abstracts

- Monday, 28 September 1998
  Dana Ron, MIT
  A sublinear bipartite tester for bounded-degree graphs
  Abstract: We present a sublinear-time algorithm for testing whether a bounded-degree graph is bipartite or far from being bipartite. The algorithm is given oracle access to the incidence list of the graph, and should determine with high probability whether the braph is bipartite or ϵ-far from bipartite for any given distance parameter ϵ. Distance between graphs is defined to be the fraction of entries on which the graphs differ in their incidence-lists representation.
  Our testing algorithm has query complexity and running time $\poly((\log N)/\epsilon)\cdot\sqrt N$ where N is the number of graph vertices. The query complexity almost matches a previously known lower bound.
  A major part of our analysis is showing that any graphy can be partitioned into subsets such that each subset exhibits a certain rapid-mixing property, and the total number of edges between the subsets is small.
  This is joint work with Oded Goldreich.

- Monday, 5 October 1998
  Luca Trevisan, Columbia University and DIMACS
  Construction of near-optimal extractors using pseudo-random generators
  Abstract: We introduce a new approach to construct extractors--combinatorial objects akin to expander graphs that have several applications. Our approach is based on error correcting codes and on the Nisan-Wigderson pseudorandom generator.
  An application of our approach yields a construction that is simple to describe and analyze, does not utilize any of the standard techniques used in related results, and improves or subsumes almost all the previous constructions.

- Monday, 12 October 1998
  Avi Wigderson, IAS/Hebrew University
  Resolution made simple
  Abstract: Resolution is perhaps the simplest nontrivial propositional proof system. It forms the basis of the most popular automated theorem proving procedures. Proving lower bounds on the length of proofs in this system for simple tautologies seemed difficult and very technical.

We consider the width (= maximal clause size) of resolution proofs. We prove a simple relation between this parameter and classical proof size. This yields much simpler proofs of exponential proof-size lower bound for simple tautologies, such as the pigeonhole principle, Tseitin graph tautologies and random k-cnf's. It also motivates a new simple automatic theorem prover, that can be tremendously more efficient than Davis-Putnam. This is joint work with Eli Ben-Sasson.

- Monday, 19 October 1998
  Paul Seymour, Princeton University
  Digraph Minors
  Abstract: The ``graph minors'' project was very fruitful, with all kinds of nice results. But what about digraphs, is there any hope of an analogue of the graph minors theorems for digraphs? What is a minor of a digraph anyway? How hard is it to test if one digraph is a minor of another? Is there such a thing as the treewidth of a digraph? Some desirable things are not true, but others are true and provable; and some (even nicer) things can be conjectured. This talk is a survey of preliminary investigations into these questions, partly joint with Matt DeVos, Thor Johnson, Bruce Reed, Neil Robertson and Robin Thomas.

- Monday, 26 October 1998
  Dmitry Kozlov, IAS
  On the action of the symmetric group on the selected-type partition lattices
  Abstract: In the first half of the talk I will outline from scratch the principal objects of study and tools of topological combinatorics, such as the nerve functor, the Goresky-MacPherson theorem about the cohomology groups of the complement of a subspace arrangement, various combinatorial techniques and some aspects of the group actions on posets.
  After that, I will present an example of a computation of the Betti numbers of a combinatorially defined cell complex C. C is the quotient of X - the order complex of a partition lattice (or a selected-type sublattice) - by the action of the symmetric group. The partition lattice is the intersection lattice of the braid arrangement and its order complex is important in the computation of the cohomology groups of the complement of the braid arrangement and also appears in Vassiliev's work on knot invariants. One interesting feature of the Betti numbers of C is that they measure the multiplicity of the trivial character in t he induced representation of the symmetric group on the homology groups of X.
  If time permits, I will discuss similar problems, where the constant sheaf is replaced by an arbitrary cellular sheaf.

- Monday, 9 November 1998
  Michael Saks, Rutgers University
  Time-Space tradeoffs for Boolean branching programs
  Abstract: 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 tradoff 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, a 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.

- Monday, 16 November 1998
  Bela Bollobas, Memphis State University and Cambridge University
  Dependent percolation in two dimensions
  Abstract: We shall present some recent results obtained jointly with Paul Balister and Alan Stacey on Peter Winkler's problem about dependent site percolation defined by two independent random sequences. Similar results have been obtained independently by Peter Winkler. Numerous open problems remain.

