abstracts

Monday, September 22, 2003

Yehuda Lindell, IBM Watson

Impossibility Results for the Composition of Secure Two-Party Protocols
Abstract:
Until recently, the security of protocols was studied in the stand-alone model where a single protocol was executed a single time. However, this model of computation does not realistically model the security concerns of modern networks where many sets of parties run many different protocols at the same time.
Furthermore, security in the classic stand-alone model does /not/ imply security in this more complex setting of protocol composition. It is therefore of great importance to understand whether or not it is possible to obtain protocols that remain secure under composition (i.e., in such multi-execution settings) and under what assumptions.

In this talk, we present impossibility results from four recent papers relating to the concurrent composition of secure two-party computation. Our results relate to
/general composition/ (where secure protocols are run concurrently with arbitrary other protocols), /self composition/ (where a single secure protocol is run many times) and /universal composability/ (a security definition that guarantees security under general composition). In short, we provide very broad impossibility results, demonstrating that secure protocols cannot be obtained for very large classes of functionalities. Our results are for the plain model, where no trusted setup phase is assumed (and so, for example, a common reference string is not allowed). In order to obtain our results, we also show interesting (and sometimes surprising) equivalences between different notions of composition.

The talk will be self contained and prior knowledge is not assumed.

Tuesday, September 23, 2003

Boaz Barak, IAS

Lower Bounds for Non-Black-Box Zero-Knowledge
Abstract:
Even after 20 years of very fruitful research, there are still some fascinating and important

open questions about
zero-knowledge proof systems. In particular, there are still relatively few lower bounds and impossibility results
on the types of zero-knowledge systems that can exist for proving non-trivial statements. Furthermore, many of these
known lower bounds hold only for *black-box* zero-knowledge. However, recent results show that such lower bounds do
not imply corresponding lower bounds for the general (i.e., *non-black-box*) case.

In this talk I will discuss the open problems, and show new lower bounds and impossibility results for general
(i.e., non-black-box) zero-knowledge proofs and arguments. The results are that, under reasonable complexity
assumptions:

1. There does not exist a constant-round zero-knowledge *strong* proof (or argument) of knowledge (as defined by
Goldreich (2001)) for a nontrivial language.

2. There does not exist a two-round zero-knowledge *proof* system with perfect completeness for an NP-complete language.

The previous impossibility result for two-round zero knowledge, by Goldreich and Oren (J. Cryptology, 1994) was only
for the case of *auxiliary-input* zero-knowledge proofs and arguments.

3. There does not exist a constant-round public-coin *proof* system for a nontrivial language that is *resettable* zero
knowledge. This result also extends to *bounded* resettable zero knowledge.

In contrast, we show that under reasonable assumptions, there does exist such a (computationally sound) *argument* system
that is bounded-resettable zero knowledge. The complexity assumptions we use are not commonly used in cryptography. However,
in all cases, we show that assumptions like ours are necessary for the above results.

Joint work with Yehuda Lindell (IBM T.J. Watson) and Salil Vadhan (Harvard).

Monday, September 29, 2003

Michael Kearns, University of Pennsylvania

Network Models for Game Theory and Economics

Abstract:

Over the last several years, a number of authors have developed graph-theoretic or network

models for large-population
game theory and economics. In such models, each player or organization is represented by a vertex in a graph, and payoffs
and transactions are restricted to obey the topology of the graph. This allows the detailed specification of rich structure
(social, technological, organizational, political, regulatory) in strategic and economic systems.

In this talk, I will survey these models and the attendant algorithms for certain basic computations, including Nash,
correlated, and Arrow-Debreu equilibria. Connections to related topics, such as Bayesian and Markov networks for probabilistic
modeling and inference, will be discussed.

Tuesday, September 30, 2003

Farid Ablayev, Kazan State University

On the Computational Power of Classical and Quantum Branching Programs

Abstract:
We describe a model of a Quantum Branching Program (QBP) and study its computational power. We define several natural
restrictions on this model, including bounded-width and read-once QBPs.

First, we present upper and lower bounds on the power of quantum and stochastic branching programs of constant width.
We show that any language in NC^1 can be recognized with exact acceptance by a width-2 quantum branching program of
polynomial length, in contrast to the classical case where width 5 is necessary unless $NC^1=ACC$. This separates width-2
quantum programs from width-2 doubly stochastic programs as we show the latter cannot compute the middle bit of the
multiplication function. We also show that constant-width quantum and stochastic branching programs can be simulated by
classical branching programs of larger but bounded width, so the languages accepted by them are in NC^1.

For read-once QBPs, we give a symmetric Boolean function which is computable by a read-once QBP with $O(\log n)$ width,
but not by a deterministic read-once QBP with $O(n)$ width, or by a classical randomized read-once BP with $O(n)$ width
whose transitions on each level are permanent. Finally, we present a general lower bound on the width of read-once QBPs,
showing that our $O(\log n)$ upper bound for this symmetric function is almost tight.

Tuesday, October 7, 2003

Russell Impagliazzo, IAS

Memorization and DPLL: Formula Caching Proof Systems

Abstract:

The DPLL algorithm, is a simple back-tracking approach
to solving Satisfiability. Strictly speaking,
DPLL is a family of algorithms, rather than a single
algorithm, since there is a non-specified sub-routine
for picking the next variable to branch on.
Experiments show that some variants of DPLL seem
to perform quite well, but that the branching rule
is critical to success.

A successful theoretical approach to understanding DPLL has been to view the branching rule as *non-deterministic*, allowing
the algorithm to make the optimal choice at each step. Since a non-deterministic algorithm for disproving satisfiability is
equivalent to a proof system, this moves the problem into the realm of proof complexity. Researchers have then obtained lower
bounds on any version of DPLL by proving lower bounds on the corresponding *tree-like resolution* proof system.

We consider extensions of the DPLL approach that add some version of *memorization*, remembering formulas the algorithm has
previously shown unsatisfiable. Various versions of such *formula caching* algorithms have been suggested for satisfiability
and stochastic satisfiability (Majercik and Lipton; Bacchus, Dalmao, and Pitassi). We formalize this method, and characterize
the strength of various versions in terms of proof systems. These proof systems seem to be both new and simple, and have a
rich structure. We compare their strength to several studied proof systems: tree-like resolution, regular resolution, general
resolution, and $Res(k)$. We give both simulations and separations.

Memoization takes the tree-like structure of a back-tracking algorithm and makes it a DAG by combining paths with identical
sub-problems. Regular and general resolution take a tree-like proof and make it a DAG by merging paths leading to the same
clause. Our initial view was that adding memoization to DPLL would yield either regular or general resolution, or something
in between. Suprisingly, none of our systems seem to be between general and regular resolution, and for most, we have either
shown that they cannot simulate regualr resolution or that general resolution cannot simulate them.

Joint work with Paul Beame, Toniann Pitassi and Nathan Segerlind.

Monday, October 20, 2003

Grigori Mints, Stanford University

Propositional Logic of Continuous Transformations

Abstract:
Dynamical topological logic studies models of the form (X,T), where X is a topological space, T a transformation on X.
Propositional formulas are constructed from variables (atomic formulas) by Boolean connectives, necessity [] and a monadic
operation o. Variables are interpreted by subsets of X, Boolean connectives act in a natural way, [] is the interior and o
is the pre-image under operation T. Under this interpretation the axiom schema

o[] A->[] oA called (C)

expresses continuity of T. J. Davoren proved completeness of S4+C [including o(A&B)=oA&oB, o(~A)=~oA] for the class of all
topological spaces, in particular for finite spaces derived from Kripke models.

