Computational Complexity

Paul Beame (University of Washington),
Russell Impagliazzo (UCSD),
Valentine Kabanets (SFU),
Toni Pitassi (University of Toronto),
Avi Wigderson (IAS)

September 4 – 9, 2016

1 Overview of the Field

Computational Complexity Theory studies the inherent costs of algorithms for solving mathematical problems. Its major goal is to identify the limits of what is efficiently computable in natural computational models. Computational complexity emerged from the combination of logic, combinatorics, information theory, and operations research. It coalesced around the central problem of ”P versus NP” (one of the seven open problems of the Clay Institute). While this problem remains open, the field has grown both in scope and sophistication. Currently, some of the most active research areas in computational complexity are

- circuit complexity (both Boolean and arithmetic circuit models),
- pseudorandomness,
- proof complexity and the connections with circuit complexity and search heuristics.

Complexity theory has often been using (and contributing to) a number of different areas of mathematics: logic, combinatorics, information theory, algebra, geometry, and analysis, to name just a few.

2 Recent Developments and Open Problems

Below we sample a few exciting recent developments in complexity.

Arithmetic circuit complexity. Building on the work by Valiant et al. [55], Agrawal and Vinay [1] observed that even very shallow (depth 4) arithmetic circuits are powerful enough to simulate general arithmetic circuits (with a certain increase in size). Later this result has been improved by Koiran [38] and Tavenas [52], showing that a lower bound $n^{\omega(\sqrt{n})}$ for a very special kind of arithmetic depth 4 circuits (homogeneous depth 4 circuits).would imply a superpolynomial lower bound for general arithmetic circuits.

These results clarify why it's so difficult to prove lower bounds even for small constant-depth arithmetic circuits. On the other hand, they also provide a potentially successful approach to proving general lower bounds by focusing on shallow arithmetic circuits of certain particular form. There have been a number of recent exciting results in that direction, culminating with the work of Kayal et al. [37], and Kumar and Saraf [43], that gives a lower bound $n^{\Omega(\sqrt{n})}$ for depth 4 homogeneous arithmetic circuits, bringing us tanatalizingly close to the holy grail of arithmetic circuit complexity — proving that Permanent requires superpolynomial arithmetic circuit complexity.

Polynomial Identity Testing (PIT). Determining if a given arithmetic circuit computes an identically zero polynomial (Polynomial Identity Testing problem) is a central problem in derandomization and in (arithmetic) circuit complexity. While devising an efficient deterministic algorithm for PIT remains a major open question in complexity theory, there has been some for restricted classes of arithmetic circuits, often with the help of geometric ideas (e.g., Sylvester-Gallai analogs).

Very recently, Kopparty et al. [41] proved that PIT is polynomial-time equivalent to the problem of deterministic multivariate polynomial factorization, another outstanding open problem in derandomization.

Boolean circuit complexity and SAT algorithms. A breakthrough circuit lower bound was proved by Williams [59], showing that nondeterministic exponential-time computable problems require superpolynomial constant-depth circuits with arbitrary counting (mod m) gates; previously, only the case of prime moduli m was known. Recently, this bound was strengthened by Williams [58] to a more powerful class of circuits.

These and related lower bound results exploit a deep connection between circuit lower bounds and SAT algorithms (or pseudorandom generators) for the same class of circuits. In particular, the classical lower bounds for constant-depth circuits (AC^0) and n^3-size de Morgan formulas have been recently shown to yield improved (better than naive brute force) SAT algorithms for the same class of circuits [30, 50] as well as new pseudorandom generators [31].

There has also been more progress on strengthening classical lower bounds for formulas. For example, Komargodski et al. [39] get a polynomial-time computable function that can't be computed well on average by any de Morgan formula of size below n^3.

Very recently, Gavinsky et al. [23] made a step towards resolving the KRW Composition Conjecture (about the depth complexity of the composition of two boolean functions), which would imply a new circuit lower bound (against the class NC^1 of $O(\log n)$-depth boolean circuits). They proved a special case of the KRW conjecture [35] for the composition of a boolean function and a universal relation. The proof relies on the information-theoretic techniques that have been very useful in a number of recent results in communication complexity.

