

MATHEMATISCHES FORSCHUNGSINSTITUT OBERWOLFACH

Report No. 54/2015

DOI: 10.4171/OWR/2015/54

Complexity Theory

Organised by
Peter Bürgisser, Berlin
Oded Goldreich, Rehovot
Madhu Sudan, Cambridge MA
Salil Vadhan, Cambridge MA

15 November – 21 November 2015

ABSTRACT. Computational Complexity Theory is the mathematical study of the intrinsic power and limitations of computational resources like time, space, or randomness. The current workshop focused on recent developments in various sub-areas including arithmetic complexity, Boolean complexity, communication complexity, cryptography, probabilistic proof systems, pseudorandomness and randomness extraction. Many of the developments are related to diverse mathematical fields such as algebraic geometry, combinatorial number theory, probability theory, representation theory, and the theory of error-correcting codes.

Mathematics Subject Classification (2010): 68-06, 68Q01, 68Q10, 68Q15, 68Q17, 68Q25, 94B05, 94B35.

Introduction by the Organisers

The workshop *Complexity Theory* was organized by Peter Bürgisser (TU Berlin), Oded Goldreich (Weizmann Institute), Madhu Sudan (Harvard), and Salil Vadhan (Harvard). The workshop was held on November 15th–21st 2015, and attended by approximately 50 participants spanning a wide range of interests within the field of Computational Complexity. The plenary program, attended by all participants, featured fifteen long lectures and five short (8-minute) reports by students and postdocs. In addition, intensive interaction took place in smaller groups.

The Oberwolfach Meeting on Complexity Theory is marked by a long tradition and a continuous transformation. Originally starting with a focus on algebraic and Boolean complexity, the meeting has continuously evolved to cover a wide variety

of areas, most of which were not even in existence at the time of the first meeting (in 1972). While inviting many of the most prominent researchers in the field, the organizers try to identify and invite a fair number of promising young researchers.

Computational complexity (a.k.a. complexity theory) is a central field of computer science with a remarkable list of celebrated achievements as well as a vibrant research activity. The field is concerned with the study of the *intrinsic complexity* of computational tasks, and this study tends to *aim at generality*: it focuses on natural computational resources, and considers the effect of limiting these resources on the class of problems that can be solved. Computational complexity is related to and has substantial interaction with other areas of mathematics such as algebra, analysis, combinatorics, geometry, number theory, optimization, probability theory, and quantum computation.

The workshop focused on several sub-areas of complexity theory and its nature may be best illustrated by a brief survey of some of the meeting's highlights.

Randomness Extraction. The problem of extracting almost perfect randomness from sources of highly defected randomness is of great theoretical and practical importance, since perfect randomness is essential to cryptography and has numerous applications in algorithmic design, whereas the natural sources of randomness are quite defected. One important setting of the problem refers to the case in which one is given samples drawn from two independent sources of defected randomness, where the level of defect is captured by a lower bound on the probability that the outcome equals any specific value. The logarithm of the reciprocal of this probability, called *min-entropy*, is a main parameter in these studies.

While it is easy to prove the existence of two-source (randomness) extractors for sources of logarithmic min-entropy, the explicit construction of extractors that can handle min-entropy rate below half was open since 1985. In 2005, Jean Bourgain obtained an explicit construction for min-entropy rate slightly below half (i.e., 0.499), but no progress on this problem has been reported till July 2015, when Eshan Chattopadhyay and David Zuckerman announced a construction that can handle poly-logarithmic min-entropy.

The workshop's actual program started with a special session devoted to this breakthrough, with both authors present. David Zuckerman presented the history and wide context of the problem of constructing two-source extractors, and Gil Cohen presented an overview of the construction. One informal specialized session, which took place on a later day, featured more detailed descriptions of two components of the construction. Specifically, Eshan Chattopadhyay presented constructions of "non-malleable extractors" and Raghu Meka presented constructions of "resilient functions".

One interesting point regarding the construction of Chattopadhyay and Zuckerman is that its analysis makes explicit use of a celebrated result about the computational limitations of bounded-depth Boolean circuits (which was presented by Mark Braverman in the 2009 complexity meeting at Oberwolfach). This is remarkable because these two areas of complexity theory did not seem related

before and their history did not register any actual interaction so far. Another peculiar connection is the use of Uri Feige's leader election protocol for the construction of non-malleable extractors, whereas this protocol was discovered in the 1998 complexity meeting at Oberwolfach, following the presentation of a different protocol by David Zuckerman (which in turn drew on ideas from the area of pseudorandomness).

Boolean Circuit Lower Bounds. The project of establishing circuit lower bounds calls for presenting explicit functions that cannot be computed within limited computational resources. One direction of research is aimed at better understanding of very restricted computation devices such as (unbounded fan-in) bounded-depth circuits and formulae.

Ben Rossman outlined his proof that shows that the simple conversion of circuits of size s and depth d into formulae of size s^d and depth d is essentially the best possible. Specifically, he showed that the parity of n variables, which can be computed by a depth d circuit of size $\exp(n^{1/(d-1)})$, requires depth d formula of size $\exp(\Omega(d \cdot n^{1/d}))$.

Avishay Tal addressed the problem of presenting explicit functions that require depth-three circuits of size $\exp(\omega(\sqrt{n}))$. He presented a proof of such a result in a *restricted model* of depth three circuits, which arises from a natural model of multi-linear circuits for computing multi-linear functions, by studying the "rigidity" of random Boolean Toeplitz matrices. Specifically, he showed that such a random matrix disagrees with any rank r matrix on at least $\tilde{\Omega}(n^3/r^2)$ entries, which improves over the previously known bound of $\Omega(n^2/r)$ when $r < n/\log^2 n$.

Fine-grained complexity. A relatively recent direction of research refers to the study of problems that are known to have polynomial-time algorithms, where the aim is to provide evidence that the known algorithms are actually the best possible. Ryan Williams surveyed research in this direction, known as fine-grained complexity, while highlighting the connection between it and the study of the exact complexity of problems that seem to require exponential-time such as 3SAT.

Doubly-efficient interactive proof systems. The invention of interactive proof systems and the exploration of their power are among the greatest success stories of computational complexity. While research in the 1980s referred to polynomial-time verification aided by a computationally unbounded prover, the term doubly-efficient refers to almost linear-time verification aided by a polynomial-time prover. Clearly, only polynomial-time solvable problems can have such a proof system, even if the soundness condition is relaxed to hold only with respect to polynomial-time cheating provers (who attempt to prove false claims).

This upper bound (on the complexity of problems having doubly-efficient interactive proof systems) is met by a result presented by Ron Rothblum, which uses only one round of communication and relies on standard intractability assumptions. A different system, presented by Rothblum in a specialized session, achieves information theoretic soundness (in a larger constant number of rounds)

for any problem that can be solved in polynomial time and space $n^{o(1)}$. (The space bound is the best possible, up to a constant power.)

Two-server PIR with improved communication complexity. While the computational assumption used by Rothblum refers to one-server computational *Private Information Retrieval* (PIR) schemes, two-server PIRs offer information theoretic security. Specifically, one can retrieve any desired bit in an n -bit long string, held by each of the two servers, by exchanging $O(n^{1/3})$ bits of communication with each server such that no single server gets information about the identity of the desired bit. The simple scheme, invented in 1995, stood unimproved for two decades. Zeev Dvir presented a vast improvement on this simple scheme, by building on results of Yekhanin and Efremenko, which were presented in past Oberwolfach meetings (in 2007 and 2009, resp). The new scheme uses $\exp(\tilde{O}(\sqrt{\log n})) = n^{o(1)}$ bits of communication, and relies on a construction of “matching vectors” family over a finite ring.

High-rate locally-testable and locally-correctable codes. The aforementioned results of Yekhanin and Efremenko refer to the construction of codes that support the recovery of any bit in a corrupted codeword based on a constant number of random probes (i.e., it achieves constant “locality”). These known results refer to codes of sub-exponential length (i.e., the codeword has length that is sub-exponential in the length of the message), and it is also known that such level of locality cannot be supported by codes of almost linear length. In his presentation, Or Meir considered the opposite extreme of the length-vs-locality trade-off: The case in which one requires the code to have linear length (or even length that is optimal with respect to its distance), and tries to minimize the number of probes that suffices for recovering a single bit. The new result asserts $\exp(\tilde{O}(\sqrt{\log n})) = n^{o(1)}$ proves suffice to the n -bit codeword, whereas the prior bound was $n^{1/O(1)}$.

With respect to local testability (i.e., testing whether a string is a valid codeword or far from it by making few queries), the results are better. In the constant-probe regime codes of almost-linear length are known, whereas the new work present linear-length codes that are testable by a quasi-poly-logarithmic number of probes (i.e., the number of probes is $(\log n)^{O(\log \log n)}$).

Computational assumptions in cryptography. Modern cryptography is based on computational assumptions, since its most basic primitive such as secure encryption and unforgeable signatures imply the existence of *one-way functions* (OWF), which in turn is a very strong version of the famous conjecture by which $\mathcal{P} \neq \mathcal{NP}$. In recent years, far stronger computational assumptions became popular in cryptographic research. One such assumption, known as the IO conjecture, postulates that it is feasible to obfuscate computer programs such that the obfuscations of functionally equivalent programs cannot be distinguished. Vinod Vaikuntanathan presented a unified framework in which a wide spectrum of cryptographic assumptions, ranging from the (very minimal) assumption by which OWF exist to the highly speculative IO conjecture. He also noted that the IO conjecture does not imply OWF (nor does it even imply $\mathcal{P} \neq \mathcal{NP}$).

Preventing false discovery in interactive data analysis. It may seem weird that such a title fits in a complexity theoretic workshop, but it turns out that a natural formulation of adaptive (or interactive) data analysis yields a natural computational problem. As explained by Jon Ullman, *interactive data analysis* refers to a setting in which first one obtains a sample of the data, and then one conducts a study of this sample by issuing queries and examining the answers (e.g., testing various hypotheses regarding the data). The point is that these queries are selected adaptively based on prior answers, and the problem is to avoid (false) discoveries that are tailored on the sample but do not reflect the original data. One key observation is that avoiding such a phenomenon is closely related to devising a “privacy preserving mechanism” for answering statistical queries to the data, whereas the design of such mechanisms is related to complexity theory. In particular, it was shown that if one-way functions exist, then false discoveries cannot be prevented when the researcher makes more than $\tilde{O}(n^2)$ queries to a sample of size n .

Additional surveys of wide areas. In addition to the aforementioned survey on fine-grained complexity, the meeting featured a large number of surveys of wide areas. These included:

- *A survey on lower bounds for low-depth arithmetic circuits.* The survey, presented by Neeraj Kayal, visited some of the main themes and techniques in this area, starting from the observation that sufficiently strong lower bound on the size of depth four circuits would yield such lower bounds for general arithmetic circuits (of unbounded depth).
- *Two surveys of recent directions in communication complexity.* The first survey, given by Mark Braverman, focused on the gap between the total length of the messages exchanged between two parties and the information contents of their interaction, raising the question of the extent by which an interactive communication can be compressed to its information contents. It is known that, in general, the best compression is to an exponential amount, but a quadratic amount is possible when the distribution of each input is independent of the distribution of the other input.

The relation between multi-party communication complexity and distributed computing was the focus of Rotem Oshman’s presentation, which highlighted the difference between the “local” model (where messages of unbounded length are allowed in each round) and the “congest” model (in which only short messages are allowed in each round). In both models, in each round, each party can only communicate with its neighbors in the fixed communication network.

- *Machine learning and complexity theory.* Rocco Servedio surveyed some of the known algorithms and lower bounds on the complexity of machine learning. He concluded his presentation suggesting to lower the expectations; that is, aim at better-than-obvious algorithms rather at algorithms that meet or approach the information-theoretic bound.

- *Random CSP instances and complexity theory.* Ryan O'Donnell surveyed the state-of-art regarding the complexity of solving random CSP instances, focusing on the use of the conjecture that it is hard to solve random instances of density that is close to the satisfiability threshold.

Informal specialized sessions. In addition to the formal plenary program, intense interaction between the participants took place in smaller groups. Part of these took place in the form of specialized sessions, some of which were already mentioned above. Other specialized sessions featured the following presentations.

- Amir Shpilka provided an inspiring exposition of a very recent construction of a (deterministic) quasi-NC algorithm for the bipartite matching problem. The said result by Fenner, Gurjar, and Thierauf was posted on ECCC a few days before the meeting (see TR15-177).
- Peter Bürgisser organized a specialized session on geometric complexity theory. This started by an outline of the geometric complexity theory program by him and then was followed by a report of Christian Ikenmeyer on recent advances in our understanding of the complexity of Kronecker coefficients. Klim Efremenko sketched the main ideas of his result on the limits of the method of shifted partial derivatives, which lead to an intense discussion with Neeraj Kayal and Pascal Koiran.
- In a specialized session on Coding theory, Venkatesan Guruswami reported recent advances on recovery of Reed-Solomon codes, and Amir Shpilka showed that Reed-Muller codes achieve the capacity of certain channels and gave a decoding algorithm from random errors in these codes.
- Or Meir organized a special session on open problems in Boolean circuit complexity. He presented a open question, which asks whether solving a computational problem on one of several distinct instances is easier than solving a single instance (this question is a close variant of a question posed by Beimel, Ben-Daniel, Kushilevitz, and Weinreb). Pascal Koiran presented an open problem in arithmetic circuit complexity which concerns finding an explicit polynomial that is hard to compute by polynomials of very restricted form. Oded Goldreich presented a line of research that concerns derandomization of randomized algorithms with very small error, and in particular, with respect to constant-depth circuits. Prasad Raghavendra presented an observation regarding a connection between the circuit complexity of a function and the properties of related polytopes.
- Avi Wigderson described (including extensive historical comments) his recent deterministic polynomial time algorithm for noncommutative rational identity testing (with Garg, Gurvits, and Oliveira). He highlighted the fact that questions and methods from very different origins (including invariant and representation theory, quantum information theorem, and optimization) interconnect and naturally combine for the solution of this problem.
- Boaz Barak talked about a Sum-of-Squares lower bound for the planted clique problem. He outlined the ideas behind a work in progress which

has still not been fully verified (with Sam Hopkins, Jon Kelner, Pravesh Kothari, Ankur Moitra and Aaron Potechin) showing that for every constant degree d and $\epsilon > 0$, the degree d Sum-of-Squares algorithm cannot certify that a random Erős-Rényi graph on n vertices does not contain a clique of size $n^{1/2-\epsilon}$.

- Separations in query and communication complexity: The aim of this session was to showcase some striking recent results (both for their strength and simplicity) giving separations, often tight, between various notions of query complexity for decision trees, and the surprising lifting of these bounds to similar separations between various models of communication complexity. These constitute progress on 30 year old questions in complexity theory. The first talk (given by Venkat Guruswami) discussed results for query complexity and the second talk (given by Raghu Meka) discussed the lifting approach for rectangle based measures of communication complexity.
 - The first talk was titled "Pointer Functions and Query complexity," given by Venkat Guruswami. It discussed a clever Boolean function construction of Goos-Pitassi-Watson which gives an optimal separation between nondeterministic and unambiguous decision tree complexities. It then discussed subsequent work by other authors showing that this function was also very useful in giving optimal quadratic separations between randomized and deterministic decision tree complexities and refuting an old 1986 conjecture by Saks and Wigderson on the largest possible gap between these models. The new developments lead to many more separations, such as between quantum query complexity and classical models, but this wasn't discussed in the talk.
 - Raghu Meka described the general method to transform query lower bounds into communication lower bounds for "composed functions". This is based on his recent works (with Mika Goos, Shachar Lovett, Thomas Watson, and David Zuckerman, and with Pravesh Kothari and Prasad Raghavendra). He presented ideas of the main structure theorem, which states that each rectangle in the communication matrix of the composed function can be simulated by a nonnegative combination of juntas. Consequently, this allows a characterization of the complexity of the composed functions in most known one-sided zero-communication models (capturing NP, co-NP, lower-bound measures such as corruption, smooth-rectangle bound, relaxed partition bound, *etc*) by a corresponding query complexity measure.
- Li-Yang Tan talked about his joint work with Ben Rossman and Rocco Servedio, where they proved an average-case depth hierarchy theorem for Boolean circuits over the standard basis of AND, OR, and NOT gates. The hierarchy theorem says that for every $d \geq 2$, there is an explicit n -variable Boolean function f , computed by a linear-size depth- d formula, which is

such that any depth- $(d - 1)$ circuit that agrees with f on $(1/2 + o_n(1))$ fraction of all inputs must have size $\exp(n^{\Omega(1/d)})$. This answers an open question posed by Håstad in his Ph.D. thesis.