We prove completeness of S4+C for Cantor space. The proof uses a continuous and open map W from the Cantor space onto a
suitable Kripke model (K,T) [with T continuous in order topology on K] and a continuous map S on Cantor space satisfying
condition: WS=TW.

P. Kremer pointed out that the real line is not complete for S4+C. No previous knowledge of non-classical logic is assumed.

Tuesday, October 21, 2003

Eyal Rozenman, The Hebrew University

A New Explicit Construction of Constant-Degree Expander Cayley Draphs

Abstract:
Expander graphs are widely used in Computer Science and Mathematics. A graph is expanding if the second eigenvalue of the standard random walk on this graph is bounded away from 1 (equivalently, the smallest eigenvalue of the Laplacian is strictly larger than 0).

Several explicit constructions of infinite families of constant degree expanders are Cayley graphs, namely graphs defined by groups and generators. All these constructions start with a single infinite "mother" group, plus a special finite set of generators, and constructs

the infinite family by taking larger and rager finite quotients of the "mother" group.

Recently, a new approach was introduced by Reingold et al (Annals '01) for explicitly constructing expander graphs. This approach starts with a single finite "seed" graph, and iteratively constructs larger and larger graphs of the same constant degree, via a new graph product (called the "zig-zag" product). They prove that if the "seed" graph is a good enough expander, so are all graphs in the infinite family.

We prove a similar theorem for a family of Cayley graphs. However, in contrast, the expansion of the "seed" graph in our family is an open question (see below).
The proof relies on the connection between zig-zag product of graphs and semi-direct product in groups of Alon et al (FOCS '01). The groups in the family are (essentially) the automorphism groups of a k-regular tree of every finite depth (very different than previously used groups). The generators are constructed by efficient algorithms for solving certain equations over the symmetric group (implicit in Nikolov '02). An essential part is the analysis of the expansion of a Cayley graph on two copies of a group GxG with (perfectly correlated) set of generators of the form (g, g^{-1}), when all elements of G are commutators.

The main open problem is to prove the conjecture below, which would establish the expansion of the "seed" Cayley graph (and hence by the theorem of all others in the family).

Conjecture: For some finite k, there exists a set of k^{1/5} permutations in S_k, such that the resulting Cayley graph is an expander.

No special background will be assumed.
Joint work with Avi Wigderson (IAS).

Wednesday, October 22, 2003

Ahuva Mu'alem, The Hebrew University

Towards a Characterization of Truthful Combinatorial Auctions

Abstract:
This paper analyzes incentive compatible (truthful) mechanisms over restricted domains of preferences, the leading example
being combinatorial auctions.

Our work generalizes the characterization of Roberts (1979) who showed that truthful mechanisms over {\em unrestricted}
domains with at least 3 possible outcomes must be ``affine maximizers''. We show that

truthful mechanisms for combinatorial
auctions (and related restricted domains) must be ``almost affine maximizers'' if they also satisfy an additional requirement
of ``independence of irrelevant alternatives (IIA)''. This requirement is without loss of generality for unrestricted domains
as well as for auctions between two players where all goods must be allocated.

This implies unconditional results for these cases, including a new proof of Roberts' theorem. The computational implications
of this characterization are severe, as reasonable ``almost affine maximizers'' are shown to be as computationally hard as exact
optimization. This implies the near-helplessness of such truthful polynomial-time auctions in all cases where exact optimization
is computationally intractable.

This is joint work with Ron Lavi and Noam Nisan.

Monday, October 27, 2003

Nicholas Pippenger, Princeton University

Probability Theory and Covering Problems

Abstract:
Many of the combinatorial problems arising in the theory of computation can be expressed as covering problems: given a
collection of sets of points, how few sets can be chosen to cover all the points? One way to obtain bounds for such problems
is to consider choosing the sets at random. Since its introduction in the 1950s, this method has grown steadily in
sophistication through the addition of various auxiliary probabilistic techniques. The goal of this talk is to survey this
development through the presentation of some illuminating examples.

Monday, November 3, 2003

Scott Aaronson, University of California Berkeley

Multilinear Formulas and Skepticism of Quantum Computing

Abstract:

Several researchers, including Leonid Levin, Gerard 't Hooft, and Stephen Wolfram, have argued that quantum mechanics will
break down before the factoring of large numbers becomes possible. If this is true, then there should be a natural "Sure/Shor
separator" -- that is, a set of quantum states that can account for all experiments performed to date, but not for Shor's
factoring algorithm. We propose as a candidate the set of states expressible by a polynomial number of additions and tensor
products. Using a recent lower bound on multilinear formula size due to Raz, we then show that states arising in quantum
error-correction require n^{Omega(log n)} additions and tensor products even to approximate, which incidentally yields the
first superpolynomial gap between general and multilinear formulas. More broadly, we introduce a complexity classification
of pure quantum states, and prove many basic facts about this classification. Our goal is to refine vague ideas about a
breakdown of quantum mechanics into specific hypotheses that might be experimentally testable in the near future.

Tuesday, November 4, 2003

Russell Impagliazzo, IAS

Priority Algorithms: Greedy Graph Algorithms, and Beyond

Abstract:

Borodin, Nielsen, and Rackoff introduced the notion of priority algorithms. The priority algorithm model is an abstract
model which captures the intrinsic power and limitations of greedy algorithms. The original paper was limited to scheduling
problems; but subsequent work extended the model to other domains. We generalize their notion to general optimization problems,
and apply it, in particular, to graph problems. We characterize the best performance of algorithms in each model in terms of a
combinatorial game between a Solver and an Adversary.

We prove bounds on the approximation ratio achievable by such algorithms for basic graph problems such as shortest path,
metric Steiner trees, independent set and vertex cover. For example, we show no fixed priority algorithm can achieve any
approximation ratio (even a function of the graph size). In contrast, the well-known Dijkstra's algorithm shows that an
adaptive priority algorithm can find optimal paths. We prove that the approximation ratio for vertex cover achievable by
adaptive priority algorithms is exactly 2, matching the known upper bound.

The above is joint work with Sashka Davis.

The priority model, in addition to capturing the limits of greedy methods for approximation algorithms, can be used
as a starting point to address the limits of standard algorithmic techniques in a variety of settings.

We will also discuss how to extend the priority framework to model back-tracking and dynamic programming algorithms.
We'll also give priority models for k-SAT. We can then pose some open problems: for what densities of random SAT formulas
can greedy heuristics find solutions? (Note that most of the lower bounds for the threshold for satisfiability involve analyzing
such heuristics.) How well do standard DPLL-like techniques work on satisfiable instances?

This is work in progress with Angelopolous, Beame, Borodin, Buresh-Oppenheim, Davis, Pitassi, and whoever else we can get
interested in it.

Monday, November 10, 2003

Erik Demaine, MIT

Folding and Unfolding in Computational Geometry

Abstract:
When can a linkage of rigid bars be untangled or folded into a desired configuration? What polyhedra can be cut along
their surface and unfolded into a flat piece of paper without overlap? What shapes can result by folding a piece of paper
flat and making one complete straight cut? Folding and unfolding is a branch of discrete and computational geometry that
addresses these and many other intriguing questions. I will give a taste of the many results that have been proved in the
past few years, as well as the several exciting open problems that remain open. Many folding problems have applications in
areas including manufacturing, robotics, graphics, and protein folding.