Proof complexity, circuit complexity, and PIT. Recently, Grochow and Pitassi [25] introduced a new proof system, the Ideal Proof System (IPS), that has tight connections to (algebraic) circuit complexity. Namely, super-polynomial lower bounds on any Boolean tautology in IPS implies that the permanent does not have polynomial-size algebraic circuits (VNP is not equal to VP). As a corollary to the proof, super-polynomial lower bounds on the number of lines in Polynomial Calculus proofs (a previously studied proof system) implies the Permanent versus Determinant Conjecture. Prior to this work, there was no proof system for which lower bounds on an arbitrary tautology implied any computational lower bound. This work begins to shed light on why proof complexity lower bounds have been so difficult to obtain, and highlights the polynomial identity testing (PIT) problem as central to this issue. More specifically, IPS is polynomially equivalent to standard Frege systems if a natural set of propositional axioms, satisfied by any Boolean circuit computing PIT, has subexponential-size Frege proofs. This work raises many exciting questions about connections between PIT, algebraic circuit complexity and proof complexity.

3 Presentation Highlights

A number of exciting breakthrough results have been discovered in complexity theory in the last couple of years. Some of these were presented at the workshop.

- Perhaps one of the most interesting results is a new construction of two-source extractors (related to bipartite Ramsey graphs) due to Chattopadhyay and Zuckerman [6]. This result and some of the follow-up work has been presented in the Extractors session of the workshop (see Section 3.1 below).

- Wigderson and Garg gave two talks discussing some exciting recent progress on the complexity of algebraic computation in the non-commutative case (see Section 3.2).
- Recent progress on locally decodable and locally testable error-correcting codes was reported by Saraf and Kopparty (see Section 3.3).
- Finally, Regev gave a talk on the very recent (yet unpublished) result showing the reverse of the classical Minkowski theorem for lattices (see Section 3.4).

3.1 Randomness extractors

Randomness extractors theory is a field with origins attributed to von Neumann [47] who considered the problem of efficiently generating random bits. Ideally, a randomness extractor would have been a function $\mathsf{Ext}\colon \{0,1\}^n \to \{0,1\}^m$ with the property that for any random variable X with min-entropy k, $\mathsf{Ext}(X) \approx_\varepsilon U_m$, where recall that X has min-entropy k if $\forall x\ \mathbf{Pr}[X = x] \le 2^{-k}$. Unfortunately, such a function does not exist for entropy $k < n-1$ and $\varepsilon < 1/2$. Thus, several restricted types of randomness extractors were introduced, including multi-source extractors [7, 2], seeded-extractors [48], and non-malleable extractors [16].

Although their original motivation was for the specific task of "purifying randomness", extractors have found dozens of applications for error correcting codes, circuit lower bounds, data structures, communication complexity, and so forth. Randomness extractors strengthen hash functions, expander graphs, and list-decodable codes, and are very much related to pseudo-random generators [53]. This myriad of applications and deep connections makes randomness extractors a central, vibrant and exciting field.

The most well understood type of extractors are seeded extractors. A function $\mathsf{Ext}\colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ is a (k,ε)-*seeded extractor* if for any random variable X with min-entropy k, $(\mathsf{Ext}(X,S),S) \approx_\varepsilon U_m \times S$ for S that is uniform independent of X. The goal in to design efficiently computable extractors that support min-entropy k that is as low as possible with respect to the sample length n and the desired error guarantee ε. With respect to k, the extractor should output as many bits as possible. For the case of seeded extractors, it is easy to prove the existence of a seeded extractor that supports min-entropy $k = \Omega(\log(1/\varepsilon))$, has seed length $d = O(\log(n/\varepsilon))$, and outputs $m = k - O(\log(1/\varepsilon))$ bits. In a long line of work, that have accumulated to [28, 17], seeded extractors with comparable parameters were obtained.

Motivated by the problem of privacy amplification, Dodis and Wichs [16] introduced the notion of a non-malleable extractor. A function $\mathsf{Ext}\colon \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ is a (k,ε) *non-malleable extractor* if for any random variable X with min-entropy k, and for any function $\mathcal{A}\colon \{0,1\}^d \to \{0,1\}^d$ with no fixed points, $(\mathsf{Ext}(X,S), \mathsf{Ext}(X,\mathcal{A}(S)), S) \approx_\varepsilon U_m \times (\mathsf{Ext}(X,\mathcal{A}(S)), S)$.

