abstracts

Monday, September 9, 2002

Valentine Kabanets, University of California, San Diego

Derandomizing Polynomial Identity Tests Means Proving Circuit Lower Bounds
Abstract:
We show that derandomizing the Polynomial Identity Testing is, essentially, equivalent to proving circuit lower bounds
for NEXP. More precisely, we prove that if one can test in polynomial time (or, even, nondeterministic subexponential
time, infinitely
often) whether a given arithmetic circuit over integers computes an identically zero polynomial, then either (i) NEXP
is not in P/poly or (ii) Permanent is not computable by polynomial-size arithmetic circuits. We also prove a (partial)
converse: If Permanent requires superpolynomial-size arithmetic circuits, then one can test in subexponential time
whether a given arithmetic formula computes an identically zero polynomial.

Since the Polynomial Identity Testing is a coRP problem, we obtain the following corollary: If RP=P (or, even,
coRP is in NSUBEXP, infinitely often), then NEXP is not computable by polynomial-size arithmetic circuits. Thus,
establishing that RP=coRP or BPP=P would require proving superpolynomial lower bounds for Boolean or arithmetic circuits.

This is joint work with Russell Impagliazzo.

Monday, September 9, 2002

Russell Impagliazzo, University of California, San Diego

A Switching Lemma for Small Restrictions and Lower ounds for K-DNF Resolution
Abstract:
We prove a new switching lemma
that works for restrictions that set only a small
fraction of the variables
and is applicable to
DNFs with small bottom fan-in.
We use this to
prove lower bounds for the Res(k) propositional

proof system, an extension of resolution
whose lines are k-DNFs instead of clauses.
We also obtain an exponential separation between
depth d circuits of bottom fan-in k and depth
d circuits of bottom fan-in k+1.

Our results for Resk proofs are:

1. The 2n to n weak pigeonhole principle requires exponential size to refute in Res(k), for k up to a power of log n.

2. For each constant k, there exists a constant w>k so
that random w-CNFs require exponential size to refute in Res(k).

3. For each constant k,
there are sets of clauses which have polynomial size
Res(k+1) refutations, but which
require exponential size Res(k) refutations.

Joint Work with: Nathan Segerlind and Samuel Buss

Tuesday, September 17, 2002

Dror Weitz, University of California, Berkeley

Mixing in Time and Space on the Integer Lattice - A Combinatorial View

Abstract:
We consider spin systems on the d-dimensional integer lattice $\Z^d$ with
nearest-neighbor interactions and prove
a sharp equivalence between exponential decay with distance of spin correlations (a spatial property of the equilibrium
state) and "super-fast'' mixing time of the Glauber dynamics (a temporal property of a Markov chain Monte Carlo algorithm).
While such an equivalence is already known in various forms, we give proofs that are purely combinatorial and avoid the
functional analysis machinery employed in previous proofs.

In the talk, I will define and explain the above notions of temporal and spatial mixing, discuss why the equivalence
between the them is interesting, and describe the main ideas of our proofs.

Joint work with Martin Dyer, Alistair Sinclair and Eric Vigoda.

Monday, September 23, 2002

Joel Spencer, Courant Institute

Phase Transitions for Random Processes
Abstract:

The Erdos-Renyi evolution of the random graph G(n,p) undergoes a fundamental change near p=1/n when a "giant component"
rapidly appears. We view this change as a phase transition. We explore a variety of random processes which exhibit
similar behavior at appropriate critical probabilities. We especially search for the right notion of the critical
window, in which the movement from the precritical to the postcritical phase is best observed.

Monday, September 23, 2002

Jiri Matousek, Charles University, Prague

Topological Lower Bounds for the Chromatic Number: A Hierarchy
Abstract:
The Lovasz-Kneser theorem states that for n>2k, the k-element subsets of {1,2,...,n} cannot be colored by
fewer than n-2k+2 colors so that no two disjoint k-tuples receive the same color. This is a remarkable result
and it provides an important example of highly chromatic graphs. Over the years, several proofs have been found
(all based on the Borsuk-Ulam theorem in topology or on variations of it); many of them imply general lower
bounds for the chromatic number of graphs. The plan is to present an overview of these proofs and lower bounds,
and indicate how they fall into an almost linearly ordered hierarchy.
(Joint work with G. M. Ziegler)

Tuesday, September 24, 2002

Sanjeev Arora, Princeton University

Proving Integrality Gaps Without Knowing the Linear Porgram

Abstract:
Many approximation algorithms for NP-hard optimization problems are designed using a linear program relaxation
of the problem. The integrality gap of the relaxation is the worst-case ratio of the cost of the optimum (integer)
solution to the optimum value of the linear program. If we can show that the integrality gap is large, it rules out
using that linear program for computing good approximations.

Proving integrality gaps is a difficult task and usually undertaken on a case-by-case basis. We initiate a more
systematic approach that proves integrality gaps for large families of linear programs. We prove an integrality gap
of 2-o(1) for three families of relaxations for vertex cover, including those obtained from the Lovasz-Schrijver
"lift-and-project" proof system. Our methods seem relevant to other problems as well.

(Joint work with Bela Bollobas and Laszlo Lovasz)

Monday, September 30, 2002

Moses Charikar, Princeton University

Dimension Reduction in the 1_1 Norm

Abstract:

The Johnson-Lindenstrauss Lemma shows that any set of n points in Euclidean space can be mapped linearly down
to O((log n)/eps^2) dimensions such that all pairwise distances are distorted by at most 1+eps. We study the following basic question: Does there exist an
analogue of the Johnson-Lindenstrauss Lemma for the l_1 norm?

Note that the Johnson-Lindenstrauss Lemma gives a linear embedding which is independent of the point set.
For the l_1 norm, we show that one cannot hope to use linear embeddings as a dimensionality reduction tool
for general point sets, even if the linear embedding is chosen as a function of the given point set.
In particular, we construct a set of O(n) points in l_1 such that any linear embedding into l_1^d must incur
a distortion of Omega(sqrt{n/d}). This bound is tight up to a log n factor. We then initiate a systematic
study of general classes of l_1 embeddable metrics that admit low dimensional, small distortion embeddings.
In particular, we show dimensionality reduction theorems for circular-decomposable metrics, tree metrics,
and metrics supported on K_{2,3}-free graphs, giving embeddings into l_1^{O(log^2 n)} with constant distortion.
Finally, we also present lower bounds on dimension reduction techniques for other l_p norms.

Our work suggests that the notion of a stretch-limited embedding, where no distance is stretched by more than a
factor d in any dimension, is important to the study of dimension reduction for l_1. We use such stretch limited
embeddings as a tool for proving lower bounds for dimension reduction and also as an

algorithmic tool for proving
positive results.

This is joint work with Amit Sahai and will appear in FOCS '02.

Tuesday, October 1, 2002

Hartmut Klauck, IAS

Quantum Security in the Bounded Storage Model
Abstract:

In the bounded storage model introduced by Maurer (1992) a public source of random bits is used to expand
a secret key shared by two parties to a much longer key that is almost uniformly distributed, even given
the information available to an eavesdropper that can store a constant fraction of the information in the
transmitted random bits and later also obtains access to the secret key. The long derived key can then be
used to encode messages as in a one-time pad. While a strong security proof has been obtained this year by
Dziembowski and Maurer, no security proof known so far addresses the case that an eavesdropper might also
use some quantum storage. We analyze the performance of the scheme of Dziembowski and Maurer under this
condition and essentially retain the results given there in the stronger model of a quantum eavesdropper.
Our proof is, however, very different, and simpler, using ideas from quantum information theory.

Joint work with Harry Buhrman.