Tuesday, November 11, 2003

Dorit Aharonov, Hebrew University

Approximating the Shortest and Closest Vector in a Lattice to within Sqrt(n) Lie in NP Intersect coNP

Abstract:

I will describe this new result with Oded Regev. The result (almost) subsumes the three mutually-incomparable previous
results regarding these lattice problems: Lagarias, Lenstra and Schnorr [1990], Goldreich and Goldwasser [2000], and
Aharonov and Regev [2003]. Our technique is based on a simple fact regarding succinct approximation of functions using
their Fourier transform over the lattice. This technique might be useful elsewhere--I will demonstrate this by giving a
simple and efficient algorithm for one other lattice problem (CVPP) improving on previous results.

Monday, November 17, 2003

Claire Kenyon, &#65533;cole Polytechnique and Institut Universitaire de France

Approximation Algorithms for Packing

Abstract:
This talk will discuss design and analysis techniques for
bin-packing type problems. We will review classical bin-packing
algorithms, focusing on average-case analysis, and present the
"Sum-of-Squares" algorithm, a simple online algorithm with great
average-case performance. In terms of worst-case performance, we will
recall the classical approximation scheme of de la Vega and Lueker and
explore how it
can be extended to several higher-dimensional settings, and in particular
present an asymptotic (i.e., as OPT/(maximum rectangle
height) goes to infinity) approximation scheme for dynamic storage
allocation.

Tuesday, November 18, 2003

Mikhail Alekhnovitch, IAS

Joint with Edward A. Hirsch and Dmitry Itsykson

Exponential Lower Bounds for the Running Time of DPLL Algorithms on Satisfiable Formulas

Abstract:
DPLL (for \emph{Davis}, \emph{Putnam}, \emph{Logemann}, and \emph{Loveland})
algorithms form the largest family of contemporary
algorithms for SAT (the propositional satisfiability problem) and are widely used in
applications. The recursion trees of DPLL

algorithm executions on unsatisfiable formulas are equivalent to tree-like resolution proofs. Therefore, lower bounds for
tree-like resolution (which are known since 1960s) apply to them.

However, these lower bounds say nothing about the behavior of such algorithms on satisfiable formulas. Proving exponential
lower bounds for them in the most general setting would be equivalent to proving $\mathbf{P}\neq\mathbf{NP}$. In this paper,
we give exponential lower bounds for two families of DPLL algorithms: generalized "myopic" algorithms (that read upto
$n^{1-\epsilon}$ of clauses at each step and see the remaining part of the formula without negations) and "drunk"
algorithms (that choose a variable using any complicated rule and then pick its value at random).

Monday, November 24, 2003

Adam Kalai, TTI

Boosting in the Presence of Noise

Abstract:
Boosting algorithms are procedures that "boost" low-accuracy weak learning algorithms to achieve arbitrarily high accuracy.
Over the past decade, boosting has been widely used in practice and has become a major research topic in computational
learning theory. One of the difficulties associated with boosting in practice is noisy data. We study boosting in the
presence of random classification noise, giving an algorithm and a matching lower bound.

This is joint work with Rocco Servedio.

Tuesday, November 25, 2003

Omer Reingold, AT & T and IAS

PCP Testers: Towards a Cominatorial Proof of the PCP Theorem

Abstract:
In this work we look back into the proof of the PCP
theorem, with the goal of finding new proofs that are either simpler or "more combinatorial". For that we introduce the notion
of a PCP-Tester. This natural object is a strengthening of the standard PCP verifier, and enables simpler composition that is
truly modular (i.e. one can compose two testers without any assumptions on how they are

constructed, as opposed to the current
proof in which one has to construct the two PCPs with respect to "compatible encoding" of their variables). Based on this notion,
we present two main results:

1. The first is a new proof of the PCP theorem. This proof relies on a very weak PCP given as a "black box". From this, we
construct combinatorially the full PCP, relying on composition and a new combinatorial aggregation technique.

2. Our second construction is a "standalone" combinatorial construction showing NP subset PCP[polylog, 1]. This implies, for
example, that max-SAT is quasi-NP-hard.

Joint work with Irit Dinur

Monday, December 1, 2003

Sanjeev Arora, Princeton University

Expander Flows and a sqrt{log n}-Approximation for Graph Expansion/Sparsest Cut

Abstract:
In graph separation problems such as EXPANSION, BALANCED CUT etc., the goal is to divide the graph into two roughly equal
halves so as to minimize the number of edges in this cut. They arise in a variety of settings, including divide-and- conquer
algorithms and analysis of Markov chains. Finding optimal solutions is NP-hard.

Classical algorithms for these problems include eigenvalue approaches (Alon and Millman 1985) and multicommodity flows
(Leighton and Rao 1988). The latter yields a O(log n)-approximation, and it has been an open problem to improve this ratio.
We describe a new O(\sqrt{log n})-approximation algorithm.

The algorithm relies on semidefinite relaxations that use the triangle inequality constraints. Analysing these relaxations
has been an important open problem as well and our work may be seen as significant progress in that direction.

We also introduce the notion of {\em expander flows}, which constitute an interesting ``certificate'' of a graph's expansion.
We use them to prove a surprising new theorem about graph embeddings:

for every n-node graph it is possible to embed an n-node expander graph in it with appropriate dilation and congestion.
We think this embedding theorem may have other applications.

Our techniques are an interesting mix of graph theory and high-dimensional geometry. The talk will be self-contained.

(Joint work with Satish Rao and Umesh Vazirani)

Tuesday, December 2, 2003

Avi Wigderson, IAS

Derandomized "low degree" tests via "epsilon-biased" sets, with Applications

to short Locally Testable Codes and PCPs

Abstract:
We present the first explicit construction of Probabilistically Checkable Proofs (PCPs) and Locally Testable Codes (LTCs)
of fixed constant query complexity which have almost-linear size. Such objects were recently shown to exist (nonconstructively)
by Goldreich and Sudan.

The key to these constructions is a nearly optimal randomness-efficient version of the Rubinfeld-Sudan low degree test.
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 showed 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 randomness-efficient version of the Blum-Rubinfeld-Sudan linearity test (which is used, for instance,
in locally testing the Hadamard code).

Both derandomizations are obtained through epsilon-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 epsilon-biased set.

The analysis of the derandomized tests rely on alternative views of epsilon-biased sets --- as generating sets of Cayley
expander graphs for the low degree test, and as defining good linear error-correcting codes for the linearity test.

Joint work with Eli Ben-Sasson, Madhu Sudan and Salil Vadhan

Monday, December 8, 2003

Valentine Kabanets, Simon Fraser University

Complexity of Succinct Zero-sum Games

Abstract:
We study the complexity of solving succinct zero-sum games, i.e., the games whose payoff matrix M is given implicitly by a
Boolean circuit C such that M(i,j)=C(i,j). We complement the known EXP-hardness of computing the exact value of a succinct
zero-sum game by several results on approximating the value.

1. We prove that approximating the value of a succinct zero-sum game to within an additive factor is complete for the class
promise-S_2^p, the ``promise'' version of S_2^p. To the best of our knowledge, it is the first natural problem shown complete
for this class.

We describe a ZPP^{NP} algorithm for constructing approximately optimal strategies, and hence for approximating the value,
of a given succinct zero-sum game. As a corollary, we obtain, in a uniform fashion, several complexity-theoretic results, e.g.,
a ZPP^{NP} algorithm for learning circuits for SAT~\cite{BCGKT96} and a recent result by Cai~\cite{Cai01} that
S_2^p\subseteq ZPP^{NP}.
Joint work with Lance Fortnow, Russell Impagliazzo, and Chris Umans.

