Recent Advances in Complexity Theory

Stephen Cook (University of Toronto),
Arvind Gupta (Simon Fraser University),
Russell Impagliazzo (University of California, San Diego),
Valentine Kabanets (Simon Fraser University),
Madhu Sudan (Massachusetts Institute of Technology),
Avi Wigderson (Institute for Advanced Study, Princeton)

August 26–31, 2006

Computational Complexity Theory is the field that studies the inherent costs of algorithms for solving mathematical problems. Its major goal is to identify the limits of what is efficiently computable in natural computational models. Computational complexity ranges from quantum computing to determining the minimum size of circuits that compute basic mathematical functions to the foundations of cryptography and security.

Computational complexity emerged from the combination of logic, combinatorics, information theory, and operations research. It coalesced around the central problem of "P versus NP" (one of the seven open problems of the Clay Institute). While this problem remains open, the field has grown both in scope and sophistication. Currently, some of the most active research areas in computational complexity are the following:

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

Many new developments in these areas were presented by the participants of the workshop. These new results will be described in the following sections of this report, grouped by topic. For each topic, we give a brief summary of the presented results, followed by the abstracts of the talks.

1 Computational Randomness

Computational randomness, or pseudorandomness, is the area concerned with explicit constructions of various "random-like" combinatorial objects. New constructions of one type of such objects, *randomness extractors*, have been reported by Anup Rao, Ronen Shaltiel, and David Zuckerman. The work by Shaltiel and his co-authors also yields a new explicit construction of a (bipartite)

Ramsey graph with better parameters than those of the construction by Frankl and Wilson. Chris Umans reported on a new construction of lossless condensers based on derandomized curve samplers.

Anup Rao, **Extractors for a Constant Number of Polynomially Small Min-Entropy Independent Sources**

We consider the problem of randomness extraction from independent sources. We construct an extractor that can extract from a constant number of independent sources of length n, each of which have min-entropy n^γ for an arbitrarily small constant $\gamma > 0$. Our extractor is obtained by composing seeded extractors in simple ways. We introduce a new technique to *condense* independent somewhere-random sources which looks like a useful way to manipulate independent sources. Our techniques are different from those used in recent work [BIW04, BKS$^+$05, Raz05, Bou05] for this problem in the sense that they do not rely on any results from additive number theory.

Using Bourgain's extractor [Bou05] as a black box, we obtain a new extractor for 2 independent block-sources with few blocks, even when the min-entropy is as small as polylog(n). We also show how to modify the 2 source disperser for linear min-entropy of Barak et al. [BKS$^+$05] and the 3 source extractor of Raz [Raz05] to get dispersers/extractors with exponentially small error and linear output length where previously both were constant.

In terms of Ramsey Hypergraphs, for every constant $1 > \gamma > 0$ our construction gives a family of explicit $O(1/\gamma)$-uniform hypergraphs on N vertices that avoid cliques and independent sets of size $2^{(\log N)^\gamma}$.

Ronen Shaltiel, **2-Source Dispersers for $n^{o(1)}$ Entropy, and Ramsey Graphs Beating the Frankl-Wilson Construction** (joint work with B. Barak, A. Rao, and A. Wigderson)

We present an explicit disperser for two independent sources on n bits, each of entropy $k = n^{o(1)}$. Put differently, setting $N = 2^n$ and $K = 2^k$, we construct explicit $N \times N$ Boolean matrices for which no $K \times K$ submatrix is monochromatic. Viewed as adjacency matrices of bipartite graphs, this gives an explicit construction of K-Ramsey *bipartite* graphs of size N.

This greatly improves the previous the previous bound of $k = o(n)$ of Barak, Kindler, Shaltiel, Sudakov and Wigderson [BKS$^+$05]. It also significantly improves the 25-year record of $k = \tilde{O}(\sqrt{n})$ on the very special case of Ramsey graphs, due to Frankl and Wilson [FW81].

The construction uses (besides "classical" extractor ideas) almost all of the machinery developed in the last couple of years for extraction from independent sources, including:

- Bourgain's extractor for 2 independent sources of some entropy rate $< 1/2$ [Bou05]
- Raz' extractor for 2 independent sources, one of which has any entropy rate $> 1/2$ [Raz05]
- Rao's extractor for 2 independent block-sources of entropy $n^{\Omega(1)}$ [Rao06]
- The "Challenge-Response" mechanism for detecting "entropy concentration" of [BKS$^+$05].

The main novelty comes in a bootstrap procedure which allows the Challenge-Response mechanism of [BKS$^+$05] to be used with sources of less and less entropy, using recursive calls to itself. Subtleties arise since the success of this mechanism depends on restricting the given sources, and so recursion constantly changes the original sources. These are resolved via a new construct, in between a disperser and an extractor, which behaves like an extractor on sufficiently large subsources of the given ones.

David Zuckerman, **Deterministic Extractors For Small Space Sources** (joint work with Jesse Kamp, Anup Rao, and Salil Vadhan)

We give explicit deterministic extractors for sources generated in small space, where we model space s sources on $\{0,1\}^n$ by width 2^s branching programs. We give extractors which extract almost all of the randomness from sources with constant entropy rate, when the space s is a small enough constant times n. We can extract from smaller min-entropies assuming efficient algorithms to find large primes. Previously, nothing was known for entropy rate less than $1/2$, even for space 0.

