Mathematisches Forschungsinstitut Oberwolfach
Tagungsbericht 44/1998

Complexity Theory

15–21 November 1998

Complexity theory is concerned with the study of the intrinsic difficulty of computational tasks. It is a central field of theoretical computer science.

The 13$^{\text{th}}$ Oberwolfach Conference on Complexity Theory was organized by Joachim von zur Gathen (Paderborn), Oded Goldreich (Rehovot), and Claus Peter Schnorr (Frankfurt).

The 1998 Meeting on Complexity Theory brought together 46 scientists with a major interest in this field. Among them were many prominent senior researchers as well as promising young researchers. They came from various countries in Europe, America, and Asia.

The meeting consisted of five general sessions and ten special sessions. The general sessions hosted 17 talks, ranging from high-level expositions of various sub-areas and research directions in complexity theory to highly technical talks dedicated to a single result.

The topics of the special sessions were:

- Algebraic Complexity I,II
- Circuit and Proof Complexity I,II
- Cryptography
- Extractors and Lattices
- Graph Problems and Complexity
- Lattice Theory
- Miscellaneous I,II

Some of the special sessions took place in parallel and were marked by talks intended for experts in the sub-area. Some sessions had very detailed technical expositions of new results and one had an open-panel discussion in which many attendees took part. In total, there were 30 talks and one discussion in the specialized sessions.

1 Abstracts of general session talks

Oded Goldreich:
Trevisan's Construction of Extractors Using Pseudo-Random Generators

We presented a new approach, due to Luca Trevisan (currently in DIMACS), for constructing extractors. Extractors are functions that transform a weakly random distribution into an almost uniform distribution. Explicit constructions of extractors have a variety of important applications, but tend to be very difficult to achieve.

The main result is the following unsuspected connection: Every pseudorandom generator based on a hard predicate (provided the proof of security uses a "black box" argument) yields an extractor. In particular, the Impagliazzo-Wigderson pseudorandom generator, yields an extractor that beats most known constructions and resolves an important open question in the area. Furthermore, using the simpler Nisan-Wigderson generator and standard error-correcting codes, one can construct even better extractors with the additional advantage that both the construction and the analysis are extremely simple and admit a short self-contained treatment (without reference to results on pseudorandomness).

Daniele Micciancio:
The Hardness of the Shortest Vector Problem

We prove that the Shortest Vector Problem is NP-hard to approximate within any constant factor less than $\sqrt{2}$. The inapproximability result is proved by randomized many-one reduction from the Closest Vector Problem, which is known to be hard to approximate within any constant factor. The key to the reduction is the construction of a lattice packing in dimension n such that there exists a sphere of radius less than the minimum distance of the lattice by a factor almost $\sqrt{2}$ containing exponentially many lattice points. The NP-hardness of SVP within the same constant factor can be proved also by deterministic many-one reductions, using a reasonable number theoretic conjecture concerning the distribution of smooth numbers.

Erich Kaltofen:
Algebraic Complexity and Algorithms: Recent Advances and New Open Problems

We cover four areas of algebraic questions.

A) sensitivity analysis: here an algebraic problem is given with rational coefficients (thought as fixed precision floating point numbers) that lacks a desired property. One wishes to compute in polynomial time in the input size the nearest problem that does. Nearness can be measured by different norms of the vector of floating point coefficients, for example by the maximum of the perturbations in each coefficient (infinity norm).

Problem	Norm	Solution	
Singular matrix	Matrix norm	Poly-time	Eckart & Young 1936
Singular matrix	Infinity norm	NP-hard	Poljak & Rohn 1990
Matrix with one eigenvalue on a given curve	Matrix norm	Poly-time	Hitz, Kaltofen & Lakshman 1998
Polynomial with a root with a positive real part	Infinity norm	Poly-time	Kharitonov 1978
Polynomial with real root	Infinity norm, Euclidean norm	Poly-time	Hitz & Kaltofen 1998
Polynomials with a common root(GCD)	Euclidean norm	Poly-time	Karmarkar & Lakshman 1996
	any norm	Open problem	
Polynomial in 2 variables that factors over $\mathbb{C}$	any norm	Open problem	

B) black box linear algebra: there have been many recent results on doing linear algebra problems simultaneously in $O(n)$ matrix times vector products that are viewed as oracle calls for the matrix, $n^{2+o(1)}$ arithmetic steps and $O(n)$ arithmetic space (n the matrix dimensions). For instance, with Las Vegas randomization one can certify a Diophantine linear system with a black box coefficient matrix inconsistent. An open problem is to compute the characteristic polynomial.

C) Mathematical constants in SC$^*(b)$: The formula

$$\pi = \sum_{i=0}^{\infty} \frac{1}{16^i} \left(\frac{4}{8\,i+1} - \frac{2}{8\,i+4} - \frac{1}{8\,i+5} - \frac{1}{8\,i+6} \right)$$

by Bailey, Borwein, and Plouffe (1996) allows the computation of the n-th digit of π in radix $b = 2$ (more precisely, of an approximation of π within precision $2^{(-n-1000)}$— there could be a very long sequence of ones in the binary expansion of π) in $n(\log n)^{O(1)}$ time and simultaneously $(\log n)^{O(1)}$ space. This defines the class of constants SC$^*(2)$. It is open whether $\pi \in SC^*(10)$.

D) Transposition principle: by a theorem from circuit theory it is known that a linear algorithm for a matrix-times-vector product can be transformed to one for the transposed matrix-times-vector product without asymptotic loss of time. The open problem is to simultaneously preserve space (asymptotically).

