

FOCS 2011— LIST OF ACCEPTED PAPERS WITH ABSTRACTS

(By submitted order)

1. 104. The Grothendieck constant is strictly smaller than Krivine's bound

Mark Braverman, Konstantin Makarychev, Yury Makarychev, Assaf Naor

Abstract: We prove that $K_G < \frac{\pi}{2 \log(1 + \sqrt{2})}$, where K_G is the Grothendieck constant.

2. 106. The minimum k-way cut of bounded size is fixed-parameter tractable

Mikkel Thorup and Ken-ichi Kawarabayashi

Abstract: We consider the minimum k-way cut problem for unweighted undirected graphs with a size bound s on the number of cut edges allowed. Thus we seek to remove as few edges as possible so as to split a graph into k components, or report that this requires cutting more than s edges. We show that this problem is fixed-parameter tractable (FPT) with the standard parameterization in terms of the solution size s . More precisely, for $s = O(1)$, we present a quadratic time algorithm. Moreover, we present a much easier linear time algorithm for planar graphs and bounded genus graphs.

Our tractability result stands in contrast to known $W[1]$ hardness of related problems. Without the size bound, Downey et al. [2003] proved that the minimum k-way cut problem is $W[1]$ hard with parameter k , and this is even for simple unweighted graphs. Downey et al. asked about the status for planar graphs. We get linear time with fixed parameter k for simple planar graphs since the minimum k-way cut of a planar graph is of size at most $6k$. More generally, we get FPT with parameter k for any graph class with bounded average degree.

A simple reduction shows that vertex cuts are at least as hard as edge cuts, so the minimum k-way vertex cut is also $W[1]$ hard with parameter k . Marx [2004] proved that finding a minimum k-way vertex cut of size s is also $W[1]$ hard with parameter s . Marx asked about the FPT status with edge cuts, which we prove tractable here. We are not aware of any other cut problem where the vertex version is $W[1]$

hard but the edge version is FPT, e.g., Marx [2004] proved that the k -terminal cut problem is FPT parameterized by cut size, both for edge and vertex cuts.

3.

110. Randomness buys depth for approximate counting

Emanuele Viola

Abstract: We show that the promise problem of distinguishing n -bit strings of hamming weights $1/2 \pm \Omega(1/\log^{d-1} n)$ can be solved by explicit, randomized (unbounded-fan-in) $\text{poly}(n)$ -size depth- d circuits with error $\leq 1/3$, but cannot be solved by deterministic $\text{poly}(n)$ -size depth- $(d+1)$ circuits, for every $d \geq 2$; and the depth of both is tight. Previous results bounded the depth to within at least an additive 2.

Our sharper bounds match Ajtai's simulation of randomized depth- d circuits by deterministic depth- $(d+2)$ circuits (Ann. Pure Appl. Logic; '83), and provide an example where randomization (provably) buys resources.

\bigskip

\emph{Techniques:}

To rule out deterministic circuits we combine the switching lemma with an earlier depth-3 lower bound by the author (Comp. Complexity 2009).

To exhibit randomized circuits we combine recent analyses by Amano (ICALP '09) and Brody and Verbin (FOCS '10) with derandomization. To make these circuits explicit -- which we find important for the main message of this paper -- we construct a new pseudorandom generator for certain combinatorial rectangle tests. Based on expander walks, the generator for example fools tests $A_1 \times A_2 \times \dots \times A_{\lceil \lg n \rceil}$ for $A_i \subseteq [n]$, $|A_i| = n/2$ with error $1/n$ and seed length $O(\lg n)$, improving on the seed length $\Omega(\lg n \lg \lg n)$ of previous constructions.

4.

117. Local Distributed Decision

Pierre Fraigniaud and Amos Korman and David Peleg

Abstract: A central theme in distributed network algorithms concerns understanding and coping with the issue of locality. Despite considerable progress, research efforts in this direction have not yet resulted in a solid basis in the form of a fundamental computational complexity theory for locality. Inspired by sequential complexity theory, we focus on a complexity theory for distributed decision problems. In the context of locality, solving a decision problem requires the processors to independently inspect their local neighborhoods and then collectively decide whether a given global input instance belongs to some specified language.

We consider the standard LOCAL model of computation and define $\text{LD}(t)$ (for local decision) as the class of decision problems that can be solved in t communication rounds. We first study the intriguing question of whether randomization helps in local distributed computing, and to what extent.

Specifically, we define the corresponding randomized class $\text{BPLD}(t,p,q)$, containing all languages for which there exists a randomized algorithm that runs in t rounds, accepts correct instances with probability at least p and rejects incorrect ones with probability at least q . We show that $p^2+q = 1$ is a threshold for the containment of $\text{LD}(t)$ in $\text{BPLD}(t,p,q)$. More precisely, we show that there exists a language that does not belong to $\text{LD}(t)$ for any $t=o(n)$ but does belong to $\text{BPLD}(0,p,q)$ for any $p,q \in (0,1]$ such that $p^2+q \leq 1$. On the other hand, we show that, restricted to hereditary languages, $\text{BPLD}(t,p,q) = \text{LD}(O(t))$, for any function t and any $p,q \in (0,1]$ such that $p^2+q > 1$.

In addition, we investigate the impact of non-determinism on local decision, and establish some structural results inspired by classical computational complexity theory.

Specifically, we show that non-determinism does help, but that this help is limited, as there exist languages that cannot be decided non-deterministically. Perhaps surprisingly, it turns out that it is the combination of randomization with non-determinism that enables to decide all languages in constant time. Finally, we introduce the notion of local reduction, and establish some completeness results.

5.

120. A Small PRG for Polynomial Threshold Functions of Gaussians

Daniel M. Kane

Abstract: We develop a new pseudo-random generator for fooling arbitrary degree- d polynomial threshold functions with respect to the Gaussian distribution. Our generator fools such functions to within ϵ with a generator of seed length $\log(n)2^{O(d)}\epsilon^{-4-c}$, where c is an arbitrarily small positive constant.

6.

126. Evolution with Recombination

Varun Kanade

Abstract: Valiant (2007) introduced a computational model of evolution and suggested that Darwinian evolution be studied in the framework of computational learning theory. Valiant describes evolution as a restricted form of learning where exploration is limited to a set of possible mutations and feedback is received by the survival of the fittest mutant. In subsequent work Feldman (2008) showed that evolvability in Valiant's model is equivalent to learning in the correlational statistical query (CSQ) model. We extend Valiant's model to include genetic recombination and show that in certain cases, recombination can significantly speed-up the process of evolution in terms of the number of generations, though at the expense of population size. This follows by a reduction from parallel -CSQ algorithms to evolution with recombination. This gives an exponential speed-up (in terms of the number of generations) over previous known results for evolving conjunctions and halfspaces with respect to restricted distributions.

7.

128. Extractors for circuit sources

Emanuele Viola

Abstract: We obtain the first deterministic extractors for sources generated (or sampled) by small circuits of bounded depth. Our main results are:

(1) We extract $k (k/nd)^{O(1)}$ bits with exponentially small error from n -bit sources of min-entropy k that are generated by functions $f : \{0,1\}^{\ell} \rightarrow \{0,1\}^n$ where each output bit depends on $\leq d$ input bits. In particular, we extract from NC^0 sources, corresponding to $d = O(1)$.

(2) We extract $k (k/n^{1+\gamma})^{O(1)}$ bits with super-polynomially small error from n -bit sources of min-entropy k that are generated by $\text{poly}(n)$ -size AC^0 circuits, for any $\gamma > 0$.

As our starting point, we revisit the connection by Trevisan and Vadhan (FOCS 2000) between circuit lower bounds and extractors for sources generated by circuits. We note

that such extractors (with very weak parameters) are equivalent to lower bounds for generating distributions (FOCS 2010; with Lovett, CCC 2011). Building on those bounds, we prove that the sources in (1) and (2) are (close to) a convex combination of high-entropy "bit-block"

sources. Introduced here, such sources are a special case of affine ones. As extractors for (1) and (2) one can use the extractor for low-weight affine sources by Rao (CCC 2009).

Along the way, we exhibit an explicit boolean function $b : \{0,1\}^n \rightarrow \{0,1\}$ such that $\text{poly}(n)$ -size AC^0 circuits cannot generate the distribution $(x, b(x))$, solving a problem about the complexity of distributions.

Independently, De and Watson (ECCC TR11-037) obtain a result similar to (1) in the special case $d = o(\lg n)$.

8.

133. The Second-Belief Mechanism.

Jing Chen and Silvio Micali

Abstract: In settings of incomplete information, we put forward a very conservative ---indeed, purely set-theoretic--- model of the knowledge that the players may have about the types of their opponents. Yet, we prove that such knowledge can be successfully and robustly leveraged by means of a solution concept relying on very weak assumptions: in essence, via extensive-form mechanisms under "mutual" knowledge of rationality.

We demonstrate the potential of our approach in auctions of a single good by

1. considering a new revenue benchmark, always lying between the highest and second-highest valuation,
2. proving that no classical mechanism can even slightly approximate it in any robust way, and
3. providing a new mechanism that perfectly and robustly achieves it, with the extra property that the good will always be sold out at the end of the auction.

Our impossibility result for robustly implementing our revenue benchmark applies not only to implementation in dominant strategies, but also to any implementation "at equilibrium", as well as to implementation in undominated strategies.

9.

140. New extension of the Weil bound for character sums with applications to coding

Tali Kaufman and Shachar Lovett

Abstract: The Weil bound for character sums is a deep result in Algebraic Geometry with many applications both in mathematics and in the theoretical computer science. The Weil bound states that for any polynomial $f(x)$ over a finite field \mathbb{F} and any additive character $\chi: \mathbb{F} \rightarrow \mathbb{C}$, either $\chi(f(x))$ is a constant function or it is distributed close to uniform. The Weil bound is quite effective as long as $\deg(f) \ll \sqrt{|\mathbb{F}|}$, but it breaks down when the degree of f exceeds $\sqrt{|\mathbb{F}|}$. As the Weil bound plays a central role in many areas, finding extensions for polynomials of larger degree is an important problem with many possible applications.

In this work we develop such an extension over finite fields \mathbb{F}_{p^n} of small characteristic: we prove that if $f(x) = g(x) + h(x)$ where $\deg(g) \ll \sqrt{|\mathbb{F}|}$ and $h(x)$ is a sparse polynomial of arbitrary degree but bounded weight degree, then the same conclusion of the classical Weil bound still holds: either $\chi(f(x))$ is constant or its distribution is close to uniform. In particular, this shows that the subcode of Reed-Muller codes of degree $\omega(1)$ generated by traces of sparse polynomials is a code with near optimal distance, while Reed-Muller of such a degree has no distance (i.e. $\Theta(1)$ distance); this is one of the few examples where one can prove that sparse polynomials behave differently from non-sparse polynomials of the same degree.

As an application we prove new general results for affine invariant codes. We prove that any affine-invariant subspace of quasi-polynomial size is (1) indeed a code (i.e. has good distance) and (2) is locally testable. Previous results for general affine invariant codes were known only for codes of polynomial size, and of length 2^n where n needed to be a prime. Thus, our techniques are the first to extend to general families of such codes of super-polynomial size, where we also remove the requirement from n to be a prime. The proof is based on two main ingredients: the extension of the Weil bound for character sums, and a new Fourier-analytic approach for estimating the weight distribution of general codes with large dual distance, which may be of independent interest.

10.

144. A Two Prover One Round Game with Strong Soundness

Subhash Khot and Muli Safra

Abstract: We show that for any fixed prime $q \geq 5$ and constant $\zeta > 0$, it is NP-hard to distinguish whether a two prover one round game with q^6 answers has value at least $1 - \zeta$ or at most $\frac{4}{q}$.

The result is obtained by combining two techniques: (i) An Inner PCP based on the *point versus subspace* test for linear functions. The test is analyzed Fourier analytically. (ii) The Outer/Inner PCP composition that relies on a certain *sub-code covering* property for Hadamard codes. This is a new and essentially black-box method to translate a *codeword test* for Hadamard codes to a *consistency test*, leading to a full PCP construction.

As an application, we show that unless NP has quasi-polynomial time deterministic algorithms, the Quadratic Programming Problem is inapproximable within factor $(\log n)^{1/6 - o(1)}$.

11.

145. Optimal testing of multivariate polynomials over small prime fields

Elad Haramaty and Amir Shpilka and Madhu Sudan

Abstract: We consider the problem of testing if a given function $f: F_q^n \rightarrow F_q$ is close to a n -variate degree d polynomial over the finite field F_q of q elements. The natural, low-query, test for this property would be to pick the smallest dimension $t = t_{q,d} \sim d/q$ such that every function of degree greater than d reveals this feature on some t -dimensional affine subspace of F_q^n and to test that f when restricted to a random t -dimensional affine subspace is a polynomial of degree at most d on this subspace. Such a test makes only q^t queries, independent of n .

Previous works, by Alon et al. (AKKLR), and Kaufman & Ron and Jutla et al., showed that this natural test rejected functions that were $\Omega(1)$ -far from degree d -polynomials with probability at least $\Omega(q^{-t})$

(the results of Kaufman & Ron hold for all fields F_q , while the results of Jutla et al. hold only for fields of prime order). Thus to get a constant probability of detecting functions that were at constant distance from the space of degree d polynomials, the tests made q^{2t} queries. Kaufman & Ron also noted that when q is prime, then q^t queries are necessary. Thus these tests were off by at least a quadratic factor from known lower bounds.

It was unclear if the soundness analysis of these tests were tight and this question relates closely to the task of understanding the behavior of the Gowers Norm. This motivated the work of Bhattacharyya et al., who gave an optimal analysis for the case of the binary field and showed that the natural test actually rejects functions that were $\Omega(1)$ -far from degree d -polynomials with probability at least $\Omega(1)$.

In this work we give an optimal analysis of this test for all fields showing that the natural test does indeed reject functions that are $\Omega(1)$ -far from degree d polynomials with $\Omega(1)$ -probability. Our analysis thus shows that this test is optimal (matches known lower bounds) when q is prime. (It is also potentially best possible for all fields.) Our approach extends the proof technique of Bhattacharyya et al., however it has to overcome many technical barriers in the process. The natural extension of their analysis leads to an $O(q^d)$ query complexity, which is worse than that of Kaufman and Ron for all q except 2! The main technical ingredient in our work is a tight analysis of the number of "hyperplanes" (affine subspaces of co-dimension 1) on which the restriction of a degree d polynomial has degree less than d . We show that the number of such hyperplanes is at most $O(q^{t_{q,d}})$ - which is tight to within constant factors.

12.

150. Fully dynamic maximal matching in $O(\log n)$ update time

Surender Baswana and Manoj Gupta and Sandeep Sen

Abstract: We present an algorithm for maintaining maximal matching in a graph under addition and deletion of edges. Our data structure is randomized that takes $O(\log n)$ expected amortized time for each edge update where n is the number of vertices in the graph. While there is a trivial $O(n)$ algorithm for edge update, the previous best known result for this problem for a graph with n vertices and m edges is $O((n+m)^{0.7072})$ which is sub-linear only for a sparse graph. To the best of our knowledge this

is the first polylog update time for maximal matching that implies an exponential improvement from the previous results.