Tuesday, December 9, 2003

Ryan O'Donnell, IAS

Learning Mixtures of Product Distributions

Abstract:
In this talk we give a general method for learning mixtures of unknown product distributions on R^n. In particular we give:

1. A (weakly) polynomial time algorithm for learning a mixture of any constant number of axis-aligned Gaussians in R^n
(the Gaussians need not be spherical). Our algorithm constructs a highly accurate approximation to the unknown mixture

of Gaussians, and unlike previous algorithms, makes no assumptions about the minimum separation between the centers of the
Gaussians.

A poly(n) time algorithm for learning a mixture of any constant number of product distributions over the boolean cube
{0,1}^n. Previous efficient algorithms could only learn a mixture of two such product distributions. We also give evidence
that no polynomial time algorithm can learn a mixutre of a superconstant number of such distributions.

This is joint work with Jon Feldman and Rocco Servedio of Columbia University.

Monday, January 19, 2004

Thomas Hayes, Toyota Technological Institute, Chicago

Randomly Sampling Graph Colorings

Abstract:
For a given graph G, and a positive integer k, we consider the problem of sampling proper k-colorings of G almost-uniformly at random. When k is larger than the maximum degree of G, there is a greedy algorithm for _constructing_ such a coloring in linear time. However, even with this many colors, it is not known whether colorings can be sampled nearly uniformly in polynomial time.

We recently showed that, assuming G has high girth and high maximum degree, that when $k > (1+\epsilon) \times$max degree, k-colorings can be sampled efficiently. The factor $(1+\epsilon)$ improves the previously best known factor 1.489..., which also required similar assumptions on the graph. The best known factor which does not require any assumptions on the graph is 11/6.

The coloring algorithm we study is a very simple Markov chain on the space of proper graph colorings. A main focus of this talk will be new extensions to the coupling technique for proving rapid "mixing" of such chains. This is joint work with Eric Vigoda, of Univ. of Chicago.

Tuesday, January 20, 2004

Ran Raz, Weizmann Institute

Multi-Linear Formulas for Permanent and Determinant are of Super-Polynomial Size

Abstract:

Arithmetic formulas for computing the determinant and the permanent of a matrix have been studied since the 19th century.
Are there polynomial size formulas for these functions ? Although the determinant and the permanent are among the most
extensively studied computational problems, polynomial size formulas for these functions are not known. An outstanding
open problem in complexity theory is to prove that polynomial size formulas for these functions do not exist. Note, however,
that super-polynomial lower bounds for the size of arithmetic formulas are not known for any explicit function and that
questions of this type are considered to be among the most challenging open problems in theoretical computer science.
I will talk about a recent solution of this problem for the subclass of multi-linear formulas.

An arithmetic formula is multi-linear if the polynomial computed by each of its sub-formulas is multi-linear, that is, in
each of its monomials the power of every input variable is at most one. Multi-linear formulas are restricted, as they do not
allow the intermediate use of higher powers of variables in order to finally compute a certain multi-linear function. Note,
however, that for many multi-linear functions, formulas that are not multi-linear are very counter-intuitive, as they require
a "magical" cancellation of all high powers of variables. Note also that both the determinant and the permanent are multi-linear
functions and that many of the well known formulas for these functions are multi-linear formulas.

We prove that any multi-linear arithmetic formula for the determinant or the permanent of an n dimensional matrix is of size
super-polynomial in n. Previously, no lower bound was known (for any explicit function) even for the special case of
multi-linear formulas of constant depth.

Monday, January 26, 2004

Mario Szegedy, Rutgers University

Spectra of Quantized Walks and a $sqrt{\delta\epsilon}$-Rule

Abstract:
We introduce quantized bipartite walks, compute their spectra, generalize the algorithms of Grover and Ambainis and interpret
them as quantum walks with memory. We compare the performance of walk based classical and quantum algorithms and show that the
latter run much quicker in general. Let P be a symmetric Markov chain with transition probabilities

$P[i,j]$, $(i ,j\in [n])$. Some elements of the state space are marked. We are promised that the set of marked elements
has size either zero or at least ϵn. The goal is to find out with great certainty which of the above two cases holds.
Our model is a black box that can answer certain yes/no questions and can generate random elements picked from certain
distributions. More specifically, by request the black box can give us a uniformly distributed random element for the cost
of $\wp_{0}$. Also, when ``inserting'' an element i into the black box we can obtain a random element j, where j is
distributed according to $P[i,j]$. The cost of the latter operation is $\wp_{1}$. Finally, we can use the black box to test
if an element i is marked, and this costs us $\wp_{2}$. If δ is the eigenvalue gap of P, there is a simple
classical algorithm with cost $O(\wp_{0} + (\wp_{1}+\wp_{2})/\delta\epsilon)$
that solves the above promise problem.
(The algorithm is efficient if $\wp_{0}$
is much larger than $\wp_{1}+\wp_{2}$.) In contrast,
we show that for the ``quantized'' version of the algorithm
it costs only $O(\wp_{0} + (\wp_{1}+\wp_{2})/\sqrt{\delta\epsilon})$
to solve the problem. We refer to this as the $\sqrt{\delta\epsilon}$ rule. Among the technical contributions we
give a formula for the spectrum of the product of two general reflections.

Tuesday, January 27, 2004

James R. Lee, Berkeley University

Metric Decomposition: Coping with Boundaries

Abstract:
In recent years, "stochastic" metric decomposition has become a significant tool in the study of discrete metric spaces.
In this talk, I will survey the emerging structural theory, with an eye toward applications in mathematics and computer science.

In particular, I will describe some recent work with Assaf Naor on a classical problem in geometry and analysis, that of
extending Lipschitz functions. This shows that metric decomposition also yields important insights in the continuous
setting.

In the finite setting, I will sketch a new proof of Bougain's embedding theorem based on metric decomposition and a
technique we call "measured descent," which answers some open problems in the field and lends new insights.

(Joint work with R. Krauthgamer, M. Mendel, and A. Naor; a more detailed exposition of this result will be given at
Princeton during the preceding week, but I think it is instructive to see a sketch in the broader context.)

Monday, February 2, 2004

Aner Shalev, Hebrew University

Probabilistic Generation of Finite Simple Groups, Random Walks, Fuchsian Groups and

Abstract:
We use character theory and probabilistic methods to solve several seemingly unrelated problems of combinatorial and geometric flavor involving finite and infinite groups.

These include random generation of finite simple groups, determining the mixing time of certain random walks on
symmetric groups and matrix groups, finding the subgroup growth of Fuchsian groups, and giving a probabilistic proof to a conjecture of Higman on their finite quotients.

If time allows I will also discuss further applications to representation varieties, and to counting branched coverings of Riemann surfaces.

A main tool in the proofs is the study of the so called Witten's zeta function encoding the character degrees of certain groups.

Tuesday, February 3, 2004

Noga Alon, Tel-Aviv University

CutNorm, Grothendieck's Inequality, and Approximation Algorithms for Dense Graphs

Abstract:
The cut-norm of a real matrix A is the maximum absolute value of the sum of all elements in a

submatrix of it.
This concept plays a major role in the design of efficient approximation algorithms for dense graph and matrix problems.
After briefly explaining this role I will show that the problem of approximating the cut-norm of a given real matrix is
computationally hard, and describe an efficient approximation algorithm. This algorithm finds, for a given matrix A,
a submatrix A' so that the absolute value of the sum of all entries of A' is at least c times the cut-norm of A, where c > 0.56. The algorithm combines semidefinite programming with a novel rounding technique based on Grothendieck's Inequality.

Joint work with Assaf Naor.