Monday, October 7, 2002

Tim Roughgarden, Cornell

The Elusiveness of Braess's Paradox: Designing Networks for Selfish Users is Hard
Abstract:
Given a network with congestion-dependent edge delays and a prescribed source-destination pair and traffic rate,
which subnetwork will exhibit the best performance when used selfishly? Braess's Paradox is a famous example that
shows that the trivial algorithm of choosing the whole network is a suboptimal heuristic; we show that it is the
best possible polynomial-time heuristic, unless P=NP. We also give an infinite family of

examples generalizing
Braess's Paradox, providing the first demonstration that the severity of the paradox grows with the network size.
Time permitting, we will conclude with very recent progress (joint with Richard Cole and Yevgeniy Dodis) showing
that edge tolls can radically improve the performance of selfish routing, even when edge deletions cannot.

Tuesday, October 8, 2002

Ran Canetti, IBM Watson Research Center

Universally Composable Security: Overview of the Paradigm and some Constructions

Abstract:
Recently, a new paradigm for defining security of cryptographic protocols was proposed. The salient property of
notions of security that follow this paradigm (called universally composable (UC) security) is that they guarantee
security even when a protocol is used as a component within an arbitrary system. In particular, UC notions guarantee
security even when an unbounded number of protocol instances are running concurrently in an adversarially controlled
manner, they guarantee non-malleability with respect to arbitrary protocols, and more. Such properties are crucial
for arguing the security of cryptographic protocols in complex and unpredictable environments such as the Internet.

The first part of the talk reviews the general methodology for formulating UC notions of security, and the composition
theorem that underlies the strong security properties provided. We then quickly survey a number of works that use the UC
paradigm in a variety of settings. The second part describes in more detail some of these works, demonstrating how to solve
"any cryptographic protocol problem" in a universally composable way.

Monday, October 14, 2002

Irit Dinur, NEC Research

The Hardness of 3-Uniform Hypergraph Coloring

Abstract:

We prove that coloring a 3-uniform 2-colorable hypergraph with any constant number of colors is NP-hard. The best known

algorithm [KNS] colors such a graph using O(n^{1/5}) colors. Our result immediately implies that for any constants k>2
and c_2>c_1>1, coloring a k-uniform c_1-colorable hypergraph with c_2 colors is NP-hard; leaving completely open only
the k=2 graph case.

We are the first to obtain a hardness result for approximately-coloring a 3-uniform hypergraph that is colorable with a
constant number of colors. For k-uniform hypergraphs with k\ge 4 such a result has been shown by [GHS], who also discussed
the inherent difference between the k=3 case and k > 3.

Our proof presents a new connection between the Long-Code and the Kneser graph, and relies on the high chromatic numbers
of the Kneser graph and the Schrijver graph. Joint work with Oded Regev and Clifford Smyth.

Tuesday, October 15, 2002

Xiaodong Sun, IAS

Time-space Tradeoff Lower Bounds for Randomized Computation

Abstract:
We prove the first time-space lower bound tradeoffs for randomized computation of decision problems. The bounds hold
even in the case that the computation is allowed to have arbitrary probability of error on a small fraction of inputs.
Our techniques are extension of those used by Ajtai and by Beame, Jayram, and Saks that applied to deterministic branching
programs. Our results also give a quantitative improvement over the previous results.

Joint work with Paul Beame, Mike Saks and Erik Vee.

Tuesday, October 21, 2002

Jan Krajicek, Czech Academy of Sciences, Prague

Free and Pseudo-Surjective Functions, and Provability of Circuit Lower Bounds

Abstract:
Let g be a polynomial-time computable function
extending n bits to m(n) > n bits. The notions of
freeness and pseudo-surjectivity of g w.r.t. a
propositional proof system P formalize a situation
when it is consistent to think in P that g is
surjective (the two notions differ in parameters only).

I show that if there is any free (resp. pseudo-surjective) function for P then the truth table function is such
a function. The truth table function takes as an input a circuit with k inputs and size bounded above by some c^k,
c < 2 a constant, and outputs the truth table of the function computed by the circuit (i.e. 2^k bits). In particular,
if there is a free/pseudo-surjective function for P then P does not prove any exponential circuit lower bound.
The complexity of g, together with the ratio m/n, determine a circuit class for which P proves no lower bounds.

Further, for such g only finite NP-sets are P-provably disjoint from Rng(g), and this property alone implies
(is equivalent to, in fact) that all tau-formulas from g are hard for P. Hence functions g of interest are
those behaving like a hitting set generator good w.r.t. NP-properties.

Tuesday, October 22, 2002

Xiaodong Sun, IAS

Time-space Tradeoff Lower Bounds for Randomized Computation (Continued)
Abstract:

Last time we discussed time-space tradeoff lower bounds for decision problems over large domain. We continue
to discuss lower bounds for boolean functions which require more careful counting of embedded rectangles using
probabilistic argument.

Joint work with Paul Beame, Mike Saks and Erik Vee.

Monday, October 28, 2002

Ravindran Kannan, Yale University

Random Sub-Problems of a Given Problem
Abstract:
We consider a class of problems typcial of which is the MAX-r-SAT problem of satisfying as many clauses as possible
among given clauses with r literals each (r fixed). Our main result that if we pick at random a small subset of the
variables, the answer to the sub-problem consisting of only the clauses involving the picked variables gives us an
estimate of the answer to the whole problem.

Our methods are in a sense purely "linear algebraic". Our starting point is the formulation of the problem by means
of r dimensional arrays of reals and a simple approximation of these arrays by the sum of "rank 1" arrays. A central
result we prove and use is an upper bound on a certain norm of a random sub-array of an r dimensional array in terms
of the same norm of the whole array. The norm is chosen to model these combinatorial problems.

Joint work with N. Alon, F.W. dela Vega and M. Karpinski

Tuesday, October 29, 2002

Manindra Agarwal, IIT Kanpur, India

Derandomizing Special Polynomial Identities via Cyclotomic Rings
Abstract:

It is well known that any succinctly represented polynomial identity can be verified in randomized polynomial time.
One of the algorithms for this randomly selects a low degree polynomial and verifies the identity modulo the random
polynomial. A possible ways of derandomizing this algorithm is to construct a sample space of a "few" low degree
polynomials. Polynomials of the form x^r-1 (r small) are very good candidates for such a space since they have
excellent properties. Indeed, in the recent AKS primality testing algorithm, such a sample space was used (with a slight
generalization). In this talk, we discuss the AKS primality algorithm with this view. We also discuss some other identities
for which this approach may be useful.

Monday, November 4, 2002

Assaf Naor, Microsoft Research

Non-Linear Versions of Dvoretzky's Theorem
Abstract:
A classical theorem due to A. Dvoretzky states that for every $D>1$, any n-dimensional normed space contains
an $\Omega_D(\log n)$ dimensional subspace which is D-isomorphic to a Hilbert space.
In this talk we will discuss the following metric version of Dvoretzky's
theorem: Given an integer n and $D\ge 1$, what is the maximal k such that any n-point metric space
contains a k-point subset which may be embedded with distortion D in Hilbert space. We present an

asymptotic calculation of k (as a function of n and D), and show in particular that the behavior
of k exhibits a clear phase transition at $D=2$. We will also discuss the analogous problem in several
particular examples such as the discrete cube and expander graphs.

Tuesday, November 5, 2002

Amit Chakrabarti, IAS

A Lower Bound for Approximate Nearest Neighbor Searching

Abstract:
In the Approximate Nearest Neighbor Searching problem (ANNS), we are required to preprocess a database S of points
from a metric space so that, given a query point q, we can quickly find a point x in S such that dist(q,x) is within
a small factor of dist(q,S).

We consider this problem over the d-dimensional Hamming cube and establish a lower bound of Omega(log log d / log log log d)
on the query time of any deterministic algorithm that uses polynomial storage. This holds for approximation factors as large
as 2^{(log d)^{1 - epsilon}} for any fixed positive epsilon.

