anari, Stanford
Eshan Chattopadhyay, Cornell
Gil Cohen, Tel Aviv U.
Parikshit Gopalan, VmWare Research
Prahladh Harsha, Tata Institute of Fundamental Research
Sam Hopkins, UC Berkeley
Valentine Kabanets, Simon Fraser Un.
Gautam Kamath, U. Waterloo
Tali Kaufman, Bar-Ilan U.
Yin-Tat Lee, U. Washington
Sepideh Mahabadi, TTIC Chicago
Raghu Meka, UCLA, PC chair
Jelani Nelson, UC Berkeley
Ryan O'Donnell, Carnegie Mellon
Ilya Razenshteyn, Microsoft Research
Barna Saha, UC Berkeley
Tselil Schramm, Stanford
Madhu Sudan, Harvard
Avishay Tal, UC Berkeley
Eric Vigoda, Georgia Tech
Mary Wootters, Stanford

Brief synopses of the papers in the Special Issue

- "Reaching a Consensus on Random Networks: The Power of Few" by Linh Tran and Van Vu.

  This paper studies a classical question in population dynamics: Imagine we have n people connected by a graph, and each node starts with a red or blue color. Each subsequent day, every person revises their color based on some function of the colors of their neighbors. Let us say that a color wins if all people are of that color. How sensitive is the winning color to the initial configuration? The paper shows, surprisingly, that under the natural 'majority rule' (a vertex adopts the color of the majority of its neighbors), if the graph is a random one, then flipping a constant number of vertices near the decision boundary will change the color with a high probability. That is, if there are $n/2+c$ nodes of one color at the beginning, then that color is very likely to win. The main point is that the advantage needed to make one color the overwhelming favorite does not depend on$n$. For instance, if there are $n/2+5$ red nodes initially, red wins within four days with probability greater than .9. The key difficulty in analyzing the dynamics is that after one step, the colors of the nodes and the random graph are dependent. Hence, it is difficult to apply concentration inequalities. The authors find a clever way to overcome such dependency issues and obtain a tight analysis of the dynamics.

- "Pinning Down the Strong Wilber 1 Bound for Binary Search Trees" by Parinya Chalermsook, Julia Chuzhoy and Thatchaphol Saranurak

  The Dynamic Optimality conjecture of Sleator and Tarjan is relatively easy to state, but remains one of the most intriguing open problems in the area of data structures. Here it is: we want to maintain a dynamic binary search tree (BST) as some set of elements is accessed via a request sequence. Since the access sequence may have some structure, we are allowed to change the BST in some local ways (via "rotations") to reduce the access cost. The Dynamic Optimality conjecture says: for each sequence of accesses, the total cost incurred by the Splay Tree data structure is within a constant factor of the optimal cost (the cost incurred by the best dynamic BST for *that* sequence). We can also ask a weaker question: does there exist any algorithm that performs almost optimally on each access sequence? We do not yet know. In order to get such a universally optimal algorithm, these questions force us to understand the power and inherent trade-offs of basic data structures like BSTs.

  What is the optimal cost of a dynamic BST on some sequence? Previous attempts to reason about this optimal cost use one of two bounds proved by Wilbur in 1986. And until recently it seemed conceivable that either of these bounds was within a constant factor of the optimal cost. This paper shows that the first bound (WB-1) can be a factor of nearly $\Omega(\log\log n)$ away from the optimal cost on some instances. [Lecomte and Weinstein independently proved the same gap.] Moreover, the paper gives another bound, called the Guillotine bound, which may be useful in resolving the conjecture. Finally, they ask the algorithmic question: given a sequence, how to compute the optimal cost on it? The final result of this paper is an algorithm that smoothly translates between an $O(\log\log n)$-approximation in polynomial time, and an exact algorithm in exponential time. We do not know how to compute a constant-approximation in polynomial time.

  The Dynamic Optimality conjecture still remains tantalizingly open, but this paper deepens our understanding of this fascinating question.

GUEST EDITORS

Jarosław Byrka
University of Wrocław
Poland
jby@cs.uni.wroc.pl
https://sites.google.com/cs.uni.wroc.pl/jby

Raghu Meka
UCLA
Los Angeles, CA, USA
raghum@cs.ucla.edu
https://samueli.ucla.edu/people/raghu-meka/