- Monday, 23 November 1998
  Dorit Aharonov, IAS
  Shor's Quantum Factorization Algorithm
  Abstract: Quantum computation is a computation model which uses quantum physical systems as computational devices. As for now, it is the only computational model which is not known to be polynomially reducible to a classical randomized Turing machine. There is evidence that this model can exponentially speed up certain computations.
  The most powerful quantum algorithm known today is Shor's algorithm, which factorizes an integer N in time and space which are polynomial in log(N), the size of the input. The best known classical algorithm is exponential. This algorithm is potentially

extremely important from the practical point of view, since the security of the widely used RSA cryptographic systems relies on the hardness of factorization.
Shor's algorithm is based on Fourier transform over the group Z_Q, which can be performed very quickly on a quantum computer. Apparently, the Fourier transform is the source of the exponential speed up in all known quantum algorithms.
I will define the model of quantum computation, and describe Shor's algorithm. I will also try to explain what the origins for the possible extra power of quantum computation are, using Shor's algorithm as an example.

- Monday, 30 November 1998
  Vijay Vazirani, Georgia Tech
  On the Bidirected cut relaxation for the metric Steiner tree problem
  Abstract: The Steiner tree problem was defined by Gauss in a letter to Schumacher -- today it occupies a central place in the emerging theory of approximation algorithms. Interest in this problem arises not only because of its rich mathematical structure, but also because it has arisen repeatedly in diverse applications.
  None of the currently known algorithms with proven approximation guarantees is suitable as the core algorithmic idea for solving large instances in practice, such as those arising in the VLSI design industry. Perhaps the most promising avenue is a remarkable LP-relaxation that has its origins in the work of Edmonds on branchings. This relaxation is conjectured to have integrality gap close to 1. However, even though this relaxation has been known for decades, there has been no success in designing algorithms using it or upper bounding its integrality gap.
  In this work, we restrict our attention to quasi-bipartite graphs -- graphs that do not have edges connecting pairs of Steiner vertices. This enables us to finesse the difficulty caused by such edges and address the rest of the aspects of the problem. We give a $3/2 + \epsilon$ factor algorithm for this class of graphs, for any $\epsilon > 0$. The algorithm involves extending the primal-dual schema in two important ways: for the first time, the dual growth process is not greedy, and the algorithm does not use the usual mechanism of relaxing complementary slackness conditions.
  Our algorithm yields a natural heuristic for general graphs. Preliminary experimental comparisons with other heuristics, on benchmarks obtained from the VLSI design industry, are quite promising. This is joint work with Sridhar Rajagopalan; experiments performed by Ion Mandoiu.

- Monday, 7 December 1998
  Jeff Kahn, Rutgers University
  Combinatorial uses of entropy
  Abstract: Entropy has turned out to be a useful tool for various discrete problems, but is perhaps not as widely known as it should be. In this talk we give minimal basics and some illustrative applications.

- Monday, 14 December 1998
  Jean Bourgain, IAS
  Arithmetic progressions in finite sets
  Abstract: We show that every set A of at least cn (log log n/ log n)^{1/2} integers between 1 and n contains a three term arithmetic progression. This improves previous estimates of Roth, Heath Brown and Szemeredi. The proof is based on the circle method, and the main new idea is that if A does not contain roughly as many progressions as a random set of the same density then it has a higher density in an appropriately defined "Bohr set".

