IP1
Distributed Algorithms in Ant Colonies

Like many distributed systems, both natural and engineered, ant colonies operate without any central control. No ant can assess what needs to be done. Each ant responds to its interactions with other ants nearby and in the aggregate, these stochastic, dynamical networks of interaction regulate colony behavior. Ant species use distributed algorithms that differ according to environmental constraints such as operating costs, time pressure, patchy distribution of resources in space, risk of rupture, and the intensity of competition. I will discuss algorithms used to regulate foraging activity, nestmate recognition and colony security, collective search, and stable trail systems.

Deborah M. Gordon
Stanford University
dmgordon@stanford.edu

IP2
Beyond Planarity: On Geometric Intersection Graphs

Many efficient algorithms had been developed over the years for planar graphs and more general graphs such as low genus graphs. Intersection graphs of geometric objects (in low dimensions) with some additional proprieties, such as fatness or low density, provide yet another family of graphs for which one can design better algorithms. This family is a vast extension of planar graphs, and yet is still algorithmically tractable for many problems. In this talk, we will survey this class of graphs, and some algorithms and intractability results known for such graphs, and outline open problems for further research.

Sariel Har-Peled
Department of Computer Science
University of Illinios, Urbana-Champaign
sariel@uiuc.edu

IP3
Combinatorial Statistics

Combinatorial Statistics studies combinatorial algorithms for inference of statistical models that are parametrized by discrete parameters such as trees, graphs or permutations. The talk will outline some key examples of results and techniques that are used in this area as well as future research directions and open problems.

Elchanan Mossel
University of Pennsylvania
mossel@wharton.upenn.edu

CP0
A Novel Dual Ascent Algorithm for Solving the Min-Cost Flow Problem

We present a novel algorithm for the min-cost flow problem that is competitive with recent third-party implementations of well-known algorithms for this problem and even outperforms them on certain realistic instances. We formally prove correctness of our algorithm and show that the worst-case running time is in $O(||b||_1(m + n \log n))$ where b is the vector of demands. Combined with standard scaling techniques, this pseudo-polynomial bound can be made polynomial in a straightforward way. Furthermore, we evaluate our approach experimentally. Our empirical findings indeed suggest that the running time does not significantly depend on the costs and that a linear dependence on $||b||_1$ is overly pessimistic.

Ruben Becker, Maximilian Fickert, Andreas Karrenbauer
Max Planck Institute for Informatics
ruben.becker@mpi-inf.mpg.de, maximilian.fickert@mpi-inf.mpg.de, andreas.karrenbauer@mpi-inf.mpg.de

CP0

Computing Top-k Closeness Centrality Faster in Unweighted Graphs

Centrality indices are widely used analytic measures for the importance of nodes in a network. Closeness centrality is very popular among these measures. For a single node v, it takes the sum of the distances of v to all other nodes into account. The currently best algorithms in practical applications for computing the closeness for all nodes exactly in unweighted graphs are based on breadth-first search (BFS) from every node. Thus, even for sparse graphs, these algorithms require quadratic running time in the worst case, which is prohibitive for large networks. In many relevant applications, however, it is unnecessary to compute closeness values for all nodes. Instead, one requires only the k nodes with the highest closeness values in descending order. Thus, we present a new algorithm for computing this top-k ranking in unweighted graphs. Following the rationale of previous work, our algorithm significantly reduces the number of traversed edges. It does so by computing upper bounds on the closeness and stopping the current BFS search when k nodes already have higher closeness than the bounds computed for the other nodes. In our experiments with real-world and synthetic instances of various types, one of these new bounds is good for small-world graphs with low diameter (such as social networks), while the other one excels for graphs with high diameter (such as road networks). Combining them yields an algorithm that is faster than the state of the art for top-k computations for all test instances, by a wide margin for high-diameter graphs. Finally, we prove that the quadratic worst-case complexity cannot be improved on directed, disconnected graphs, under reasonable complexity assumptions.

Elisabetta Bergamini
Institute of Theoretical Informatics
Karlsruhe Institute of Technology (KIT)
elisabetta.bergamini@kit.edu

Michele Borassi
IMT Institute for Advanced Studies Lucca
michele.borassi@imtlucca.it

Pierluigi Crescenzi
Dipartimento di Ingegneria dellInformazione
Universita‘ di Firenze
pierluigi.crescenzi@unifi.it

Andrea Marino
Dipartimento di Informatica
Universita‘ di Pisa
marino@di.unipi.it

Henning Meyerhenke
Karlsruhe Institute of Technology
Institute for Theoretical Informatics, Parallel Computing