Toni Pitassi:

Proof Complexity, a Survey and New Results

We introduce algebraic proof systems for proving unsolvability of a family of (polynomial) equations over $\{0,1\}$. In particular, a polynomial calculus (PC) refutation of a sequence of equations $Q_1, \ldots, Q_r$ is a sequence of polynomials $s_1, \ldots, s_m$ such that each s_i is either an initial equation or follows from two earlier ones by either: $cQ_1 + dQ_2$, or xQ_i. The degree of such a proof is the maximal degree of any polynomial in the sequence. We prove nearby optimal lower bounds on the degree of any PC refutation of the equations describing various counting principles modulo some p.

Paul Beame:

Optimal Bounds for the Predecessor Problem

The predecessor problem is the classic problem solved by binary search given a stored set S from some ordered universe U, find the element of S that immediately precedes x, if any. In the comparison model, binary search is optimal $O(\log n)$ time but in the unit-cost RAM model with bounded word-size, which is representative of modern computer instruction sets, one can solve the problem more efficiently using data structures such as the binary trie data structures of van Emde Boas, Willard, and the fusion trees of Fredman & Willard in only $O(\sqrt{\log n})$ time. We give optimal upper and lower bounds for the problem, improving the time bound to $\Theta(\sqrt{\log n/\log\log n})$. The lower bounds are proved in the strong asymmetric communication complexity model that generalizes the cell-probe model which is itself more general than the unit-cost RAM model in which the algorithms are designed.

Joint work with Faith Fich.

Salil Vadhan:

Statistical Zero-Knowledge: A Survey of Recent Developments

Zero-knowledge proofs, introduced by Goldwasser, Micali, and Rackoff, are fascinating constructs which enable one party to convince another of an assertion without revealing anything other than the validity of the assertion. *Statistical* zero-knowledge proofs are a particular type of such proofs in which the condition that the verifier learns nothing is interpreted in a strong statistical sense. In this talk, we survey a number of recent results which have given us a much more refined understanding of statistical zero-knowledge proofs and the class SZK of languages ("assertions") which possess such proofs.

Particular items of focus in this survey are

- The role of Okamoto's theorem (1996) that any SZK proof can be converted into a "public coin" one in facilitating these recent improvements in our understanding of SZK.

- The use of "complete problems" to obtain new characterizations of the class and to reduce the study of the class to a single problem (as first seen in [Sahai, Vadhan, 1997]).

We illustrate the benefits of these two tools, by surveying some of the results that have been obtained:

- Strong boolean closure properties of SZK [Sahai, Vadhan, 1997].
- Converting honest verifier SZK proofs to any verifier SZK proofs [Goldreich, Sahai, Vadhan, 1998].
- Extending the theory to "noninteractive" SZK proofs [De Santis, Di Crescenzo, Persiano, Yung, 1998] and using this to relate SZK to "noninteractive" SZK [Goldreich, Sahai, Vadhan, 1998].

David Zuckerman:
Perfect Information Leader Election

Leader election is a generalization of collective coin flipping. In this latter problem, n players wish to generate a random bit. The difficulty is that some subset of players collude to bias the output of the bit. In the perfect information model, all communication is by broadcast, and the bad players have unlimited computational power and may wait to see the inputs of the good players. Thus, for example, PARITY can be broken by any coalition of 1 player, while in MAJORITY no coalition of $O(\sqrt{n})$ players can force a particular output to occur with probability $1 - o(1)$.

In leader election, the goal is to elect a good leader with probability bounded away from 0. We give a simple leader election protocol that is resilient against coalitions of size βn, for any $\beta < 1/2$, and takes $\log^* n + O(1)$ rounds.

We also give a lower bound for both collective coin-flipping and leader election in the case where each player can broadcast only 1 bit per round. In particular, we show that any protocol with linear resilience must take at least $[1/2 - o(1)] \log^* n$ rounds. The primary component of this result is a new bound on the influence of random sets of variables on Boolean functions.

This is joint work with Alex Russell and Mike Saks.

Michael Saks:
Time-Space Tradeoffs for Branching Programs

The branching program model is a well-studied combinatorial model that allows one to study the relationship between the time and space complexity of a computational problem. Here a computation is modeled by a digraph, where the computation time is lower bounded by the depth of the digraph, and the computation space is lower bounded by the logarithm of the number of nodes of the digraph. There has been considerable success in the past in proving time-space tradeoff lower bounds for multi-output functions such as sorting, and also in comparison based models.

In the case of single-output boolean functions, space lower bounds were known in a restricted model (the so-called syntactic read-k-times model) but essentially nothing was known for unrestricted models. If one ignores size, then every n-variable boolean function has a branching program of depth n (just take a decision tree for computing the function), and this bound is known to be tight for "most" functions and for many explicit ones. However, the depth n branching program typically requires exponential size. A major research direction is to study the power of polynomial size branching programs. A modest, but thus far elusive, goal is to prove, for some explicit boolean function (in, say, complexity class P), that any polynomial size branching program requires superlinear depth. Prior to the present work, no such lower bound on depth greater than $n + o(n)$ was known. Here we prove the first (barely) nontrivial bound of this type by exhibiting an explicit function in P for which any subexponential size branching program requires depth at least $1.0178n$.

This is joint work with Paul Beame and Jayram Thatachar.