Our results are obtained by a reduction to a new class of sources that we call independent symbol sources, which generalize both the well-studied models of independent sources and symbol-fixing

sources. These sources consist of a string of n independent symbols over a d symbol alphabet with min-entropy k. We give deterministic extractors for such sources when k is as small as polylog(n), for small enough d.

Chris Umans, **Better lossless condensers through derandomized curve samplers** (joint work with Amnon Ta-Shma)

Lossless condensers are unbalanced expander graphs, with expansion close to optimal. Equivalently, they may be viewed as functions that use a short random seed to map a source on n bits to a source on many fewer bits while preserving all of the min-entropy. It is known how to build lossless condensers when the graphs are slightly unbalanced [CRVW02]. The highly unbalanced case is also important but the only known construction does not condense the source well. We give explicit constructions of lossless condensers with condensing close to optimal, and using near-optimal seed length.

Our main technical contribution is a randomness-efficient method for sampling F^D (where F is a field) with low-degree curves. This problem was addressed before [BSSVW03, MR06] but the solutions apply only to degree one curves, i.e., lines. Our technique is new and elegant. We use subsampling and obtain our curve samplers by composing a sequence of low-degree manifolds, starting with high-dimension, low-degree manifolds and proceeding through lower and lower dimension manifolds with (moderately) growing degrees, until we finish with dimension-one, low-degree manifolds, i.e., curves. The technique may be of independent interest.

2 Cryptography and Quantum Communication Complexity

Cryptography aims to develop protocols that will hide sensitive information from any unauthorized observer. One of the famous examples of such protocols is a "zero knowledge" protocol, which allows one one to convince an untrusting party of the truth of some statement without revealing any sensitive information about the statement. Salil Vadhan gave a survey and reported some new exciting results on zero-knowledge proofs. Adi Akavia presented the results showing that (essentially) one needs much more than P≠NP in order to build any cryptographic protocols. Scott Aaronson explained how one might be able to copy-protect quantum-computer software. Paul Valiant talked about a new notion of incrementally verifiable computation. Finally, Dmitry Gavinsky explained some differences between classical shared random string and quantum shared entanglements in the setting of communication complexity.

Salil Vadhan, **The Complexity of Zero Knowledge**

I will survey our efforts in the complexity-theoretic study of zero-knowledge proofs, where we have characterized the classes of problems having various types of zero-knowledge proofs, established general theorems about these classes, and minimized (indeed, often eliminated) complexity assumptions in the study of zero knowledge. In particular, I will discuss our most recent result, showing that all of NP has "statistical zero-knowledge arguments" under the (minimal) assumption that one-way functions exist, which resolves an open problem posed by Naor, Ostrovsky, Venkatesan, and Yung in 1992 [NOV$^+$].

The talk covers joint works with Minh Nguyen, Shien Jin Ong, and others, focusing on the papers [Vad04, NV06, NOV06].

Adi Akavia, **On Basing One-Way Function on NP-Hardness** (joint work with Oded Goldreich, Shafi Goldwasser and Dana Moshkovitz)

One-way functions are the cornerstone of modern cryptography. Informally speaking, one-way functions are functions that are easy to compute but are hard to invert (on the average case). There are several candidate functions, such as RSA or discrete-log, that are believed to be one-way, nonetheless, to date, no function was proved to be one-way. A puzzling question of fundamental nature is what are the minimal assumptions required for proving that a function is one-way. A necessary condition is that P does not equal NP (or more precisely, BPP does not equal NP, namely, that there is a problem in NP that cannot be solved by any probabilistic polynomial time algorithm). We ask whether this is also a sufficient condition. Namely, we ask whether there can be an efficient

reduction from NP (that is, from the task of deciding an NP-complete language on the worst case) to a one-way function (that is, to the task of inverting a one-way function on the average case).

We proved two results on the impossibility of reducing NP to a one-way function; both results hold under the (widely believed) complexity assumption that coNP is not contained in AM. 1. There cannot be a reduction (not even an adaptive reduction) from NP to a "size verifiable" one-way function; where we call f size-verifiable if, given y, the number of pre-image $|f^{-1}(y)|$ is efficiently computable, or, more generally, efficiently verifiable via an AM protocol. 2. There cannot be a non-adaptive reduction from NP to any one-way function (be it size-verifiable or not).

Our results improve on previously known negative results of [FF93, BT03] by (i) handling adaptive reductions (whereas previous works were essentially confined to non-adaptive reductions), and by (ii) relying on a seemingly weaker complexity assumption.

In the course of proving the above results, we designed a new constant round interactive protocol for proving upper bounds on the sizes of NP sets. We believe this protocol may be of independent interest.

Scott Aaronson, **Quantum Copy-Protection**

In the classical world, copy-protecting software is trivially impossible (not that that's stopped numerous companies from trying). But what if your computer program were a quantum state? In this talk, I'll present evidence that there exist quantum states that (1) can be used to evaluate some function f, but (2) can't be used to efficiently prepare more states with which to evaluate f. Indeed, in the black-box model, *any* function at all can be quantumly copy-protected, except in the degenerate case that one can efficiently learn the function by querying it. The proof of this result uses several new ideas that might be of interest on their own. These include an explicit construction of "d-wise independent quantum states," and a common generalization of the No-Cloning Theorem and the quantum search lower bound.