- Gil Cohen presented his recent work on improved explicit constructions of Ramsey graphs. Erdős, in 1947 proved the existence of $2 \log n$ -Ramsey graphs on n vertices, and matching this result with a constructive proof is considered a central problem in combinatorics. The new result achieves an exponential improvement over previous results, and provides explicit $\exp((\log \log n)^c)$ -Ramsey graphs.
- Or Meir presented a new proof for a special case of the Karchmer, Raz, and Wigderson conjecture. If this conjecture is proved in full generality, it will imply super-polynomial formula lower bounds which is one of major challenges of the research in circuit complexity. While this case was already proved implicitly in Håstad's work on random restrictions, the new proof uses an entirely different approach based on communication complexity, and seems more likely to be generalizable to other cases of the conjecture.
- Alexander Razborov described recent work on Continuous Combinatorics as well as its context. He noted that Combinatorics was conceived, and then developed over centuries as a discipline about finite structures. However, currently, its applications increasingly pertain to structures that, although finite, are extremely large (e.g., the Internet network, social networks, statistical physics, to name just a few). Moreover, the numerical characteristics that researchers are normally interested in are "continuous" in the sense that small perturbations in the structure do not change the output very much. This makes it very natural to try to think of the "limit theory" of such objects by pretending that "very large" actually means "infinite". It turns out that this mathematical abstraction is very useful and instructive and leads to unexpected connections with many other things, both in mathematics and computer science. Two complementing approaches to constructing such a theory and applying it elsewhere are known as *graph limits* and *flag algebras*, and some of this theory was reviewed.
- Prasad Raghavendra presented exciting new results on lower bounds for linear programs and semidefinite programs based on his work (with Lee and Steurer).
- Zvika Brakerski described progress in the study of Fully Homomorphic Encryption (FHE) in the past couple of years. FHE is an encryption scheme that allows to compute arbitrary function "underneath" the encryption; that is, to go from $\text{Enc}(x)$ to $\text{Enc}(f(x))$ for all f , without any knowledge of the key. This allows to "outsource" computation to a third party without foregoing privacy. In particular, he focused on the "approximate eigenvector approach" based on work by Gentry, Sahai and Waters, and optimizing its performance via "sequentialization" based on joint his work with Vaikuntanathan.

- Raghu Meka described progress on constructing pseudorandom generators for small-space computation. His talk was based on his work with Gopalan and Daniel Kane.
- Zeev Dvir reviewed the Brascamp-Lieb Inequality and described its proof given by Franck Brathe. The proof introduces a normalization technique which allows one to apply a change of basis that puts a set of directions in radial isotropic positions. This technique has found several applications including Foster’s sign-rank lower bound and recent work on Sylvester-Gallai type theorems and Locally-Correctable Codes.
- Salil Vadhan, following up on Ryan O’Donnell’s plenary survey talk, described the proof of Daniely, Linial, and Shalev-Shwartz that polynomial-time PAC learning of DNF formulas (a long-standing open problem) is impossible if random k -SAT formulas on $n^{f(k)}$ clauses are hard to “refute” for some $f(k) \rightarrow \infty$. This led to an intensive small-group discussion on directions for obtaining stronger hardness results, and better understanding the relationship between learning and refutation.

Acknowledgement: The MFO and the workshop organizers would like to thank the National Science Foundation for supporting the participation of junior researchers in the workshop by the grant DMS-1049268, “US Junior Oberwolfach Fellows”. Moreover, the MFO and the workshop organizers would like to thank the Simons Foundation for supporting Luca Trevisan in the “Simons Visiting Professors” program at the MFO.

Workshop: Complexity Theory**Table of Contents**

David Zuckerman (joint with Eshan Chattopadhyay) <i>Two-Source Randomness Extractors: History and Context</i>	3061
Gil Cohen (reporting on the work of Chattopadhyay and Zuckerman) <i>Explicit Two-Source Extractors and Resilient Functions</i>	3064
Jonathan Ullman <i>Preventing False Discovery in Interactive Data Analysis</i>	3066
Ryan Williams (joint with Russell Impagliazzo, Daniel Marx, Mohan Paturi, and Virginia Vassilevska Williams) <i>Recent Work in Fine-Grained Complexity</i>	3069
Benjamin Rossman <i>Lower bounds for bounded-depth formulas</i>	3071
Ryan O'Donnell <i>Recent results concerning random kSAT</i>	3073
Mark Braverman <i>(Non)-compressibility of interactive communication: progress and challenges</i>	3076
Rotem Oshman (joint with Mark Braverman, Andrew Drucker, Fabian Kuhn) <i>The Role of Communication Complexity in Distributed Computing</i>	3078
Neeraj Kayal (joint with Chandan Saha, Ramprasad Saptharishi) <i>Lower bounds for low-depth arithmetic circuits</i>	3080
Or Meir (joint with Swastik Kopparty, Noga Ron-Zewi, and Shubhangi Saraf) <i>Recent developments in high-rate locally-testable and locally-decodable codes.</i>	3082
Avishay Tal (joint with Oded Goldreich) <i>Rigidity of Random Toeplitz Matrices with an Application to Depth-Three Circuits</i>	3085
Ron Rothblum (joint with Yael Tauman Kalai, Ran Raz) <i>How to Delegate Computations: the Power of No-Signaling Proofs</i>	3088
Zeev Dvir (joint with Sivakanth Gopi) <i>2-Server PIR with sub-polynomial communication</i>	3090

Rocco A. Servedio

*The complexity of learning Boolean functions: past progress and future
frontiers* 3092

Abstracts

Two-Source Randomness Extractors: History and Context

DAVID ZUCKERMAN

(joint work with Eshan Chattopadhyay)

The area of randomness extraction deals with the problem of obtaining nearly uniform bits from sources that are only weakly random. This is motivated by the ubiquitous use of randomness in various branches of computer science like algorithms, cryptography, and more. Further, most applications require truly random, uncorrelated bits, but most easily-obtainable sources of randomness do not satisfy these conditions. In particular, pseudorandom generators in practice try to accumulate entropy by using thermal noise or clock drift, but then this needs to be purified before using it to seed a pseudorandom generator.

We model a weak source on n bits using min-entropy. A source X on n bits is said to have min-entropy at least k if for any x , $\Pr[X = x] \leq 2^{-k}$. Any source X on $\{0, 1\}^n$ with min-entropy at least k is called an (n, k) -source.

An extractor $\text{Ext} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is a deterministic function that takes input from a weak source with sufficient min-entropy and produces nearly uniform bits. Unfortunately, a simple argument shows that it is impossible to design an extractor to extract even 1 bit for sources with min-entropy $n - 1$. To get past this difficulty, Santha and Vazirani [SV86], and Chor and Goldreich [CG88] suggested designing extractors for two or more independent sources, each with sufficient min-entropy. When the extractor has access to just two sources, it is called a two-source extractor. An efficient two-source extractor could be quite useful in practice, if just two independent sources of entropy can be found.

We measure the error of the extractor by statistical distance, or variation distance. We say $D_1 \approx_\epsilon D_2$ if $|D_1 - D_2| = \frac{1}{2} \sum_x |\Pr[D_1 = x] - \Pr[D_2 = x]|$.

We can now define a two-source extractor. A function $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ is called a two-source extractor for min-entropy k and error ϵ if for any independent (n, k) -sources X and Y , we have $\text{Ext}(X, Y) \approx_\epsilon U_m$, where U_m is the uniform distribution on m bits. Further, Ext is strong in Y if it also satisfies $|\langle \text{Ext}(X, Y), Y \rangle - \langle U_m, Y \rangle| \leq \epsilon$, where U_m is independent from Y . Note that for $m = 1$, this corresponds to an $N \times N$ matrix with entries in $\{0, 1\}$ such that every $K \times K$ submatrix has $1/2 \pm \epsilon$ fraction of 1's, where $N = 2^n$ and $K = 2^k$.

A simple probabilistic argument shows the existence of 2-source extractors for min-entropy $k \geq 2 \log n + 10 \log(1/\epsilon)$. However, in computer science, it is important to construct such functions explicitly, and this has drawn a lot of attention in the last three decades. Chor and Goldreich [CG88] used Lindsey's Lemma to show that the inner-product function is a 2-source extractor for min-entropy more than $n/2$. However, no progress was made on this problem for around 20 years, when Bourgain [Bou05] broke the "half-barrier" for min-entropy, and constructed a 2-source extractor for min-entropy $0.499n$. This remained the best known result prior to this work. Raz [Raz05] obtained an improvement in terms of total min-entropy,

and constructed 2-source extractors requiring one source with min-entropy more than $n/2$ and the other source with min-entropy $O(\log n)$.

The lack of progress on constructing two-source extractors motivated researchers to use more than two sources. Several researchers managed to construct excellent extractors using a constant number of sources, culminating in Li's construction of a 3-source extractor for polylogarithmic min-entropy [Li15c].

However, despite much attention and progress over the last 30 years, it remained open to explicitly construct two-source extractors for min-entropy rate significantly smaller than $1/2$. Our main result is an explicit two-source extractor for polylogarithmic min-entropy.

Main theorem: There exists a constant $C > 0$ such that for all $n \in \mathbb{N}$, there exists a polynomial time computable construction of a 2-source extractor $2\text{Ext} : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ for min-entropy at least $\log^C(n)$ and error $n^{-\Omega(1)}$. ($C=75$ suffices.)

By an argument of Barak, every 2-source extractor is also a strong 2-source extractor with similar parameters.

Note that an improvement of the output length of the above extractor to $c \log n$ bits, for a large enough constant c , will immediately allow one to extract $\Omega(k)$ bits using a standard trick of composition with a strong-seeded extractor.

Furthermore, improving the error to negligible while outputting many bits would have applications in cryptography and distributed computing.

Subsequent Work: Recently, Li [Li15b] extended our construction to achieve an explicit strong 2-extractor with output length k^α bits, for some small constant α . By our observation above, this immediately implies a 2-source extractor for min-entropy $k \geq \log^{C'} n$, for some large enough constant C' , with output length $\Omega(k)$; in fact, the output can be k bits.

Li also used our construction to build an affine extractor for polylogarithmic min-entropy [Li15a].

Ramsey Graphs: A key application is to Ramsey graphs. A graph on N vertices is called a K -Ramsey graph if does not contain any independent set or clique of size K . It was shown by Erdős in one of the first applications of the probabilistic method that there exists K -Ramsey graphs for $K = (2 + o(1)) \log N$. Since then, many researchers have tried to construct such Ramsey graphs explicitly. By explicit, we mean a polynomial-time algorithm that determines whether there is an edge between two nodes, i.e., the running time should be polylogarithmic in the number of nodes. The best construction before our work was due to Barak, Rao, Shaltiel, and Wigderson [BRSW12], who achieved $K = 2^{2^{\log^{1-\alpha}(\log N)}}$, for some absolute constant α .

In fact, these graphs are bipartite Ramsey graphs, which are harder to construct. A bipartite graph with N left vertices and N right vertices is called a bipartite K -Ramsey graph if it does not contain any complete $K \times K$ -bipartite sub-graph or empty $K \times K$ sub-graph. Given any bipartite K -Ramsey graph, a simple reduction gives a $K/2$ -Ramsey graph on N vertices.

It is easy to show that a 2-source extractor gives a bipartite Ramsey graph. Thus, as an immediate consequence of our Main Theorem, we obtain the following result.

Explicit Ramsey Graphs: There exists a constant $C > 0$ such that for all large enough $n \in \mathbb{N}$, there exists an explicit construction of a bipartite K -Ramsey graph on $2N$ vertices, where $N = 2^n$ and $K = 2^{(\log \log N)^C}$. (As before, $C=75$ suffices.)

In independent work¹, Cohen [Coh15] obtained an alternate explicit constructions of bipartite-Ramsey graphs with $K = 2^{(\log \log N)^{O(1)}}$.

Resilient Functions: A key ingredient in our construction is an explicit construction of a monotone, almost-balanced boolean function on n bits that is resilient to coalitions of size $n^{1-\delta}$, for any $\delta > 0$. In fact, our construction is stronger in that it gives an explicit extractor for a generalization of non-oblivious bit-fixing sources on n bits, where some unknown $n - q$ bits are chosen almost polylog(n)-wise independently, and the remaining $q = n^{1-\delta}$ bits are chosen by an adversary as an arbitrary function of the $n - q$ bits. The best previous construction, by Viola [Vio14], achieved $q = n^{1/2-\delta}$. Subsequently, Meka improved our result to $q = cn/\log^2 n$ [Mek15], which improves the constant C in our two-source extractor to 18 or 10 for polynomial and constant error, respectively.

REFERENCES

- [BRSW12] Boaz Barak, Anup Rao, Ronen Shaltiel, and Avi Wigderson. 2-source dispersers for $n^{o(1)}$ entropy, and Ramsey graphs beating the Frankl-Wilson construction. *Annals of Mathematics*, 176(3):1483–1543, 2012.
- [Bou05] J. Bourgain. More on the sum-product phenomenon in prime fields and its applications. *International Journal of Number Theory*, 01(01):1–32, 2005.
- [CG88] Benny Chor and Oded Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM Journal on Computing*, 17(2):230–261, 1988.
- [Coh15] Gil Cohen. Two-source dispersers for polylogarithmic entropy and improved Ramsey graphs. *Electronic Colloquium on Computational Complexity (ECCC)*, 2015.
- [Li15a] Xin Li. Extractors for affine sources with polylogarithmic entropy. Technical Report TR15-121, ECCC, 2015.
- [Li15b] Xin Li. Improved constructions of two-source extractors. *Electronic Colloquium on Computational Complexity (ECCC)*, 2015.
- [Li15c] Xin Li. Three-source extractors for polylogarithmic min-entropy. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science*, 2015.
- [Mek15] R. Meka. Explicit resilient functions matching Ajtai-Linial. Technical Report TR15-144, ECCC, 2015.
- [Raz05] Ran Raz. Extractors with weak random seeds. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 11–20, 2005.
- [SV86] Miklos Santha and Umesh V. Vazirani. Generating quasi-random sequences from semi-random sources. *Journal of Computer and System Sciences*, 33:75–87, 1986.
- [Vio14] Emanuele Viola. Extractors for circuit sources. *SIAM J. Comput.*, 43(2):655–672, 2014.

¹Cohen's work appeared before ours. When his paper appeared, we had an outline of the proof but had not filled in the details.

Explicit Two-Source Extractors and Resilient Functions

GIL COHEN

Generally speaking, an *extractor* is an algorithm that produces truly random bits given a sample from one or more “defective” sources of randomness. Many natural types of extractors are obtained by instantiating the notion of a defective source. One prominent example is a *multi-source extractor*. In this setting, the extractor is given s samples from independent sources, each supported on n -bit strings, and has min-entropy k (that is, no point is sampled by the source with probability higher than 2^{-k}). The goal is to design efficient extractors for small values of s, k . Computational aspects aside, one can prove the existence of an extractor for $s = 2$ sources with min-entropy $k = \log n + O(1)$. We refer to an extractor for $s = 2$ sources as a two-source extractor.