Monday, February 9, 2004

Nathan Segerlind, IAS

Using Hypergraph Homomorphisms to Guess Three Secrets

Abstract:
The problem of ``guessing k secrets'' is the following two-player game between an adversary and a seeker: The adversary has
k binary strings and the seeker wishes to learn as much as he can by asking boolean questions about strings. For each question, the adversary chooses one of his secret strings and provides an answer for that
particular string. In particular, if the seeker asks a question such that the adversary has a string for which the answer YES
and a string for which the answer is NO, the adversary may provide either answer. This problem was introduced by Chung,
Graham and Leighton to model difficulties encountered in internet routing by Akamai.

We present the first polynomial-time strategy for solving the problem for guessing three secrets (previous work had solved
only the case for two secrets). We give both adaptive and oblivious strategies. The queries of the oblivious strategy are
built from a special separating system that we call a prefix-separating system. The algorithm for recovering the solution
from the queries is based on the homomorphism structure of three-uniform, intersecting hypergraphs. Along the way we prove
some combinatorial facts about three-uniform, intersecting hypergraph cores that may be of independent interest (in this
setting a core is a hypergraph that admits no proper endomorphism).

Tuesday, February 10, 2004

Peter Winkler, Bell Labs and IAS

Some Vexing Combinatorial and Mixing Problems

Abstract:
Included: Euler tours, Thorp shuffles, particles in space,
and a conjecture motivated by optical networking which generalizes theorems of two Halls.

Monday, February 16, 2004

Christos Papadimitriou, University California Berkeley

Nash Equilibria and Complexity

Abstract:
Using the Nash equilibrium problem as a departure point, we explore the intricate and largely mysterious interplay
between computational complexity and existence proofs in combinatorics. We present polynomial-time algorithms and
complexity results for congestion games.

Tuesday, February 17, 2004

Maria Chudnovsky, Princeton, CMI and IAS

The Structure of Clawfree Graphs

Abstract:
A graph is said to be clawfree if it has no induced subgraph isomorphic to $K_{1,3}$. Line graphs are one well-known
class of clawfree graphs, but there others, such as circular arc graphs and subgraphs of the Schl\"{a}fli graph.
It has been an open question to describe the structure of all clawfree graphs. Recently, in joint work with Paul
Seymour, we were able to prove that all clawfree graphs can be constructed from basic pieces (which include the
graphs mentioned above, as well as a few other ones) by gluing them together in prescribed ways. This talk will
survey the main ideas of the proof, as well as some examples of claw-free graphs that turned out to be important
in the course of this work, and some applications.

Wednesday, February 18, 2004

Roy Meshulam, Technion, Haifa

Laplacians, Homology and Hypergraph Matching

Abstract:
We'll discuss some relations between the expansion of a graph and the topology of certain complexes associated with the graph.
Applications include a lower bound on the homological connectivity of the independence complex, in terms of a new graph
parameter defined via certain vector representations of the graph. This in turn implies Hall type theorems for matchings
in hypergraphs. Joint work with R. Aharoni and E. Berger.

Monday, February 23 2004

Ravi Kumar, IBM Almaden Reseaerch Center

On Separating Nondeterminism and Randomization inCommunication Complexity

Abstract:
In the two-party communication complexity model, we show that the tribes function on n inputs has two-sided error randomized complexity Omega(n), while its nondeterminstic complexity and co-nondeterministic complexity are both Theta(sqrt(n)). This separation between randomized and nondeterministic complexity is the best possible and it settles an open problem posed by Beame--Lawry and Kushilevitz--Nisan.

Our lower bound is obtained using generalizations of information complexity, which quantifies the minimum amount of information that will have to be revealed about the inputs by every correct communication protocol.

(Joint work with T. S. Jayram and D. Sivakumar)

Monday, February 23, 2004

Mark Braverman, University of Toronto

On the Computability of Julia Sets

Abstract:
While the computer is a discrete device, it is often used to solve problems of a continuous nature. The field of Real Computation addresses the issues of computability in the continuous setting. We will discuss different models of computation for subsets of R^n.
The main definition we use has a computer graphics interpretation (in the

case n=2), as well as a deeper mathematical meaning.
The Julia sets are particularly well studied sets arising from complex dynamics. In the talk we will present the basic facts about Julia sets and some computability results for them. Our computability results come in contrast to the Julia sets noncomputability results presented by Blum/Cucker/Shub/Smale. This discrepancy follows from the fact that we are using a different computability model.

Tuesday, February 24, 2004

Stephen Cook, University of Toronto

Making Sense of Bounded Arithmetic

Abstract:
We present a unified treatment of logical theories for each of the major complexity classes between AC0 and P, and give simple translations into the quantified propositional calculus.

Monday, March 1, 2004

Yonatan Bilu, Hebrew University

Quasi-Ramanujan 2-lifts - A New Construction of Expander Graphs

Abstract:
The adjacency matrix of a graph on n vertices is a 0-1 nxn matrix, with the (i,j) entry being 1, iff there is an edge between vertices i and j. Often, understanding the eigenvalues of the adjacency matrix sheds light on combinatorial properties of a graph. In a d-regular graph, the largest eigenvalue is d. If the absolute value all other eigenvalues is small, we say that such a graph is an expander. Optimal constructions of such graphs are known, but they rely on group representation theory. In this work we describe a simple, combinatorial construction that is close to being optimal.
A useful property of expander graphs is the so-called Expander Mixing Lemma, which bounds the deviation of the number of edges between two sets of vertices from what is expected in a random graph. We show a converse to this lemma, namely, that if the deviation is small then the graph is an expander. The implication is surprisingly strong, in comparison to other combinatorial properties (edge expansion, vertex expansion) which also imply a bound on the eigenvalues.

The key tool in the construction is a signing of the edges of a d-regular graph by +1 and -1. This yields a 2-lift of a graph - graph on twice as many vertices which is also d-regular. Getting an expander graph in this way reduces to finding a signing of a graph such that the spectral radius of the signed adjacency matrix is small. We show that for all d-regular graphs such a signing exists, and that if the graph we start from is an expander, then such a signing can be found efficiently. This is joint work with Nati Linial.

Tuesday, March 2, 2004

Toniann Pitassi, IAS

On a Model for Backtracking

Abstract:
Most known efficient algorithms fit into a few basic paradigms: divide-and-conquer, dynamic programming, greedy algorithms, hill-climbing, and linear programming.
There has been a growing interest in trying to formalize precise models for these and other algorithmic paradigms, most notably in the context of approximation algorithms. For example, models of local search algorithms have been studied by several researchers; more recently, Borodin, Nielsen and Rackoff introduce priority algorithms to model greedy algorithms, and Arora, Bollobas and Lovasz provide integrality gaps for a wide class of LP formulations.

We continue in this direction by studying a new model for backtracking (BT). Informally, a BT algorithm is a levelled tree where each level corresponds to the probing of an input item and branches correspond to different possible irrevocable decisions that can be made about this item. The complexity of a BT algorithm is the number of nodes in the tree. We study several classic problems within this framework (knapsack, interval selection and satisfiability), and present a number of upper and lower bounds, and inapproximability results on the power of BT algorithms for these problems.
Finally we discuss connections with other models (i.e. dynamic programming, DPLL), and many open problems.

This is work in progress with Allan Borodin, Russell Impagliazzo, Josh Buresh-Oppenheim, and Avner Magen.

Monday, March 8, 2004

Abraham Neyman, Institute of Mathematics, Hebrew University

Online Concealed Correlation by Boundedly Rational Players

Abstract:
Joint paper with Gilad Bavly (Hebrew University)

In a repeated game with perfect monitoring,
correlation among a group of players may evolve in the common
course of play (called, online correlation). Such a correlation
may be `concealed' from a boundedly rational player. We show that
``strong'' players, i.e., players whose strategic complexity is
less stringently bounded, can orchestrate online correlation of
the actions of ``weak'' players, in a manner that is concealed
from an opponent of ``intermediate'' strength.