In this talk, we shall prove this result in detail. The proof involves a novel combinatorial construction inside the
Hamming cube inspired by the techniques of Ajtai in his work on the predecessor problem. Joint work with Bernard Chazelle,
Benjamin Gum, and Alexey Lvov.

Tuesday, November 11, 2002

Vijay V. Vazirani, Georgia Tech

How Intractable is the "Invisible Hand": Polynomial Time Algorithms for Market Equilibria

Abstract:

Although the study of market equilibria has occupied center stage within Mathematical Economics for over a century,
polynomial time algorithms for such questions have so far evaded researchers. It has been customary to relegate such
efficiency issues to the ``Invisible Hand'' of the market, a notion propounded by Adam Smith (Wealth of Nations, 1776).

In the context of Algorithmic Game Theory, a nascent area attempting to address new issues arising from the Internet,
we provide the first polynomial time algorithm for the linear version of a problem defined by Irving Fisher in 1891.

Fisher's original problem was defined for concave utility functions, which model the fact that buyers get satiated
with goods. Along these lines, we give a different generalization of the linear case and extend the algorithm to to it.
Our algorithms are modeled after Kuhn's primal-dual algorithm for bipartite matching.

(First result joint with Devanur, Papadimitriou, and Saberi, and available at http://www.cc.gatech.edu/fac/Vijay.Vazirani)

Monday, November 25, 2002

Christian Borgs, Microsoft Research

Erdos-Renyi Scaling for the n-Cube and Beyond
Abstract:
It is well-know that the random graph G_{n,p} has a scaling window of width n^{-1/3} inside of which the largest
component has size of order n^{2/3}. In this talk, I formulate conditions under which general finite graphs exhibit
analogous scaling.

Our techniques are a combination of methods from mathematical physics, in particular the lace expansion, and methods
from combinatorics. The key ingredient is the tight relationship between critical exponents and finite-size scaling.

This work is in collaboration with J. Chayes, R. van der Hofstad, G. Slade, and J. Spencer.

Monday, November 25, 2002

Michael Langberg, Weizmann Institute

Graphs with Tiny Vector Chromatic Numbers and Huge Chromatic Numbers
Abstract:
Karger, Motwani and Sudan (JACM 1998) introduced the notion of a vector coloring of a graph. In particular they show
that every k-colorable
graph is also vector k-colorable, and that for constant k, graphs
that are vector k-colorable can be colored by roughly Delta^{1 - 2/k}
colors. Here "Delta" is the maximum degree in the graph. Their results
play a major role in the best approximation algorithms for coloring and

for maximal independent set.

We show that for every positive integer k there are graphs that are vector k-colorable but do not have independent sets significantly larger than n/Delta^{1 - 2/k} (and hence cannot be colored with significantly less that Delta^{1 - 2/k} colors). For k = O(log(n)/loglog(n)) we show vector k-colorable graphs that do not have independent sets of size (log(n))^c, for some constant c. This shows that the vector chromatic number does not approximate the chromatic number within factors better than n/polylog(n).

As part of our proof, we analyze ``property testing'' algorithms that distinguish between graphs that have an independent set of size n/k, and graphs that are ``far'' from having such an independent set. Our bounds on the sample size improve previous bounds of Goldreich, Goldwasser and Ron (JACM 1998) for this problem.

Joint work with U. Feige and G. Schechtman.

Monday, November 26, 2002

Luke Pebody, IAS

How Combinatorial Reconstruction Using Invariant Polynomials
Abstract:
A reconstruction problem studied by Alon, Caro, Krasikov and Roditty takes a group action G:X and asks for the smallest integer k such that all subsets S of X are specified (up to G-isomorphism) by the G-isomorphism classes of all the subsets of S of size at most k.

The authors looked at this problem in many different cases, including the case of a finite cyclic group acting on itself (the so-called "necklace" problem). They showed that the answer to the problem was at most of the order of the logarithm of the size of G.

Looking specifically at the necklace problem, Radcliffe and Scott were able to greatly improve upon this bound, proving that for the cyclic group of n elements, k was at most 9 times the number of distinct prime factors of n, and was at most 3 if k was prime.

Here, I will demonstrate an improvement of both these results, that the reconstruction number for the necklace problem is at most 6. Further, I will completely evaluate the reconstruction number for all abelian groups of arbitrary cardinality. Further I will point directions for further development.

Monday, December 9, 2002

Leonid A. Levin, Boston University

Forbidden Information

Abstract:

There appears to be a loophole in Goedel Incompleteness Theorem. Closing this loophole does not seem obvious and involves
Kolmogorov complexity. (This is unrelated to, well studied before, complexity quantifications of the usual Goedel effects.)
Similar problems and answers apply to other unsolvability results for tasks where required solutions are not unique, such as,
e.g., non-recursive tilings.

D.Hilbert asked if the formal arithmetic can be consistently extended to a complete theory. The question was somewhat vague
since an obvious answer was `yes': just add to the axioms of Peano Arithmetic
(PA) a maximal consistent set, clearly existing albeit hard to find. K.Goedel formalized this question as existence among
such extensions of recursively enumerable ones and gave it a negative answer (apparently, never accepted by Hilbert). Its
mathematical essence is the lack of total recursive extensions of universal partial recursive predicate.

As is well known, the absence of algorithmic solutions is no obstacle when the requirements do not make a solution unique.
A notable example is generating strings of linear Kolmogorov complexity, e.g., those that cannot be compressed to half their
length. Algorithms fail, but a set of dice does a perfect job!

Thus, while r.e. sets of axioms cannot complete PA,
the possibility of completion by other simple means remained open. In fact, it is easy to construct an r.e. theory R that,
like PA, allows no consistent completion with r.e. axiom sets. Yet, it allows a recursive set of PAIRS of axioms such that
random choice of one in each pair assures such completion with probability 99%.

The reference to randomized algorithms seems rather special. However, the impossibility of a task can be formulated more
generically. In 1965 Kolmogorov defined a concept of Mutual Information in two finite strings. It can be refined and extended
to infinite sequences, so that it satisfies conservation laws: cannot be increased by deterministic algorithms or in random
processes or with any combinations of both. In fact, it seems reasonable to assume that no physically realizable process
can increase information about a specific sequence.

In this framework one can ask if the Hilbert-Goedel task of a consistent completion of a formal system is really possible for
PA, as it is for an artificial system R just mentioned. A negative answer follows from the

existence of a specific sequence
that has infinite mutual information with ALL total extensions of a universal partial recursive predicate. It plays a role
of a password: no substantial information about it can be guessed, no matter what methods are allowed. This "robust" version
of Incompleteness Theorem is, however, trickier to prove than the old one.

[The full article appears in FOCS-2002; A preprint is also available at
http://arXiv.org/abs/cs.CC/0203029 .]

Tuesday, December 10, 2002

Michael Elkin, IAS

Inapproximability and Instance Complexity of the Distributed Minimum Spanning Tree Problem

Abstract:

We study the **minimum-weight
spanning tree** (MST) problem in the **distributed** model of computation. In this model each vertex of the input n-vertex
graph G hosts a processor, and each processor has infinite computational power, but its initial knowledge of the graph is
very limited. Specifically, each vertex v knows the identities of its neighbors, and the weights of the edges that are incident
to v. The vertices are allowed to communicate through the edges of the graph. The communication is synchronous, and occurs
in **discrete rounds**. The goal is to minimize the number of rounds of distributed computation, which is henceforth referred
as the **running time**.

