13w5010 Computational Complexity

Paul Beame (University of Washington),
Russell Impagliazzo (University of California, San Diego)
Valentine Kabanets (Simon Fraser University),
Toni Pitassi (University of Toronto),
Avi Wigderson (Institute for Advanced Study)

July 7–12, 2013

1 Overview of the Field

Computational Complexity Theory is the field that 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 ranges from quantum computing to determining the minimum size of circuits that compute basic mathematical functions to the foundations of cryptography and security.

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:

- the study of hardness of approximation of various optimization problems (using probabilistically checkable proofs), and the connections to coding theory,
- the complexity of problems on lattices, and the connections to cryptography (e.g., homomorphic encryption),
- the study of the role of randomness in efficient computation, and explicit constructions of "random-like" combinatorial objects,
- the study of pseudorandomness in computer science and in mathematics,
- the study of the power of various proof systems of logic, and the connections with circuit complexity and search heuristics,
- the study of the power of quantum computation.

The present workshop focussed on many of these areas, in particular, on Communication Complexity, Hardness of Approximation, Probabiistically Checkable Proofs and Error-Correcting Codes, and Pseudorandomness. Below we describe these areas in more detail, and explain the progress made in the meeting.

2 Communication Complexity

Communication complexity (first introduced by Yao) has proven to be one of the most fruitful areas of complexity theory. In this model, two or more parties each know of a different input, and wish to communicate some joint function of all of their inputs. The complexity of the function is the number of bits they must exchange in order to determine the value of the function. This model has proven to be extremely useful because on the one hand, the model is very simple, and so many computational processes (such as circuits, streaming computation and many others) can be viewed as executing communication protocols. On the other hand, over the years we have found powerful techniques that can be used to prove lowerbounds on the communication complexity of different types of functions, and so obtain lowerbounds on other computational processes.

At this workshop, we discussed a few fundamental recent results about communication complexity.

2.1 Presentation Highlights

2.1.1 Information Complexity

Mark Braverman gave an overview of information complexity. Information complexity is closely related to communication complexity, but instead of considering the number of bits exchanged, it studies the amount of information (in Shannon's Information Theory sense) that the parties must exchange to solve the problem. Information complexity has found many applications within communication complexity and related areas. The information complexity of problems has been shown to demonstrate some of the attractive properties of Shannon's entropy. For example, it is additive over independent problems. This additivity allows one to obtain, among other things, tight bounds on the communication complexity of problems. In the talk, Mark discussed the information complexity of the AND function (where Alice and Bob each get a bit and need to output the AND of their inputs). Through known connection, understanding the information complexity of the AND function allows one to get a tight formula of $0.4827\ldots n + o(n)$ for the communication complexity of the Set Disjointness problem: where Alice and Bob are given subsets X and Y of $\{1, \ldots, n\}$ and need to decide whether $X \cap Y = \emptyset$. The talk was largely based on [8].

2.1.2 Direct Products

Anup Rao gave a talk on Direct Products in Communication Complexity [9], about joint work with Mark Braverman, Omri Weinstein and Amir Yehudayoff. The talk was about the following question, in the setting of randomized communication complexity. If C bits of communication are required to compute a function, then how many bits are required to compute the same function on many different inputs? They used methods relying on information theory to answer this question. If $\mathsf{suc}(f, C)$ denotes the maximum success probability that can be achieved for computing f with C bits of communication, and f^n denotes the function that computes n copies of f, they showed (ignoring constants) that $\mathsf{suc}(f^n, nC/\mathsf{polylog}(nC))$ is bounded by $\mathsf{suc}(f, C)^n$. Prior to their work, it was known that $\mathsf{suc}(f^n, C) < \mathsf{suc}(f, C)^n$ [26], and that $\mathsf{suc}(f^n, nC/\mathsf{polylog}(n)) < \mathsf{suc}(f, C)$, proved by [1]. They gave the first exponentially small bounds when the communication is allowed to increase.

2.1.3 Set Disjointness

Alexander Sherstov presented his work [29] on the *set disjointness problem*. In this communication problem, the goal is for k parties to determine whether the given k sets $S_1, S_2, \ldots, S_k \subseteq \{1, 2, \ldots, n\}$ have nonempty intersection, where no party knows all the k sets. This question has an intuitive interpretation in terms of k parties trying to schedule a meeting, where each party's availability is a subset of the time slots $\{1, 2, \ldots, n\}$ and no single party knows everyone's availability. Set disjointness plays the role of SAT in communication complexity theory and has been a focal point of communication complexity research since the late 1980s.