meyerhenke@kit.edu

CP0

Sampling Weighted Perfect Matchings on the Square-Octagon Lattice

We consider perfect matchings of the square-octagon lattice, also known as "fortresses'. There is a natural local Markov chain on the set of perfect matchings that is known to be ergodic. However, unlike seemingly related Markov chains on perfect matchings on the square and hexagonal lattices, corresponding to domino and lozenge tilings, respectively, the Markov chain on the square-octagon lattice appears to converge slowly. To understand why, we consider a weighted version of the problem. As with domino and lozenge tilings, it will be useful to view perfect matchings on the square-octagon lattice in terms of sets of paths and cycles on a corresponding lattice region; here, the paths and cycles lie on the Cartesian lattice and are required to turn left or right at every step. For input parameters λ and μ, we define the weight of a configuration to be $\lambda^{E(\sigma)}\mu^{V(\sigma)}$, where $E(\sigma)$ is the total number of edges on the paths and cycles of σ and $V(\sigma)$ is the number of vertices that are not incident to any of the paths or cycles in σ. Weighted paths already come up in the reduction from perfect matchings to turning lattice paths, corresponding to the case when $\lambda = 1$ and $\mu = 2$. First fixing $\mu = 1$, we show that there are choices of λ for which the chain converges slowly and another for which it is fast, suggesting a phase change in the mixing time. More precisely, the chain requires exponential time (in the size of the lattice region) when $\lambda < 1/(2\sqrt{e})$ or $\lambda > 2\sqrt{e}$, while it is polynomially mixing at $\lambda = 1$. When $\mu > 1$, we show that the Markov chain is slowly mixing if $\lambda < \sqrt{\mu}/(2\sqrt{e})$ or $\lambda > 2\mu\sqrt{e}$. These are the first rigorous proofs explaining why the natural local Markov chain can be slow for the weighted fortress model, or perfect matchings on the square-octagon lattice.

Prateek Bhakta, Dana Randall
Georgia Institute of Technology
pbhakta@gatech.edu, randall@cc.gatech.edu

CP0

Fast Algorithms for Pseudoarboricity

We show that the pseudoarboricity p of a graph can be computed in $\mathcal{O}(|E|^{3/2}\sqrt{\log\log p})$ time, and give better runtime estimates for certain asymptotic bounds on p. These results are achieved by accelerating a binary search with an approximation scheme, and a runtime analysis of Dinitz's algorithm on 'almost unit capacity' networks. We experimentally compare several algorithms for the problem.

Markus Blumenstock
Institute of Computer Science
Johannes Gutenberg University
mablumen@uni-mainz.de

CP0

Graphs with Degree Constraints

Given a set D of nonnegative integers, we derive the asymptotic number of graphs with a given number of vertices, edges, and such that the degree of every vertex is in D. This generalizes existing results, such as the enumeration of graphs with a given minimum degree, and establishes new ones, such as the enumeration of Euler graphs, *i.e.* where all vertices have an even degree. Those results are derived using analytic combinatorics.

Elie De Panafieu
Sorbonne Universités, UPMC Univ Paris 06, CNRS, LIP6 UMR, 4 place Jussieu 75005 Paris, France
depanafieuelie@gmail.com

Lander Ramos
Universitat Politècnica de Catalunya, Omega 412, Jordi Girona 1-3, 08034 Barcelona
lander.ramos@upc.edu

CP0

An Empirical Comparison of Graph Laplacian Solvers

Solving Laplacian linear systems is an important task in a variety of practical and theoretical applications. This problem is known to have solutions that perform in linear times polylogarithmic work in theory, but these algorithms are difficult to implement in practice. We examine existing solution techniques in order to determine the best methods currently available and for which types of problems they are useful.

Kevin Deweese
UC Santa Barbara
kdeweese@cs.ucsb.edu

John R. Gilbert
Dept of Computer Science
University of California, Santa Barbara
gilbert@cs.ucsb.edu

Erik G. Boman
Center for Computing Research
Sandia National Labs
egboman@sandia.gov

CP0

Betweenness Centrality in Random Trees

The betweenness centrality of a node is a quantity that is frequently used in network theory to measure how "central' a node v is. It is defined as the summed proportion of shortest paths between pairs of vertices that pass through v. We study the average and limiting distribution of the betweenness centrality in random trees and, related, subcritical graph families.

Kevin Durant
Stellenbosch University
kdurant@outlook.com

Stephan Wagner
Stellenbosch University, South Africa
swagner@sun.ac.za

CP0

Experimental Evaluation of Distributed Node Coloring Algorithms for Wireless Networks

In this paper we evaluate distributed node coloring algorithms for wireless networks using the network simulator Sinalgo. All considered algorithms operate in the realistic signal-to-interference-and-noise-ratio (SINR) model of interference. We evaluate two recent coloring algorithms,

RAND4DCOLOR and COLORRED, proposed by Fuchs and Prutkin in [SIROCCO'15], the MW-Coloring algorithm introduced by Moscibroda and Wattenhofer [SPAA'15] and transferred to the SINR model by Derbel and Talbi [ICDCS'10], and a variant of the coloring algorithm of Yu et al. [ALGOSENSORS'10]. We additionally consider several practical improvements to the algorithms and evaluate their performance in both static and dynamic scenarios. Our experiments show that RAND4DCOLOR is very fast, computing a valid (4Δ)-coloring in less than one third of the time slots required for local broadcasting, where Δ is the maximum node degree in the network. Regarding other $\mathcal{O}(\Delta)$-coloring algorithms RAND4DCOLOR is at least 4 to 5 times faster. Additionally, the algorithm is robust even in networks with mobile nodes and an additional listening phase at the start of the algorithm makes RAND4DCOLOR robust against the late wake-up of large parts of the network. Regarding $(\Delta + 1)$-coloring algorithms, we observe that COLORRED it is significantly faster than the considered variant of the Yu et al. coloring algorithm, which is the only other $(\Delta + 1)$-coloring algorithm for the SINR model. Further improvement can be made with an error-correcting variant that increases the runtime by allowing some uncertainty in the communication and afterwards correcting the introduced conflicts.