The MST problem is one of the most fundamental and well-studied problems in the area of distributed computing. Nearly matching
upper and lower bounds are known for this problem. We improve both upper and lower bounds, and furthermore, we initiate the
study of the **approximability** of **distributed** problems. Specifically, we show that for any $0 < \eps < 1$ computing
$n^{1-\eps}$-approximate MST requires $\Omega(n^{\eps/2})$ time even on graphs of small unweighted diameter. Denoting
the approximation ratio by H and the time complexity by T, the result can be interpreted as a lower bound on the
time-approximation tradeoff, $T^2 \cdot H = \Omega(n)$. We also generalize this tradeoff to the MST problem restricted
to graphs with diameter at most (a possibly constant) Lambda (the exponent of T becomes roughly $2 + 1/Lambda$ instead of 2).

We also study the **instance complexity** of the MST problem, and identify an **explicit** graph parameter, that we call
MST-diameter, that reflects (up to a small constant factor) the instance complexity of the problem. In particular, we
devise a **nearly instance-optimal algorithm for the MST problem, and show a **nearly tight lower bound** on the possible
running time of **any** correct algorithm. The running time of our algorithm matches (up to the same factor) the lower
bound on **every instance** of the MST problem.

The presentation will be based on a very recent manuscript called "Inapproximability and Instance Complexity of the
Distributed Minimum Spanning Tree Problem", authored by the speaker.

Monday, December 16, 2002

Eli Ben-Sasson, Harvard University

Derandomizing Low Degree Tests via Epsilon-Biased Spaces

Abstract:
We present the first explicit construction of Locally Testable Codes
(LTCs) of fixed constant query complexity which have almost-linear (= n^{1+o(1)}) size. Such objects were recently shown
to exist (nonconstructively) by Goldreich and Sudan.

The key to this construction is a near-optimal derandomization of the low degree test of Rubinfeld and Sudan. The original
test uses a random line in the given vector space. The number of such lines is quadratic in the size of the space, which
implied a similar blow up in previous constructions of LTCs. Goldreich and Sudan show that there exists a nearly linear
sized sample space of lines such that running the low-degree test on a random line from this collection is a good test.
We give an explicit sample space with this property.

In a similar way we give a near-optimal derandomization of the Blum Luby Rubinfeld linearity test (which is used, for
instance, in locally testing the Hadamard code).

The explicit constructions use $\eps$-biased sets for vector spaces over finite fields. The sample space consists of
the lines defined by the edges of the Cayley expander graph generated by the $\eps$-biased set. The analysis of both
tests heavily relies on the relation between the algebraic structure of the expander used and that of the code being tested.

Joint work with Madhu Sudan, Salil Vadhan and Avi Wigderson

Tuesday, December 17, 2002

Daniel Rockmore, IAS

Path Algebras for FFTs on Groups

Abstract:

This talk will focus on the "separation of variables" approach
to computing Fast Fourier Transforms (FFTs) for groups. This technique, which makes use of a path algebra indexing formalism
has yielded the most efficient FFT algorithms for many classes of finite groups, including symmetric groups and their wreath
products, matrix groups over finite fields, and others. This
is mostly joint work with David Maslen.

Monday, January 13, 2003

Amir Shpilka, Harvard University and MIT

Lower Bounds for Matrix Multiplication

Abstract:
We prove lower bounds on the number of product gates in bilinear and quadratic circuits that compute the product of two nxn
matrices over finite fields. In particular we obtain the following results:

1. We show that the number of product gates in any bilinear circuit that computes the product of two nxn matrices over GF(2)
is at least 3n^2 - o(n^2)$.

2. We show that the number of product gates in any bilinear circuit that computes the product of two nxn matrices over GF(q)
is at least (2.5 + {1.5}/{q^3 -1})n^2 -o(n^2).

These results improve the former results of Bshouty ('89) and Blaser ('99) who proved lower bounds of 2.5n^2 - o(n^2).

Tuesday, January 14, 2003

Oded Regev, IAS

New Lattice Based Cryptographic Constructions

Abstract:
We introduce the use of methods from harmonic analysis as an integral part of a lattice based construction. The tools we
develop provide an elegant description of certain Gaussian distributions around lattice points. Our results include two
cryptographic constructions which are based on the worst-case hardness of the unique shortest vector problem. The main
result is a new public key cryptosystem whose security guarantee is considerably stronger than previous results
(O(n^{1.5}) instead of O(n^7)). This provides the first alternative to Ajtai and Dwork's original 1996 cryptosystem.
Our second result is a collision resistant hash function which, apart from improving the security in terms of the unique
shortest vector problem, is also the first example of an analysis which is not based on Ajtai's iterative step. Surprisingly,
the two results are derived from the same tool which presents two indistinguishable distributions on the segment [0,1). It
seems that this tool can have further applications and as an example we mention how it can be used to solve an open problem
related to quantum computation.

Tuesday, January 20, 2003

Peter Winkler, Bell Labs

a Second Threshold for the Hard-Core Model

Abstract:
The threshold of greatest interest in statistical physics models has been the cutoff determined by uniqueness of Gibbs measure.
For a number of models, however, especially on the Bethe lattice (Cayley tree), a second threshold of considerable interest is
hit when the unique "nice" Gibbs measure ceases to be extremal.

In a sense, the first threshold tells you whether boundary information *can* exert long-range influence; the second, whether
boundary information actually *does* so. We will discuss this second threshold for the Ising and Potts models leading to a
recent result (with Graham Brightwell) for the hard-core model, i.e. for random independent sets in a Cayley tree.

Tuesday, January 21, 2003

Michael Elkin, IAS

Inapproximability of the Distributed Minimum Spanning Tree Problem

Abstract:
We study the **minimum-weight spanning tree**
(MST) problem in the **distributed** model of computation.
In this model each vertex of the input n-vertex graph G
hosts a processor, and each processor has infinite computational power,
but its initial knowledge of the graph is very limited. Specifically, each vertex v knows the identities of its neighbors,
and the weights of the edges that are incident to v. The vertices are allowed to communicate through the edges of the graph.
The communication is synchronous, and occurs in **discrete rounds**. The goal is to minimize the number of rounds of distributed
computation, which is henceforth referred as the **running time**.

The MST problem is one of the most fundamental and well-studied
problems in the area of distributed computing.
Nearly matching upper and lower bounds are known for this problem. We improve both upper and lower bounds, and furthermore,
we initiate the study of the **approximability** of **distributed** problems. Specifically, we show that for any $0 < \eps < 1$
computing
$n^{1-\eps}$-approximate MST requires $\Omega(n^{\eps/2})$ time even on graphs of small unweighted diameter. Denoting
the approximation ratio by H and the time complexity by T, the result can be interpreted as an **unconditional lower bound**
on the
time-approximation tradeoff, $T^2 \cdot H = \Omega(n)$.
We also generalize this tradeoff to the MST problem restricted to graphs
with diameter at most (a possibly constant) Λ
(the exponent of T becomes roughly $2 + 1/\Lambda$ instead of 2).

The presentation is based on a very recent manuscript called "Inapproximability and Every-Case Complexity of the Distributed
Minimum Spanning Tree Problem". Although this talk is a continuation of the talk from Dec. 10, it will be self-contained. The
focus of this talk will be on
inapproximability, while the focus of the previous talk was on every-case complexity.

Monday, January 27, 2003

Peter Keevash, Princeton University

The Exact Turan Number of the Fano Plane

Abstract:
The Fano plane is the unique hypergraph with 7 triples on 7 vertices in which every pair of vertices is contained in a unique
triple. Its edges correspond to the lines of the projective plane over the field with two elements. The Turan problem is to find the maximum number
of edges in a 3-uniform hypergraph on n vertices not containing a Fano plane.

Noting that the Fano plane is not 2-colourable, but becomes so if one deletes an edge, a natural candidate is the largest
2-colourable 3-uniform hypergraph on n vertices. This is obtained by partitioning the vertices into two parts, of sizes differing by at most one, and taking all the
triples which intersect both of them. Denote this hypergraph by $H_2(n)$.

