08381 Abstracts Collection
Computational Complexity of Discrete Problems
— Dagstuhl Seminar —

Peter Bro Miltersen[1], Rüdiger Reischuk[2], Georg Schnitger[3] and Dieter van Melkebeek[4]

[1] University of Aarhus, Denmark
bromille@daimi.au.dk
[2] Universität Lübeck, Germany
reischuk@tcs.uni-luebeck.de
[3] Universität Frankfurt, Germany
georg@thi.informatik.uni-frankfurt.de
[4] University of Wisconsin - Madison, USA
dieter@cs.wisc.edu

Abstract. From the 14th of September to the 19th of September, the Dagstuhl Seminar 08381 "Computational Complexity of Discrete Problems" was held in the International Conference and Research Center (IBFI), Schloss Dagstuhl. During the seminar, several participants presented their current research, and ongoing work as well as open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this report. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords:
computational complexity, discrete problems, Turing machines, circuits, proof complexity, pseudorandomness, derandomization, cryptography, computational learning, communication complexity, query complexity, hardness of approximation

08381 Executive Summary - Computational Complexity of Discrete Problems We briefly describe the state of the art concerning the complexity of discrete problems. After describing the formal organization of the seminar we describe the different topics that have been discussed and mention some of the major achievements. The summary closes with an outlook on the future development of discrete computational complexity.

Joint work of: Miltersen, Peter Bro; Reischuk, Rüdiger; Schnitger, Georg; van Melkebeek, Dieter

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2008/1778

Dagstuhl Seminar Proceedings 08381
Computational Complexity of Discrete Problems
http://drops.dagstuhl.de/opus/volltexte/2008/1779

Univariate solutions for low-degree multivariate problems

Eli Ben-Sasson (Technion, Haifa, Israel)

We present three results regarding problems related to multivariate polynomials of low degree over a p-element finite field F:

1. Explicit construction of a seedless disperser for affine subspaces of F^n of dimension greater than $2n/5 + 10$.
2. Explicit construction of a family of functions f mapping F^n to F that have exponentially small correlation with degree-d multivariate polynomials over F.
3. An improved "global" list-decoding algorithm for Reed-Muller codes.

The common theme underlying these results is their method of proof. We translate the problems into questions about univariate polynomials over F_{p^n}. Functions represented by low-degree multivariate polynomials in $F[x_1, \ldots, x_n]$, when viewed as univariate polynomials over F_{p^n}, have a very special structure. We use various properties of this special structure to obtain the above mentioned results.

Keywords: Affine dispersers, derandomization, list decoding, reed-muller codes

Joint work of: Ben-Sasson, Eli; Kopparty, Swastik

Randomness efficient zero testing of blackbox polynomials

Markus Bläser (Saarland University, Germany)

Our main result is an efficient construction of a hitting set generator against the class of polynomials of degree d_i in the i-th variable.

The seed length of this generator is $\log D + O(\log^{1/2} D)$. Here, $\log D = \sum_i \log(d_i + 1)$ is a lower bound on the seed length of any hitting set generator against this class.

Our construction is the first to achieve asymptotically optimal seed length for every choice of the parameters d_i.

In fact, we present a nearly linear time construction with this asymptotic guarantee. Furthermore, our results extend to classes of polynomials parameterized by upper bounds on the number of nonzero terms in each variable. Underlying our constructions is a general and novel framework that exploits the product structure common to the classes of polynomials we consider. This framework allows us to obtain efficient and asymptotically optimal hitting set generators from primitives that need not be optimal or efficient by themselves.

As our main corollary, we obtain the first blackbox polynomial identity tests with an asymptotically optimal randomness consumption.

Keywords: Randomized algorithms, polynomial identity testing

Joint work of: Bläser, Markus; Hardt, Moritz; Steurer, David

On the OBDD complexity of the most significant bit of integer multiplication

Beate Bollig (Technische Universität Dortmund, Germany)

Integer multiplication as one of the basic arithmetic functions has been in the focus of several complexity theoretical investigations.

Ordered binary decision diagrams (OBDDs) are one of the most common dynamic data structures for boolean functions.

Analyzing the limits of symbolic graph algorithms for the reachability problem Sawitzki (2006) has presented the first exponential lower bound on the pi-OBDD size for the most significant bit of integer multiplication according to one predefined variable order pi. Since the choice of the variable order is a main issue to obtain OBDDs of small size, the investigation is continued. As a result a new upper bound method and the first non-trivial upper bound on the size of OBDDs according to an arbitrary variable order is presented. Furthermore, Sawitzki's lower bound is improved.

