

Track A

1. **Best paper in Track A**

Mark Bun and Justin Thaler. *Dual Lower Bounds for Approximate Degree and Markov-Bernstein Inequalities*

Abstract: The ϵ -approximate degree of a Boolean function $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ is the minimum degree of a real polynomial that approximates f to within ϵ in the ℓ_∞ norm. We prove several lower bounds on this important complexity measure by explicitly constructing solutions to the dual of an appropriate linear program. Our first result resolves the ϵ -approximate degree of the two-level AND-OR tree for any constant $\epsilon > 0$. We show that this quantity is $\Theta(\sqrt{n})$, closing a line of incrementally larger lower bounds. The same lower bound was recently obtained independently by Sherstov using related techniques. Our second result gives an explicit dual polynomial that witnesses a tight lower bound for the approximate degree of any symmetric Boolean function, addressing a question of V{S}palek. Our final contribution is to reprove several Markov-type inequalities from approximation theory by constructing explicit dual solutions to natural linear programs. These inequalities underlie the proofs of many of the best-known approximate degree lower bounds, and have important uses throughout theoretical computer science.

See [this](#) blog post by the authors for a discussion of details of the paper.

2. **Best student paper in Track A**

Radu Curticapean. *Counting matchings of size k is #W[1]-hard*

Abstract: We prove #W[1]-hardness of the following parameterized counting problem: Given a simple undirected graph G and a parameter k , compute the number of matchings of size k in G .

It is known from (Bläser, Curticapean 2012) that, given an edge-weighted graph G , computing a particular weighted sum over the matchings in G is #W[1]-hard. In the present paper, we exhibit a reduction that does not require weights.

This solves an open problem from (Flum, Grohe 2006) and adds a natural parameterized counting problem to the scarce list of #W[1]-hard problems. Since the classical version of this problem has been thoroughly studied, we believe that our result facilitates future #W[1]-hardness proofs for other problems.

3. **Oren Weimann and Raphael Yuster.** *Approximating the Diameter of Planar Graphs in Near Linear Time*

Abstract: We present a $(1+\epsilon)$ -approximation algorithm running in $\tilde{O}(\epsilon^{-1} n)$ time for finding the diameter of an undirected planar graph with non-negative edge lengths.

4. **Vladimir Braverman, Rafail Ostrovsky and Dan Vilenchik.** *How Hard is Counting Triangles in the Streaming Model*

Abstract: The problem of (approximately) counting the number of triangles in a graph is one of the basic problems in graph theory. In this paper we study the problem in the streaming model. We study the amount of memory required by a randomized algorithm to solve this problem. In case the algorithm is allowed one pass over the stream, we present a best possible lower bound of $\Omega(m)$ for graphs G with m edges on n vertices. If a constant number of passes is allowed, we show a lower bound of $\Omega(m/T)$, T the number of triangles. We match, in some sense, this lower bound with a 2-pass $O(m/T^{1/3})$ -memory algorithm that solves the problem of distinguishing graphs with no triangles from graphs with at least T triangles. We present a new graph parameter $\rho(G)$ -- the triangle density, and conjecture that the space complexity of the triangles problem is $\Omega(m/\rho(G))$. We match this by a second algorithm that solves the distinguishing problem using $O(m/\rho(G))$ -memory.

5. **Tom Gur and Ran Raz.** *Arthur-Merlin Streaming Complexity*

Abstract: We study the power of Arthur-Merlin probabilistic proof systems in the data stream model. We show a canonical AM streaming algorithm for a wide class of data stream problems. The algorithm offers a tradeoff between the length of the proof and the space complexity that is needed to verify it.

As an application, we give an AM streaming algorithm for the *Distinct Elements* problem. Given a data stream of length m over alphabet of size n , the algorithm uses $\tilde{O}(s)$ space and a proof of size $\tilde{O}(w)$, for every s, w such that $s \cdot w \geq n$ (where \tilde{O} hides a $\text{polylog}(m, n)$ factor). We also prove a lower bound, showing that every MA streaming algorithm for the *Distinct Elements* problem that uses s bits of space and a proof of size w , satisfies $s \cdot w = \Omega(n)$.

As a part of the proof of the lower bound for the *Distinct Elements* problem, we show a new lower bound of $\Omega(\sqrt{n})$ on the MA communication complexity of the *Gap Hamming Distance* problem, and prove its tightness.

6. **Vladimir Kolmogorov.** *The power of linear programming for finite-valued CSPs: a constructive characterization*

Abstract: A class of valued constraint satisfaction problems (VCSPs) is characterised by a valued constraint language, a fixed set of cost functions on a finite domain. An instance of the problem is specified by a sum of cost functions from the language with the goal to minimise the sum. We study which classes of finite-valued languages can be solved exactly by the basic linear programming relaxation (BLP). Thapper and Zivny showed [20] that if BLP solves the language then the language admits a binary commutative fractional polymorphism. We prove that the converse is also true. This leads to a necessary and a sufficient condition which can be checked in polynomial time for a given language. In contrast, the previous necessary and sufficient condition due to [20] involved infinitely many inequalities.

More recently, Thapper and Zivny [21] showed (using, in particular, a technique introduced in this paper) that core languages that do not satisfy our condition are NP-hard. Taken together, these results imply that a finite-valued language can either be solved using Linear Programming or is NP-hard.

7. **Eun Jung Kim, Alexander Langer, Christophe Paul, Felix Reidl, Peter Rossmanith, Ignasi Sau and Somnath Sikdar.** *Linear kernels and single-exponential algorithms via protrusion decompositions*

Abstract: We present a linear-time algorithm to compute a decomposition scheme for graphs G that have a set $X \subseteq V(G)$, called a *treewidth-modulator*, such that the treewidth of $G-X$ is bounded by a constant. Our decomposition, called a *protrusion decomposition*, is the cornerstone in obtaining the following two main results.

Our first result is that any parameterized graph problem (with parameter k) that has *finite integer index* and such that positive instances have a treewidth-modulator of size $O(k)$ admits a linear kernel on the class of H -topological-minor-free graphs, for any fixed graph H . This result partially extends previous meta-theorems on the existence of linear kernels on graphs of bounded genus and H -minor-free graphs.

Let \mathcal{F} be a fixed finite family of graphs containing at least one planar graph. Given an n -vertex graph G and a non-negative integer k , *Planar- \mathcal{F} -Deletion* asks whether G has a set $X \subseteq V(G)$ such that $|X| \leq k$ and $G-X$ is H -minor-free for every $H \in \mathcal{F}$. As our second application, we present the first *single-exponential* algorithm to solve *Planar- \mathcal{F} -Deletion*. Namely, our algorithm runs in time $2^{O(k)} \cdot n^2$, which is asymptotically optimal with respect to k . So far, single-exponential algorithms were only known for special cases of the family \mathcal{F} .

8. **Martin Hirt and Pavel Raykov.** *On the Complexity of Broadcast Setup*

Abstract: Byzantine broadcast is a distributed primitive that allows a specific party (called "sender") to consistently distribute a value v among n parties in the presence of potential misbehavior of up to t of the parties. Broadcast requires that correct parties always agree on the same value and if the sender is correct, then the agreed value is v . Broadcast without a setup (i.e., from scratch) is achievable from point-to-point channels if and only if $t < n/3$. In case $t \geq n/3$ a trusted setup is required. The setup may be assumed to be given initially or generated by the parties in a setup phase.

It is known that generating setup for protocols with cryptographic security is relatively simple and only consists of setting up a public-key infrastructure. However, generating setup for information-theoretically secure protocols is much more involved. In this paper we study the complexity of setup generation for information-theoretic protocols using point-to-point channels and temporarily available broadcast channels. We optimize the number of rounds in which the temporary broadcast channels are used while minimizing the number of bits broadcast with them. We give the first information-theoretically secure broadcast protocol tolerating $t < n/2$ that uses the temporary broadcast channels during only 1 round in the setup phase. Furthermore, only $O(n^3)$ bits need to be broadcast with the temporary broadcast channels during that round, independently of the security parameter employed. The broadcast protocol presented in this paper allows to construct the first information-theoretically secure MPC protocol which uses a broadcast channel during only one round. Additionally, the presented broadcast protocol supports refreshing, which allows to broadcast an a priori unknown number of times given a fixed-size setup.

9. **Karl Bringmann and Tobias Friedrich.** *Exact and Efficient Generation of Geometric Random Variates and Random Graphs*

Abstract: The standard algorithm for fast generation of Erdos-Renyi random graphs only works in the Real RAM model. The critical point is the generation of geometric random variates $\text{Geo}(p)$, for which there is no algorithm that is both exact and efficient in any bounded precision machine model. For a RAM model with word size $w = O(\log\log(1/p))$, we show that this is possible and present an exact algorithm for sampling $\text{Geo}(p)$ in optimal expected time $O(1 + \log(1/p)/w)$. We also give an exact algorithm for sampling $\min\{n, \text{Geo}(p)\}$ in optimal expected time $O(1 + \log(\min\{1/p, n\})/w)$. This yields a new exact algorithm for sampling Erdos-Renyi and Chung-Lu random graphs of n vertices and m (expected) edges in optimal expected runtime $O(n + m)$ on a RAM with word size $w = T(\log n)$.

10. **Joseph Cheriyan, Zhihan Gao, Konstantinos Georgiou and Sahil Singla.** *On Integrality Ratios for Asymmetric TSP in the Sherali-Adams Hierarchy*

Abstract: We study the ATSP (Asymmetric Traveling Salesman Problem), and our focus is on negative results in the framework of the Sherali-Adams (iSA) Lift and Project method.

Our main result pertains to the standard LP (linear programming) relaxation of ATSP, due to Dantzig, Fulkerson, and Johnson. For any fixed integer $t \geq 0$ and small ϵ , $0 < \epsilon \ll 1$, there exists a digraph G on $n = n(t, \epsilon) = O(t/\epsilon)$ vertices such that the integrality ratio for level t of the iSA system starting with the standard LP on G is $\geq 1 + \frac{1 - \epsilon}{2t + 3} \approx \frac{43}{65}, \frac{87}{\dots}$. Thus, in terms of the input size, the result holds for any $t = 0, 1, \dots, \Theta(n)$ levels. Our key contribution is to identify a structural property of digraphs that allows us to construct fractional feasible solutions for any level t of the iSA system starting from the standard LP. Our hard instances are simple and satisfy the structural property.