We show that for sufficiently large n, the unique largest 3-uniform hypergraph on n vertices not containing a Fano plane
is $H_2(n)$, thus proving a conjecture of V. Sos raised in 1976.

This is joint work with Benny Sudakov.

Tuesday, January 28, 2003

Paul Seymour, Princeton University

Perfect Graphs

Abstract:
A graph G is perfect if for every induced subgraph H of G, the chromatic number and clique number of H are equal; that is,
if the largest clique of H has k vertices then H can be coloured using only k colours. It is known that perfect graphs contain
many special classes of interest, and in general have pretty properties; for instance, Lovasz showed in 1972 that the complement
of any perfect graph is also perfect. In 1961, Claude Berge proposed the so-called "strong perfect graph conjecture", that a
graph G is perfect if and only if no induced subgraph is an odd cycle of length >4 or the complement of one. In joint work with
Maria Chudnovsky, Neil Robertson and Robin Thomas, we have at last proved Berge's conjecture.

Another open question about these graphs was, how can one test in polynomial time whether a graph is perfect? This we have also
solved, in joint work with Chudnovsky. This talk will sketch both results.

Monday, February 3, 2003

Janos Pach, city College, NY and Renyi Institute, Budapest

The Number of Directions Determined by n Points in Space

Abstract:
Erdos pointed out the following immediate consequence of
the celebrated Gallai-Sylvester theorem on ordinary lines:
n non-collinear points in the plane determine at least n different connecting lines. Equality is attained
if and only if all but one of the points are collinear.

In the same spirit, Scott posed two similar questions in
1970:

1. Is it true that the number of different
   directions assumed by the connecting lines of n > 3 non-collinear points in the plane is at least n-1?

Is it true that the number of different
directions assumed by the connecting lines of n > 5 non-coplanar points in 3-space is at least 2n-3?
The first question was answered in the affirmative
by Ungar in 1982, using allowable sequences. We solve
the second problem of Scott.

Joint work with Rom Pinchasi and Micha Sharir.

Tuesday, February 4, 2003

Noga Alon, Tel Aviv University and IAS

Testing Large Directed Graphs

Abstract:
A digraph G on n vertices is epsilon-far from satisfying
a property P, if one has to add to or delete from G at least epsilon n2 directed edges to obtain a digraph satisfying P.
An epsilon-tester for P is a randomized algorithm that,given n, an access to a digraph G on n vertices, and the ability to ask
queries of the form "is (i,j) a directed edge?" can distinguish, with high probability, between the case thatG satisfies P and
the case that G is epsilon-far from satisfying P. The tester is a one-sided tester if it does not err on graphs satisfying P.
For a fixed digraph H, let P(H) denote the property that G is H-free. We show that for every fixed H, there is an epsilon-tester

for P(H) whose query complexity is independent of n. We further characterize all digraphs H for which this complexity
(for one-sided error, or two-sided error testers) is polynomial in (1/epsilon), and show that for all of them there is a
two-sided tester sampling only O(1/epsilon) vertices, though one sided testers must sample at least (1/\epsilon)^{\Omega(d)}
vertices, where d is the average degree of H.

Joint work with Asaf Shapira.

Monday, February 10, 2003

Ehud Friedgut, Hebrew University of Jerusalem

Coloring Products of Graphs, a Fourier Approach

Abstract:
Assume that at a given road junction there are n three-position switches that control the red-yellow-green position of the
traffic light. You are told that no matter what the switch configuration is, if you change the position of every single one of
the switches then the color of the light changes. Can you prove that in fact the light is controlled by only one of the switches?
What if the above information holds for only 99.99% of the configurations? The above question deals with a special case of
coloring a graph which is a product of smaller complete graphs (in this case k_3^n). In the talk I will present some results
about independent sets and colorings of such graphs, including stability versions (see the "99.99%" question above.)
Our approach is based on Fourier analysis on Abelian groups.

Joint work with Irit Dinur.

Tuesday, February 11, 2003

Benny Sudakov, Princeton University and IAS

Set-Systems with Restricted k-wise Intersections

Abstract:
A large variety of problems and results in Extremal Set Theory are dealing with estimates of the size of a family of

sets with some restrictions on the intersections of its members. Notable examples of such results, among others, are
the celebrated theorems of Fischer, Ray-Chaudhuri- Wilson and Frankl-Wilson on set systems with restricted pairwise
intersections.

In this talk we discuss an extension of these
results when the restrictions apply to k-wise intersections, $k>2$. We obtain asymptotically tight bounds
for this problem. Our proofs combine tools from linear algebra with some combinatorial arguments.

Part of this is joint work with V. Grolmusz and part was done jointly with Z. Furedi.

Tuesday, February 18, 2003

Hartmut Klauck, IAS

Quantum Time-Space Tradeoffs for Sorting

Abstract:
We investigate the complexity of sorting in the model of sequential quantum circuits. While it is known that in general a
quantum algorithm based on comparisons alone cannot outperform classical sorting algorithms by more than a constant factor
in time complexity, this is wrong in a space bounded setting. We observe that for all storage bounds n\log n>S>\log3 n one
can devise a quantum algorithm that sorts n numbers (using comparisons only) in time T=O(n^{3/2}\log^{3/2} n/\sqrt S). We
then show the following lower bound on the time-space tradeoff for sorting n numbers from a polynomial size range in a
general sorting algorithm (not necessarily based on comparisons): TS=\Omega(n^{3/2}). Hence for small values of S the
upper bound is almost tight. Classically the time-space tradeoff for sorting is TS=\Theta(n2).

Monday, February 24, 2003

Marek Karpinski, University of Bonn

Approximation Complexity of MIN-BISECTION Problems

Abstract:
We present some recent results on the complexity of approximating optimal solutions of instances of the Minimum
Bisection and some related Partitioning problems. We formulate also some intriguing and still

open questions on that problem.

Tuesday, February 25, 2003

Alexander Razborov, IAS

Systems of Linear Equations Hard for k-DNF Resolution

Abstract:
Let A be an (m x n) (0-1) matrix, and b be a fixed (0-1) vector of length n such that the system of linear equations AX=b
over GF[2] is inconsistent. We are interested in showing that for some class of matrices A this fact of inconsistency is
hard to prove in the propositional logic (all results of this kind known so far require that the bipartite graph associated
with A has certain expansion-like properties). Whenever we have such a hardness result for some propositional proof system
P, we can show, via a chain of reductions, that P also does not possess efficient proofs of super-polynomial circuit lower
bounds, and, in particular, does not prove efficiently that $NP\not\subseteq P/poly$.

We solve this problem for the proof system $Res(\epsilon\log n)$ operating with ($\epsilon\log n$)-DNF (ϵ a small
constant), and in the talk we will elaborate a little bit on the above mentioned application, as well as give some proof ideas
of this result.

Monday, March 3, 2003

Ryan O'Donnell, MIT

Coin Flipping From a Cosmic Source; or, On Error Correction of Truly Random Bits

Abstract:
We study a new problem related to collective coin flipping, coding theory, and noise sensitivity, with motivations from
cryptography.

Suppose there is a "cosmic" source x of n uniformly random bits, and k non-communicating players who want to use
these bits to agree on a shared random coin toss. Further assume the players only receive ϵ-corrupted versions
of x; that is, each player i independently sees only the string y^i, where y^i is given by flipping each bit of

x independently with probability ϵ. The players want to preselect \emph{balanced} functions
$f_i : \{0,1\}^n \to \{0,1\}$ such that $\Pr[f_1(y^1) = \cdots f_k(y^k)]$ is maximized. That is, each player
wants to toss a fair coin, and the players want to maximize the probability they all agree. The functions f_i
may be thought of as an error correcting procedure for the source x.

When $k=2,3$ no error correction is possible, as the optimal protocol is given by having all players use the first-bit
function $f(y) = y_1$. On the other hand, for large values of k better protocols exist. We study general properties
of the optimal protocols and the asymptotic behavior of the problem with respect to k, n, and ϵ. Several
interesting and unexpected open problems are uncovered.