Fabian Fuchs
Karlsruhe Institute of Technology (KIT)
fabian.fuchs@kit.edu

CP0

Scaling Limit of Random k-Trees

We consider a random k-tree $\mathsf{G}_{n,k}$ that is uniformly selected from the class of labelled k-trees with $n+k$ vertices. Since 1-trees are just trees, it is well-known that $\mathsf{G}_{n,1}$ (after scaling the distances by $1/(2\sqrt{n})$) converges to the Continuum Random Tree T_e. Our main result is that for $k \neq 1$, the random k-tree $\mathsf{G}_{n,k}$, scaled by $(kH_{k-1}+1)/(2\sqrt{n})$ where H_{k-1} is the $(k-1)$-th Harmonic number, converges to the Continuum Random Tree T_e, too. In particular this shows that the diameter as well as the typical distance of two vertices in a random k-tree $\mathsf{G}_{n,k}$ are of order $\sqrt{n}$.

Michael Drmota
TU Vienna
michael.drmota@tuwien.ac.at

Emma Yu Jin
Institute of Discrete Mathematics and Geometry
Vienna University of Technology
yu.jin@tuwien.ac.at

CP0

A General Framework for Dynamic Succinct and Compressed Data Structures

Succinct data structures are becoming increasingly popular in big data processing applications due to their low memory consumption. In this paper we design, implement, and test a general framework that allows for practical dynamic succinct structures. We first describe implementations of compressed modifiable bit vectors, and extended compressed random access memory. Then, we implement and test our data structures using several popular compression libraries, and both synthetic data and a real-world temporal graph.

Patrick K. Nicholson
Bell Labs
pat.nicholson@gmail.com

Patrick Klitzke
Universitat des Saarlandes, Saarbrucken, Germany
philologos14@gmail.com

CP0

Randomized Strategies for Cardinality Robustness in the Knapsack Problem

We address randomized strategies for the following zero-sum game related to the knapsack problem. Given an instance of the knapsack problem, Alice chooses a knapsack solution and Bob, knowing Alice's solution, chooses a cardinality k. Then, Alice obtains a payoff equal to the ratio of the profit of the best k items in her solution to that of the best solution of size at most k. We first show an instance such that the payoff of an arbitrary randomized strategy is both $\mathrm{O}(\log\log\mu/\log\mu)$ and $\mathrm{O}(\log\log\rho/\log\rho)$, where μ is the exchangeability of the independence system and ρ is the ratio of the size of a maximum feasible set to that of minimum infeasible set minus one. We then design two randomized strategies with payoff $\Omega(1/\log\mu)$ and $\Omega(1/\log\rho)$ guaranteed, respectively, which substantially improve upon that of deterministic strategies and almost attain the above upper bounds.

Yusuke Kobayashi
University of Tsukuba
kobayashi@sk.tsukuba.ac.jp

Kenjiro Takazawa
Research Institute for Mathematical Sciences,
Kyoto University
takazawa@kurims.kyoto-u.ac.jp

CP0

Engineering Oracles for Time-Dependent Road Networks

We implement and experimentally evaluate landmark-based oracles for min-cost paths in large-scale time-dependent road networks. We exploit parallelism and lossless compression, combined with a novel travel-time approximation technique, to severely reduce preprocessing space and time. We significantly improve the FLAT oracle, improving the previous query-time by 30% and doubling the Dijkstra-rank speedup. We also implement and experimentally evaluate a novel oracle (HORN), based on a landmark hierarchy, achieving even better performance wrt to FLAT.

Spyros Kontogiannis
Computer Technology Institute and Press
kontog@cti.gr

George Michalopoulos
University of Patras, Greece
michalog@ceid.upatras.gr

Georgia Papastavrou
University of Ioannina, Greece
gioulycs@gmail.com

Andreas Paraskevopoulos
University of Patras, Greece
paraskevop@ceid.upatras.gr

Dorothea Wagner
Karlsruhe Institute of Technology
Institute of Theoretical Informatics
dorothea.wagner@kit.edu

Christos Zaroliagis
University of Patras, Greece
zaro@ceid.upatras.gr

CP0

An Algorithm for Online K-Means Clustering

Abstract not available.

Edo Liberty
Yahoo Labs
edo@yahoo-inc.com

CP0

More Analysis of Double Hashing for Balanced Allocations

With double hashing, for a key x, one generates two hash values $f(x)$ and $g(x)$, and then uses combinations $(f(x) + ig(x)) \bmod n$ for $i = 0, 1, 2, \ldots$ to generate multiple hash values in the range $[0, n-1]$ from the initial two. For balanced allocations, keys are hashed into a hash table where each bucket can hold multiple keys, and each key is placed in the least loaded of d choices. It has been shown previously that asymptotically the performance of double hashing and fully random hashing is the same in the balanced allocation paradigm using fluid limit methods. Here we extend a coupling argument used by Lueker and Molodowitch to show that double hashing and ideal uniform hashing are asymptotically equivalent in the setting of open address hash tables to the balanced allocation setting, providing further insight into this phenomenon. We also discuss the potential for and bottlenecks limiting the use this approach for other multiple choice hashing schemes.