There is a further relaxation of the standard LP called the balanced LP, and our methods simplify considerably when the starting LP for the iSA system is the balanced LP; in particular, the relevant structural property (of digraphs) simplifies such that it is satisfied by the digraphs given by the well-known construction of Charikar, Goemans and Karloff (CGK). Consequently, the CGK digraphs serve as hard instances, and we obtain an integrality ratio of $1 + \frac{1 - \epsilon}{t + 1}$ for any level t of the iSA system, where $0 < \epsilon \ll 1$.

Also, our results for the standard LP extend to the path-ATSP (find a min cost Hamiltonian dipath from a given source vertex to a given sink vertex).

11. **Nikos Leonardos.** *An improved lower bound for the randomized decision tree complexity of recursive majority*

Abstract: We prove that the randomized decision tree complexity of the recursive majority-of-three is $\Omega(2.55^d)$, where d is the depth of the recursion. The proof is by a bottom up induction, which is same in spirit as the one in the proof of Saks and Wigderson in their 1986 paper on the complexity of evaluating game trees. Previous work includes an $\Omega((7/3)^d)$ lower bound, published in 2003 by Jayram, Kumar, and Sivakumar. Their

proof used a top down induction and tools from information theory. In 2011, Magniez, Nayak, Santha, and Xiao, improved the lower bound to $\Omega(2.5^d)$ and the upper bound to $O(2.64946^d)$.

12. **Petr Golovach, Pinar Heggernes, Dieter Kratsch and Yngve Villanger.** *An incremental polynomial time algorithm to enumerate all minimal edge dominating sets*

Abstract: We show that all minimal edge dominating sets of a graph can be generated in incremental polynomial time. We present an algorithm that solves the equivalent problem of enumerating minimal (vertex) dominating sets of line graphs in incremental polynomial, and consequently output polynomial, time. Enumeration of minimal dominating sets in graphs has very recently been shown to be equivalent to enumeration of minimal transversals in hypergraphs. The question whether the minimal transversals of a hypergraph can be enumerated in output polynomial time is a fundamental and challenging question; it has been open for several decades and has triggered extensive research. To obtain our result, we present a flipping method to generate all minimal dominating sets of a graph. Its basic idea is to apply a flipping operation to a minimal dominating set D^* to generate minimal dominating sets D such that $G[D]$ contains more edges than $G[D^*]$. We show that the flipping method works efficiently on line graphs, resulting in an algorithm with delay $O(n^2 m^{|\mathcal{L}|})$ between each output minimal dominating set, where n and m are the numbers of vertices and edges of the input graph, respectively, and \mathcal{L} is the set of already generated minimal dominating sets. Furthermore, we are able to improve the delay to $O(n^2 m^{|\mathcal{L}|})$ on line graphs of bipartite graphs. Finally we show that the flipping method is also efficient on graphs of large girth, resulting in an incremental polynomial time algorithm to enumerate the minimal dominating sets of graphs of girth at least 7.

13. **Philip Bille, Inge Li Gørtz, Gad Landau and Oren Weimann.** *Tree Compression with Top Trees*

Abstract: We introduce a new compression scheme for labeled trees based on top trees~\cite{TopTrees}. Our compression scheme is the first to simultaneously take advantage of internal repeats in the tree (as opposed to the classical DAG compression that only exploits rooted subtree repeats) while also supporting fast navigational queries directly on the compressed representation. We show that the new compression scheme achieves close to optimal worst-case compression, can compress exponentially better than DAG compression, is never much worse than DAG compression, and supports navigational queries in logarithmic time.

14. **Claire Mathieu and Hang Zhou.** *Graph reconstruction via distance oracles*

Abstract: We study the problem of reconstructing a hidden graph given only access to a distance oracle. We first give an algorithm to reconstruct degree bounded graphs using $\tilde{O}(n^{3/2})$ queries on average. Next, we focus on degree bounded outerplanar graphs. We give an algorithm to reconstruct such a graph using $\tilde{O}(n)$ queries on average. Finally, we consider the approximate reconstruction problem on general graphs. We give a near-optimal algorithm whose query complexity coincides with the query lower bound up to a logarithmic factor.

15. **Andrei Negoescu and Gabriel Moruz.** *Improved Space Bounds for Strongly Competitive Randomized Paging Algorithms*

Abstract: Paging is a prominent problem in the field of online algorithms. While in the deterministic setting there exist simple and efficient strongly competitive algorithms, in the randomized setting a tradeoff between competitiveness and memory is still not settled. In this paper we address the conjecture by Bein et al., that there exist strongly competitive randomized paging algorithms using $o(k)$ bookmarks, i.e. pages not in cache that the algorithm keeps track of. We prove tighter bounds for Equitable2, showing that it requires less than k bookmarks, more precisely approximately $0.62k$. We then give a lower bound for Equitable2 showing that it cannot both be strongly competitive and use $o(k)$ bookmarks. Our main result proves the conjecture that there exist strongly competitive paging algorithms using $o(k)$ bookmarks. We propose an algorithm, denoted Partition2, which is a variant of the Partition algorithm by McGeoch and Sleator. While Partition is unbounded in its space requirements, Partition2 uses $\Theta(k/\log k)$ bookmarks.

16. **Ran Duan and Kurt Mehlhorn.** *A Combinatorial Polynomial Algorithm for the Linear Arrow-Debreu Market*

Abstract: We present the first combinatorial polynomial time algorithm for computing the equilibrium of the Arrow-Debreu market model with linear utilities. Our algorithm views the allocation of money as flows and iteratively improves the balanced flow as in [Devanur et al. 2008] for Fisher's model. We develop new methods to carefully deal with the flows and surpluses during price adjustments. In our algorithm, we need $O(n^6 \log(nU))$ maximum flow computations, where n is the number of persons and U is the maximum integer utility, and the length of the numbers is at most $O(n \log(nU))$ to guarantee an exact solution. Previously, [Jain 2007, Ye 2007] have given polynomial time algorithms for this problem, which are based on solving convex programs using the ellipsoid algorithm and interior-point method.

17. **Anna Gilbert, Hung Ngo, Ely Porat, Atri Rudra and Martin Strauss.** *l2/l2-foreach sparse recovery with low risk*

Abstract: In this paper, we consider the "foreach" sparse recovery problem with failure probability p . The goal of the problem is to design a distribution over $m \times N$ matrices Φ and a decoding algorithm algo such that for every $v \in \mathbb{R}^N$, we have the following error guarantee with probability at least $1-p$ $\|v - \text{algo}(\Phi v)\|_2 \leq C \|v - v_k\|_2$ where C is a constant (ideally arbitrarily close to 1) and v_k is the best k -sparse approximation of v .

Much of the sparse recovery and compressive sensing literature has focused on the case of either $p = 0$ or $p = \Omega(1)$. We address this problem for the entire range of failure probability. Our two main results are as follows. (1) We prove a lower bound on m , the number measurements, of $\Omega(k \log(n/k) + \log(1/p))$ for $2^{-\Theta(N)} \leq p < 1$. Cohen, Dahmen, and DeVore [CDD2007:NearOptimal] prove that this bound is tight. (2) We prove nearly matching upper bounds that also admit sub-linear time decoding. Previous such results were obtained only when $p = O(1)$. One corollary of our result is an extension of Gilbert et al. [GHR12:SimpleSignals] results for information-theoretically bounded adversaries.

18. **Massimo Lauria, Pavel Pudlák, Vojtěch Rödl and Neil Thapen.** *The complexity of proving that a graph is Ramsey*

Abstract: We say that a graph with n vertices is c -Ramsey if it does not contain either a clique or an independent set of size $c \log n$. We define a CNF formula which expresses this property for a graph G . We show a superpolynomial lower bound on the length of resolution proofs that G is c -Ramsey, for every graph G . Our proof makes use of the fact that every Ramsey graph must contain a large subgraph with some of the statistical properties of the random graph.

19. **Reut Levi and Dana Ron.** *A Quasi-Polynomial Time Partition Oracle for Graphs with an Excluded Minor*

Abstract: Motivated by the problem of testing planarity and related properties, we study the problem of designing efficient *partition oracles*. A *partition oracle* is a procedure that, given access to the incidence lists representation of a bounded-degree graph $G = (V, E)$ and a parameter ϵ , when queried on a vertex $v \in V$, returns the part (subset of vertices) which v belongs to in a partition of all graph vertices. The partition should be such that all parts are small, each part is connected, and if the graph has certain properties, the total number of edges between parts is at most $\epsilon |V|$. In this work we give a partition oracle for graphs with excluded minors whose query complexity is quasi-polynomial in $1/\epsilon$, thus improving on the result of Hassidim et al. (*Proceedings of FOCS 2009*) who gave a partition oracle with query complexity exponential in $1/\epsilon$. This improvement implies corresponding improvements in the complexity of testing planarity and other properties that are characterized by excluded minors as well as sublinear-time approximation algorithms that work under the promise that the graph has an excluded minor.

20. **Heng Guo and Tyson Williams.** *The Complexity of Planar Boolean #CSP with Complex Weights*

Abstract: We prove a complexity dichotomy theorem for symmetric complex-weighted Boolean #CSP when the constraint graph of the input must be planar. The problems that are #P-hard over general graphs but tractable over planar graphs are precisely those with a holographic reduction to matchgates. This generalizes a theorem of Cai, Lu, and Xia for the case of real weights. We also obtain a dichotomy theorem for a symmetric arity 4 signature with complex weights in the planar Holant framework, which we use in the proof of our #CSP dichotomy. In particular, we reduce the problem of evaluating the Tutte polynomial of a planar graph at the point $(3,3)$ to counting the number of Eulerian orientations over planar 4-regular graphs to show the latter is #P-hard. This strengthens a theorem by Huang and Lu to the planar setting. Our proof techniques combine new ideas with refinements and extensions of existing techniques. These include planar pairings, the recursive unary construction, the anti-gadget technique, and pinning in the Hadamard basis.