Moreover, it is shown that the OBDD complexity of the most significant bit is exponential answering an open question posed by Wegener (2000).

Keywords: Computational complexity

Joint work of: Bollig, Beate; Klump, Jochen

See also: Bollig, B. (2008). On the OBDD complexity of the most significant bit of integer multiplication. Proc. of TAMC 2008, LNCS 4978, 306-317.
Bollig, B. and Klump, J. (2008). New results on the most significant bit of integer multiplication. Accepted for Proc. of ISAAC 2008.

The complexity of simulating Brownian Motion

Mark Braverman (Microsoft Research NE, USA)

We analyze the complexity of the Walk on Spheres algorithm for simulating Brownian Motion in a domain $\Omega \subseteq I\!R^d$. The algorithm, which was first proposed in the 1950s, produces samples from the hitting probability distribution of the Brownian Motion process on boundary of Ω within an error of ϵ. The algorithm is used as a building block for solving a variety of differential equations, including the Dirichlet Problem.

The WoS algorithm simulates a BM starting at a point $X_0 = x$ in a given bounded domain Ω until it gets ϵ-close to the boundary of Ω. At every step, the algorithm measures the distance d_k from its current position X_k to the boundary

of Ω and jumps a distance of $d_k/2$ in a uniformly random direction from X_k to obtain X_{k+1}. The algorithm terminates when it reaches X_n that is ϵ-close to the boundary of Ω.

It is not hard to see that the algorithm requires at least $\Omega(\log 1/\epsilon)$ steps to converge. Only partial results with respect to upper bounds existed. In 1959 M. Motoo established an $O(\log 1/\epsilon)$ bound on the running time for convex domains. The results were later generalized for a wider, but still very restricted, class of planar and 3-dimensional domains by G.A. Mikhailov (1979). In our earlier work (2007), we established an upper bound of $O(\log^2 1/\epsilon)$ on the rate of convergence of WoS for arbitrary planar domains.

We introduce subharmonic energy functions to obtain very general upper bounds on the convergence of the algorithm. Special instances of the upper bounds yield the following results for bounded domains Ω:

- if Ω is a planar domain with connected exterior, the WoS converges in $O(\log 1/\epsilon)$ steps.
- if Ω is a domain in $I\!R^3$ with connected exterior, the WoS converges in $O(\log^2 1/\epsilon)$ steps.
- for $d > 2$, if Ω is a domain in $I\!R^d$, the WoS converges in $O((1/\epsilon)^{2-4/d})$ steps.
- for $d > 3$ if Ω is a domain in $I\!R^d$ with connected exterior, the WoS converges in $O((1/\epsilon)^{2-4/(d-1)})$ steps.
- for any d if Ω is a domain in $I\!R^d$ bounded by a smooth surface, the WoS converges in $O(\log 1/\epsilon)$ steps.

We also demonstrate that the bounds are tight, i.e. we construct a domain from each class for which the upper bound is exact. Our results give the optimal upper bound of $O(\log 1/\epsilon)$ in many cases for which only a bound polynomial in $1/\epsilon$ was previously known.

Keywords: Random walk, energy function

Tight bounds for blind search on the integers

Martin Dietzfelbinger (TU Ilmenau, Germany)

We analyze a simple random process in which a token is moved in the interval $A = \{0, 1, \ldots, n\}$. Fix a probability distribution μ over $\{1, \ldots, n\}$. Initially, the token is placed in a random position in A. In round t, a random value d is chosen according to μ .

If the token is in position $a \geq d$, then it is moved to position $a - d$. Otherwise it stays put. Let T be the number of rounds until the token reaches position 0. We show tight bounds for the expectation of T for the optimal distribution μ . More precisely, we show that $\min_\mu(E_\mu(T)) = \Omega((\log n)^2)$. For the proof, a novel potential function argument is introduced. The research is motivated by the problem of approximating the minimum of a continuous function over $[0, 1]$ with a "blind" optimization strategy.

Keywords: Randomized search heuristic, black box optimization, lower bound

Joint work of: Dietzfelbinger, Martin; Rowe, Jonathan; Wegener, Ingo; Woelfel, Philipp;

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1348

See also: Martin Dietzfelbinger, Jonathan E. Rowe, Ingo Wegener, Philipp Woelfel: Tight Bounds for Blind Search on the Integers. STACS 2008: 241-252

Program Equilibria and Discounted Computation Time

Lance Fortnow (Northwestern - Evanston, USA)

Consider the Prisoner's Dilemma game played for one round. The only Nash Equilibrium is for both players to defect even though they would both gain if they both cooperated. Tennenholtz (GEB 2004) consider program equilibria where each player provides a program that can look at the other player's program. Tennenholtz shows that cooperation is now possible in this model.