Michael Mitzenmacher
Harvard University
michaelm@eecs.harvard.edu

CP0

A New Approach to Analyzing Robin Hood Hashing

Robin Hood hashing is a variation on open addressing hashing designed to reduce the maximum search time as well as the variance in the search time for elements in the hash table. While the case of insertions only using Robin Hood hashing is well understood, the behavior with deletions has remained open. Here we show that Robin Hood hashing with random hash functions can be analyzed under the framework of finite-level finite-dimensional jump Markov chains. This framework allows us to re-derive some past results for the insertion-only case with some new insight, as well as provide new analyses for a standard deletion model, where we alternate between deleting a random old key and inserting a new one.

Michael Mitzenmacher
Harvard University
michaelm@eecs.harvard.edu

CP0

The Clt Analogue for Cyclic Urns

A cyclic urn is an urn model for balls of types $0, \ldots, m-1$ where in each draw the ball drawn, say of type j, is returned to the urn together with a new ball of type $j+1$ mod m. The case $m = 2$ is the well-known Friedman urn. The composition vector, i.e., the vector of the numbers of balls of each type after n steps is, after normalization, known to be asymptotically normal for $2 \leq m \leq 6$. For $m \geq 7$ the normalized composition vector does not converge. However, there is an almost sure approximation by a periodic random vector. In this paper the asymptotic fluctuations around this periodic random vector are identified. We show that these fluctuations are asymptotically normal for all $m \geq 7$. However, they are of maximal dimension $m-1$ only when 6 does not divide m. For m being a multiple of 6 the fluctuations are supported by a two-dimensional subspace.

Ralph Neininger
J.W. Goethe University, Frankfurt
neiningr@math.uni-frankfurt.de

CP0

Geometry Helps to Compare Persistence Diagrams

Exploiting geometric structure to improve the asymptotic complexity of discrete assignment problems is a well-studied subject. In contrast, the practical advantages of using geometry for such problems have not been explored. We implement geometric variants of the Hopcroft–Karp algorithm for bottleneck matching (based on previous work by Efrat el al.), and of the auction algorithm by Bertsekas for Wasserstein distance computation. Both implementations use k-d trees to replace a linear scan with a geometric proximity query. Our interest in this problem stems from the desire to compute distances between persistence diagrams, a problem that comes up frequently in topological data analysis. We show that our geometric matching algorithms lead to a substantial performance gain, both in running time and in memory consumption, over their purely combinatorial counterparts. Moreover, our implementation significantly outperforms the only other implementation available for comparing persistence diagrams.

Michael Kerber
University of Technology, Graz, Austria
kerber@tugraz.at

Dmitriy Morozov
Lawrence Berkeley National Lab
dmitriy@mrzv.org

Arnur Nigmetov
Graz University of Technology
Institute of Geometry
nigmetov@tugraz.at

CP0

Generating Massive Scale Free Networks under Resource Constraints

Preferential attachment is a popular building block in generators for scale-free random graphs. We present the first two I/O-efficient node sampling primitives for preferen-

tial attachment with linear fitness functions and evaluate their performance for Barabási-Albert graphs. One approach is parallelizable: Our GPGPU-accelerated generator is 14.4 times faster than a sequential linear-time scheme by Batagelj and Brandes for large graphs in main memory and scales well beyond the RAM size. Both schemes support arbitrary seed graphs exceeding main memory, random node insertion patterns, directed edges and uniform node sampling.

Manuel Penschuck, Ulrich Meyer
Goethe University Frankfurt, Germany
mpenschuck@ae.cs.uni-frankfurt.de, umeyer@cs.uni-frankfurt.de

CP0

Real-Time k-Bounded Preemptive Scheduling

We consider a variant of the classic real-time scheduling problem that has natural applications in cloud computing. The input consists of a set of jobs, and an integer parameter $k \geq 1$. Each job is associated with a processing time, a release time, a due-date and a positive weight. The goal is to feasibly schedule a subset of the jobs of maximum total weight on a single machine, such that each of the jobs is preempted at most k times. Our theoretical results for the *real-time k-bounded preemptive scheduling* problem include hardness proofs, as well as algorithms for subclasses of instances, for which we derive constant-ratio performance guarantees. We bridge the gap between theory and practice through a comprehensive experimental study, in which we also test the performance of several heuristics for general instances on multiple parallel machines. The experiments employ a linear programming relaxation to upper bound the optimal solution for a given instance. Our results show that while k-bounded preemptive scheduling is hard to solve already on highly restricted instances, simple priority-based heuristics yield almost optimal schedules for realistic inputs and arbitrary values of k.