21. **Nicole Megow and José Verschae.** *Dual techniques for scheduling on a machine with varying speed*

Abstract: We study scheduling problems on a machine of varying speed. Assuming a known speed function (given through an oracle) we ask for a cost-efficient scheduling solution. Our main result is a PTAS for minimizing the total weighted completion time on a machine of

varying speed. This implies also a PTAS for the closely related problem of scheduling to minimize generalized global cost functions. The key to our results is a re-interpretation of the problem within the well-known *{\em two-dimensional Gantt chart}*: instead of the standard approach of scheduling in the *{\em time-dimension}*, we construct scheduling solutions in the *{\em weight-dimension}*.

We also consider a dynamic problem variant in which deciding upon the speed is part of the scheduling problem and we are interested in the tradeoff between scheduling cost and speed-scaling cost, which is typically the energy consumption. We obtain two insightful results: (1) the optimal scheduling order is independent of the energy consumption and (2)~the problem can be reduced to the setting where the speed of the machine is fixed, and thus admits a PTAS. We also give complexity results, more efficient algorithms for special cases, and a simple $(2+\epsilon)$ -approximation for speed-scaling of parallel processors and jobs with release dates.

22. **Thomas Bläsius, Ignaz Rutter and Dorothea Wagner.** *Optimal Orthogonal Graph Drawing with Convex Bend Costs*

Abstract: Traditionally, the quality of orthogonal planar drawings is quantified by the total number of bends, or the maximum number of bends per edge. However, this neglects that in typical applications, edges have varying importance. We consider the problem *OptimalFlexDraw* that is defined as follows. Given a planar graph G on n vertices with maximum degree 4 and for each edge e a cost function $\text{cost}_e : \mathbb{N}_0 \rightarrow \mathbb{R}$ defining costs depending on the number of bends e has, compute an orthogonal drawing of G of minimum cost.

In this generality *OptimalFlexDraw* is NP-hard. We show that it can be solved efficiently if 1) the cost function of each edge is convex and 2) the first bend on each edge does not cause any cost. Our algorithm takes time $O(n * T_{\text{flow}}(n))$ and $O(n^2 * T_{\text{flow}}(n))$ for biconnected and connected graphs, respectively, where $T_{\text{flow}}(n)$ denotes the time to compute a minimum-cost flow in a planar network with multiple sources and sinks. Our result is the first polynomial-time bend-optimization algorithm for general 4-planar graphs optimizing over all embeddings. Previous work considers restricted graph classes and unit costs.

23. **Hans Bodlaender, Marek Cygan, Stefan Kratsch and Jesper Nederlof.** *Deterministic Single Exponential Time Algorithms for Connectivity Problems Parameterized by Treewidth*

Abstract: It is well known that many local graph problems, like Vertex Cover and Dominating Set, can be solved in $2^{O(tw)}|V|^{O(1)}$ time for graphs $G=(V,E)$ with a given tree decomposition of width tw . However, for nonlocal problems, like the fundamental class of connectivity problems, for a long time it was unknown how to do this faster than $tw^{O(tw)}|V|^{O(1)}$ until recently, when Cygan et al. (FOCS 2011) introduced the Cut and Count technique that gives randomized algorithms for a wide range of connectivity problems running in time $c^{tw}|V|^{O(1)}$ for a small constant c .

In this paper, we show that we can improve upon the Cut and Count technique in multiple ways, with two new techniques. The first technique (rank-based approach) gives deterministic algorithms with $O(c^{tw}|V|)$ running time for connectivity problems (like Hamiltonian Cycle and Steiner Tree) and for weighted variants of these; the second technique

(determinant approach) gives deterministic algorithms running in time $c^{tw} |V|^{O(1)}$ for counting versions, e.g., counting the number of Hamiltonian cycles for graphs of treewidth tw .

The rank-based approach introduces a new technique to speed up dynamic programming algorithms which is likely to have more applications. The determinant-based approach uses the matrix tree theorem for deriving closed formulas for counting versions of connectivity problems; we show how to evaluate those formulas via dynamic programming.

24. **Anindya De, Ilias Diakonikolas and Rocco Servedio.** *A robust Khintchine inequality, and algorithms for computing optimal constants in Fourier analysis and high-dimensional geometry*

Abstract: This paper makes two contributions towards determining some well-studied optimal constants in Fourier analysis of Boolean functions and high-dimensional geometry.

`\begin{enumerate}`

`\item` It has been known since 1994 `\cite{GL:94}` that every linear threshold function has squared Fourier mass at least $1/2$ on its degree-0 and degree-1 coefficients. Denote the minimum such Fourier mass by $w^{\leq 1}$, where the minimum is taken over all n -variable linear threshold functions and all $n \geq 0$. Benjamini, Kalai and Schramm `\cite{BKS:99}` have conjectured that the true value of $w^{\leq 1}$ is $2/\pi$.

We make progress on this conjecture by proving that $w^{\leq 1} \geq 1/2 + c$ for some absolute constant $c > 0$. The key ingredient in our proof is a "robust" version of the well-known Khintchine inequality in functional analysis, which we believe may be of independent interest.

`\item` We give an algorithm with the following property: given any $\epsilon > 0$, the algorithm runs in time $2^{O(\text{poly}(1/\epsilon))}$ and determines the value of $w^{\leq 1}$ up to an additive error of $\pm \epsilon$. We give a similar $2^{O(\text{poly}(1/\epsilon))}$ -time algorithm to determine `\emph{Tomaszewski's constant}` to within an additive error of $\pm \epsilon$; this is the minimum (over all origin-centered hyperplanes H) fraction of points in $\{-1, 1\}^n$ that lie within Euclidean distance 1 of H . Tomaszewski's constant is conjectured to be $1/2$; lower bounds on it have been given by Holzman and Kleitman `\cite{HK92}` and independently by Ben-Tal, Nemirovski and Roos `\cite{BNR02}`. Our algorithms combine tools from anti-concentration of sums of independent random variables, Fourier analysis, and Hermite analysis of linear threshold functions.

`\end{enumerate}`

25. **Michael Lampis.** *Model Checking Lower Bounds for Simple Graphs*

Abstract: A well-known result by Frick and Grohe shows that deciding FO logic on trees involves a parameter dependence that is a tower of exponentials. Though this lower bound is tight for Courcelle's theorem, it has been evaded by a series of recent meta-theorems for other graph classes. Here we provide some additional non-elementary lower bound results, which are in some senses stronger. Our goal is to explain common traits in these recent meta-theorems and identify barriers to further progress.

More specifically, first, we show that on the class of threshold graphs, and therefore also on any union and complement-closed class, there is no model-checking algorithm with elementary parameter dependence even for FO logic. Second, we show that there is no

model-checking algorithm with elementary parameter dependence for MSO logic even restricted to paths (or equivalently to unary strings), unless $\text{EXP}=\text{NEXP}$. As a corollary, we resolve an open problem on the complexity of MSO model-checking on graphs of bounded max-leaf number. Finally, we look at MSO on the class of colored trees of depth d . We show that, assuming the ETH, for every fixed $d \geq 1$ at least $d+1$ levels of exponentiation are necessary for this problem, thus showing that the $(d+1)$ -fold exponential algorithm recently given by Gajarsky and Hlineny is essentially optimal.

26. **S. Anand, Karl Bringmann, Tobias Friedrich, Naveen Garg and Amit Kumar.**
Minimizing maximum (weighted) flow-time on related and unrelated machines

Abstract: In this paper we initiate the study of job scheduling on related and unrelated machines so as to minimize the maximum flow time or the maximum weighted flow time (when each job has an associated weight). Previous work for these metrics considered only the setting of parallel machines, while previous work for scheduling on unrelated machines only considered $L_p, p < 8$ norms. Our main results are:

- (i) We give an $O(e-3)$ -competitive algorithm to minimize maximum weighted flow time on related machines where we assume that the machines of the online algorithm can process $1 + e$ units of a job in 1 time-unit (e speed augmentation).
- (ii) For the objective of minimizing maximum flow time on unrelated machines we give a simple $2/e$ -competitive algorithm when we augment the speed by e . For m machines we show a lower bound of $O(m)$ on the competitive ratio if speed augmentation is not permitted. Our algorithm does not assign jobs to machines as soon as they arrive. To justify this "drawback" we show a lower bound of $O(\log m)$ on the competitive ratio of immediate dispatch algorithms. In both these lower bound constructions we use jobs whose processing times are in $\{1, 8\}$, and hence they apply to the more restrictive subset parallel setting.
- (iii) For the objective of minimizing maximum weighted flow time on unrelated machines we establish a lower bound of $O(\log m)$ on the competitive ratio of any online algorithm which is permitted to use $s = O(1)$ speed machines. In our lower bound construction, job j has a processing time of p_j on a subset of machines and infinity on others and has a weight $1/p_j$. Hence this lower bound applies to the subset parallel setting for the special case of minimizing maximum stretch.

27. **Marcin Mucha and Maxim Sviridenko.** *No-Wait Flowshop Scheduling is as Hard as Asymmetric Traveling Salesman Problem*

Abstract: In this paper we study the classical no-wait flowshop scheduling problem with makespan objective ($F_{\text{no-wait}}|C_{\text{max}}$ in the standard three-field notation). This problem is well-known to be a special case of the asymmetric traveling salesman problem (atasp) and as such has an approximation algorithm with logarithmic performance guarantee. In this work we show a reverse connection, we show that any polynomial time α -approximation algorithm for the no-wait flowshop scheduling problem with makespan objective implies the existence of a polynomial-time $\alpha(1+\epsilon)$ -approximation algorithm for the atasp , for any $\epsilon > 0$. This in turn implies that all non-approximability results for the atasp (current or future) will carry over to its special case. In particular, it follows that no-wait flowshop problem is APX-hard, which is the first non-approximability result for this problem.