Our analysis uses tools from probability, Fourier analysis, convexity, isoperimetry, and discrete symmetrization.

Tuesday, March 4, 2003

Xiaodong Sun, IAS

Exponential Lower Bound for 2-Query Locally Decodable Codes via a Quantum Argument

(on the paper by Iordanis Kerenidis and Ronald de Wolf)

Abstract:
A locally decodable code (LDC) encodes n-bit strings x in m-bit
codewords C(x), in such a way that one can recover any bit x_i from a corrupted codeword by querying only a few bits
of that word. Locally decodable codes are interesting due to its relation to PCP theory and private information retrieval(PIR).

For the first result, the authors use a quantum argument to prove that LDCs with 2 classical queries need exponential length:
m=2^{Omega(n)} by showing that a 2-query LDC can be decoded with only 1 quantum query, and then proves an exponential lower
bound for such 1-query locally quantum-decodable codes. This gives new classical lower bounds for the setting of private
information retrieval. This is one of the few results that use quantum argument to prove lower bounds for classical computation.

In addition to that, the authors show that quantum 2-server PIR is stronger than the classical counterpart. In particular,
they exhibit a quantum 2-server PIR scheme with O(n^{3/10}) qubits of communication, improving upon the O(n^{1/3}) bits of

communication of the best known classical 2-server PIR.

Monday, March 10, 2003

Rocco Servedio, Columbia University

Learning Juntas

Abstract:
In this talk we consider a fundamental (and seemingly quite difficult) learning problem: learning an arbitrary Boolean
function which depends on an unknown set of k out of n Boolean variables. This problem was posed by Avrim Blum in 1994 and
is an important special case of several notorious open problems such as learning decision trees or DNF formulas.

We give an algorithm for learning such functions given only a source of random uniformly distributed labeled examples.
Our algorithm achieves the first polynomial factor improvement on a naive n^k time bound (which can be achieved via
exhaustive search). The algorithm and its analysis are based on new structural properties which hold for all Boolean
functions and may be of independent interest.

Joint work with Elchanan Mossel (UC Berkeley) and Ryan O'Donnell (MIT).

Tuesday, March 11, 2003

Van Vu, University of California at San Diego

Long Arithmetic Progressions in Sumsets and Erd&#65533;s-Folkman Conjecture

Abstract:
Let A be a (finite or infinite) set of positive integers, S(A) denotes the collection of finite subset sums of A.
For instance, if A={1,2, 9}, then S(A)={1,2,3, 9, 10,11, 12}. Recently, we proved Theorem 1: There is a constant c
such that for any set A containing cn^{1/2} different integers between 1 and n, S(A) contains an arithmetic progression
of length n. Using this theorem, we confirmed a long standing conjecture of Erdos and Folkman, posed in the sixties.

Conjecture 2: There is a constant c such that for any increasing sequence A of density cn^{1/2}, S(A) contains an
infinite arithmetic progression.

The proof of Theorem 1 introduces a new method for proving the existence of long arithmetic progressions. In this
talk we shall discuss some essential points of this proof.

Joint work with E. Szemeredi.

Monday, March 17, 2003

Alexander Soifer, DIMACS, Rutgers University and University of Colorado

Chromatic Number of the Plane and its Relatives: History, Problems, Results

Abstract:
Define Unit Distance Plane U2 as a graph on the set of all points of the plane R2 as its vertex set, with two vertices adjacent iff they are distance 1 apart. The chromatic number x of U2 is called chromatic number of the plane. Surprisingly, the past 52 years have not brought about any improvement in general case: x = 4, or 5, or 6, or 7. The problem CNP of finding chromatic number of the plane has been credited in print to P. Erdos, M. Gardner, H. Hadwiger, L. Moser and E. Nelson. Who created CNP?

A plane set S is said to realize distance d if S contains two points distance d apart. In 1958 Paul Erdos posed the following problem: what is the smallest number of colors sufficient for coloring the plane in such a way that no color realizes all distances? D. Raiskii (1970) and D. Woodall (1973) obtained striking results. A gap between this problem and CNP was filled by a notion of almost chromatic number of the plane, i.e., the minimal number of colors sufficient for coloring the plane in such a way that all but one color forbid distance 1, and the remaining color forbids a distance [Soifer, 1992]. A continuum of 6-colorings of the plane has been constructed, in which first five colors forbid distance 1 and the 6th color forbids a distance [Soifer, 1994].

In 1976 P. Erdo"s asked whether forbidding 3-cycles would limit chromatic number of a unit distance graph to 3. Three years later N. Wormald constructed an example of a triangle-free unit distance 4-chromatic graph. His huge graph had 6448 vertices. The problem got a new life when I presented these questions in 1991, and a few young colleagues entered the race for the smallest triangle-free unit distance 4-chromatic graph. Many new "world records" were set by K. Chilakamarri, R. Hochberg and P. O'Donnell on the pages (and covers) of Geombinatorics.
The talk is dedicated to the memory of Paul Erdos on occasion of his 90th birthday.

Monday, March 18, 2003

Amit Chakrabarti, IAS

Lower Bounds for Multi-Party Set Disjointness

Abstract:
In the multi-party set disjointness problem, we have t >= 2 players, each of whom holds a subset of {1,...,n}. The players
would like to determine (via a randomised communication protocol) whether all t sets have a common element.

For t = 2, a classic result in communication complexity says that such a protocol must communicate at least Omega(n) bits.
We show a lower bound of Omega(n/(t log t)) for general t. Our bound also holds for a very restricted promise version of
the disjointness problem and is essentially optimal for this restriction; the restriction has a connection with the space
complexity of approximating frequency moments in the data stream model. Our bound improves to an optimal Omega(n/t) if the
communication model is restricted to be one-way (the connection with the data stream model still remains).

The above lower bounds are proven using the information complexity technique, recently introduced by Chakrabarti, Shi, Wirth,
and Yao, and generalised by Bar-Yossef, Jayram, Kumar, and Sivakumar. In the talk, I shall describe the information complexity
technique, and give an outline of the proofs of our results. This is joint work with Subhash Khot and Xiaodong Sun.

Monday, March 24, 2003

Daniele Micciancio, University of California, San Diego

Generalized Compact Knapsacks, Cyclic Lattices, and Efficient One-Way Functions from Worst-Case Complexity Assumptions

Abstract:
We study a generalization of the compact knapsack problem for arbitrary rings: given m = O(log n) ring elements
a_1,...,a_m in R and a target value b in R, find coefficients x_1,...,x_m in X (where X is a subset of R of size 2^n)
such that sum(a_i x_i) = b. The computational complexity of this problem depends on the choice of the ring R and set of
coefficients X. This problem is known to be solvable in quasi polynomial time when R is the ring of the integers and X is
the set of small integers {0,...,2^n - 1}. We show that if R is an appropriately chosen ring of modular polynomials and X

is the subset of polynomials with small coefficients, then the compact knapsack problem is as hard to solve on the average
as the worst case instance of approximating the covering radius (or the length of the shortest vector, or various other well
known lattice problems) of any cyclic lattice within a polynomial factor. Our proof adapts, to the cyclic lattice setting,
techniques initially developed by Ajtai for the case of general lattices.

Monday, March 24, 2003

Laslo Lovasz, Microsoft Reseaerch

Discrete Analytic Functions and Global Information from Local Observation

Abstract:
We observe a certain random process on a graph "locally", i.e., in the neighborhood of a node, and would like to derive information about "global" properties of the graph. For example, what can we know about a graph based on observing the returns of a random walk to a given node? This can be considered as a discrete version of "Can you hear the shape of a drum?"