Baruch Schieber
IBM T.J. Watson Research Center
sbar@watson.ibm.com

Sivan Albagli-Kim, Hadas Shachnai
Computer Science Department, Technion
sivanalg@cs.technion.ac.il, hadas@cs.technion.ac.il

Tami Tamir
School of Computer Science, The Interdisciplinary Center
tami@idc.ac.il

CP0

k-Way Hypergraph Partitioning Via n-Level Recursive Bisection

We develop a multilevel algorithm for hypergraph partitioning that contracts the vertices one at a time. Using several caching and lazy-evaluation techniques during coarsening and refinement, we reduce the running time by up to two-orders of magnitude compared to a naive n-level algorithm that would be adequate for ordinary graph partitioning. The overall performance is even better than the widely used hMetis hypergraph partitioner that uses a classical multilevel algorithm with few levels. Aided by a portfolio-based approach to initial partitioning and adaptive budgeting of imbalance within recursive bipartitioning, we achieve very high quality. We assembled a large benchmark set with 310 hypergraphs stemming from application areas such VLSI, SAT solving, social networks, and scientific computing. Experiments indicate that our algorithm is the method of choice for a wide range of hypergraph partitioning tasks. The algorithm presented in this work forms the basis of our hypergraph partitioning framework *KaHyPar* (**Ka**rlsruhe **Hy**pergraph **Par**titioning).

Sebastian Schlag
Karlsruhe Institute of Technology
Institute for Theoretical Informatics, Algorithmics II
sebastian.schlag@kit.edu

Vitali Henne, Tobias Heuer
Karlsruhe Institute of Technology
vitali.henne@gmail.com, tobias.heuer@gmx.net

Henning Meyerhenke
Karlsruhe Institute of Technology
Institute for Theoretical Informatics, Parallel Computing
meyerhenke@kit.edu

Peter Sanders
Karlsruhe Institue of Technology
Institute of Theoretical Informatics
sanders@kit.edu

Christian Schulz
Institute of Theoretical Informatics
Karlsruhe Institute of Technology
christian.schulz@kit.edu

CP0

Total Variation Discrepancy of Deterministic Random Walks for Ergodic Markov Chains

Motivated by a derandomization of Markov chain Monte Carlo, this paper investigates *deterministic random walks*, which is a deterministic process analogous to a random walk. This paper gives upper bounds of the *total variation discrepancy* between the expected number of tokens in a Markov chain and the number of tokens in its corresponding deterministic random walk, in terms of the mixing time of the Markov chain. We also present some lower bounds.

Takeharu Shiraga, Yukiko Yamauchi, Shuji Kijima, Masafumi Yamashita
Kyushu University
takeharu.shiraga@inf.kyushu-u.ac.jp, yamauchi@inf.kyushu-u.ac.jp, kijima@inf.kyushu-u.ac.jp, mak@inf.kyushu-u.ac.jp

CP0

Scalable Transfer Patterns

We consider the problem of Pareto-optimal route planning in public-transit networks of a whole country, a whole continent, or even the whole world. On such large networks, existing approaches suffer from either a very large space consumption, a very long preprocessing time or slow query processing. Transfer Patterns, a state-of-the-art technique for route planning in transit networks, achieves excellent query times, but the space consumption is large and the preprocessing time is huge. In this paper, we introduce a new scheme for the Transfer Pattern precomputation and query graph construction that reduces both the necessary preprocessing time and space consumption by an order of magnitude and more. Average query times are below 1ms for local queries, independent of the size of the network,

around 30ms for non-local queries on the complete transit network of Germany, and an estimated 200ms for a fictitious transit network covering the currently available data of the whole world.

Hannah Bast
Albert-Ludwigs-Universität Freiburg
bast@informatik.uni-freiburg.de

Matthias Hertel
Albert-Ludwigs Universität Freiburg
hertelm@informatik.uni-freiburg.de

Sabine Storandt
Albert-Ludwigs-Universitaet Freiburg
storandt@informatik.uni-freiburg.de

CP0

Finding Near-Optimal Independent Sets at Scale

The independent set problem is NP-hard and particularly difficult to solve in large sparse graphs which typically take exponential time to solve exactly using the best-known exact algorithms. In this work, we develop an advanced evolutionary algorithm, which incorporates kernelization techniques to compute large independent sets in huge sparse networks. A recent exact algorithm has shown that large networks can be solved exactly by employing a branch-and-reduce technique that recursively kernelizes the graph and performs branching. However, one major drawback of their algorithm is that, for huge graphs, branching still can take exponential time. To avoid this problem, we recursively choose vertices that are likely to be in a large independent set (using an evolutionary approach), then further kernelize the graph. We show that identifying and removing vertices likely to be in large independent sets opens up the reduction space—which not only speeds up the computation of large independent sets drastically, but also enables us to compute high-quality independent sets on much larger instances than previously reported in the literature.

Sebastian Lamm
Institute of Theoretical Informatics
Karlsruhe Institute of Technology
lamm@ira.uka.de

Peter Sanders
Karlsruhe Institue of Technology
Institute of Theoretical Informatics
sanders@kit.edu

Christian Schulz
Institute of Theoretical Informatics
Karlsruhe Institute of Technology
christian.schulz@kit.edu

Darren Strash
Institute of Theoretical Informatics
Karlsruhe Institute of Techology
strash@kit.edu

Renato F. Werneck
Unaffiliated
rwerneck@acm.org

CP0

Graph Bisection with Pareto-Optimization

We introduce FlowCutter, a novel algorithm to compute a set of edge cuts or node separators that optimize cut size and balance in the Pareto-sense. Using the computed Pareto-set we can identify cuts with a particularly good trade-off between cut size and balance that can be used to compute contraction and minimum fill-in orders, which can be used in Customizable Contraction Hierarchies, a speed-up technique for shortest path computations.