In a striking 1994 paper, Grolmusz [16] proved that the problem can be solved using roughly $O(n/2^k)$ bits of communication. The best lower bound [28] for $k \geq 3$ parties was until now $\Omega(n/4^k)^{1/4}$, valid regardless of which party knows which inputs. In this talk, Sherstov presented an improved lower bound of $\Omega(\sqrt{n}/2^k k)$. This result is tight for quantum communication protocols, and by Grolmusz it is within a square of tight for classical protocols. The proof is a combination of Fourier-theoretic and combinatorial techniques.

The main notion in the proof is that of *directional derivatives* of multiparty communication protocols, used in this work to approximate the acceptance probabilities of protocols by low-degree real polynomials.

2.1.4 Log-Rank Conjecture

Shachar Lovett gave a talk about the log-rank conjecture. The log-rank conjecture is one of the fundamental open problems in communication complexity. Formulated in 1982 by Lovász and Saks [22], it speculates that the two-party communication complexity of any function is governed by the log of rank of the matrix describing this function. More formally, if $f(x, y)$ is a boolean function whose matrix $M_{x,y} = f(x, y)$ has rank r, then the conjecture claims that f can be computed by deterministic protocols using $\log^{O(1)}(r)$ bits. This holds in all known examples. In general, a trivial upper bound on the amount of communication is r. Despite much effort, the best improvements[21, 20] until recently were $O(r)$, and more recently [7] it was improved to $O(r/\log r)$ assuming a number-theoretic conjecture. In his talk, Shachar described a new work [23] which improves the upper bound to $O(\sqrt{r}\log(r))$, and which is based on relating the discrepancy of a function to efficient protocols for it.

2.2 Recent Developments and Open Problems

Open problems:

Information vs communication Although we have made significant progress in understanding the relationship between the information complexity and the communication complexity of computing functions using two parties, it is still a major open problem to understand this. Barak, Braverman, Chen and Rao showed that any protocol with information I and communication C can be simulated by a protocol with communication roughly $\sqrt{I \cdot C}$, but is this tight? Is it possible that such a simulation can be carried out using $O(I)$ bits of communication? This remains a very interesting open question.

The AM complexity of disjointness In the AM model of communication complexity, Alice, Bob, and the prover Merlin get access to a shared source of randomness R. Alice gets an input X and Bob gets an input Y. Merlin tries to convince Alice and Bob that $F(X, Y) = 1$ for some function F by sending them a message $M(X, Y, R)$, he succeeds if both Alice and Bob accept the proof. An AM protocol is one in which if $F(X, Y) = 1$ Merlin has a strategy that succeeds with high probability, and if $F(X, Y) = 0$, Merlin will always fail with high probability. The open question is whether there is a sub-linear proof of disjointness. In other words, Alice and Bob are given sets $X, Y \subset \{1, \ldots, n\}$, and Merlin is trying to convince them that $X \cap Y = \log n$ (note that Merlin can convince them that $X \cap Y \neq \emptyset$ in $\sim \log n$ bits, if that is the case). Does Merlin have a strategy which uses $o(n)$. The answer is conjectured to be negative.

An information complexity approach to 3-party disjointness (breaking the $\sqrt{n}$ barrier?) At the meeting, Sherstov presented a lowerbound of $\sqrt{n}$ for the multiparty communication complexity of disjointness. These methods also prove lowerbounds for quantum protocols, for which $\sqrt{n}$ is tight. It remains open what the true classical communication complexity of multiparty set disjointness is. Sherstov's methods (unlike the methods used to prove a linear lowerbound for two party set disjointness) are not based on information complexity, and one could hope that there is a proof that does use the framework of information complexity to give a lowerbound that beats the $\sqrt{n}$ barrier.