Dodis and Wichs proved that non-malleable extractors with parameters comparable to those of standard seeded extractors exist, though they left the problem of constructing non-malleable extractors for future research. This proved to be a challenging task. The many successful techniques and ideas that were used for the construction of seeded extractors fail to satisfy non-malleability. In a long, intense, line of work, new techniques were developed in the form of new pseudo-random primitives such as correlation breakers [8, 5] and independence-preserving mergers [10] which are at the core of existing constructions of non-malleable extractors. The state of the art constructions of non-malleable extractors [11, 44] have seed length $d = O(\log n) + \tilde{O}(\log(1/\varepsilon))$, support any min-entropy $k = \Omega(d)$, and output $m = 0.49k$ output bits.

Although existing constructions of non-malleable extractors are optimal up to $\log\log(1/\varepsilon)$ factors, an important open problem would be to remove these excessive factors. This would imply improved constructions of two-source extractors [6, 4] and in particular, significantly improved constructions of Ramsey graphs [49, 19]. By designing better correlation breakers, one can obtain such improvements. Thus, future progress lies on the fundamental problem of breaking undesired correlations between random variables.

3.2 Deterministic polynomial time algorithms for non-commutative rank of symbolic matrices

Avi Wigderson and Ankit Garg gave two talks on recent exciting results on deterministic efficient algorithms for computing non-commutative rank of *symbolic matrices* [22, 33, 32]. Symbolic matrices are matrices whose entries are linear functions in variables $\{x_1, x_2, \ldots, x_m\}$ over a field $\mathbb{F}$. Any such matrix can be expressed as a linear combination of the variables with matrix coefficients

$$L = x_1A_1 + x_2A_2 + \cdots + x_mA_m$$

where $A_1, A_2 \ldots, A_m$ are $n \times n$ matrices over $\mathbb{F}$.

The main computational problem of interest (known as SINGULAR) regarding symbolic matrices is determining whether a given symbolic matrix is invertible or not (over the field of fractions in the given variables). This problem has a dual life, depending on whether the variables commute or don't commute. In the *commutative* case this problem has an illustrious history and significance. It was first explicitly stated by Edmonds [18], and shown to have a randomized polynomial time algorithm by Lovasz [45]. The completeness of determinant for arithmetic formulas by Valiant [54] means that singularity captures the celebrated Polynomial Identity Testing (PIT) problem. The derandomization of the latter probabilistic algorithm for PIT is a major open problem in complexity theory as it would imply nontrivial arithmetic or Boolean lower bounds, due to a result of Kabanets and Impagliazzo [34].

In the *non-commutative* case even the meaning of this problem SINGULAR is unclear. It took decades to fully define and understand the related notion of a "field of fractions" for non-commutative polynomials, namely the *free skew field* over which we (attempt to) invert the matrix. However, this non-commutative SINGULAR problem has many intuitive and clean equivalent formulations (some entirely commutative!). It captures a non-commutative version of identity testing for polynomials and rational functions, provides a possible avenue to attack the notorious commutative PIT version, and quite surprisingly, its different formulations arise naturally in diverse areas of mathematics, revealing surprising connections between them.

The algorithm in [22] is a *quantum* generalization of Sinkhorn's *matrix scaling* algorithm [51] (and works only over subfields of $\mathbb{C}$). It utilizes a non-trivial connection between non-commutative symbolic matrices and completely positive operators. A completely positive operator is an operator $\widetilde{L} : M_n(\mathbb{C}) \to M_n(\mathbb{C})$ of the form $\widetilde{L}(P) = \sum_{i=1}^m A_iPA_i^\dagger$. It turns out that the symbolic matrix L has full non-commutative rank iff the corresponding operator $\widetilde{L}$ is rank non-decreasing $\left(\text{Rank}(\widetilde{L}(P)) \geq \text{Rank}(P) \text{ for all } P \succeq 0\right)$. This is a non-trivial theorem which follows from the work of Cohn [9]. This rank non-decreasing property can then be checked by testing the convergence of a natural iterative procedure (first suggested by Gurvits [27]) which tries to bring the operator $\widetilde{L}$ into a *doubly-stochastic* form (operator and its dual both map identity to identity). The analysis of this algorithm again relies on the different formulations of this SINGULAR problem. Particularly, some invariant theoretic bounds due to Derksen [13] play an important role. This operator scaling algorithm also solves a non-convex optimization problem which has found applications in approximating optimal constants in Brascamp-Lieb inequalities [21].