Michael Hamann
Karlsruhe Institute of Technology
michael.hamann@kit.edu

Ben Strasser
Karlsruhe Institute of Technology
Institute of Theoretical Informatics
strasser@kit.edu

CP0

Prime Factorization of the Kirchhoff Polynomial: Compact Enumeration of Arborescences

We study the problem of enumerating all arborescences of a directed graph G. They can be expressed as the Kirchhoff polynomial, a polynomial over the edge labels of G, in which each monomial represents a single arborescence. We show how to compute, in linear time, a compact representation of the Kirchhoff polynomial – its prime factorization, and how it relates to combinatorial properties of digraphs such as strong connectivity and vertex domination.

Przemyslaw Uznanski
Aalto University
przemyslaw.uznanski@aalto.fi

Matus Mihalak
Maastricht University
matus.mihalak@maastrichtuniversity.nl

Pencho Yordanov
ETH Zurich
pencho.yordanov@bsse.ethz.ch

CP0

Ricci-Ollivier Curvature of the Rooted Phylogenetic Subtree-Prune-Regraft Graph

We investigated the Ricci-Ollivier curvature of the graph induced by subtree-prune-regraft moves on rooted tree topologies, such as those explored by statistical phylogenetic inference methods. We developed fast new algorithms for constructing and sampling these graphs. We confirm using simulation that access time distributions of random walks on these graphs depend on distance, degree, and curvature. We then study degree changes in these graphs theoretically and give bounds on the curvature.

Chris Whidden, Frederick Matsen
Fred Hutchinson Cancer Research Center

cwhidden@fredhutch.org, matsen@fredhutch.org

CP0

On Terminal Nodes and the Degree Profile of Preferential Dynamic Attachment Circuits

We investigate terminal nodes and the degree profile in preferential dynamic attachment circuits. We study the distribution of the number of terminal nodes, which are the nodes that have not recruited other nodes, as the circuit ages. The expectation and variance of the number of terminal nodes are both linear with respect to the age of the circuit. We show via martingale that the number of terminal nodes asymptotically follows a Gaussian law. We also study the exact distribution of the degree of a specific node as the circuit grows. The exact expectation and variance of the degree of a node are determined via a series of Pólya–Eggenberger urn models with "hiccups' in between and recurrence methods. Phase transitions of these degrees are discussed briefly.

Panpan Zhang
The George Washington University
pzhang@gwu.edu

CP0

On Connectivity in a General Random Intersection Graph

There has been growing interest recently in studies of general random intersection graphs. In this paper, we consider a general random intersection graph $\mathbb{G}(n, \overrightarrow{a}, \overrightarrow{K_n}, P_n)$ defined on a set $\mathcal{V}_n$ comprising n vertices, where $\overrightarrow{a}$ is a probability vector $(a_1, a_2, \ldots, a_m)$ and $\overrightarrow{K_n}$ is $(K_{1,n}, K_{2,n}, \ldots, K_{m,n})$. Suppose there is a pool $\mathcal{P}_n$ consisting of P_n distinct objects. The n vertices in $\mathcal{V}_n$ are divided into m groups $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_m$. Each vertex v is independently assigned to exactly a group according to the probability distribution with $\mathbb{P}[v \in \mathcal{A}_i] = a_i$, where $i = 1, 2, \ldots, m$. Afterwards, each vertex in group $\mathcal{A}_i$ independently chooses $K_{i,n}$ objects uniformly at random from the object pool $\mathcal{P}_n$. Finally, an undirected edge is drawn between two vertices in $\mathcal{V}_n$ that share at least one object. This graph model $\mathbb{G}(n, \overrightarrow{a}, \overrightarrow{K_n}, P_n)$ has applications in secure sensor networks and social networks. We investigate connectivity in this general random intersection graph $\mathbb{G}(n, \overrightarrow{a}, \overrightarrow{K_n}, P_n)$ and present a sharp zero-one law.

Jun Zhao
Carnegie Mellon University
junzhao@cmu.edu

Panpan Zhang
George Washington University
pzhang@emailgwu.edu

CP1

Undirected Graph Exploration with $\Theta(\log \log N)$ Pebbles

We consider the fundamental problem of exploring an undirected and initially unknown graph by an agent with little memory. The vertices of the graph are unlabeled, and the edges incident to a vertex have locally distinct labels. In this setting, it is known that $\Theta(\log n)$ bits of memory are necessary and sufficient to explore any graph with at most n vertices. We show that this memory requirement can be decreased significantly by making a part of the memory distributable in the form of pebbles. A pebble is a device that can be dropped to mark a vertex and can be collected when the agent returns to the vertex. We show that for an agent $\mathcal{O}(\log \log n)$ distinguishable pebbles and bits of memory are sufficient to explore any bounded-degree graph with at most n vertices. We match this result with a lower bound exhibiting that for any agent with sub-logarithmic memory, $\Omega(\log \log n)$ distinguishable pebbles are necessary for exploration.