Ingo Wegener:
Branching Programs and Binary Decision Diagrams — Complexity and Algorithms

BPs and BDDs are used as complexity theoretical model for space complexity but also as representations of Boolean functions in CAD applications. One looks for variants which allow a compact representation of many (important) functions and efficient algorithms for synthesis, satisfiability test, equivalence test, minimization, and replacement by constants. Several variants (BPs, DTs, OBDDs, π-OBDDs, FBDDs, G-FBDDs, $(1,+k)$BPs, BPks, nondet. OBDDs, partitioned OBDDs, EXOR-OBDDs, and randomized BDDs) are investigated. Upper and lower bound techniques are discussed and efficient manipulation algorithms are presented.

Noam Nisan:
Algorithms for Selfish Agents

One interesting aspect of distributed algorithms designed for the Internet environment is that the computers that are to participate belong to different users. It is likely that these computers may not follow the algorithm but rather manipulate it according to their self interest. The algorithm designer should thus attempt to design the distributed algorithm in a way that will ensure that the participants' self interest is maximized by behaving correctly.

In this talk I will present a model for studying such algorithms, based on the theory of mechanism design. In this model the algorithm is augmented with payments to the participants. These carefully chosen payments should motivate the players to act "correctly".

Most of the talk will be devoted to a self-contained thinly-disguised introduction to the field of mechanism design — from a distributed computation perspective. I will then present some new results and suggest directions for further research.

The new results in this talk are joint work with Amir Ronen.

Adi Shamir:
Attacks on Algebraic Cryptosystems

Several multivariate algebraic cryptographic schemes have been proposed in recent years, but most of them have been broken by exploiting the fact that their secret keys are low rank algebraic structures. In this talk I survey the proposed schemes and the developed attacks, and in particular show a novel attack on Patarin's "oil and vinegar" signature scheme, which does not contain low rank structures.

Ran Raz:
Exponential Separation of Quantum and Classical Communication Complexity

Communication complexity has become a central complexity model. In that model, we count the amount of communication bits needed between two parties in order to solve certain computational problems. We show that for certain communication complexity problems quantum communication protocols are exponentially faster than classical ones. More explicitly, we give an example for a communication complexity relation (or promise problem) $\mathcal{P}$ such that:

1. The quantum communication complexity of $\mathcal{P}$ is $O(\log m)$.
2. The classical probabilistic communication complexity of $\mathcal{P}$ is $\Omega(m^{1/4}/\log m)$.

(where m is the length of the inputs). This gives an exponential gap between quantum communication complexity and classical probabilistic communication complexity. Only a quadratic gap was previously known.

Our problem $\mathcal{P}$ is of geometrical nature, and is a finite precision variation of the following problem: Player I gets as input a unit vector $x \in \mathbb{R}^n$ and two orthogonal subspaces $M_0, M_1 \subset \mathbb{R}^n$. Player II gets as input an orthogonal matrix $T : \mathbb{R}^n \to \mathbb{R}^n$. Their goal is to answer 0 if $T(x) \in M_0$ and 1 if $T(x) \in M_1$, (and any answer in any other case). We give an almost tight analysis for the quantum communication complexity and for the classical-probabilistic communication complexity of this problem.

Avi Wigderson:
Randomness vs. Time

We give the first nontrivial deterministic simulation of BPP under a uniform complexity assumption weaker that one-way functions. Specifically, we show that if BPP$\neq$EXP, then every language in BPP has a deterministic subexponential time algorithm that is correct on all but a negligible fraction of inputs.

Joint work with Russell, Impagliazzo (UC San Diego).

Allan Borodin:
Lower Bounds for Geometric Search Problems

The *curse of dimensionality* describes the phenomenon whereby (in spite of extensive and continuing research) for various geometric search problems we only have algorithms with performance that grows exponentially in the dimension. Recent results [1, 2, 3] show that in some sense it is possible to avoid the curse of dimensionality for the approximate nearest neighbor search problem. But must the exact nearest neighbor search problem suffer this curse? We provide some evidence in support of the curse. Specifically we investigate the exact nearest neighbor search problem and the related problem of exact partial match within the asymmetric communication model first used by Miltersen [4] to study data structure problems. We derive non-trivial asymptotic lower bounds for the exact problem that stand in contrast to known algorithms for approximate nearest neighbor search. Specifically, in the d-dimensional Hamming cube, we consider the nearest neighbour decision problem (i.e. does there exist a vector in a preprocessed data base that is within a specified distance λ of the query point) and the partial match decision problem (i.e. given a query with both "exposed" and don't care coordinates, is there a vector in the preprocessed data base that exactly matches the query on the exposed coordinates). We assume that the dimension d is relatively large and growing with respect to the number of points; in particular, we assume $\log n << d = d(n) << n^{\kappa}$ where $\kappa > 0$ is arbitrarily small and n is the size of the data base. Using the richness technique of Miltersen, Nisan, Safra and Wigderson [5] we show that either the query player sends a total of $\Omega(\log d \log n)$ bits or the data base player sends a total of $n^{1-\epsilon}$ bits where ϵ is an arbitrarily small constant. In contrast to the approximate nearest neighbour problem, this implies that any constant round cell probe algorithm for these exact problems must either use a non polynomial (i.e. $n^{\log d}$) number of storage cells or else must transmit cells of size $n^{1-\epsilon}$.

References

[1] J. Kleinberg. Two algorithms for nearest-neighbor search in high dimensions. In *Proc. of 29th STOC*, pp. 599–6

[2] P. Indyk and R. Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proc. of 30th STOC*, 1998.

[3] E. Kushilevitz, R. Ostrovsky, and Y. Rabani. Efficient search for approximate nearest neighbor in high dimensional spaces. In *Proc. of 30th STOC*, 1998.