Our main result concerns a graph embedded in an orientable surface with genus g, and a process, consisting of random excitations of edges and random balancing around nodes and faces. It is shown that by observing the process locally in a "small" neighborhood of any node sufficiently (but only polynomially) long, we can determine the genus of the surface. The result depends on the notion of "discrete analytic functions" on graphs embedded in a surface, and extensions of basic results on analytic functions to such discrete objects; one of these is the fact that such functions are determined by their values in a "small" neighborhood of any node.

This is joint work with Itai Benjamini.

Tuesday, March 25, 2003

Madhu Sudan, MIT

Algebraic Constraint Satisfaction Problems

Abstract:
The theory of probabilistically checkable proofs (PCPs) has introduced a lot of algebraic tools and techniques that are

potentially of general interest, but are often buried too deep inside proofs to be readily visible to a novice. In this
talk I'll try to motivate a class of problems called Algebraic Constraint Satisfaction Problems and show how many of the
important ingredients in the construction of probabilistically checkable proofs are based on more accessible results about
algebraic constraint satisfaction problems.

Based on joint work with Eli Ben-Sasson and Prahladh Harsha

Tuesday, March 25, 2003

Luca Trevisan, Berkeley University

List-Decoding Using the XOR Lemma

Abstract:

We show that Yao's XOR Lemma and its essentially equivalent rephrasing as a "Concatenation Lemma," when properly
re-interpreted, can be seen as a way of obtaining error-correcting codes with good list-decoding algorithms from
error-correcting codes having weak unique-decoding algorithms. To get codes with good rate and efficient list
decoding algorithms one needs a proof of the Concatenation Lemma that is, respectively, "very derandomized," and
that uses very small advice.

We show how to reduce advice in Impagliazzo's proof of the Concatenation Lemma for pairwise independent inputs,
which leads to error-correcting codes with O(n2 polylog n) encoding length and encoding time, and probabilistic
O(n polylog n) list-decoding time. (Note that the decoding time is sub-linear in the length of the encoding.)

Back to complexity theory, our small-advice proof of Impagliazzo's ``hard-core set'' results yields a (weak)
uniform version of O'Donnell results on amplification of hardness in NP. We show that if there is a problem in
NP that cannot be solved by BPP algorithms on more than a 1-1/(log n)^c fraction of inputs, then there is a
problem in NP that cannot be solved by BPP algorithms on more than a 3/4 + 1/\poly(n) fraction of inputs,
where c>0 is an absolute constant.

Monday, March 31, 2003

Bo Brinkman, Princeton University

On the Impossibility of Dimension Reduction in l1

Abstract:
The Johnson-Lindenstrauss Lemma shows that any n points in Euclidean space (with distances measured by the ℓ_2 norm) may be mapped down to $O((\log n)/\epsilon^2)$ dimensions such that no pairwise distance is distorted by more than a $(1+\epsilon)$ factor. Determining whether such dimension reduction is possible in ℓ_1 has been an intriguing open question. Charikar and Sahai recently showed lower bounds for dimension reduction in ℓ_1 that can be achieved by linear projections, and positive results for shortest path metrics of restricted graph families. However the question of general dimension reduction in ℓ_1 was still open. For example, it was not known whether it is possible to reduce the number of dimensions to $O(\log n)$ with $1+\epsilon$ distortion.

We show strong lower bounds for general dimension reduction in ℓ_1. We give an explicit family of n points in ℓ_1 such that any embedding with distortion δ requires $n^{\Omega(1/\delta^2)}$ dimensions. This proves that there is no analog of the Johnson-Lindenstrauss Lemma for ℓ_1; in fact embedding with any constant distortion requires n^ϵ dimensions. Our proof establishes this lower bound for shortest path metrics of series-parallel graphs.

Joint work with Moses Charikar

Tuesday, April 1, 2003

Omer Reingold, IAS

Extractors - Optimal Up To Constant Factors

Abstract:

Randomness "extractors" are functions that extract almost-uniform bits from sources of biased and correlated bits, using a small number of additional uniform bits (known as the "seed") as a catalyst. Extractors play a fundamental role in the theory of pseudorandomness and have a wide variety of applications. Thus coming up with explicit constructions has been the focus of a large body of work over the past decade.

In this talk, we will describe a new construction of extractors from joint work with Chi-Jen Lu, Salil Vadhan, and Avi Wigderson (to appear in STOC

03). These extractors can extract any constant fraction of the min-entropy from an n-bit source, using a seed of length O(log n). This is the first explicit construction of extractors that works for all min-entropies and is simultaneously optimal up to constant factors in both the seed length and output length.

Monday, April 7, 2003

Bela Bollobas, University of Memphis and University of Cambridge

Scale-free Random Graphs

Abstract:
In 1998, Watts and Strogatz observed that many large-scale real-world networks, including neural networks, power grids, collaboration graphs, and the internet, have numerous common features that resemble properties of random graphs.
It was also realized that the standard mean-field and lattice-based random graphs are not appropriate models of these large-scale networks, so we should look for other classes of random graphs. One of the main features demanded of these new random graphs is that they should be scale-free. The first such model was introduced by Barabasi and Albert in 1999; by now, numerous models of scale-free random graphs have been proposed and studied, mostly by computer simulations and heuristic analysis.

In the talk we shall review a number of these models, and present several recent rigorous results obtained jointly with Oliver Riordan.

Monday, April 14, 2003

Michael Krivelevich, Tel Aviv University, Israel

Adding Random Edges to Dense Graphs

Abstract:
In this talk I will discuss a novel model of random graphs. In this model, a large graph H=(V,E) on n vertices, usually
with bounded minimum degree, is given, and a set R of m edges is chosen uniformly at random from the set of non-edges
of H to form a random graph G=(V,E+R). The question is then: how large should be the value of m to guarantee the almost
sure appearance of various monotone properties in G?

Here we treat the case where the minimum degree of H is linear in n. We consider such properties as existence of a
fixed size clique, diameter, k-connectivity, and Ramsey properties, and obtain tight results for

most of the problems.

A joint work with Tom Bohman, Alan Frieze, Ryan Martin (Carnegie Mellon University), and with Prasad Tetali (GeorgiaTech).

Monday, April 21, 2003

Muli Safra, Tel Aviv University

Analysis of Boolean Functions and Various Applications

Abstract:
Representing a Boolean function as a polynomial is only natural. It turns out that this representation, along with some
related technology -- for example the study of the Influence of variables on Boolean functions -- gives insight to many a
spects of such functions. This field was founded in a paper by Kahn, Kalai and Linial from '89, and has since shown
applications in a wide array of fields, including Game Theory and Social Choice, Economics, Percolation, and Complexity theory.

The talk will survey the methodology and some of its applications, to Mechanism Design, Graph Properties and Complexity Theory.
We would then consider some further applications, show the state of art in terms of known results in the field; and suggest
open problems with their relevant applications.

Monday, April 28, 2003

Bruce Reed, McGill University

Partial Results on the Total Colouring Conjecture

Abstract:
A total colouring of a graph is a colouring of its vertices
and edges so that no two incident edges receive the same colour, no two adjacent vertices receive the same colour,
and no edge gets the same colour as one of its endpoints. Clearly, if a graph has maximum degree Δ then every
total colouring requires $\Delta+1$ colours. In the 1960s, Behzad and Vizing independently conjectured that $\Delta+2$
colours always suffice to totally colour such a graph. We provide evidence for this conjecture and discuss some related
problems.