Improved bounds for matrix rigidity The result presented by Lovett seems to have potential to obtain improved bounds for matrix rigidity. The challenge is to provide an explicit matrix which provably cannot be decomposed as the sum of a low rank matrix and a sparse matrix. This problem is well studied and is tightly related to several central problems in complexity, such as arithmetic circuits lower bounds and separation of the communication complexity analog of PH and PSPACE. Specifically, the following conjecture was made: if A is an $n \times n$ matrix with rank r and at most a constant fraction of nonzero elements (say, $n^2/2$ nonzero elements), then A has a zero sub-matrix with at least $n^2 \cdot \exp(-c\sqrt{r})$ many elements, where $c > 0$ is an absolute constant. If true, this is known to be best possible.

3 Hardness of Approximation

Many well-known optimization problems are NP-hard, i.e. assuming P $\neq$ NP, there is no efficient algorithm to solve these problems exactly. A natural approach is to design efficient algorithms that compute approximate solutions along with a guarantee on the quality of approximation. This has been a very successful approach and much research has been devoted to designing good approximation algorithms as well as showing hardness results, i.e. results showing that there is no good approximation algorithm under standard complexity theoretic assumptions such as P $\neq$ NP. In early 90s, a celebrated complexity theoretic result known as the *PCP Theorem* completely revolutionized the field of hardness results. The theorem states that there is a way of writing proofs for NP-statements such that the correctness of the proof can be probabilistically checked by querying only a few bits in the proof! Specifically, the verifier is probabilistic and queries only a constant number of bits from the proof; every true statement has a proof that is accepted with probability 1, and every proof of a false statement is accepted with only a tiny error probability. The PCP Theorem can be equivalently viewed as a hardness result result for the subclass of NP-hard problems called *constraint satisfaction problems*. Since its discovery, there has been a tremendous progress on hardness-of-approximation results for several fundamental NP-hard problems.

3.1 Presentation Highlights

3.1.1 Analytical Approach to Parallel Repetition

Label Cover is a canonical NP-hard problem that is used as a starting point for most hardness reductions. The problem is known to be hard to approximate via a combination of the PCP Theorem and Raz's Parallel Repetition Theorem. David Steurer talked about his co-authored work [13] that proposes an analytical framework to study parallel repetition, in contrast to previous approaches (including Raz's) based on information theory. In this framework, given a game G with value $\mathrm{val}(G)$, the authors consider a *relaxation* of the value denoted as $\mathrm{val}_+(G)$ and prove that for *projection* games, the relaxed value is both *multiplicative* (under parallel repetition) and a good *approximation* for the optimal value. These two properties imply a parallel repetition bound as

$$\mathrm{val}(G^{\otimes k}) \approx \mathrm{val}_+(G^{\otimes k}) = \mathrm{val}_+(G)^k \approx \mathrm{val}(G)^k.$$

The framework leads to a new and significantly simpler proof for the hardness of Label Cover. In addition, the framework allows one to show new parallel repetition bounds in regimes where previously no non-trivial bounds were available, leading to stronger hardness results for Set Cover and other problems.

3.1.2 Parallel Repetition For Low Degree Testing

One of the major open problems in hardness of approximation is whether Label Cover is hard up to a polynomial factor in the size of the input. As a PCP, this corresponds to a 2-query PCP that uses $O(\log n)$ random bits, a polynomial size alphabet, and error probability that is inverse polynomial. This is open even when a constant number of queries is allowed (instead of 2) and is known as the *Sliding Scale Conjecture*.

Dana Moshkovitz [24] talked about her approach to resolving the Sliding Scale Conjecture. *Low degree testing* is an important module in PCP constructions, and Moshkovitz formalized and analyzed parallel repetition of low degree testing. She showed that an appropriate repetition decreases the error probability exponentially. The work is inspired by a previous work of Impagliazzo, Kabanets and Wigderson, who had a similar soundness amplification technique for direct product testing, however without achieving exponentially small soundness error. The work shows further that a *derandomization* of the Parallel Repetition Theorem therein, together with an appropriate *base test*, would imply the Sliding Scale Conjecture. In contrast, standard parallel repetition of two-prover games achieves exponentially small error probability but cannot be derandomized, and a certain variant of parallel repetition has a derandomization but also high error probability.

3.1.3 Approximation Resistance

Constraint satisfaction problems (CSPs) are some of the most well-studied NP-hard problems. Given a predicate $f : \{-1,1\}^k \mapsto \{0,1\}$, an instance of CSP($f$) consists of n ± 1-valued variables and m constraints, where each constraint is the predicate f applied to an ordered subset of k variables, possibly in negated form.

