THEORY OF COMPUTING, Volume 16 (19), 2020, pp. 1–5
www.theoryofcomputing.org

SPECIAL ISSUE: CCC 2019

Guest Editor's Foreword

Yuval Filmus

December 31, 2020

This collection comprises the expanded and fully refereed versions of selected papers presented at the 34th Computational Complexity Conference (CCC 2019), held July 18–20, 2019 in New Brunswick, NJ, USA. These papers were selected by the Program Committee from among the 32 papers that appeared in the conference proceedings. Preliminary versions of the papers were presented at the conference, and the extended abstracts appeared in the proceedings of the conference, published by Dagstuhl Publishing, LIPIcs.

The CCC Program Committee selected 32 out of 99 submissions for presentation at the conference; of these, the six described below were invited to this Special Issue. These six papers were refereed in accordance with the rigorous standards of Theory of Computing.

- "UG-hardness to NP-hardness by losing half" by Amey Bhangale and Subhash Khot.

  Many hardness of approximation results in the literature are conditioned on Khot's celebrated Unique Games Conjecture, which states that it is NP-hard to distinguish instances of Unique Label Cover which are $(1-\varepsilon)$-satisfiable from those which are only ε-satisfiable. Recently, in a breakthrough result, Khot, Minzer and Safra (building on previous work together with Dinur and Kindler) proved the 2-to-1 conjecture, which states that it is NP-hard to distinguish instances of Unique Label Cover which are $(1/2-\varepsilon)$-satisfiable from those which are ε-satisfiable.

  It might seem that many hardness of approximation results which are conditional on the Unique Games Conjecture can be adapted to use the 2-to-1 theorem instead, losing a half in the resulting approximation ratio. However, it turns out that such results need a stronger promise in the $(1/2-\varepsilon)$-satisfiable case: instead of just requiring a $1/2-\varepsilon$ fraction of the constraints to be satisfiable, they

ACM Classification: F, F.2

AMS Classification: 68Qxx

Key words and phrases: foreword, special issue, CCC 2019

 DOI: 10.4086/toc.2020.v016a019

require the existence of a $1/2-\varepsilon$ fraction of the vertices (on one of the sides), all of whose adjacent constraints can be simultaneously satisfied. The main result of the paper is this stronger form of the 2-to-1 theorem. As applications, the authors prove the following hardness of approximation results:

1. Independent set is $O(d/\log^2 d)$-hard to approximate in d-regular graphs. This is tight up to constants due to a matching algorithm of Bansal, Gupta, and Guruganesh, and improves the previous known hardness of $O(d/\log^4 d)$, due to Chan.
2. Maximum acyclic subgraph is $2/3$-hard to approximate, improving on the previous known hardness of $14/15$, due to Austrin, Manokaran, and Wenner. The Unique Games Conjecture itself implies the optimal hardness of $1/2$, as shown by Guruswami, Manokaran, and Raghavendra, matching a trivial algorithm.
3. For any predicate P, if P-CSP has a (c,s)-SDP integrality gap, then P-CSP is $(c/2,s)$-hard to approximate. Raghavendra famously showed that the Unique Games Conjecture implies the optimal (c,s)-hardness.

- "Sherali–Adams strikes back" by Ryan O'Donnell and Tselil Schramm.

  How does linear programming compare with semidefinite programming? A case in point is the fundamental problem of Maximum Cut. Goemans and Williamson famously gave an 0.878-approximation algorithm based on semidefinite programming, which is tight assuming the Unique Games Conjecture. In contrast, linear programming relaxations are much weaker: Charikar, Makarychev, and Makarychev showed that the integrality gap of *any* linear programming relaxation of size 2^{n^δ} has an integrality gap of $2-\varepsilon$ on a random Δ-regular graph, for constant Δ.