The result is illustrated in two models, each captures another
aspect of bounded rationality. In the first, players use bounded
recall strategies. In the second, players use strategies that are
implementable by finite automata.

Tuesday, March 9, 2004

Boaz Barak, IAS

Extracting Randomness from Few Independent Sources

Abstract:

We consider the problem of extracting truly
random bits from several independent sources of data that contains
entropy. The best previously known explicit constructions
extracted randomness from two independent samples of distributions
over $\bits^n$ such that each has min-entropy at least $n/2$. The
optimal, non-explicit construction only requires the min-entropy
to be more than $\log n$.

In this work, we manage to go beyond this $n/2$ ``barrier'' and
give an explicit construction for extracting randomness from
distributions over $\bits^n$ with δn entropy for every
constant $\delta>0$. The number of samples we require is a
constant (depending polynomially on $1/\delta$).

Our main tools are results from additive number theory and in
particular a recent result by Bourgain, Katz and Tao (GAFA, to

appear).

We also consider the related problem of constructing randomness dispersers, and construct an almost optimal disperser that requires the input distribution only to have min-entropy at least $\Omega(\log n)$, with the caveat that all the samples have to come from the \emph{same} distribution. The main tool we use is a variant of the ``stepping-up lemma'' used in establishing lower bound on the Ramsey number for hypergraphs (Erdos and Hajnal, 71).

Joint work with Russell Impagliazzo and Avi Wigderson.

Monday, March 15, 2004

Amir Shpilka, The Weizmann Institute

Locally Testable Cyclic Codes

Abstract:
Cyclic codes are codes which are invariant under a cyclic shift of their coordinates. Locally testable codes are, roughly, codes for which we can verify, by querying a small number of positions, whether a given word is in the code or far from being a code word.

It is a long standing open problem in coding theory whether there exist good cyclic codes. It is a more recent question, rising from the study of probabilistically checkable proofs, whether there exist good locally testable codes. In this talk we address the intersection of the two problems and show that there are no good locally testable cyclic codes.

In particular our results imply that for certain block lengths there are no good cyclic codes, without any local testability assumption.

The techniques we use are algebraic in nature and rely on the beautiful connection between cyclic codes and cyclotomic polynomials. We will try to give as many details from the proof as possible.

This is a joint work with Laci Babai and Daniel Stefankovic from the university of Chicago.

Tuesday, March 16, 2004

Subhash Khot, IAS

BCH Codes, Augmented Tensor Products and Hardness of theShortest Vector Problem in Lattices

Abstract:
The Shortest Vector Problem in lattices has been studied by mathematicians
for two centuries. Given a basis for an n-dimensional lattice, the problem is
to find the shortest non-zero vector in the lattice. The approximation version
of the problem asks for a non-zero lattice vector that is guaranteed to be
within a certain factor of the shortest vector.

There is a rich set of results associated with SVP. For example,

1. Gauss' algorithm that works for 2-dimensional lattices.

2. Minkowski's Convex Body Theorem that shows existence of a short lattice vector.

The famous LLL algorithm of Lenstra, Lenstra and Lovasz that achieves 2^n
approximation to SVP. It was improved to 2^{o(n)} by Schnorr. This algorithm
has numerous applications in mathematics, computer science and cryptography.

Ajtai's reduction from worst-case hardness of approximating SVP to its average
case hardness.

Ajtai-Dwork's public-key cryptosystem based on (conjectured) worst case hardness
of approximating SVP. Also, a recent alternate construction by Regev.

Results showing that a gap-version of SVP with factor n or \sqrt{n} is "unlikely"
to be NP-hard, e.g. Lagarias, Lenstra and Schnorr; Goldreich and Goldwasser;
and recently Aharonov and Regev.
However, there has been very little progress in actually proving hardness of
approximation results for SVP. Even the NP-hardness of exact version of SVP came
only in 1998 (by Ajtai). It was strengthened to a hardness of approximation
result with factor \sqrt{2} by Micciancio. This left a huge gap of \sqrt{2} vs
2^{o(n)} between the best hardness result and the best algorithmic result.

In this talk, we greatly improve the hardness factor. We show that

assuming
NP \not\in BPP, there is no constant factor approximation for SVP (in
polytime).
Assuming NP \not\in BPTIME(2^{polylog n}}, we show that SVP has no
approximation
with factor 2^{(\log n)^{1/2-\eps}} where \eps > 0 is an arbitrarily
small constant.
In our opinion, this gives evidence that there is no efficient algorithm
that
achieves polynomial factor approximaiton to SVP, a major open problem in
algorithms.

We first give a new (randomized) reduction from Closest Vector
Problem (CVP)
to SVP that achieves *some* constant factor hardness. The reduction is based
on BCH Codes. Its advantage is that the SVP instances produced by the
reduction
behave well under the augmented tensor product, a new variant of tensor
product
that we introduce. This enables us to boost the hardness factor to an
arbitrarily
large constant assuming NP \not\in BPP, and to factor 2^{(\log
n)^{1/2-\epsilon}}
assuming the stronger complexity assumption.

Monday, March 22, 2004

Moni Naor, Weizmann Institute of Science

Spam and Pebbling

Abstract:

Consider the following simple technique for combating spam:

> If I don't know you, and you want your e-mail to appear in my inbox,
> then you must attach to your message an easily verified
> "proof of computational effort", just for me and just for this message.

To apply this approach one needs to be able to come up with computational problems where solving them requires significant expenditure of resources while verifying a solution can be done easily. Recent work dealt with the choice of computational problems for which most of the work is in retrieving information from memory. In this talk I will describe this approach and describe the connection to pebbling problems.

The talk is based on two papers:
Cynthia Dwork, Andrew Goldberg and Moni Naor:
On Memory-Bound Functions for Fighting Spam.
Cynthia Dwork, Moni Naor and Hoeteck Wee: work in progress

Tuesday, March 23, 2004

Manindra Agrawal, IAS

Efficient Primality Testing

Abstract:
Since the discovery of polynomial time primality test, a number of
improvements have been made to the original algorithm. Amongst the most notable are:

1. An O(log^{10.5} n) time algorithm with a completely elementary
   proof. This eliminates
   the dependence on a difficult analytic number theory lemma used
   in the original proof.

An O(log6 n) time deterministic algorithm. This matches the
conjectured running time
of the original algorithm.
An O(log4 n) time randomized algorithm that produces primality
certificates. This is
the fastest provable algorithm of its kind.
In this talk, I will discuss details of these algorithms and their proofs.

Monday, March 29, 2004

Mike Capalbo, DIMACS

Graph Products are (almost!) Practical

Abstract:
We are given an infinite
family of expanders. We would like to use this family of
expanders to construct routing networks with N inputs, N outputs,
bounded degree, and O(N log N) edges, where

the routing can be done in a distributed fashion in O(log N) time
(using vertex-disjoint paths).

In the early 1990's, Pippenger devised a method to construct such routing networks from an arbitrary family of expanders. The drawback to this method is that the constants hidden under the O-notation are very, very large. Here, using a graph product, we present a new method to construct such routing networks, where the constants hidden behind the O-notation are much, much smaller, even using the simple Gabber-Galil expanders.

Tuesday, March 30, 2004

Andris Ambainis, IAS

Search by Quantum Walks

Abstract:
I will present two new quantum algorithms:

- O(N^{2/3}) quantum algorithm for element distinctness;
- O(\sqrt{N\log N}) quantum algorithm for search on 2-dimensional grid.