Monday, May 5, 2003

Leonid Khachiyan, Rutgers University

Dual-Bounded Monotone Properties

Abstract:
Given a monotone property M over the subsets of a finite set V, we consider the problem of incrementally generating the family
F of all minimal subsets of V satisfying M. For a number of interesting monotone properties, F turns out to be uniformly
dual-bounded in the sense that for any nonempty subfamily F' of F, the number of all maximal infeasible sets for M that are
also maximal independent sets for F' can be bounded by a (quasi) polynomial in the size of F' and the length of the input
description of M. For instance, the number of maximal infeasible vectors for a monotone system of m linear inequalities in
n binary (or integer) variables does not exceed the number of minimal feasible solutions, to within a factor of nm. When the
input monotone property M can be evaluated for any subset of V in (quasi) polynomial time, the uniform dual- boundedness of
F implies that all elements of F can be generated in incremental quasi-polynomial time. Examples include the generation of
all minimal feasible integer or binary solutions to a monotone system of linear inequalities, efficient generation of minimal
infrequent item sets for a database, maximal independent sets in the intersection of any number of matroids, minimal spanning
collections of subspaces from a given list, and more generally, minimal solutions to systems of polymatroid inequalities of
polynomial range. In contrast, for all of the above examples, it is NP-hard to incrementally generate all maximal infeasible
sets for M.

Talk is based on joint papers with E. Boros, K. Elbassioni,
V. Gurvich and K. Makino.

Monday, May 12, 2003

Edith Elkind, Princeton University

Frugality in Path Auctions

Abstract:

We consider the problem of picking (buying) an inexpensive s-t path in a graph where edges are owned by independent (selfish)
agents, and the cost of an edge is known to its owner only. We focus on minimizing the buyer's total payments.

First, we show that any mechanism with (weakly) dominant strategies (or, equivalently, any truthful mechansim) for the agents
can force the buyer to make very large payments. Namely, for every such mechanism, the buyer can be forced to pay
c(P) + n/2*(c(Q)-c(P)), where c(P) is the cost of the shortest path, c(Q) is the cost of the second-shortest path, and
n is the number of edges in P. This extends the previous work of Archer and Tardos, who showed a similar lower bound for
a subclass of truthful mechanisms called min-function mechanisms. Our lower bounds have no such limitations on the mechanism.

Motivated by this lower bound, we study a wider class of mechanisms for this problem, namely ones that admit Nash equilibrium
strategies.In this class, we identify the optimal mechanism with regard to total payment. We then demonstrate a separation in
terms of average overpayments between the classical VCG mechanism and the optimal mechanism showing that under various natural
distributions of edge costs, the optimal mechanism pays at most logarithmic factor more than the actual cost, whereas VCG pays
$\sqrt{n}$ times the actual cost. On the other hand, we also show that the optimal mechanism does incur at least a constant
factor overpayment in natural distributions of edge costs. Since our mechanism is optimal, this gives a lower bound on all
mechanisms with Nash equilibria.

Joint work with Amit Sahai and Ken Steiglitz.

Monday, June 2, 2003

Graham Cormode, DIMACS

Tracking Frequent Items Dynamically

Abstract:

Most database management systems maintain statistics on the underlying relation. One of the important statistics is that of
the "hot items" in the relation: those that appear many times (most frequently, or more than some threshold). For example,
end-biased histograms keep the hot items as part of the histogram and are used in selectivity estimation. Hot items are used
as simple outliers in data mining, and in anomaly detection in networking applications.

I will describe some novel solutions to this problem based on viewing this as a group testing problem. These approaches use
both adaptive and non-adaptive group testing. The algorithms maintain a small space data structure that monitors the
transactions on the relation, and when required, quickly output all hot items, without rescanning the relation in the
database. With user-specified probability, it is able to report all hot items. I will then go on to describe some work
in progress on extensions to this model, when the goal is to find items whose frequency has changed by a significant amount,
either in absolute or relative terms; and also finding related items which together are "hot" but individually are not.

Joint work with S. Muthukrishnan

Tuesday, June 3, 2003

Hartmut Klauck, IAS

On the Rectangle Bound in Communication Complexity

Abstract:
We investigate the power of the most important lower bound technique in randomized communication complexity, which is based
on an evaluation of the maximal size of approximately monochromatic rectangles. While it is known that the 0-error version
of this bound is polynomially tight for deterministic communication, nothing in this direction is known for constant error
and randomized communication complexity. We relate this bound to Arthur Merlin communication complexity, and also give a
combinatorial characterization of the value of the lower bound method. This characterization is done by what we call a
bounded error uniform threshold cover. We then study different relaxations of bounded error uniform threshold covers
related to other lower bound methods for randomized communication complexity.

Thursday, June 5, 2003

Guy Kindler, Tel-Aviv University

Noise-Resistant Boolean Functions are Juntas

Abstract:
Let f be a Boolean functions over n Boolean variables. We say that f is noise-resistant if, roughly, when evaluated on two

random inputs which differ on a small fraction of their coordinates it yields the same values with high probability. In the
case where the distribution of each of the two random inputs is uniform, Bourgain showed that in order for a function f to be
noise-resistant, it must essentially depend on only a constant number of variables (such a function is called a junta).

In many cases, one is more interested in the behavior of a Boolean functions when its inputs are p-biased, namely when "1"
occurs in each coordinate with some probability p. We show a conceptually simple proof for the aforementioned result (with
somewhat weaker parameters), that easily extends to the case where the random inputs are distributed according to the p-biased
distribution.

Joint work with Muli Safra.

Monday, June 16, 2003

Jozsef Balogh, Ohio State University

Optimal Integer Arrangement on the Square Grid

Abstract:
Assume that we have a discrete information source uniformly and randomly emitting integers within some very long interval. We want to transmit the information over two channels in such a way that

(1) if someone receives both channels, they can exactly reconstruct the input; and

(2) if they get only one channel, they can come up with a guess as to what the integer was with reasonable accuracy.
Such situations might arise in sending information over a packet-switched
network.

We consider the smallest possible distortion achievable with two identical channels of various throughputs, with the distortion measured in the absolute or in the mean squared sense, and we discuss combinatorial

problems arising from there. An example of such a problem is:

Put all the integers into the plane square grid, using each integer exactly once, so that each row and each column contains exactly N integers. Let $d=d(N)$ be the biggest difference between numbers in the same row or column. What is the smallest possible value for d? We have solved this problem.

There are several variations of this problem,
if instead of optimizing the maximal difference we would like to optimize the maximal variance of the numbers
occurring in the same line. We get a much harder problem, where there are still many open questions left. As a byproduct we have got a "Shearer entropy"- like inequality for the Laplacians of graph products.

This is joint work with Janos Csirik and partly with Clifford Smyth (CMU).

Sponsored by:

National Science Foundation

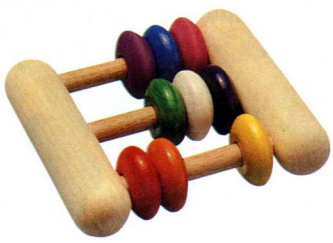

Abacus

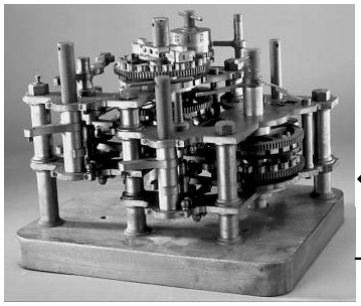

"Analytical Engine"

by Charles Babbage

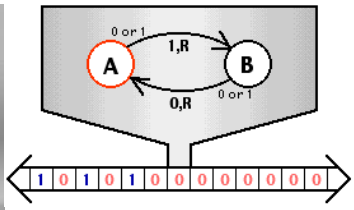

Turing Machine

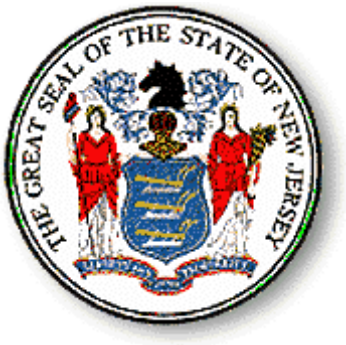

State of New Jersey