Paul Valiant, **Incrementally Verifiable Computation**

The probabilistically checkable proof (PCP) system enables proofs to be verified in time polylogarithmic in the length of a classical proof. Computationally sound proofs improve upon PCPs by additionally shortening the length of the transmitted proof to be polylogarithmic in the length of the classical proof. In this paper we explore the limits of such non-interactive proof systems. We present a proof system that in addition to the above properties allows proofs to be constructed in space polynomial in the space that it takes to classically accept the language, and time that is essentially linear in the time to classically accept. Our proof system is also *incremental*, a new notion that allows proofs of partial results to be composed together so that the length of the composition is no more than that of each part. Our construction relies on the hypothesized existence of a *proof of knowledge* system that reduces the length of classical proofs by a constant factor.

Dmitry Gavinsky, **On the role of shared entanglement**

Despite the apparent similarity between shared randomness and shared entanglement in the context of Communication Complexity, our understanding of the latter is not as good as of the former. In particular, there is no known "entanglement analogue" for the famous theorem by Newman [New91, NS96], saying that the number of shared random bits required for solving any communication problem can be at most logarithmic in the input length (i.e., using more than $O(log(n))$ shared random bits would not reduce the complexity of an optimal solution).

We prove that the same is not true for entanglement. We establish a wide range of tight (up to a logarithmic factor) entanglement vs. communication tradeoffs for relational problems. The "low-end" is: for any $t > 2$, reducing shared entanglement from $log^t(n)$ to $o(log^{t-1}(n))$ qubits can increase the communication required for solving a problem almost exponentially, from $O(log^t(n))$ to $\omega(\sqrt{n})$. The "high-end" is: for any $\epsilon > 0$, reducing shared entanglement from $n^{1-\epsilon}\log(n)$ to $o(n^{1-\epsilon})$ can increase the required communication from $O(n^{1-\epsilon}\log(n))$ to $\omega(n^{1-\epsilon/2})$.

3 Circuit complexity

Classical complexity theory aims to understand the power and limitations of efficient computation. One way to understand the limitations is to prove circuit lower bounds. While no strong circuit lower bounds are known for the general circuit model, there are some results for weaker models as well as there are some connections between circuit lower bounds and other areas of complexity, e.g., pseudorandomness. Eric Allender reported on new connections among arithmetic circuit complexity, real computation, and derandomization. Toni Pitassi described new constructions of small monotone circuits for computing the Majority function. Pierre McKenzie discussed lower bounds for a special case of branching programs. Dieter van Melkebeek presented results on time hierarchy for probabilistic complexity classes. Rahul Santhanam showed a new circuit lower bound for the "promise" version of complexity class MA. Amnon Ta-Shma discussed limitations of "black-box" reductions. Finally, Josh Buresh-Oppenheim explained how one could construct computationally hard Boolean functions via "hardness condensing".

Eric Allender, **Arithmetic Circuits, Real Numbers, and the Counting Hierarchy**

Arithmetic circuit complexity is the object of intense study in three different subareas of theoretical computer science:

1. **Derandomization.** The problem of determining if two arithmetic circuits compute the same function is known as $ACIT$ (arithmetic circuit identity testing). $ACIT$ is the canonical example of a problem in BPP that is not known to have a deterministic polynomial-time algorithm. Kabanets and Impagliazzo showed that the question of whether or not $ACIT$ is in P very tightly linked to the question of proving circuit size lower bounds [KI03].

2. **Computation over the Reals.** The Blum-Shub-Smale model of computation over the reals is an algebraic model that has received wide attention [BCS$^+$98].

3. **Valiant's Classes VP and VNP.** Valiant characterized the complexity of the permanent in two different ways. Viewed as a function mapping n-bit strings to binary encodings of Natural numbers, the permanent is complete for the class CP [Val79b]. Viewed as an n-variate polynomial, the permanent is complete for the class VNP [Val79a].

The general thrust of these three subareas has been in three different directions, and the questions addressed seem quite different from those addressed by work in the numerical analysis community, such as that surveyed by Demmel and Koev [DK03].

This talk will survey some recent work that ties all of these areas together in surprising ways. Most of the results that will be discussed can be found in [ABK$^+$05, Bur06], but I will also discuss some more recent progress.

Toniann Pitassi, **Monotone circuits for MAJORITY** (joint work with Shlomo Hoory and Avner Magen)

First I discuss what is currently known: the constructions by Ajtai, Komlos, and Szemeredy and by Valiant. Then I give our new results. We get smaller monotone circuits for MAJORITY; the size is roughly n^2 (rather than Valiant's $n^5.3$), while the depth is still $O(\log n)$. The circuit construction is also partially derandomized. The second phase which solves the promise problem uses belief propagation algorithm, and is derandomized and optimal; the first phase is still randomized.

Pierre McKenzie, **Incremental branching programs** (joint work with Anna Gál and Michal Koucký)

We propose a new model of restricted branching programs which we call *incremental branching programs*. We show that *syntactic* incremental branching programs capture previously studied structured models of computation for the problem GEN, namely marking machines [Coo74] and Poon's extension [Poo93] of jumping automata on graphs [CR80]. We then prove exponential size lower bounds for our syntactic incremental model, and for some other restricted branching program models as well. We further show that nondeterministic syntactic incremental branching programs are provably stronger than their deterministic counterpart when solving a natural NL-complete

GEN subproblem. It remains open if syntactic incremental branching programs are as powerful as unrestricted branching programs for GEN problems.