For the related problem of maximum matching, Onak and Rubinfeld \cite{onak} designed a randomized data structure that achieves $O(\log^2 n)$ amortized time for each update for maintaining a c -approximate maximum matching for some large constant c .

In contrast, we can maintain a factor two approximate maximum matching in $O(\log n)$ expected time per update as a direct corollary of the maximal matching scheme. This in turn also implies a two approximate vertex cover maintenance scheme that takes $O(\log n)$ expected time per update.

13.

151. Optimal bounds for quantum bit commitment

André Chailloux and Iordanis Kerenidis

Abstract: Bit commitment is a fundamental cryptographic primitive with numerous applications. Quantum information allows for bit commitment schemes in the information theoretic setting where no dishonest party can perfectly cheat. The previously best-known quantum protocol by Ambainis achieved a cheating probability of at most $3/4$. On the other hand, Kitaev showed that no quantum protocol can have cheating probability less than $1/\sqrt{2}$ (his lower bound on coin flipping can be easily extended to bit commitment). Closing this gap has since been an important open question.

In this paper, we provide the optimal bound for quantum bit commitment. First, we show a lower bound of approximately 0.739, improving Kitaev's lower bound. For this, we present some generic cheating strategies for Alice and Bob and conclude by proving a new relation between the trace distance and fidelity of two quantum states. Second, we present an optimal quantum bit commitment protocol which has cheating probability arbitrarily close to 0.739. More precisely, we show how to use any weak coin flipping protocol with cheating probability $1/2 + \epsilon$ in order to achieve a quantum bit commitment protocol with cheating probability $0.739 + O(\epsilon)$. We then use the optimal quantum weak coin flipping protocol described by Mochon. Last, in order to stress the fact that our protocol uses quantum effects beyond the weak coin flip, we show that any classical bit commitment protocol with access to perfect weak (or strong) coin flipping has cheating probability at least $3/4$.

14.

152. SHARP MIXING TIME BOUNDS FOR SAMPLING RANDOM SURFACES

PIETRO CAPUTO AND FABIO MARTINELLI AND FABIO LUCIO
TONINELLI

Abstract: We analyze the mixing time of a natural local Markov Chain (Gibbs sampler) for two commonly studied models of random surfaces: (i) discrete monotone surfaces in Z^3 with “almost planar” boundary conditions and (ii) the one-dimensional discrete Solid-on-Solid (SOS) model. In both cases we prove the first almost optimal bounds $O(L^2 \text{polylog}(L))$ where L is the size of the system. Our proof is inspired by the so-called “mean curvature” heuristic: on a large scale, the dynamics should approximate a deterministic motion in which each point of the surface moves according to a drift proportional to the local inverse mean curvature radius. Key technical ingredients are monotonicity, coupling and an argument due to D. Wilson [17] in the framework of lozenge tiling Markov Chains. The novelty of our approach with respect to previous results consists in proving that, with high probability, the dynamics is dominated by a deterministic evolution which, apart from $\text{polylog}(L)$ corrections, follows the mean curvature prescription. Our method works equally well for both models despite the fact that their equilibrium maximal deviations from the average height profile occur on very different scales ($\log L$ for monotone surfaces and $\Theta(L)$ for the SOS model).

15.

154. Solving connectivity problems parameterized by treewidth in single exponential time

Marek Cygan and Jesper Nederlof and Marcin Pilipczuk and Michał Pilipczuk and Johan M. M. van Rooij and Jakub Onufry Wojtaszczyk

Abstract: For the vast majority of local problems on graphs of small treewidth (where by local we mean that a solution can be verified by checking separately the neighbourhood of each vertex), standard dynamic programming techniques give $c^{tw} |V|^{O(1)}$ time algorithms, where tw is the treewidth of the input graph $G = (V;E)$ and c is a constant. On the other hand, for problems with a global requirement (usually connectivity) the best-known algorithms were naive dynamic programming schemes running in at least tw^{tw} time.

We breach this gap by introducing a technique we named Cut&Count that allows to produce $c^{tw} |V|^{O(1)}$ time Monte Carlo algorithms for most connectivity-type problems, including HAMILTONIAN PATH, STEINER TREE, FEEDBACK VERTEX SET and CONNECTED DOMINATING SET. These results have numerous consequences in various fields, like parameterized complexity, exact and approximate algorithms on planar and H-minor-free graphs and exact algorithms on graphs of bounded degree. In all these fields we are able to improve the best-known results for some problems. Also, looking from a more theoretical perspective, our results are surprising since the equivalence relation that partitions all partial solutions with respect to extendability to global solutions seems to consist of at least tw^{tw} equivalence classes for all these problems. Our results answer an open problem raised by Lokshantov, Marx and Saurabh [SODA'11].

In contrast to the problems aiming to minimize the number of connected components that we solve using Cut&Count as mentioned above, we show that, assuming the Exponential Time Hypothesis, the aforementioned gap cannot be breached for some problems that aim to maximize the number of connected components like CYCLE PACKING.

The constant c in our algorithms is in all cases small (at most 4 for undirected problems and at most 6 for directed ones), and in several cases we are able to show that improving those constants would cause the Strong Exponential Time Hypothesis to fail.

16.

161. How to Play Unique Games Against a Semi-Random Adversary

Alexandra Kolla and Konstantin Makarychev and Yury Makarychev

Abstract: In this paper, we study the average case complexity of the Unique Games problem.

We propose a natural semi-random model, in which a unique game instance is generated in several steps. First an adversary selects a completely satisfiable instance of Unique Games, then she chooses an epsilon fraction of all edges, and finally replaces ("corrupts") the constraints corresponding to these edges with new constraints. If all steps are adversarial, the adversary can obtain

any (1-epsilon) satisfiable instance, so then the problem is as hard as in the worst case. We show that known algorithms for unique games (in particular, all algorithms that use the standard SDP relaxation) fail to solve semi-random instances of Unique Games.

We present an algorithm that with high probability finds a solution satisfying a (1-delta) fraction of all constraints in semi-random instances (we require that the average degree of the graph is $\Omega(\log k)$). To this end, we consider a new non-standard SDP program for Unique Games, which is not a relaxation for the problem, and show how to analyze it. We present a new rounding scheme that simultaneously uses SDP and LP solutions, which we believe is of independent interest.

Finally, we study semi-random instances of Unique Games that are at most (1-epsilon) satisfiable. We present an algorithm that distinguishes between the case when the instance is a semi-random instance and the case when the instance is an (arbitrary) (1-delta)-satisfiable instance if $\epsilon > c \delta$ (for some absolute constant c).

17.

162. Near-Optimal Column-Based Matrix Reconstruction

Christos Boutsidis and Petros Drineas and Malik Magdon-Ismail

Abstract: We consider low-rank reconstruction of a matrix using its columns and we present asymptotically optimal algorithms for both spectral norm and Frobenius norm reconstruction. The main tools we introduce to obtain our results are: (i) the use of fast approximate SVD-like decompositions for column reconstruction, and (ii) two deterministic algorithms for selecting rows from matrices with orthonormal columns, building upon the sparse representation theorem for decompositions of the identity that appeared in [BSS09](#).

18.

172. Tight lower bounds for 2-query LCCs over finite fields

Arnab Bhattacharyya and Zeev Dvir and Shubhangi Saraf and Amir Shpilka

Abstract: A Locally Correctable Code (LCC) is an error correcting code that has a probabilistic self-correcting algorithm that, with high probability, can correct any coordinate of the codeword by looking at only a few other coordinates, even if a fraction δ of the coordinates are corrupted. LCC's are a stronger form of LDCs (Locally Decodable Codes) which have received a lot of attention recently due to their many applications and surprising constructions.

In this work we show a separation between 2-query LDCs and LCCs over finite fields of prime order. Specifically, we prove a lower bound of the form $p^{\Omega(\delta d)}$ on the length of linear 2 -query LCCs over \mathbb{F}_p , that encode messages of length d . Our bound improves over the known bound of $2^{\Omega(\delta d)}$ [GKST06, KdW04, DS07] which is tight for LDCs. Our proof makes use of tools from additive combinatorics which have played an important role in several recent results in Theoretical Computer Science.

We also obtain, as corollaries of our main theorem, new results in incidence geometry over finite fields. The first is an improvement to the Sylvester-Gallai theorem over finite fields [SS10] and the second is a new analog of Beck's theorem over finite fields.

19.

173. Separator Theorems for Minor-Free and Shallow Minor-Free Graphs with Applications

Christian Wulff-Nilsen

Abstract: Alon, Seymour, and Thomas generalized Lipton and Tarjan's planar separator theorem and showed that a K_h -minor free graph with n vertices has a separator of size at most $h^{3/2}\sqrt{n}$. They gave an algorithm that, given a graph G with m edges and n vertices and given an integer $h \geq 1$, outputs in $O(\sqrt{hn}m)$ time such a separator or a K_h -minor of G . Plotkin, Rao, and Smith gave an $O(hm\sqrt{n\log n})$ time algorithm to find a separator of size $O(h\sqrt{n\log n})$. Kawarabayashi and Reed improved the bound on the size of the separator to $h\sqrt{n}$ and gave an algorithm that finds such a separator in $O(n^{1+\epsilon})$ time for any constant $\epsilon > 0$, assuming h is constant. This algorithm has an extremely large dependency on h in the running time (some power tower of h whose height is itself a function of h), making it impractical even for small h . We are interested in a small polynomial time dependency on h and we show how to find an $O(h\sqrt{n\log n})$ -size separator or report that G has a K_h -minor in $O(\text{poly}(h)n^{5/4+\epsilon})$ time for any constant $\epsilon > 0$. We also present the first $O(\text{poly}(h)n)$ time algorithm to find a separator of size $O(n^c)$ for a constant $c < 1$. As corollaries of our results, we get improved algorithms for shortest paths and maximum matching. Furthermore, for integers ℓ and h , we give an $O(m + n^{2+\epsilon/\ell})$ time algorithm that either produces a K_h -minor of depth $O(\ell\log n)$ or a separator of size at most $O(n^{1/\ell} + \ell h^2\log n)$. This improves the shallow minor algorithm of Plotkin, Rao, and Smith when $m = \Omega(n^{1+\epsilon})$. We get a similar running time improvement for an approximation algorithm for the problem of finding a largest K_h -minor in a given graph.

20.

175. 3-SAT Faster and Simpler - Unique-SAT Bounds for PPSZ Hold in General

Timon Hertli

Abstract: The PPSZ algorithm by Paturi, Pudl'ak, Saks, and Zane [1998] is the fastest known algorithm for Unique k-SAT, where the input formula does not have more than one satisfying assignment. For $k \geq 5$ the same bounds hold for general k-SAT. We show that this is also the case for $k=3,4$, using a slightly modified PPSZ algorithm. We do the analysis by defining a cost for satisfiable CNF formulas, which we prove to decrease in each PPSZ step by a certain amount. This improves our previous best bounds with Moser and Scheder [2011] for 3-SAT to $O(1.308^n)$ and for 4-SAT to $O(1.469^n)$.

21.

177. On Range Searching in the Group Model and Combinatorial Discrepancy

Kasper Green Larsen

Abstract: In this paper we establish an intimate connection between dynamic range searching in the group model and combinatorial discrepancy. Our result states that, for a broad class of range searching data structures (including all known upper bounds), it must hold that $t_u t_q = \Omega(\text{disc}^2 / \lg n)$ where t_u is the worst case update time, t_q the worst case query time and disc is the combinatorial discrepancy of the range searching problem in question. This relation immediately implies a whole range of exceptionally high and near-tight lower bounds for all of the basic range searching problems. We list a few of them in the following:

$\begin{itemize}$

\item For halfspace range searching in d -dimensional space, we get a lower bound of $t_u t_q = \Omega(n^{1-1/d} / \lg n)$. This comes within a $\lg n \lg \lg n$ factor of the best known upper bound.

\item For orthogonal range searching in d -dimensional space, we get a lower bound of $t_u t_q = \Omega(\lg^{d-2+\mu(d)} n)$, where $\mu(d) > 0$ is some small but strictly positive function of d .

\item For ball range searching in d -dimensional space, we get a lower bound of $t_u t_q = \Omega(n^{1-1/d} / \lg n)$.

$\end{itemize}$

We note that the previous highest lower bound for any explicit problem, due to Patrascu [STOC'07], states that $t_q =$

$\Omega((\lg n \lg(\lg n + t_u))^2)$, which does however hold for a less restrictive class of data structures.

Our result also has implications for the field of combinatorial discrepancy. Using textbook range searching solutions, we improve on the best known discrepancy upper bound for axis-aligned rectangles in dimensions $d \geq 3$.

22.

179. Coin Flipping with Constant Bias Implies One-Way Functions

Iftach Haitner and Eran Omri

Abstract: It is well known (cf Impagliazzo and Luby [FOCS '89]) that the existence of almost all "interesting" cryptographic applications, i.e. ones that cannot hold information theoretically, implies one-way functions. An important exception where the above implication is not known, however, is the case of coin-flipping protocols. Such protocols allow honest parties to mutually flip an unbiased coin, while guaranteeing that even a cheating (efficient) party cannot bias the output of the protocol by much. While Impagliazzo and Luby proved that coin-flipping protocols that are safe against negligible bias do imply one-way functions, and, very recently, Maji, Prabhakaran, Sahai and Schreiber [FOCS '10] proved the same for constant-round protocols (with any non-trivial bias). For the general case, however, no such implication was known.

We make a significant step towards answering the above fundamental question, showing that coin-flipping protocols safe against a constant bias (concretely, $\frac{\sqrt{2} - 1}{2}$) imply one-way functions.

23.

187. The Complexity of the Homotopy Method, Equilibrium Selection, and Lemke-Howson Solutions.

Paul W. Goldberg and Christos H. Papadimitriou and Rahul Savani

Abstract: We show that the widely used homotopy method for solving fixpoint problems, as well as the Harsanyi-Selten equilibrium selection process for games, are PSPACE-complete to implement. Extending our result for the Harsanyi-Selten process, we show that several other homotopy-based algorithms for solving games are also

PSPACE-complete to implement. A further application of our techniques yields the result that it is PSPACE-complete to compute any of the equilibria that could be found via the classical Lemke-Howson algorithm, a complexity-theoretic strengthening of the result in [24]. These results show that our techniques can be widely applied and suggest that the PSPACE-completeness of implementing homotopy methods is a general principle.

24.

193. Information Equals Amortized Communication

Mark Braverman and Anup Rao

Abstract: We show how to efficiently simulate the sending of a message M to a receiver who has partial information about the message, so that the expected number of bits communicated in the simulation is close to the amount of additional information that the message reveals to the receiver. This is a generalization and strengthening of the Slepian-Wolf theorem, which shows how to carry out such a simulation with low *amortized* communication in the case that M is a deterministic function of X . A caveat is that our simulation is interactive.

As a consequence, we obtain new relationships between the randomized amortized communication complexity of a function, and its information complexity. We prove that for any given distribution on inputs, the internal information cost (namely the information revealed to the parties) involved in computing any relation or function using a two party interactive protocol is *exactly* equal to the amortized communication complexity of computing independent copies of the same relation or function. Here by amortized communication complexity we mean the average per copy communication in the best protocol for computing multiple copies, with a bound on the error in each copy (i.e. we require only that the output in each coordinate is correct with good probability, and do not require that all outputs are simultaneously correct). This significantly simplifies the relationships between the various measures of complexity for average case communication protocols, and proves that if a function's information cost is smaller than its communication complexity, then multiple copies of the function can be computed more efficiently in parallel than sequentially.