The algorithms are based on using a quantum walk (a quantum process similar to a
random walk) to search graphs. This improves over the standard quantum search
by using the structure of a graph.

The talk will be self-contained and no previous knowledge of quantum computation
will be assumed.

The second algorithm is a joint work with Julia Kempe and Alexander Rivosh.

Monday, April 5, 2004

Van Vu, University of California, Dan Diego

A Near Optimal Bound on Erdos Distinct Distances in high Dimensions

Abstract:
One of the oldest and most well known problems of Erdos in discrete geometry is the following: what is the minimum number of distances between

n points in R^d ? (here d is fixed and n tends to infinity)

In this talk, I will give a brief review about the history of the problem and discuss a recent joint result with Solymosi, which proves a near sharp bound in high dimensions. Our proof uses tools from theoretical computer science, in particular a cutting lemma by Chazelle et al.

Tuesday, April 6, 2004

Jan Krajicek, IAS

Strong Proof Systems and Hard Tautologies

Abstract:
I shall discuss what we know about strong (propositional) proof systems, and describe a new method how to construct them. In particular, given a proof system P I define a possibly much stronger proof system iP. The system iP operates with an exponentially long P-proof described ``implicitly'' by a polynomial size circuit and with an ordinary P-proof of the "correctness" of the long proof. I will give some evidence that iP is indeed much stronger than P, discuss the iteration of the construction, and describe some applications in proof-complexity.

An issue important for potential lower bounds is to have "reasonable" candidates of tautologies that could be hard for strong proof systems. I recall some known facts and then show that implicit proof systems are relevant also in this context. In particular, I will show that lower bounds for proofs of implicit formulas in implicit versions of "weak" proof systems can give, in principle, lower bounds for ordinary proofs of ordinary formulas in "strong" proof systems.

Monday, April 12, 2004

Benny Sudakov, Princeton University

Solving Extremal Problems Using Stability Approach

Abstract:
In this talk we discuss a "stability approach" for solving extremal problems. Roughly speaking, it can be described as follows. In order to show that given configuration is a unique optimum for an extremal

problem, we first prove an approximate structure theorem for all constructions whose value is close to the optimum and then use this theorem to show that any imperfection in the structure must lead to a suboptimal configuration. To illustrate this strategy, we discuss three recent results in which stability approach was used to answer a question of Erdos-Rothschild and to resolve two conjectures of Sos and Frankl.

All the results in this talk are co-authored with P. Keevash and the first one is in addition co-authored with N. Alon and J. Balogh.

Tuesday, April 13, 2004

Alexander Razborov, IAS

Guessing More Secrets via List Decoding

Abstract:
We consider the following game introduced by Chung, Graham and Leighton. One player, A picks k>1 secrets from a universe of N possible secrets, and another player, B tries to gain as much information about this set as possible by asking binary questions f:{1,2,..,N}-->{0,1}. Upon receiving such a question f, A adversarially chooses one of her k secrets, and answers f according to it.

For any constant number of secrets k we present an explicit set of $O(\log N)$ questions along with an $O(\log^2 N)$ recovery algorithm that achieve B's goal in this game (previously such algorithms were known only for k=2,3) thus solving one of the central open problems in this area.

Our strategy is based on the list decoding of Reed-Solomon codes, and it extends and generalizes ideas introduced earlier by Alon, Guruswami, Kaufman and Sudan.

Monday, April 19, 2004

Andrew Yao, Princeton University

Some Optimality Results in Bounded-Storage Cryptography

Abstract:
We study the amount of storage needed for two parties to agree on a secret key in the bounded-storage model proposed by Maurer. Assume that a public random string of length n is available and any adversary has a storage of at most νn bits, for some constant $\nu<1$. We show that to have a secure protocol, one of the legitimate

parties must have a storage of $\Omega(\sqrt{n})$ bits. This can be generalized to the case where two parties can share a private key beforehand and then try to agree on a longer secret. We show that with a private key of length r, one of the legitimate parties now needs a storage of $\Omega(\sqrt{n/2^r})$ bits. Our lower bounds are optimal within constant factors as we have protocols with storage requirements matching these bounds. This is joint work with C.J. Lu and D.W. Wang.

Tuesday, April 20, 2004

Andris Ambainis, IAS

Search by Quantum Walks II

Abstract:
I will continue my talk from Tuesday, March 30. I will present an O(\sqrt{N\log N}) quantum algorithm for spatial search and then describe the common mathematical structure behind the two algorithms (element distinctness, presented on March 30 and spatial search).

Monday, April 26, 2004

Jon Kleinberg, Cornell University

Network Failure Detection and Graph Connectivity

Abstract:
Measuring the properties of a large, unstructured network can be difficult: one may not have full knowledge of the network topology, and detailed global measurements may be infeasible. A valuable approach to such problems is to take measurements from selected locations within the network and then aggregate them to infer large-scale properties. One sees this notion applied in settings that range from Internet topology discovery tools to remote software agents that estimate the download times of popular Web pages. Some of the most basic questions about this type of approach, however, are largely unresolved at an analytical level. How reliable are the results? How much does the choice of measurement locations affect the aggregate information one infers about the network?

We describe algorithms that yield provable guarantees for a problem of this type: detecting a network failure.
In particular, we provide methods for placing a small set of ``agents'' at nodes in a network, in such a way that any significant partition of the network will be detected

by the separation of some pair of these agents.
We find that the number of agents required can be bounded
in terms of certain natural parameters of the failure, and
independently of the size of the network itself.
These bounds establish connections between graph separators
and the notion of VC-dimension, employing results related
to non-bipartite matchings and the disjoint paths problem.
In recent joint work with Mark Sandler and Alex Slivkins we have
obtained further improvements to these bounds in terms of the
underlying edge-connectivity, making use of connections to the cactus
representation of the minimum cuts in a graph.

Tuesday, April 27, 2004

Ryan O'Donnell, IAS

Optimal Inapproximability Results for MAX-CUT and Other 2-Variable CSPs

Abstract:
In this paper we give evidence that it is hard to polynomial-time
approximate MAX-CUT to within a factor of alpha_GW + eps, for all eps > 0.
Here alpha_GW denotes the approximation ratio achieved by the
Goemans-Williamson algorithm [GW95], alpha_GW = .878567. We show that the
result follows from two conjectures: a) the Unique Games conjecture of
Khot [Khot02]; and, b) a widely-believed Fourier-theoretic conjecture we
call the Majority Is Stablest conjecture. Our results suggest that the
naturally hard ``core'' of MAX-CUT is the set of instances in which the
graph is embedded on a high-dimensional Euclidean sphere and the weight of
an edge is given by the squared distance between the vertices it connects.

The same two conjectures also imply that it is hard to
(beta + eps)-approximate MAX-2SAT, where beta = .943943 is the minimum of
[2 + (2/pi) theta]/[3 - cos theta] on (pi/2, pi). Motivated by our proof
techniques, we show that if the MAX-2CSP and MAX-2SAT problems are
slightly restricted --- in a way that seems to retain all their hardness
--- then they have polynomial-time (alpha_GW - eps)- and
(beta - eps)-approximation algorithms, respectively.

Although we are unable to prove the Majority Is Stablest conjecture, we
give some partial results and indicate possible directions of attack.
Our partial results are enough to imply that MAX-CUT is hard to
(3/4 + 1/2pi + eps)-approximate (about .909155) assuming only the Unique
Games conjecture.