Dieter van Melkebeek, **Time Hierarchies for Semantic Models of Computation** (joint work with Konstantin Pervyshev)

A basic question in computational complexity asks whether somewhat more time allows us to solve strictly more decision problems on a given model of computation. Despite its fundamental nature, the question remains unanswered for many models of interest. Essentially, time hierarchies are known for every syntactic model of computation but open for everything else, where we call a model syntactic if there exists a computable enumeration consisting exactly of the machines in the model.

There has been significant progress in recent years, namely in establishing time hierarchies for non-syntactic models with small advice. In this talk, we survey these results and present a generic theorem that captures and strengthens all of them. We show that for virtually any semantic model of computation and for any rationals $1 \leq c \leq d$, there exists a language computable in time n^d with one bit of advice but not in time n^c with one bit of advice, where we call a model semantic if there exists a computable enumeration that contains all machines in the model but may also contain others.

Our result implies the first such hierarchy theorem for randomized machines with zero-sided error, quantum machines with one- or zero-sided error, unambiguous machines, symmetric alternation, Arthur-Merlin games of any signature, etc. Our argument also yields considerably simpler proofs of earlier hierarchy theorems with one bit of advice for randomized or quantum machines with two-sided error.

Rahul Santhanam, **Circuit Lower Bounds for Promise-MA**

We show that for each $k > 0$, $MA/1$ doesn't have circuits of size n^k. This implies the first super-linear circuit lower bounds for the promise versions of the classes $MA, AM, ZPP_{\|}^{NP}$ and BPP_{path}.

We extend our lower bound to the average-case setting, i.e., we show that $MA/1$ is not approximable by circuits of size n^k. Earlier, it was not even known if there is a language computable in Σ_2 with sublinear advice which is inapproximable by linear-size circuits.

Amnon Ta-Shma, **New connections between derandomization, worst-case complexity and average-case complexity** (joint work with Dan Gutfreund)

There has been a long line of research trying to explain our failure in proving worst-case to average-case reductions within NP [FF93, Vio03, BT03, AGGM06]. The bottom line of this research is, essentially, that under plausible assumptions black-box techniques cannot prove such results. A simple generalization of [BT03] shows:

Theorem 1 *Suppose that there is a language $L \in NP$ and a distribution D sampleable in time $n^{\log n}$ such that there is a* black-box *and non-adaptive reduction from solving SAT on the worst-case to solving L on the average with respect to D. Then every language in $coNP$ can be computed by a family of nondeterministic Boolean circuits of size $n^{polylog(n)}$.*

In particular, assuming no unexpected collapse occurs for the polynomial time hierarchy, the above worst-case to average-case reduction cannot be obtained via a black-box, non-adaptive reduction.

On the other hand, we show that the reduction of Gutfreund, Shaltiel and Ta-Shma [GSTS05] breaks the above lower bound. Specifically,

Theorem 2 *There exists a distribution D sampleable in time $n^{\log n}$, such that there is a non-adaptive reduction from solving SAT on the worst-case to solving SAT on the average with respect to D.*

In particular, the [GSTS05] reduction bypasses the black-box limitation imposed by Theorem 1 (if the above collapse does not happen), and indeed the [GSTS05] reduction is non black-box.

As it turns out, the [GSTS05] reduction is black-box in the reduction function (mapping an algorithm good on average to a worst-case algorithm), and this reduction is simply the search to decision reduction. However, it is not black-box in the proof. Instead, the proof of correctness only shows that any *efficient* algorithm to the average-case problem, is mapped to an efficient algorithm

for the worst-case problem. We call such reductions *class black-box*. We believe such reductions are often as useful as black-box reductions, and yet, our work demonstrates that they can break black-box limitations.

Finally, we are now in a position where there are no negative results to stop us. How far can we go? Given the techniques of [GSTS05] a natural goal is to answer the following Open Question: Does $NP \not\subseteq BPP$ imply the existence of a language in $QNP = NTIME(n^{O(\log n)})$ that is hard on average for BPP?

Using the [IL90] reduction we show such a result, but only using some weak, *unproven* derandomization assumption. Resolving this Open Question without any assumptions remains a challenge.

Joshua Buresh-Oppenheim, **Making Hard Problems Harder** (joint work with Rahul Santhanam)

Proving circuit lower bounds for explicit Boolean functions is one of the most fundamental and challenging questions in theoretical computer science. We consider an approach to this question which aims to improve hardness rather than give a direct proof of hardness. We define "hardness extractors," which are procedures taking in a Boolean function as input together with a relatively small advice string, and outputting a Boolean function on a smaller number number of bits which has greater hardness when measured in terms of its input length. We show a construction of a hardness extractor with linear advice extracting deterministic hardness from non-deterministic hardness. As a consequence, we obtain a "gap" theorem for E with linear advice: if E with linear advice requires exponential non-uniform space, then E with linear advice requires non-uniform space $2^n/n$.

We also define a natural class of "relativizing" hardness extractors and give lower bounds on the advice required by such extractors. This indicates that hardness extraction without advice and extraction of deterministic hardness from deterministic hardness in general will require novel techniques. On the other hand, we show two special cases where we can extract from deterministic hardness without advice: biased functions and functions that are hard on average.

4 Error-correcting codes and PCPs

Error-correcting codes play a major role in modern complexity theory. Many important results in complexity (e.g., the famous PCP theorem) are best viewed as constructions of special error-correcting codes. Once this connection between coding and complexity theory is realized, both areas enjoy mutual benefits by using ideas and insights from the other area. Venkatesan Guruswami gave an explicit construction of list-decodable codes with optimal rate. Ran Raz presented a construction of a very efficient low-degree test (useful for PCPs). Ragesh Jaiswal talked about some error-correcting codes directly inspired by complexity-theoretic questions. Eli Ben-Sasson discussed some limitations of list-decoding. Finally, Oded Regev showed an improved hardness of approximation result for the problem of finding a shortest vector in a lattice.