The density of the predicate $\rho(f) = \frac{|f^{-1}(1)|}{2^k}$ is the probability that a uniformly random assignment to its variables satisfies the predicate. An instance of CSP(f) is called α-satisfiable if there is an assignment that satisfies at least an α fraction of the constraints. The predicate f is called *approximation resistant* if given a $(1 - o(1))$-satisfiable instance of CSP(f), it is computationally hard to find an assignment such that the fraction of constraints satisfied is at least $\rho(f) + \Omega(1)$. There has been great progress in showing that some specific predicates are approximation resistant, but a complete characterization of approximation resistance (i.e. a necessary and sufficient condition for a predicate to be approximation resistant) remains elusive.

Subhash Khot talked about his co-authored work [19] making progress on the topic. This work considers the closely related notion of *strong approximation resistance*. The predicate f is called *strongly approximation resistant* if given a $(1 - o(1))$-satisfiable instance of CSP(f), it is hard to find an assignment such that the fraction of constraints satisfied is outside the range $[\rho(f) - \Omega(1), \rho(f) + \Omega(1)]$. The work gives, among other results, a complete characterization of strong approximation resistance in terms of existence of a probability measure on a natural polytope associated with the predicate. (The characterization also depends on the so-called *Unique Games Conjecture*.)

Subhash Khot also talked about a remarkable recent result of Chan [10] that shows approximation resistance of a specific predicate called the Hypergraph Linearity Test predicate. The result was known earlier under the Unique Games Conjecture but Chan obtained an NP-hardness result. The result implies several other NP-hardness-of-approximation results, e.g. for Label Cover, Max-k-CSP, independent sets in bounded degree graphs, graph coloring etc.

3.1.4 LP Relaxations for CSPs

Prasad Raghavendra talked about his co-authored work [11] showing that no polynomial-sized linear programming relaxation can achieve better than a $\frac{1}{2}$-approximation for MAX-CUT, a $\frac{7}{8}$-approximation for MAX-3SAT, or a $\frac{3}{4}$-approximation for MAX-2SAT. This is accomplished by bringing together two formerly disparate lines of research. On one hand, there has been a recent sequence of exciting lower bounds on the size of extended formulations for various polytopes that arise in combinatorial optimization. On the other hand, researchers have extensively studied the power of specific Linear Programming hierarchies for approximating NP-hard problems. The authors show that for CSPs, general polynomial-sized LPs have exactly the same power as the LPs arising from a constant number of *rounds* of a specific family of LPs known as the Sherali–Adams hierarchy.

During the meeting, Prasad Raghavendra and David Steurer pursued some promising approaches towards extending these lower bounds to the more powerful computational model of semidefinite programs.

4 Probabilistically Checkable Proofs

Probabilistically Checkable Proofs (PCPs) are proofs that admit a very efficient verification procedure, one that queries a vanishing fraction of entries in the proof. Locally testable codes are error-correcting codes with the feature that membership in the code can be verfied using a tester that again queries a vanishing fraction of co-ordinates in the codeword. Locally testable codes are thought to be the *combinatorial core* of PCPs, and indeed progress on the two often goes hand-in-hand.

In recent years, there has been much progress on constructing new and more efficient locally testable codes and PCPs. This progress has enabled new applications to approximation algorithms and hardness of approximation. We have even come to the point where PCPs are efficient enough to be practically implemented in real-world cryptographic primitives. Some talks at this workshop were devoted to the most recent progress.

4.1 Recent Developments and Open Problems

There has been considerable interest in obtaining the right tradoeffs between rate, distance and query complexity for locally testable codes. A closely related theme in the PCP literature is to construct shorter proofs that admit query-efficient verifiers (the original constructions were rather inefficient). More recently, there has been work aimed at coming up with constructions of both objects that are efficient enough to be used in practice.

4.2 Presentation Highlights

4.2.1 Locally Testable Codes and Cayley Graphs (Parikshit Gopalan)

This talk presented two graph-theoretic formulations of locally testable codes.