Finally, we show that the only way to prove a strong direct sum theorem for randomized communication complexity is by solving a particular variant of the pointer jumping problem that we define. If this problem has a cheap communication protocol, then a strong direct sum theorem must hold. On the other hand, if it does not, then the problem itself gives a counterexample for the direct sum question. In the process we show that a strong direct sum theorem for communication complexity holds if and only if efficient compression of communication protocols is possible.

25.

194. Graph Connectivities, Network Coding, and Expander Graphs

Ho Yee Cheung and Lap Chi Lau and Kai Man Leung

Abstract: In this paper we present a new algebraic formulation to compute edge connectivities in a directed graph, using the ideas developed in network coding. This reduces the problem of computing edge connectivities to solving systems of linear equations, thus allowing us to use tools in linear algebra to design new algorithms. Using the algebraic formulation we obtain faster algorithms for computing single source edge connectivities and all pairs edge connectivities, in some settings the amortized time to compute the edge connectivity for one pair is sublinear. Through this connection, we have also found an interesting use of expanders and superconcentrators to design fast algorithms for some graph connectivity problems.

26.

202. Limitations of Randomized Mechanisms for Combinatorial Auctions

Shaddin Dughmi and Jan Vondrak

Abstract: The design of computationally efficient and incentive compatible mechanisms that solve or approximate fundamental resource allocation problems is the main goal of algorithmic mechanism design. A central example in both theory and practice is welfare-maximization in combinatorial auctions. Recently, a randomized mechanism has been discovered for combinatorial auctions that is truthful in expectation and guarantees a $(1-1/e)$ -approximation to the optimal social welfare when players have coverage valuations \cite{DRY11}. This approximation ratio is the best possible even for non-truthful algorithms, assuming $P \neq NP$ \cite{KLMM05}.

Given the recent sequence of negative results for combinatorial auctions under more restrictive notions of incentive compatibility \cite{DN07,BDFKMPSSU10,Dobzin11}, this development raises a natural question: Are truthful-in-expectation mechanisms compatible with polynomial-time approximation in a way that deterministic or universally truthful mechanisms are not? In particular, can polynomial-time truthful-in-expectation mechanisms guarantee a near-optimal approximation ratio for more general variants of combinatorial auctions?

We prove that this is not the case. Specifically, the result of \cite{DRY11} cannot be extended to combinatorial auctions with submodular valuations in the value oracle model.

(Absent strategic considerations, a $(1-1/e)$ -approximation is still achievable in this setting \cite{V08}.) More precisely, we prove that there is a constant $\gamma > 0$ such that there is no randomized mechanism that is truthful-in-expectation --- or even approximately truthful-in-expectation --- and guarantees an $m^{-\gamma}$ -approximation to the optimal social welfare for combinatorial auctions with submodular valuations in the value oracle model.

We also prove an analogous result for the flexible combinatorial public projects (CPP) problem, where a truthful-in-expectation $(1-1/e)$ -approximation for coverage valuations has been recently developed \cite{Dughmi11}. We show that there is no truthful-in-expectation --- or even approximately truthful-in-expectation --- mechanism that achieves an $m^{-\gamma}$ -approximation to the optimal social welfare for combinatorial public projects with submodular valuations in the value oracle model. Both our results present an unexpected separation between coverage functions and submodular functions, which does not occur for these problems without strategic considerations.

27.

204. How to Store a Secret on Continually Leaky Devices

Yevgeniy Dodis and Allison Lewko and Brent Waters and Daniel Wichs

Abstract: We consider the question of how to store a value secretly on devices that continually leak information about their internal state to an external attacker. If the secret value is stored on a single device, and the attacker can leak even a single predicate of the internal state of that device, then she may learn some information about the secret value itself. Therefore, we consider a setting where the secret value is shared between multiple devices (or multiple components of one device), each of which continually leaks arbitrary adaptively chosen predicates of its individual state. Since leakage is continual, each device must also continually update its state so that an attacker cannot just leak it entirely one bit at a time. In our model, the devices update their state individually and asynchronously, without any communication between them. The update process is necessarily randomized, and its randomness can leak as well.

As our main result, we construct a sharing scheme for two devices, where a constant fraction of the internal state of each device can leak in between and during updates. Our scheme has the structure of a public-key encryption, where one share is a secret key and the other is a ciphertext. As a contribution of independent interest, we also get public-key encryption in the continual leakage model, introduced by Brakerski et al. and Dodis et al. (FOCS '10). This scheme tolerates continual leakage on the secret key and the updates, and simplifies the recent construction of Lewko, Lewko and Waters (STOC '11). For our main result, we also show how to update the ciphertexts of the encryption scheme so that the message remains hidden even if an attacker interleaves leakage on secret key and

ciphertext shares. The security of our scheme is based on the linear assumption in prime-order bilinear groups.

We also provide an extension to general access structures realizable by linear secret sharing schemes across many devices. The main advantage of this extension is that the state of some devices can be compromised entirely, while that of the all remaining devices is susceptible to continual leakage.

Lastly, we show impossibility of information theoretic sharing schemes in our model, where continually leaky devices update their state individually.

28.

215. A Polylogarithmic-Competitive Algorithm for the k -Server Problem

Nikhil Bansal and Niv Buchbinder and Aleksander Madry and Seffi Naor

Abstract: We give the first polylogarithmic-competitive randomized algorithm for the k -server problem on an arbitrary finite metric space.

In particular, our algorithm achieves a competitive ratio of $\widetilde{O}(\log^3 n \log^2 k)$ for any metric space on n points.

This improves upon the $(2k-1)$ -competitive algorithm of Koutsoupias and Papadimitriou (J.ACM.'95) whenever n is sub-exponential in k .

29.

217. Minimum Weight Cycles and Triangles: Equivalences and Algorithms

Liam Roditty and Virginia Vassilevska Williams

Abstract: We consider the fundamental algorithmic problem of finding a cycle of minimum weight in a weighted graph.

In particular, we show that the minimum weight cycle problem in an undirected n -node graph with edge weights in $\{1, \dots, M\}$ or in a directed

n -node graph with edge weights in $\{-M, \dots, M\}$ and no negative cycles can be efficiently reduced to finding a minimum weight $\{\text{triangle}\}$ in an $\Theta(n)$ -node $\{\text{undirected}\}$ graph with weights in $\{1, \dots, O(M)\}$. Roughly speaking, our reductions imply the following surprising phenomenon: a minimum cycle with an arbitrary number of weighted edges can be "encoded" using only $\{\text{three}\}$ edges within roughly the same weight interval!

This resolves a longstanding open problem posed in a seminal work by Itai and Rodeh [SIAM J. Computing 1978 and STOC'77] on minimum cycle in unweighted graphs.

A direct consequence of our efficient reductions are $\tilde{O}(Mn^{\omega}) \leq \tilde{O}(Mn^{2.376})$ -time algorithms using fast matrix multiplication (FMM) for finding a minimum weight cycle in both undirected graphs with integral weights from the interval $[1, M]$ and directed graphs with integral weights from the interval $[-M, M]$. The latter seems to reveal a strong separation between the all pairs shortest paths (APSP) problem and the minimum weight cycle problem in directed graphs as the fastest known APSP algorithm has a running time of $O(M^{0.681}n^{2.575})$ by Zwick [J. ACM 2002].

In contrast, when only combinatorial algorithms are allowed (that is, without FMM) the only known solution to minimum weight cycle is by computing APSP. Interestingly, any separation between the two problems in this case would be an amazing breakthrough as by a recent paper by Vassilevska W. and Williams [FOCS'10], any $O(n^{3-\epsilon})$ -time algorithm ($\epsilon > 0$) for minimum weight cycle immediately implies a $O(n^{3-\delta})$ -time algorithm ($\delta > 0$) for APSP.

30.

218. Streaming Algorithms via Precision Sampling

Alexandr Andoni and Robert Krauthgamer and Krzysztof Onak

Abstract: A technique introduced by Indyk and Woodruff [STOC 2005] has inspired several recent advances in data-stream algorithms.

We show that a number of these results follow easily from the application of a single probabilistic method called *Precision Sampling*.

Using this method, we obtain simple data-stream algorithms that maintain a randomized sketch of an input vector $x = (x_1, \dots, x_n)$, which is useful for the following applications:

$\begin{itemize}$

\item

Estimating the F_k -moment of x , for $k > 2$.

\item

Estimating the ℓ_p -norm of x , for $p \in [1, 2]$, with small update time.

\item

Estimating cascaded norms $\ell_p(\ell_q)$ for all $p, q > 0$.

\item

ℓ_1 sampling, where the goal is to produce an element i with probability (approximately) $|x_i|/|x|_1$. It extends to similarly defined ℓ_p -sampling, for $p \in [1, 2]$.

For all these applications the algorithm is essentially the same: pre-multiply the vector x entry-wise by a well-chosen random vector, and run a heavy-hitter estimation algorithm on the resulting vector. Our sketch is a linear function of x , thereby allowing general updates to the vector x .

Precision Sampling itself addresses the problem of estimating a sum $\sum_{i=1}^n a_i$ from weak estimates of each real $a_i \in [0, 1]$. More precisely, the estimator first chooses a desired precision $u_i \in (0, 1]$ for each $i \in [n]$, and then it receives an estimate of every a_i within additive u_i . Its goal is to provide a good approximation to $\sum a_i$ while keeping a tab on the cost $\sum_i (1/u_i)$. Here we refine previous work [Andoni, Krauthgamer, and Onak, FOCS 2010] which shows that as long as $\sum a_i = \Omega(1)$, a good multiplicative approximation can be achieved using total precision of only $O(n \log n)$.

31.

219. A Parallel Approximation Algorithm for Positive Semidefinite Programming

Rahul Jain and Penghui Yao

Abstract: Positive semidefinite programs are an important subclass of semidefinite programs in which all matrices involved in the specification of the problem are positive semidefinite and all scalars involved are non-negative. We present a parallel algorithm, which given an instance of a positive semidefinite program of size N and an approximation factor $\epsilon > 0$, runs in (parallel) time $\text{poly}(1/\epsilon) \text{polylog}(N)$, using $\text{poly}(N)$ processors, and outputs a value which is within multiplicative factor of $(1+\epsilon)$ to the optimal. Our result generalizes analogous result of Luby and Nisan [1993] for positive linear programs and our algorithm is inspired by their algorithm.

32.

222. Planar Graphs: Random Walks and Bipartiteness Testing

Artur Czumaj and Morteza Monemizadeh and Krzysztof Onak and Christian Sohler

Abstract: We initiate the study of the testability of properties in arbitrary planar graphs. We prove that bipartiteness can be tested in constant time. The previous bound for this class of graphs was $O(\sqrt{n})$, and the constant-time testability was only known for planar graphs with bounded degree. Previously used transformations of unbounded-degree sparse graphs into bounded-degree sparse graphs cannot be used to reduce the problem to the testability of bounded-degree planar graphs. Our approach extends to arbitrary minor-free graphs.

Our algorithm is based on random walks. The challenge is here to analyze random walks for a class of graphs that has good separators, i.e., bad expansion. Standard techniques that use a fast convergence to a uniform distribution do not work in this case. Roughly speaking, our analysis technique self-reduces the problem of finding an odd length cycle in a multigraph G induced by a collection of cycles to another multigraph G' induced by a set of shorter odd-length cycles, in such a way that when a random walk finds a cycle in G' with probability $p > 0$, then it does so with probability $\lambda(p) > 0$ in G . This reduction is applied until the cycles collapse to self-loops that can be easily detected.

33.

223. Pseudorandomness for read-once formulas

Andrej Bogdanov and Periklis Papakonstantinou and Andrew Wan

Abstract: We give an explicit construction of a pseudorandom generator for read-once formulas whose inputs can be read in arbitrary order. For formulas in n inputs and arbitrary gates of fan-in at most $d = O(n/\log n)$, the pseudorandom generator uses $(1 - \Omega(1))n$ bits of randomness and produces an output that looks $2^{-\Omega(n)}$ -pseudorandom to all such formulas.

Our analysis is based on the following lemma. Let $pr = Mz + e$, where M is the parity-check matrix of a sufficiently good binary error-correcting code of constant rate, z is a random string, e is a small-bias distribution, and all operations are modulo 2. Then for every pair of functions $f, g: \mathbb{B}^{n/2} \rightarrow \mathbb{B}$ and every equipartition $(I,$

$J)$ of $\{n\}$, the distribution ρ is pseudorandom for the pair $(f(x|_I), g(x|_J))$, where $x|_I$ and $x|_J$ denote the restriction of x to the coordinates in I and J , respectively.

34.

226. Approximating Graphic TSP by Matchings

Tobias Moemke and Ola Svensson

Abstract: We present a framework for approximating the metric TSP based on a novel use of matchings. Traditionally, matchings have been used to add edges in order to make a given graph Eulerian, whereas our approach also allows for the removal of certain edges leading to a decreased cost.

For the TSP on graphic metrics (graph-TSP), the approach yields a 1.461-approximation algorithm with respect to the Held-Karp lower bound. For graph-TSP restricted to a class of graphs that contains degree three bounded and claw-free graphs, we show that the integrality gap of the Held-Karp relaxation matches the conjectured ratio $4/3$. The framework allows for generalizations in a natural way and also leads to a 1.586-approximation algorithm for the traveling salesman path problem on graphic metrics where the start and end vertices are prespecified.

35.

230. Efficient Distributed Medium Access

Devavrat Shah and Jinwoo Shin and Prasad Tetali

Abstract: Consider a wireless network of n nodes represented by a graph $G=(V, E)$ where an edge (i,j) models the fact that transmissions of i and j interfere with each other, i.e. simultaneous transmissions of i and j become unsuccessful. Hence it is required that at each time instance a set of non-interfering nodes (corresponding to an independent set in G) access the wireless medium. To utilize wireless resources efficiently, it is required to arbitrate the access of medium among interfering nodes properly. Moreover, to be of practical use, such a mechanism is required to be totally distributed as well as simple.

As the main result of this paper, we provide such a medium access algorithm. It is randomized, totally distributed and simple: each node attempts to access medium at each time with probability that is a function of its local information. We establish efficiency of the algorithm by showing that the corresponding network Markov chain is positive recurrent as long as the demand imposed on the network can be supported by the wireless network (using any algorithm). In that sense, the proposed algorithm is optimal in terms of utilizing wireless resources. The algorithm is oblivious to the network graph structure, in contrast with the so-called 'polynomial back-off' algorithm by Hastad-Leighton-Rogoff (STOC '87, SICOMP '96) that is established to be optimal for the complete graph and bipartite graphs (by Goldberg-MacKenzie (SODA '96, JCSS '99)).

36.

232. Near Linear Lower Bound for Dimension Reduction in L1

Alexandr Andoni and Moses S. Charikar and Ofer Neiman and Huy L. Nguyen

Abstract: Given a set of n points in ℓ_1 , how many dimensions are needed to represent all pairwise distances within a specific distortion ?

This dimension-distortion tradeoff question is well understood for the ℓ_2 norm, where $O((\log n)^\epsilon)$ dimensions suffice to achieve $1+\epsilon$ distortion. In sharp contrast, there is a significant gap between upper and lower bounds for dimension reduction in ℓ_1 .