Venkatesan Guruswami, **List Decoding with Optimal Rate: Folded Reed-Solomon Codes** (joint work with Atri Rudra)

Suppose you want to communicate a message of k packets on a noisy communication channel. So you judiciously encode it as a redundant collection of $n = ck$ packets and transmit these. What is the fewest number of correct packets one needs to receive in order to have any hope of recovering the message?

Well, clearly one needs at least k correct packets. In this talk, I will describe an encoding scheme that attains this information-theoretic limit: for any desired $\epsilon > 0$, it enables recovery of the message as long as at least $k(1 + \epsilon)$ packets are received intact. The location of the correct packets and the errors on the remaining packets can be picked adversarially by the channel.

This achieves the optimal trade-off (called "capacity") between redundancy and error-resilience for a malicious noise model where the channel can corrupt the transmitted symbols arbitrarily subject to a bound on the total number of errors. These results are obtained in an error-recovery model called list decoding. The talk will introduce and motivate the problem of list decoding, and then give a peek into the algebraic ideas and constructions that lead to the above result.

Ran Raz, **Sub-Constant Error Low Degree Test of Almost Linear Size** (joint work with Dana Moshkovitz)

Given a function $f : F^m \to F$ over a finite field F, a *low degree tester* tests its agreement with an m-variate polynomial of total degree at most d over F. The tester is usually given access to an oracle A providing the *supposed* restrictions of f to affine subspaces of constant dimension (e.g., lines, planes, etc.). The tester makes very few (probabilistic) queries to f and to A (say, one query to f and one query to A), and decides whether to accept or reject based on the replies.

We wish to minimize two parameters of a tester: its *error* and its *size*. The *error* bounds the probability that the tester accepts although the function is far from a low degree polynomial. The *size* is the number of bits required to write the oracle replies on all possible tester's queries.

Low degree testing is a central ingredient in most constructions of probabilistically checkable proofs (PCPs) and locally testable codes (LTCs). The error of the low degree tester is related to the soundness of the PCP and its size is related to the size of the PCP (or the length of the LTC).

We design and analyze new low degree testers that have both *sub-constant error* $o(1)$ and *almost-linear size* $n^{1+o(1)}$ (where $n = |F|^m$). Previous constructions of *sub-constant error* testers had *polynomial size* [AS03, RS97]. These testers enabled the construction of PCPs with *sub-constant soundness*, but *polynomial size* [AS03, RS97, DFK$^+$99]. Previous constructions of *almost-linear size* testers obtained only *constant error* [GS02, BSSVW03]. These testers were used to construct *almost-linear size* LTCs and *almost-linear size* PCPs with *constant soundness* [GS02, BSSVW03, BSGH$^+$04, BSS05, Din06].

Ragesh Jaiswal, **Approximately list-decoding direct product codes and uniform hardness amplification** (joint work with Russell Impagliazzo and Valentine Kabanets)

We consider the problem of locally list-decoding *direct product* codes. For a parameter k, the k-wise direct product encoding of an N-bit message msg is an N^k-length string over the alphabet $\{0,1\}^k$ indexed by k-tuples $(i_1, \ldots, i_k) \in \{1, \ldots, N\}^k$ so that the symbol at position $(i_1, \ldots, i_k)$ of the codeword is $msg(i_1) \ldots msg(i_k)$. Such codes arise naturally in the context of hardness amplification of Boolean functions via Yao's Direct Product Lemma (and closely related Yao's XOR Lemma), where typically $k \ll N$ (e.g., $k = poly \log N$).

We describe an efficient randomized algorithm for approximate local list-decoding of direct product codes. Given oracle access to a word which agrees with a k-wise direct product encoding of some message msg in at least ϵ fraction of positions, our algorithm outputs a list of $poly(1/\epsilon)$ Boolean circuits computing N-bit strings (viewed as truth tables of $\log N$-variable Boolean functions) such that at least one of them agrees with msg in at least $1-\delta$ fraction of positions, for $\delta = O(\frac{\log(1/\epsilon)}{k} + k^{-0.1})$, provided that $\epsilon = \Omega(poly(1/k))$; the running time of the algorithm is polynomial in $\log N$ and $1/\epsilon$. When $\epsilon > e^{-k^\alpha}$ for a certain constant $\alpha > 0$, we get a randomized approximate list-decoding algorithm that runs in time quasipolynomial in $1/\epsilon$ (i.e., $(1/\epsilon)^{poly \log 1/\epsilon}$).

By concatenating the k-wise direct product codes with Hadamard codes, we obtain locally list-decodable codes over the binary alphabet, which can be efficiently approximately list-decoded from fewer than $1/2 - \epsilon$ fraction of corruptions as long as $\epsilon = \Omega(poly(1/k))$. As an immediate application, we get *uniform* hardness amplification for $P^{NP}_{\|}$, the class of languages reducible to NP through one round of parallel oracle queries: If there is a language in $P^{NP}_{\|}$ that cannot be decided by any BPP algorithm on more that $1 - 1/n^{\Omega(1)}$ fraction of inputs, then there is another language in $P^{NP}_{\|}$ that cannot be decided by any BPP algorithm on more that $1/2 + 1/n^{\omega(1)}$ fraction of inputs.