  Should we, then, give up on linear programming? The main result of this paper shows that $n^{O(c/\log\Delta)}$ rounds of Sherali–Adams (a linear programming hierarchy) can certify, with high probability, that the maximum cut in a random Δ-regular graph is at most $1/2+2^{-c}$. This refutes the common assumption that certifying such bounds requires linear programs of truly exponential size (that is, $2^{\Omega(n)}$). The result is actually stronger — it suffices that all nontrivial eigenvalues of the graph be small enough, a condition that a random Δ-regular graph satisfies with high probability.

  Using similar techniques, the authors prove analogous results for random CSPs and for Independent Set and Vertex Cover on regular graphs.

- "Limits on the universal method for matrix multiplication" by Josh Alman. This paper received the Best Student Paper award.
  "Barriers for fast matrix multiplication from irreversibility" by Matthias Christandl, Peter Vrána, and Jeroen Zuiddam.

  We discuss these two papers together, since their main results are essentially identical.

  The area of fast matrix multiplication has recently seen a resurgence, with work of Stothers, Vassilevska-Willams, and Le Gall progressively improving the decades-old algorithm of Coppersmith and Winograd, from $\omega < 2.376$ to $\omega < 2.373$. Can this approach prove that $\omega = 2$? Ambainis, Filmus and Le Gall ruled this out, proving a barrier of 2.3725 for this particular approach.

The algorithm of Coppersmith and Winograd and its recent improvements all rely on an analysis of the so-called Coppersmith–Winograd tensor using Strassen's laser method. The result of Ambainis et al. shows that applying the laser method to this tensor cannot give a bound better than $\omega < 2.3725$. Alman and Vassilevska-Williams subsequently showed that even if the laser method is replaced by monomial degeneration (a vast generalization of the laser method), then the Coppersmith–Winograd tensor cannot be used to prove that $\omega = 2$.

In the present two papers, similar results are proved for an even more general method, degeneration. In this method, one takes a large tensor power of the starting tensor, and then degenerates it into a matrix multiplication tensor. The papers consider three different classes of tensors:

1. Coppersmith–Winograd tensors: these are the tensors used by Coppersmith and Winograd, as well as in all recent work. They are parametrized by an integer q.
2. Simple (or small) Coppersmith–Winograd tensors: these were used as a "warm-up" in the work of Coppersmith and Winograd, and are also parametrized by an integer q.
3. Cyclic group tensors, also known as reduced polynomial multiplication tensors: these arise from the group-theoretic approach to fast matrix multiplication, due to Cohn and Umans.

Proving $\omega = 2$ is ruled out using all of these tensors, except for the simple Coppersmith–Winograd tensor in the case $q = 2$ (this exception corresponds to an observation of Coppersmith and Winograd).

Alman proves his results using a parameter known as *asymptotic rank*, whereas Christandl et al. bound *asymptotic subrank*, which is a priori tighter. However, Strassen showed that for the tensors considered in the papers, the two parameters coincide, and moreover are given by a third parameter, the *minimum quantum functional*, which is easy to compute, and is used in both papers.

- "Fourier and circulant matrices are not rigid" by Zeev Dvir and Allen Liu.

  Matrix rigidity was invented by Valiant as an avenue toward proving superlinear size lower bounds on circuits of logarithmic depth for explicit functions. He showed that if a matrix A is $(n/\log\log n, n^{1+\varepsilon})$-rigid (Valiant-rigid for short), meaning that reducing its rank below $O(n/\log\log n)$ requires changing $\Omega(n^{1+\varepsilon})$ entries, then Ax cannot be computed using a circuit of size $O(n)$ and depth $O(\log n)$. Valiant showed that a random matrix is $(r,(n-r)^2)$-rigid, and in particular Valiant-rigid, but the best explicit constructions are only $(r,(n^2/r)\log(n/r))$-rigid.