Tennenholtz and others use a limited computation model that, for example, lacks universal simulation. We look at a new model that allows for arbitrary Turing machines but discounts the payoffs based on running time. We recover the results to Tennenholtz and others in this model and show our model also handles some games that have no program equilibria in the old sense. Finally we show how discounted time may allow us to understand strategies used in computationally difficult games like chess.

Keywords: Game Theory; Computational Complexity

Descriptive Complexity and Graphs with Excluded Minors

Martin Grohe (HU Berlin, Germany)

The central open problem in descriptive complexity theory is the question whether there exists a logic capturing PTIME. A good part of my talk will be an introduction into the background of this question for non-experts (and non-logicians). Then I will explain a recent result stating that fixed-point logic with counting captures polynomial time on all classes of graphs with excluded minors. A consequence of this result is that the Weisfeiler-Lehman algorithm, a simple and generic combinatorial algorithm, can be used as a polynomial time isomorphism test for such graph classes with excluded minors.

Keywords: Descriptive complexity, graph minors, graph isomorphism

Beating the random ordering is hard

Venkatesan Guruswami (University of Washington, USA)

Given a directed acyclic graph, it is easy to "topological sort" its vertices so all directed edges go forward in the ordering. But what if there is some noise and the graph is only nearly acyclic, say 1% of the edges need to be removed to make it acyclic. In this case, it was not known how to efficiently find an ordering of the vertices where even 51% of the edges go forward. (It is trivial to get half the edges going forward by picking a random ordering, or taking either the forward or backward edges in an arbitrary ordering.)

We prove that finding such an ordering is Unique-Games hard.

Specifically, for any constant $\epsilon > 0$, given a directed graph G that has an acyclic subgraph consisting of a fraction $(1 - \epsilon)$ of its edges, finding an acyclic subgraph of G with more than $(1/2 + \epsilon)$ of its edges is Unique-Games hard. This is the first tight hardness of approximation result for an ordering problem. Our result implies a super-constant factor inapproximability result (under the UGC) for the Feedback Arc Set problem.

Our proof uses two main ingredients: a directed acyclic graph constructed by Charikar, Makarychev, and Makarychev which looks "pseudorandom" (in terms of near equal forward/backward edges) at every "scale"; and Raghavendra's reduction to convert integrality gaps for CSPs into matching UG-hardness results.

Joint work of: Guruswami, Venkatesan; Manokaran, Rajsekar; Raghavendra, Prasad

Lower Bounds on Streaming Algorithms for Approximating the Length of the Longest Increasing Subsequence

Anna Gál (Univ. of Texas at Austin, USA)

We show that any deterministic data-stream algorithm that makes a constant number of passes over the input and gives a constant factor approximation of the length of the longest increasing subsequence in a sequence of length n must use space $\Omega(\sqrt{n})$. This proves a conjecture made by Gopalan, Jayram, Krauthgamer and Kumar.

Our results yield asymptotically tight lower bounds for all approximation factors.

Our proof is based on analyzing a related communication problem and proving a direct sum property for it.

Joint work of: Gál, Anna; Gopalan, Parikshit

Depth Reduction for Circuits with a Single Layer of Modular Counting Gates

Kristoffer Arnsfelt Hansen (University of Aarhus, Denmark)

We consider the class of constant depth AND/OR circuits augmented with a layer of modular counting gates at the bottom layer, i.e $\mathbf{AC}^0 \circ \mathbf{MOD}_m$ circuits. We show that the following holds for several types of gates G: by adding a gate of type G at the output, it is possible to obtain an equivalent randomized depth 2 circuit of quasipolynomial size consisting of a gate of type G at the output and a layer of modular counting gates, i.e $G \circ \mathbf{MOD}_m$ circuits. The types of gates G we consider are modular counting gates and threshold-style gates. For all of these, strong lower bounds are known for (deterministic) $G \circ \mathbf{MOD}_m$ circuits.

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1782

Some results on approximation resistance

Johan Håstad (KTH Stockholm, Sweden)

Assuming the Unique Games Conjecture we establish some general results on approximation resistance.

There is a universal constant $c < 1$, such that any predicate that accepts at least a fraction c of the inputs is approximation resistant.