- Monday, 18 January 1999
  Michael Krivelevich, DIMACS/Rutgers University
  Testing regular languages
  Abstract: In 1996 Goldreich, Goldwasser and Ron launched a systematic study of combinatorial property testing. A very general setting they considered is, for a given property P (which has usually a certain combinatorial meaning) and an input function f, to determine quickly and reliably whether f belongs to P or it is far from any function in P in a certain given metric. An algorithm is allowed to query the value of f on x, chosen randomly from the domain of f or deterministically. A complexity of the algorithm is measured by the number of queries it asks. Already in their original paper Goldreich et al. were able to get several remarkable results on testing graph properties such as k-colorability, Max Clique, Max Cut.
  In this talk we explore yet another avenue of property testing by turning our attention to regular languages. A language L over the binary alphabet A={0,1} is a subset of all words of finite length over A. Using graph theoretic terms, we can define informally a regular language as a language which can be described by a directed multigraph with edges labeled by symbols 0,1. Thus, regular languages form about the simplest possible class of languages.
  Our main result states that all regular languages are testable. More precisely, we prove the following theorem:
  For a given regular language L, large enough integer n and small enough parameter epsilon>0, there exists an algorithm which, for an input word w of length n, produces the following output: a) if w\in L, then the algorithm always outputs "YES"; b) if w is at least epsilon n bits far from any word in L, the algorithm outputs "NO" with probability at least 2/3. The algorithm queries, up to logarithmic in 1/epsilon factors, 1/epsilon bits of w.
  Thus we are able we solve this testing problem extremely quickly.
  I will indicate main ideas of the proof of this result. I will not assume any familiarity with regular languages or combinatorial property testing.
  This is a joint work with Noga Alon (Tel Aviv), Mario Szegedy (AT&T Research) and Ilan Newman (Haifa).

- Monday, 25 January 1999

Miklos Bona, IAS

A combinatorial proof of the log-concavity of the numbers of permutations with k runs

Abstract: The theory of permutations with a given number of runs has been studied by Don Knuth in connection with sorting and searching. In this talk, using a new lattice path interpretation, we combinatorially prove that the number $R(n,k)$ of permutations of length n having k runs is a log-concave sequence in k, for all n. We also give a new combinatorial proof for the log-concavity of the Eulerian numbers.
This is joint work with Richard Ehrenborg.

- Monday, 1 February 1999

Yehuda Shalom, Princeton University

Groups and expanders

Abstract: Expander graphs are finite graphs with a strong connectivity property. Their existence is known by general counting methods, but the explicit constructions use deep tools from representation and number theory, and are all essentially Cayley graphs of (arithmetic) families of finite groups, with respect to very special sets of generators. The dependence of the property on the choice of generators, as well as other fundamental questions in the subject, will be discussed, along with several new results.

- Monday, 8 February 1999

Van Vu, IAS

Small complete arcs in projective planes

Abstract: An arc of a projective plane is a set of points with no three on a line. The arc is complete if no other point from the plane could be added to it without violating this property. The notion of arcs and complete arcs was developed by B. Segre in the 50's and 60's.
Given a projective plane, determining the size of the smallest complete arc is one of the main open questions in finite geometry. Let $n(P)$ denote the size of a smallest complete arc. For any projective plane of order q, a lower bound $n(P) > \sqrt {2q}$ was already shown in the 50's by Lunelli and Sce, but no close upper bound has been known. For a Galois plane, which is a special projective plane, the best upper bound was $n(P) < c q^{3/4}$. The proof (due to Sz\H onyi) made use of Hesse-Weil's theorem and therefore depends on the structure of the fields.
Practically nothing has been known for general planes. In this paper, we will show that there is some constant c such that $n(P) < q^{1/2} \log ^c q$ for any projective plane P. This matches the lower bound within a polylogarithmic factor.
The proof uses a probabilistic method known as the semi-random method or R\"odl nibble. The core of this is a new and very useful concentration result. This is a quite surprising application of a probabilistic method in this area where algebra seems to dominate. >From the algorithmic point of view, the proof gives a randomized algorithm which terminates in $polylog(q)$ number of steps, (each step has polynomially

bounded number of basic operations) and outputs a small almost complete arc with probability close to 1. This almost complete arc can be completed in a very quick manner.

- Monday, 22 February 1999
  Meir Katchalski, Rutgers/Technion, Israel
  Touching and representing convex sets
  Abstract: Two seemingly unrelated problems on families of convex sets will be discussed:
  1) Given a planar family A of n convex sets we wish to contract each set into a convex polygon in such a way that the total number of vertices is minimal and the family of polygons has the same intersection pattern as that of the original family. Bounds on the the number of vertices needed when no three members of A intersect are obtained.
  2) Two planar sets touch if they intersect and there is a straight line that separates one from the other. The planar families A and B are called touching if every member of A touches every member of B. It will be shown that there is a $0 < c < 1$ such that if A and B are touching families, each containing n convex sets, then there is a point contained either in at least cn members of A or in at least cn members of B.
  The results are joint work with Liu and Pach.