[4] P.B. Miltersen. On the cell probe complexity of polynomial evaluation. *Theoretical Computer Science*, 143:167–174, 1995.

[5] P.B. Miltersen, N. Nisan, S. Safra, and A. Wigderson. On data structures and asymmetric communication complexity. In *Proc. of 27th STOC*, pp. 103-111, 1995.

Uri Feige:

Heuristics for Finding Large Independent Sets, with applications to coloring semi-random graphs

We study a semi-random graph model for finding independent sets. For $\alpha > 0$, an n-vertex graph with an independent set S of size αn is constructed by blending random and adversarial decisions. Randomly and independently with probability p, each pair of vertices, such that one is in S and the other is not, is connected by an edge. An adversary can then add edges arbitrarily (provided that S remains an independent set). The smaller p is, the larger the control the adversary has over the semi-random graph. We design heuristics that with high probability recover S when $p > (1+\epsilon)\ln n/|S|$, for any constant $\epsilon > 0$. We show that when $p < (1-\epsilon)\ln n/|S|$, an independent set of size $|S|$ cannot be recovered, unless NP $\subseteq$ BPP.

We use our results to obtain greatly improved coloring algorithms for the model of k-colorable semi-random graphs introduced by Blum and Spencer.

Joint work with Joe Kilian of NEC Research Institute.

Johan Håstad:

The security of all RSA bits

We prove that if RSA encryption is secure then any individual bit is hard to predict. To be more exact let $E_{N,e}(x)$ be encryption of x under RSA with modulus N and exponent e. We prove that if bit i of x can be predicted with probability $1/2 + n^{-c}$ in probabilistic polynomial time from $E_{N,e}(x)$, then all of x can be recovered in probabilistic polynomial time. This is true for any c and any i such that bit i is unbiased. If bit i is biased then $1/2$ in the above formula should be replaced by the success of the best constant guessing algorithm.

Claus Peter Schnorr:

On the Generic Group Model

Let G be a cyclic group of prime order q with generator g. In the generic model we count arbitrary group operations at unit costs. All other operations are for free, e.g. equality tests for group elements and computations on non-group elements. It is known that the generic complexity of the discrete logarithm is $\Theta(\sqrt{q})$ (Nechaer, Shoup). We show that the generic security of the j least significant discrete log bits is $\Omega(\frac{1}{n^2 j^2}\sqrt{q/2^j})$. We also show that subsets $H \subset G$ of size $|H| = \sqrt{q}$ are hardcore for the discrete log, except for a negligible fraction of H. For $|H| < \sqrt{q}$ the generic complexity of the discrete log, restricted to H, is $\frac{m}{2} + o(m)$.

2 Abstracts of special session talks

Avi Wigderson:

Short Proofs Are Narrow

We introduce width of Resolution proofs, and relate it to traditional complexity measures of size and tree-like size. These relations give a unified simple way to derive most exponential size lower bounds. It also yields simple new automated theorem prover for Resolution.

Joint work Eli Ben-Sasson (Hebrew University).

Georg Schnitger:

Neural Circuits and Efficient Associative Memory

Arbitrary patterns $x_1, \ldots, x_m \in \{0,1\}^n$ are given. For a query $y \in \{0,1\}^n$ estimates of the Hamming distance between y and the closest pattern x_i have to be provided. If additionally small errors are allowed, then significant memory result. It is for instance possible to solve this relaxed version of the Hamming distance problem with $O(m \cdot \log_2 n)$ memory, provided constant factor approximations are allowed with constant error probability. This should be compared to the obvious $O(m \cdot n)$ memory bound.

Neural circuits are designed that simultaneously achieve near optimal memory complexity as well as near optimal depth and number of edges. Under some conditions a near optimal number of gates can be reached as well.

Dima Grigoriev:

Tseitin's Tautologies and Lower Bounds for Nullstellensatz Proofs

A linear lower bound for the degree of a Nullstellensatz proof was shown for the system of binomials arising from Tseitin's tautologies. Earlier a sublinear lower bound was known for a polynomial system arising from the pigeon-hole principle.

Ingo Wegener:

Relating Branching Program Size and Formula Size over the Full Binary Basis

The classical complexity measures for Boolean functions are circuit size C, branching program size BP, formula size L, and formula size over $\{\wedge, \vee, \neg\}$ L^*. The following trivial relations are known and optimal: $C \leq 3BP, C \leq L \leq L^*, L^* \leq L^2$, and $BP \leq L^*$. This implies $BP \leq L^2$. This inequality is improved to $BP \leq 1.36L^{\log_4(3+\sqrt{5})}$ where $\log_4(3+\sqrt{5}) \approx 1.19$. The transformation leads for read-once formulas to optimal OBDDs.

Johannes Blömer:

Complexity of Short Linear Independent Vectors

Motivated by Ajtai's worst-case to average-case reduction for lattice problems, we study the complexity of computing short linearly independent vectors (short basis) in a lattice. We show that approximating the length of a shortest set of linearly independent vectors (shortest basis) within any constant factor is NP-hard. Under the assumption that problems in NP cannot be solved in $\mathsf{DTIME}(n^{\mathrm{polylog}(n)})$ we show that no polynomial time algorithm can approximate the length of a shortest set of linearly independent vectors (shortest basis) within a factor of $2^{\log^{1-\epsilon}(n)}$, $\epsilon > 0$ arbitrary, but fixed. Finally, we obtain results on the limits of non-approximability for computing short linearly independent vectors (short basis). Our strongest result in this direction states that under reasonable complexity-theoretic assumptions, approximating the length of a shortest set of linearly independent vectors (shortest basis) within a factor of $n/\sqrt{\log(n)}$ is not NP-hard.