There is a constant D such that a random predicate on k Boolean inputs that accepts Dk^2 of its input string, is, with high probability approximation resistant.

Keywords: Approximation resistance, unique games conjecture

On Parallel Repetition

Thomas Holenstein (Microsoft - Mountain View, USA)

In this talk, we review the consistent sampling lemma, which was recently used in the context of parallel repetition. We show how Ran Raz uses this lemma to get a counterexample to the strong parallel repetition conjecture, and then review how it is used in the proof of the parallel repetition theorem.

Keywords: Consistent Sampling, Parallel Repetition

Entropy of Operators or Nechiporuk for Depth-2 Circuits

Stasys Jukna (Universität Frankfurt, Germany)

We consider unbounded fanin depth-2 circuits with *arbitrary* boolean functions as gates.

We define the entropy of an operator $f : \{0,1\}^n \to \{0,1\}^m$ as the logarithm of the maximum number of vectors distinguishable by at least one special subfunction of f.

Main result: Every depth-2 circuit for f requires at least entropy(f) wires.

This is reminiscent of Nechiporuk's lower bound on the formula size, and gives an information-theoretic explanation of why some operators require many wires.

As a direct corollary this implies that n^3 wires are necessary to multiply two $n \times n$ matrices using depth-2 circuits with arbitrary gates.

Previously known lower bound for this operator was $n^2 \log n$.

Keywords: Entropy, depth-2 circuits, quadratic forms, matrix product

Full Paper:
http://www.springerlink.com/content/r048k436r53024v4/

Uniform Direct Product Theorems: Simplified, Optimized, and Derandomized

Valentine Kabanets (Simon Fraser University - Burnaby, Canada)

Direct Product Theorems are formal statements of the intuition: "if solving one instance of a problem is hard, then solving multiple instances is even harder". For example, a Direct Product Theorem with respect to bounded size circuits computing a function is a statement of the form: "if a function f is hard to compute on average for small size circuits, then $f^k(x_1, ..., x_k) = f(x_1), ..., f(x_k)$ is even harder on average for certain smaller size circuits". The proof of the such a statement is by contradiction: we start with a circuit which computes f^k on some non-negligible fraction of the inputs and then use this circuit to construct another circuit which computes f on almost all inputs.

As observed by Impagliazzo and Trevisan, such a constructive proof yields a decoding algorithm for the Direct-Product code, where the encoding of a message $f : \{0,1\}^n \to \{0,1\}$ (viewed as a truth table of length 2^n) is the (truth table of) the direct-product function f^k.

For the parameters of interest (when a received word has only tiny agreement with the correct codeword), it is impossible to decode the message uniquely, and so one needs to settle for list-decoding, where one outputs a list of possible messages. In this case, the list size is the main parameter to try to minimize.

We achieve an information-theoretically optimal value of the list size, which is a substantial improvement compared to previous proofs of the Direct-Product Theorem. In particular, this new version can be applied to uniform models of computation (e.g., randomized algorithms) whereas all previous versions applied only to nonuniform models (e.g., circuits).

Moreover, both our new list-decoding algorithm and its analysis are extremely simple. Finally, we also get a certain derandomized version of the direct-product code.

Keywords: Direct product, list-decoding, hardness amplification

Joint work of: Impagliazzo, Russell; Jaiswal, Ragesh; Kabanets, Valentine; Wigderson, Avi

Amplifying Lower Bounds by Means of Self-Reducibility - Part 2

Michal Koucký (Academy of Sciences - Prague, Czech Republic)

The lower bound amplification using downward self-reducibility shows that proving $n^{1+\epsilon}$ lower bounds for TC^0 circuits computing the usual NC^1 complete problems implies separation of TC^0 from NC^1. Similar results can be obtained for other circuit classes such as ACC^0, CC^0, etc. In this talk I will discuss the implications of these results for natural proofs. Further using downward self-reducibility I will show a lower bound on the size of a reduction between NP-complete problems. Time permitting I may demonstrate that ACC^0 is almost identical with CC^0.

Keywords: Bounded depth circuits, natural proofs, circuit lower bounds

Traffic Analysis and Complexity of Anonymous Communication

Miroslaw Kutylowski (Wroclaw University of Technology, Poland)

Investigating anonymity degree offered by communication protocols is one of the most challenging problems in computer security.

This requires providing adequate models capturing essential properties of the real protocol execution, finding appropriate quantitative measures, and designing analytic techniques for determining anonymity level expressed with these measures.