28. **David Avis and Hans Raj Tiwary.** *On the extension complexity of combinatorial polytopes*

Abstract: In this paper we extend recent results of Fiorini et al. on the extension complexity of the cut polytope and related polyhedra. We first describe a lifting argument to show exponential extension complexity for a number of NP-complete problems including subset-sum and three dimensional matching. We then obtain a relationship between the extension complexity of the cut polytope of a graph and that of its graph minors. Using this we are able to show exponential extension complexity for the cut polytope of a large number of graphs, including those used in quantum information and suspensions of cubic planar graphs.

29. **Maxim Sviridenko and Justin Ward.** *Large Neighborhood Local Search for the Maximum Set Packing Problem*

Abstract: In this paper we consider the classical maximum set packing problem where set cardinality is upper bounded by k . We show how to design a variant of a polynomial-time local search algorithm with performance guarantee $(k+2)/3$. This local search algorithm is a special case of a more general procedure that allows to swap up to $\Theta(\log n)$ elements per iteration. We also design problem instances with locality gap $k/3$ even for a wide class of exponential time local search procedures, which can swap up to cn elements for a constant c . This shows that our analysis of this class of algorithms is almost tight.

30. **Roberto Grossi, Rajeev Raman, Srinivasa Rao Satti and Rossano Venturini.** *Dynamic Compressed Strings with Random Access*

Abstract: We consider the problem of storing a string S in dynamic compressed form, while permitting operations directly on the compressed representation of S : access a substring of S ; replace, insert or delete a symbol in S ; count how many occurrences of a given symbol appear in any given prefix of S (called rank operation) and locate the position of the i th occurrence of a symbol inside S (called select operation). We discuss the time complexity of several combinations of these operations along with the entropy space bounds of the corresponding compressed indexes. In this way, we extend or improve the bounds of previous work by Ferragina and Venturini [TCS, 2007], Jansson et al. [ICALP, 2012], Nekrich and Navarro [SODA, 2013].

31. **Brett Hemenway, Rafail Ostrovsky and Mary Wootters.** *Local Correctability of Expander Codes*

Abstract: In this work, we present the first local-decoding algorithm for expander codes. This yields a new family of constant-rate codes that can recover from a constant fraction of errors in the codeword symbols, and where any symbol of the codeword can be recovered with high probability by reading N^ϵ symbols from the corrupted codeword, where N is the block-length of the code.

Expander codes, introduced by Sipser and Spielman, are formed from an expander graph $G = (V, E)$ of degree d , and an inner code of block-length d over an alphabet Σ . Each edge of the expander graph is associated with a symbol in Σ . A string in Σ^E will be a codeword if for each vertex in V , the symbols on the adjacent edges form a codeword in the inner code.

In this work, we show that if the inner code has a smooth reconstruction algorithm in the noiseless setting, then the corresponding expander code has an efficient local-correction algorithm in the noisy setting. Instantiating our construction with inner codes based on finite geometries, we obtain a new family of locally-decodable codes with constant rate. This provides an alternative to the multiplicity codes of Kopparty, Saraf and Yekhanin (STOC '11) and the lifted codes of Guo, Kopparty and Sudan (ITCS '13). Our codes are the first constant-rate locally decodable codes that have a local-decoding procedure that is linear in the number of queries.

32. **Karl Wimmer and Yuichi Yoshida.** *Testing Linear-Invariant Function Isomorphism*

Abstract: A function $f: \mathbb{F}_2^n \rightarrow \{-1, 1\}$ is called linear-isomorphic to g if $f = g \circ A$ for some non-singular matrix A . In the g -isomorphism problem, we want a randomized algorithm that distinguishes whether an input function f is linear-isomorphic to g or far from being so.

We show that the query complexity to test g -isomorphism is essentially determined by the spectral norm of g . That is, if g is close to having spectral norm s , then we can test g -isomorphism with $\text{poly}(s)$ queries, and if g is far from having spectral norm s , then we cannot test g -isomorphism with $o(\log s)$ queries. The upper bound is almost tight since there is indeed a function g close to having spectral norm s whereas testing g -isomorphism requires $\Omega(s)$ queries. As far as we know, our result is the first characterization of this type for functions. Our upper bound is essentially the Kushilevitz-Mansour learning algorithm, modified for use in the implicit setting.

Exploiting our upper bound, we show that any property is testable if it can be well-approximated by functions with small spectral norm. We also extend our algorithm to the setting where A is allowed to be singular.

33. **Telikepalli Kavitha and Nithin Varma.** *Small Stretch Pairwise Spanners*

Abstract: Let $G = (V, E)$ be an undirected unweighted graph on n vertices. A subgraph H of G is called an (a, b) spanner of G if for each pair of vertices (u, v) , the u - v distance in H is at most $a \cdot \delta(u, v) + b$, where $\delta(u, v)$ is the u - v distance in G . The following is a natural relaxation of the above problem: we care for only certain distances, these are captured by the set $P \subseteq V \times V$ and the problem is to construct a sparse subgraph H , called an (a, b) P -spanner, where for every $(u, v) \in P$, the u - v distance in H is at most $a \cdot \delta(u, v) + b$. In this paper we show algorithms to construct the following:

-- a $(1, 2)$ P -spanner of size $\tilde{O}(n \cdot |P|^{1/3})$ for general P -- a $(1, 2)$ P -spanner of size $\tilde{O}(n \cdot |P|^{1/4})$ when $P = S \times V$ for some $S \subseteq V$.

Our $(1, 2)$ P -spanner is sparser than the $(1, 2)$ spanner and the P -preserver (a subgraph where distances between pairs in P are exactly preserved) for a large range of values of $|P|$. Our $(1, 2)$ $(S \times V)$ -spanner leads to a simple construction of a $(1, 4)$ spanner.

Let a D -spanner refer to a P -spanner when P is described implicitly via a distance threshold D as: $P = \{(u, v): \delta(u, v) \geq D\}$. An $O(n^2/D)$ bound is known for a D -preserver. For a given D , we show algorithms to construct the following:

-- a (1,4) D-spanner of size $\tilde{O}(n^{1.5}/\{D^{0.25}\})$ -- for $D \geq 2$, a (1,4log D) D-spanner of size $\tilde{O}(n.\sqrt{n/D})$.

34. **Yaron Velner.** *The Complexity of Infinitely Repeated Alternating Move Games*

Abstract: We consider infinite duration alternating move games. These games were previously studied by Roth, Balcan, Kalai and Mansour. They presented an FPTAS for computing an approximated equilibrium, and conjectured that there is a polynomial algorithm for finding an exact equilibrium. We extend their study in two directions: (1)~We show that finding an exact equilibrium, even for two-player zero-sum games, is polynomial time equivalent to finding a winning strategy for a (two-player) mean-payoff game on graphs. The existence of a polynomial algorithm for the latter is a long standing open question in computer science. Our hardness result for two-player games suggests that two-player alternating move games are harder to solve than two-player simultaneous move games, while the work of Roth et al., suggests that for $k \geq 3$, k -player games are easier to analyze in the alternating move setting. (2)~We show that optimal equilibriums (with respect to the social welfare metric) can be obtained by pure strategies, and we present an FPTAS for computing a pure approximated equilibrium that is δ -optimal with respect to the social welfare metric. This result extends the previous work by presenting an FPTAS that finds a much more desirable approximated equilibrium. We also show that if there is a polynomial algorithm for mean-payoff games on graphs, then there is a polynomial algorithm that computes an optimal exact equilibrium, and hence, (two-player) mean-payoff games on graphs are inter-reducible with k -player alternating move games, for any $k \geq 2$.

35. **Klaus Jansen and Kim-Manuel Klein.** *A Robust AFPTAS for Online Bin Packing with Polynomial Migration*

Abstract: In this paper we develop general LP and ILP techniques to find an approximate solution with improved objective value close to an existing solution. The task of improving an approximate solution is closely related to a classical theorem of

Cook et al. in the sensitivity analysis for LPs and ILPs. This result is often applied in designing robust algorithms for online problems. We apply our new techniques to the online bin packing problem, where it is allowed to reassign a certain number of items, measured by the migration factor. The migration factor is defined by the total size of reassigned items divided by the size of the arriving item. We obtain a robust asymptotic fully polynomial time approximation scheme (AFPTAS) for the online bin packing problem with migration factor bounded by a polynomial in $\frac{1}{\epsilon}$. This answers an open question stated by Epstein and Levin in the affirmative. As a byproduct we prove an approximate variant of the sensitivity theorem by Cook et al. for linear programs.

36. **Cecilia Bohler, Rolf Klein, Chih-Hung Liu, Evanthia Papadopoulou, Panagiotis Cheilaris and Maksym Zavershynskiy.** *On the Complexity of Higher Order Abstract Voronoi Diagrams*

Abstract: Voronoi diagrams [17,18] are based on bisecting curves enjoying simple combinatorial properties, rather than on the geometric notions of sites and circles. They serve as a unifying concept. Once the bisector system of any concrete type of Voronoi diagram is

shown to fulfill the AVD properties, structural results and efficient algorithms become available without further effort. For example, the first optimal algorithms for constructing nearest Voronoi diagrams of disjoint convex objects, or of line segments under the Hausdorff metric, have been obtained this way [22].

In a concrete order- k Voronoi diagram, all points are placed into the same region that have the same k nearest neighbors among the given sites. This paper is the first to study abstract Voronoi diagrams of arbitrary order k . We prove that their complexity is upper bounded by $2k(n-k)$. So far, an $O(k(n-k))$ bound has been shown only for point sites in the Euclidean and L_p plane [20,21], and, very recently, for line segments [25]. These proofs made extensive use of the geometry of the sites.

Our result on AVDs implies a $2k(n-k)$ upper bound for a wide range of cases for which only trivial upper complexity bounds were previously known, and a slightly sharper bound for the known cases.