A long line of research was designated to match this existential result with an explicit construction, starting with the influential paper by Chor and Goldreich [6] who gave an explicit and simple construction of a two-source extractor for min-entropy $k > n/2$. Unlike extractors for $s > 2$ sources, constructing two-source extractors has witnessed very little progress, where only two improvements upon [6] were made. Raz [11] constructed a two-source extractor that is guaranteed to work assuming one source has min-entropy larger than $n/2$ while the second can have min-entropy as low as $O(\log n)$. Bourgain [3] designed a two-source extractor for min-entropy $(1/2 - \delta) \cdot n$ for some (small) universal constant $\delta > 0$. While the construction of Raz was based on small-biased sample spaces, the extractor of Bourgain applied deep results from additive combinatorics. Both techniques seem inadequate for supporting logarithmic or even poly-logarithmic min-entropy.

In this talk we present a striking recent result by Chattopadhyay and Zuckerman [5] who construct a two-source extractor for poly-logarithmic min-entropy. The construction makes use of three other types of extractors – non-malleable extractors, extractors for non-oblivious bit-fixing (NOBF) sources, and (the more familiar) strong seeded extractors. We start by giving a short account of the first two types of extractors.

Dodis and Wichs [8] introduced the notion of a *non-malleable extractor* as a tool for devising privacy amplification protocols against active adversaries. Informally speaking, a non-malleable extractor is a randomized algorithm that gets a single sample from a source with sufficient min-entropy, and has the following guarantee. With high probability over its internal randomness, the output of a non-malleable extractor is uniform even conditioned on the output obtained using an adversarial choice of internal randomness. After a fairly extensive study, we now have close-to-optimal explicit non-malleable extractors [4, 7].

A second tool used in the construction of [5] is an extractor for NOBF sources. This is a function that outputs a low biased bit given a sample from a distribution on n -bit strings having the following guarantee. Some $n - q$ of the bits of the sample are uniform and independent whereas the remaining q “bad” bits can be arbitrary functions of the $n - q$ “good” bits.

Constructing extractors for NOBF sources is related to the classical problem of collective coin flipping. Ajtai and Linial [1] proved the *existence* of an extractor for NOBF sources for $q = O(n/\log^2 n)$, and this is tight up to a logarithmic factor, as implied by [9]. As for *explicit* constructions, the largest value for q supported prior to [5] is $n^{0.63}$, which is attained by the recursive majority of three function.

One contribution of [5] is an explicit construction of an extractor for NOBF sources, denoted by **NOBFExt**, with $q = n^{1-\delta}$ for any desired constant $\delta > 0$. Further, their construction, which can be viewed as a derandomization of [1], has two important properties: (1) It is a monotone function, and (2) it is computable by a constant depth Boolean circuit of polynomial size in the input length. We mention here that a followup work by Meka [10] matches [1] with an explicit construction that also has these two properties.

Using these building blocks, we can present the two-source extractor of [5]. Let $x \sim X$, $y \sim Y$ be the two samples from the independent n -bit sources X, Y . First, one takes an optimal strong seeded extractor **Ext** and iterates over all $N = \text{poly}(n)$ seeds $i = 1, \dots, N$ to collect all N possible outputs $\{\text{Ext}(x, i)\}_i$. We set D , the output length of **Ext**, to equal the seed length for a non-malleable extractor **nmExt** with one output bit and error guarantee $2^{-\log^c n}$ for some suitable choice of a constant c . We then compute a length N string, which is denoted by $z = z(x, y)$, where $z_i = \text{nmExt}(y, \text{Ext}(x, i))$.

The analysis proceeds by showing that the random variable $Z = Z(X, Y)$ is close to a random variable with the following structure: There exists some small universal constant $\delta > 0$ such that $N - N^{1-\delta}$ of the bits of Z are jointly t -wise independent, where t is some constant that depends on our choice of c . Moreover, the identity of these $N - N^{1-\delta}$ “good” bits is a function of the underlining distributions X, Y and is fixed with respect to the samples x, y . The remaining $N^{1-\delta}$ bits of Z are arbitrary functions of the good bits. This almost amount to saying that Z is a NOBF source with $q = N^{1-\delta}$, though not quite – we are only guaranteed that the joint distribution of any t -tuple is uniform, though there might be strong correlations between these good bits. One can control the value of t but that will affect the entropy required by the sources. In particular, to support poly-logarithmic entropy in n , t can be taken to be no larger than poly-logarithmic in n .

To overcome this problem, Chattopadhyay and Zuckerman make use of the two properties of their extractor **NOBFExt** to show that this structure of Z suffices to conclude that **NOBFExt**(Z) has low bias. The reason for this lies in a fundamental result by Braverman [2] that shows that a constant depth circuit of polynomial-size cannot distinguish a truly uniform string on N bits from one that is sampled by a $(\log N)^b$ -wise independent distribution, where the value of b depends on the depth of the circuit. Using also the fact that their extractor is monotone, [5] shows that the bias of **NOBFExt**(Z) is close to the bias of **NOBFExt**(Z') where Z' is the distribution obtained by replacing the good bits of Z by truly uniform bits. This concludes the analysis.

REFERENCES

- [1] M. Ajtai, N. Linial, *The influence of large coalitions*, *Combinatorica* **2** (1993), 129–145.
- [2] M. Braverman, *Polylogarithmic independence fools AC0 circuits*, *Journal of the ACM* **5** (2010).
- [3] J. Bourgain, *More on the sum-product phenomenon in prime fields and its applications*, *International Journal of Number Theory* **1** (2005), 1–32.
- [4] E. Chattopadhyay, V. Goyal, X. Li, *Non-malleable extractors and codes, with their many tampered extensions*, *Electronic Colloquium on Computational Complexity (ECCC)* **75** (2015).
- [5] E. Chattopadhyay, D. Zuckerman, *Explicit two-source extractors and resilient functions*, *Electronic Colloquium on Computational Complexity (ECCC)* **119** (2015).
- [6] B. Chor, O. Goldreich, *Unbiased bits from sources of weak randomness and probabilistic communication complexity*, *SIAM Journal on Computing* **17** (1988), 230–261.
- [7] G. Cohen, *Non-malleable extractors – new tools and improved constructions*, *Electronic Colloquium on Computational Complexity (ECCC)* **183** (2015).
- [8] Y. Dodis, D. Wichs, *Non-malleable extractors and symmetric key cryptography from weak secrets*, *Proceedings of the forty-first annual ACM Symposium on Theory of Computing (2009)*, 601–610.
- [9] J. Kahn, G. Kalai, N. Linial, *The influence of variables on Boolean functions*, *29th Annual Symposium on Foundations of Computer Science (1988)*, 68–80.
- [10] R. Meka, *Explicit resilient functions matching Ajtai-Linial*, *arXiv preprint arXiv:1509.00092* (2015).
- [11] R. Raz, *Extractors with weak random seeds*, *Proceedings of the thirty-seventh annual ACM Symposium on Theory of Computing (2005)*, 11–20.

Preventing False Discovery in Interactive Data Analysis

JONATHAN ULLMAN

Multiple hypothesis testing is a ubiquitous task in empirical research. A finite sample of data is drawn from some unknown population, and several analyses are performed on that sample. The outcome of an analysis is deemed significant if it is unlikely to have occurred by chance alone, and a “false discovery” occurs if the analyst incorrectly declares an outcome to be significant. False discovery has been identified as a substantial problem in the scientific community (see e.g. [7, 5]). This problem persists despite decades of research by statisticians on methods for preventing false discovery.

In this extended abstract we briefly summarize a recent attempt by the theoretical computer science community to understand the role of *interactive data analysis*—repeatedly querying the same dataset in a way that depends on previous interactions with the dataset—in false discovery. This has quickly grown into a rich literature, giving new algorithms for preventing false discovery, and showing inherent computational and information-theoretic barriers to preventing false discovery. The problem of interactive data analysis was formalized by Dwork, Feldman, Hardt, Pitassi, Reingold, and Roth [3] and of Hardt and Ullman [6] as follows.

The object of study is a *population* P over some domain \mathcal{X} . We would like to estimate the answer to some sequence of *queries*

q_1, \dots, q_k on P . The queries are interactive in the sense that they are chosen online and q_i may depend on the answers to q_1, \dots, q_{i-1} . We would like an *accurate answer* a_i to each query q_i such that $a_i \approx q_i(P)$.

Of course P cannot be queried directly, so we assume that we have access to a *sample* S consisting of n independent draws from P . Can we design a (possibly randomized) procedure $M(S)$ such that for any interactive sequence of queries, with high probability over S and the randomness of M , $M(S)$ outputs accurate answers?

There are many types of queries that are considered in the literature [1], but for this discussion we focus on Kearns' *statistical queries model* [8]. These queries are specified by an efficiently computable predicate $\phi : \mathcal{X} \rightarrow \{\pm 1\}$ and $q_\phi(P)$ is defined to be $\mathbb{E}_{z \leftarrow P}[\phi(z)]$.

The most natural procedure $M(S)$ is to answer each query $q_\phi(P)$ with the *empirical answer* $q_\phi(S)$. If the queries were chosen non-interactively, then this approach answers accurately¹ as long as $k = 2^{o(n)}$. However, when the queries are chosen interactively it can fail to answer accurately even when $k = O(n)$, an exponential gap.

In our opinion, the three most relevant messages of this literature are:

- (1) Answering interactive queries is intimately related to *differential privacy* [4].
- (2) There are much better procedures than the naïve mechanism.
- (3) There are serious computational barriers to solving this formulation of the problem.

Of course, we refer the reader to many papers written on this subject, but we will now give a brief personally summary of this area.

We start with the definition of differential privacy. Informally, an algorithm is differentially private if changing one element of the sample does not change the *distribution* of the algorithm's outputs substantially. This can be viewed as a much stronger version of the types of stability conditions that are known to prevent false discovery in non-interactive data analysis.

Definition 1 ([4]). *A randomized algorithm $A : \mathcal{X}^n \rightarrow \mathcal{R}$ is (ϵ, δ) -differentially private if for every pair of samples S, S' that differ on at most one element, and every $R \subseteq \mathcal{R}$*

$$\mathbb{P}[A(S) \in R] \leq e^\epsilon \mathbb{P}[A(S') \in R] + \delta.$$

The key lemma to the whole endeavor states that differentially private algorithms cannot output queries that distinguish the sample from the population.

Informal Theorem 1 ([3, 1]). *Suppose A is a (ϵ, δ) -differentially private algorithm that takes a sample $S \in \mathcal{X}^n$ and outputs a statistical query q . Then*

$$\mathbb{P}[|q(P) - q(S)| \leq O(\epsilon)] \geq 1 - O(\delta/\epsilon),$$

where the probability is taken over the choice of S and the randomness of A .

¹For brevity, we are intentionally suppressing some parameters such as the precise definition of accuracy.

Once this lemma is in hand, it is not too hard to show that any procedure that is simultaneously *accurate with respect to its sample S* and *differentially private* must also give accurate answers with respect to the population P . Surprisingly, there is a rich literature on differential privacy giving algorithms with exactly those properties! This leads to the following “corollary.”

Informal Corollary 1 ([3, 1]). *There are procedures that prevent false discovery better than the naïve procedure. Specifically.*

- (1) *There is a procedure M that runs in time $\text{poly}(n, \log |\mathcal{X}|)$ per query and is accurate for $k \gtrsim n^2$ queries.*
- (2) *There is a procedure M that runs in time $\text{poly}(n, |\mathcal{X}|)$ per query and is accurate for $k \gtrsim 2^{n/\log^{1/2} |\mathcal{X}|}$ queries.*

We believe this is a strong result. However, the second bullet leads to some natural questions. Given that $|\mathcal{X}|$ can be huge—imagine the setting where $\mathcal{X} = 2^d$ where d is the *dimensionality* of the data—the second bullet point leads to an algorithm that is both *computationally inefficient* and *useless when n is smaller than the square root of the dimensionality of the data*. Can we remove these limitations? Unfortunately the answer was shown to be no.

Informal Theorem 2 ([6, 9]). *There are two strong limitations to preventing false discovery in interactive data analysis.*

- (1) *If one-way functions exist², then there is no procedure that runs in time $\text{poly}(n, \log |\mathcal{X}|)$ per query and answers significantly more than n^2 interactive queries.*
- (2) *Unconditionally, there is no procedure that answers significantly more than n^2 interactive queries when $n \ll \sqrt{\log |\mathcal{X}|}$, regardless of the running time of the procedure.*

We conclude with a remark about the way we formulated the problem. In practice, there are many common sources of interaction that likely capture a majority of relevant scenarios, so it may seem overly challenging for us to ask for this sort of “universal” procedure M to prevent false discovery. Certainly positive results in this challenging model are very strong, but more generally this model is tailored to asking questions about the inherent cost of allowing interactivity in data analysis. My personal hope is that researchers will use these results as a starting point to guide a more detailed investigation into the most useful procedures for interactive data analysis. In particular, we would like to highlight the work of Blum and Hardt [2], which investigates the application to data science competitions more closely, as an encouraging step in this direction.

²Equivalently, if secure private key cryptography exists.

REFERENCES

- [1] Raef Bassily, Kobbi Nissim, Adam D. Smith, Thomas Steinke, Uri Stemmer, and Jonathan Ullman. Algorithmic stability for adaptive data analysis. *CoRR*, abs/1511.02513, 2015.
- [2] Avrim Blum and Moritz Hardt. The ladder: A reliable leaderboard for machine learning competitions. *CoRR*, abs/1502.04585, 2015.
- [3] Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Roth. Preserving statistical validity in adaptive data analysis. In *ACM Symposium on the Theory of Computing (STOC)*. ACM, June 2015.
- [4] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *TCC*, pages 265–284. Springer, March 4-7 2006.
- [5] Andrew Gelman and Eric Loken. The statistical crisis in science. *American Scientist*, 102(6):460, 2014.
- [6] Moritz Hardt and Jonathan Ullman. Preventing false discovery in interactive data analysis is hard. In *FOCS*. IEEE, October 19-21 2014.
- [7] John P. A. Ioannidis. Why most published research findings are false. *PLoS Medicine*, 2(8):124, August 2005.
- [8] Michael J. Kearns. Efficient noise-tolerant learning from statistical queries. In *STOC*, pages 392–401. ACM, May 16-18 1993.
- [9] Thomas Steinke and Jonathan Ullman. Between pure and approximate differential privacy. *CoRR*, abs/1501.06095, 2015.

Recent Work in Fine-Grained Complexity

RYAN WILLIAMS

(joint work with Russell Impagliazzo, Daniel Marx, Mohan Paturi, and Virginia Vassilevska Williams)

This survey talk reported on some of the concepts and motivations behind the current semester-long program at the Simons Institute (UC Berkeley), entitled “Fine-Grained Complexity and Algorithm Design.” Here, I outline some of the main points from the survey talk.

A Focus on Time Complexity. In Fine-Grained Complexity as it is presently studied, one commits firmly to studying the *running time complexity* of problems in LOGSPACE, P, NP, PSPACE, and other well-studied classes. Since we believe all these complexity classes are different, we have different expectations for the time complexities. For example:

- For NP and PSPACE problems, we generally expect $\Theta(c^n)$ time complexity, for some constant $c > 1$. We attempt to minimize c as much as possible, by developing asymptotically faster algorithms. We observe the pure time complexity of some problems do not always match the original intuitions from complexity: while it is believed that the NP-complete SAT problem requires $\Omega(c^n)$ time for all $c < 2$ ([3]), the PSPACE-complete QBF problem can be solved in $O(1.8^n)$ time when the number of quantifier alternations is sufficiently high ([5]).
- For LOGSPACE and P problems, we expect $\Theta(n^c)$ time complexity, and attempt to minimize c similarly. But even some LOGSPACE problems

are probably very difficult to solve. For example, according to parameterized complexity theory, the k -clique problem, easily solvable in LOGSPACE for every constant k , is likely to require $n^{\Omega(k)}$ time to solve on n -node graphs [4].