We overview recent advances in this area: information theoretic approach due to Berman, Fiat and Ta-Shma, as well as Markov chain approach and results by Gogolewski, Luczak and ourselves.

Keywords: Anonymous communication, traffic analysis

Joint work of: Klonowski, Marek; Kutylowski, Miroslaw

Hidden Subset Identifiers

Miroslaw Kutylowski (Wroclaw University of Technology, Poland)

We present a few results concerning authentication mechanisms for low-end devices.

While they are severely limited in computational power and communication complexity, an interactive protocol should deliver a high level of security against malicious adversaries.

Keywords: Authentication, RFID

Approximation norms and duality for communication complexity lower bounds

Troy Lee (Rutgers Univ. - Piscataway, USA)

We will discuss a general norm based framework for showing lower bounds on communication complexity. An advantage of this approach is that one can use duality theory to obtain a lower bound quantity phrased as a maximization problem, which can be more convenient to work with in showing lower bounds.

We discuss two applications of this approach.

1. The approximation rank of a matrix A is the minimum rank of a matrix close to A in ℓ_∞ norm. The logarithm of approximation rank lower bounds quantum communication complexity and is one of the most powerful techniques available, albeit difficult to compute in practice. We show that an approximation norm known as γ_2 is polynomially related to approximation rank.
   This results in a polynomial time algorithm to approximate approximation rank, and also shows that the logarithm of approximation rank lower bounds quantum communication complexity even with entanglement which was previously not known.
2. By means of an approximation norm which lower bounds multiparty number-on-the-forehead complexity, we show non-trivial lower bounds on the complexity of the disjointness function for up to $c \log \log n$ players, $c < 1$.

Keywords: Communication complexity, lower bounds

Joint work of: Lee, Troy; Shraibman, Adi

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1776

Understanding space in resolution: optimal lower bounds and exponential trade-offs

Jakob Nordström (MIT - Cambridge, USA)

We continue the study of tradeoffs between space and length of resolution proofs and focus on two new results:

1. We show that length and space in resolution are uncorrelated. This is proved by exhibiting families of CNF formulas of size $O(n)$ that have proofs of length $O(n)$ but require space $\Omega(n/\log n)$. Our separation is the strongest possible since any proof of length $O(n)$ can always be transformed into a proof in space $O(n/\log n)$, and improves previous work reported in [Nordström 2006, Nordström and Håstad 2008].
2. We prove a number of trade-off results for space in the range from constant to $O(n/\log n)$, most of them superpolynomial or even exponential. This is a dramatic improvement over previous results in [Ben-Sasson 2002, Hertel and Pitassi 2007, Nordström 2007].

The key to our results is the following, somewhat surprising, theorem:

Any CNF formula F can be transformed by simple substitution transformation into a new formula F' such that if F has the right properties, F' can be proven in resolution in essentially the same length as F but the minimal space needed for F' is lower-bounded by the number of variables that have to be mentioned simultaneously in any proof for F. Applying this theorem to so-called pebbling formulas defined in terms of pebble games over directed acyclic graphs and analyzing black-white pebbling on these graphs yields our results.

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1781

Joint work of: Ben-Sasson, Eli; Nordström, Jakob

Finding duplicates in a data stream

Jaikumar Radhakrishnan (TIFR Bombay, India)

Given a string of length n over an alphabet $[m]$ where $n > m$, we consider the problem of finding a repeated element in a single pass. We given a randomized algorithm for this problem that uses $O((\log m)^4)$ space. This is the first sublinear space one-pass algorithm for this problem, answering a question raised by Muthukrishnan and Tarui.

In contrast, we show that any deterministic one-pass algorithm for finding a repeated element needs space $\Omega(m)$, even if n is arbitrarily larger than m. However, we show that when we have access to a clock that records the current position in the stream, there is a deterministic algorithm using space $\log m + O(1)$ if $n \geq 2^m$; thus for this problem non-uniformity adds considerable power to deterministic algorithms.

Keywords: Isolation, Pairwise Independence, Datastream

Rounding Parallel Repetitions of Unique Games

Oded Regev (Tel Aviv University, Israel)

We show a connection between the semidefinite relaxation of unique games and their behavior under parallel repetition.