Joint work with Jean-Pierre Seifert.

Jean-Pierre Seifert:

SVP Is Not Harder Than CVP

We show that for any poly-bounded function $f : \mathbb{N} \to \mathbb{N}$, such that f is non-decreasing, SVP_f is poly-time reducible to CVP_f. In a later talk by D. Micciancio the construction was extremely simplified and extended to arbitrary non-decreasing functions $f : \mathbb{N} \to \mathbb{N}$.

Claus Peter Schnorr:

Fast LLL-like Reduction

The famous LLL-algorithm transforms a given lattice basis into a LLL-reduced basis. For this it performs $O(n^4 \log M)$ arithmetic steps on integers of bit length $O(n \log M)$, where n is the dimension of the lattice and M is the maximal length of the basis vectors. We propose a variant of LLL-reduction that uses merely $O(n^3 \log M)$ arithmetic steps. For this we partition the basis $b_1, \dots, b_n$ with $n = k \cdot m$ in blocks of size k. We perform a standard LLL-reduction within the blocks. We relax the condition of LLL-reduced bases at the border of the blocks. For $k = O(\sqrt{n})$ the algorithm runs in $O(n^3 \log M)$ steps. This algorithm is a practical and useful variant of the algorithm of Schönhage (1984) for semi reduction.

Daniele Micciancio:

Proof of the Technical Lemma

See D. Micciancio's talk on "The Hardness of the Shortest Vector Problem".

Rüdiger Reischuk:

Optimal Lower Bounds for the Average Case Complexity of PARITY

We investigate the average time complexity (delay) of parallel prefix functions defined over a groupoid for the unbounded fan-in circuit model — the parity function is one such example. Properties of the corresponding transition graph, confluent and strictly diffluent, are defined. They are mutually exclusive and complete, that is exactly one of them occurs. Based on these notions we can show that a function has unbounded fan-in circuits with constant average time complexity iff it is confluent. In case of strict diffluence the average case complexity for circuits of size s is $\Theta(\log n/\log\log s)$. Thus the average case complexity of parity is as large as its worst case complexity.

Joint work with Andreas Jakoby.

Pavel Pudlak:

On Generalized Tseitin Tautologies

Let H be a k-uniform hypergraph (ie. edges have size k), with k even and suppose c is a mapping which assigns 0's and 1's to vertices of H so that the number of 1's is odd. The generalized Tseitin tautology expresses the trivial fact that for each assignment of 0's and 1's to the edges of the hypergraph H there exists a vertex v such that $c(v)$ is different from the parity of the values assigned to the edges incident to v.

We consider a game associated to H in which first player chooses edges and the second player chooses values of the edges. The goal of the first player is to find a vertex where there is the disagreement of parities. Our aim is to design a nondeterministic strategy for the second player such that in each play there are many nondeterministic moves (ie. the second player can choose any value). A lower bound r on the number of such moves implies a lower bound 2^r on the size of a decision tree solving the search problem of the first player.

We present a strategy which can potentially be used for obtaining bounds close to 2^m, where m is the number of edges. This will imply the same lower bound on the tree-resolution proofs of the generalized tautology and running time of any DLL algorithm for l-CNF's, where l is the maximal degree of H.

Ran Raz:

Separation of the Monotone NC Hierarchy

We prove tight lower bounds, of up to n^ϵ, for the monotone depth of functions in monotone-P. As a result we achieve the separation of the following classes.

1. monotone-NC $\neq$ monotone-P.
2. $\forall i \geq 1$, monotone-$NC^i \neq$ monotone-NC^{i+1}.

3. More generally: For any integer function $D(n)$, up to n^ϵ (for some $\epsilon > 0$), we give an explicit example of a monotone Boolean function, that can be computed by polynomial size monotone Boolean circuits of depth $D(n)$, but that cannot be computed by **any** (fan-in 2) monotone Boolean circuits of depth less than $Const \cdot D(n)$ (for some constant $Const$).

Only a separation of monotone-NC^1 from monotone-NC^2 was previously known.

Our argument is more general: we define a new class of communication complexity search problems, referred to below as DART games, and we prove a tight lower bound for the communication complexity of every member of this class. As a result we get lower bounds for the monotone depth of many functions. In particular, we get the following bounds:

1. For st-connectivity, we get a tight lower bound of $\Omega(\log^2 n)$. That is, we get a new proof for Karchmer-Wigderson's theorem, as an immediate corollary of our general result.

2. For the k-clique function, with $k \leq n^\epsilon$, we get a tight lower bound of $\Omega(k \log n)$. Only a bound of $\Omega(k)$ was previously known.

Peter Bürgisser:
The Computational Complexity to Evaluate Immanents and Representations of General Linear Groups

We describe a fast algorithm to evaluate irreducible matrix representations of general linear groups $\mathrm{GL}_m(\mathbb{C})$ with respect to a symmetry adapted basis (Gelfand-Tsetlin basis). This is complemented by a lower bound, which shows that our algorithm is optimal up to a factor m^2 with regard to nonscalar complexity. Our algorithm can be used for the fast evaluation of special functions: for instance, we obtain an $O(\ell \log \ell)$ algorithm to evaluate all associated Legendre functions of degree ℓ.