Also, our proof shows that the reasons for this bound are combinatorial properties of certain permutation sequences.

37. **Christian Konrad and Adi Rosén.** *Approximating Semi-Matchings in Streaming and in Two-Party Communication*

Abstract: We study the communication complexity and streaming complexity of approximating unweighted semi-matchings. A semi-matching in a bipartite graph $G = (A, B, E)$, with $n = |A|$, is a subset of edges $S \subseteq E$ that matches all A vertices to B vertices with the goal usually being to do this as fairly as possible. While the term `semi-matching` was coined in 2003 by Harvey et al. [WADS 2003], the problem had already previously been studied in the scheduling literature under different names.

We present a deterministic one-pass streaming algorithm that for any $0 \leq \epsilon \leq 1$ uses space $O(n^{1+\epsilon})$ and computes an $O(n^{(1-\epsilon)/2})$ approximation to the semi-matching problem. Furthermore, with $O(\log n)$ passes it is possible to compute an $O(\log n)$ approximation with space $O(n)$.

In the one-way two-party communication setting, we show that for every $\epsilon > 0$, deterministic communication protocols for computing an $O(n^{\frac{1}{1+\epsilon}})$ approximation require a message of size more than $\Omega(n)$ bits. We present two deterministic protocols communicating n and $2n$ edges that compute an $O(\sqrt{n})$ and an $O(n^{1/3})$ approximation respectively.

Finally, we improve on results of Harvey et al. [Journal of Algorithms 2006] and prove new links between semi-matchings and matchings. As a by-product, we obtain a new and purely structural proof of the $\lceil \log(n+1) \rceil$ competitive ratio of the deterministic greedy online semi-matching algorithm of Azar et al. [Journal of Algorithms 1995].

38. **Yuval Filmus, Massimo Lauria, Mladen Miksa, Jakob Nordstrom and Marc Vinyals.** *Towards an Understanding of Polynomial Calculus: New Separations and Lower Bounds (Extended Abstract)*

Abstract: During the last decade, an active line of research in proof complexity has been into the space complexity of proofs and how space is related to other measures. By now these aspects of resolution are fairly well understood, but many open problems remain for the related but stronger polynomial calculus (PC/PCR) proof system. For instance, the space complexity of many standard "benchmark formulas" is still open, as well as the relation of space to size and degree in PC/PCR.

We prove that if a formula requires large resolution width, then making XOR substitution yields a formula requiring large PCR space, providing some circumstantial evidence that degree might be a lower bound for space. More importantly, this immediately yields formulas that are very hard for space but very easy for size, exhibiting a size-space separation similar to what is known for resolution. Using related ideas, we show that if a graph has good expansion and in addition its edge set can be partitioned into short cycles, then the Tseitin formulas over this graph requires large PCR space. In particular, Tseitin formulas over random 4-regular graphs almost surely require space at least $\Omega(\sqrt{n})$.

Our proofs use techniques recently introduced in [Bonacina-Galesi '13]. Our final contribution, however, is to show that these techniques provably cannot yield non-constant space lower bounds for the functional pigeonhole principle, delineating the limitations of this framework and suggesting that we are still far from characterizing PC/PCR space.

39. **Jan Bulánek, Michal Koucký and Michael Saks.** *On Randomized Online Labeling with Polynomially Many Labels*

Abstract: In this paper we prove an optimal lower bound on the complexity of randomized algorithms for the online labeling problem with polynomially many labels. All previous work on this problem (both upper and lower bounds) only applied to deterministic algorithms, so this is the first paper addressing the (im)possibility of faster randomized algorithms. Our lower bound $\Omega(n \log(n))$ matches the complexity of a known deterministic algorithm for this setting of parameters so it is asymptotically optimal.

In the online labeling problem with parameters n and m we are presented with a sequence of n keys from a totally ordered universe U and must assign each arriving key a label from the label set $\{1, 2, \dots, m\}$ so that the order of labels (strictly) respects the ordering on U . As new keys arrive it may be necessary to change the labels of some items; such changes may be done at any time at unit cost for each change. The goal is to minimize the total cost. An alternative formulation of this problem is the file maintenance problem, in which the items, instead of being labeled, are maintained in sorted order in an array of length m , and we pay unit cost for moving an item.

40. **Irit Dinur and Elazar Goldenberg.** *Clustering in the Boolean Hypercube in a List Decoding Regime*

Abstract: We consider the following clustering with outliers problem: Given a set of points $X \subseteq \{-1, 1\}^n$, such that there is some point $z \in \{-1, 1\}^n$ for which $\Pr_{x \in X}[\langle x, z \rangle > \epsilon] > \delta$, find z . We call such a point z a (δ, ϵ) -center of X .

In this work we give lower and upper bounds for the task of finding a (δ, ϵ) -center. Our main upper bound shows that for values of ϵ and δ that are larger than

$1/\text{polylog}(n)$, there exists a polynomial time algorithm that finds a $(\delta - o(1), \epsilon - o(1))$ -center. Moreover, it outputs a list of centers explaining all of the clusters in the input. Our main lower bound shows that given a set for which there exists a (δ, ϵ) -center, it is hard to find even a $(\delta/n^c, \epsilon)$ -center for some constant c and $\epsilon = 1/\text{poly}(n)$, $\delta = 1/\text{poly}(n)$.

</x,z>

41. **Mrinal Kumar, Gaurav Maheshwari and Jayalal Sarma.** *Arithmetic Circuit Lower Bounds via MaxRank*

Abstract: We introduce the polynomial coefficient matrix and identify maximum rank of this matrix under variable substitution as a complexity measure for multivariate polynomials. We use our techniques to prove super-polynomial lower bounds against several classes of non-multilinear arithmetic circuits. In particular, we obtain the following results:

As our main result, we prove that any homogeneous depth-3 circuit for computing the product of d matrices of dimension $n \times n$ requires $\Omega(n^{(d-1)/2^d})$ size. This improves the lower bounds by Nisan and Wigderson(1995) when $d = \omega(1)$.

As our second main result, we show that there is an explicit polynomial on n variables and degree at most $n/2$ for which any depth-3 circuit C of product dimension at most $n/10$ (dimension of the space of affine forms feeding into each product gate) requires size $2^{\Omega(n)}$. This generalizes the lower bounds against diagonal circuits proved by Saxena(2007). Diagonal circuits are of product dimension 1.

We prove a $n^{\Omega(\log n)}$ lower bound on the size of product-sparse formulas. By definition, any multilinear formula is a product-sparse formula. Thus, our result extends the known super-polynomial lower bounds on the size of multilinear formulas by Raz(2006).

We prove a $2^{\Omega(n)}$ lower bound on the size of partitioned arithmetic branching programs. This result extends the known exponential lower bound on the size of ordered arithmetic branching programs given by Jansen(2008).

42. **Marek Cygan and Marcin Pilipczuk.** *Faster exponential-time algorithms in graphs of bounded average degree*

Abstract: We first show that the Traveling Salesman Problem in an n -vertex graph with average degree bounded by d can be solved in $O^*(2^{(1-\epsilon_d)n})$ time and exponential space for a constant ϵ_d depending only on d , where the O^* -notation suppresses factors polynomial in the input size. Thus, we generalize the recent results of Bjorklund et al. [TALG 2012] on graphs of bounded degree.

Then, we move to the problem of counting perfect matchings in a graph. We first present a simple algorithm for counting perfect matchings in an n -vertex graph in $O^*(2^{n/2})$ time and polynomial space; our algorithm matches the complexity bounds of the algorithm of Bjorklund [SODA 2012], but relies on inclusion-exclusion principle instead of algebraic transformations. Building upon this result, we show that the number of perfect matchings in an n -vertex graph with average degree bounded by d can be computed in $O^*(2^{(1-\epsilon_{2d})n/2})$ time and exponential space, where ϵ_{2d} is the constant obtained by us

for the Traveling Salesman Problem in graphs of average degree at most $2d$.

Moreover we obtain a simple algorithm that counts the number of perfect matchings in an n -vertex bipartite graph of average degree at most d in $O^*(2^{\{(1-1/(3.55d))n/2\}})$ time, improving and simplifying the recent result of Izumi and Wadayama [FOCS 2012].

43. **Tobias Brunsch and Heiko Röglin.** *Finding Short Paths on Polytopes by the Shadow Vertex Algorithm*

Abstract: We show that the shadow vertex algorithm can be used to compute a short path between any pair of vertices of a polytope $P = \{x \in \mathbb{R}^n : Ax \leq b\}$ along the edges of P , where $A \in \mathbb{R}^{m \times n}$. Both, the length of the path and the running time of the algorithm, are polynomial in m , n , and a parameter $1/\delta$ that is a measure for the flatness of the vertices of P . For integer matrices $A \in \mathbb{Z}^{m \times n}$ we show a connection between δ and the largest absolute value Δ of any sub-determinant of A , yielding a bound of $O(\Delta^4 m n^4)$ for the length of the computed path. This bound is expressed in the same parameter Δ as the recent non-constructive bound of $O(\Delta^2 n^4 \log(n \Delta))$ by Bonifas et al.

For the special case of totally unimodular matrices, the length of the computed path simplifies to $O(mn^4)$, which significantly improves the previously best known constructive bound of $O(m^{16} n^3 \log^3(mn))$ by Dyer and Frieze.

44. **Christian Glasser, Dung T. Nguyen, Christian Reitwießner, Alan L. Selman and Maximilian Witek.** *Autoreducibility of Complete Sets for Log-Space and Polynomial-Time Reductions*

Abstract: We investigate the autoreducibility and mitoticity of complete sets for several classes with respect to different polynomial-time and logarithmic-space reducibility notions.

Previous work in this area focused on polynomial-time reducibility notions. Here we obtain new mitoticity and autoreducibility results for the classes EXP and NEXP with respect to some restricted truth-table reductions (e.g., polynomial-time 2-tt, ctt, dtt reducibility).