Fine-Grained Reductions. Another defining characteristic of this subject is the notion of a “Fine-Grained reduction” ([6, 7]). Informally, the setup is as follows. Suppose we have two problems A and B and we believe that the time complexity of both problems is $t(n)^{1-o(1)}$. Very roughly speaking, a Fine-Grained reduction from A to B is an oracle reduction which for all $\varepsilon > 0$, there is some $\delta > 0$ such that, assuming B runs in $t(n)^{1-\varepsilon}$ time, the oracle reduction solves A in $t(n)^{1-\delta}$ time. (The actual definition is a little more technical but captures exactly this intuition, allowing many calls to B , even a variable adaptive number.)

That is, our philosophy is to *allow for the strongest possible kind of reduction from A to B* , such that if there is a significantly faster-than- $t(n)^{1-o(1)}$ algorithm for B , then the oracle reduction yields a faster algorithm for A . Because our reducibility notion is designed to be as strong as possible, we *maximize our chances* of proving a relationship between two problems A and B with similar running times. By choosing the most relaxed reducibility notion that is still sensible, we are able to prove reductions (and in some cases, *equivalences*) between many problems which would not be considered equivalent under weaker reducibility notions. See, for instance, [6, 2].

Cornerstones. Another aspect of Fine-Grained complexity is its focus on specific basic problems that have received considerable study, and have defied many attempts to be solved significantly faster:

- **SAT.** The SAT problem is the canonical NP-complete problem. Two key versions of the SAT problem are k -SAT and Circuit-SAT. While it is known that k -SAT is in $2^{n(1-1/O(k))}$ time for all constant k , it is conjectured that there is no $\delta > 0$ such that k -SAT is in $2^{n(1-\delta)}$ time for *all* constant k ; this is called the Strong Exponential Time Hypothesis (SETH) [3]. SETH has many intriguing consequences to the solvability of problems in both P and NP (see [1, 7]). There are also rather surprising consequences of faster Circuit-SAT algorithms to the area of circuit complexity: roughly speaking, solving Circuit-SAT on circuits of s size and n inputs in $O(2^n \cdot s/n^{\log n})$ time implies the circuit complexity lower bound $\text{NEXP} \not\subseteq \text{P/poly}$ [9].
- **3SUM.** The 3SUM problem asks whether we can find three numbers in a given set of n numbers which sum to 0. An $O(n^2)$ time algorithm is well-known, and it is conjectured that 3SUM cannot be solved in $n^{2-\varepsilon}$ time for every $\varepsilon > 0$.
- **OV.** The Orthogonal Vectors (OV) problem asks whether we can find a pair of vectors among n Boolean vectors (of d dimensions) which are orthogonal. An $O(n^2 d)$ time algorithm is obvious, and it is conjectured that OV cannot be solved in $n^{2-\varepsilon} \cdot 2^{o(d)}$ time for every $\varepsilon > 0$. (In fact it is known that a refutation of this conjecture would also refute SETH [8].)

- APSP. The all-pairs shortest paths (APSP) problem asks to compute the shortest distance between all pairs of nodes in a given weighted graph. An $O(n^3)$ time algorithm is well-known, and the exponent of 3 is conjectured to be optimal.

In Fine-Grained Complexity, one conjectures that basic problems such as the above cannot be solved much faster than their best known algorithms; from these problems and a large tapestry of Fine-Grained reductions among the problems, many fascinating consequences can be derived from the conjectures. The survey [7] gives many examples.

REFERENCES

- [1] Amir Abboud and Virginia Vassilevska Williams. Popular conjectures imply strong lower bounds for dynamic problems. In *FOCS*, pages 434–443, 2014.
- [2] Amir Abboud, Fabrizio Grandoni, and Virginia Vassilevska Williams. *Subcubic equivalences between graph centrality problems, APSP and diameter*. In Proceedings of SODA (2015), 1681–1697.
- [3] Chris Calabro, Russell Impagliazzo, and Ramamohan Paturi. *The complexity of satisfiability of small depth circuits*. In Proceedings of Parameterized and Exact Complexity (IWPEC) (2009), 75–85.
- [4] Rodney G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science, Springer (2013).
- [5] Rahul Santhanam and Ryan Williams. *Beating Exhaustive Search for Quantified Boolean Formulas and Connections to Circuit Complexity*. In Proceedings of SODA (2015) 231–241.
- [6] Virginia Vassilevska Williams and Ryan Williams. *Subcubic Equivalences between Path, Matrix and Triangle Problems*. In Proceedings of FOCS (2010) 645–654.
- [7] Virginia Vassilevska Williams. *Hardness of easy problems: Basing hardness on popular conjectures such as the strong exponential time hypothesis*. In Proceedings of Parameterized and Exact Computation (2015), 16–28.
- [8] Ryan Williams. *A new algorithm for optimal 2-constraint satisfaction and its implications*. Theor. Comput. Sci. **348(2-3)**(2005), 357–365.
- [9] Ryan Williams. *Improving exhaustive search implies superpolynomial lower bounds*. SIAM J. Comput. **42(3)** (2013), 1218–1244.

Lower bounds for bounded-depth formulas

BENJAMIN ROSSMAN

In this talk, I present a result (from a paper that appeared in FOCS 2015) showing that unbounded fan-in boolean formulas of depth $d + 1$ and size s have average sensitivity $O(\frac{1}{d} \log s)^d$. In particular, this gives a tight $2^{\Omega(d(n^{1/d}-1))}$ lower bound on the size of depth $d + 1$ formulas computing the PARITY function. These results strengthen the corresponding bounds for circuits due to Håstad (1986) and Boppana (1997).

Lower bounds against bounded-depth circuits were first proved in the 1980s [1, 3, 8, 4], culminating in a tight size-depth tradeoff for circuits computing the PARITY function. The technique, based on random restrictions, applies more generally to boolean functions with high average sensitivity.

Theorem 1 (Håstad [4]). Depth $d + 1$ circuits computing PARITY have size $2^{\Omega(n^{1/d})}$.

Theorem 2 (Boppana [2]). Depth $d + 1$ circuits of size s have average sensitivity $O(\log s)^d$.

Our results give stronger versions of Theorems 1 and 2 for bounded-depth formulas:

Theorem 3. Depth $d + 1$ formulas computing PARITY have size $2^{\Omega(d(n^{1/d}-1))}$.

Theorem 4. Depth $d + 1$ formulas of size s have average sensitivity $O(\frac{1}{d} \log s)^d$.

Theorems 3 and 4 directly strengthen Theorems 1 and 2 in light of the elementary fact that every depth $d + 1$ circuit of size s is equivalent to a depth $d + 1$ formula of size at most s^d . Moreover, Theorems 1,2,3,4 are asymptotically tight, since PARITY is computable by depth $d + 1$ circuits (resp. formulas) of size $n2^{O(n^{1/d})}$ (resp. $2^{O(d(n^{1/d}-1))}$).

The main tool in the proof of Theorems 3 and 4 is Håstad's Switching Lemma [4]. The Switching Lemma states that every small-width CNF or DNF simplifies, with high probability under a random restriction, to a small-depth decision tree. This yields lower bounds against bounded-depth *circuits* via a straightforward depth-reduction argument. In this paper we show how the Switching Lemma can be applied more efficiently to bounded-depth *formulas*, though in a less straightforward manner.

In more detail: for independent uniformly distributed random $\sigma \in \{0, 1\}^n$ ("assignment") and $\tau \in [0, 1]^n$ ("timestamp"), we consider the family of restrictions $\{R_p^{\sigma, \tau}\}_{0 \leq p \leq 1}$ (i.e. functions $[n] \rightarrow \{0, 1, *\}$ representing partial assignments to input variables x_1, \dots, x_n) where $R_p^{\sigma, \tau}$ sets the variable x_i to σ_i if $\tau_i < p$ and leaves x_i unset if $\tau_i \geq p$. In the usual application of the Switching Lemma to circuits of depth $d + 1$, all subcircuits of depth $k + 1$ are hit with the restriction $R_{p_k}^{\sigma, \tau}$ for a fixed sequence $p_1 > \dots > p_d$ (typically $p_k = n^{-k/(d+1)}$). In this paper we achieve sharper bounds against formulas by hitting each subformula Φ with the restriction $R_{\mathbf{q}(\Phi)}^{\sigma, \tau}$ where the parameter $\mathbf{q}(\Phi)$ ($= \mathbf{q}^{\sigma, \tau}(\Phi)$) is defined inductively, according to a random process indexed by subformulas of Φ . Our technical main theorem is a tail bound on $\mathbf{q}(\Phi)$, viewed as a random variable determined by σ and τ .

REFERENCES

- [1] Miklós Ajtai. Σ_1^1 formulae on finite structures. *Annals of Pure and Applied Logic*, 24:1–48, 1983.
- [2] Ravi B. Boppana. The average sensitivity of bounded-depth circuits. *Information Processing Letters*, 63(5):257–261, 1997.
- [3] Merrick L. Furst, James B. Saxe, and Michael Sipser. Parity, circuits, and the polynomial-time hierarchy. *Mathematical Systems Theory*, 17:13–27, 1984.
- [4] Johan Håstad. Almost optimal lower bounds for small depth circuits. In *18th Annual ACM Symposium on Theory of Computing*, pages 6–20, 1986.
- [5] V.M. Khrapchenko. Complexity of the realization of a linear function in the case of Π -circuits. *Math. Notes Acad. Sciences*, 9:21–23, 1971.

- [6] Benjamin Rossman. Formulas vs. circuits for small distance connectivity. In *46th Annual ACM Symposium on Theory of Computing*, pages 203–212, 2014.
- [7] P.M. Spira. On time-hardware complexity tradeoffs for Boolean functions. In *4th Hawaii Symposium on System Sciences*, pages 525–527, 1971.
- [8] Andrew C.C. Yao. Separating the polynomial-time hierarchy by oracles. In *26th Annual IEEE Symposium on Foundations of Computer Science*, pages 1–10, 1985.

Recent results concerning random k SAT

RYAN O'DONNELL

In this talk I gave a summary of some recent research concerning random constraint satisfaction problems (CSPs). As a prototypical example, consider the random 4SAT problem, with n variables and m independent random constraints (each chosen uniformly from the $\binom{n}{4}2^4$ possibilities). Naturally, the larger m is, the more like the CSP is to be unsatisfiable. As a rather trivial bound, suppose $m = 11n$. Then for each fixed $x \in \{0, 1\}^n$, the probability it satisfies all constraints is $(\frac{15}{16})^{11n} \approx .49^n$. Thus by union-bounding over all 2^n x 's, we see that the random 4SAT instances is unsatisfiable with overwhelming probability provided $m \geq 11n$. On the other hand, for $m \leq .99n$, say, it is not hard to show that the random 4SAT instance will be satisfied even “as 4XOR” with high probability. Thus the transition from satisfiability to unsatisfiability occurs somewhere in the range $m \in [.99n, 11n]$. Somewhat sharper bounds are known. For example, Frieze and Suen [1] gave an efficient algorithm that finds a satisfying assignment with high probability provided $m \leq 5.54n$.

Circa 2002–2005, sophisticated methods from statistical physics were used to (conjecturally) determine the exact threshold for satisfiability for random k SAT as a function of k , using the 1-Step Replica Symmetry Breaking (1RSB) method. [2, 3]. For example, the threshold for 4SAT is $m = 9.93n$. Here the number 9.93... is not an empirical estimation but is rather empirically determined: it is the least zero of a certain function related to a certain recurrence relation on probability distributions on $[0, 1]$. As mentioned, although these physics methods are highly sophisticated, they are ultimately conjectural. However, in a recent breakthrough, the 1RSB-based threshold predictions were rigorously mathematically verified by Ding, Sly, and Sun [4] (partly building on [5]) for all sufficiently large k .

When m is less than the satisfiability threshold, the natural algorithmic task is to efficiently find a satisfying assignment (with high probability). As mentioned, for 4SAT there is a rigorously proven efficient algorithm for finding a satisfying assignment for m as large as $5.54n$. Furthermore, certain statistical physics methods work for larger values of m , and there are sophisticated analytic methods used by the physicists to (conjecturally) determine the largest m for which they apply. For example, the Belief Propagation Guided Decimation method is “known” (by physics standards) to work for m as large as $9.05n$ [6]. Statistical physicists have also suggested heuristic algorithms — such as Survey Inspired Decimation [7] — that seem to work for even larger m , but which currently cannot be analyzed using physics methods. Instead, the statistical physicists analyze them

experimentally. A very notable recent example is the work of Marino, Parisi, and Ricci-Tersenghi [8], proposing the Backtracking Survey Propagation (BSP) algorithm. For 4SAT, this algorithm seems to efficiently find satisfying assignments for m as large as $9.90n$.

The general belief of physicists seems to be that efficient algorithms can work for m as large as, but not beyond, the “freezing threshold”; this is the least value of m for which all solutions belong to solution-“clusters” in which a constant fraction of the variables are all “frozen” to the same value. (There is also a slightly lower threshold called the “rigidity threshold”, at which *almost all* solutions belong to frozen clusters. The speaker admits to being slightly confused about the difference between the freezing and rigidity thresholds.) It seems that the BSP algorithm works beyond the rigidity threshold (by finding the “subdominant”, non-frozen clusters) but possibly not beyond the freezing threshold. We remark that Achlioptas and Ricci-Tersenghi have rigorously shown [9] that for sufficiently large k , the freezing threshold for k SAT is strictly below the satisfiability threshold. Thus it is expected that efficient algorithms can not work up to the satisfiability threshold for k SAT, k sufficiently large. On the other hand, for the most famous case of $k = 3$, it is a matter of debate whether or not the freezing threshold and the satisfiability threshold coincide. In particular, when applied to 3SAT, the BSP algorithm appears to work for m up to “essentially” the satisfiability threshold of $4.2667n$ (though it is difficult to be 100% certain of such empirically-drawn conclusions). In any case, it is certainly true that practical SAT-solvers do extremely well at solving random 3SAT for m very close to the threshold [10].

For m larger than the satisfiability threshold, the natural algorithmic task is to efficiently find a (with high probability) a *refutation* — i.e., a proof of unsatisfiability. In contrast to solving satisfiable instances just below the threshold, refuting random k SAT instances just above the threshold seems very difficult, both theoretically and practically. (E.g., SAT-solving competitions for this problem have been canceled due to lack of successful entrants.) Indeed, for 4SAT, it is only known how to efficiently refute random instances (with high probability) once $m \gg n^2$. (Herein the notation \gg glosses over log-factors.) This was first shown by Goerdt and Krivelevich [11], using spectral methods. They also showed that in general, random k SAT can be efficiently refuted once $m \gg n^{\lceil k/2 \rceil}$; and, for $k = 3$ specifically, Friedman and Goerdt [12] showed that $m \gg n^{3/2}$ suffices. Very recently, it was shown [13] that $m \gg n^{k/2}$ suffices for all k . Intriguingly, Feige, Kim, and Ofek [14] showed that one can go below $n^{3/2}$ for random 3SAT by using a *nondeterministic* algorithm. Specifically, once $m \gg n^{1.4}$ a random 3SAT formula has a polynomial-size refutation with high probability (but it is not known how to find it efficiently). Unpublished recent work of Feige and Witmer generalized this to k SAT for higher k ; e.g., for 4SAT, polynomial-size refutations exist with high probability once $m \gg n^{2-1/6}$, and for k SAT, once $m \gg n^{k/2-1/2+o_k(1)}$.

Feige’s “R3SAT Hypothesis” [15] essentially states that for $m = n^{1.01}$, there is no efficient refutation algorithm for random 3SAT. (Actually, the hypothesis is for

$m \geq Cn$ for all sufficiently large C , and he prefers to conjecture that there is no algorithm that can prove a random instance is at most $(1 - \epsilon)$ -satisfiable.) One may also consider the stronger, nondeterministic conjecture, that short refutations with high probability *do not exist*. The R3SAT Hypothesis is known to have many applications in inapproximability. Recently, Daniely and Shalev–Shwartz [16] made the related “RSAT Hypothesis”, that there exists some function $c(k) = \omega_{k \rightarrow \infty}(1)$ such that there is no efficient refutation algorithm for random k SAT when $m = n^{c(k)}$. Using this hypothesis, they showed a remarkable hardness-of-learning consequence: there is no efficient PAC-learning algorithm (even an improper one) for the class of DNF of size $\omega(\log n)$. (The work [16] was based on earlier work with Linial [17], which obtained stronger learning hardness results based on a stronger hypothesis about hardness of refuting random CSP; unfortunately, the stronger hypothesis was refuted [13].) One might even make an extremely strong conjecture: a random k SAT instance with $m = n^{4k}$ (say) does not have a polynomial-size refutation, with high probability. Perhaps this conjecture might have further consequences for inapproximability or hardness of learning.

