THEORY OF COMPUTING, Volume 18 (14), 2022, pp. 1–4
www.theoryofcomputing.org

SPECIAL ISSUE: CCC 2020

Guest Editors' Foreword

Zeev Dvir   Avishay Tal

June 12, 2022

This collection comprises the expanded and fully refereed versions of selected papers presented at the 35th Computational Complexity Conference (CCC 2020), held July 28–30, 2020, online. These papers were selected by the Program Committee from among the 36 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 36 out of 96 submissions for presentation at the conference; of these, the five described below were invited to this Special Issue. These five papers were refereed in accordance with the rigorous standards of Theory of Computing.

- "Sign rank vs. Discrepancy" by Kaave Hosseini, Hamed Hatami and Shachar Lovett.

  The main result of the paper is an optimal separation between the sign-rank of a matrix, which is the minimum rank of a matrix of the same sign pattern, and the discrepancy of a matrix, which is the maximal correlation it has with a rectangle. More specifically, the authors exhibit a matrix of sign-rank 3 and discrepancy $\exp(-n)$. Prior to this work, the largest known separation was sign-rank $O(\log n)$ and discrepancy $\exp(-n)$. The authors further observe that sign-rank of 3 is lowest possible, as any matrix of sign-rank 1 or 2 has discrepancy at least $1/\mathrm{poly}(n)$. As it is known that sign-rank characterizes the *unbounded-error randomized communication complexity* model (i. e., error probability smaller than half), while discrepancy characterizes the *weakly bounded error* model (the

ACM Classification: F, F.2

AMS Classification: 68Qxx

Key words and phrases: foreword, special issue, CCC 2020

 DOI: 10.4086/toc.2022.v018a014

communication cost includes a “penalty” depending on how close to half is the error probability), the result also implies an optimal separation between these classes. Moreover, by known relations between discrepancy and *approximate rank*, the above also gives an optimal separation between approximate and sign ranks.

- “Hitting Sets Give Two-Sided Derandomization of Small Space” by Kuan Cheng and William Hoza.

  Pseudorandom generators and hitting set generators naturally allow the derandomization of two-sided and one-sided error probabilistic algorithms, respectively. Motivated to derandomize small-space computation and prove that $\mathsf{BPL} = \mathsf{L}$ or $\mathsf{RL} = \mathsf{L}$, researchers have been designing and analyzing pseudorandom generators (PRGs) and hitting set generators (HSGs) for the model of polynomial-width read-once oblivious branching programs. Indeed, space-efficient PRGs with logarithmic seed length for this model suffice to prove that $\mathsf{BPL} = \mathsf{L}$, whereas the weaker objects, HSGs, suffice to prove the weaker result $\mathsf{RL} = \mathsf{L}$. Both challenges remain elusive to date. The surprising result in this article is that hitting set generators with logarithmic seed length suffice for the derandomization of BPL! The proof relies on the technique of local consistency: obtaining estimates for the probabilities that sub-computations yield specific values, and checking that these estimates are locally consistent. This white-box reduction relies on the structure of the program. A second result gives a black-box reduction that works with only query access to the program but relies on the program having only *constant-width*. A third result shows that any generic reduction of the form *“HSGs imply PRGs with similar parameters”* would imply new unconditional PRGs and would put BPL in $\mathsf{L}^{1+\varepsilon}$ for any constant $\varepsilon > 0$.

- “Non-Disjoint Promise Problems from Meta-Computational View of Pseudorandom Generator Constructions” by Shuichi Hirahara.

  Promise problems are relaxations of decision problems. To solve a promise problem, a device needs to output YES on a set A of inputs and NO on another set B of inputs. The device may answer arbitrarily on inputs outside $A \cup B$. Naturally, A and B are assumed to be disjoint, for if some string is both in A and B, then no device can satisfy the above requirements.

  The paper introduces a new concept called “non-disjoint promise problems.” These are promise problems that under standard complexity assumptions would be non-disjoint and thus ill-defined and cannot be solved by any device. For example, a well-known complexity assumption, which suffices for constructing optimal pseudo-random generators, is that E, the class of problems solved in simply exponential time, requires exponential size circuits.