The algorithm in [33, 32] is completely different and is an algebraic analogue of the augmenting paths algorithm for perfect matching in bipartite graphs (and works over finite fields as well). The role of augmenting paths is played by *Wong sequences*. This algorithm and its analysis also draw upon the rich mathematical structure around the non-commutative SINGULAR problem. Another way of phrasing Cohn's result is that the symbolic matrix $L = \sum_{i=1}^m x_iA_i$ is singular (in the non-commutative case) iff $A_1, A_2, \ldots, A_m$ admit a shrunk subspace i.e. there exists $V \subseteq \mathbb{F}^n$ s.t. there exists a $W \subseteq \mathbb{F}^n$ s.t. $\dim(W) < \dim(V)$ and $A_i(V) \subseteq W$ for all i. This can be thought of as an algebraic analogue of Hall's marriage theorem (in fact Hall's theorem is a special case when the A_i's are elementary matrices corresponding to the edges of a bipartite graph). The authors in [33, 32] use Wong sequences to keep producing evaluations of L (by specializing x_i's to matrices of dimension roughly n) of higher and higher rank until the process stops and produces a shrunk subspace witnessing the maximum rank. Along the way, they develop several fundamental mathematical results, for example, regularity of rank of blow-ups of matrix spaces, which have led to remarkable progress in degree bounds for semi-invariants of quiver representations (see [14]).

Now we list some of the interesting future directions and open problems. One is designing *black box* algorithms for the non-commutative rank problem i.e. an algorithm that doesn't even look at the matrices $A_1, A_2, \ldots, A_m$. The only allowed access to the symbolic matrix L is via rank of evaluations at matrix points i.e. queries of the form

$$\text{Rank}\left(\sum_{i=1}^m B_i \otimes A_i\right)$$

for matrices $B_1, B_2, \ldots, B_m$ of polynomial dimension. Both the algorithms in [22, 33, 32] are *white box*. Another interesting direction is whether the techniques and results for the non-commutative rank problem can lead to progress for the commutative rank problem. Non-commutative rank is always within a factor 2 of the commutative rank ([20]), so the algorithms for non-commutative rank yield a factor 2 approximation for commutative rank (in deterministic polynomial time). In a very recent development, a greedy algorithm was designed to give a deterministic PTAS for commutative rank [3] and *Wong sequences* played an important role in the analysis of the algorithm! It is likely that other techniques and results from the non-commutative rank problem will come into play as well. Finally, can one find more applications in optimization of non-commutative rank problem. It already captures as special cases, bipartite matching, matroid intersection (even in a *hidden* sense) [27] and Brascamp-Lieb inequalities/polytopes [21]. One interesting challenge is if matching in general graphs can be reduced to non-commutative rank.

3.3 Error-correcting codes

An error-correcting code is given by an *encoding map* $E : \{0,1\}^k \to \{0,1\}^n$, which "encodes" strings of length k into strings of length n (everything can also be done with $\{0,1\}$ replaced by any finite set Σ; we stick to $\{0,1\}$ for this discussion). The image of this map is called the code, which we will denote by $\mathcal{C}$, and its elements are called codewords. The main measures of the quality of an error-correcting code are its rate R and its minimum distance δ. The rate R is defined to be k/n, which measures the redundancy/wastage introduced in the encoding. The minimum distance $\delta \in (0,1)$ is defined to the be the smallest (fractional) Hamming distance ($\Delta(\cdot,\cdot)$) between two distinct elements of $\mathcal{C}$.

Since the late 1980s, error-correcting codes and the paradigms for constructing them found great impact in theoretical computer science. In particular, error-correcting codes based on polynomials played a central role in the development of Interactive Proofs, Probabilistically Checkable Proofs (PCPs), cryptographic hard-core bits, hardness amplifiers and pseudorandom generators. The centerpiece of all these developments was the fact that a multivariate low-degree polynomial over a finite field could be locally interpolated at a point x by looking at the values taken by that polynomial on all other points of any line passing through x. This endowed the evaluation table of a low-degree multivariate polynomial with some local robustness; errors can be corrected by only looking at a few other entries of the table.

Motivated by this, one can define a *locally decodable code* as an error-correcting code equipped with a (randomized) decoding algorithm, which when given as input a received word $r \in \{0,1\}^n$ which is with distance $0 < \delta_0 < \frac{\delta}{2}$ of a codeword, and a message coordinate $i \in [k]$, the algorithm looks only at $o(k)$ entries of r and returns the "correct" message bit m_i with high probability (i.e., if m is the unique codeword such that $\Delta(E(m), r) < \delta_0$, then the algorithm returns m_i with high probability). Similarly one can define *locally testable codes*, which come with a testing algorithm that with high probability distinguishes, using few queries, between a given received word being within distance ϵ_1 of some codeword, and being further than distance ϵ_2 of every codeword. This ability to work with error-correcting codes in *sublinear-time* formed the conceptual heart of the various developments in theoretical computer science.