For studying the complexity to evaluate single entries of the representation matrix, we investigate the complexity of immanents. These matrix functions contain the permanent and determinant as special cases. While the determinant can be computed in polynomial time, a well-known completeness result due to Valiant indicates that the permanent is computationally intractable. We obtain an algorithm to evaluate immanents, which is faster than previous algorithms due to Hartmann and Barvinok. Finally, we show that the problem to evaluate certain immanents corresponding to hook diagrams or rectangular diagrams is complete in Valiant's sense.

Marc Giusti:
Applications of Efficient Geometric Solving

We show two applications of the geometric resolution of a system of polynomial equations. This notion goes back to the geometric description of the associated algebraic geometry by Kronecker's elimination theory : the zero locus of an eliminating polynomial and rational parametrizations (also called generalized shape lemma or rational univariate representation in computer algebra).

When the input equations form a secant family outside of an hypersurface, there exists an efficient algorithm constructing incrementally a geometric resolution in any dimension [G – Haegele – Heintz – Montaãna – Morais – Pardo]. The running time is polynomial in the number of variables, the total degree of the equations, and a suitably defined notion of degree of the system (at most the Bézout number) and linear in the length of a straight-line program evaluating the equations and the inequation.

We apply first the algorithm to approximate inconsistent non-linear systems. Given a rational (inconsistent) approximation of a real consistent over-determined system, we recover an approximation of its solution(s) under mild assumptions. A practical example consists in the over-determined static Stewart's platform and is treated successfully on experimental data: recover the mutual position of two solids from the knowledge of seven (and not six) distances between pairs of points [G – Schost].

Eventually we consider real solving: find a point on every connected component of a real smooth compact hypersurface given by a regular equation. This is achieved through a systematic use of polar varieties, and yields an efficient theoretical algorithm. The complexity inherits the good behaviour of the previous one, the degree becoming the maximal degree of all successive polar varieties [Bank – G – Heintz – Mbakop]. The result improves the up to now known complexities on the subject.

Volker Weispfenning:
Mixed Real-Integer Linear Quantifier Elimination

We prove some positive and some negative results concerning effective quantifier elimination (QE) and decidability of linear fragments of the elementary theory of the real numbers $\mathbf{R}$ with the set $\mathbf{Z}$ of integers a definable subset.

Theorem 1 *The elementary theory of* $\mathbf{R}$ *in the language* $L_1 = \{0, 1, +, -, <, \{\underset{n}{\equiv}\}_{n\in\mathbf{N}}, \mathbf{Z}(.)\}$ *is decidable, but does not admit* QE*. Here congruences are defined by* $a \underset{n}{\equiv} b$ *iff* $a - b \in n\mathbf{Z}$.

Theorem 2 *The elementary theory of* $\mathbf{R}$ *in the language* $L_2 = \{0, 1, +, -, <, \{\underset{n}{\equiv}\}_{n\in\mathbf{N}}, [.]\}$ *admits effective* QE *and hence is decidable. Here* $[.]$ *is the integer-part function.*

Theorem 3 *The elementary theory of* $\mathbf{R}$ *in the language* $L_3 = \{0, 1, +, -, |\}$ *is undecidable. Here the divisibility relation is defined by* $a|b$ *iff* $b \in \mathbf{Z}a$.

Theorem 4 *The positive existential theory of* $\mathbf{R}$ *in the language* $L_4 = \{0, 1, +, -, |'\}$ *is decidable. Here the modified divisibility relation is defined by* $a|b$ *iff* $a \in \mathbf{Z}$ *and* $b \in \mathbf{Z}a$.

By known results on Presburger arithmetic, QE for $\mathbf{R}$ in L_2 is at least triply exponential in the worst case, and the decision problem for this theory requires at least doubly exponential time on an alternating Turing machine with a linear number of alternations. We conjecture that the upper complexity bounds are of the same type.

As a consequence of Theorem 2, the L_2-definable subsets of $\mathbf{R}$ are composed of finite unions of rational intervals in an ultimately periodic fashion.

Omer Reingold:
Trevisan's Extractors: The Next Generation

We give explicit constructions of extractors which work for a source of any min-entropy k on strings of length n. These constructions are based on the recent work of Trevisan and they reduce the number of uniformly distributed bits used by Trevisan's extractors in two cases:

1. When trying to extract all the min-entropy of the weak source (or a constant fraction of it).

2. When the output of the extractor should be very close to uniform.

In particular to output $m = \Omega(k)$ bits that are ϵ-close to uniform the number of truly random bits used by our extractors is $O(\min\{(\log(n/\epsilon))^2, (\log n)^2\epsilon\})$.

In addition, for any extractor that outputs the entire min-entropy, we show a way to reduce its entropy-loss to $2\log(1/\epsilon) + O(1)$ bits which is optimal up to a constant additive term.

Joint work with Ran Raz and Salil Vadhan.

Ammon Ta-Shma:
Almost Optimal Dispersers

We present an explicit construction of efficient dispersers. A $G = (V, W, E)$ bipartite graph is a $(K, epsilon)$-disperser if for any subset X of W of cardinality K, the neighbors of X cover at least $(1-\epsilon)$ fraction of W. The goal is to build such graphs with small degree D. For any K and ϵ we show a construction where $D = \mathrm{poly}(\log|V|, 1/\epsilon)$, which is almost optimal.

Salil Vadhan:
Pseudorandom Generators without the XOR-Lemma

Impagliazzo and Wigderson have recently shown that if there exists a decision problem solvable in time $2^{O(n)}$ and having circuit complexity $2^{\Omega(n)}$ (for all but finitely many n) then P=BPP. This result is a culmination of a series of works showing connections between the existence of hard predicates and the existence of good pseudorandom generators.