Moreover, we start a systematic study of logarithmic-space autoreducibility and mitoticity which enables us to also consider P and smaller classes. Among others, we obtain the following results: - Regarding log-space many-one, 2-tt, dtt and ctt reducibility, complete sets for PSPACE and EXP are mitotic, and complete sets for NEXP are autoreducible. - All log-space 1-tt-complete sets for NL and P are log-space 2-tt-autoreducible, and all log-space btt-complete sets for NL, P and the Δ -levels of the PH are log-space Turing-autoreducible. - There is a log-space 3-tt-complete set for PSPACE that is not even log-space btt-autoreducible. Using the last result, we conclude that some of our results are hard or even impossible to improve.

45. **Ryan O'Donnell and Li-Yang Tan.** *A composition theorem for the Fourier Entropy-Influence conjecture*

Abstract: The Fourier Entropy-Influence (FEI) conjecture of Friedgut and Kalai seeks to relate two fundamental measures of Boolean function complexity: it states that $H[f] \leq C \cdot$

$\text{Inf}[f]$ holds for every Boolean function f , where $H[f]$ denotes the spectral entropy of f , $\text{Inf}[f]$ is its total influence, and $C > 0$ is a universal constant. Despite significant interest in the conjecture it has only been shown to hold for a few classes of Boolean functions.

Our main result is a composition theorem for the FEI conjecture. We show that if g_1, \dots, g_k are functions over disjoint sets of variables satisfying the conjecture, and if the Fourier transform of F taken with respect to the product distribution with biases $E[g_1], \dots, E[g_k]$ satisfies the conjecture, then their composition $F(g_1(x^1), \dots, g_k(x^k))$ satisfies the conjecture. As an application we show that the FEI conjecture holds for read-once formulas over arbitrary gates of bounded arity, extending a recent result which proved it for read-once decision trees. Our techniques also yield an explicit function with the largest known ratio of $C \geq 6.278$ between $H[f]$ and $\text{Inf}[f]$, improving on the previous lower bound of 4.615.

46. **Endre Boros, Khaled Elbassioni, Vladimir Gurvich and Kazuhisa Makino.** *A Pseudo-Polynomial Algorithm for Mean Payoff Stochastic Games with Perfect Information and a Few Random Positions*

Abstract: We consider two-person zero-sum stochastic mean payoff games with perfect information, or BWR-games, given by a digraph $G = (V = V_B \cup V_W \cup V_R, E)$, with local rewards $r: E \rightarrow \mathbb{R}$, and three types of vertices: black V_B , white V_W , and random V_R . It is a long-standing open question whether a polynomial time algorithm for BWR-games exists, or not. In fact, a pseudo-polynomial algorithm for these games would already imply their polynomial solvability. In this paper, we show that BWR-games with a constant number of random nodes can be solved in pseudo-polynomial time. That is, for any such game with a few random nodes $|V_R| = O(1)$, a saddle point in pure stationary strategies can be found in time polynomial in $|V_W| + |V_B|$, the maximum absolute local reward R , and the common denominator of the transition probabilities.

47. **Reinhard Bauer, Tobias Columbus, Ignaz Rutter and Dorothea Wagner.** *Search-Space Size in Contraction Hierarchies*

Abstract: Contraction hierarchies are a speed-up technique to improve the performance of shortest-path computations, which works very well in practice. Despite convincing practical results, there is still a lack of theoretical explanation for this behavior.

In this paper, we develop a theoretical framework for studying search space sizes in contraction hierarchies. We prove the first bounds on the size of search spaces that depend solely on structural parameters of the input graph, that is, they are independent of the edge lengths. To achieve this, we establish a connection with the well-studied elimination game. Our bounds apply to graphs with treewidth $\sim k$, and to any minor-closed class of graphs that admits small separators. For trees, we show that the maximum search space size can be minimized efficiently, and the average size can be approximated efficiently within a factor of ~ 2 .

We show that, under a worst-case assumption on the edge lengths, our bounds are comparable to the recent results of Abraham et al. (ICALP'11), whose analysis depends also on the edge lengths. As a side result, we link their notion of highway dimension (a parameter that is conjectured to be small, but is unknown for all practical instances) with the notion of

pathwidth. This is the first relation of highway dimension with a well-known graph parameter.

48. **Chandra Chekuri, Guylain Naves and Bruce Shepherd.** *Maximum Edge-Disjoint Paths in k -Sums of Graphs*

Abstract: We consider the approximability of the maximum edge-disjoint paths problem (MEDP) in undirected graphs, and in particular, the integrality gap of the natural multicommodity flow based relaxation for it. The integrality gap is known to be $\Omega(\sqrt{n})$ even for planar graphs \cite{GargVY97} due to a simple topological obstruction and a major focus, following earlier work \cite{KleinbergT98}, has been understanding the gap if some constant congestion is allowed. In planar graphs the integrality gap is $O(1)$ with congestion 2 \cite{SeguinS11,CKS-planar-constant}. In general graphs, recent work has shown the gap to be $\text{polylog}(n)$ \cite{Chuzhoy12,ChuzhoyL12} with congestion 2 . Moreover, the gap is $\log^{\Omega(c)} n$ in general graphs with congestion c for any constant $c \geq 1$ \cite{andrews2010inapproximability}.

It is natural to ask for which classes of graphs does a constant-factor constant-congestion property hold. It is easy to deduce that for given constant bounds on the approximation and congestion, the class of "nice" graphs is minor-closed. Is the converse true? Does every proper minor-closed family of graphs exhibit a constant factor, constant congestion bound relative to the LP relaxation? We conjecture that the answer is yes. One stumbling block has been that such bounds were not known for bounded treewidth graphs (or even treewidth 3). In this paper we give a polytime algorithm which takes a fractional routing solution in a graph of bounded treewidth and is able to integrally route a constant fraction of the LP solution's value. Note that we do not incur any edge congestion. Previously this was not known even for series parallel graphs which have treewidth 2 . The algorithm is based on a more general argument that applies to k -sums of graphs in some graph family, as long as the graph family has a constant factor, constant congestion bound. We then use this to show that such bounds hold for the class of k -sums of bounded genus graphs.

49. **Karl Bringmann, Benjamin Doerr, Adrian Neumann and Jakub Sliacan.** *Online Checkpointing with Improved Worst-Case Guarantees*

Abstract: In the online checkpointing problem, the task is to continuously maintain a set of k checkpoints that allow to rewind an ongoing computation faster than by a full restart. The only operation allowed is to remove an old checkpoint and to store the current state instead. Our aim are checkpoint placement strategies that minimize rewinding cost, i.e., such that at all times T when requested to rewind to some time $t < T$ improving over earlier work showing $1 + 1/k \leq p_k \leq 2$, we show that p_k can be chosen less than 2 uniformly for all k . More precisely, we show the uniform bound $p_k \leq 1.7$ for all k , and present algorithms with asymptotic performance $p_k \leq 1.59 + o(1)$ valid for all k and $p_k \leq \ln(4) + o(1) \leq 1.39 + o(1)$ valid for k being a power of two. For small values of k , we show how to use a linear programming approach to compute good checkpointing algorithms. This gives performances of less than 1.53 for $k \leq 10$.

On the more theoretical side, we show the first lower bound that is asymptotically more than one, namely $p_k \geq 1.30 - o(1)$. We also show that optimal algorithms (yielding the infimum performance) exist for all k .

</t>

50. **Gregory Kucherov and Yakov Nekrich.** *Full-fledged Real-Time Indexing for Constant Size Alphabets*

Abstract: In this paper we describe a data structure that supports pattern matching queries on a dynamically arriving text over an alphabet of constant size. Each new symbol can be prepended to T in $O(1)$ worst-case time. At any moment, we can report all occurrences of a pattern P in the current text in $O(|P|+k)$ time, where $|P|$ is the length of P and k is the number of occurrences. This resolves, under assumption of constant-size alphabet, a long-standing open problem of existence of a real-time indexing method for string matching (see \cite{AmirN08}).

51. **Dimitris Fotakis and Christos Tzamos.** *On the Power of Deterministic Mechanisms for Facility Location Games*

Abstract: We investigate the approximability of K -Facility Location by deterministic strategyproof mechanisms. Our main result is an elegant characterization of deterministic strategyproof mechanisms with a bounded approximation ratio for 2-Facility Location on the line. Specifically, we show that for instances with $n \geq 5$ agents, any such mechanism either admits a unique dictator, or always places the facilities at the two extremes. As a consequence, we obtain that the best approximation ratio achievable by deterministic strategyproof mechanisms for 2-Facility Location on the line is precisely $n-2$. Employing a technical tool developed for the characterization, we show that for every $K \geq 3$, there do not exist any deterministic anonymous strategyproof mechanisms with a bounded approximation ratio for K -Facility Location on the line, even for simple instances with $K+1$ agents. Moreover, building on the characterization for the line, we show that there do not exist any deterministic mechanisms with a bounded approximation ratio for 2-Facility Location in more general metric spaces, which is true even for simple instances with 3 agents located in a star.

52. **Hannes Uppman.** *The Complexity of Three-element Min-Sol and Conservative Min-Cost-Hom*

Abstract: Thapper and Zivny [STOC'13] recently classified the complexity of VCSP for all finite-valued constraint languages. However, the complexity of VCSPs for constraint languages that are not finite-valued remains poorly understood. In this paper we study the complexity of two such VCSPs, namely Min-Cost-Hom and Min-Sol. We obtain a full classification for the complexity of Min-Sol on domains that contain at most three elements and for the complexity of conservative Min-Cost-Hom on arbitrary finite domains. Our results answer a question raised by Takhanov [STACS'10, COCOON'10].

53. **Per Austrin, Petteri Kaski, Mikko Koivisto and Jussi Määttä.** *Space-Time Tradeoffs for Subset Sum: An Improved Worst Case Algorithm*

Abstract: The technique of Schroepel and Shamir (SICOMP, 1981) has long been the most efficient way to trade space against time for the Subset Sum problem. In the random-instance setting, however, improved tradeoffs exist. In particular the recently discovered dissection method of Dinur et al. (CRYPTO 2012) yields a significantly improved space--time tradeoff

curve for instances with strong randomness properties. Our main result is that these strong randomness assumptions can be removed, obtaining the same space--time tradeoffs in the worst case. Our strategy for dealing with arbitrary instances is, essentially, to instead inject the required randomness into the dissection process itself by working over a carefully selected but random composite modulus, and to introduce explicit space--time controls into the algorithm by means of a "bailout mechanism".