3.3.1 High rate error correcting codes

This writeup is on codes in the constant/high rate regime. This is the setting with huge potential applications to data storage, and is also the setting where the most activity has been in recent years. For a long time, the only known family of local codes in this regime were constant-variable Reed-Muller codes over large fields, which are codes based on low degree polynomials. These codes could achieve any constant rate $R < 1/2$, constant distance δ, and local testability/correctability with $O(n^{\epsilon_R})$ queries, where n is the length of the code. In recent years, there has been a flurry of activity in this area [42, 26, 29, 40, 56, 57] resulting in many new constructions of local codes of constant rate. These codes improved the rate-distance tradeoffs achievable by locally testable codes and locally decodable codes, while also reducing the query complexity significantly.

Today the best bounds on query complexity that we know for high rate LDCs and LTCs are from the work of [40] and the bounds obtained are as follows:

- LTCs which can achieve any constant rate $R < 1$, any constant distance $\delta < 1 - R$, and with query complexity
$$O((\log n)^{O(\log \log n)}),$$

- LDCS (and LCCs) which can achieve any constant rate $R < 1$, any constant distance $\delta < 1 - R$, and with query complexity
$$O(2^{\sqrt{\log n \log\log n}}).$$

On the other hand, it is known that LDCs with constant rate require query complexity at least $\Omega(\log n)$. For LTCs, there is no lower bound on query complexity known, and for all we know LTCs with constant rate, constant distance and constant query complexity could exist!

Even more recently, in [24] it was shown how to construct locally testable codes and locally decodable codes that have rate versus distance tradeoff achieving the *Gilbert-Varshamov bound* over small alphabets. In particular, this means that the best rate-distance tradeoff known for classical (non-local) codes can also be achieved by codes that support local algorithms. For local testing the query complexity is quasipolylogarithmic, but for local decoding the query complexity is only polynomially small.

3.3.2 Open Questions

The main question that remains is to understand the optimal query complexity that can be supported by high rate LDCs and LTCs. In particular

Question 3.1. *Do locally decodable codes of constant rate and polylogarithmic decoding complexity exist?*

Note that the running time of any local decoder cannot be smaller than logarithmic (since it needs to be able to index coordinates of the received word it has access to). It turns out that this is also a lower bound on the query complexity [36] (without considering running time).

Question 3.2. *Do locally testable codes of constant rate and polylogarithmic query complexity, or even constant query complexity exist?*

For codes achieving the GV bound, one very interesting question is the following.

Question 3.3. *Do there exist locally decodable codes over the binary alphabet with rate-distance tradeoff approaching the Gilbert-Varshamov bound and subpolynomial query complexity?*

3.4 Reverse Minkowski's Theorem

Oded Regev from Courant Institute of Mathematical Sciences, New York University, gave a talk on a version of Minkowski's theorem for lattices, based on joint works with Daniel Dadush (CWI, Amsterdam) and Noah Stephens-Davidowitz (Courant, NYU).

A lattice is the set of integer linear combinations of linearly independent basis vectors $\mathbf{B} = (\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n)$. The determinant of the lattice $\det(\mathcal{L}) = |\det(\mathbf{B})|$ is a measure of its global density in the sense that

$$\det(\mathcal{L}) = \lim_{r \to \infty} \frac{\mathrm{vol}(rB_2^n)}{|\mathcal{L} \cap rB_2^n|} ,$$

where rB_2^n denotes the Euclidean ball of radius r. A celebrated theorem due to Minkowski from 1891 shows that a lattice with small determinant must also have many lattice points in a relatively small ball [46]. Informally, it says that if a lattice is globally dense, then it must also be locally dense.

Theorem (Minkowski's Theorem (for ℓ_2 balls))**.** *For any lattice $\mathcal{L} \subset \mathbb{R}^n$ with $\det(\mathcal{L}) \leq 1$ and $r > 0$,*

$$|\mathcal{L} \cap rB_2^n| \geq 2^{-n} \cdot \mathrm{vol}(rB_2^n) \geq (r/\sqrt{n})^n .$$

This theorem is one of the foundational results in the study of lattices and the geometry of numbers, and it has innumerable applications in algebraic number theory, sphere packing, computational complexity, cryptography, and more. It is therefore quite natural to ask whether a converse of Minkowski's Theorem holds. In particular, if a lattice has sufficiently many points in a sufficiently small ball, does it necessarily have small determinant? I.e., does local density imply global density?