REFERENCES

- [1] Frieze, Alan; Suen, Stephen. *Analysis of two simple heuristics on a random instance of k -SAT*, J. Algorithms **20** (1996), no. 2, 312–355.
- [2] Mézard, Marc; Parisi, Giorgio; Zecchina, Riccardo. *Analytic and algorithmic solution of random satisfiability problems*, Science, **297** (2002) 812–815.
- [3] Mertens, Stephan; Mézard, Marc; Zecchina, Riccardo. *Threshold values of random K -SAT from the cavity method*, Random Structures & Algorithms **28** (2006), no. 3, 340–373.
- [4] Ding, Jian; Sly, Allan; Sun, Nike. *Proof of the satisfiability conjecture for large k* , arXiv:1411.0650 (2014).
- [5] Coja-Oghlan, Amin. *The asymptotic k -SAT threshold*, arXiv:1310.2728 (2013).
- [6] Ricci-Tersenghi, Federico; Semerjian, Guilhem. *On the cavity method for decimated random constraint satisfaction problems and the analysis of belief propagation guided decimation algorithms*, J. Statistical Mechanics, **09** (2009).
- [7] Braunstein, Alfredo; Mzard, Marc; Zecchina, Riccardo. *Survey propagation: an algorithm for satisfiability*, Random Structures & Algorithms **27** (2005), no. 2, 201–226.
- [8] Marino, Raffaele; Parisi, Giorgia; Ricci-Tersenghi, Federico. *The Backtracking Survey Propagation algorithm for solving random K -SAT problems*, arXiv:1508.05117 (2014).
- [9] Achlioptas, Dimitris; Ricci-Tersenghi, Federico. *Random formulas have frozen variables*, SIAM J. Computing **39** (2009), no. 1, 260–280.
- [10] Gableske, Oliver. *Dimetheus*, Proceedings of SAT Competition (2014).
- [11] Goerdt, Andreas; Krivelevich, Michael. *Efficient recognition of random unsatisfiable k -SAT instances by spectral methods*, Proceedings of STACS (2001), 294–304.
- [12] Friedman, Joel; Goerdt, Andreas. *Recognizing more unsatisfiable random 3-SAT instances efficiently*, Proceedings of ICALP (2001), 310–321.
- [13] Allen, Sarah; O’Donnell, Ryan; Witmer, David. *How to refute a random CSP*, arXiv:1505.04383 (2015).
- [14] Feige, Uriel; Kim, Jeong Han; Ofek, Eran. *Witnesses for non-satisfiability of dense random 3CNF formulas*, Proceedings of FOCS (2006), 497–508.
- [15] Feige, Uriel. *Relations between average case complexity and approximation complexity*, Proceedings of STOC (2002), 534–543.

- [16] Daniely, Amit; Shalev-Shwartz, Shai. *Complexity theoretic limitations on learning DNFs*, arXiv:1404.3378 (2014).
- [17] Daniely, Amit; Linial, Nati; Shalev-Shwartz, Shai. *From average case complexity to improper learning complexity*, Proceedings of STOC (2014), 441–448.

(Non)-compressibility of interactive communication: progress and challenges

MARK BRAVERMAN

In this talk we discussed the following question, which has recently received quite a bit of attention:

“Alice and Bob communicate by executing a protocol π on their inputs. When can π be compressed — i.e. simulated by a shorter protocol — and by how much?”

This interactive compression question turns out to be equivalent to the direct sum and direct product problems for randomized two-party communication complexity [4, 5, 7]. The question is formalized in terms of Shannon’s information theory. A detailed discussion of the question as well as the definitions and background can be found in the recent survey [12].

Alice and Bob are given inputs (X, Y) distributed according to a joint distribution μ . The information cost of a protocol π is the amount of information Alice and Bob learn about each other’s inputs from executing π . π is allowed to use both public randomness (available to both Alice and Bob) and private randomness (generated privately by each party, and not accessible to the other party). Let Π denote the random variable representing the transcript of the execution of π , including its public but not private randomness. Then the (internal) information cost of π is given by

$$(1) \quad IC(\pi, \mu) := I(X; \Pi|Y) + I(Y; \Pi|X).$$

The first term in the sum represents the (expected) amount of information observing Π reveals to Bob (who knows Y) about X . The second term represents the amount of information Alice learns about Y . Note that this quantity does not only depend on π but also on μ . For example, if μ is a distribution supported on a single value $\{(x, y)\}$, then $IC(\pi, \mu) = 0$ for any protocol π .

The *external* information cost of π is defined as the amount of information observing Π reveals to an external observer:

$$(2) \quad IC_{ext}(\pi, \mu) := I(XY; \Pi).$$

Denote by $|\pi|$ the communication cost of π , that is, the maximum number of bits exchanged by the participants of π . It is not hard to prove that

$$IC(\pi, \mu) \leq IC_{ext}(\pi, \mu) \leq |\pi|.$$

Moreover, the first \leq is an equality when $\mu = \mu_x \times \mu_y$ is a product distribution over the inputs.

The interactive compression problem can be formulated as follows: given a protocol π with information cost, external information cost, and communication

cost I , I_{ext} , and C , respectively, what is the communication cost of the shortest protocol π' that can simulate π to within error, say, $1/3$?

Classical results of Shannon and Huffman, which show that a single message X can be transmitted using at most $H(X) + 1$ bits can be viewed as a special case of this problem where $Y = \perp$ is an empty input.

General state-of-the-art compression results are as follows [1, 2]:

- communication can be compressed to $\tilde{O}(\sqrt{I \cdot C})$;
- communication can be compressed to $O(I_{ext} \cdot \text{polylog}(C))$;
- communication can be compressed to $2^{O(I)}$.

These compression results left the question of whether compression to $O(I)$ communication is possible open. In a series of recent papers by Ganor, Kol, and Raz [8, 9, 10], a negative answer has been given to this question. Specifically, these papers show that:

- There exists a protocol whose communication cost is $C = 2^{2^k}$, and whose information cost is k , but which are not compressible to less than $2^{O(k)}$ communication. This means that the last upper bound above cannot be improved, and the first upper bound cannot be improved to anything better than $O(I \cdot \log C)$. We still do not know whether $O(I \cdot \log C)$ or $O(\sqrt{I \cdot C})$ is the right answer.
- There exists a protocol whose communication cost is $C = 2^{2^{2^k}}$, and whose *external* information cost is k , but which are not compressible to less than $2^{O(k)}$ communication. This means that the second upper bound above cannot be improved to better than $O(I_{ext} \cdot \log \log C)$.

In the talk we discussed the proof of the second result, which was very recently given in [10]. It relies on analyzing a variant of the ‘‘Hidden Layers Game’’ proposed in [3]. The hardness of this game relies, among other things, on the fact that solving the Greater Than function on n -bit integers requires $\Omega(\log n)$ communication [11, 6]. What is remarkable about the [10] reduction is that it is ‘‘protocol dependent’’: it does not use the protocol as a black box, but rather uses the assumption that the too-good-to-be-true simulation protocol exists to turn it into a too-good-to-be-true protocol for a different problem for which a lower bound is known.

REFERENCES

- [1] B. Barak, M. Braverman, X. Chen, A. Rao. How to Compress Interactive Communication. *SIAM J. Comput.* **42**(3): 1327-1363, 2013.
- [2] M. Braverman. Interactive information complexity, *STOC'12*.
- [3] Mark Braverman. A hard-to-compress interactive task? *Allerton (2013)*: 8-12.
- [4] M. Braverman, A. Rao. Information equals amortized communication, *IEEE Transactions on Information Theory*, **60**(10): 6058-6069, 2014.
- [5] M. Braverman, A. Rao, O. Weinstein, A. Yehudayoff. Direct products in communication complexity, *FOCS'13*.
- [6] M. Braverman, O. Weinstein. A Discrepancy Lower Bound for Information Complexity, *APPROX-RANDOM'12*.

- [7] M. Braverman, O. Weinstein. An interactive information odometer with applications, *STOC'15*.
- [8] A. Ganor, G. Kol, R. Raz. Exponential Separation of Information and Communication, *FOCS'14*.
- [9] A. Ganor, G. Kol, R. Raz. Exponential Separation of Information and Communication for Boolean Functions, *STOC'15*.
- [10] A. Ganor, G. Kol, R. Raz. Exponential Separation of Communication and External Information. *Electronic Colloquium on Computational Complexity (ECCC)* **22**: 88 (2015).
- [11] E. Viola. The communication complexity of addition. *Electronic Colloquium on Computational Complexity (ECCC)* **18**: 152 (2011).
- [12] O. Weinstein. Information Complexity and the quest for interactive compression. *SIGACT News* **46**(2): 41-64, 2015.

The Role of Communication Complexity in Distributed Computing

ROTEM OSHMAN

(joint work with Mark Braverman, Andrew Drucker, Fabian Kuhn)

Distributed algorithms must deal with several challenges that sequential algorithms do not have to face. First, the input to the computation is usually divided between the participants, with no single participant having a global picture; for example, in distributed graph algorithms, no single node knows the entire graph, and yet the nodes wish to compute some globally consistent solution (e.g., a spanning tree or a coloring). Second, communication between the participants is costly and sometimes restricted. Congestion is a serious issue, so we want to reduce the number of messages and bits sent; in addition, it is expensive to *synchronize* the participants, so we want to use as few *communication rounds* as possible. These restrictions have led to distributed algorithms that strongly focus on *communication efficiency*.

One popular model that brings out the communication aspect is the CONGEST model, defined as follows: we have a network of n nodes, modelled as a (usually undirected) graph G . The edges of G represent communication links between the nodes. The nodes have unique identifiers, but they initially do not know the graph G . Various goals are considered, including testing properties of G , finding various subgraphs in G (e.g., a spanning tree or a Hamiltonian cycle), or computing functions of local inputs to the nodes. The computation proceeds in synchronous rounds: in each round, each node sends at most B bits on each of its edges in G , receives the messages sent by its neighbors, and then the next round begins. Lower bounds in this model typically seek to show that certain tasks require many rounds — usually even when the diameter of the network graph G is small — because of the communication restriction.

Many lower bounds for the CONGEST model proceed by reduction from Yao's *two-party communication model* [1]. In Yao's classical model, there are two players, Alice and Bob, who receive private inputs X, Y respectively, and wish to compute a joint function $f(X, Y)$ of their inputs. We study how many bits Alice and Bob must exchange in order to compute f , possibly using randomization and with some

probability of error; this is called the *communication complexity* of f . Crucially, Alice and Bob are not restricted in computation power.

A typical reduction from 2-party communication complexity to CONGEST proceeds as follows. We fix some “hard” function f , for which we know a high communication complexity lower bound. We assume for the sake of contradiction that there is a distributed algorithm A for some task T using a small number of rounds in the CONGEST model. We then construct a two-player protocol P_A from the protocol A , where the players construct some network graph using their inputs to f , *simulate* the execution of A in this network, and finally use the output of A to compute the output to f . If the reduction is designed cleverly, simulating the execution of A does not require much communication between the two players, so we get a protocol for f that does not have high enough communication complexity, violating the lower bound for f .

This type of reduction has led to many lower bounds in the CONGEST model (see, e.g., [2]), but it has some inherent weaknesses: by “splitting” the network between only two players, we may give away too much of what makes the original distributed problem hard. Therefore, recent work has started looking at *multi-party communication complexity*, where instead of two players we have a larger number $k > 2$, and we directly study the amount of communication or the number of rounds required to solve various problems. We consider both *point-to-point* communication, where participants send each other individual messages (as in the CONGEST model), and *broadcast* communication, where each participant can send a single message to all the other participants together (this is the classical *shared blackboard* model).

The *symmetrization* technique, introduced by Phillips, Verbin and Zhang in [3], allows us to *lift* two-part lower bounds to k -party lower bounds in some scenarios. However, for some lower bounds, symmetrization is not sufficient; proving a lower bound on the Set Disjointness problem, where the players receive inputs X_1, \dots, X_k and wish to check whether $X_1 \cap \dots \cap X_k = \emptyset$, required the use of *information complexity*, extending notions from classical information theory to the multi-party interactive setting [4, 5].

We also consider some novel ways in which the input can be partitioned between the participants; for example, in some cloud applications, it can be beneficial to *duplicate* parts of the input across different nodes, so that more than one node can look at each part of the input at the same time. Having nodes share parts of their input can make it harder to prove lower bounds, as can be witnessed by the notorious difficulty of proving lower bounds in the number-on-forehead model (where each player sees the entire input except one small part). Nevertheless, in our setting only a small part of the input is shared between any two nodes, and we hope that this will make lower bounds more tractable. In this context, one simple question that remains open is the hardness of *triangle detection*, in a model where vertex of the graph is “assigned” to some player, and this player receives all the edges adjacent to the vertex. Because we are interested in undirected graphs, each

edge of the graph is known to two players (those assigned its endpoints). What is the communication complexity of randomized triangle detection in this model?

REFERENCES

- [1] A. Yao, *Some complexity questions related to distributive computing (Preliminary Report)*, Proc. 11th Symp. on Theory of Comp. (STOC), 209–213.
- [2] A. Das Sarma, S. Holzer, L. Kor, A. Korman, D. Nanongkai, G. Pandurangan, D. Peleg, R. Wattenhofer, *Distributed Verification and Hardness of Distributed Approximation*, SIAM J. Comput. Vol. 41 No. 5 (2012), 1235–1265.
- [3] J. Phillips, E. Verbin, Q. Zhang, *Lower bounds for number-in-hand multiparty communication complexity, made easy*, Proc. 23rd Symp. on Discrete Algorithms (SODA), 486–501.
- [4] M. Braverman, F. Ellen, R. Oshman, V. Vaikuntanathan, T. Pitassi, *A Tight Bound for Disjointness in the Message-Passing Model*, Proc. 54th Symp. on Found. of Comp. Science (FOCS), 668–677.
- [5] M. Braverman, R. Oshman, *On Information Complexity in the Broadcast Model*, Proc. 2015 Symp. on Princ. of Dist. Comp. (PODC), 355–364.

Lower bounds for low-depth arithmetic circuits

NEERAJ KAYAL

(joint work with Chandan Saha, Ramprasad Saptharishi)

An arithmetic circuit computes a polynomial function over some underlying field \mathbb{F} via a sequence of operations involving $+$ and \times starting from its inputs x_1, x_2, \dots, x_N . We typically allow arbitrary constants from \mathbb{F} on the incoming edges to a $+$ gate so that a $+$ gate can in fact compute an arbitrary \mathbb{F} -linear combination of its inputs. The complexity of a circuit is measured in terms of its size (the number of edges in the corresponding graph) and depth (the maximum length of a path in the corresponding graph). A central open problem in this area is to prove arithmetic circuit lower bounds (for some explicit family of polynomials). In this talk we highlight an ongoing effort towards proving lower bounds for (subclasses of) arithmetic circuits via proving *strong enough* lower bounds for low-depth arithmetic circuits. We illustrate this via a superpolynomial lower bound for a subclass of arithmetic circuits that we refer to as *regular arithmetic formulas* (based on [1]).

Definition (Regular Arithmetic Formulas). We say that an arithmetic circuit is a regular formula if:

- (1) The underlying graph is a tree consisting of alternating layers of $+$ and \times gates, and
- (2) all the nodes at a layer have the same fanin, and
- (3) the formal degree of the output node is at most a constant factor (say twice) more than d , the degree of the polynomial computed by the formula.