Yann Disser, Jan Hackfeld, Max Klimm
TU Berlin
disser@math.tu-berlin.de, hackfeld@math.tu-berlin.de, klimm@math.tu-berlin.de

CP1

New Directions in Nearest Neighbor Searching with Applications to Lattice Sieving

To solve the approximate nearest neighbor search problem (NNS) on the sphere, we propose a method using locality-sensitive filters (LSF), with the property that nearby vectors have a higher probability of surviving the same filter than vectors which are far apart. We instantiate the filters using spherical caps of height $1-\alpha$, where a vector survives a filter if it is contained in the corresponding spherical cap, and where ideally each filter has an independent, uniformly random direction. For small α, these filters are very similar to the spherical locality-sensitive hash (LSH) family previously studied by Andoni et al. For larger α bounded away from 0, these filters potentially achieve a superior performance, provided we have access to an efficient oracle for finding relevant filters. Whereas existing LSH schemes are limited by a performance parameter of $\rho \geq 1/(2c^2-1)$ to solve approximate NNS with approximation factor c, with spherical LSF we potentially achieve smaller asymptotic values of ρ, depending on the density of the data set. For sparse data sets where the dimension is super-logarithmic in the size of the data set, we asymptotically obtain $\rho = 1/(2c^2-1)$, while for a logarithmic dimensionality with density κ we obtain asymptotics of $\rho \sim 1/(4\kappa c^2)$. To instantiate the filters, we replace the independent filters by filters taken from certain structured random product codes. We show that the additional structure in these concatenation codes allows us to decode efficiently using techniques similar to lattice enumeration, and we can find the relevant filters with low overhead, while at the same time not significantly changing the collision probabilities of the filters. We finally apply spherical LSF to sieving algorithms for solving the shortest vector problem (SVP) on lattices, and show that this leads to a time complexity for solving SVP in dimension n of $(3/2)^{n/2+o(n)}$.

Anja Becker
EPFL, Lausanne, Switzerland
anja.becker@epfl.ch

Leo Ducas
CWI, Amsterdam, The Netherlands
l.ducas@cwi.nl

Nicolas Gama
UVSQ, Versailles, France
gama.nicolas@free.fr

Thijs Laarhoven
Technische Universiteit Eindhoven

t.m.m.laarhoven@tue.nl

CP1

Locality-Sensitive Hashing Without False Negatives

We consider a new construction of locality-sensitive hash functions for Hamming space that is *covering* in the sense that is it guaranteed to produce a collision for every pair of vectors within a given radius r. The construction is *efficient* in the sense that the expected number of hash collisions between vectors at distance cr, for a given $c > 1$, comes close to that of the best possible data independent LSH without the covering guarantee, namely, the seminal LSH construction of Indyk and Motwani (FOCS '98). The efficiency of the new construction essentially *matches* their bound if $cr = \log(n)/k$, where n is the number of points in the data set and $k \in \mathbf{N}$, and differs from it by at most a factor $\ln(4) < 1.4$ in the exponent for general values of cr. As a consequence, LSH-based similarity search in Hamming space can avoid the problem of false negatives at little or no cost in efficiency.

Rasmus Pagh
IT University of Copenhagen
pagh@itu.dk

CP1

Phase Transitions in Group Testing

We study the fundamental limits of any group testing procedure regardless of its computational complexity. In the noiseless case with the number of defective items k scaling with the total number of items p as $O(p^\theta)$ ($\theta \in (0,1)$), we show that the probability of reconstruction error tends to one when $n \leq k \log_2 \frac{p}{k}(1 + o(1))$, but vanishes when $n \geq c(\theta)k \log_2 \frac{p}{k}(1 + o(1))$, for some explicit constant $c(\theta)$. For $\theta \leq \frac{1}{3}$, we show that $c(\theta) = 1$, thus proving a phase transition. Analogous conditions are derived for noisy and partial recovery settings.

Volkan Cevher
EPFL
volkan.cevher@epfl.ch

Jonathan Scarlett
École Polytechnique Fédérale de Lausanne
jmscarlett@gmail.com

CP1

Tight Bounds for the Distribution-Free Testing of Monotone Conjunctions

We study the distribution-free testing of monotone conjunctions. We present an $\tilde{O}(n^{1/3}/\epsilon^5)$-query algorithm that tests whether an unknown f is a monotone conjunction versus ϵ-far from monotone conjunctions with respect to an unknown distribution. This improves the $\tilde{O}(n^{1/2}/\epsilon)$ upper bound by Dolev and Ron. We also prove a lower bound of $\tilde{\Omega}(n^{1/3})$, improving the $\tilde{\Omega}(n^{1/5})$ lower bound by Glasner and Servedio. Our bounds are tight, up to a polylogarithmic factor, when ϵ is a constant.