- They are equivalent to Cayley graphs on F_2^n with "nice" spectral properties called derandomized hypercubes. These graphs share many of the nice properties of the Boolean hypercube (small eigenvalue gap, small-set expansion, cuts have influential co-ordinates) but have much fewer vertices [2, 15].
- They are equivalent to certain Cayley graphs on F_2^n whose shortest path metric emebeds into L_1 with constant distortion [15].

The graph-theoretic view has proved useful in the construction of integraility gap examples for various problems in hardness of approximation including Unique Games, Max-Cut and Sparsest Cut. The speaker raised the question of whether it can be used to shed light on rate-distance tradeoffs for locally testbale codes, either through new lower bounds or through better constructions.

In the discussions following the talk, the questions of whether Cayley graphs over other groups give rise to interesting combinatorial objects related to error-correcting codes was raised. Also, it was noted that the definition of LTCs used here is somewhat stronger than the one that is standard in the literature. While it is possible that the two definitions are equivalent (up to some small slack in the parameters), no formal equivalence is known.

4.2.2 Linear-length PCPs for Circuit-SAT with sublinear query complexity (Swastik Kopparty)

This talk presented a recent result [6], which constructed the first PCPs of linear length with nontrivial query complexity. Specifically, it was shown that for every ϵ, Circuit-SAT (and hence for 3-SAT) instances with n gates, there are non-uniform PCPs of length $O(n)$ which can be tested with $O(n^\epsilon)$ queries.

The main method here is to construct an algebraic PCP, along the lines of the original proofs of the PCP theorem, but to use **transitive algebraic-geometric codes** instead of the classically-used polynomial codes. Moving from polynomial codes to transitive algebraic-geometric codes has the advantage that good transitive AG codes can be defined over much smaller fields than polynomial codes can, while still retaining the two key features of polynomials:

1. Algebraicness: which enables the code to support multiplication, which is necessary for supporting the boolean operations of circuit-SAT.
2. Transitivity: which enables the code to support routing, which is necessary for the efficiency of the reductions from circuit-SAT.

The transitive AG codes that get used were constructed by Stichtenoth in the appendix to [6].

This result brings PCP parameters in line with the known locally testable code parameters. It is a major open question whether locally testable codes and PCPs with length $O(n)$ and $n^{o(1)}$ query complexity exist.

A big point of discussion during the workshop was whether the answer to this open question is believed to be yes or no. There seemed to be a general consensus that whatever the answer is, the first progress would most likely be on the locally testable code question, where the combinatorics is clearer (as in Parikshit Gopalan's talk). Another issue that came up is the non-uniformity of the above-constructed PCP. It seems likely that the above PCP can be made uniform, but it would require some deeper understanding of Stichtenoth's codes. This talk also ties in with the talk about practical implementations of PCPs (as in Eli Ben-Sasson's talk), because having short length of PCPs is of high importance for practical applications.

4.2.3 Succinct Computational Integrity and Privacy Research — SCIPR (Eli Ben-Sasson)

This talk presented research performed over the past 3 years in our "scipr-lab" (www.scipr-lab.org), cf. [3, 4, 5]. The work described relies on theoretical constructions of various proof systems (e.g., zero knowledge, probabilistically checkable, transparent, holographic, computationally sound, to name a few) and builds practical systems that compile programs written in high-level languages (such as C) into executables that sign

the integrity of their computation, in a way that is succinctly verifiable and zero-knowldege. (Succinctly verifiable means verifier running time is negligible compared to the time needed to execute the program.)

The implementations use a wide range of tools from computational complexity, algebra and cryptography, like quasilinear PCPs, polylogarithmically verifiable proofs, FFTs over finite fields, succinct noninteractive arguments of knowledge (SNARKs), quadratic span programs, and bilinear group (also known as pairing-based) cryptography.

This talk discussed the project goals, the systems built so far, and the theoretical challenges that lie ahead. Specific theoretical open problems were described, ones that are needed to improve the efficiency and soundness of our systems. One such problem is making the constant-rate PCP construction described in Section 4.2.2 efficient, and a starting point is to get an explicit construction of the transitive AG-codes used there.

5 Pseudorandomness

One fundamental issue in computational complexity is the power of randomness. Can randomness significantly speed up computation for some problems, or does every efficient randomized algorithm have an efficient deterministic simulation ("derandomization")? Many researchers now believe that the latter is the case; indeed, this is known under plausible complexity assumptions, e.g. [18]. However, obtaining *unconditional* derandomization results remains a very active line of research, with many open problems.