In this talk we presented a $n^{\Omega(\log n)}$ lower bound for regular arithmetic formulas. The proof highlights the theme of trying to prove lower bounds for large-depth circuits/formulas via proving lower bounds for low-depth circuits/formulas. The proof consists of the following four steps.

- (1) **Depth Reduction.** In the case of regular formulas one reduces to depth four - if Φ is a regular formula of size $2^{o(\log^2 d)}$ computing a polynomial f_d of degree d then for some $t = \Omega(\log d)$ there exists a representation of f of the form

$$f = T_1 + T_2 + \dots + T_s,$$

where each T_i is a product of $O(\frac{d}{t})$ -many polynomials of degree t and $s = 2^{o(\frac{d}{t} \cdot \log d)}$.

- (2) **Identifying a Geometric Property π .** One then tries to identify a *weakness* of such representations by pinpointing interesting geometric properties of the geometric variety of a term T_i . Recall that the variety corresponding to a polynomial T , denoted $\mathbb{V}(T)$, is the set of all zeroes of the polynomial T , i.e.

$$\mathbb{V}(T) = \{\mathbf{a} \in \mathbb{F}^n : T(\mathbf{a}) = 0\}.$$

In our case, we note that when T is a product of many polynomials then $\mathbb{V}(T)$ has lots of high-order singularities. This is the geometric property that we use.

- (3) **Translating the property π into smallness of rank of a matrix.** We then try to associate a matrix $M(f)$ to any polynomial f such that the following two properties hold:
- (a) **Linearity.** For any two polynomials f and g and any two constants $\alpha, \beta \in \mathbb{F}$, it holds that $M(\alpha f + \beta g) = \alpha M(f) + \beta M(g)$, and
 - (b) **Smallness of rank.** If the variety of any polynomial f has the property π identified above then the rank of $M(f)$ is significantly smaller than its size.

In our case, we do this by looking at the set of polynomials $(\mathbf{x}^{\ell}) \cdot (\partial^k f)$, view polynomial as a row of an appropriate matrix and look at the rank of the matrix.

- (4) **Finding an explicit polynomial f such that rank of $M(f)$ is large.** We finally find an explicit polynomial f such that $M(f)$ has large rank. The matrix $M(f)$ is typically very huge but remarkably we are able to prove lower bounds on rank of $M(f)$ via two simple tools:
- (a) **Via existence of a large triangular submatrix.** If $M(f)$ contains an upper-triangular square submatrix U (with nonzero entries on the diagonal) then the size of U is a lower bound on the rank of $M(f)$.
 - (b) **Via near-orthogonality of the columns of the matrix.** A beautiful lemma commonly attributed to Noga Alon intuitively says that if the columns (or the rows) of a matrix are *almost orthogonal* then the matrix has *nearly full rank*. Specifically for any matrix M over the real numbers:

$$\text{rank}(M) \geq \frac{\text{Tr}(M^T \cdot M)^2}{\text{Tr}((M^T \cdot M)^2)}.$$

In our case we construct an explicit family of polynomials based on Nisan-Wigderson designs for which rank $M(f)$ can be shown to be large using either of the two observations above.

Discussion. We are able to successfully implement the proof strategy outlined above in the case of regular formulas but we hope that this can be successfully implemented for more general and interesting classes of circuits. It is one of the ongoing lines of research in the area.

REFERENCES

- [1] N. Kayal, C. Saha, R. Saptharishi, *A super-polynomial lower bound for regular arithmetic formulas*, STOC **46** (2014), 146–153.

Recent developments in high-rate locally-testable and locally-decodable codes.

OR MEIR

(joint work with Swastik Kopparty, Noga Ron-Zewi, and Shubhangi Saraf)

Locally-decodable codes [2, 17] and locally-testable codes [9, 20, 10] are error-correcting codes that admit local algorithms for decoding and testing respectively. More specifically, those algorithms are only allowed to make a small number of queries to their input, and this number is called the *query complexity*. Clearly, such algorithms must be randomized, and err with some probability.

LDCs and LTCs were originally studied in the setting where the query complexity was either constant or poly-logarithmic. In those settings, it is believed that LDCs and LTCs must be very redundant. Hence, we do not expect such codes to achieve a high rate. In particular, in the setting of constant query complexity, it is known that LDCs cannot have constant rate [17, 22, 23], and that LTCs with certain restrictions cannot have constant rate [7, 4]. On the other hand, the best-known constant-query LDCs have super-polynomial length [24, 8, 16, 5], and the best-known constant-query LTCs have quasi-linear length (see e.g. [3, 6, 21]).

It turns out that the picture is completely different when allowing the query complexity to be much larger. In this setting, it has long been known that one can have LDCs and LTCs with constant rate and query complexity $O(n^\beta)$ for constant $\beta > 0$ [2, 20]. More recently, it has been discovered that both LDCs [18, 11, 14] and LTCs [21, 11] can simultaneously achieve rates that are arbitrarily close to 1 and query complexity $O(n^\beta)$ for an arbitrary constant $\beta > 0$. This is in contrast with the general belief that local correctability and testability require much redundancy.

In this work, we show that there are LDCs and LTCs with constant rate (which can in fact be taken to be arbitrarily close to 1) and constant relative distance, whose associated local algorithms have $n^{o(1)}$ query complexity and running time. We find it quite surprising in light of the fact that there were several quite different constructions of LDCs and LTCs [2, 20, 18, 21, 11, 14] with constant rate and constant relative distance, all of which had $\Omega(n^\beta)$ query complexity. Specifically:

- For LDCs, we obtain query complexity and running time $\exp(\sqrt{\log n \cdot \log \log n})$.
- For LTCs, we obtain query complexity and running time $(\log n)^{O(\log \log n)}$.

Furthermore, we show that such codes can achieve stronger trade-offs between the rate and relative distance than was known before. Specifically:

- Over the binary alphabet, our codes meet the Zyablov bound, which is a much better trade-off than what seemed achievable in the past.
- Over large alphabets (of constant size), our codes approach the Singleton bound: they achieve a tradeoff between rate and distance which is essentially as good as possible for general error-correcting codes. This means that, remarkably, local correctability and local testability with $n^{o(1)}$ queries over large alphabets is not only possible with constant rate and constant relative distance, but it also does not require “paying” anything in terms of rate and relative distance.

Such results were previously not known for any $o(n)$ query complexity.

Our constructions are based on a technique of Alon, Edmonds, and Luby [1]. We observe that this technique can be viewed as a method for distance amplification. This distance amplifier, based on a d -regular expander, converts an error-correcting code with relative distance $\gg 1/d$ into an error-correcting code with larger relative distance δ , while reducing the rate only by a factor of $\approx (1 - \delta)$. Thus for a large enough constant d , if we start with a code of rate $1 - \varepsilon$ and relative distance $\gg 1/d$, where $\varepsilon \ll \delta$, then after distance amplification with a d -regular expander, we get a code with rate $(1 - \delta)(1 - \varepsilon) \approx (1 - \delta)$ and relative distance δ .

This technique was used in [1, 12] to construct linear-time codes, in [13] to construct list-decodable codes with small alphabet, and (following our work) in [15] to construct linear-time list-recoverable codes. All those constructions shared a similar outline: One first constructs codes with high rate with some (possibly very small) constant relative distance and a certain desirable property. Then, applying distance amplification with a (possibly very large) constant-degree expander, one obtains a code with a much better tradeoff between its rate and relative distance. Finally one shows that the distance amplification with a constant degree expander preserves the desirable property.

The first main observation of this paper is that the distance-amplification technique also preserves the property of being an LDC or an LTC. Specifically, if we start with an LDC or LTC with query complexity q , and then apply distance amplification with a d -regular expander, then the resulting code is an LDC/LTC with query complexity $q \cdot \text{poly}(d)$.

The next main observation is that this connection continues to hold even if we take d to be super-constant, and take the LDC or LTC to have sub-constant relative distance $\Theta(1/d)$. This is potentially useful, since we only blow up the query complexity by a factor of $\text{poly}(d)$, and perhaps LDCs/LTCs with high rate and sub-constant relative distance can have improved query complexity over their constant relative distance counterparts. As far as we are aware, there have been no

previous uses of the distance-amplification technique using an expander of super-constant degree.

We construct our LDCs by showing that an existing family of high-rate LDCs can achieve sub-polynomial query complexity if we only require them to have sub-constant relative distance. Specifically, multiplicity codes [18] in a super-constant number of variables give us the desired LDCs. In order to obtain LTCs with query complexity $(\log n)^{O(\log \log n)}$, we use an iterative construction that combines tensor products and the AEL distance-amplification in a sophisticated way, following [19].

REFERENCES

- [1] N. Alon, J. Edmonds, and M. Luby, *Linear time erasure codes with nearly optimal recovery*, in *proceedings of the 36th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, IEEE Computer Society (1995), 512–519.
- [2] L. Babai, L. Fortnow, L. A. Levin, and M. Szegedy, *Checking computations in polylogarithmic time*, in *STOC* (1991), 21–31.
- [3] E. Ben-Sasson and M. Sudan, *Short PCPs with polylog query complexity*, *SIAM J. Comput.* 38 (2008)(2), 551–607, preliminary version in *STOC* 2005.
- [4] E. Ben-Sasson and M. Viderman, *Towards lower bounds on locally testable codes via density arguments*, *Computational Complexity* 21 (2012)(2), 267–309.
- [5] Y. M. Chee, T. Feng, S. Ling, H. Wang, and L. F. Zhang, *Query-efficient locally decodable codes of subexponential length*, *Computational Complexity* 22 (2013)(1), 159–189.
- [6] I. Dinur, *The PCP theorem by gap amplification*, *Journal of ACM* 54 (2007)(3), 241–250, preliminary version in *STOC* 2006.
- [7] I. Dinur and T. Kaufman, *Dense locally testable codes cannot have constant rate and distance*, in *APPROX-RANDOM* (2011), 507–518.
- [8] K. Efremenko, *3-query locally decodable codes of subexponential length*, *SIAM J. Comput.* 41 (2012)(6), 1694–1703.
- [9] K. Friedl and M. Sudan, *Some improvements to total degree tests*, in *ISTCS* (1995), 190–198.
- [10] O. Goldreich and M. Sudan, *Locally testable codes and PCPs of almost linear length*, *Journal of ACM* 53 (2006)(4), 558–655, preliminary version in *FOCS* 2002, pages 13–22.
- [11] A. Guo, S. Kopparty, and M. Sudan, *New affine-invariant codes from lifting*, in *ITCS* (2013), 529–540.
- [12] V. Guruswami and P. Indyk, *Linear-time encodable/decodable codes with near-optimal rate*, *IEEE Transactions on Information Theory* 51 (2005)(10), 3393–3400.
- [13] V. Guruswami and A. Rudra, *Explicit codes achieving list decoding capacity: Error-correction with optimal redundancy*, *IEEE Transactions on Information Theory* 54 (2008)(1), 135–150.
- [14] B. Hemenway, R. Ostrovsky, and M. Wootters, *Local correctability of expander codes*, in *ICALP (1)* (2013), 540–551.
- [15] B. Hemenway and M. Wootters, *Linear-time list recovery of high-rate expander codes*, in *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I* (2015), 701–712.
- [16] T. Itoh and Y. Suzuki, *Improved constructions for query-efficient locally decodable codes of subexponential length*, *IEICE Transactions* 93-D (2010)(2), 263–270.
- [17] J. Katz and L. Trevisan, *On the efficiency of local decoding procedures for error-correcting codes*, in *STOC* (2000), 80–86.
- [18] S. Kopparty, S. Saraf, and S. Yekhanin, *High-rate codes with sublinear-time decoding*, *J. ACM* 61 (2014)(5), 28.
- [19] O. Meir, *Combinatorial construction of locally testable codes*, *SIAM J. Comput.* 39 (2009)(2), 491–544.

- [20] R. Rubinfeld and M. Sudan, *Robust characterization of polynomials with applications to program testing*, SIAM Journal of Computing 25 (1996)(2), 252–271.
- [21] M. Viderman, *A combination of testability and decodability by tensor products*, Random Struct. Algorithms 46 (2015)(3), 572–598.
- [22] S. Wehner and R. de Wolf, *Improved lower bounds for locally decodable codes and private information retrieval*, in *Automata, Languages and Programming, 32nd International Colloquium, ICALP 2005, Lisbon, Portugal, July 11-15, 2005, Proceedings* (2005), 1424–1436.
- [23] D. P. Woodruff, *New lower bounds for general locally decodable codes*, Electronic Colloquium on Computational Complexity (ECCC) 14 (2007)(006).
- [24] S. Yekhanin, *Towards 3-query locally decodable codes of subexponential length*, J. ACM 55 (2008)(1).

Rigidity of Random Toeplitz Matrices with an Application to Depth-Three Circuits

AVISHAY TAL

(joint work with Oded Goldreich)

This talk concerns the construction of rigid matrices, a central open problem posed by Valiant [7], and its application to lower bounds on *canonical* depth-three Boolean circuits (where “canonical” is as defined by Goldreich and Wigderson [2]). In particular, we improve the known lower bound on matrix rigidity, but the improvement is for a range of parameters that is not the one motivated by Valiant’s problem, but rather the one that arises from [2]. Indeed, this improvement resolves open problems posed by Goldreich and Wigderson [2].

Matrix Rigidity. The “Matrix Rigidity Problem” (i.e., providing explicit matrices of high rigidity) is one of the most alluring problems in arithmetic circuits lower bounds. Introduced in 1977 by Valiant [7], the problem was originally motivated by proving lower bounds for the computation of linear transformations. A matrix A over a field \mathbb{F} has rigidity s for rank r if every matrix of rank at most r (over \mathbb{F}) disagrees with A on more than s entries.

Valiant showed that any matrix with rigidity $n^{1+\delta}$ for rank $\omega(n/\log \log n)$, where δ is some constant greater than 0, cannot be computed by a linear circuit of size $O(n)$ and depth $O(\log n)$. Valiant also proved that almost all n -by- n matrices, over a finite field \mathbb{F} (e.g., the two-element field \mathbb{F}_2), have rigidity $\Omega((n-r)^2/\log n)$ for rank r . Since then, coming up with an explicit (i.e., matrices that can be computed by a polynomial time algorithm) rigid matrix has remained a challenge. The best techniques to date provide explicit n -by- n matrices of rigidity $\frac{n^2}{r} \log \frac{n}{r}$ for rank r (see [4] for a survey about matrix rigidity).

To the best of our knowledge, this state of affairs also holds for “simple” randomized constructions that use $O(n)$ random bits. The common belief is that rigidity bounds for such randomized constructions can be used for proving lower bounds for explicit computational problems that are related to the original ones. For example, an adequate rigidity lower bound for random Toeplitz (or Hankel)

matrices¹ would yield a lower bound on the complexity of computing explicit bilinear transformations. Indeed, this is analogous to Andreev's proof of formula lower bounds [1], where a lower bound for a randomized function is transformed into a lower bound for an explicit function (which takes the random bits of the construction as part of its input). Our main result is the following

Theorem 1 (on the rigidity of random Toeplitz/Hankel matrices). *Let $A \in \mathbb{F}_2^{n \times n}$ be a random Toeplitz/Hankel matrix. Then, for every $r \in [\sqrt{n}, n/32]$, with probability $1 - o(1)$, the matrix A has rigidity $\Omega(\frac{n^3}{r^2 \log n})$ for rank r .*

Our bounds are asymptotically better than $\Omega(\frac{n^2}{r} \log \frac{n}{r})$ for rank $r = o(\frac{n}{\log n \cdot \log \log n})$, alas Valiant's original motivation refers to $r > n/\log \log n$. For rank $r = n^{0.5+\varepsilon}$, where $\varepsilon \in (0, 0.5)$, our bound yields a significant improvement (i.e., $\frac{n^3}{r^2} = n^{2-2\varepsilon} \gg n^{1.5-\varepsilon} = \frac{n^2}{r}$), and this is actually the range that is relevant for the project of [2].