54. **Marcin Bieńkowski, Jarek Byrka, Marek Chrobak, Neil Dobbs, Tomasz Nowicki, Maxim Sviridenko, Grzegorz Świrszcz and Neal Young.** *Approximation Algorithms for the Joint Replenishment Problem with Deadlines*

Abstract: The Joint Replenishment Problem (JRP) is a fundamental optimization problem in supply-chain management, concerned with optimizing the flow of goods over time from a supplier to retailers. Over time, in response to demands at the retailers, the supplier sends shipments, via a warehouse, to the retailers. The objective is to schedule shipments to minimize the sum of shipping costs and retailers' waiting costs.

We study the approximability of JRP with deadlines, where instead of waiting costs the retailers impose strict deadlines. We study the integrality gap of the standard linear-program (LP) relaxation, giving a lower bound of 1.207, and an upper bound and approximation ratio of 1.574. The best previous upper bound and approximation ratio was 1.667; no lower bound was previously published. For the special case of equal-length demand periods, we give an upper bound of 1.5, a lower bound of 1.2, and show APX-hardness.

55. **Dániel Marx and László A. Végh.** *Fixed-parameter algorithms for minimum cost edge-connectivity augmentation*

Abstract: We consider connectivity-augmentation problems in a setting where each potential new edge has a nonnegative cost associated with it, and the task is to achieve a certain connectivity with at most p new edges of minimum total cost. The main result is that the minimum cost augmentation of edge-connectivity from $k-1$ to k with at most p new edges is fixed-parameter tractable parameterized by p and admits a polynomial kernel. We also prove the fixed-parameter tractability of increasing edge-connectivity from 0 to 2, and increasing node-connectivity from 1 to 2.

56. **Martin Aumüller and Martin Dietzfelbinger.** *Optimal Partitioning for Dual Pivot Quicksort*

Abstract: Dual pivot quicksort refers to variants of classical quicksort where in the partitioning step two pivots are used to split the input into three segments. This can be done in different ways, giving rise to different algorithms. Recently, a dual pivot algorithm due to Yaroslavskiy received much attention, because it replaced the well-engineered quicksort algorithm in Oracle's Java 7 runtime library. Nebel and Wild (ESA 2012, best paper) analyzed this algorithm and showed that on average it uses $1.9n \ln n + O(n)$ comparisons to sort an input of size n , beating standard quicksort, which uses $2n \ln n + O(n)$. We introduce a model that captures all dual pivot algorithms, give a unified analysis, and identify a new dual pivot algorithm that minimizes the expected number of key comparisons among all possible algorithms. This minimum is $1.8n \ln n + o(n \ln n)$. If pivots are chosen from a sample of size s , it is $1.623n \ln n + o(n \ln n)$.

57. **Alexandr Andoni, Huy Nguyen, Yury Polyanskiy and Yihong Wu.** *Tight Lower Bound for Linear Sketches of Moments*

Abstract: The problem of estimating frequency moments of a data stream has attracted a lot of attention since the onset of streaming algorithms [Alon-Matias-Szegedy '99]. While the space complexity for computing the p^{th} moment, for $p \in (0, 2]$ has been settled [Kane-Nelson-Woodruff '10], for $p > 2$ the exact complexity remains open. For $p > 2$ the current best algorithm uses $O(n^{1-2/p} \log n)$ words of space [Andoni-Krauthgamer-Onak '11, Braverman-Ostrovsky '10], whereas the lower bound is of $\Omega(n^{1-2/p})$ [Bar-Yossef-Jayram-Kumar-Sivakumar '04].

In this paper, we show a tight lower bound of $\Omega(n^{1-2/p} \log n)$ for the class of algorithms based on linear sketches, which store only a sketch Ax of input vector x and some (possibly randomized) matrix A . We note that all known algorithms for this problem are linear sketches.

58. **Xin Li, Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, Manoj Prabhakaran, Amit Sahai and David Zuckerman.** *Robust Pseudorandom Generators*

Abstract: Let $G: \{0, 1\}^n \rightarrow \{0, 1\}^m$ be a pseudorandom generator. We say that a circuit implementation of G is (k, q) -robust if for every set S of at most k wires anywhere in the circuit, there is a set T of at most $q|S|$ outputs, such that conditioned on the values of S and T the remaining outputs are pseudorandom. We initiate the study of robust PRGs, obtaining explicit and non-explicit constructions in which k is close to n , q is constant, and $m \gg n$. These include unconditional constructions of robust r -wise independent PRGs and small-bias PRGs, as well as conditional constructions of robust cryptographic PRGs.

In addition to their general usefulness as a more resilient form of PRGs, our study of robust PRGs is motivated by cryptographic applications in which an adversary has a local view of a large source of secret randomness. We apply robust r -wise independent PRGs towards reducing the randomness complexity of private circuits and protocols for secure multiparty computation, as well as improving the "black-box complexity" of constant-round secure two-party computation.

59. **Markus Blaeser.** *Noncommutativity makes determinants hard*

Abstract: We consider the complexity of computing the determinant over arbitrary finite-dimensional algebras. We first consider the case that A is fixed. We obtain the following dichotomy: If $A/\text{rad } A$ is noncommutative, then computing the determinant over A is hard. "Hard" here means $\#P$ -hard over fields of characteristic 0 and $\text{Mod } P$ -hard over fields of characteristic $p > 0$. If $A/\text{rad } A$ is commutative and the underlying field is perfect, then we can compute the determinant over A in polynomial time.

We also consider the case when A is part of the input. Here the hardness is closely related to the nilpotency index of the commutator ideal of A . The commutator ideal $\text{com}(A)$ of A is the ideal generated by all elements of the form $xy - yx$ with $x, y \in A$. We prove that if the nilpotency index of $\text{com}(A)$ is linear in n , where $n \times n$ is the format of the given matrix, then computing the determinant is hard. On the other hand, we show the following upper bound: Assume that there is an algebra $B \subseteq A$ with $B = A/\text{rad } A$. (If the underlying field is

perfect, then this is always true.) The center $Z(A)$ of A is the set of all elements that commute with all other elements. It is a commutative subalgebra. We call an ideal J a complete ideal of noncommuting elements if $B + Z(A) + J = A$. If there is such a J with nilpotency index $o(n/\log n)$, then we can compute the determinant in subexponential time. Therefore, the determinant cannot be hard in this case, assuming the counting version of the exponential time hypothesis.

Our results answer several open questions posed by Chien et al.

60. **T-H. Hubert Chan, Mingfei Li, Li Ning and Shay Solomon.** *New Doubling Spanners: Better and Simpler*

Abstract: In a seminal STOC'95 paper, Arya et al. conjectured that spanners for low-dimensional Euclidean spaces with constant maximum degree, hop-diameter $O(\log n)$ and lightness $O(\log n)$ (i.e., weight $O(\log n) w(\text{MST})$) can be constructed in $O(n \log n)$ time. This conjecture, which became a central open question in this area, was resolved in the affirmative by Elkin and Solomon in STOC'13. In fact, Elkin and Solomon proved that the conjecture of Arya et al. holds even in doubling metrics. However, Elkin and Solomon's spanner construction is complicated. In this work we present a simpler construction of spanners for doubling metrics with the above guarantees. Moreover, our construction extends in a simple and natural way to provide k -fault tolerant spanners with maximum degree $O(k^2)$, hop-diameter $O(\log n)$ and lightness $O(k^2 \log n)$.

61. **Arnab Bhattacharyya and Yuichi Yoshida.** *An Algebraic Characterization of Testable CSP's*

Abstract: Given an instance I of a CSP, a tester for I distinguishes assignments satisfying I from those which are far from any assignment satisfying I . The efficiency of a tester is measured by its query complexity, the number of variable assignments queried by the algorithm. In this paper, we characterize the hardness of testing Boolean CSPs in terms of the algebra generated by the relations used to form constraints. In terms of computational complexity, we show that if a non-trivial Boolean CSP is sublinear-query testable (resp., not sublinear-query testable), then the CSP is in NL (resp., P-complete, parityL-complete or NP-complete) and that if a sublinear-query testable Boolean CSP is constant-query testable (resp., not constant-query testable), then counting the number of solutions of the CSP is in P (resp., #P-complete).

Also, we conjecture that a CSP instance is testable in sublinear time if its Gaifman graph has bounded treewidth. We confirm the conjecture when a near-unanimity operation is a polymorphism of the CSP.

62. **Amir Abboud and Kevin Lewi.** *Exact Weight Subgraphs and the k -Sum Conjecture*

Abstract: We consider the Exact-Weight- H problem of finding a (not necessarily induced) subgraph H of weight 0 in an edge-weighted graph G . We show that for every H , the complexity of this problem is strongly related to that of the infamous k -Sum problem. In particular, we show that under the k -Sum Conjecture, we can achieve tight upper and lower bounds for the Exact-Weight- H problem for various subgraphs H such as matching, star, path, and cycle.

One interesting consequence is that improving on the $O(n^3)$ upper bound for Exact-Weight-4-Path or Exact-Weight-5-Path will imply improved algorithms for 3-Sum, 5-Sum, All-Pairs Shortest Paths and other fundamental problems. This is in sharp contrast to the minimum-weight and (unweighted) detection versions, which can be solved easily in time $O(n^2)$. We also show that a faster algorithm for any of the following three problems would yield faster algorithms for the others: 3-Sum, Exact-Weight-3-Matching, and Exact-Weight-3-Star.