The main tool used to attack this question is the pseudorandom generator. A pseudorandom generator (PRG) takes as input a short random seed, and outputs a long string which appears random to a large class of algorithms. Formally, let U_s denote the uniform distribution on $\{0,1\}^s$. We say that a PRG $G : \{0,1\}^\ell \to \{0,1\}^n$ *fools* a class $\mathcal{M}$ of functions to within error ε if, for every $M \in \mathcal{M}$,

$$|\Pr[M(U_n) \text{ accepts}] - Pr[M(G(U_\ell)) \text{ accepts}]| \leq \varepsilon.$$

What are the best unconditional PRGs for natural classes of functions? Efficiently computable PRGs imply circuit lower bounds, which are notoriously difficult. We therefore restrict to natural classes of functions where lower bounds are known. Such classes include functions computable by efficient formulas (perhaps of a special form, like DNFs or CNFs), polynomial threshold functions, branching programs, and combinatorial rectangles of certain sizes. All of these classes were tackled by papers presented at the workshop.

5.1 Recent Developments and Meeting Highlights

5.1.1 Pseudorandomness from Shrinkage

In one talk, David Zuckerman presented "Pseudorandomness from Shrinkage," about joint work with Russell Impagliazzo and Raghu Meka [17]. The hardness vs. randomness paradigm has led to the construction of PRGs from lower bound assumptions ([25] and many more); however, these suffer a polynomial loss in parameters and yield nothing nontrivial for known polynomial lower bounds. Impagliazzo, Meka, and Zuckerman instead provide a way to construct PRGs from lower bounds which have only a negligible loss in parameters, provided the lower bound is proved by the well-known method of shrinkage from random restrictions [30]. (In fact, they need something a bit stronger.) In particular, their methods yield efficient PRGs for several circuit classes which give the smallest seed length possible without improving the corresponding lower bound, up to $n^{o(1)}$ factors. More specifically, they give PRGs fooling the following circuits of size s: de Morgan formulas (seed length $s^{1/3+o(1)}$); formulas over an arbitrary basis (seed length $s^{1/2+o(1)}$); read-once de Morgan formulas (seed length $s^{.234\cdots}$); and branching programs (seed length $s^{1/2+o(1)}$). Previously there were no known nontrivial PRGs for these classes.

5.1.2 Pseudorandom Restrictions

Salil Vadhan gave a talk describing two works that also utilize (pseudo)random restrictions to obtain improved unconditional PRGs. The first was a paper "Better Pseudorandom Generators via Milder Pseudorandom Restrictions" joint with Gopalan, Meka, Reingold, and Trevisan [14], and the second was a paper "Pseudorandomness for Branching Programs via Fourier Analysis," joint with Reingold and Steinke [27].

The use of pseudorandom restrictions in these works is unusual in that not enough bits are set to simplify the function with nonzero probability, but instead it is only argued that the real-valued function obtained by *averaging* over the restricted bits is simpler (which suffices for the pseudorandom generators constructions). Using this approach, improved pseudorandom generators were constructed for read-once CNF formulas [14], combinatorial rectangles [14], constant-width read-once permutation branching programs that can read their bits in any order [27], and large-width read-once branching programs that can read their bits in any order [27], as well as an improved hitting-set generator for width 3 branching programs [14].

Collectively, the above presentations of Zuckerman and Vadhan as well as the recent work of Trevisan and Xue [31], show that (pseudo)random restrictions are a powerful technique for constructing pseudorandom generators, and may enable even more substantial advances on the long-standing open problems in this area.

5.1.3 Deterministic Approximate Counting

The talk by Rocco Servedio was an illustration of how better derandomizations of certain algorithms can be obtained by using algorithmic techniques other than pseudorandom generators. Specifically, he spoke about the paper "Deterministic Approximate Counting for Degree-2 Polynomial Threshold Functions," joint with Anindya De and Ilias Diakonikolas [12]. In this work, they exhibit a deterministic algorithm for approximately computing the fraction of Boolean assignments that satisfy a degree-2 polynomial threshold function. Specifically, given a degree-2 input polynomial $p(x_1, ..., x_n)$ and a parameter $\varepsilon > 0$, the algorithm approximates $\Pr[p(x) \geq 0]$ to within an additive $\pm\varepsilon$ in time $\mathrm{poly}(n, 2^{\mathrm{poly}(1/\varepsilon)})$. Previous algorithms based on pseudorandom generators did not run in fixed polynomial time for constant $\varepsilon > 0$ (but rather the degree of the polynomial depended on ε).