Eli Ben-Sasson, **Subspace Polynomials and List Decoding of Reed-Solomon Codes** (joint work with Swastik Kopparty and Jaikumar Radhakrishnan)

We show combinatorial limitations on efficient list decoding of Reed-Solomon codes beyond the Johnson and Guruswami-Sudan bounds. In particular, we show that for arbitrarily large fields $F_N, |F_N| = N$, for any $\delta \in (0,1)$, and $K = N^\delta$:

- **Existence:** there exists a received word that agrees with a super-polynomial number of distinct degree K polynomials on approximately $N^{\sqrt{\delta}}$ points each;
- **Explicit:** there exists a polynomial time constructible received word that agrees with a super-polynomial number of distinct degree K polynomials, on approximately $2^{\sqrt{\log N}} K$ points each.

In both cases, our results improve upon the previous state of the art, which was about N^δ/δ for the existence case and about $2N^\delta$ for the explicit one. Furthermore, for δ close to 1 our bound approaches the Guruswami-Sudan bound (which is $\sqrt{NK}$) and implies limitations on extending their efficient RS list decoding algorithm to larger decoding radius.

Our proof method is surprisingly simple. We work with polynomials that vanish on subspaces of an extension field viewed as a vector space over the base field. These subspace polynomials are a subclass of linearized polynomials that were studied by Ore in the 1930s and by coding theorists. For us their main attraction is their sparsity and abundance of roots, virtues that recently won them pivotal roles in probabilistically checkable proofs of proximity and sub-linear proof verification.

Oded Regev, **Tensor-based Hardness of the Shortest Vector Problem to within Almost Polynomial Factors** (joint work with Ishay Haviv)

We show that unless $NP \subseteq RTIME(2^{poly(\log n)})$, the Shortest Vector Problem (SVP) on n-dimensional lattices in the ℓ_p norm ($1 \leq p < \infty$) is hard to approximate in polynomial-time to within a factor of $2^{(\log n)^{1-\epsilon}}$ for any $\epsilon > 0$. This improves the previous best factor of $2^{(\log n)^{1/2-\epsilon}}$ under the same complexity assumption due to Khot [Kho05]. Under the stronger assumption $NP \not\subseteq RSUBEXP$, we obtain a hardness factor of $n^{c/\log\log n}$ for some $c > 0$. Our proof starts with SVP instances from [Kho05] that are hard to approximate to within some constant. To boost the hardness factor we simply apply the standard tensor product of lattices. The main novel part is in the analysis, where we show that the lattices of [Kho05] behave nicely under tensorization. At the heart of the analysis is a certain matrix inequality which was first used in the context of lattices by de Shalit [deS06].

5 Computational Learning

The area of computational learning is concerned with the problems of learning a function, given a number of samples drawn according to some distribution on the inputs to the function. Ryan O'Donnell explained how to test, using very few samples, whether a given function is a Boolean halfspace. Adam Klivans presented results on learning halfspaces. Finally, Scott Aaronson showed how to learn quantum states.

Ryan O'Donnell, **Testing Halfspaces** (joint work with Kevin Matulef (MIT), Ronitt Rubinfeld (MIT), and Rocco Servedio (Columbia))

In this talk we describe work showing that the class of Boolean halfspaces – -i.e., functions $f : \{-1,1\}^n \to \{-1,1\}$ representable as $f(x) = \mathrm{sgn}(c_1x_1 + \cdots + c_nx_n - \theta)$ — has a property testing algorithm making only $\mathrm{poly}(1/\epsilon)$ queries.

Adam Klivans, **Agnostically Learning Halfspaces** (joint work with A. Kalai, Y. Mansour, and R. Servedio)

We give the first algorithm that efficiently learns halfspaces (under distributional assumptions) in the notoriously difficult agnostic framework of Kearns, Schapire, and Sellie. In this model, a learner is given arbitrarily labeled examples from a fixed distribution and must output a hypothesis competitive with the optimal halfspace hypothesis.

Our algorithm constructs a hypothesis whose error rate on future examples is within an additive ϵ of the optimal halfspace in time $poly(n)$ for any constant $\epsilon > 0$ under the uniform distribution over $\{0,1\}^n$ or the unit sphere in R^n, as well as under any log-concave distribution over R^n. It also agnostically learns Boolean disjunctions in time $2^{\tilde{O}(\sqrt{n})}$ with respect to *any* distribution. The new algorithm, essentially L_1 polynomial regression, is a noise-tolerant arbitrary-distribution generalization of the "low-degree" Fourier algorithm of Linial, Mansour, and Nisan. Our Fourier-type algorithm over the unit sphere makes use of approximation properties of various classes of orthogonal polynomials.

Scott Aaronson, **The Learnability of Quantum States**

Traditional quantum state tomography requires a number of measurements that grows exponentially with the number of qubits n. But using ideas from computational learning theory, we show

that "for most practical purposes" one can learn a state using a number of measurements that grows only linearly with n. Besides possible implications for experimental physics, our learning theorem has two applications to quantum computing: first, a new simulation of quantum one-way protocols, and second, the use of trusted classical advice to verify untrusted quantum advice.