A recent result shows that distortion $1+\epsilon$ can be achieved with n^ϵ dimensions.

On the other hand, the only lower bounds known are that distortion δ requires $n^{\Omega(1/\delta^2)}$ dimension and that distortion $1+\epsilon$ requires $n^{1/2 - O(\epsilon \log(1/\epsilon))}$ dimensions.

In this work, we show the first near linear lower bounds for dimension reduction in ℓ_1 .

In particular, we show that $1+\epsilon$ distortion requires at least $n^{1 - O(1/\log(1/\epsilon))}$ dimensions.

Our proofs are combinatorial, but inspired by linear programming. In fact, our techniques lead to a simple combinatorial argument that is equivalent to the LP based proof of Brinkman-Charikar for lower bounds on dimension reduction in ℓ_1 .

37.

233. Maximum Edge-Disjoint Paths in Planar Graphs with Congestion 2

Loïc S\'equin-Charbonneau and F. Bruce Shepherd

Abstract: We study the maximum edge-disjoint path problem (medp) in planar graphs $G=(V,E)$. We are given a set of terminal pairs s_i, t_i , $i=1,2,\dots,k$ and wish to find a maximum routable subset of demands. That is, a subset of demands that can be connected by edge-disjoint paths. It is well-known that there is an integrality gap of $\Omega(\sqrt{n})$ for this problem even on a grid-like graph, and hence in planar graphs (Garg et al.). In contrast, Chekuri et al. show that for planar graphs, if sc LP is the optimal solution to the natural LP relaxation for medp , then there is a subset which is routable in $2G$ that is of size $\Omega(\text{opt}) / O(\log n)$. Subsequently they showed that $\Omega(\text{opt})$ is possible with congestion 4 (i.e., in $4G$) instead of 2 . We strengthen this latter result to show that a constant approximation is possible also with congestion 2 (and this is tight via the integrality gap grid example). We use a basic framework from work by Chekuri et al. At the heart of their approach is a 2-phase algorithm that selects an Okamura-Seymour instance. Each of their phases incurs a factor 2 congestion. It is possible to reduce one of the phases to have congestion 1. In order to achieve an overall congestion 2, however, the two phases must share capacity more carefully. For the Phase 1 problem, we extract a problem called rooted clustering that appears to be an interesting problem class in itself.

38.

237. Min-Max Graph Partitioning and Small-Set Expansion

Nikhil Bansal and Uriel Feige and Robert Krauthgamer and Konstantin Makarychev and Viswanath Nagarajan and Joseph (Seffi) Naor and Roy Schwartz

Abstract: We study graph partitioning problems from a min-max perspective, in which an input graph on n vertices should be partitioned into k parts, and the objective is to minimize the maximum number of edges leaving a single part. The

two main versions we consider are where the k parts need to be of equal-size, and where they must separate a set of k given terminals. We consider a common generalization of these two problems, and design for it an $O(\sqrt{\log n \log k})$ -approximation algorithm. This improves over an $O(\log^2 n)$ approximation for the second version due to Svitkina and Tardos [ST04], and roughly $O(k \log n)$ approximation for the first version that follows from other previous work. We also give an improved $O(1)$ -approximation algorithm for graphs that exclude any fixed minor.

Along the way, we study the ρ -Unbalanced Cut problem, whose goal is to find, in an input graph $G=(V,E)$, a set $S \subseteq V$ of size $|S|=\rho n$ that minimizes the number of edges leaving S . We provide a bicriteria approximation of $O(\sqrt{\log n \log(1/\rho)})$; when the input graph excludes a fixed-minor we improve this guarantee to $O(1)$. Note that the special case $\rho = 1/2$ is the well-known Minimum Bisection problem, and indeed our bounds generalize those of Arora, Rao and Vazirani [ARV08] and of Klein, Plotkin, and Rao [KPR93]. Our algorithms work also for the closely related Small Set Expansion (SSE) problem, which asks for a set $S \subseteq V$ of size $0 < |S| \leq \rho n$ with minimum edge-expansion, and was suggested recently by Raghavendra and Steurer [RS10]. In fact, our algorithm handles more general, weighted, versions of both problems. Previously, an $O(\log n)$ true approximation for both ρ -Unbalanced Cut and Small Set Expansion follows from Racke [Racke08].

39.

240. An FPTAS for #Knapsack and Related Counting Problems

Parikshit Gopalan and Adam Klivans and Raghu Meka and Daniel Stefankovic and Santosh Vempala and Eric Vigoda

Abstract: Given n elements with non-negative integer weights w_1, \dots, w_n and an integer capacity C , we consider the counting version of the classic knapsack problem: find the number of distinct subsets whose weights add up to at most C . We give the first deterministic, fully polynomial-time approximation scheme (FPTAS) for estimating the number of solutions to any knapsack constraint (our estimate has relative error $1 \pm \epsilon$). Our algorithm is based on dynamic programming. Previously, randomized polynomial-time approximation schemes (FPRAS) were known first by

Morris and Sinclair via Markov chain Monte Carlo techniques, and subsequently by Dyer via dynamic programming and rejection sampling.

In addition, we present a new method for deterministic approximate counting using read-once branching programs. Our approach yields an FPTAS for several other counting problems, including counting solutions for the multidimensional knapsack problem with a constant number of constraints, the general integer knapsack problem, and the contingency tables problem with a constant number of rows.

40.

246. Improved Mixing Condition on the Grid for Counting and Sampling Independent Sets

Ricardo Restrepo and Jinwoo Shin and Prasad Tetali and Eric Vigoda and Linji Yang

Abstract: The hard-core model has received much attention in the past couple of decades as a lattice gas model with hard constraints in statistical physics, a multicast model of calls in communication networks, and as a weighted independent set problem in combinatorics, probability and theoretical computer science.

In this model, each independent set I in a graph G is weighted proportionally to $\lambda^{|I|}$, for a positive real parameter λ . For large λ , computing the partition function (namely, the normalizing constant which makes the weighting a probability distribution on a finite graph) on graphs of maximum degree $\Delta \geq 3$, is a well known computationally challenging problem. More concretely, let $\lambda_c(\mathbb{T}_\Delta)$ denote the critical value for the so-called uniqueness threshold of the hard-core model on the infinite Δ -regular tree; recent breakthrough results of Dror Weitz (2006) and Allan Sly (2010) have identified $\lambda_c(\mathbb{T}_\Delta)$ as a threshold where the hardness of estimating the above partition function undergoes a computational transition.

We focus on the well-studied particular case of the square lattice \mathbb{Z}^2 , and provide a new lower bound for the uniqueness threshold, in particular taking it well above $\lambda_c(\mathbb{T}_4)$. Our technique refines and builds on the tree of self-avoiding walks approach of Weitz, resulting in a new technical sufficient criterion (of wider applicability) for establishing strong spatial mixing (and hence uniqueness) for the hard-core model. Applying our technique to \mathbb{Z}^2 we prove that strong spatial mixing holds for all $\lambda < 2.3882$, improving upon the work of Weitz that held for $\lambda < 27/16 = 1.6875$. Our results imply a fully-polynomial (deterministic) approximation algorithm for estimating the partition function, as well as rapid mixing of the associated Glauber dynamics to sample from the hard-core distribution. While we

focus here on the notoriously difficult hard-core model, our approach can also be applied to any 2-spin model, such as the Ising model.

41.

250. Balls and Bins: Smaller Hash Families and Faster Evaluation

L. Elisa Celis and Omer Reingold and Gil Segev and Udi Wieder

Abstract: A fundamental fact in the analysis of randomized algorithm is that when n balls are hashed into n bins independently and uniformly at random, with high probability each bin contains at most $O(\log n / \log \log n)$ balls. In various applications, however, the assumption that a truly random hash function is available is not always valid, and explicit functions are required.

In this paper we study the size of families (or, equivalently, the description length of their functions) that guarantee a maximal load of $O(\log n / \log \log n)$ with high probability, as well as the evaluation time of their functions. Whereas such functions must be described using $\Omega(\log n)$ bits, the best upper bound was formerly $O(\log^2 n / \log \log n)$ bits, which is attained by $O(\log n / \log \log n)$ -wise independent functions. Traditional constructions of the latter offer an evaluation time of $O(\log n / \log \log n)$, which according to Siegel's lower bound [FOCS '89] can be reduced only at the cost of significantly increasing the description length.

We construct two families that guarantee a maximal load of $O(\log n / \log \log n)$ with high probability. Our constructions are based on two different approaches, and exhibit different trade-offs between the description length and the evaluation time. The first construction shows that $O(\log n / \log \log n)$ -wise independence can in fact be replaced by "gradually increasing independence", resulting in functions that are described using $O(\log n \log \log n)$ bits and evaluated in time $O(\log n \log \log n)$. The second construction is based on derandomization techniques for space-bounded computations combined with a tailored construction of a pseudorandom generator, resulting in functions that are described using $O(\log^{3/2} n)$ bits and evaluated in time $O(\sqrt{\log n})$. The latter can be compared to Siegel's lower bound stating that $O(\log n / \log \log n)$ -wise independent functions that are evaluated in time $O(\sqrt{\log n})$ must be described using $\Omega(2^{\sqrt{\log n}})$ bits.

42.

251. Multiple-Source Multiple-Sink Maximum Flow in Directed Planar Graphs in Near-Linear Time

Glencora Borradaile and Philip N. Klein and Shay Mozes and Yahav Nussbaum and Christian Wulff-Nilsen

Abstract: We give an $O(n \log^3 n)$ algorithm that, given an n -node directed planar graph with arc capacities, a set of source nodes, and a set of sink nodes, finds a maximum flow from the sources to the sinks. Previously, the fastest algorithms known for this problem were those for general graphs.

43.

254. Quantum query complexity of state conversion

Troy Lee and Rajat Mittal and Ben W. Reichardt and Robert Spalek and Mario Szegedy

Abstract: State-conversion generalizes query complexity to the problem of converting between two input-dependent quantum states by making queries to the input. We characterize the complexity of this problem by introducing a natural information-theoretic norm that extends the Schur product operator norm. The complexity of converting between two systems of states is given by the distance between them, as measured by this norm.

In the special case of function evaluation, the norm is closely related to the general adversary bound, a semi-definite program that lower-bounds the number of input queries needed by a quantum algorithm to evaluate a function.

We thus obtain that the general adversary bound characterizes the quantum query complexity of any function whatsoever. This generalizes and simplifies the proof of the same result in the case of boolean input and output. Also in the case of function evaluation, we show that our norm satisfies a remarkable composition property, implying that the quantum query complexity of the composition of two functions is at most the product of the query complexities of the functions, up to a constant. Finally, our result implies that discrete and continuous-time query models are equivalent in the bounded-error setting, even for the general state-conversion problem.

44.

255. A constant factor approximation algorithm for unsplittable flow on paths

Paul Bonsma and Jens Schulz and Andreas Wiese

Abstract: In this paper, we present the first constant-factor approximation algorithm for the unsplittable flow problem on a path. This represents a large improvement over the previous best polynomial time approximation algorithm for this problem in its full generality, which was an $O(\log n)$ -approximation algorithm; it also answers an open question by Bansal et al. [SODA'09]. The approximation ratio of our algorithm is $7+\epsilon$ for any $\epsilon > 0$. In the unsplittable flow problem on a path, we are given a capacitated path P and n tasks, each task having a demand, a profit, and start and end vertices. The goal is to compute a maximum profit set of tasks such that the total demand of the selected tasks does not exceed the capacity of any edge on P . This is a well-studied problem that occurs naturally in various settings, and therefore it has been studied under alternative names, such as resource allocation, bandwidth allocation, resource constrained scheduling and temporal knapsack. Polynomial time constant factor approximation algorithms for the problem were previously known only under the no-bottleneck assumption (in which the maximum task demand must be no greater than the minimum edge capacity).

We introduce several novel algorithmic techniques, which might be of independent interest: a framework which reduces the problem to instances with a bounded range of capacities, and a new geometrically inspired dynamic program which solves a special case of the maximum weight independent set of rectangles problem to optimality. We also give the first proof that the problem is strongly NP-hard; we show that this is true even if all edge capacities are equal and all demands are either 1, 2, or 3.

45.

258. The Graph Minor Algorithm with Parity Conditions

Ken-ichi Kawarabayashi and Bruce Reed and Paul Wollan

Abstract: We generalize the seminal Graph Minor algorithm of Robertson and Seymour to the parity version. We give polynomial time algorithms for the following problems:

\begin{cases}

item

the parity H -minor (Odd K_k -minor) containment problem, and

\item

the disjoint paths problem with k terminals and the parity condition for each path,
 \end{enumerate}
 as well as several other related problems.

We present an $O(m \alpha(m,n) n)$ time algorithm for these problems for any fixed k , where n, m are the number of vertices and the number of edges, respectively, and the function $\alpha(m,n)$ is the inverse of the Ackermann function (see Tarjan \cite{tarjan}).

Note that the first problem includes the problem of testing whether or not a given graph contains k disjoint odd cycles (which was recently solved in \cite{tony,oddstoc}), if H consists of k disjoint triangles. The algorithm for the second problem generalizes the Robertson Seymour algorithm for the k -disjoint paths problem.

As with the Robertson-Seymour algorithm for the k -disjoint paths problem for any fixed k , in each iteration, we would like to either use the presence of a huge clique minor, or alternatively exploit the structure of graphs in which we cannot find such a minor. Here, however, we must take care of the parity of the paths and can only use an "odd clique minor". This requires new techniques to describe the structure of the graph when we cannot find such a minor.

We emphasize that our proof for the correctness of the above algorithms does not depend on the full power of the Graph Minor structure theorem \cite{RS16}. Although the original Graph Minor algorithm of Robertson and Seymour does depend on it and our proof does have similarities to their arguments, we can avoid the structure theorem by building on the shorter proof for the correctness of the graph minor algorithm in \cite{kw}. Consequently, we are able to avoid the much of the heavy machinery of the Graph Minor structure theory. Our proof is less than 50 pages.

46.

263. Mutual Exclusion with $O(\log^2 \log n)$ Amortized Work

Michael A. Bender and Seth Gilbert

Abstract: This paper gives a new algorithm for mutual exclusion in which each passage through the critical section costs amortized $O(\log^2 \log n)$ RMRs with high probability. The algorithm operates in a standard asynchronous, local spinning, shared-

memory model. The algorithm works against an oblivious adversary and guarantees that every process enters the critical section with high probability. The algorithm achieves its efficient performance by exploiting a connection between mutual exclusion and approximate counting. A central aspect of the work presented here is the development and application of efficient approximate-counting data structures.

Our mutual-exclusion algorithm represents a departure from previous algorithms in terms of techniques, adversary model, and performance. Most previous mutual exclusion algorithms are based on tournament-tree constructions. The most efficient prior algorithms require $O(\log n / \log \log n)$ RMRs and work against an adaptive adversary. In this paper, we focus on an oblivious model, and the algorithm in this paper is the first (for any adversary model) that can beat the $O(\log n / \log \log n)$ RMR bound.

47.

265. How Bad is Forming Your Own Opinion?

David Bindel and Jon Kleinberg and Sigal Oren

Abstract: A long-standing line of work in economic theory has studied models by which a group of people in a social network, each holding a numerical opinion, can arrive at a shared opinion through repeated averaging with their neighbors in the network. Motivated by the observation that consensus is rarely reached in real opinion dynamics, we study a related sociological model in which individuals' intrinsic beliefs counterbalance the averaging process and yield a diversity of opinions.