5.1.4 Locally Testable Codes

Parikshit Gopalan's presentation "Locally Testable Codes and Derandomized Hypercubes" described constructions of combinatorial objects that are "partly" pseudorandom, in that some of the properties required hold for randomly generated objects (such as large distance in an error-correcting codes) but others are strongly violated by randomly generated objects (such as "local testability" of an error-correcting code). Specifically, he described results that relate "locally testable error-correcting codes" to Cayley graphs with certain properties. These results come from the papers "Making the long code shorter, with applications to the Unique Games Conjecture" joint with Barak, Håstad, Raghavendra, and Steurer [2] and "Locally testable codes and Cayley graphs" joint with Vadhan and Zhou [15]. The former paper uses locally testable codes to construct Cayley graphs that have many large eigenvalues but are still "small-set expanders," shedding light on the complexity of the Small-Set Expansion Problem and the Unique Games Conjecture. The latter paper presents exact equivalences between locally testable codes and Cayley graphs with certain metric embedding or spectral properties, with the hope that the Cayley graph formulations will help in approaching some of the long-standing open problems about locally testable codes (which in turn are closely related to open problems about probabilistically checkable proofs).

5.1.5 Rounding Group Actions

Amir Yehudayoff's presentation "Rounding group actions" considered generalizations of Cayley graphs, obtained by "rounding" the action of a group on a set. He described results and open problems about when such constructions can and cannot yield expander graphs (which play a fundamental role in the theory of pseudorandomness).

5.2 Outcome of the Meeting

In addition to the above speakers, the workshop attendees included many experts on pseudorandomness and the problems discussed above, including Eli Ben-Sasson, Mark Braverman, Zeev Dvir, Russell Impagliazzo, Valentine Kabanets, Daniel Kane, Swastik Kopparty, Shachar Lovett, Dana Moshkovitz, Ryan O'Donnell, Prasad Raghavendra, Anup Rao, Michael Saks, Shubhangi Saraf, Ronen Shaltiel, David Steurer, Madhu Sudan, Luca Trevisan, and Avi Wigderson. There were many informal discussions among these attendees

about the possibilities raised by the above works and other recent developments, and these discussions are likely to contribute to significant advances in the field during the coming years.

References

[1] Boaz Barak, Mark Braverman, Xi Chen, and Anup Rao. How to compress interactive communication. In *Proceedings of the 2010 ACM International Symposium on Theory of Computing*, pages 67–76, 2010.

[2] Boaz Barak, Parikshit Gopalan, Johan Håstad, Raghu Meka, Prasad Raghavendra, and David Steurer. Making the long code shorter, with applications to the Unique Games Conjecture. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science*, 2012.

[3] E. Ben-Sasson, A. Chiesa, D. Genkin and E. Tromer *Fast Reductions from RAMs to Delegatable Succinct Constraint Satisfaction Problems*, 4th Symposium on Innovations in Theoretical Computer Science (ITCS 2013)

[4] E. Ben-Sasson, A. Chiesa, D. Genkin and E. Tromer *On the Concrete Efficiency of Probabilistically-Checkable Proofs*, 45th ACM Symposium on the Theory of Computing (STOC 2013)

[5] E. Ben-Sasson, A. Chiesa, D. Genkin, E. Tromer and M. Virza *SNARKs for C: Verifying Program Executions Succinctly and in Zero Knowledge*, 33rd International Cryptology Conference (CRYPTO 2013)

[6] E. Ben-Sasson, Y. Kaplan, S. Kopparty, O. Meir. *Constant-rate PCPs for Circuit-SAT with sublinear query complexity*, (with an appendix by Henning Stichtenoth), Proc. IEEE Symposium on Foundations of Computing, 2013, to appear.

[7] E. Ben-Sasson, S. Lovett, and N. Ron-Zewi. An additive combinatorics approach relating rank to communication complexity. In *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*, pages 177–186, 2012.