Finally, we discuss the possibility of equivalence between the Unique
Games conjecture and the hardness of distinguishing (1 - eps)-satisfiable
and eps-satisfiable instances of two-variable linear equations mod q, for
large q.

This is joint work with Subhash Khot (IAS), Guy Kindler (DIMACS), and Elchanan Mossel (Berkeley).

Monday, May 3, 2004

Sean Hallgren, NEC Research, Princeton

Fast Quantum Algorithms for Computing the Unit Group and Class Group of a Number Field

Abstract:
Computing the unit group and class group of a number field are two of the main tasks in computational algebraic number theory. Factoring integers reduces to a special case of computing the unit group, but a reduction in the other direction is not known and appears more difficult. We give polynomial-time quantum algorithms for computing the unit group and class group when the number field has constant degree.

Tuesday, May 4, 2004

Subhash Khot, IAS

Ruling Out PTAS for Graph Min-Bisection

Abstract:
Graph Min-Bisection is the following problem : Given a graph, partition it into two equal parts so as to minimize the number of crossing edges. The problem arises as a subroutine in many graph algorithms that rely on divide-and-conquer strategy. Feige and Krauthgamer gave an O(log2 n) approximation algorithm for this problem. On the other hand, no inapproximability result was known. It was one of the central open questions in (in)approximability theory whether Min-Bisection has a Polynomial Time Approximation Scheme (i.e. (1+\eps)-approximation algorithm for every \eps > 0).
The result is this talk resolves the above question, ruling out PTAS for Min-Bisection and two other important problems called Densest Subgraph and Bipartite Clique. Recently, Feige ruled out a PTAS for these problems assuming a certain conjecture about average-case hardness of Random 3SAT. Our result needs only a (standard) assumption that NP has no subexponential time algorithms.
The result follows via a new construction of a PCP where the queries of the verifier "look random". To be precise, the verifier makes q queries and for any set of half the bits in the proof, the probability that all

queries fall into this set is essentially 1/2^q. We introduce several new ideas and techniques, in addition to using variations/generalizations of algebraic techniques used to prove the PCP Theorem. I will try to make the talk as self-contained as possible.

Monday, May 10, 2004

Manoj Prabhakaran, Princeton University

New Notions of Security: Universal Composability without Trusted Setup

Abstract:
We propose a modification to the framework of Universally Composable (UC) security [Canetti'01], which enables us to give secure protocols for tasks for which no secure protocol is possible in the original UC framework (except with trusted setup).

Our new notion, involves comparing the protocol executions with an ideal execution involving ideal functionalities (just as in UC-security), but allowing the environment and adversary access to some super-polynomial computational power. We argue the meaningfulness of the new notion, which in particular subsumes many of the traditional notions of security.

We generalize the Universal Composition theorem of [Canetti] to the new setting. Then under new computational assumptions, we realize secure multiparty computation (for static adversaries), without a common reference string or any other setup assumptions, in the new framework. This is known to be impossible under the UC framework.

Joint work with Amit Sahai.

Monday, May 11, 2004

Yuval Peres, University of California, Berkeley

Two Topics on the Interface of Probability and Algorithms

Abstract:

1. Unbiasing and simulation given a coin with unknown bias: Suppose that we are given a coin with an unknown probability p of heads. By tossing this coin repeatedly, a classical trick shows we can simulate an unbiased coin. How about simulating a coin with probability 2p of

heads (when p<1/2)? and with probability f(p) of heads? Solutions to these questions, that start with von Neumann (1952), have led to connections with automata theory, Polya's theorem on positive polynomials, approximation theory and complex analysis. There's an intriguing open problem relating pushdown automata to Algebraic functions. (based on joint works with E. Mossel, S. Nacu.)

2. Let F be a random k-SAT formula on n variables, formed by selecting uniformly and independently m = rn out of all possible k-clauses. It is well-known that if r>2^k ln 2, then the formula F is unsatisfiable with probability that tends to 1 as n grows. We prove that if r < 2^k ln 2 - O(k), then the formula F is satisfiable with probability that tends to 1 as n grows. E.g., When k=10 our lower bound is 704.94 while the upper bound is 708.94. Related methods (with harder analysis) also lead to good bounds for random Max-k-sat. These results raise the question: "Are there different thresholds for the existence of solutions to random satisfiability problems, and for existence of solutions that can be found by polynomial-time algorithms?" (This part based on work with D. Achlioptas and A. Naor).

Monday, May 17, 2004

Yuval Rabani, Technion, on Sabbatical at Cornell University

Two Topics on the Interface of Probability and Algorithms

Abstract:
The study of low distortion embeddings into ℓ_1 is closely related to the study of the integrality gaps of conventional relaxations for some optimization problems on graphs. One obvious way to strengthen the relaxations is to add valid constraints on small subsets of points. We study the effect of such constraints by examining the distortion of embedding metrics that satisfy them into ℓ_1 and other classes of metrics.

Joint work with Ilan Newman and Yuri Rabinovich.

Monday, May 24, 2004

Dana Moshkovitz

Algorithmic Construction of Sets for k-Restrictions

Abstract:

In k-restriction problems one wishes to find a small set of strings that satisfies the following property: if one observes any k indices, and seeks for some specific restriction on them (out of a large set of possible restrictions given as input), then at least one of the strings meets this restriction.

Problems of this type arise in many fields in Computer Science. A few prominent examples are group testing and generalized hashing.

The standard approach for deterministically solving such problems is via almost k-wise independence or k-wise approximations for other distributions.

We offer a generic algorithmic method that yields considerably smaller constructions. Specifically it allows us to derive substantial improvements for the problems mentioned above.

To this end, we generalize a previous work of Naor, Schulman and Srinivasan.

Among other results, we greatly enhance the combinatorial objects in the heart of their method, called /splitters/, using the topological /Necklace Splitting Theorem/.

We are also able to derive an improved inapproximability result for /Set-Cover/ under the assumption P != NP.

Joint work with Noga Alon and Muli Safra

Tuesday, May 18, 2004

Peter Winkler, Bell Labs and IAS

Tournaments, Boxes and Non-Transitive Dice

Abstract:
Many gymnasts compete in 2k-1 events, and are ranked in each with no ties. One competitor is deemed to have "beaten" another if she outscores the other in a majority of the events.

Afterward the organizers are embarrassed to discover that no matter how they award the prizes, there will always be a gymnast who didn't get a prize but beat every gymnast who did.

How many prizes should the organizers have had on hand to be sure this wouldn't happen?

(A meandering blackboard presentation based on recent work with N. Alon, G. Brightwell, H. Kierstead, and A. Kostochka.)

Sponsored by:

[National Science Foundation]

[State of New Jersey]

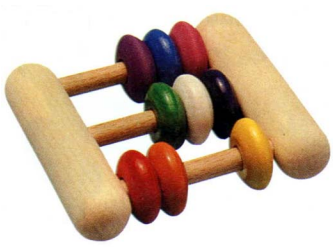

Abacus

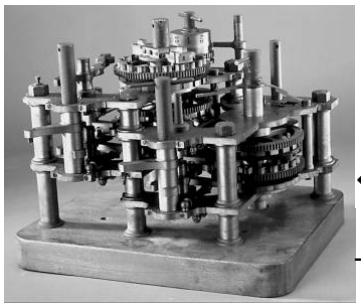

"Analytical Engine"

by Charles Babbage

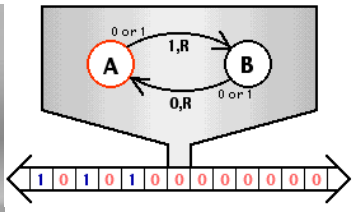

Turing Machine