By interpreting the repeated averaging as best-response dynamics in an underlying game with natural payoffs, and the limit of the process as an equilibrium, we are able to study the cost of disagreement in these models relative to a social optimum. We provide a tight bound on the cost at equilibrium relative to the optimum; our analysis draws a connection between these agreement models and extremal problems for generalized eigenvalues. We also consider a natural network design problem in this setting, where adding links to the underlying network can reduce the cost of disagreement at equilibrium.

48.

268. The Complexity of Renaming

Dan Alistarh and James Aspnes and Seth Gilbert and Rachid Guerraoui

Abstract: We study the complexity of renaming, a fundamental problem in distributed computing in which a set of processes need to pick distinct names from a given namespace. We prove a local lower bound of $\Omega(k)$ process steps for deterministic renaming into any namespace of size sub-exponential in k , where k is the number of participants. This bound is tight: it draws an exponential separation between deterministic and randomized solutions, and implies tight bounds for deterministic fetch-and-increment registers, queues and stacks. The proof of the bound is interesting in its own right, for it relies on the first reduction from renaming to another fundamental problem in distributed computing: mutual exclusion. We complement our local bound with a global lower bound of $\Omega(k \log(k/c))$ on the total step complexity of renaming into a namespace of size ck , for any $c \geq 1$. This applies to randomized algorithms against a strong adversary, and helps derive new global lower bounds for randomized approximate counter and fetch-and-increment implementations, all tight within logarithmic factors.

49.

271. On the Power of Adaptivity in Sparse Recovery

Piotr Indyk and Eric Price and David Woodruff

Abstract: The goal of (stable) sparse recovery is to recover a k -sparse approximation x^* of a vector x from linear measurements of x . Specifically, the goal is to recover x^* such that

$$\| \|x - x^*\|_p \leq C \min_{\text{support } x'} \|x - x'\|_q$$

for some constant C and norm parameters p and q . It is known that, for $p=q=1$ or $p=q=2$, this task can be accomplished using $m=O(k \log(n/k))$ non-adaptive measurements [CRT06:Stable-Signal] and that this bound is tight [DIPW,FPRU].

In this paper we show that if one is allowed to perform measurements that are adaptive, then the number of measurements can be considerably reduced. Specifically, for $C=1+\epsilon$ and $p=q=2$ we show:

* A scheme with $m = O\left(\frac{1}{\epsilon} k \log \log (n\epsilon/k)\right)$ measurements that uses $O\left(\sqrt{\log k} \cdot \log \log (n\epsilon/k)\right)$ rounds. This is a significant improvement over the *best possible* non-adaptive bound.

* A scheme with $m = O\left(\frac{1}{\epsilon} k \log (k/\epsilon) + k \log (n/k)\right)$ measurements that uses *two* rounds. This improves over the *best known* non-adaptive bound.

To the best of our knowledge, these are the first results of this type.

50.

274. Enumerative Lattice Algorithms in any Norm via M-ellipsoid Coverings

Daniel Dadush and Chris Peikert and Santosh Vempala

Abstract: We give a novel algorithm for enumerating lattice points in any convex body, and give applications to several classic lattice problems, including the Shortest and Closest Vector Problems (SVP and CVP, respectively) and Integer Programming (IP). Our enumeration technique relies on a classical concept from asymptotic convex geometry known as the *M-ellipsoid*, and uses as a crucial subroutine the recent algorithm of Micciancio and Voulgaris (STOC 2010) for lattice problems in the ℓ_2 norm. As a main technical contribution, which may be of independent interest, we build on the techniques of Klartag (Geometric and Functional Analysis, 2006) to give an expected $2^{O(n)}$ -time algorithm for computing an M-ellipsoid for any n -dimensional convex body.

As applications, we give deterministic $2^{O(n)}$ -time and -space algorithms for solving exact SVP, and exact CVP when the target point is sufficiently close to the lattice, on n -dimensional lattices *in any (semi-)norm* given an M-ellipsoid of the unit ball. In many norms of interest, including all ℓ_p norms, an M-ellipsoid is computable in deterministic $\text{poly}(n)$ time, in which case these algorithms are fully deterministic. Here our approach may be seen as a derandomization of the "AKS sieve" for exact SVP and CVP (Ajtai, Kumar, and Sivakumar; STOC 2001 and CCC 2002).

As a further application of our SVP algorithm, we derive an expected $O(f^*(n))^n$ -time algorithm for Integer Programming, where $f^*(n)$ denotes the optimal bound in the so-called "flatness theorem," which satisfies $f^*(n) = O(n^{4/3} \text{polylog}(n))$ and is conjectured to be $f^*(n) = \Theta(n)$. Our runtime improves upon the previous best of $O(n^2)^n$ by Hildebrand and Kopp (2010).

51.

276. Efficient and Explicit Coding for Interactive Communication

Ran Gelles and Ankur Moitra and Amit Sahai

Abstract: In this work, we study the fundamental problem of reliable interactive communication over a noisy channel. In a breakthrough sequence of papers published in 1992 and 1993, Schulman gave non-constructive proofs of the existence of general methods to emulate any two-party interactive protocol such that: (1) the emulation protocol only takes a constant-factor longer than the original protocol, and (2) if the emulation protocol is executed over any discrete memoryless noisy channel with constant capacity, then the probability that the emulation protocol fails to perfectly emulate the original protocol is exponentially small in the total length of the protocol. Unfortunately, Schulman's emulation procedures either only work in a nonstandard model with a large amount of shared randomness, or are non-constructive in that they rely on the existence of "absolute" tree codes. The only known proofs of the existence of absolute tree codes are non-constructive, and finding an explicit construction remains an important open problem. Indeed, randomly generated tree codes are not absolute tree codes with overwhelming probability. In this work, we revisit the problem of reliable interactive communication, and obtain the first fully explicit (randomized) efficient constant-rate emulation procedure for reliable interactive communication. Our protocol works for any discrete memoryless noisy channel with constant capacity, and our protocol's probability of failure is exponentially small in the total length of the protocol. We accomplish this goal by obtaining the following results: We introduce a new notion of goodness for a tree code, and define the notion of a potent tree code. We believe that this notion is of independent interest. We prove the correctness of an explicit emulation procedure based on any potent tree code. (This replaces the need for absolute tree codes in the work of Schulman.) We show that a randomly generated tree code (with suitable constant alphabet size) is an efficiently decodable potent tree code with overwhelming probability. Furthermore we are able to partially derandomize this result by means of epsilon-biased distributions using only $O(n)$ random bits, where n is the depth of the tree.

These (derandomized) results allow us to obtain our main result. Our results also extend to the case of interactive multi-party communication among a constant number of parties.

52.

286. Dispersers for affine sources with sub-polynomial entropy

Ronen Shaltiel

Abstract: We construct an explicit disperser for affine sources over \mathbb{F}_2^n with entropy $k=2^{\log^{0.9} n}=n^{o(1)}$. This is a polynomial time computable function $\text{Disp}:\mathbb{F}_2^n \rightarrow \mathbb{B}$ such that for every affine space V of \mathbb{F}_2^n that has dimension at least k , $\text{Disp}(V)=\{0,1\}$. This improves the best previous construction of [BK](#) that achieved $k = \Omega(n^{4/5})$.

Our technique follows a high level approach that was developed in [BKSSW,BRSW](#) in the context of dispersers for two independent general sources. The main steps are:

- Adjust the high level approach to make it suitable for affine sources.
- Implement a "challenge-response game" for affine sources (in the spirit of [BKSSW,BRSW](#) that introduced such games for two independent general sources).
- In order to implement the game, we construct extractors for affine block-wise sources. For this we use ideas and components from [Rao09](#).
- Combining the three items above, we obtain dispersers for affine sources with entropy that larger than \sqrt{n} , and we use a recursive win-win analysis in the spirit of [RSW](#) to get affine dispersers with entropy less than \sqrt{n} .

53.

290. Approximation Algorithms for Correlated Knapsacks and Non-Martingale Bandits

Anupam Gupta and Ravishankar Krishnaswamy and Marco Molinaro and R. Ravi

Abstract: In the stochastic knapsack problem, we are given a knapsack of size B , and a set of items whose sizes and rewards are drawn from a known probability distribution. However, the only way to know the actual size and reward is to schedule the item—when it completes, we get to know these values. The goal is to schedule these

items (possibly making adaptive decisions based on the sizes seen thus far) to maximize the expected total reward of items which successfully pack into the knapsack. We know constant-factor approximations when (i) the rewards and sizes are independent of each other, and (ii) we cannot prematurely cancel items after we schedule them. What can we say when either or both of these assumptions are relaxed?

Related to this are other stochastic packing problems like the multi-armed bandit (and budgeted learning) problems; here one is given several arms which evolve in a specified stochastic fashion with each pull, and the goal is to (adaptively) decide which arms to pull, in order to maximize the expected reward obtained after B pulls in total. Much recent work on this problem focus on the case when the evolution of each arm follows a martingale, i.e., when the expected reward from one pull of an arm is the same as the reward at the current state. What can we say when the rewards do not form a martingale? In this paper, we give constant-factor approximation algorithms for the stochastic knapsack problem with correlations and/or cancellations. Extending ideas we develop for this problem, we also give constant-factor approximations for MAB problems without the martingale assumption. Indeed, we can show that previously proposed linear programming relaxations for these problems have large integrality gaps. So we propose new time-indexed LP relaxations; using a decomposition and “gap-filling” approach, we convert these fractional solutions to distributions over strategies, and then use the LP values and the time ordering information from these strategies to devise randomized adaptive scheduling algorithms. We hope our LP formulation and decomposition methods may provide a new way to address other stochastic optimization problems with more general contexts.

54.

293. Lasserre Hierarchy, Higher Eigenvalues, and Approximation Schemes for Graph Partitioning and Quadratic Integer Programming with PSD Objectives

Venkatesan Guruswami and Ali Kemal Sinop

Abstract: We present an approximation scheme for optimizing certain Quadratic Integer Programming problems with positive semidefinite objective functions and global linear constraints. This framework includes well known graph problems such as Uniform sparsest cut, Minimum graph bisection, and Small Set expansion, as well as the Unique Games problem. These problems are notorious for the existence of huge gaps between the known algorithmic results and NP-hardness results.

Our algorithm is based on rounding semidefinite programs from the Lasserre hierarchy, and the analysis uses bounds for low-rank approximations of a matrix in Frobenius norm using columns of the matrix.

For all the above graph problems, we give an algorithm running in time $n^{O(r/\epsilon^2)}$ with approximation ratio $\frac{1+\epsilon}{\min\{1, \lambda_r\}}$, where λ_r is the r 'th smallest eigenvalue of the normalized graph Laplacian L . In the case of graph bisection and small set expansion, the number of vertices in the cut is within lower-order terms of the stipulated bound. Our results imply $(1+O(\epsilon))$ factor approximation in time $n^{O(r^{\text{ast}})}$ where r^{ast} is the number of eigenvalues of L smaller $1-\epsilon$. This perhaps gives some indication as to why even showing mere APX-hardness for these problems has been elusive, since the reduction must produce graphs with a slowly growing spectrum (and classes like planar graphs which are known to have such a spectral property often admit good algorithms owing to their nice structure).

For Unique Games, we give a factor $(1+\frac{2+\epsilon}{\lambda_r})$ approximation for minimizing the number of unsatisfied constraints in $n^{O(r/\epsilon)}$ time. This improves an earlier bound for solving Unique Games on expanders, and also shows that Lasserre SDPs are powerful enough to solve well-known integrality gap instances for the basic SDP. We also give an algorithm for independent sets in graphs that performs well when the Laplacian does not have too many eigenvalues bigger than $1+o(1)$.

55.

296. A Unified Continuous Greedy Algorithm for Submodular Maximization

Moran Feldman and Joseph (Seffi) Naor and Roy Schwartz

Abstract: The study of combinatorial problems with a submodular objective function has attracted much attention in recent years, and is partly motivated by the importance of such problems to economics, algorithmic game theory and combinatorial optimization.

Classical works on these problems are mostly combinatorial in nature. Recently, however, many results based on continuous algorithmic tools have emerged. The main bottleneck of such continuous techniques is how to approximately solve a non-convex relaxation for the submodular problem at hand.

Thus, the efficient computation of better fractional solutions immediately implies improved approximations for numerous applications.

A simple and elegant method, called "continuous greedy", successfully tackles this issue for monotone submodular objective functions, however, only much more complex tools are known to work for general non-monotone submodular objectives.

In this work we present a new unified continuous greedy algorithm which finds approximate fractional solutions for both the non-monotone and monotone cases, and improves on the approximation ratio for many applications.

For general non-monotone submodular objective functions, our algorithm achieves an improved approximation ratio of about $1/e$.

For monotone submodular objective functions, our algorithm achieves an approximation ratio that depends on the density of the polytope defined by the problem at hand, which is always at least as good as the previously known best approximation ratio of $1 - 1/e$.

Some notable immediate implications are an improved $1/e$ -approximation for maximizing a non-monotone submodular function subject to a matroid or $O(1)$ -knapsack constraints, and information-theoretic tight approximations for Submodular Max-SAT and Submodular Welfare with k players, for any number of players k .

A framework for submodular optimization problems, called the contention resolution framework, was introduced recently by Chekuri et al. The improved approximation ratio of the unified continuous greedy algorithm implies improved approximation ratios for many problems through this framework. Moreover, via a parameter called stopping time, our algorithm merges the relaxation solving and re-normalization steps of the framework, and achieves, for some applications, further improvements. We also describe new monotone balanced contention resolution schemes for various matching, scheduling and packing problems, thus, improving the approximations achieved for these problems via the framework.

56.

297. Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers

Saeed Alaei

Abstract: For Bayesian combinatorial auctions, we present a general framework for reducing the mechanism design problem for many buyers to the mechanism design problem for one buyer. Our generic reduction works for any separable objective (e.g., welfare, revenue, etc) and any space of valuations (e.g. submodular, additive, etc) and any distribution of valuations as long as valuations of different buyers are distributed independently (not necessarily identically). Roughly speaking, we present two generic n -buyer mechanisms that use 1 -buyer mechanisms as black boxes. We show that if we have an α -approximate 1 -buyer mechanism for each buyer^{footnote{Note that we can use different 1 -buyer mechanisms for different buyers.}} then our generic n -buyer mechanisms are $\frac{1}{2}\alpha$ -approximation of the optimal n -buyer mechanism. Furthermore, if we have several copies of each item and no buyer ever needs more than

$\frac{1}{k}$ of all copies of each item then our generic n -buyer mechanisms are $\frac{1}{\alpha}$ -approximation of the optimal n -buyer mechanism where $\frac{1}{\alpha} \geq 1 - \frac{1}{\sqrt{k+3}}$. Observe that $\frac{1}{\alpha}$ is at least $\frac{1}{2}$ and approaches 1 as k increases.