It is easy to see that the answer is actually no. Consider, for example, the lattice generated by the vectors $(\boldsymbol{e}_1/s, s^2\boldsymbol{e}_2)$ for some arbitrarily large s. This lattice has as many points as we like in any ball around the origin (about $2sr + 1$ points in a ball of radius r), but it has arbitrarily large determinant s. Notice, however, that this lattice contains a *sublattice* generated by $\boldsymbol{e}_1/s$ that does have small determinant (1/s). This leads us to a more refined question: if a lattice has sufficiently many points in a sufficiently small ball, does it necessarily have a small-determinant *sublattice*? I.e., does local density imply global density *in a subspace*?

Prior to our work, the best known results were quantitatively quite weak. But, to our knowledge, the worst example is simply the integer lattice $\mathbb{Z}^n$, which has roughly n^{r^2} points in a ball of radius $r > 0$ for $r \ll \sqrt{n}$. Dadush conjectured that any lattice whose sublattices all have determinant at least one must have fewer than $2^{O(\operatorname{poly}\log(n) r^2)}$ points in a ball of radius $r > 0$ [12, 15]. Our main result is a proof of this conjecture.

Theorem 3.4 (Reverse Minkowski Theorem)**.** *For any lattice $\mathcal{L} \subset \mathbb{R}^n$ with $\det(\mathcal{L}') \geq 1$ for all sublattices $\mathcal{L}' \subseteq \mathcal{L}$,*

$$\sum_{\boldsymbol{y} \in \mathcal{L}} e^{-\pi t^2 \|\boldsymbol{y}\|^2} \leq \frac{3}{2} \,,$$

where $t := 10 \log n + 1$.

In particular, the theorem implies that

$$|\mathcal{L} \cap r B_2^n| \leq 1 + \frac{1}{2} \cdot e^{\pi t^2 r^2} \,,$$

for any $r > 0$, as desired. The proof follows a framework suggested by Shapira and Weiss and relies on techniques from convex geometry, including the ℓ position of convex bodies and Bobkov's observation of a "maximal Gaussian mass position".

As described in [15], the theorem has many applications, including to computational complexity, cryptography, Brownian motion on flat tori (answering a question of Saloff-Coste), Integer Programming (through a conjecture of Kannan and Lovasz), additive combinatorics, Mikowski's conjecture (through work of Shapira and Weiss), and more.

Some of the remaining open questions include:

- Can we show a tight bound, namely, that $\mathbb{Z}^n$ is the densest lattice?
- Can we extend the result to other convex bodies, thereby resolving the general Kannan-Lovasz conjecture?
- Can the techniques be used to solve Minkowski's conjecture?

4 Outcome of the meeting

The meeting provided a great overview of the field, with a number of excellent talks describing some recent exciting results in complexity. There were also many group discussions among leading researchers as well as young students and postdocs. It is quite certain that some new interesting results will emerge as a result of the workshop.

References

[1] Manindra Agrawal and V. Vinay. Arithmetic circuits: A chasm at depth four. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA*, pages 67–75. IEEE Computer Society, 2008.

[2] B. Barak, R. Impagliazzo, and A. Wigderson. Extracting randomness using few independent sources. *SIAM Journal on Computing*, 36(4):1095–1118, 2006.

[3] Markus Blaser, Gorav Jindal, and Anurag Pandey. Greedy strikes again: A deterministic ptas for commutative rank of matrix spaces. *http://eccc.hpi-web.de/report/2016/145/*, 2016.

[4] A. Ben-Aroya, D. Doron, and A. Ta-Shma. Explicit two-source extractors for near-logarithmic min-entropy. In *Electronic Colloquium on Computational Complexity (ECCC)*, 2016.

[5] E. Chattopadhyay, V. Goyal, and X. Li. Non-malleable extractors and codes, with their many tampered extensions. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 285–298. ACM, 2016.

[6] E. Chattopadhyay and D. Zuckerman. Explicit two-source extractors and resilient functions. *Electronic Colloquium on Computational Complexity (ECCC)*, 2015.

[7] B. Chor and O. Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM Journal on Computing*, 17(2):230–261, 1988.

[8] G. Cohen. Local correlation breakers and applications to three-source extractors and mergers. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 845–862. IEEE, 2015.