Specifically, denoting by $\mathrm{val}(G)$ the value of a two-prover unique game G, and by $\mathrm{sval}(G)$ the value of a natural semidefinite program to approximate $\mathrm{val}(G)$, we prove that for every $\ell \in I\!N$, if $\mathrm{sval}(G) \geq 1 - \delta$, then $\mathrm{val}(G^\ell) \geq 1 - \sqrt{s\ell\delta}$ Here, G^ℓ denotes the ℓ-fold parallel repetition of G, and $s = O(\log(k/\delta))$, where k denotes the alphabet size of the game. For the special case where G is an XOR game (i.e., $k = 2$), we obtain the same bound but with s as an absolute constant.

Our bounds on s are optimal up to a factor of $O(\log(1/\delta))$.

For games with a significant gap between the quantities $\mathrm{val}(G)$ and $\mathrm{sval}(G)$, our result implies that $\mathrm{val}(G^\ell)$ may be much larger than $\mathrm{val}(G)^\ell$, giving a counterexample to the strong parallel repetition conjecture. In a recent breakthrough, Raz (FOCS '08) has shown such an example using the max-cut game on odd cycles. Our results are based on a generalization of his techniques.

Keywords: Parallel repetition, unique games

Joint work of: Barak, Boaz; Hardt, Moritz; Haviv, Ishay; Rao, Anup; Regev, Oded; Steurer, David

Universal Steganography

Rüdiger Reischuk (Universität Lübeck, Germany)

In order to embed secret messages reliably and without being detectable into unsuspicious covertexts, a stegosystem has to draw samples from a covertext source/channel to estimate the distribution. For stegosystems that use a black-box sampler (there is no a priori knowledge of the covertext distribution), an exponential lower bound (with respect to the length of the secret message) has been shown for the query complexity of the sampling procedure.

However, it is assumed that the attacker has complete knowledge of the covertext distribution.

We consider a more fair and realistic situation where the stegoencoder and the attacker have the same state of knowledge concerning the covertext distribution.

Both have to learn the distribution by sampling.

It is investigated how algorithmic learning techniques can be used to design secure, reliable and computationally efficient stegosystems.

Positive results are obtained one the one hand for covertext channels with simple descriptions (concepts and hypothesis spaces) and on the other hand for pseudorandom channels.

Joint work of: Liskiewicz, Maciej; Reischuk, Rüdiger; Wölfel, Ulrich

Unconditional weak derandomization of weak algorithms: Explicit versions of Yao's Lemma.

Ronen Shaltiel (Haifa University, Israel)

The (easy direction) of Yao's minmax lemma says that if there is a randomized algorithm A which solves some problem (meaning that for every input, A succeeds with high probability) then there is a deterministic algorithm B of "roughly the same complexity" that solves the problem well on average (meaning that B succeeds with high probability on a random input). This can be viewed as "weak derandomization" and the statement follows by an averaging argument: there exist a fixed value r for A's random coins such that hardwiring r into A produces the deterministic algorithm B. Note that this averaging argument does not provide an explicit way to find r.

Recently, Zimand (building on an approach by Goldreich and Wigderson) proved an explicit version of the implication for randomized decision trees which toss "few" random coins. In this work, we consider weak derandomization of various classes of randomized algorithms.

We develop a new proof technique that applies to any class of randomized algorithms as long as one can explicitly construct an appropriate randomness extractor. Using this approach we prove unconditional weak derandomization results for communication games, constant depth circuits and streaming algorithms. More precisely we show that:

1. Given a randomized communication protocol that tosses few random coins and assuming that this protocol is explicitly constructible (in the sense that players can compute their strategy in polynomial time).
   Then, there is an explicitly constructible deterministic communication protocol of comparable communication complexity that simulates the randomized protocol correctly on most inputs.
2. Given a randomized algorithm that can be implemented by a uniform family of poly-size constant depth circuits we construct a uniform family of deterministic poly-size constant depth circuits that succeed on most inputs. (A classic result by Nisan and Wigderson gives a deterministic circuit that succeeds on all inputs but has quasi-polynomial size).

Our techniques follow the approach of Goldreich and Wigderson in the sense that we also "extract randomness from the input". However, in contrast to previous papers we use seedless extractors rather than seeded ones. We use extractors for bit-fixing sources (for decision trees) 2-source extractors (for communication games and streaming algorithms) and PRG based extractors (for constant depth circuits).

Keywords: Randomized algorithms, Derandomization, Yao's lemma, Average case complexity

The Pattern Matrix Method for Lower Bounds on Bounded-Error Communication

Alexander Sherstov (University of Texas - Austin, USA)

We develop a novel and powerful technique for lower bounds on bounded-error communication complexity, the pattern matrix method.