The construction of Impagliazzo and Wigderson goes through three phases of "hardness amplification" (a multivariate polynomial encoding, a first derandomized XOR Lemma, and a second derandomized XOR Lemma) that are composed with the Nisan–Wigderson generator. In this work, we present two different approaches to proving of the main result of Impagliazzo and Wigderson. In developing each approach, we introduce new techniques and prove new results that could be useful in future improvements and/or applications of hardness-randomness trade-offs.

Our first result is that when (a modified version of) the Nisan–Wigderson generator construction is applied with a "mildly" hard predicate, the result is a generator that produces a distribution indistinguishable from having large min-entropy. An extractor can then be used to produce a distribution computationally indistinguishable from uniform. This is the first construction of a pseudorandom generator that works with a mildly hard predicate without doing hardness amplification.

Our second result, not described in this talk, is that in the Impagliazzo–Wigderson construction only the first hardness-amplification phase (encoding with multivariate polynomial) is necessary, since it already gives the required average-case hardness. We prove this result by (i) establishing a connection between the hardness-amplification problem and a list-decoding problem for error correcting codes based on multivariate polynomials; and (ii) presenting a list-decoding algorithm that improves and simplifies a previous one by Arora and Sudan.

Joint work with Madhu Sudan and Luca Trevisan.

Oded Goldreich:
GapCVP within sqrt n in coAM (On Limitation of Non-Approximability)

We show simple constant-round interactive proof systems for problems capturing the approximability, to within a factor of $\sqrt{n}$, of optimization problems in integer lattices; specifically, the closest vector problem (CVP), and the shortest vector problem (SVP). These interactive proofs are for the "coNP direction"; that is, we give an interactive protocol showing that a vector is "far" from the lattice (for CVP), and an interactive protocol showing that the shortest-lattice-vector is "long" (for SVP).

We conclude that approximating CVP (resp., SVP) within a factor of $\sqrt{n}$ is in $\mathrm{NP} \cap \mathrm{coAM}$. Thus, it seems unlikely that approximating these problems to within a $\sqrt{n}$ factor is NP-hard. Previously, for the CVP (resp., SVP) problem, Lagarias et al showed that the gap problem corresponding to approximating CVP within $n^{1.5}$ (resp., approximating SVP within n) is in $\mathrm{NP} \cap \mathrm{coNP}$. On the other hand, Arora et al showed that the gap problem corresponding to approximating CVP within $2^{\log^{0.499} n}$ is quasi-NP-hard.

Joint work with Shafi Goldwasser.

Daniele Micciancio:

Approximating shortest lattice vectors is not harder than approximating closest lattice vectors

We show that given oracle access to a subroutine which returns approximate closest vectors in a lattice, one may find in polynomial-time approximate shortest vectors in a lattice. The level of approximation is maintained; that is, for any function f, the following holds: Suppose that the subroutine, on input a lattice $\mathcal{L}$ and a target vector $\mathbf{t}$ (not necessarily in the lattice), outputs $\mathbf{v} \in \mathcal{L}$ such that $\|\mathbf{v}-\mathbf{t}\| \leq f(n)\|\mathbf{u}-\mathbf{t}\|$ for any $\mathbf{u} \in \mathcal{L}$. Then, our algorithm, on input a lattice $\mathcal{L}$, outputs a nonzero vector $\mathbf{v} \in \mathcal{L}$ such that $\|\mathbf{v}\| \leq f(n)\|\mathbf{u}\|$ for any nonzero vector $\mathbf{u} \in \mathcal{L}$.

This result estalishes the widely believed conjecture by which the shortest vector problem is not harder than the closest vector problem. The proof can be easily adapted to establish an analogous result for the corresponding problems for linear codes.

Joint work with Oded Goldreich, Muli Safra, and Jean-Pierre Seifert.

Dima Grigoriev:

Exponential Lower Bounds on the Size of Depth-3 Arithmetic Formulae for the Determinant

We prove exponential lower bounds on the size of depth-3 arithmetic formulae for the determinant over a finite field treating it in

(a) the polynomial ring (this is a joint result with M.Karpinski);

(b) the algebra of functions (even admitting some errors) over a field (this is a joint result with A.Razborov)

Markus Bläser:

Lower Bounds for the Multiplicative Complexity of Matrix Multiplication

We prove a lower bound of $km + mn + k - m + n - 3$ for the multiplicative complexity of the multiplication of $k \times m$–matrices with $m \times n$–matrices using substitution method ($k, n \geq 2$). For the multiplication of 3×3–matrices, this yields the improved lower bound of 18.

Peter Bürgisser:

Some New Results in Valiant's Algebraic Model of NP-completeness

Some new results in Valiant's algebraic theory of NP-completeness
In 1979 Valiant proposed an algebraic analogue of the theory of NP-completeness for computations of polynomials over a field, in connection with his famous

hardness result for the permanent. We have further developed this theory and present some of our new results.

Over finite fields, we exhibit a *specific* example of a family of polynomials (cut polynomials), which is neither complete nor p-computable, provided the polynomial hierarchy does not collapse. In the classical setting, no such specific examples are known.

We establish the following connection between the classical (nonuniform) P-NP hypothesis and Valiant's hypothesis. If the nonuniform versions of NC and NP are different, then Valiant's hypothesis is true over infinite fields of characteristic zero. Our proof is based on tools from algebraic geometry and number theory, and relies on the generalized Riemann hypothesis.

Moni Naor:
What is a Cryptographic Assumption? Discussion.