[9] P. M. Cohn. *Skew Fields, Theory of General Division Rings*. Cambridge University Press, 1995.

[10] G. Cohen and L. Schulman. Extractors for near logarithmic min-entropy. In *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2016.

[11] Gil Cohen. Two-source extractors for quasi-logarithmic min-entropy and improved privacy amplification protocols. In *Electronic Colloquium on Computational Complexity (ECCC)*, number 114, 2016.

[12] Daniel Dadush. Private communication, 2012.

[13] Harm Derksen. Polynomial bounds for rings of invariants. *Proceedings of the American Mathematical Society*, 129(4):955–964, 2001.

[14] Harm Derksen and Visu Makam. Polynomial degree bounds for matrix semi-invariants. *arXiv preprint arXiv:1512.03393*, 2015.

[15] Daniel Dadush and Oded Regev. Towards strong reverse Minkowski-type inequalities for lattices. In *FOCS*, 2016. `http://arxiv.org/abs/1606.06913`.

[16] Y. Dodis and D. Wichs. Non-malleable extractors and symmetric key cryptography from weak secrets. In *Proceedings of the forty-first annual ACM Symposium on Theory of Computing*, pages 601–610. ACM, 2009.

[17] Z. Dvir, S. Kopparty, S. Saraf, and M. Sudan. Extensions to the method of multiplicities, with applications to Kakeya sets and mergers. In *50th Annual IEEE Symposium on Foundations of Computer Science*, pages 181–190. IEEE, 2009.

[18] Jack Edmonds. Systems of distinct representatives and linear algebra. *Journal of research of the National Bureau of Standards*, 71(241-245), 1967.

[19] P. Erdős. Some remarks on the theory of graphs. *Bulletin of the American Mathematical Society*, 53(4):292–294, 1947.

[20] Marc Fortin and Christophe Reutenauer. Commutative/noncommutative rank of linear matrices and subspaces of matrices of low rank. 2004.

[21] Ankit Garg, Leonid Gurvits, Rafael Oliveira, and Avi Wigderson. Algorithmic aspects of brascamp-lieb inequalities. *arXiv preprint arXiv:1607.06711*, 2016.

[22] Ankit Garg, Leonid Gurvits, Rafael Oliveira, and Avi Wigderson. A deterministic polynomial time algorithm for non-commutative rational identity testing with applications. *FOCS*, 2016.

[23] Dmitry Gavinsky, Or Meir, Omri Weinstein, and Avi Wigderson. Toward better formula lower bounds: an information complexity approach to the KRW composition conjecture. In STOC'2014, pages 213–222.

[24] Sivakanth Gopi, Swastik Kopparty, Rafael Mendes de Oliveira, Noga Ron-Zewi, and Shubhangi Saraf. Locally testable and locally correctable codes approaching the gilbert-varshamov bound. 2017.

[25] Joshua A. Grochow and Toniann Pitassi. Circuit complexity, proof complexity, and polynomial identity testing. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 110–119.

[26] Alan Guo, Swastik Kopparty, and Madhu Sudan. New affine-invariant codes from lifting. In *proceedings of the Innovations in Theoretical Computer Science Conference (ITCS)*, pages 529–540. ACM Press, 2013.

[27] Leonid Gurvits. Classical complexity and quantum entanglement. *Journal of Computer and System Sciences*, 69(3):448–484, 2004.

[28] V. Guruswami, C. Umans, and S. Vadhan. Unbalanced expanders and randomness extractors from Parvaresh–Vardy codes. *Journal of the ACM*, 56(4):20, 2009.

[29] Brett Hemenway, Rafail Ostrovsky, and Mary Wootters. Local correctability of expander codes. *Inf. Comput*, 243:178–190, 2015.

[30] Russell Impagliazzo, William Matthews, and Ramamohan Paturi. A satisfiability algorithm for ac^0. In Yuval Rabani, editor, *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 961–972. SIAM, 2012.

[31] Russell Impagliazzo, Raghu Meka, and David Zuckerman. Pseudorandomness from shrinkage. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 111–119. IEEE Computer Society, 2012.

[32] Gábor Ivanyos, Youming Qiao, and K. V. Subrahmanyam. Constructive noncommutative rank computation in deterministic polynomial time over fields of arbitrary characteristics. *arXiv preprint arXiv:1512.03531*, December 2015.

[33] Gabor Ivanyos, Youming Qiao, and K. V. Subrahmanyam. Non-commutative edmonds' problem and matrix semi-invariants. *arXiv preprint arXiv:1508.00690*, 2015.