It works not only in the classical model but also in the quantum model, regardless of prior entanglement. Specifically, fix an arbitrary function $f : \{0,1\}^{n/4} \to \{0,1\}$ and let A be the matrix whose columns are each an application of f to some subset of the variables $x_1, x_2, \ldots, x_n$. We prove that A has bounded-error communication complexity $\Omega(d)$, where d is the approximate degree of f.

To illustrate our technique, we give a new proof of Razborov's breakthrough quantum lower bounds (2003) for all functions of the form $f(x,y) = D(|xANDy|)$, where $D : \{0,1,...,n\} \to \{0,1\}$ is a given predicate. Our paper also establishes a large new class of functions whose quantum communication complexity (regardless of prior entanglement) is only polynomially smaller than their classical complexity.

One of the ingredients of our proof is a certain equivalence of approximation and orthogonality in Euclidean n-space, which follows by linear-programming duality. Another ingredient is a construction of suitably structured matrices with low spectral norm, the pattern matrices, which we realize using matrix analysis and the Fourier transform over $(Z_2)^n$.

The method of this paper has recently enabled important progress in multiparty communication complexity.

Keywords: Bounded-error communication complexity, quantum lower bounds, classical/quantum gaps, pattern matrix method

Full Paper:
http://www.cs.utexas.edu/∼sherstov/publications/pdf/quantum-disc.pdf

See also: Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC), pp. 85–94, 2008

Explicit construction of a small epsilon-net for linear threshold function

Amir Shpilka (Technion - Haifa, Israel)

We give an explicit construction of ϵ-net for the family of linear threshold functions. More formally, for every $\epsilon > 0$ we give an explicit construction of a set $S \subset \{-1,+1\}^n$ of size $poly(n, 1/\epsilon)$, such that for any linear threshold function $f : \{-1,+1\}^n \to \{-1,+1\}$ if $\Pr[f(x) = 1] > \epsilon$ then there exists $x \in S$ with $f(x) = 1$. To the best of our knowledge no such construction was previously

known. Our result matches, up to a polynomial factor in the size of S, the parameters achieved by a random set.

As a consequence of our techniques we also get the following result: we give an explicit construction of a set $S \subset \{-1,+1\}^n$ of polynomial size such that for every $x \in \{-1,+1\}^n$ there exists $s \in S$ with $H(x,s) < n/2 - \sqrt{n \log n}$, where H is the Hamming distance. This improves on the well known construction of dual BCH codes that only guarantee covering radius of $n/2 - \sqrt{n}$.

The construction uses tools such as k-wise independent distributions, random walks on expander graphs and families of perfect hash functions.

Keywords: Threshold functions, epsilon-net, derandomization, explicit construction

Joint work of: Rabani, Yuval; Shpilka, Amir

Reconstruction of Depth-3 Arithmetic Circuits

Amir Shpilka (Technion - Haifa, Israel)

Depth-3 arithmetic circuits compute polynomials that can be represented as sums of products of linear functions. In spite of its simple structure, we are far from understanding this model.

In this talk we will focus on the reconstruction problem for depth-3 circuits, that have a constant number of multiplication gates. I.e., we have access to some unknown polynomial, that can be represented as a sum of a constant number of products of linear functions, and by asking for its values on (a small number of) inputs of our choice we would like to find a depth-3 circuit computing it.

We will show how to reconstruct such depth-3 circuits in time exp(polylog(s)), where s is the size of a depth-3 circuit computing the unknown polynomial. Our techniques rely on a theorem on the structure of zero depth-3 circuits and on a zero testing algorithm that it implies.

Keywords: Reconstruction, Arithmetic Circuits

Joint work of: Shpilka, Amir; Karnin, Zohar

A Combinatorial Construction of Almost-Ramanujan Graphs Using the Zig-Zag Product

Amnon Ta-Shma (Tel Aviv University, Israel)

Reingold, Vadhan and Wigderson introduced the graph zig-zag product.

This product combines a large graph and a small graph into one graph, such that the resulting graph inherits its size from the large graph, its degree from the small graph and its spectral gap from both. Using this product they gave the first fully-explicit combinatorial construction of expander graphs. They showed

how to construct D-regular graphs having spectral gap $1-O(D^{1/3})$. In the same paper, they posed the open problem of whether a similar graph product could be used to achieve the almost-optimal spectral gap $1-O(D^{-1/2})$.

In this paper we propose a generalization of the zig-zag product that combines a large graph and several small graphs. The new product gives a better relation between the degree and the spectral gap of the resulting graph. We use the new product to give a fully-explicit combinatorial construction of D-regular graphs having spectral gap $1-D^{-1/2+o(1)}$.

