THEORY OF COMPUTING, Volume 17 (3), 2021, pp. 1–4
www.theoryofcomputing.org

SPECIAL ISSUE: APPROX-RANDOM 2019

Guest Editors' Foreword

Ivona Bezáková Noga Ron-Zewi László A. Végh

September 18, 2021

This collection contains the expanded and fully refereed versions of selected papers presented at the 23rd International Conference on Randomization and Computation (RANDOM 2019), and the 22nd International Conference on Approximation Algorithms for Combinatorial Optimization Problems (APPROX 2019). The conferences were colocated at the Massachusetts Institute of Technology (MIT), in Boston, MA, United States, September 20–22, 2019. In retrospect these were the last APPROX and RANDOM conferences held physically before the COVID-19 pandemic period, and one of the last chances to meet members of the ToC community in person!

The proceedings of the meetings was published in the Dagstuhl LIPIcs series. APPROX published 33 contributed papers and RANDOM published 39 contributed papers, selected by the program committees. The total number of submissions was 66 for each conference.

In consultation with the RANDOM PC, guest editors Ivona Bezáková and Noga Ron-Zewi invited the following four papers to the special issue.

- "Improved pseudorandom generators from pseudorandom multi-switching lemmas" by Rocco Servedio and Li-Yang Tan
- "The Large-Error Approximate Degree of AC^0" by Mark Bun and Justin Thaler
- "Deterministic Approximation of Random Walks in Small Space" by Jack Murtagh, Omer Reingold, Aaron Sidford, and Salil Vadhan.
- "Optimal Convergence Rate of Hamiltonian Monte Carlo for Strongly Logconcave Distributions" by Zongchen Chen and Santosh S. Vempala

ACM Classification: G.3, F.2

AMS Classification: 68Q25

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

 DOI: 10.4086/toc.2021.v017a003

Hamiltonian Monte Carlo (HMC) is the most widely used algorithm for sampling log-concave densities in both computer science and statistics. In general, its running time depends on how close the log-density function is to the quadratic functions. For the standard setting where the density $\exp(-f(x))$ satisfies $0 \prec \alpha \preceq \nabla^2 f(x) \preceq \beta$ for all x, it has been conjectured that the relaxation time of HMC is $O(\beta/\alpha)$. In 2017, Mangoubi and Smith gave the bound of $O((\beta/\alpha)^2)$ and this was later improved to $O((\beta/\alpha)^{3/2})$. This paper confirms the conjecture by proving the optimal bound $O(\beta/\alpha)$.

In consultation with the APPROX PC, PC chair and guest editor László Végh invited the following three papers to the special issue.

- "Max-Min Greedy Matching" by Alon Eden, Uriel Feige, and Michal Feldman.
  The online bipartite matching problem was introduced in a classical 1990 paper by Karp, Vazirani, and Vazirani. For a bipartite graph $G(V,U;E)$, the vertices in U are revealed one-by-one along with the incident edges, and every vertex in V can be irrevocably matched at the time it shows up. They showed that if we pick a random permutation π of V, and match each vertex to its first available neighbour in π (if there exists one), then in expectation we can obtain a matching of size at least $(1-1/\mathrm{e})$-times the size of the maximum matching in the graph.
  The paper introduces a natural and interesting two-player game variant of this problem, motivated by resource allocation and pricing problems in mechanism design. Given the bipartite graph $G(V,U;E)$, maximizer chooses a permutation π of V; subsequently, minimizer selects a permutation σ of U. The vertices in U then arrive in the order σ, and are matched to their first available neighbour in the order π.
  The matching obtained is always maximal and hence at least half the size of the maximum matching. It is not at all obvious if a better ratio can be achieved. The main result of the paper shows that maximizer can attain a ratio slightly above 0.51. The proof is via a careful combinatorial argument involving a maximal path cover in a directed graph associated with the bipartite graph. To motivate the argument, the authors first outline three natural approaches that fail for the problem.
- "Maximizing Covered Area in a Euclidean Plane with Connectivity Constraint" by Chien-Chung Huang, Mathieu Mari, Claire Mathieu, Joseph S. B. Mitchell, and Nabil H. Mustafa.
  The paper discusses a special case of the maximum area connected subset problem: Given a set of unit disks in the plane and an integer k, select k disks whose union is connected and has maximum area. This problem is motivated by applications such as router deployment and cancer genome studies. The authors first show that the problem is NP-hard and that the natural greedy algorithm performs poorly. Then, they present a clever greedy approach that adds two disks at a time. This yields a 1/2-approximation, and the analysis is tight. Further, a bicriteria PTAS is given by an adaptation of the classical dynamic programming of Arora and Mitchell. Additional results on geometric shapes other than the unit disks are also presented, inspiring interesting open problems.
- "Fast and Deterministic Approximations for k-Cut" by Kent Quanrud.
  In the minimum k-cut problem, we need to break an undirected graph into at least k connected components by removing a minimum-cost subset of edges. This is NP-hard if k is part of the input, and 2-approximation is the best one can hope for under the small-set expansion hypothesis. Previous results include a deterministic 2-approximation algorithm, as well as randomized $\tilde{O}(km)$-

approximation algorithms; the latter is obtained by $O(k)$ calls to Karger's minimum-cut algorithm. The main result of the paper is a deterministic $\tilde{O}(m/\varepsilon^2)$ approximation algorithm for computing a $(2+\varepsilon)$-approximate minimum weight k-cut. Apart from the dependence on ε, this is the best one could hope for.

This is achieved by using a multiplicate weights update (MWU) framework. Interesting new ideas are needed to adapt the framework to this particular setting. A new LP formulation is introduced; even though this has a substantially larger number of constraints than the classical one by Naor and Rabani, it is purely a covering LP, and therefore the MWU framework can be applied on the dual problem. Another nontrivial challange is to update the weight function implicitly, using a dynamic data structure.

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 providing timely reviews of high quality during the COVID-19 pandemic, often under continuous lockdowns and other related challenges. It was a pleasure to edit this Special Issue for Theory of Computing. We wish all health and safety, and we look forward to seeing you in person at various conferences.

GUEST EDITORS

Ivona Bezáková
Rochester Institute of Technology
Rochester, NY, U.S.A.
ib@cs.rit.edu

Noga Ron-Zewi
University of Haifa
Haifa, Israel
noga@cs.haifa.ac.il

László A. Végh
London School of Economics and Political Science
London, UK
l.vegh@lse.ac.uk