[34] Valentine Kabanets and Russell Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Computational Complexity*, 13:1–46, 2004.

[35] Mauricio Karchmer, Ran Raz, and Avi Wigderson. Super-logarithmic depth lower bounds via the direct sum in communication complexity. *Computational Complexity*, 5(3/4):191–204, 1995.

[36] Jonathan Katz and Luca Trevisan. On the efficiency of local decoding procedures for error-correcting codes. In *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 80–86. ACM Press, 2000.

[37] Neeraj Kayal, Nutan Limaye, Chandan Saha, and Srikanth Srinivasan. An exponential lower bound for homogeneous depth four arithmetic formulas. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 61–70.

[38] Pascal Koiran. Arithmetic circuits: The chasm at depth four gets wider. *Theor. Comput. Sci.*, 448:56–65, 2012.

[39] Ilan Komargodski, Ran Raz, and Avishay Tal. Improved average-case lower bounds for demorgan formula size. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 588–597. IEEE Computer Society, 2013.

[40] Swastik Kopparty, Or Meir, Noga Ron-Zewi, and Shubhangi Saraf. High-rate locally-correctable and locally-testable codes with sub-polynomial query complexity. In Daniel Wichs and Yishay Mansour, editors, *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 202–215. ACM, 2016.

[41] Swastik Kopparty, Shubhangi Saraf, and Amir Shpilka. Equivalence of polynomial identity testing and polynomial factorization. *Computational Complexity*, 24(2):295–331, 2015.

[42] Swastik Kopparty, Shubhangi Saraf, and Sergey Yekhanin. High-rate codes with sublinear-time decoding. *J. ACM*, 61(5):28, 2014.

[43] Mrinal Kumar and Shubhangi Saraf. On the power of homogeneous depth 4 arithmetic circuits. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 364–373.

[44] Xin Li. Improved non-malleable extractors, non-malleable codes and independent source extractors. In *Electronic Colloquium on Computational Complexity (ECCC)*, number 115, 2016.

[45] Laszlo Lovasz. On determinants, matchings, and random algorithms. *Fundamentals of Computation Theory*, pages 565–574, 1979.

[46] Hermann Minkowski. *Geometrie der Zahlen*. B.G. Teubner, 1910.

[47] J. V. Neumann. Various techniques used in connection with random digits. 12(1):36–38, 1951.

[48] N. Nisan and D. Zuckerman. Randomness is linear in space. *Journal of Computer and System Sciences*, 52(1):43–52, 1996.

[49] F. P. Ramsey. On a problem of formal logic. *Proceedings of the London Mathematical Society*, 30(4):338–384, 1928.

[50] Rahul Santhanam. Fighting perebor: New and improved algorithms for formula and QBF satisfiability. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 183–192. IEEE Computer Society, 2010.

[51] R. Sinkhorn. A relationship between arbitrary positive matrices and doubly stochastic matrices. *The Annals of Mathematical Statistics*, 35:876–879, 1964.

[52] Sébastien Tavenas. Improved bounds for reduction to depth 4 and depth 3. In Krishnendu Chatterjee and Jirí Sgall, editors, *Mathematical Foundations of Computer Science 2013 - 38th International Symposium, MFCS 2013, Klosterneuburg, Austria, August 26-30, 2013. Proceedings*, volume 8087 of *Lecture Notes in Computer Science*, pages 813–824. Springer, 2013.

[53] Salil P. Vadhan. Pseudorandomness. *Foundations and Trends in Theoretical Computer Science*, 2011.

[54] Leslie Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8:189–201, 1979.

[55] Leslie G. Valiant, Sven Skyum, S. Berkowitz, and Charles Rackoff. Fast parallel computation of polynomials using few processors. *SIAM J. Comput.*, 12(4):641–644, 1983.

[56] Michael Viderman. A note on high-rate locally testable codes with sublinear query complexity. *Electronic Colloquium on Computational Complexity (ECCC)*, 17:171, 2010.

[57] Michael Viderman. A combination of testability and decodability by tensor products. *Random Struct. Algorithms*, 46(3):572–598, 2015.

[58] Ryan Williams. New algorithms and lower bounds for circuits with linear threshold gates. In STOC'2014, pages 194–202.

[59] Ryan Williams. Nonuniform ACC circuit lower bounds. *J. ACM*, 61(1):2:1–2:32, 2014.