Keywords: Expanders

Joint work of: Ben-Aroya, Avi; Ta-Shma, Amnon

Full Paper:
http://portal.acm.org/citation.cfm?id=1374424

See also: STOC 2008

Fast polynomial factorization and modular composition

Christopher Umans (CalTech - Pasadena, USA)

We obtain randomized algorithms for factoring degree n univariate polynomials over F_q requiring $O(n^{1.5+o(1)}\log^{1+o(1)} q + n^{1+o(1)}\log^{2+o(1)} q)$ bit operations.

When $\log q < n$, this is asymptotically faster than the best previous algorithms (von zur Gathen & Shoup (1992) and Kaltofen & Shoup (1998)); for $\log q \geq n$, it matches the asymptotic running time of the best known algorithms.

The improvements come from new algorithms for modular composition of degree n univariate polynomials, which is the asymptotic bottleneck in fast algorithms for factoring polynomials over finite fields. The best previous algorithms for modular composition use $O(n^{(\omega+1)/2})$ field operations, where ω is the exponent of matrix multiplication (Brent & Kung (1978)), with a slight improvement in the exponent achieved by employing fast rectangular matrix multiplication (Huang & Pan (1997)).

We show that modular composition and multipoint evaluation of multivariate polynomials are essentially equivalent, in the sense that an algorithm for one achieving exponent α implies an algorithm for the other with exponent $\alpha+o(1)$, and vice versa. We then give two new algorithms that solve the problem optimally (up to lower order terms): an algebraic algorithm for fields of characteristic at most $n^{o(1)}$, and a non-algebraic algorithm that works in arbitrary characteristic.

The latter algorithm works by lifting to characteristic 0, applying a small number of rounds of *multimodular reduction*, and finishing with a small number of multidimensional FFTs. The final evaluations are reconstructed using the Chinese Remainder Theorem. As a bonus, this algorithm produces a very efficient data structure supporting polynomial evaluation queries, which is of independent interest.

Our algorithms use techniques which are commonly employed in practice, so they may be competitive for real problem sizes. This contrasts with all previous subquadratic algorithms for these problems, which rely on fast matrix multiplication.

Keywords: Modular composition; polynomial factorization; multipoint evaluation; Chinese Remaindering

Joint work of: Kedlaya, Kiran; Umans, Christopher

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1777

A generalization of the classification noise model of PAC learning and its applications to protein-protein interaction

Stephan Waack (Universität Göttingen, Germany)

We generalize the classification noise model of PAC learning introduced by Angluin and Laird (1988) to a model, where the noise can depend on the input vector. To be able to devise efficient learning algorithms, we utilize statistical queries due to Kearns (1998). Haussler's covering method is our vehicle for demonstrating how these ideas work in theory and practice. Having modifying its cover phase, we apply it to the problem of finding a powerful discriminator between protein-protein interfaces and random sets of pairs of protein surface residues. The hypothesis learned is convincing with respect to the classification rate as well as to its biological interpretation.

Keywords: Algorithms and data structures, algorithmic learning theory, bioinformatics

Joint work of: Waack, Stephan; Brodag, Thomas

Separating Deterministic from Randomized Multiparty Communication Complexity

Philipp Woelfel (University of Calgary, Canada)

We solve some fundamental problems in the number-on-forehead (NOF) k-player communication model. We show that there exists a function which has at most logarithmic communication complexity for randomized protocols with a one-sided error probability of 1/3 but which has linear communication complexity for deterministic protocols. The result is true for $k = n^{O(1)}$ players, where n is the number of bits on each players' forehead. This separates the analogues of RP and P in the NOF communication model.

We also show that there exists a function which has constant randomized complexity for public coin protocols but at least logarithmic complexity for private coin protocols. No larger gap between private and public coin protocols is possible. Our lower bounds are existential and we do not know of any explicit function which allows such separations. However, for the 3-player case we exhibit an explicit function which has $\Omega(\log \log n)$ randomized complexity for private coins but only constant complexity for public coins.

Keywords: Communication complexity, number on the forehead

Joint work of: Beame, Paul; David, Matei; Pitassi, Toni; Woelfel, Philipp

Full Paper:
http://pages.cpsc.ucalgary.ca/∼woelfel/paper/multiparty-RPvsP/paper.pdf

See also: Automata, Languages and Programming, 34th International Colloquium (ICALP 2007), pp. 134-145, volume 4596 of Lecture Notes in Computer Science, Springer, 2007.