6 Research Emerging from Workshop

The goal of the workshop was to bring some of the best researchers in the area of computational complexity to discuss the current state of the art in the area, and point out further directions of research. The workshop has been very successful from the point of view of many fruitful interactions among various groups of the workshop participants. Some ideas first discussed during the workshop already found their way into research papers. One example is the following paper *Extractors and Condensers from Univariate Polynomials* by Venkatesan Guruswami, Chris Umans, and Salil Vadhan, which was posted on *Electronic Colloquium on Computational Complexity*, October 2006 [GUV06]. Below is the description of this work provided by the authors.

Context and Genesis: There is a long body of work in theoretical computer science on constructions of both *randomness extractors* — functions that extract almost-uniform bits from sources of biased and correlated bits, and *expander graphs* — graphs that are sparse but highly connected. These two kinds of objects are closely related, and both have a wide variety of applications in theoretical computer science. The paper [GUV06] resulting from the BIRS workshop presents new constructions of both extractors and expanders (described as 'lossless condensers') that significantly improve previous work, while also being simpler and more direct.

The work began with a conversation between the participants Guruswami, Umans, and Vadhan in the Corbett Lounge after Guruswami had presented his new work [GR06] on capacity-achieving error-correcting codes. Indeed, for a few years, it has been known that randomness extractors can be viewed as a generalization of "list-decodable" error-correcting codes. Because of this connection and similarities between the Guruswami–Rudra codes and a previous extractor construction of Shaltiel and Umans [SU01], it seemed natural to explore whether the ideas underlying the Guruswami–Rudra codes and their predecessors [PV05] could be applied to construct better extractors and expander graphs. The participants pursued this idea via email after the workshop, and within a few weeks, the new results had emerged.

Abstract: We give new constructions of randomness extractors and lossless condensers that are optimal to within constant factors in both the seed length and the output length. For extractors, this matches the parameters of the current best known construction [LRV$^+$03]; for lossless condensers, the previous best constructions achieved optimality to within a constant factor in one parameter only at the expense of a polynomial loss in the other.

Our constructions are based on the Parvaresh-Vardy codes [PV05], and our proof technique is inspired by the list-decoding algorithm for those codes. The main object we construct is a condenser that loses *only* the entropy of its seed plus one bit, while condensing to entropy rate $1 - \alpha$ for any desired constant $\alpha > 0$. This construction is simple to describe, and has a short and completely self-contained analysis. Our other results only require, in addition, standard uses of randomness-efficient hash functions (to obtain a lossless condenser) or expander walks (to obtain an extractor).

Our techniques also show for the first time that a natural construction based on univariate polynomials (i.e., Reed-Solomon codes) yields a condenser that retains a $1 - \alpha$ fraction of the source min-entropy, for any desired constant $\alpha > 0$, while condensing to constant entropy rate and using a seed length that is optimal to within constant factors.

References

[ABK+05] E. Allender, P. Bürgisser, Johann Kjeldgaard-Pedersen, and Peter Bro Miltersen. On the complexity of numerical analysis. In *Proc. 21st Ann. IEEE Conf. on Computational Complexity (CCC '06)*, pages 331–339, 2006.

[AGGM06] A. Akavia, O. Goldreich, S. Goldwasser, and D. Moshkovitz. On basing one-way functions on NP-hardness. In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing*, pages 701–710, 2006.

[AS03] S. Arora and M. Sudan. Improved low-degree testing and its applications. *Combinatorica*, 23(3):365–426, 2003.

[BCS+98] L. Blum, F. Cucker, M. Shub, and S. Smale. *Complexity and Real Computation*. Springer, 1998.

[Bur06] P. Bürgisser. On defining integers in the counting hierarchy and proving lower bounds in algebraic complexity. Technical Report TR06-113, Electronic Colloquium on Computational Complexity, 2006.

[BIW04] B. Barak, R. Impagliazzo, and A. Wigderson. Extracting randomness using few independent sources. In *Proceedings of the Forty-Fifth Annual IEEE Symposium on Foundations of Computer Science*, 2004.

[BKS+05] B. Barak, G. Kindler, R. Shaltiel, B. Sudakov, and A. Wigderson. Simulating independence: New constructions of condensers, Ramsey graphs, dispersers, and extractors. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 1–10, 2005.

[Bou05] J. Bourgain. More on the sum-product phenomenon in prime fields and its applications. *International Journal of Number Theory*, 1:1–32, 2005.

[BSGH+04] E. Ben-Sasson, O. Goldreich, P. Harsha, M. Sudan, and S. Vadhan. Robust pcps of proximity, shorter pcps and applications to coding. In *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing*, pages 1–10, 2004.

[BSS05] E. Ben-Sasson and M. Sudan. Simple pcps with poly-log rate and query complexity. In *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*, pages 266–275, 2005.

[BSSVW03] E. Ben-Sasson, M. Sudan, S. Vadhan, and A. Wigderson. Randomness-efficient low degree tests and short pcps via epsilon-biased sets. In *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, pages 612–621, 2003.

[BT03] A. Bogdanov and L. Trevisan. On worst-case to average-case reductions for NP problems. In *Proceedings of the Forty-Fourth Annual IEEE Symposium on Foundations of Computer Science*, pages 308–317, 2003.

[Coo74] S.A. Cook. An observation on time-storage trade-off. *Journal of Computer and System Sciences*, 9(3):308–316, 1974.

[CR80] S.A. Cook and C.W. Rackoff. Space lower bounds for maze threadability on restricted machines. *SIAM Journal on Computing*, 9:636–652, 1980.

[CRVW02] M. Capalbo, O. Reingold, S. Vadhan, and A. Wigderson. Randomness conductors and constant-degree expansion beyond the degree/2 barrier. In *STOC*, pages 659–668, 2002.