  It has long been conjectured that certain familiar classes of matrices such as the Walsh–Hadamard matrices are Valiant-rigid. Recently, Alman and Williams surprisingly refuted this conjecture, showing that the Walsh–Hadamard matrices are not Valiant-rigid. The current paper shows that several broad classes of matrices over $\mathbb{C}$ are also not Valiant-rigid:

  1. Discrete Fourier Transform (DFT) matrices over any finite abelian group G. When $G = \mathbb{Z}_2^n$, this is the Walsh–Hadamard matrix. When G is a cyclic group, we get the classical DFT matrix.
  2. Circulant matrices. These are matrices whose rows consist of all possible rotations of the first row.

3. Toeplitz matrices. These are matrices in which all diagonals are constant.
4. G-circulant matrices for any finite abelian group G. These are $|G| \times |G|$ matrices in which the (x,y) entry $(x,y \in G)$ depends only on $x-y$. When G is a cyclic group, we get the classical circulant matrices.

Analogous results are also obtained over finite fields $\mathbb{F}_q$ under the assumption that $\gcd(|G|,q) = 1$.

- "Hardness magnification near state-of-the-art lower bounds" by Igor Carboni Oliveira, Ján Pich, and Rahul Santhanam.

  Hardness magnification is the phenomenon in which a weak lower bound on a specific problem implies a much stronger lower, often beyond the reach of current techniques. The paper proves several such results:

  1. If distinguishing Boolean functions whose circuit complexity is at least $2^{\beta n}$ from those whose circuit complexity is at most $2^{\beta n}/cn$ is hard for circuits of size $N^{1+\varepsilon}$ (where $N = 2^n$ is the input size), then NP does not have polynomial-size circuits.
  2. If distinguishing Boolean functions whose Kt complexity (a kind of Kolmogorov complexity taking into account the running time) is at least $2^{\beta n} + cn$ from those whose Kt complexity is at most $2^{\beta n}$ is hard for de Morgan formulas of size $N^{2+\varepsilon}$ or for U_2-formulas (also allowing XOR) of size $N^{3+\varepsilon}$, then EXP does not have polynomial-size formulas.

  Many other such results concerning Kt complexity are proved in the paper. The results on Kt complexity follow from techniques in a previous paper of the authors, while the result on circuit complexity is the main technical contribution.

  These hardness magnification results are complemented by several lower bounds of a similar form but involving different parameters:

  1. Neither de Morgan formulas of size $N^{2-\varepsilon}$ nor U_2-formulas of size $N^{3-\varepsilon}$ can distinguish functions whose Kt complexity is at least 2^{n-1} from those whose Kt complexity is at most $2^{(1-\delta)n}$.
  2. U_2-formulas of size $N^{2-\varepsilon}$ cannot distinguish functions whose Kt complexity is at least $2^{(\varepsilon/2)n-2}$ from those whose Kt complexity is at most Cn^2.
  3. Such formulas also cannot distinguish functions whose circuit complexity is at least $2^{(\varepsilon/2-o(1))n}$ from those whose circuit complexity is at most $n^{1+\delta}$.

I would like to thank the authors for their contributions, the CCC program committee for their initial reviews, Amir Shpilka for his fantastic work as chair of the program committee, László Babai for his advice on matters related to *Theory of Computing,* and the anonymous referees for their hard work. It was a pleasure to edit this Special Issue for Theory of Computing.

CCC 2019 Program Committee

Andrej Bogdanov (Chinese University of Hong Kong)

Irit Dinur (Weizmann Institute of Science)

Yuval Filmus (Technion — Israel Institute of Technology)

Pavel Hrubeš (Czech Academy of Sciences)

Valentine Kabanets (Simon Fraser University)

Gillat Kol (Princeton University)

Troy Lee (University of Technology Sydney)

Raghu Meka (University of California at Los Angeles)

Ramprasad Saptharishi (Tata Institute of Fundamental Research)

Amir Shpilka (Tel Aviv University) (Chair)

Madhu Sudan (Harvard University)

GUEST EDITOR

Yuval Filmus
The Henry and Marilyn Taub Faculty of Computer Science
Technion — Israel Institute of Technology
Haifa, Israel.
yuvalfi@cs.technion.ac.il
https://yuvalfilmus.cs.technion.ac.il/