Applications of our main theorem include the following improvements on results from the literature. For each of the following models we construct a $\frac{1}{\alpha}$ -buyer mechanism and then apply our generic expansion: For revenue maximization in combinatorial auctions with hard budget constraints, \cite{BGM10} presented a $\frac{1}{4}$ -approximate BIC mechanism for additive/correlated valuations and an $O(1)$ -approximate \footnote{\math{O}(1)=\frac{1}{96}} sequential posted pricing mechanism for additive/independent valuations. We improve this to a $\frac{1}{\alpha}$ -approximate BIC mechanism and a $\frac{1}{\alpha} (1 - \frac{1}{e})$ -approximate sequential posted pricing mechanism respectively. For revenue maximization in combinatorial auctions with unit demand buyers, \cite{CHMS10} presented a $\frac{1}{6.75}$ -approximate sequential posted pricing mechanism. We improve this to a $\frac{1}{2} \frac{1}{\alpha}$ approximate sequential posted pricing mechanism. We also present a $\frac{1}{\alpha}$ -approximate sequential posted pricing mechanism for unit-demand multi-unit auctions (homogeneous) with hard-budget constraints. Furthermore, our sequential posted pricing mechanisms assume no control or prior information about the order in which buyers arrive.

57.

301. Extreme-Value Theorems for Optimal Multidimensional Pricing

Yang Cai and Constantinos Daskalakis

Abstract: We provide a Polynomial Time Approximation Scheme for the (multi-dimensional unit-demand pricing problem), when the buyer's values are independent (but not necessarily identically distributed.) For all $\epsilon > 0$, we obtain a $(1+\epsilon)$ -factor approximation to the optimal revenue in time polynomial, when the values are sampled from Monotone Hazard Rate (MHR) distributions, quasi-polynomial, when sampled from regular distributions, and polynomial in $n^{\text{poly}(\log r)}$, when sampled from general distributions supported on a set $[u_{\min}, r u_{\min}]$. We also provide an additive PTAS for all bounded distributions.

Our algorithms are based on novel extreme value theorems for MHR and regular distributions, and apply probabilistic techniques to understand the statistical properties of revenue distributions, as well as to reduce the size of the search space of the algorithm. As a byproduct of our techniques, we establish structural properties of optimal solutions. We show that, for all $\epsilon > 0$, $g(1/\epsilon)$ distinct prices suffice to obtain a $(1+\epsilon)$ -factor approximation to the optimal revenue for MHR distributions, where $g(1/\epsilon)$ is a quasi-linear function of $1/\epsilon$ that does not depend on the number of items. Similarly, for all $\epsilon > 0$ and $n > 0$, $g(1/\epsilon \cdot \log n)$ distinct prices suffice for regular distributions, where n is the number of items and $g(\cdot)$ is a polynomial function. Finally, in the i.i.d. MHR case, we show that, as long as the number of items is a sufficiently large function of $1/\epsilon$, a single price suffices to achieve a $(1+\epsilon)$ -factor approximation.

58.

304. Approximation Algorithms for Submodular Multiway Partition

Chandra Chekuri and Alina Ene

Abstract: We study algorithms for the `{\sc Submodular Multiway Partition}` problem (`\SubMP`). An instance of `\SubMP` consists of a finite ground set V , a subset of k elements $S = \{s_1, s_2, \dots, s_k\}$ called terminals, and a non-negative submodular set function $f: 2^V \rightarrow \mathbb{R}_+$ on V provided as a value oracle. The goal is to partition V into k sets A_1, \dots, A_k such that for $1 \leq i \leq k$, $s_i \in A_i$ and $\sum_{i=1}^k f(A_i)$ is minimized. `\SubMP` generalizes some well-known problems such as the `{\sc Multiway Cut}` problem in graphs and hypergraphs, and the `{\sc Node-weighted Multiway Cut}` problem in graphs. `\SubMP` for arbitrary submodular functions (instead of just symmetric functions) was considered by Zhao, Nagamochi and Ibaraki `\cite{ZhaoNI05}`. Previous algorithms were based on greedy splitting and divide and conquer strategies. In very recent work `\cite{ChekuriE11}` we proposed a convex-programming relaxation for `\SubMP` based on the Lov'asz-extension of a submodular function and showed its applicability for some special cases. In this paper we obtain the following results for arbitrary submodular functions via this relaxation.

`\begin{itemize}`

`\item` A 2 -approximation for `\SubMP`. This improves the $(k-1)$ -approximation from `\cite{ZhaoNI05}`.

`\item` A $(1.5 - 1/k)$ -approximation for `\SubMP` when f is `{\em symmetric}`. This improves the $2(1 - 1/k)$ -approximation from `\cite{Queyranne99,ZhaoNI05}`.

`\end{itemize}`

59.

318. Delays and the Capacity of Continuous-time Channels

Sanjeev Khanna and Madhu Sudan

Abstract: Any physical channel of communication offers two potential reasons why its capacity (the number of bits it can transmit in a unit of time) might be unbounded: (1) (Uncountably) infinitely many choices of signal strength at any given time, and (2) (Uncountably) infinitely many instances of time at which signals may be sent. However channel noise cancels out the potential unboundedness of the first aspect, leaving typical channels with only a finite capacity per instant of time. The latter source of infinity seems less extensively studied. A potential source of unreliability that might restrict the capacity also from the second aspect is "delay": Signals transmitted by the sender at a given point of time may not be received with a predictable delay at the receiving end. In this work we examine this source of uncertainty by considering a simple discrete model of delay errors. In our model the communicating parties get to subdivide time as finely as they wish, but still have to cope with communication delays that are variable. The continuous process becomes the limit of our process as the time subdivision becomes infinitesimal. We analyze the limits of such channels and reach somewhat surprising conclusions: The capacity of a physical channel is finitely bounded only if at least one of the two sources of error (signal noise or delay noise) is adversarial. If both error sources are stochastic, or the adversarial source is noise that is independent of the stochastic delay, the capacity of the associated physical channel is infinite!

60.

320. Fully Homomorphic Encryption without Squashing Using Depth-3 Arithmetic Circuits

Craig Gentry and Shai Halevi

Abstract: All currently known fully homomorphic encryption (FHE) schemes use the same blueprint from [Gentry 2009]: First construct a somewhat homomorphic encryption (SWHE) scheme, next "squash" the decryption circuit until it is simple enough to be handled within the homomorphic capacity of the SWHE scheme, and finally "bootstrap" to get a FHE scheme. In all existing schemes, the squashing technique induces an additional assumption: that the sparse subset sum problem (SSSP) is hard.

We describe a *new approach* that constructs FHE as a hybrid of a SWHE scheme and a multiplicatively homomorphic encryption (MHE) scheme, such as Elgamal. Our construction eliminates the need for the squashing step, and thereby also removes the need to assume the SSSP is hard. We describe a few concrete instantiations of the new method, obtaining the following results:

1. A "simple" FHE scheme where we replace SSSP with Decision Diffie-Hellman.
2. The first FHE scheme based entirely on worst-case hardness: Specifically, we describe a "leveled" FHE scheme whose security can be quantumly reduced to the approximate shortest independent vector problem over ideal lattices (ideal-SIVP).
3. Some efficiency improvements for FHE: While at present our new method does not improve computational efficiency, we do provide an optimization that reduces the ciphertext length. For example, at one point, the entire FHE ciphertext may consist of a single Elgamal ciphertext!

Our new method does not eliminate the bootstrapping step. Whether this can be done remains an intriguing open problem. As in the previous blueprint, we can get "pure" (non-leveled) FHE by assuming circular security.

Our main technique is to express the decryption function of SWHE schemes as a depth-3 arithmetic circuit of a particular form. When evaluating this circuit homomorphically, as needed for bootstrapping, we temporarily switch to a MHE scheme, such as Elgamal, to handle the product part of the circuit. We then translate the result back to the SWHE scheme by homomorphically evaluating the decryption function of the MHE scheme. (Due to the special form of the circuit, switching to the MHE scheme can be done without having to evaluate anything homomorphically.) Using our method, the SWHE scheme only needs to be capable of evaluating the MHE scheme's decryption function, not its own decryption function. We thereby avoid the circularity that necessitated squashing in the original blueprint.

61.

323. A Randomized Rounding Approach to the Traveling Salesman Problem

Shayan Oveis Gharan and Amin Saberi and Mohit Singh

Abstract: For some positive constant ϵ_0 , we give a $\frac{3}{2} - \epsilon_0$ -approximation algorithm for the following problem: given a graph $G_0 = (V, E_0)$, find the shortest tour that visits every vertex at least once. This is a special case of the metric traveling salesman problem when the underlying metric is defined by shortest path distances in G_0 . The result improves on the $\frac{3}{2}$ -approximation algorithm due to Christofides for this special case.

Similar to Christofides, our algorithm finds a spanning tree whose cost is upper bounded by the optimum, then it finds the minimum cost Eulerian augmentation (or T-join) of that tree. The main difference is in the selection of the spanning tree. Except in certain cases where the solution of LP is nearly integral, we select the spanning tree randomly by sampling from a maximum entropy distribution defined by the linear programming relaxation.

Despite the simplicity of the algorithm, the analysis builds on a variety of ideas such as properties of strongly Rayleigh measures from probability theory, graph theoretical results on the structure of near minimum cuts, and the integrality of the T-join polytope from polyhedral theory. Also, as a byproduct of our result, we show new properties of the near minimum cuts of any graph, which may be of independent interest.

62.

329. Algorithms for the Generalized Sorting Problem

Zhiyi Huang and Sampath Kannan and Sanjeev Khanna

Abstract: We study the generalized sorting problem where we are given a set of n elements to be sorted but only a subset of all possible pairwise element comparisons is allowed. The goal is to determine the sorted order using the smallest possible number of allowed comparisons. The generalized sorting problem may be equivalently viewed as follows. Given an undirected graph $G(V,E)$ where V is the set of elements to be sorted and E defines the set of allowed comparisons, adaptively find the smallest subset $E' \subseteq E$ of edges to probe such that the directed graph induced by E' contains a Hamiltonian path.

When G is a complete graph, we get the standard sorting problem, and it is well-known that $\Theta(n \log n)$ comparisons are necessary and sufficient. An extensively studied special case of the generalized sorting problem is the nuts and bolts problem where the allowed comparison graph is a complete bipartite graph between two equal-size sets. It is known that for this special case also, there is a deterministic algorithm that sorts using $\Theta(n \log n)$ comparisons. However, when the allowed comparison graph is arbitrary, to our knowledge, no bound better than the trivial $O(n^2)$ bound is known. Our main result is a randomized algorithm that sorts any allowed comparison graph using $O(n^{3/2})$ comparisons with high probability (provided the input is sortable). We also study the sorting problem in randomly generated allowed comparison graphs, and show that when the edge probability is p , $O(\min\{\frac{n}{p^2}, n^{3/2}\sqrt{p}\})$ comparisons suffice on average to sort.

63.

334. Privacy Amplification and Non-Malleable Extractors Via Character Sums

Xin Li and Trevor D. Wooley and David Zuckerman

Abstract: In studying how to communicate over a public channel with an active adversary, Dodis and Wichs introduced the notion of a non-malleable extractor. A non-malleable extractor dramatically strengthens the notion of a strong extractor. A strong extractor takes two inputs, a weakly-random x and a uniformly random seed y , and outputs a string which appears uniform, even given y . For a non-malleable extractor nmExt , the output $\text{nmExt}(x,y)$ should appear uniform given y as well as $\text{nmExt}(x,A(y))$, where A is an arbitrary function with $A(y) \neq y$.

We show that an extractor introduced by Chor and Goldreich is non-malleable when the entropy rate is above half.

It outputs a linear number of bits when the entropy rate is $1/2 + \alpha$, for any $\alpha > 0$.

Previously, no nontrivial parameters were known for any non-malleable extractor.

To achieve a polynomial running time when outputting many bits, we rely on a widely-believed conjecture about the distribution of prime numbers in arithmetic progressions.

Our analysis involves a character sum estimate, which may be of independent interest.

Using our non-malleable extractor, we obtain protocols for "privacy amplification": key agreement between two parties who share a weakly-random secret. Our protocols work in the presence of an active adversary with unlimited computational power, and have optimal entropy loss. When the secret has entropy rate greater than $1/2$, the protocol follows from a result of Dodis and Wichs, and takes two rounds. When the secret has entropy rate δ for any constant $\delta > 0$, our new protocol takes $O(1)$ rounds. Our protocols run in polynomial time under the above well-known conjecture about primes.

64.

338. A nearly $m \log n$ time solver for SDD linear systems

Ioannis Koutis and Gary L. Miller and Richard Peng

Abstract: We present an improved algorithm for solving symmetrically diagonally dominant linear systems. On input of an $n \times n$ symmetric diagonally dominant matrix A with m non-zero entries and a vector b such that $A\bar{x} = b$ for some (unknown) vector \bar{x} , our algorithm computes a vector x such that $\|x - \bar{x}\|_A < \epsilon \|\bar{x}\|_A$ $\|\cdot\|_A$ denotes the A-norm in time $\tilde{O}(m \log n \log(1/\epsilon))$.

The solver utilizes in a standard way a 'preconditioning' chain of progressively sparser graphs. To claim the faster running time we make a two-fold improvement in the algorithm for constructing the chain. The new chain exploits previously unknown properties of the graph sparsification algorithm given in [Koutis, Miller, Peng, FOCS 2010], allowing for stronger preconditioning properties. We also present an algorithm of independent interest that constructs nearly-tight low-stretch spanning trees in time $\tilde{O}(m \log n)$, a factor of $O(\log n)$ faster than the algorithm in [Abraham, Bartal, Neiman, FOCS 2008]. This speedup directly reflects on the construction time of the preconditioning chain.

65.

341. Which Networks Are Least Susceptible to Cascading Failures?

Lawrence Blume and David Easley and Jon Kleinberg and Robert Kleinberg and Eva Tardos

Abstract: The resilience of networks to various types of failures is an undercurrent in many parts of graph theory and network algorithms. In this paper we study the resilience of networks in the presence of *cascading failures* --- failures that spread from one node to another across the network structure. One finds such cascading processes at work in the kind of contagious failures that spread among financial institutions during a financial crisis, through nodes of a power grid or communication network during a widespread outage, or through a human population during the outbreak of an epidemic disease.

A widely studied model of cascades in networks assumes that each node v of the network has a threshold $\ell(v)$, and fails if it has at least $\ell(v)$ failed neighbors. We

assume that each node selects a threshold $\ell(v)$ independently using a probability distribution μ . Our work centers on a parameter that we call the μ -risk of a graph: the maximum failure probability of any node in the graph, in this threshold cascade model parameterized by threshold distribution μ . This defines a very broad class of models; for example, the large literature on edge percolation, in which propagation happens along edges that are included independently at random with some probability p , takes place in a small part of the parameter space of threshold cascade models, and one where the distribution μ is monotonically decreasing with the threshold. In contrast we want to study the whole space, including threshold distributions with qualitatively different behavior, such as those that are sharply increasing.

We develop techniques for relating differences in μ -risk to the structures of the underlying graphs. This is challenging in large part because, despite the simplicity of its formulation, the threshold cascade model has been very hard to analyze for arbitrary graphs G and arbitrary threshold distributions μ . It turns out that when selecting among a set of graphs to minimize the μ -risk, the result depends quite intricately on μ . We develop several techniques for evaluating the μ -risk of d -regular graphs. For $d=2$ we are able to solve the problem completely: the optimal graph is always a clique (i.e. triangle) or tree (i.e. infinite path), although which graph is better exhibits a surprising non-monotonicity as the threshold parameters vary. When $d>2$ we present a technique based on power-series expansions of the failure probability that allows us to compare graphs in certain parts of the parameter space, deriving conclusions including the fact that as μ varies, at least three different graphs are optimal among d -regular graphs. In particular, the set of optimal 3-regular graphs includes one which is neither a clique nor a tree.