Friedhelm Meyer auf der Heide:
"Balls into Bins" — New Variations of an Old Game

Assume n balls and n bins are given, each ball chooses independently and uniformly a random bin. It is well known that some bin gets $\approx \log(n)/\log\log(n)$ balls, i.e., the contention is $\approx \log(n)/\log\log(n)$, w.h.p. This game often appears e.g. in protocols for task allocation or distributed data servers. In the letter case, it is assumed that many data objects are distributed randomly among n disks, and an access to n given objects has to be realized. The above results says that some disks have nothing to do, others have to answer $\approx \log(n)/\log\log(n)$ requests, w.h.p. Starting with a paper by Karp, Luby, Meyer auf der Heide (92), a series of results have come up where the following variants of the game are considered. Each ball now chooses not just one, but some number d many bins. A protocol then decides for each ball, which of the d chosen bins accepts the ball. If the balls arrive sequentially (sequential game) then contention $\approx \log\log(n)/\log(d)$ can be achieved, in case of all balls arriving concurrently, contention can be reduced to constant, w.h.p. In both cases, very fast protocols are known.

The talk presents two new results: The first is a design proposal for a data server: the n disks and the n users (issuing the requests to objects) are connected by a butterfly- like network. Then not only constant contention at the disks can be guaranteed, w.h.p., but also $\log\log(n)$ congestion at the edges of the network. (with Vöcking, Schröder (98))

The second result (by Vöcking (98)) is a variant of the sequential protocol that guarantees contention $\approx \log\log(n)/_{d}$ rather than $\log\log(n)/_{\log(d)}$.

Hans Jürgen Prömel:
On the Complexity of the Steiner Tree Problem

The Steiner tree problem in networks (SPN) is the following problem. Given a network $N = (V, E, T, \ell)$, where (V, E) is a graph, $T \subseteq V$ the set of terminals and $\ell : E \to Q^+$ a length function. Find a shortest tree in N connecting the vertices in T.

The SPN is APX-complete but the best known "non-approximability constant" so far is $1 + 1/5600$. In recent years there was considerable interest in finding good upper bounds for this constant, i.e., in finding approximation algorithms for the SPN with low performance ratio.

We present a general iterative framework for improving the performance ratio of Steiner tree approximation algorithms. By applying this framework to one specific algorithm we obtain a new polynomial time approximation algorithm for the SPN that achieves a performance ratio of 1.598 after 11 iterations. This beats the so far best known factor of 1.644. This result is joint work with S. Hougardy (Berlin).

Martin Fürer:

Combinatorial Isomorphism Tests for Graphs with Small Spectral Multiplicities?

The k-dimensional Weisfeiler-Lehman algorithm (k-dim W-L or canonical k-tuple coloring algorithm) seems to be more natural and is stronger than any proposed combinatorial graph isomorphism test. It succeeds if it is able to decompose the vertex set (or the set of k-tuples of vertices) into orbits under the automorphism group. For $k = o(n)$, it has been shown to fail badly for graphs of bounded degree. It might work for graphs of bounded genus, but is conjectured to fail for graphs of bounded eigenvalue multiplicity. For $k = o(n/\log n)$, the k-dim W-L algorithm is shown to fail even for graphs whose eigenspaces are decomposed by the 2-dim W-L algorithm into 1-dimensional canonical subspaces.

Matthias Krause:

Approximation by OBDDs, the Variable Ordering Problem, and Genetic Programming

A lot of BDD (binary decision diagram) variants are motivated by CAD applications and have let to several complexity theoretical problems and results. Here, methods from communication complexity and information theory are combined to prove that the direct storage access function and the inner product function have the following property. They have linear π-OBDD size for some variable ordering π but, for almost all variable orderings π' all functions which compute them on considerably more than half of the inputs, need exponential π'-OBDD size. These complexity theoretical lower bounds have implications for the use of OBDDs in genetic programming.

Joint work with Petr Savicky (Uni Prague), Ingo Wegener (Uni Dortmund).

Uri Feige:
A simple protocol for leader election in the full information model

Inspired by David Zuckerman's talk earlier in the workshop, a simple leader election protocol is presented. It is based on games of throwing balls into bins. When there are $(1+\epsilon)n/2$ good players, then the protocol chooses a good leader with probability $\epsilon^{O(\log 1/\epsilon)}$. The protocol has low communication complexity, and in particular can be implemented in $\log^* n + O(\log 1/\epsilon)$ rounds.

Amit Sahai:
Overview of Concurrent ZK

Adi Shamir:
How to Quickly Find Secret RSA Keys

In this talk we consider the problem of efficiently locating cryptographic keys hidden in gigabytes of data, such as the complete file system of a typical PC. We develop efficient algebraic attacks which can find the secret keys of several public key cryptosystems in a way which is much faster than testing each possible substring as a potential key.

Jean-Pierre Seifert:
Extending Wiener's Attack on RSA

Wiener has shown that when the RSA protocol is used with a decrypting exponent d such that $d < N^{1/4}$ and an encrypting exponent e such that $e \approx N$ then d can be recovered from the c.f.a. to e/N.

We extend this attack to the case when many e_i for a given N all with "small" e_i are available. For the case of two such e_i and d_i, the d_i can be as large as $N^{5/14}$ and still be recovered by means of lattice basis reduction algorithms.

As the number of encrypting exponents available increases the bound on the d_i increases slowly to $N^{1-\epsilon}$. However, the complexity of our method is exponential in the number of exponents available.

Joint work with Nicholas Howgrave-Graham.

This report was edited by Michael Nüsken.

Contents