  Based on this assumption, no device can tell apart A, truth-tables of functions in E, from B, truth-tables of functions requiring circuit-size $2^{\Omega(n)}$ (since these sets of truth-tables are not disjoint).

  Alas, settling whether A and B are disjoint is a long-standing open problem in computational complexity theory, perhaps beyond the reach of current techniques. The present paper shows that optimal hitting-set generators (the one-sided error analog of pseudo-random generators) could be attained even if A and B are disjoint, as long as it is computationally hard to tell A and B apart. This brings us closer to establishing derandomization since we can base it on weaker assumptions. In another setting, the computational hardness of supposedly non-disjoint promise problems would yield an equivalence between the worst-case hardness and average-case hardness of PH.

- "Multiparty Karchmer–Wigderson Games and Threshold Circuits" by Alexander Kozachinskiy and Vladimir Podolskii.

  Karchmer–Wigderson games are communication games whose communication complexity exactly captures the formula depth of the corresponding function. For a given Boolean function $f : \{0,1\}^n \to \{0,1\}$, Alice is given an input $x \in f^{-1}(0)$, Bob is given an input $y \in f^{-1}(1)$ and their goal is to find a coordinate $i \in [n]$ on which x_i and y_i differ. Protocols for this game translate to De Morgan formulae computing f with the same depth (and vice versa). This article generalizes the Karchmer–Wigderson framework to the number-in-hand multiparty communication setting, where k parties are given inputs, and they seek to find a coordinate on which all their inputs agree. For a certain class of functions (including the majority function), the communication cost of a protocol is shown to be equal up to a constant factor to the depth of a formula over the gate set consisting only of thresholds of arity $k+1$ (instead of AND, OR gates in De Morgan formulae). A beautiful application of this new perspective is an explicit construction of a logarithmic-depth formula, composed only of 3-input majority gates, computing the majority function. An explicit logarithmic-depth formula computing the majority function with AND, OR gates follows from Ajtai, Komlós, and Szemerédi's seminal result. Valiant used the probabilistic method to give a simpler proof for the *existence* of such a formula. His proof was adapted by Goldreich to a probabilistic construction using only 3-input majority gates, but the question of explicitness remained open. The present paper settles in the affirmative a conjecture by Cohen et al., who were motivated by applications in multi-party computation.

- "Connecting Perebor Conjectures: Towards a Search to Decision Reduction for Minimizing Formulas" by Rahul Ilango.

  Whether the *Minimum Circuit Size Problem*, MCSP, is NP-hard has been a central question in meta-complexity with far-reaching consequences since the seminal work of Cook and Levin. The present paper focuses on a "cousin" of MCSP, called *Minimum Formula Size Problem*, or MFSP. If MFSP is NP-hard, then it has a search-to-decision reduction. Thus, tackling the latter is a natural intermediate step before tackling the former. The present paper gives a $2^{O(N/\log\log N)}$ search-to-decision reduction that works for most inputs of length N, and a $2^{0.67N}$ randomized reduction that works for all such inputs. The reductions rule out a "bizarre world" where finding the smallest formula of a given truth table requires brute force, but determining its size requires only polynomial time.

We would like to thank the authors for their contributions, the CCC Program Committee for their initial reviews, Shubhangi Saraf for her 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 2020 Program Committee

Andrej Bogdanov (Chinese University of Hong Kong)

Per Austrin (KTH Royal Institute of Technology)

Zeev Dvir (Princeton University)

Prahladh Harsha (Tata Institute of Fundamental Research)

Toniann Pitassi (University of Toronto and IAS)

Noga Ron-Zewi (Haifa University)

Shubhangi Saraf (Rutgers University) (Chair)

Avishay Tal (University of California, Berkeley)

Salil Vadhan (Harvard University)

Ryan Williams (Massachusetts Institute of Technology)

Ronald de Wolf (CWI and University of Amsterdam)

Amir Yehudayoff (Technion—Israel Institute of Technology)

GUEST EDITORS

Zeev Dvir
Princeton University
zdvir@cs.princeton.edu
https://www.cs.princeton.edu/~zdvir/

Avishay Tal
University of California, Berkeley
atal@berkeley.edu
https://www.avishaytal.org/