[8] Mark Braverman, Ankit Garg, Denis Pankratov, and Omri Weinstein. From information to exact communication. In *Proceedings of the 45th annual ACM symposium on Symposium on theory of computing*, pages 151–160. ACM, 2013.

[9] Mark Braverman, Anup Rao, Omri Weinstein, and Amir Yehudayoff. Direct Products in Communication Complexity. In *Electronic Colloquium on Computational Complexity (ECCC)"*, volume 19, 2012.

[10] S.O. Chan, Approximation Resistance from Pairwise Independent Subgroups, *STOC* 2013.

[11] S.O. Chan, J. Lee, P. Raghavendra and D. Steurer, Approximate Constraint Satisfaction Requires Large LP Relaxations, *FOCS* 2013.

[12] A. De, I. Diakonikolas, and R. Servedio. Deterministic Approximate Counting for Degree-2 Polynomial Threshold Functions. Manuscript, 2013.

[13] I. Dinur and D. Steurer, Analytical Approach to Parallel Repetition, *Manuscript*.

[14] P. Gopalan, R. Meka, O. Reingold, L. Trevisan, and S. Vadhan. Better pseudorandom generators via milder pseudorandom restrictions. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS '12)*, pages 120–129. IEEE, 20–23 October 2012.

[15] P. Gopalan, S. Vadhan, Y. Zhou. Locally Testable Codes and Cayley Graphs. Manuscript, 2013.

[16] V. Grolmusz. The BNS lower bound for multi-party protocols in nearly optimal. *Inf. Comput.*, 112(1):51–54, 1994.

[17] R. Impagliazzo, R. Meka, and D. Zuckerman. Pseudorandomness from shrinkage. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science*, pages 111–119, 2012.

[18] R. Impagliazzo and A. Wigderson. P = BPP if E requires exponential circuits: Derandomizing the XOR lemma. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, pages 220–229, 1997.

[19] S. Khot, M. Tulsiani and P. Worah, A Characterization of Strong Approximation Resistance, *ECCC Report* TR13-075.

[20] A. Kotlov. Rank and Chromatic Number of a Graph. *Journal of Graph Theory 26(1)*, pages 1–8, 1997.

[21] A. Kotlov and L. Lovsz. The rank and size of graphs. *J. of Graph Theory*, 23:185–189, 1996.

[22] L. Lovász and M. Saks. Lattices, Möbius Functions and Communication Complexity. *Annual Symposium on Foundations of Computer Science*, pages 81–90, 1988.

[23] S. Lovett. Communication is bounded by root of rank. *Electronic Colloquium on Computational Complexity (ECCC)*, 20:84, 2013.

[24] D. Moshkoviz, Parallel Repetition For Low Degree Testing, *In preparation*.

[25] N. Nisan and A. Wigderson. Hardness vs. randomness. *Journal of Computer and System Sciences*, 49:149–167, 1994.

[26] Itzhak Parnafes, Ran Raz, and Avi Wigderson. Direct product results and the GCD problem, in old and new communication models. In *Proceedings of the 29th Annual ACM Symposium on the Theory of Computing (STOC '97)*, pages 363–372, New York, May 1997. Association for Computing Machinery.

[27] O. Reingold, T. Steinke, and S. Vadhan. Pseudorandomness for regular branching programs via fourier analysis. In *Proceedings of the 17th International Workshop on Randomization and Computation (RANDOM '13)*, Lecture Notes in Computer Science. Springer-Verlag, 21–23 August 2013. To appear. Full version posted as ECCC TR13-086 and arXiv:1306.3004 [cs.CC].

[28] A. A. Sherstov. The multiparty communication complexity of set disjointness. In *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing* (STOC), pages 525–544, 2012.

[29] A. A. Sherstov. Communication lower bounds using directional derivatives. In *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing* (STOC), pages 921–930, 2013.

[30] B. A. Subbotovskaya. Realizations of linear functions by formulas using +, *, -. *Sov. Math. Dokl.*, 2:110–112, 1961.

[31] L. Trevisan and T. Xue. A Derandomized Switching Lemma and an improved Derandomization of AC0. In Proc. 28th IEEE Conference on Computational Complexity (CCC '13), 2013.