- Monday, 1 March 1999
  Avi Wigderson, IAS
  Some problems in algebraic complexity
  Abstract: In this talk I will present several concrete open problems regarding the difficulty of computing some natural functions. The nature of these problems is on the border of Algebra and Combinatorics. I will also demonstrate some of the (meager) bounds which the (currently) best techniques yield.

- Monday, 8 March 1999
  Eva Tardos, Cornell
  Approximation algorithms for the k-median problem, and other clustering problems
  Abstract: The k-median problem is one of the most well-studied clustering problems, i.e., those problems in which the aim is to partition a given set of points into clusters so that the points within a cluster are relatively close with respect to some measure. For the metric k-median problem, we are given n points in a metric space. We must select k cluster centers so as to minimize the average distance between a point and the closest center. In this talk we will survey approximation algorithms for various clustering problems, and present a constant factor approximation algorithm for the k-median problem.

My work on clustering is joint with Moses Charikar, Sudipto Guha, and David Shmoys.

- Monday, 15 March 1999
  Salil Vadhan, MIT
  Statistical Zero-Knowledge: An Introduction and a Survey of Recent Developments
  Abstract: Zero-knowledge proofs, introduced by Goldwasser, Micali, and Rackoff in 1985, 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 SZK 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].

- Monday, 22 March 1999
  Mario Szegedy, AT&T
  Algebraic problems raised in complexity theory
  Abstract: In the talk we survey a number of algebraic problems, many of them open, that have led or would lead to complexity theoretic lower bounds.
  An open problem [Szegedy, Thorup]: Give an integer valued polynomial time computable map from $\{0,1\}^{n}$ to $C^{n^{2}}$ which does not factor through a polynomial map $(p_{1},\ldots,p_{n^{2}})$ from $C^{n log n}$ to $C^{n^{2}}$ such that the total degree of p_{i} ($1\le i\le n^{2}$) is at most $log^{2} n$.
  A solved problem [Beame, Saks, Thathachar]: There is a family of $n\times n$ matrices A_{n} ($n=1,...$) such that A_{n} is computable in time polynomial in n, and for any fixed $0 < \delta < 1$ there is a threshold n_{0} such that for $n>n_{0}$ the rank of each $\delta n \times \delta n$ sub-matrix of A_{n} is at least $\delta^{2} n$.
  If time allows we shall also explain how different complexity theoretic problems reduce to the algebraic problems we discuss.

- Monday, 29 March 1999
  Thor Johnson, Princeton University
  Eulerian Digraph Immersion
  Abstract: The Graph Minors series of Robertson and Seymour has yielded many interesting concepts and results. This talk concerns a recent attempt to achieve similar goals with immersion of Eulerian digraphs replacing minors of undirected graphs. The talk will begin with an explanation of immersion and why we consider this particular containment relation. Tree-width and path-width (two fruitful concepts from Graph Minors) for Eulerian digraphs will be defined, and I will discuss how certain Eulerian digraphs are tied to large tree- or path-width. I will then present a theorem which describes the structure of Eulerian digraphs which do not immerse a given Eulerian digraph. Finally, if time allows, an algorithm will be presented which, given vertices ${s1,...,sk,t1,...,tk}$, checks in polynomial time for fixed k for the existence of pairwise edge-disjoint directed paths from si to ti for all i in an Eulerian digraph.

Sponsored by:

National Science Foundation

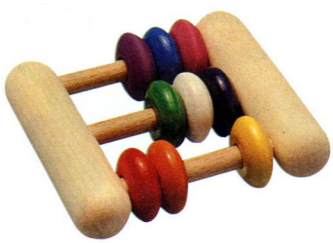

Abacus

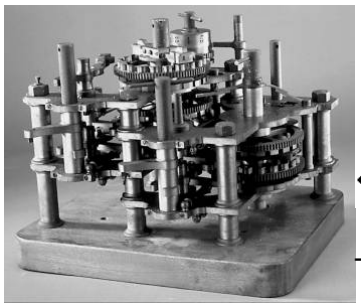

"Analytical Engine"

by Charles Babbage

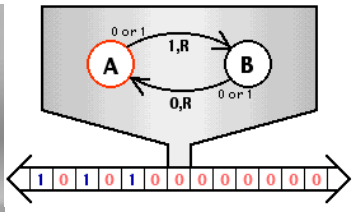

Turing Machine

State of New Jersey