66.

343. Online Node-weighted Steiner Tree and Related Problems

Joseph (Seffi) Naor and Debmalya Panigrahi and Mohit Singh

Abstract: We obtain the first online algorithms for the node-weighted Steiner tree, Steiner forest and group Steiner tree problems that achieve a poly-logarithmic competitive ratio. Our algorithm for the Steiner tree problem runs in polynomial time, while those for the other two problems take quasi-polynomial time. Our algorithms can be viewed as online LP rounding algorithms in the framework of Buchbinder and Naor; however, while the *natural* LP formulation of these problems do lead to fractional algorithms with a poly-logarithmic competitive ratio, we are unable to round these LPs online without losing a polynomial factor. Therefore, we design new LP formulations for these problems drawing on a combination of paradigms such as *spider decompositions*, *low-depth Steiner trees*, *generalized group Steiner*

problems}, etc. and use the additional structure provided by these to round the more sophisticated LPs losing only a poly-logarithmic factor in the competitive ratio. As further applications of our techniques, we also design polynomial-time online algorithms with polylogarithmic competitive ratios for two fundamental network design problems in edge-weighted graphs: the group Steiner forest problem (thereby resolving an open question raised by Chekuri {\em et al}) and the single source ℓ -vertex connectivity problem (which complements similar results for the corresponding edge-connectivity problem due to Gupta {\em et al}).

67.

344. Welfare and Profit Maximization with Production Costs

Avrim Blum and Anupam Gupta and Yishay Mansour and Ankit Sharma

Abstract: Combinatorial Auctions are a central problem in Algorithmic Mechanism Design: pricing and allocating goods to buyers with complex preferences in order to maximize some desired objective (e.g., social welfare, revenue, or profit). The problem has been well-studied in the case of limited supply (one copy of each item), and in the case of digital goods (the seller can produce additional copies at no cost). Yet the case of resources---think oil, labor, computing cycles, etc.---neither of these abstractions is just right: additional supplies of these resources can be found, but only at a cost (where the marginal cost is an increasing function).

In this work, we initiate the study of the algorithmic mechanism design problem of combinatorial pricing under increasing marginal cost. The goal is to sell these goods to buyers with unknown and arbitrary combinatorial valuation functions to maximize either the social welfare, or our own profit; specifically we focus on the setting of *posted item prices* with buyers arriving online. We give algorithms that achieve constant factor approximations for a class of natural cost functions---linear, low-degree polynomial, logarithmic---and that give logarithmic approximations for all convex marginal cost functions (along with a necessary additive loss). We show that these bounds are essentially best possible for these settings.

68.

345. On the complexity of Commuting Local Hamiltonians, and tight conditions for Topological Order in such systems

Dorit Aharonov and Lior Eldar

Abstract: The local Hamiltonian problem plays the equivalent role of SAT in quantum

complexity theory. Understanding the complexity of the intermediate case in which the constraints are quantum but all local terms in the Hamiltonian commute, is of importance for conceptual, physical and computational complexity reasons.

Bravyi and Vyalyi showed in 2003 \cite{BV}, using a clever application of the representation theory of C^* -algebras, that if the terms in the Hamiltonian are all two-local, the problem is in NP, and the entanglement in the ground states is local. The general case remained open since then.

In this paper we extend this result beyond the two-local case, to the case of three-qubit interactions.

We then extend our results even further, and show that NP verification is possible for three-wise interaction between qutrits as well, as long as the interaction graph is planar and also "nearly Euclidean" in some well-defined sense.

The proofs imply that in all such systems, the entanglement in the ground states is local.

These extensions imply an intriguing sharp transition phenomenon in commuting Hamiltonian systems: the ground spaces of 3-local "physical"

systems based on qubits and qutrits are diagonalizable by a basis whose entanglement is highly local, while more involved interactions (the particle dimensionality or the locality of the interaction is larger) can already exhibit topological order;

In particular, for those parameters, there exist Hamiltonians all of whose groundstates have entanglement which spreads over scales proportional to the size of the system, such as Kitaev's Toric Code system.

This has a direct implication to the developing theory of Topological Order, since it shows that one cannot improve on the parameters to construct topological order systems based on commuting Hamiltonians.

This is of particular interest in light of the recent proofs by Bravyi, Hastings and Michalakis

that Topological Order generated by commuting systems exhibits robustness against local perturbations of the Hamiltonian, implying the fault-tolerance of such systems;

Our results imply that one cannot hope to improve in parameters over Kitaev's seminal construction, as it is optimal in terms of parameters which allow construction of TO using commuting systems.

69.

346. Efficient Fully Homomorphic Encryption from (Standard) LWE

Zvika Brakerski and Vinod Vaikuntanathan

Abstract: We present a fully homomorphic encryption scheme that is based solely on the (standard) learning with errors (LWE) assumption. Applying known results on LWE, the security of our scheme is based on the worst-case hardness of short vector problems on arbitrary lattices. As icing on the cake, our scheme is quite efficient, and has very short ciphertexts.

Our construction improves upon previous works in two aspects:

1. We show that "somewhat homomorphic" encryption can be based on LWE, using a new *re-linearization* technique. In contrast, all previous schemes relied on complexity assumptions related to ideals in various rings.
2. More importantly, we deviate from the "squashing paradigm" used in all previous works. We introduce a new *dimension reduction* technique, which shortens the ciphertexts and reduces the decryption complexity of our scheme, without introducing additional assumptions. In contrast, all previous works required an additional, very strong assumption (namely, the sparse subset sum assumption).

Since our scheme has very short ciphertexts, we use it to construct an asymptotically-efficient LWE-based single-server private information retrieval (PIR) protocol. The communication complexity of our protocol (in the public-key model) is $k \cdot \text{polylog}(k, \log |DB|)$ bits per single-bit query, which is better than any known scheme. Previously, it was not known how to achieve a communication complexity of even $\text{poly}(k, \log |DB|)$ based on LWE.

70.

347. Testing and Reconstruction of Lipschitz Functions with Applications to Data Privacy

Madhav Jha and Sofya Raskhodnikova

Abstract: A function $f : D \rightarrow \mathbb{R}$ has Lipschitz constant c if $d_{\mathbb{R}}(f(x), f(y)) \leq c d_D(x, y)$ for all x, y in D , where $d_{\mathbb{R}}$ and d_D denote the distance functions on the range and domain of f , respectively. We say a function is Lipschitz if it has Lipschitz constant 1. (Note that rescaling by a factor of $1/c$ converts a function with a Lipschitz constant c into a Lipschitz function.) In other words, Lipschitz functions are not very sensitive to small changes in the input.

We initiate the study of testing and local reconstruction of the Lipschitz property of functions. A property tester has to distinguish functions with the property (in this case, Lipschitz) from functions that are ϵ -far from having the property, that is, differ from every function with the property on at least an ϵ fraction of the domain. A local filter reconstructs an arbitrary function f to ensure that the reconstructed function g has the desired property (in this case, is Lipschitz), changing f only when necessary. A local filter is given a function f and a query x and, after looking up the value of f on a small number of points, it has to output $g(x)$ for some function g , which has the desired property and does not depend on x . If f has the property, g must be equal to f .

We consider functions over domains $\{0,1\}^d$, $\{1,\dots,n\}$ and $\{1,\dots,n\}^d$, equipped with ℓ_1 distance. We design efficient testers of the Lipschitz property for functions of the form $f: \{0,1\}^d \rightarrow \delta \mathbb{Z}$, where $\delta \in (0,1]$ and $\delta \mathbb{Z}$ is the set of multiples of δ , and of the form $f: \{1,\dots,n\} \rightarrow \mathbb{R}$, where \mathbb{R} is (discretely) metrically convex. In the first case, the tester runs in time $O(d \min\{d,r\}^{\delta/\epsilon})$, where r is the diameter of the image of f ; in the second, in time $O((\log n)^{1/\epsilon})$. We give corresponding lower bounds of $\Omega(d)$ and $\Omega(\log n)$ on the query complexity (in the second case, only for nonadaptive 1-sided error testers). Our lower bound for functions over $\{0,1\}^d$ is tight for the case of the $\{0,1,2\}$ range and constant ϵ . The first tester implies an algorithm for functions of the form $f: \{0,1\}^d \rightarrow \mathbb{R}$ that distinguishes Lipschitz functions from functions that are ϵ -far from $(1+\delta)$ -Lipschitz. We also present a local filter of the Lipschitz property for functions of the form $f: \{1,\dots,n\}^d \rightarrow \mathbb{R}$ with lookup complexity $O((\log n+1)^d)$. For functions of the form $\{0,1\}^d$, we show that every nonadaptive local filter has lookup complexity exponential in d .

The testers that we developed have applications to programs analysis. The reconstructors have applications to data privacy. For the first application, the Lipschitz property of the function computed by a program corresponds to a notion of robustness to noise in the data. The application to privacy is based on the fact that a function f of entries in a

database of sensitive information can be released with noise of magnitude proportional to a Lipschitz constant of f , while preserving the privacy of individuals whose data is stored in the database (Dwork, McSherry, Nissim and Smith, TCC 2006). We give a differentially private mechanism, based on local filters, for releasing a function f when a Lipschitz constant of f is provided by a distrusted client. We show that when no reliable Lipschitz constant of f is given, previously known differentially private mechanisms either have a substantially higher running time or have a higher expected error for a large class of symmetric functions f .

71.

349. Lexicographic Products and the Power of Non-Linear Network Coding

Anna Blasiak and Robert Kleinberg and Eyal Lubetzky

Abstract: We introduce a technique for establishing and amplifying gaps between parameters of network coding and index coding. The technique uses linear programs to establish separations between combinatorial and coding-theoretic parameters and applies hypergraph lexicographic products to amplify these separations. This entails combining the dual solutions of the lexicographic multiplicands and proving that they are a valid dual of the product. Our result is general enough to apply to a large family of linear programs. This blend of linear programs and lexicographic products gives a recipe for constructing hard instances in which the gap between combinatorial or coding-theoretic parameters is polynomially large. We find polynomial gaps in cases in which the largest previously known gaps were only small constant factors or entirely unknown. Most notably, we show a polynomial separation between linear and non-linear network coding rates. This involves exploiting a connection between matroids and index coding to establish a previously unknown separation between linear and non-linear index coding rates. We also construct index coding problems with a polynomial gap between the broadcast rate and the trivial lower bound for which no gap was previously known.

72.

355. Efficient computation of approximate pure Nash equilibria in congestion games

Ioannis Caragiannis and Angelo Fanelli and Nick Gravin and Alexander Skopalik

Abstract: Congestion games constitute an important class of games in which computing an exact or even approximate pure Nash equilibrium is in general P^{LS} -complete. We present a surprisingly simple polynomial-time algorithm that computes

$O(1)$ -approximate Nash equilibria in these games. In particular, for congestion games with linear latency functions, our algorithm computes $(2+\epsilon)$ -approximate pure Nash equilibria in time polynomial in the number of players, the number of resources and $1/\epsilon$. It also applies to games with polynomial latency functions with constant maximum degree d ; there, the approximation guarantee is $d^{O(d)}$. The algorithm essentially identifies a polynomially long sequence of best-response moves that lead to an approximate equilibrium; the existence of such short sequences is interesting in itself. These are the first positive algorithmic results for approximate equilibria in non-symmetric congestion games. We strengthen them further by proving that, for congestion games that deviate from our mild assumptions, computing ρ -approximate equilibria is PLS -complete for any polynomial-time computable ρ .

73.

356. How to Garble Arithmetic Circuits

Benny Applebaum and Yuval Ishai and Eyal Kushilevitz

Abstract: Yao's garbled circuit construction transforms a boolean circuit $C: \{0,1\}^n \rightarrow \{0,1\}^m$ into a "garbled circuit" \mathcal{H}_C along with n pairs of k -bit keys, one for each input bit, such that \mathcal{H}_C together with the n keys corresponding to an input x reveal $C(x)$ and no additional information about x . The garbled circuit construction is a central tool for constant-round secure computation and has several other applications.

Motivated by these applications, we suggest an efficient arithmetic variant of Yao's original construction. Our construction transforms an arithmetic circuit $C: \mathbb{Z}^n \rightarrow \mathbb{Z}^m$ over integers from a bounded (but possibly exponential) range into a garbled circuit \mathcal{H}_C along with n affine functions $L_i: \mathbb{Z} \rightarrow \mathbb{Z}^k$ such that \mathcal{H}_C together with the n integer vectors $L_i(x_i)$ reveal $C(x)$ and no additional information about x . The security of our construction relies on the intractability of the decisional variant of the learning with errors (LWE) problem.

74.

360. Rounding Semidefinite Programming Hierarchies via Global Correlation

Boaz Barak and Prasad Raghavendra and David Steurer

Abstract: We show a new way to round vector solutions of semidefinite programming (SDP) hierarchies into integral solutions, based on a connection between these hierarchies and the spectrum of the input graph. We demonstrate the utility of our method by providing a new SDP-hierarchy based algorithm for Unique Games. Our

algorithm matches the performance of the recent algorithm of Arora, Barak and Steurer (FOCS 2010) in the worst-case, but is shown to run in polynomial time on a richer family of instances, thus ruling out more possibilities for hard instances for the Unique Games Conjecture.

Specifically, we give a rounding algorithm for $O(r)$ levels of the Lasserre hierarchy that finds a good integral solution as long as, very roughly speaking, the average correlation between vectors in the SDP solution is at least $1/r$. In the case of Unique Games, the latter condition is implied by having at most r large eigenvalues in the constraint graph. This improves upon prior works that required the potentially stronger condition of a bound on the number of eigenvalues in the *extended graph*.

Our algorithm actually requires less than the $n^{O(r)}$ constraints specified by the r^{th} level of the Lasserre hierarchy, and in particular r rounds of our program can be evaluated in time $2^{O(r)} \text{poly}(n)$.

75.

363. Efficient Reconstruction of Random Multilinear Formulas

Ankit Gupta and Neeraj Kayal and Satya Lokam

Abstract: In the reconstruction problem for a multivariate polynomial f , we have blackbox access to f and the goal is to efficiently reconstruct a representation of f in a suitable model of computation. We give a polynomial time randomized algorithm for reconstructing random multilinear formulas. Our algorithm succeeds with high probability when given blackbox access to the polynomial computed by a random multilinear formula according to a natural distribution. This is the strongest model of computation for which a reconstruction algorithm is presently known, albeit efficient in a distributional sense rather than in the worst-case. Previous results on this problem considered much weaker models such as depth-3 circuits with various restrictions or read-once formulas.

Our proof uses ranks of partial derivative matrices as a key ingredient and combines it with analysis of the algebraic structure of random multilinear formulas. Partial derivative matrices have earlier been used to prove lower bounds in a number of models of arithmetic complexity, including multilinear formulas and constant depth circuits. As such, our results give supporting evidence to the general thesis that mathematical properties that capture efficient computation in a model should also enable learning algorithms for functions efficiently computable in that model.

76.