63. **Andrew Goldberg, Maxim Babenko, Anupam Gupta and Viswanath Nagarajan.** *Algorithms for Hub Label Optimization*

Abstract: Cohen et al. developed an $O(\log n)$ -approximation algorithm for minimizing the total hub label size (L1 norm). We give $O(\log n)$ -approximation algorithms for related problems, including minimizing the maximum label (L-inf norm) and minimizing L_p and L_q norms simultaneously.

64. **Piotr Indyk and Ilya Razenshteyn.** *On Model-Based RIP-1 Matrices*

Abstract: The Restricted Isometry Property (RIP) is a fundamental property of a matrix enabling sparse recovery~\cite{CRT06}. Informally, an $m \times n$ matrix satisfies RIP of order k in the ℓ_p norm if $\|Ax\|_p \approx \|x\|_p$ for any vector x that is k -sparse, i.e., that has at most k non-zeros. The minimal number of rows m necessary for the property to hold has been extensively investigated, and tight bounds are known.

Motivated by signal processing models, a recent work of Baraniuk et al~\cite{BCDH10} has generalized this notion to the case where the support of x must belong to a given $\{\text{em model}\}$, i.e., a given family of supports. This more general notion is much less understood, especially for norms other than ℓ_2 .

In this paper we present tight bounds for the model-based RIP property in the ℓ_1 norm. Our bounds hold for the two most frequently investigated models: tree-sparsity and block-sparsity. We also show implications of our results to sparse recovery problems.

65. **Daniel Grier.** *Deciding the Winner of an Arbitrary Finite Poset Game is PSPACE-Complete*

Abstract: A poset game is a two-player game played over a partially ordered set (poset) in which the players alternate choosing an element of the poset, removing it and all elements greater than it. The first player unable to select an element of the poset loses. Polynomial time algorithms exist for certain restricted classes of poset games, such as the game of Nim. However, until recently the complexity of arbitrary finite poset games was only known to exist somewhere between NC1 and PSPACE. We resolve this discrepancy by showing that deciding the winner of an arbitrary finite poset game is PSPACE-complete. To this end, we give an explicit reduction from Node Kayles, a PSPACE-complete game in which players vie to choose an independent set in a graph.

66. **Mohammadhossein Bateni, Mohammadtaghi Hajiaghayi and Vahid Liaghat.** *Improved Approximation Algorithms for (Budgeted) Node-weighted Steiner Problems*

Abstract: Moss and Rabani (STOC 2001, SICOMP 2007) study constrained node-weighted

Steiner tree problems with two independent weight values associated with each node, namely, cost and prize (or penalty). They give an $O(\log n)$ -approximation algorithm for the prize-collecting node-weighted Steiner tree problem (PCST)---where the goal is to minimize the cost of a tree plus the penalty of vertices not covered by the tree. They use the algorithm for PCST to obtain a bicriteria $(2, O(\log n))$ -approximation algorithm for the Budgeted node-weighted Steiner tree problem --- where the goal is to maximize the prize of a tree with a given budget for its cost. Their solution may cost up to twice the specified budget, but collects a factor $\Omega(\frac{1}{\log n})$ of the optimal prize. We improve the results of Moss and Rabani from at least two aspects.

Our first main result is a primal-dual $O(\log h)$ -approximation algorithm for a more general problem, prize-collecting node-weighted Steiner forest (PCSF), where we have h demands each requesting the connectivity of a pair of vertices. Our approximation guarantee is tight within constant factors. Not only is this the first nontrivial approximation algorithm for the classical (node-weighted) PCSF problem, but also our algorithm as well as its analysis is much simpler than that of Moss and Rabani~\cite{MossRabani}. Incidentally, K\"onemann et al.~\cite{Sina13} have shown very recently that the prize-collecting algorithm in \cite{MossRabani} is flawed.

Our second main contribution is for the Budgeted node-weighted Steiner tree problem, which is also an improvement to Moss and Rabani~\cite{MossRabani} and Guha et al.~\cite{Guha99} (STOC'99). In the unrooted case, we improve upon an $O(\log^2 n)$ -approximation of \cite{Guha99}, and present an $O(\log n)$ -approximation algorithm without any budget violation. For the rooted case, where a specified vertex has to appear in the solution tree, we improve the bicriteria result of \cite{MossRabani} to a bicriteria approximation ratio of $(1+\epsilon, O(\log n)/\epsilon^2)$ for any positive (possibly subconstant) ϵ . That is, for any permissible budget violation $1+\epsilon$, we present an algorithm achieving a tradeoff in the guarantee for prize. Indeed, we show that this is almost tight for the natural linear-program relaxation used by us as well as in \cite{MossRabani}.

67. **Omri Weinstein, Mark Braverman, Amir Yehudayoff and Anup Rao.** *Direct product via round-preserving compression*

Abstract: We obtain a strong direct product theorem for two-party bounded round communication complexity. Let $\text{succ}_r(\mu, f, C)$ denote the maximum success probability of an r -round communication protocol that uses at most C bits of communication to compute $f(x, y)$ when $(x, y) \sim \mu$. Jain et al. \cite{JainPP12} have recently showed that if $\text{succ}_r(\mu, f, C) \leq \frac{2}{3}$ and $T \ll (C - \Omega(r^2)) \cdot \frac{1}{r}$, then $\text{succ}_r(\mu^n, f^n, T) \leq \exp(-\Omega(n/r^2))$. Here we prove that if $\text{succ}_{7r}(\mu, f, C) \leq \frac{2}{3}$ and $T \ll (C - \Omega(r \log r)) \cdot n$ then $\text{succ}_r(\mu^n, f^n, T) \leq \exp(-\Omega(n))$. Our result asymptotically matches the upper bound on $\text{succ}_{7r}(\mu^n, f^n, T)$ given by the trivial solution which applies the per-copy optimal protocol independently to each coordinate. The proof relies on a compression scheme that improves the tradeoff between the number of rounds and the communication complexity over known compression schemes.

68. **Aleksandrs Belovs, Andrew M. Childs, Stacey Jeffery, Robin Kothari and Frederic Magniez.** *Time-efficient quantum walks for 3-distinctness*

Abstract: We present an extension to the quantum walk search framework that facilitates quantum walks with nested updates. We apply it to give a quantum walk algorithm for S -Distinctness with query complexity $\tilde{O}(n^{5/7})$, matching the best known upper bound (obtained via learning graphs) up to log factors. Furthermore, our algorithm has time complexity $\tilde{O}(n^{5/7})$, improving the previous $\tilde{O}(n^{3/4})$.

69. **Matthew Patitz, Scott Summers, Damien Woods, Erik Demaine, Robert Schweller and Trent Rogers.** *The two-handed tile assembly model is not intrinsically universal*

Abstract: In this paper, we study the intrinsic universality of the well-studied Two-Handed Tile Assembly Model (2HAM), in which two "supertile" assemblies, each consisting of one or more unit-square tiles, can fuse together (self-assemble) whenever their total attachment strength is at least the global temperature τ . Our main result is that every temperature- $(\tau - 1)$ 2HAM tile system cannot simulate at least one temperature- τ 2HAM tile system. This impossibility result proves that the 2HAM is not intrinsically universal, in stark contrast to the simpler abstract Tile Assembly Model which was shown to be intrinsically universal (The tile assembly model is intrinsically universal, FOCS 2012). On the positive side, we prove that, for every fixed temperature $\tau \geq 2$, temperature- τ 2HAM tile systems are intrinsically universal: for each τ there is a single universal 2HAM tile set S that, when appropriately initialized, is capable of simulating the behavior of any temperature- τ 2HAM tile system. As a corollary of these results we find an infinite set of infinite hierarchies of 2HAM systems with strictly increasing power within each hierarchy. Finally, we show how to construct, for each τ , a temperature- τ 2HAM system that simultaneously simulates all temperature- τ 2HAM systems.

70. **Ozgur Ozkan, John Iacono, Stefan Langerman and Erik D. Demaine.** *Combining Binary Search Trees*

Abstract: We present a general transformation for combining a constant number of binary search tree data structures (BSTs) into a single BST whose running time is within a constant factor of the minimum of any "well-behaved" bound on the running time of the given BSTs, for any online access sequence. (A BST has a well-behaved bound with $f(n)$ overhead if it spends at most $\text{bigoh}\{f(n)\}$ time per access and its bound satisfies a weak sense of closure under subsequences.)

In particular, we obtain a BST data structure that is $\text{bigoh}\{\log \log n\}$ competitive, satisfies the working set bound (and thus satisfies the static finger bound and the static optimality bound), satisfies the dynamic finger bound, satisfies the unified bound with an additive $\text{bigoh}\{\log \log n\}$ factor, and performs each access in worst-case $\text{bigoh}\{\log n\}$ time.

Our result also raises the bar on necessary conditions for a counterexample to the dynamic optimality conjecture. Along the way, we develop a transformation for simulating a BST with multiple fingers using only one finger, which may be of independent interest. Using this transformation, we show how to combine the augmented data fields in each node in the tree to simulate a doubly linked list, where the next and previous nodes can be accessed in $\text{bigoh}\{1\}$ amortized time with respect to all tree operations. This can be viewed as augmenting the tree itself with a buffer storing global data about the tree, in contrast to augmenting each node with local data.

71. **Philip Bille, Johannes Fischer, Inge Li Gørtz, Tsvi Kopelowitz, Benjamin Sach and Hjalte Wedel Vildhøj.** *Sparse Suffix Tree Construction in Small Space*

Abstract: We consider the problem of constructing a sparse suffix array (or suffix array) for b suffixes of a given text T of size n , using only $O(b)$ words of space during construction time. Breaking the naive bound of $O(nb)$ time for this problem that can be traced back to the origins of string indexing in 1968. First results were only obtained in 1996, but only for the case where the suffixes were evenly spaced in T , here there is no constraint on the locations of the suffixes.

We show that the sparse suffix tree can be constructed in $O(n \log(b)^2)$ time. To achieve this we develop a technique, which may be of independent interest, that allows to efficiently answer b longest common prefix queries on suffixes of T , using only $O(b)$ space. We expect that this technique will prove useful in many other applications in which space usage is a concern.