The Project of Goldreich-Wigderson. The work of Goldreich and Wigderson [2] provides another motivation for the study of matrix rigidity. In fact, the problem of improving the rigidity bounds for random Toeplitz matrices was posed explicitly there. Specifically, proving a rigidity bound of $n^{1.5+\Omega(1)}$ for rank $n^{0.5+\Omega(1)}$ for random Toeplitz matrices was proposed there as a possible next step.

Lower Bounds for Depth Three Canonical Circuits. Håstad [3] showed that any depth-three Boolean circuit² computing the n -way parity function must be of size at least $\exp(\sqrt{n})$. Though Håstad's bound was refined during the years [6, 5], to date, $\exp(\Omega(\sqrt{n}))$ is the best lower bound for an explicit function in the model of depth-three Boolean circuits. The work of Goldreich and Wigderson [2] put forward a model of *depth three canonical circuits*, with the underlying long-term goal to exhibit better lower bounds for general depth-three Boolean circuits. Canonical circuits are restricted type of such depth-three circuits, which can be illustrated by considering the smallest known depth-three circuits for n -way parity. The latter $\tilde{O}(2^{\sqrt{n}})$ -size circuits are obtained by combining a CNF that computes a \sqrt{n} -way parity with \sqrt{n} DNFs that compute \sqrt{n} -way parities of disjoint blocks of the input bits. The construction, and its optimality by [3, 6], suggests the following scheme for obtaining Boolean circuits that compute multilinear functions. First, construct an arithmetic circuit that uses arbitrary multilinear gates of parameterized arity, and then convert it to a Boolean circuit whose size is exponential in the maximum between the arity and the number of gates in the arithmetic circuit. The arithmetic model is outlined next.

¹Recall that a Toeplitz matrix $T = (T_{i,j})$ has constant diagonals (i.e., $T_{i,j} = T_{i+1,j+1}$ for every i, j). Hankel matrices are obtained by turning Toeplitz matrices upside down; that is, a Hankel matrix $H = (H_{i,j})$ has constant skew-diagonals (i.e., $H_{i,j} = H_{i+1,j-1}$ for every i, j).

²That is, a circuit of unbounded fan-in OR and AND gates with leaves that are variables or their negations.

Lower Bounds for Multilinear Circuits. Suppose we wish to compute a t -linear function that depends on t blocks of inputs, $x^{(1)}, \dots, x^{(t)}$, each of length n ; that is, the function is linear in each of the $x^{(j)}$'s. We consider circuits that use arbitrary t -linear gates of parameterized arity. That is, the circuits are directed acyclic graphs, where each internal node computes a t -linear function of its inputs. We further restrict our circuit such that each internal gate computes a multilinear formal polynomial in the inputs $x^{(1)} \dots, x^{(t)}$. We say that such a multilinear circuit is of AN-complexity³ m if m equals the maximum between the number of the circuit gates and the maximal arity of the gates. For a t -linear function F , we denote by $\mathcal{C}(F)$ the minimal AN-complexity of a multilinear circuit which compute the function F . (We will abuse notation and refer to the AN-complexity of a tensor/matrix as the AN-complexity of the corresponding t -linear function.)

In the example of parity, we have a bottom layer of \sqrt{n} gates each taking \sqrt{n} inputs and computing their parity. Above these gates, we have a gate which takes the \sqrt{n} results and computes their parity. Overall, we got a (multi)-linear circuit of AN-complexity $\sqrt{n} + 1$.

Goldreich and Wigderson showed that any multilinear circuit of AN-complexity m yields a depth-three Boolean circuit of size $\exp(m)$ computing the same function (see [2, Prop. 2.9]). In fact, the Boolean circuits have much more structure, and are referred to by Goldreich and Wigderson as *canonical circuits*. Thus, a preliminary step towards beating the $\exp(\Omega(\sqrt{n}))$ lower bound on the size of depth-three Boolean circuits for explicit $O(1)$ -linear functions,⁴ will be to beat the $\Omega(\sqrt{n})$ AN-complexity lower bound for such functions in the model of multilinear circuits.

Again, as in Valiant's question, if we just ask about the existence of hard t -linear functions, then almost all t -linear functions cannot be computed by a multilinear circuit of AN-complexity smaller than $(nt)^{t/(t+1)}$: See [2, Thm. 4.1], which uses a counting argument. The more important and challenging problem is to come up with an explicit t -linear function for which such bounds, or even just $\omega(\sqrt{n})$ lower bounds, can be proved.

Reduction to Matrix Rigidity. Goldreich and Wigderson reduces the problem of proving lower bounds for bilinear circuits to the problem of rigidity [2, Sec. 4.2]. They show that if a bilinear circuit is of AN-complexity $m/2$, then its corresponding matrix is not m^3 rigid for rank m (i.e., it can be expressed as a sum of an m^3 -sparse matrix and a matrix of rank at most m over \mathbb{F}_2). Hence, any matrix that has rigidity m^3 for rank m corresponds to a bilinear function that cannot be computed by a bilinear circuit of AN-complexity at most $m/2$.

Open Problems in Goldreich-Wigderson. One open problem posed by Goldreich and Wigderson is proving that *random* Toeplitz matrices have rigidity m^3 for rank $m = n^{0.5+\Omega(1)}$. This would yield an AN-complexity lower bound of m for the

³where AN stands for Arity and Number of gates.

⁴Indeed, this suggestion presumes that there exist $O(1)$ -linear functions that require depth-three Boolean circuits of size $\exp(\omega(\sqrt{n}))$, which is also an open problem suggested in [2].

corresponding bilinear function (via the reduction in [2, Thm. 4.4]) as well as a similar lower bound for the following *explicit* trilinear function (via [2, Prop. 4.6]):

$$(1) \quad F_{\text{tet}}(x, y, z) = \sum_{\substack{i_1, i_2, i_3 \in [n]: \\ \sum_{j=1}^3 |i_j - n/2| \leq n/2}} x_{i_1} y_{i_2} z_{i_3} .$$

Resolving the Foregoing Open Problems. We resolve the aforementioned open problem by proving that random Toeplitz matrices have rigidity m^3 for rank $m = \Theta\left(\frac{n^{3/5}}{\log^{1/5} n}\right)$, with high probability. This follows from Theorem 1 by choosing $r = m$. This implies that the AN-complexity of a random Toeplitz matrix is $\tilde{\Omega}(n^{3/5})$, and ditto for the explicit trilinear function F_{tet} from Eq. (1).

A natural open problem is to prove better than $\exp(\sqrt{n})$ lower bounds on the size of (standard) depth-three circuits computing F_{tet} .

REFERENCES

- [1] A. E. Andreev. On a method for obtaining more than quadratic effective lower bounds for the complexity of π -schemes. *Moscow Univ. Math. Bull.*, 42:63–66, 1987. In Russian.
- [2] O. Goldreich and A. Wigderson. On the size of depth-three boolean circuits for computing multilinear functions. *Electronic Colloquium on Computational Complexity (ECCC)*, 20:43, 2013.
- [3] J. Håstad. Almost optimal lower bounds for small depth circuits. In *Proceedings of the 18th Annual STOC*, pages 6–20, 1986.
- [4] S. V. Lokam. Complexity lower bounds using linear algebra. *Foundations and Trends in Theoretical Computer Science*, 4(1-2):1–155, 2009.
- [5] R. Paturi, P. Pudlák, M. E. Saks, and F. Zane. An improved exponential-time algorithm for k -sat. *J. ACM*, 52(3):337–364, 2005.
- [6] R. Paturi, P. Pudlák, and F. Zane. Satisfiability coding lemma. *Chicago J. Theor. Comput. Sci.*, 1999, 1999.
- [7] L. G. Valiant. Graph-theoretic arguments in low-level complexity. In *Lecture notes in Computer Science*, volume 53, pages 162–176. Springer, 1977.

How to Delegate Computations: the Power of No-Signaling Proofs

RON ROTHBLUM

(joint work with Yael Tauman Kalai, Ran Raz)

The problem of delegating computation considers a setting where one party, the *delegator* (or *verifier*), wishes to delegate the computation of a function f to another party, the *worker* (or *prover*). The challenge is that the delegator may not trust the worker, and thus it is desirable to have the worker “prove” that the computation was done correctly. We require that verifying this proof is significantly easier than doing the computation itself; that is, the delegator’s running time is significantly smaller than the time complexity of f . Moreover, we require that the running time of the worker is not much larger than the time complexity of f .

The problem of delegating computation has become a central problem in cryptography, especially with the increasing popularity of cloud computing, where weak devices use cloud platforms to run their computations.

We focus on the problem of constructing *one-round* delegation protocols, where the delegator wants to verify a statement of the form $x \in \mathcal{L}$. The delegator sends x to the worker together with some query q ; then the worker computes $b = \mathcal{L}(x)$, and based on the query q provides a *non-interactive* proof π for the fact that $b = \mathcal{L}(x)$. The delegator should be able to verify the correctness of the proof π very efficiently, and the worker should run in time polynomial in the time it takes to compute f . Throughout this work (similarly to all previous works that consider the problem of one-round delegation), the security requirement is against *computationally bounded* cheating workers. Namely, we consider the computational setting, where the security (i.e., soundness) of our scheme relies on a cryptographic assumption, and the guarantee is that any cheating worker, who cannot break the underlying assumption, cannot prove the correctness of an incorrect statement.

Previously, a combination of a result of Goldwasser, Kalai and Rothblum [GKR08] with a result of Kalai and Raz [KR09] yielded a one-round delegation scheme for any function f that can be computed by a LOGSPACE-uniform circuit C of size $t = t(n)$ and depth $d = d(n)$, where the running time of the verifier is $\tilde{O}(n + d)$, and the running time of the prover is $\text{poly}(t)$ (assuming the existence of a sub-exponentially secure computational private information retrieval scheme). Note however that for circuits with large depth d this delegation scheme does not satisfy the efficiency criterion.

A fundamental question is: Do there exist efficient 1-round delegation schemes for *all* deterministic computations? There are several works that (partially) answer this question in the preprocessing model, or under *non-falsifiable* assumptions.¹ We elaborate on these prior works in the full version [KRR13].

In this work, we answer the above question positively, by constructing a 1-round delegation scheme for *every* deterministic computation, assuming a sub-exponentially secure computational private information retrieval (PIR) scheme. More specifically, we show a delegation scheme for every language computable in time $t = t(n)$, where the running time of the verifier is $n \cdot \text{polylog}(t)$, and the running time of the prover is $\text{poly}(t)$. The underlying assumption is that there exists a computational PIR scheme (or a fully homomorphic encryption scheme) that cannot be broken in time $t^{\text{polylog}(t)}$ for security parameter $k \leq \text{poly}(n)$.²

Our delegation scheme exploits a connection to the seemingly unrelated model of multi-prover interactive proof systems (\mathcal{MIP}) in which soundness holds even against *no-signaling* cheating provers. Loosely speaking, no-signaling provers are allowed to use arbitrary strategies (as opposed to local ones, where the reply of each prover is a function only of her own input), as long as their strategies cannot be used for communication between any two disjoint sets of provers.

We show that any \mathcal{MIP} that is sound against no-signaling cheating provers can be converted into a 1-round delegation scheme, using a fully-homomorphic

¹We note that under non-falsifiable assumptions, there are known positive results even for non-deterministic computations. The focus of this work is on deterministic computations.

²In particular, for languages in \mathcal{P} we only require a PIR scheme with quasi-polynomial security.

encryption scheme (FHE), or alternatively, using a computational private information retrieval (PIR) scheme. This connection is based on a heuristic that was originally suggested by Aiello *et-al* [ABOR00].

We then construct a new \mathcal{MIP} , for every deterministic language, with soundness against no-signaling cheating provers. This, together with the transformation above, gives us our 1-round delegation scheme.

REFERENCES

- [ABOR00] William Aiello, Sandeep Bhatt, Rafail Ostrovsky, and S. Raj. Rajagopalan. Fast verification of any remote procedure call: Short witness-indistinguishable one-round proofs for NP. In *ICALP: Annual International Colloquium on Automata, Languages and Programming*, 2000.
- [GKR08] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: interactive proofs for muggles. In *STOC*, pages 113–122, 2008.
- [KR09] Yael Tauman Kalai and Ran Raz. Probabilistically checkable arguments. In *CRYPTO*, pages 143–159, 2009.
- [KRR13] Yael Tauman Kalai, Ran Raz, and Ron D. Rothblum. How to delegate computations: The power of no-signaling proofs. *Electronic Colloquium on Computational Complexity (ECCC)*, 20:183, 2013.

2-Server PIR with sub-polynomial communication

ZEEV DVIR

(joint work with Sivakanth Gopi)

Private Information Retrieval (PIR) was first introduced by Chor, Goldreich, Kushilevitz and Sudan [CKGS98]. In a k -server PIR scheme, a user can retrieve the i th bit a_i of a n -bit database replicated among k servers (which do not communicate) while giving no information about i to any server. The goal is to design PIR schemes that minimize the communication cost which is the worst case number of bits transferred between the user and the servers in the protocol. The trivial solution which works even with one server is to ask a server to send the entire database, which has communication cost n .

When $k = 1$ the trivial solution cannot be improved [CKGS98]. But when $k \geq 2$, the communication cost can be brought down significantly. In [CKGS98], a 2-server PIR scheme with communication cost $O(n^{1/3})$ and a k -server PIR scheme with cost $O(k^2 \log kn^{1/k})$ were presented. The k -server PIR schemes were improved further in subsequent papers [Amb97, BI01, BIKR02]. This was the best for a long time until the breakthrough results of Yekhanin[Yek08] and Efremenko[Efr09] gave k -server PIR schemes with sub-polynomial cost for $k \geq 3$ which were slightly improved in [IS10]. These PIR schemes follow from the constructions of constant query smooth Locally Decodable Codes (LDCs) of sub-exponential length called Matching Vector Codes (MVCs)[DGY10]. A k -query LDC [KT00] is an error correcting code which allows the receiver of a corrupted encoding of a message to recover the i th bit of the message using only k (random) queries. In a *smooth* LDC, each query of the reconstruction algorithm is uniformly distributed among

the code word symbols. Given a k -query smooth LDC, one can construct a k -server PIR scheme by letting each server simulate one of the queries. Despite the advances in 3-server PIR schemes, the 2-server PIR case is still stuck at $O(n^{1/3})$ since 2-query LDCs provably require exponential size encoding [KdW03] (which translates to polynomial communication cost in the corresponding PIR schemes).

On the lower bounds side, there is very little known. The best known lower bound for the communication cost of a 2-server PIR is $5 \log n$ [WdW05] whereas the trivial lower bound is $\log n$. In [CKGS98], a lower bound of $\Omega(n^{1/3})$ is conjectured. An $\Omega(n^{1/3})$ lower bound was proved for a restricted model of 2-server PIR called bilinear group based PIR in [RY06]. This model encompasses all the previously known constructions which achieve $O(n^{1/3})$ cost for 2-server PIR. So to beat $O(n^{1/3})$, we need to go beyond this bilinear group based model and fortunately the constructions of Yekhanin [Yek08] and Efremenko [Efr09] are not captured by this model. Thus they provide us some clues to break the $O(n^{1/3})$ barrier.

Our main result gives the first 2-server PIR scheme with sub polynomial communication. Our proof augments the Matching vector based protocols of Yekhanin and Efremenko with the use of partial derivatives.