366. Markov Layout

Flavio Chierichetti and Ravi Kumar and Prabhakar Raghavan

Abstract: Consider the following problem of laying out a set of n images that match a query onto the nodes of a $\sqrt{n} \times \sqrt{n}$ grid. We are given a score for each image, as well as the distribution of patterns by which a user's eye scans the nodes of the grid and we wish to maximize the expected total score of images selected by the user. This is a special case of the *Markov layout problem*, in which we are given a Markov chain M together with a set of objects to be placed at the states of the Markov chain. Each object has a utility to the user if viewed, as well as a stopping probability with which the user ceases to look further at objects. We point out that this layout problem is prototypical in a number of applications in web search and advertising, particularly in the emerging genre of search results pages from major engines. In a different class of applications, the states of the Markov chain are web pages at a publishers website and the objects are advertisements.

In this paper we study the approximability of the Markov layout problem. Our main result is an $O(\log n)$ approximation algorithm for the most general version of the problem. The core idea behind the algorithm is to transform an optimization problem over partial permutations into an optimization problem over sets by losing a logarithmic factor in approximation; the latter problem is then shown to be submodular with two matroid constraints, which admits a constant-factor approximation. In contrast, we also show the problem is APX-hard via a reduction from *Cubic Max-Bisection*.

We then study harder variants of the problem in which no *gaps* --- states of M with no object placed on them --- are allowed. By exploiting the geometry, we obtain an $O(\log^{3/2} n)$ approximation algorithm when the digraph underlying M is a grid and an $O(\log n)$ approximation algorithm when it is a tree. These special cases are especially appropriate for our applications.

77.

368. $(1+\epsilon)$ -Approximate Sparse Recovery

Eric Price and David P. Woodruff

Abstract: The problem central to sparse recovery and compressive sensing is that of *stable sparse recovery*: we want a distribution \mathcal{A} of matrices $A \in \mathbb{R}^{m \times n}$ such that, for any $x \in \mathbb{R}^n$ and with probability $1 - \delta > 2/3$ over $A \in \mathcal{A}$, there is an algorithm to

recover \hat{x} from Ax with

$$\| \hat{x} - x \|_p \leq C \min_{\{x' \text{ } k\text{-sparse}\}} \| x - x' \|_p$$

$$\| \hat{x} - x \|_p \leq C \min_{\{x' \text{ } k\text{-sparse}\}} \| x - x' \|_p$$

$$\| \hat{x} - x \|_p \leq C \min_{\{x' \text{ } k\text{-sparse}\}} \| x - x' \|_p$$

for some constant $C > 1$ and norm $\| \cdot \|_p$.

The measurement complexity of this problem is well understood for constant $C > 1$.

However, in a variety of applications it is important to obtain $C = 1 + \epsilon$ for a small $\epsilon > 0$, and this complexity is not well understood.

We resolve the dependence on ϵ in the number of measurements required of a k -sparse recovery algorithm, up to polylogarithmic factors for the central cases of $p=1$ and $p=2$.

Namely, we give new algorithms and lower bounds that show the number of measurements required is $k \epsilon^{p/2} \text{polylog}(1/\epsilon)$. We also give matching bounds when the output is required to be k -sparse, in which case we achieve $k \epsilon^p \text{polylog}(1/\epsilon)$. This shows the distinction between the complexity of sparse and non-sparse outputs is fundamental.

78.

370. Quadratic Goldreich-Levin Theorems

Madhur Tulsiani and Julia Wolf

Abstract: Decomposition theorems in classical Fourier analysis enable us to express a bounded function in terms of few linear phases with large Fourier coefficients plus a part that is pseudorandom with respect to linear phases. The Goldreich-Levin algorithm can be viewed as an algorithmic analogue of such a decomposition as it gives a way to efficiently find the linear phases associated with large Fourier coefficients.

In the study of "quadratic Fourier analysis", higher-degree analogues of such decompositions have been developed in which the pseudorandomness property is stronger but the structured part correspondingly weaker. For example, it has previously

been shown that it is possible to express a bounded function as a sum of a few quadratic phases plus a part that is small in the U^3 norm, defined by Gowers for the purpose of counting arithmetic progressions of length 4. We give a polynomial time algorithm for computing such a decomposition.

A key part of the algorithm is a local self-correction procedure for Reed-Muller codes of order 2 (over \mathbb{F}_2^n) for a function at distance $1/2 - \epsilon$ from a codeword. Given a function $f: \mathbb{F}_2^n \rightarrow \{-1, 1\}$ at fractional Hamming distance $1/2 - \epsilon$ from a quadratic phase (which is a codeword of Reed-Muller code of order 2), we give an algorithm that runs in time polynomial in n and finds a codeword at distance at most $1/2 - \eta$ for $\eta = \eta(\epsilon)$.

This is an algorithmic analogue of Samorodnitsky's result, which gave a tester for the above problem. To our knowledge, it represents the first instance of a correction procedure for any class of codes, beyond the list-decoding radius.

In the process, we give algorithmic versions of results from additive combinatorics used in Samorodnitsky's proof and a refined version of the inverse theorem for the Gowers U^3 norm over \mathbb{F}_2^n .

79.

371. Maximizing Expected Utility for Stochastic Combinatorial Optimization Problems

Jian Li and Amol Deshpande

Abstract: We study the stochastic versions of a broad class of combinatorial problems where the weights of the elements in the input dataset are uncertain. The class of problems that we study includes shortest paths, minimum weight spanning trees, and minimum weight matchings over probabilistic graphs, and other combinatorial problems like knapsack. We observe that the expected value is inadequate in capturing different types of *risk-averse* or *risk-prone* behaviors, and instead we consider a more general objective which is to maximize the *expected utility* of the solution for some given utility function, rather than the expected weight (expected weight becomes a special case). We show that we can obtain a polynomial time approximation algorithm with *additive error* ϵ for any $\epsilon > 0$, if there is a pseudopolynomial time algorithm for the *exact* version of the problem. (This is true for the problems mentioned above). Our result generalizes several prior results on stochastic shortest path, stochastic spanning tree, and stochastic knapsack.

Our algorithm for utility maximization makes use of the separability of exponential utility and a technique to decompose a general utility function into exponential utility functions, which may be useful in other stochastic optimization problems.

80.

373. Stateless Cryptographic Protocols

Vipul Goyal and Hemanta K. Maji

Abstract: Secure computation protocols inherently involve multiple rounds of interaction among the parties where, typically a party has to keep a state about what has happened in the protocol so far and then *wait* for the other party to respond. We study if this is inherent. In particular, we study the possibility of designing cryptographic protocols where the parties can be completely stateless and compute the outgoing message by applying a single fixed function to the incoming message (independent of any state). The problem of designing stateless secure computation protocols can be reduced to the problem of designing protocols satisfying the notion of resettable computation introduced by Canetti, Goldreich, Goldwasser and Micali (FOCS'01) and widely studied thereafter.

The current state of art in resettable computation allows for construction of protocols which provide security only when a *single predetermined* party is resettable. An exception is for the case of the zero-knowledge functionality for which a protocol in which both parties are resettable was recently obtained by Deng, Goyal and Sahai (FOCS'09). The fundamental question left open in this sequence of works is, whether fully-resettable computation is possible, when:

- begin*{enumerate}
- item* An adversary can corrupt any number of parties, and
- item* The adversary can reset any party to its original state during the execution of the protocol and can restart the protocol.
- end*{enumerate}

In this paper, we resolve the above problem by constructing secure protocols realizing *any* efficiently computable multi-party functionality in the plain model under standard cryptographic assumptions. First, we construct a Fully-Resettable Simulation Sound Zero-Knowledge (ss-rs-rZK) protocol. Next, based on these ss-rs-rZK protocols, we show how to compile any semi-honest secure protocol into a protocol secure against fully resetting adversaries.

Next, we study a seemingly unrelated open question: "Does there exist a functionality which, in the concurrent setting, is impossible to securely realize using BB simulation but can be realized using NBB simulation?". We resolve the above question in the

affirmative by giving an example of such a (reactive) functionality. Somewhat surprisingly, this is done by making a connection to the existence of a fully resettable simulation sound zero-knowledge protocol.

81.

374. Steiner Shallow-Light Trees are Exponentially Lighter than Spanning Ones

Michael Elkin and Shay Solomon

Abstract: For a pair of parameters $\alpha \geq 1$, $\beta \geq 1$, a spanning tree T of a weighted undirected n -vertex graph $G = (V, E, w)$ is called an (α, β) -shallow-light tree (shortly, (α, β) -SLT) of G with respect to a designated vertex r_t in V if

- (1) it approximates all distances from r_t to other vertices up to a factor of α , and
- (2) its weight is at most β times the weight of the minimum spanning tree $\text{MST}(G)$ of G .

The parameter α (respectively, β) is called the root-distortion (resp., lightness) of the tree T .

Shallow-light trees (SLTs) constitute a fundamental graph structure, with numerous theoretical and practical applications.

In particular, they were used for constructing spanners, in network design, for VLSI-circuit design, for various data gathering and dissemination tasks in wireless and sensor networks, in overlay networks, and in the message-passing model of distributed computing.

Tight tradeoffs between the parameters of SLTs were established by Awerbuch, Baratz and Peleg, PODC'90 and Khuller, Raghavachari and Young, SODA'93. They showed that for any $\epsilon > 0$ there always exist $(1+\epsilon, O(1/\epsilon))$ -SLTs, and that the upper bound $\beta = O(1/\epsilon)$ on the lightness of SLTs cannot be improved. In this paper we show that using Steiner points one can build SLTs with logarithmic lightness, i.e., $\beta = O(\log 1/\epsilon)$.

This establishes an **exponential separation** between spanning SLTs and Steiner ones.

One particularly remarkable point on our tradeoff curve is $\epsilon = 0$.

In this regime our construction provides a **shortest-path tree** with weight at most $O(\log n) \cdot w(\text{MST}(G))$.

Moreover, we prove matching lower bounds that show that all our results are tight up to constant factors.

Finally, on our way to these results we settle (up to constant factors) a number of open questions that were raised by Khuller et al. in SODA'93.

82.

381. An algebraic proof of a robust social choice impossibility theorem

Dvir Falik and Ehud Friedgut

Abstract: An important element of social choice theory are impossibility theorem, such as Arrow's theorem and Gibbard-Satterthwaite's theorem, which state that under certain natural constraints, social choice mechanisms are impossible to construct. In recent years, beginning in Kalai, much work has been done in finding *robust* versions of these theorems, showing that impossibility remains even when the constraints are *almost* always satisfied. In this work we present an Algebraic approach for producing such results. We demonstrate it for a lesser known variant of Arrow's theorem, found in Dokow and Holzman.

83.

384. The Power of Linear Estimators

Gregory Valiant and Paul Valiant

Abstract: For a broad class of practically relevant distribution properties, which includes entropy and support size, nearly all of the proposed estimators have an especially simple form. Given a set of independent samples from a discrete distribution, these estimators tally the vector of summary statistics---the number of species seen once, twice, etc. in the sample---and output the dot product between these summary statistics, and a fixed vector of coefficients. We term such estimators *linear*. This historical proclivity towards linear estimators is slightly perplexing, since, despite many efforts over nearly 60 years, all proposed such estimators have significantly suboptimal convergence.

Our main result, in some sense vindicating this insistence on linear estimators, is that for any property in this broad class, there exists a near-optimal linear estimator. Additionally, we give a practical and polynomial-time algorithm for constructing such estimators for any given parameters.

While this result does not yield explicit bounds on the sample complexities of these estimation tasks, we leverage the insights provided by this result, to give explicit constructions of a linear estimators for three properties: entropy, L_1 distance to uniformity, and for pairs of distributions, L_1 distance.

Our entropy estimator, when given $O(\frac{n}{\epsilon \log n})$ independent samples from a distribution of support at most n , will estimate the entropy of the distribution to within accuracy ϵ , with probability of failure $o(1/\text{poly}(n))$. From recent lower bounds, this estimator is optimal, to constant factor, both in its dependence on n , and its dependence on ϵ . In particular, the inverse-linear convergence rate of this estimator resolves the main open question of [VV11], which left open the possibility that the error decreased only with the square root of the number of samples.

Our distance to uniformity estimator, on given $O(\frac{m}{\epsilon^2 \log m})$ independent samples from any distribution, returns an ϵ -accurate estimate of the L_1 distance to the uniform distribution of support m . This is the first sublinear-sample estimator for this problem, and is constant-factor optimal, for constant ϵ .

Finally, our framework extends naturally to properties of pairs of distributions, including estimating the L_1 distance and KL-divergence between pairs of distributions. We give an explicit linear estimator for estimating L_1 distance to accuracy ϵ using $O(\frac{n}{\epsilon^2 \log n})$ samples from each distribution, which is constant-factor optimal, for constant ϵ .

84.

387. The Randomness Complexity of Parallel Repetition

Kai-Min Chung and Rafael Pass

Abstract: Consider a m -round interactive protocol with soundness error $1/2$. How much extra randomness is required to decrease the soundness error to δ through parallel repetition? Previous work shows that for *public-coin* interactive protocols with *unconditional soundness*, $m \cdot O(\log(1/\delta))$ bits of extra randomness suffices. In this work, we initiate a more general study of the above question.

begin{itemize}

item We establish the first derandomized parallel repetition theorem for *public-coin* interactive protocols with *computational soundness* (a.k.a. arguments). The parameters of our result essentially matches the earlier works in the information-theoretic setting.

item We show that obtaining even a sub-linear dependency on the number of rounds m (i.e., $o(m) \cdot \log(1/\delta)$) in either the information-theoretic or computational settings requires proving $P \neq PSPACE$.

item We show that non-trivial derandomized parallel repetition for *private-coin* protocols is impossible in the information-theoretic setting, and requires proving $P \neq PSPACE$ in the computational setting.

end{itemize}

85.

391. The Complexity of Quantum States - a combinatorial approach

Dorit Aharonov and Itai Arad and Zeph Landau and Umesh Vazirani

Abstract: The classical description of quantum states is in general exponential in the number of qubits. Can we get polynomial descriptions for more restricted sets of states such as ground states of interesting subclasses of local Hamiltonians? This is the basic problem in the study of the complexity of ground states, and requires an understanding of multi-particle entanglement and quantum correlations in such states.

We propose a combinatorial approach to this question, based on a reformulation of the detectability lemma introduced by us in the context of quantum gap amplification \cite{ref:Aha09b}. We give an alternative proof of the detectability lemma which is not only simple and intuitive, but also removes a key restriction in the original statement, making it more suitable for this new context. We then provide a one page proof of Hastings' proof that the correlations in the ground states of Gapped Hamiltonians decay exponentially with the distance, demonstrating the simplicity of the combinatorial approach for those problems.

As our main application, we provide a combinatorial proof of Hastings' seminal 1D area law \cite{ref:Has07} for the special case of frustration free systems. Area laws provide a fundamental ingredient in the study of the complexity of ground states, since they offer a way to bound in a quantitative way the entanglement in such states. An intricate combinatorial analysis allows us to significantly improve the bounds achieved in Hastings proofs, and derive an exponentially better scaling in terms of the inverse spectral gap and the dimensionality of the particles. This holds out hope that the new approach might be a promising route towards resolving the 2D case and higher dimensions, which is one of the most important open questions in Hamiltonian complexity.