[deS06] E. de Shalit. On tensor products of semistable lattices. Preprint, 2006.

[DFK+99] I. Dinur, E. Fischer, G. Kindler, R. Raz, and S. Safra. Pcp characterizations of np: Towards a polynomially-small error-probability. In *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing*, pages 29–40, 1999.

[DK03] J. Demmel and P. Koev. Accurate and efficient algorithms for floating point computation. In *Proceedings of the 2003 International Congress of Industrial and Applied Mathematics*, 2003.

[Din06] I. Dinur. The PCP theorem by gap amplification. In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing*, pages 241–250, 2006.

[FF93] J. Feigenbaum and L. Fortnow. Random-self-reducibility of complete sets. *SIAM Journal on Computing*, 22(5):994–1005, 1993.

[FW81] Peter Frankl and Richard M. Wilson. Intersection theorems with geometric consequences. *Combinatorica*, 1(4):357–368, 1981.

[GR06] V. Guruswami and A. Rudra. Explicit capacity-achieving list-decodable codes. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 1–10, 2006.

[GS02] O. Goldreich and M. Sudan. Locally testable codes and pcps of almost-linear length. In *Proceedings of the Forty-Third Annual IEEE Symposium on Foundations of Computer Science*, pages 13–22, 2002.

[GSTS05] D. Gutfreund, R. Shaltiel, and A. Ta-Shma. If np languages are hard on the worst-case then it is easy to find their hard instances. In *Proceedings of the Twentieth Annual IEEE Conference on Computational Complexity*, pages 243–257, 2005.

[GUV06] V. Guruswami, C. Umans, and S. Vadhan. Extractors and Condensers from Univariate Polynomials. Technical Report TR06-134, Electronic Colloquium on Computational Complexity, October 2006.

[IL90] R. Impagliazzo and L. Levin. No better ways to generate hard np instances than picking uniformly at random. In *Proceedings of the Thirty-First Annual IEEE Symposium on Foundations of Computer Science*, pages 812–821, 1990.

[Kho05] S. Khot. Hardness of approximating the shortest vector problem in lattices. *Journal of the ACM*, 52(5):789–808, Sept. 2005. Preliminary version in FOCS 2004.

[KI03] V. Kabanets and R. Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. In *Proc. ACM Symp. Theory Comp.*, pages 355–364, 2003.

[LRV+03] C.-J. Lu, O. Reingold, S. Vadhan, and A. Wigderson. Extractors: optimal up to constant factors. In *Proceedings of the 35th ACM Symposium on Theory of Computing (STOC '03)*, pages 602–611. ACM, 2003.

[MR06] D. Moshkovitz and R. Raz. Sub-constant error low degree test of almost linear size. In *STOC*, 2006.

[New91] I. Newman. Private vs. common random bits in communication complexity. *Information Processing Letters*, 39(2):67–71, 1991.

[NOV+] M. Naor, R. Ostrovsky, R. Venkatesan, and M. Yung. Perfect Zero-Knowledge Arguments for NP Using Any One-Way Permutation. *J. Cryptology*, 11(2):87–108, 1998. Preliminary version in *CRYPTO '92*.

[NS96] I. Newman and M. Szegedy. Public vs. Private Coin Flips in One Round Communication Games. In *Proceedings of the 28th Symposium on Theory of Computing*, pages 561–570, 1996.

[NV06] M. Nguyen and S. Vadhan. Zero Knowledge with Efficient Provers. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC '06)*, pages 287–295, 21–23 May 2006.

[NOV06] M.-H. Nguyen, S. J. Ong, and S. Vadhan. Statistical Zero-Knowledge Arguments for NP from Any One-Way Function. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS '06)*, pages 3–13, Berkeley, CA, 22–24 October 2006. Full version posted as *ECCC* TR06-075.

[Poo93] C.K. Poon. Space bounds for graph connectivity problems on node-named JAGs and node-ordered JAGs. In em Proceedings of the 34th Annual IEEE Symposium on Foundations of Computer Science, pages 218–227, 1993.

[PV05] F. Parvaresh and A. Vardy. Correcting Errors Beyond the Guruswami-Sudan Radius in Polynomial Time. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 285–294, 2005.

[RS97] R. Raz and S. Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, pages 475–484, 1997.

[Rao06] A. Rao. Extractors for a constant number of polynomial min-entropy independent sources. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, 2006.

[Raz05] R. Raz. Extractors with weak random seeds. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 11–20, 2005.

[SU01] R. Shaltiel and C. Umans. Simple extractors for all min-entropies and a new pseudorandom generator. *Journal of the ACM*, 52(2):172–216, 2005. Conference version appeared in FOCS 2001.

[Vad04] S. P. Vadhan. An Unconditional Study of Computational Zero Knowledge. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS '04)*, pages 176–185, Rome, Italy, 17–19 October 2004. Full version accepted to *SIAM J. Computing* Special Issue on Randomness & Complexity.

[Val79a] L. Valiant. Completeness classes in algebra. In *Proc. ACM Symp. Theory Comp.*, pages 249–261, 1979.

[Val79b] L. Valiant. The complexity of computing the Permanent. *Theoret. Comp. Sci.*, 8:189–201, 1979.

[Vio03] E. Viola. Hardness vs. randomness within alternating time. In *Proceedings of the Eighteenth Annual IEEE Conference on Computational Complexity*, pages 53–62, 2003.