REFERENCES

- [Amb97] Andris Ambainis. Upper bound on communication complexity of private information retrieval. In *ICALP*, pages 401–407, 1997.
- [BI01] Amos Beimel and Yuval Ishai. Information-theoretic private information retrieval: A unified construction. In *ICALP*, pages 912–926, 2001.
- [BIKR02] Amos Beimel, Yuval Ishai, Eyal Kushilevitz, and Jean-François Raymond. Breaking the $\omega(n^{1/(2k-1)})$ barrier for information-theoretic private information retrieval. In *FOCS*, pages 261–270, 2002.
- [CKGS98] Benny Chor, Eyal Kushilevitz, Oded Goldreich, and Madhu Sudan. Private information retrieval. *J. ACM*, 45(6):965–981, 1998.
- [DGY10] Zeev Dvir, Parikshit Gopalan, and Sergey Yekhanin. Matching vector codes. In *FOCS*, pages 705–714, 2010.
- [Efr09] Klim Efremenko. 3-query locally decodable codes of subexponential length. In *STOC*, pages 39–44, 2009.
- [IS10] Toshiya Itoh and Yasuhiro Suzuki. Improved constructions for query-efficient locally decodable codes of subexponential length. *IEICE Transactions*, 93-D(2):263–270, 2010.
- [KdW03] Iordanis Kerenidis and Ronald de Wolf. Exponential lower bound for 2-query locally decodable codes via a quantum argument. In *STOC*, pages 106–115, 2003.
- [KT00] Jonathan Katz and Luca Trevisan. On the efficiency of local decoding procedures for error-correcting codes. In *32nd ACM Symposium on Theory of Computing (STOC)*, pages 80–86, 2000.
- [RY06] Alexander A. Razborov and Sergey Yekhanin. An $\Omega(n^{1/3})$ lower bound for bilinear group based private information retrieval. In *FOCS*, pages 739–748, 2006.
- [WdW05] Stephanie Wehner and Ronald de Wolf. Improved lower bounds for locally decodable codes and private information retrieval. In *ICALP*, pages 1424–1436, 2005.
- [Yek08] Sergey Yekhanin. Towards 3-query locally decodable codes of subexponential length. *J. ACM*, 55(1), 2008.
- [Yek12] Sergey Yekhanin. Locally decodable codes. *Foundations and Trends in Theoretical Computer Science*, 6(3):139–255, 2012.

The complexity of learning Boolean functions: past progress and future frontiers

ROCCO A. SERVEDIO

The talk surveyed computationally efficient learning algorithms for concept classes (classes of Boolean functions, i.e. functions mapping $\{0, 1\}^n$ to $\{-1, 1\}$) that are interesting and natural from a complexity theoretic perspective. It focused on algorithms in the Probably Approximately Correct (PAC) learning model that was introduced and studied by Vapnik and Chervonenkis [VC71] and Valiant [Val84], and in the closely related model of exact learning from equivalence queries [Ang88] (also known as the online mistake-bound model [Lit88]).

One technique that has proved highly effective in developing efficient learning algorithms is based on applying polynomial-time linear programming over an expanded “feature space” of all low-degree monomials. Since there are $O(n^d)$ such monomials of degree at most d , this leads to $n^{O(d)}$ time learning algorithms for concept classes which are such that every function in the class has a polynomial threshold function representation of degree at most d . This approach has been used to obtain the fastest known algorithms for learning decision trees (implicit in [Blu92]), DNF formulas [KS01], de Morgan formulas of bounded size [Lee09], and intersections of low-weight halfspaces [BRS95, KOS04]. However, strong lower bounds have been given on the polynomial threshold function degree of intersections of halfspaces [She09, She10] and AC^0 circuits [She14, She15], giving corresponding limitations on the efficiency of algorithms based on learning polynomial threshold functions for these concept classes. More generally, Razborov and Sherstov [RS10] have established limitations on the ability of linear programming based methods to efficiently learn DNF formulas regardless of the particular feature space that is employed (as long as it is fixed a priori in advance of running the learning algorithm).

Another technique, based on linear algebra, employs a “closure algorithm” to learn all AND-of-XOR-of-AND $_d$ circuits (where the bottom level AND gates have maximum fan-in d , but higher-level gates have unbounded fanin) in $n^{O(d)}$ time [FS92, HSW92]; however, it is not clear what concept classes have representations of this sort when d is constrained to be significantly less than n . Other techniques, based on an augmentation of Rivest’s [Riv87] “top-down” algorithm for learning decision lists [HS07] and based on an “agnostic boosting” technique of Kalai, Mansour and Verbin [KMV08], were also briefly surveyed.

A final question is whether learning rich concept classes (such as intersections of halfspaces, AC^0 circuits, etc) is a sufficiently hard computational task that the correct perspective on algorithms for these problems is to measure the *savings* that can be achieved over a brute-force running time of 2^n . (This seems to be the correct perspective for related computational problems, such as counting satisfying assignments, for some of these classes; see e.g. [BIS12, IPS13].) The talk concluded with a sketch of a “proof of concept” result of this sort, establishing that

any poly(n)-size, depth- d AC^0 circuit can be learned in the online mistake-bound learning model with running time at most $2^{n-n^{1/d}}$.

REFERENCES

- [Ang88] D. Angluin. Queries and concept learning. *Machine Learning*, 2:319–342, 1988.
- [BIS12] P. Beame, R. Impagliazzo, and S. Srinivasan. Approximating AC^0 by Small Height Decision Trees and a Deterministic Algorithm for $\#AC^0$ -SAT. In *CCC*, pages 117–125, 2012.
- [Blu92] Avrim Blum. Rank- r decision trees are a subclass of r -decision lists. *Information Processing Letters*, 42(4):183–185, June 1992.
- [BRS95] R. Beigel, N. Reingold, and D. Spielman. PP is closed under intersection. *Journal of Computer & System Sciences*, 50(2):191–202, 1995.
- [FS92] P. Fischer and H. Simon. On learning ring-sum expansions. *SIAM Journal on Computing*, 21(1):181–192, 1992.
- [HS07] L. Hellerstein and R. Servedio. On PAC learning algorithms for rich boolean function classes. *Theoretical Computer Science*, 384(1):66–76, 2007.
- [HSW92] D. Helmbold, R. Sloan, and M. Warmuth. Learning integer lattices. *SIAM Journal on Computing*, 21(2):240–266, 1992.
- [IPS13] Russell Impagliazzo, Ramamohan Paturi, and Scott Schneider. A satisfiability algorithm for sparse depth two threshold circuits. In *Proceedings of the 54th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 479–488. IEEE, 2013.
- [KMV08] A. Kalai, Y. Mansour, and E. Verbin. On agnostic boosting and parity learning. In *Proc. 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 629–638, 2008.
- [KOS04] A. Klivans, R. O’Donnell, and R. Servedio. Learning intersections and thresholds of halfspaces. *Journal of Computer & System Sciences*, 68(4):808–840, 2004.
- [KS01] Adam Klivans and Rocco A. Servedio. Learning DNF in time $2^{\tilde{O}(n^{1/3})}$. In *Proc. 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 258–265. ACM Press, 2001.
- [Lee09] T. Lee. A note on the sign degree of formulas. Available at <http://arxiv.org/abs/0909.4607>, 2009.
- [Lit88] N. Littlestone. Learning quickly when irrelevant attributes abound: a new linear-threshold algorithm. *Machine Learning*, 2:285–318, 1988.
- [Riv87] R. Rivest. Learning decision lists. *Machine Learning*, 2(3):229–246, 1987.
- [RS10] Alexander A Razborov and Alexander A Sherstov. The sign-rank of AC^0 . *SIAM Journal on Computing*, 39(5):1833–1855, 2010.
- [She09] A. Sherstov. The intersection of two halfspaces has high threshold degree. In *Proc. 50th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 343–362, 2009.
- [She10] A. Sherstov. Optimal bounds for sign-representing the intersection of two halfspaces by polynomials. In *Proc. 42nd ACM Symposium on Theory of Computing (STOC)*, pages 523–532, 2010.
- [She14] A. Sherstov. Breaking the Minsky-Papert barrier for constant-depth circuits. In *Proc. 46th ACM Symposium on Theory of Computing (STOC)*, pages 223–232, 2014.
- [She15] A. Sherstov. The power of asymmetry in constant-depth circuits. In *Proc. 56th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 431–450, 2015.

- [Val84] L. Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.
- [VC71] V. Vapnik and A. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory Probab. Appl.*, 16:264–280, 1971.

Participants

Dr. Boaz Barak

John A. Paulson School of Engineering
and Applied Sciences
Harvard University
33 Oxford Street
Cambridge, MA 02138
UNITED STATES

Prof. Dr. Markus Bläser

Fachbereich Informatik - FB 14
Universität des Saarlandes
Postfach 151150
66041 Saarbrücken
GERMANY

Prof. Dr. Johannes Blömer

Institut für Informatik
Fachgebiet Codes u. Kryptographie
Universität Paderborn
Fürstenallee 11
33102 Paderborn
GERMANY

Dr. Zvika Brakerski

Department of Computer Science
and Applied Mathematics
The Weizmann Institute of Science
P.O.Box 26
Rehovot 76100
ISRAEL

Dr. Mark Braverman

Department of Computer Science
Princeton University
35 Olden Street
Princeton, NJ 08544-5233
UNITED STATES

Prof. Dr. Christina Brzuska

Institut für IT-Sicherheitsanalyse
Technische Universität
Hamburg-Harburg
Am Irrgarten 3-9
21073 Hamburg
GERMANY

Prof. Dr. Peter Bürgisser

Institut für Mathematik
Technische Universität Berlin
Sekretariat MA 3-2
Straße des 17. Juni 136
10623 Berlin
GERMANY

Eshan Chattopadhyay

Department of Computer Science
University of Texas at Austin
Austin, TX 78712
UNITED STATES

Dr. Gil Cohen

Department of Computer Science
and Applied Mathematics
The Weizmann Institute of Science
P.O. Box 26
Rehovot 76100
ISRAEL

Prof. Dr. Irit Dinur

Department of Computer Science
and Applied Mathematics
The Weizmann Institute of Science
P.O.Box 26
Rehovot 76100
ISRAEL

Zeev Dvir

Department of Computer Science
Princeton University
35 Olden Street
Princeton, NJ 08544-5233
UNITED STATES

Klim Efremenko

Department of Computer Science
Tel-Aviv University
Ramat Aviv
Tel-Aviv 69978
ISRAEL

Dr. Michael A. Forbes

Department of Computer Science
Princeton University
35 Olden Street
Princeton, NJ 08540-5233
UNITED STATES

Prof. Dr. Oded Goldreich

Department of Computer Science
and Applied Mathematics
The Weizmann Institute of Science
P.O.Box 26
Rehovot 76100
ISRAEL

Prof. Dr. Shafi Goldwasser

MIT CSAIL
The Stata Center
32 Vassar Street
Cambridge MA 02139
UNITED STATES

Prof. Dr. Venkatesan Guruswami

Department of Computer Science
Carnegie Mellon University
GHC 7211
5000 Forbes Avenue
Pittsburgh PA 15213-3890
UNITED STATES

Christian Ikenmeyer

Department of Mathematics
Texas A & M University
College Station, TX 77843-3368
UNITED STATES

Dr. Neeraj Kayal

Microsoft Research India
VIGYAN Building
No. 9 Lavelle Road
Bangalore 560 001
INDIA

Dr. Hartmut Klauck

Centre for Quantum Technologies
National University of Singapore
Block S 15, 3 Science Drive 2
Singapore 117 543
SINGAPORE

Prof. Dr. Pascal Koiran

L I P
École Normale Supérieure de Lyon
46, Allée d'Italie
69364 Lyon Cedex 07
FRANCE

Prof. Dr. Gillat Kol

Institute for Advanced Study
School of Mathematics
Princeton University
Einstein Drive
Princeton NJ 08540
UNITED STATES

Dr. Or Meir

Department of Mathematics &
Computer Sciences
University of Haifa
Mount Carmel
Haifa 31905
ISRAEL

Dr. Raghu R. Meka

Department of Computer Science
University of California, Los Angeles
3732 H Boelter Hall
Los Angeles, CA 90095
UNITED STATES

Prof. Dr. Ryan O'Donnell

School of Computer Science
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213-3890
UNITED STATES

Prof. Dr. Rotem Oshman

School of Computer Science
Tel-Aviv University
Ramat-Aviv
Tel-Aviv 69978
ISRAEL

Dr. Prasad Raghavendra

Department of Computer Science
University of California, Berkeley
387 Soda Hall
Berkeley CA 94720-1776
UNITED STATES

Prof. Dr. Anup Rao

Dept. of Computer Science &
Engineering
University of Washington
P.O. Box 352350
Seattle WA 98195-2350
UNITED STATES

Prof. Dr. Alexander A. Razborov

Dept. of Mathematics & Computer
Science
The University of Chicago
Ryerson Hall
1100 East 58th Street
Chicago, IL 60637
UNITED STATES

Dr. Oded Regev

Courant Institute of
Mathematical Sciences
New York University
251, Mercer Street
New York, NY 10012-1110
UNITED STATES

Dr. Omer Reingold

Department of Computer Science
and Applied Mathematics
The Weizmann Institute of Science
P.O.Box 26
Rehovot 76100
ISRAEL

Benjamin Rossman

National Institute of Informatics
2-1-2 Hitotsubashi, Chiyoda-ku
Tokyo 101-8430
JAPAN

Ron Rothblum

Department of Computer Science
and Applied Mathematics
The Weizmann Institute of Science
P.O.Box 26
Rehovot 76100
ISRAEL

Prof. Dr. Michael Saks

Department of Mathematics
Rutgers University
Hill Center, Busch Campus
110 Frelinghuysen Road
Piscataway, NJ 08854-8019
UNITED STATES

Prof. Dr. Claus-Peter Schnorr

Institut für Mathematik
J.W.Goethe-Universität
60054 Frankfurt am Main
GERMANY

Prof. Dr. Rocco A. Servedio

Computer Science Department
Columbia University
MC 0401, Rm. 450
500 W. 120th Street
New York, NY 10027
UNITED STATES

Dr. Amir Shpilka

Department of Computer Science
Tel-Aviv University
Tel-Aviv 69978
ISRAEL

Prof. Dr. Christian Sohler

Fakultät für Informatik
Technische Universität Dortmund
Otto-Hahn-Strasse 14
44227 Dortmund
GERMANY

Dr. David Steurer

Department of Computer Science
Cornell University
319 Gates Hall
Ithaca NY 14850
UNITED STATES

Prof. Dr. Madhu Sudan

John A. Paulson School of Engineering
and Applied Sciences
Harvard University
33 Oxford Street
Cambridge, MA 02138
UNITED STATES

Avishay Tal

Department of Mathematics
The Weizmann Institute of Science
P.O. Box 26
Rehovot 76100
ISRAEL

Dr. Li-Yang Tan

Toyota Technological Institute at
Chicago
6045 S. Kenwood Avenue
Chicago, IL 60637
UNITED STATES

Dr. Sébastien Tavenas

Microsoft Research India
VIGYAN Building
No. 9 Lavelle Road
Bangalore 560 001
INDIA

Prof. Dr. Luca Trevisan

Department of Electrical Engineering
and Computer Science
University of California at Berkeley
387 Soda Hall
Berkeley, CA 94720
UNITED STATES

Prof. Dr. Jonathan Ullman

College of Computer and Information
Science
Northeastern University
Boston, MA 02115
UNITED STATES

Prof. Dr. Chris Umans

Department of Computer Science
MC 305-16, Annenberg 311
California Institute of Technology
1200 E. California Blvd.
Pasadena, CA 91125-5000
UNITED STATES

Dr. Salil Vadhan

School of Engineering & Applied
Sciences
Harvard University
Maxwell Dworkin 337
33 Oxford Street
Cambridge, MA 02138
UNITED STATES

Prof. Dr. Vinod Vaikuntanathan

MIT CSAIL
The Stata Center
32 Vassar Street
Cambridge MA 02139
UNITED STATES

Prof. Avi Wigderson

School of Mathematics
Institute for Advanced Study
1 Einstein Drive
Princeton, NJ 08540
UNITED STATES

Prof. Dr. Ryan Williams

Computer Science Department
Stanford University
353 Serra Mall
Stanford, CA 94305-2095
UNITED STATES

Prof. Dr. David Zuckerman

Department of Computer Science
University of Texas at Austin
2317 Speedway, Stop D9500
Austin TX 78712
UNITED STATES