Xi Chen
Columbia University
xichen@cs.columbia.edu

Jinyu Xie
Columbia University
Columbia University
jinyu@cs.columbia.edu

CP2

The Complexity of All-Switches Strategy Improvement

Strategy improvement is a widely-used and well-studied class of algorithms for solving graph-based infinite games. We show that two natural problems are PSPACE-complete for the all-switches variant of strategy improvement: the problem of determining whether a given edge is switched, and the problem of determining which optimal strategy is found. These results hold for parity games games, mean-payoff games, discounted-payoff games, and simple-stochastic games. We also show related results for acyclic unique sink orientations.

John Fearnley, Rahul Savani
University of Liverpool
john.fearnley@liverpool.ac.uk,
rahul.savani@liverpool.ac.uk

CP2

Characterisation of Strongly Stable Matchings

We present a characterisation of the set of all strongly stable matchings, thus solving an open problem already stated in 1989 in the book by Gusfield and Irving. Although the number of strongly stable matchings can be exponential, we show that there exists a partial order with $O(m)$ elements representing all strongly stable matchings. We give two algorithms that construct two such representations: one in $O(nm^2)$ time and the other in $O(nm)$ time.

Pratik Ghosal, Adam Kunysz
Institute of Computer Science
University of Wroclaw
pratikghosal@cs.uni.wroc.pl, aku@cs.uni.wroc.pl

Katarzyna Paluch
Institute of Computer Science, University of Wroclaw
abraka@cs.uni.wroc.pl

CP2

An Improved Combinatorial Polynomial Algorithm for the Linear Arrow-Debreu Market

We present an improved combinatorial algorithm for the computation of equilibrium prices in the linear Arrow-Debreu model. For a market with n agents and integral utilities bounded by U, the algorithm runs in $O(n^7 \log^3(nU))$ time. This improves upon the previously best algorithm of Ye by a factor of $\tilde{\Omega}(n)$. The algorithm refines the algorithm described by Duan and Mehlhorn and improves it by a factor of $\tilde{\Omega}(n^3)$. The improvement comes from a better understanding of the iterative price adjustment process, the improved balanced flow computation for nondegenerate instances, and a novel perturbation technique for achieving nondegeneracy.

Ran Duan
Tsinghua University, Beijing, China
duanran@mail.tsinghua.edu.cn

Jugal Garg

MPI Informatics
jugal@mpi-inf.mpg.de

Kurt Mehlhorn
MPI for Informatics
Saarbruecken, Germany
mehlhorn@mpi-inf.mpg.de

CP2

Designing Networks with Good Equilibria under Uncertainty

We consider the problem of designing network cost-sharing protocols with good equilibria under uncertainty, where the designer has incomplete information about the input but has prior knowledge of the underlying metric. We propose two different models, the *adversarial* and the *stochastic*. The main question we address is: *to what extent can prior knowledge of the underlying metric help in the design?*

George Christodoulou, Alkmini Sgouritsa
University of Liverpool
gchristo@liverpool.ac.uk, salkmini@liverpool.ac.uk

CP2

Learning and Efficiency in Games with Dynamic Population

We study the quality of outcomes in games when the population of players is dynamically changing, and where participants have to adapt to the dynamic environment. Price of Anarchy has originally been introduced to study the Nash equilibria of one-shot games, but has been extended since to repeated setting, assuming all players use learning strategies, and the environment, as well as the player population is stable. We show that in large classes of games (including independent item auctions and congestion games), if players use a form of learning that helps them to adapt to the changing environment, this guarantees high social welfare, even under very frequent changes.

Thodoris Lykouris, Vasilis Syrgkanis
Cornell University
teddlyk@cs.cornell.edu, vasy@microsoft.com

Eva Tardos
Cornell University
Department of Computer Science
eva@cs.cornell.edu

CP3

Scheduling Parallel Dag Jobs Online to Minimize Average Flow Time

In this work, we study the problem of scheduling parallelizable jobs online with an objective of minimizing average flow time. Each parallel job is modeled as a DAG where each node is a sequential task and each edge represents dependence between tasks. Previous work has focused on a model of parallelizability known as the arbitrary speed-up curves setting where a scalable algorithm is known. However, the DAG model is more widely used by practitioners, since many jobs generated from parallel programming languages and libraries can be represented in this model. However, little is known for this model in the online setting with multiple jobs. The DAG model and the speed-up curve models are incomparable and algorithmic results from one do not immediately imply results for the other. Previous work has left open the question of whether an online algorithm can be $O(1)$-competitive with $O(1)$-speed for average flow time in the DAG setting. In this work, we answer this question positively by giving a scalable algorithm which is $(1+\epsilon)$-speed $O(\frac{1}{\epsilon^3})$-competitive for any $\epsilon > 0$. We further introduce the *first* greedy algorithm for scheduling parallelizable jobs — our algorithm is a generalization of the shortest jobs first algorithm. Greedy algorithms are among the most useful in practice due to their simplicity. We show that this algorithm is $(2+\epsilon)$-speed $O(\frac{1}{\epsilon^4})$-competitive for any $\epsilon > 0$.

Kunal Agrawal, Jing Li, Kefu Lu, Benjamin Moseley
Washington University in St. Louis
kunal@wustl.edu, li.jing@wustl.edu, kefulu@wustl.edu, bmoseley@wustl.edu

CP3

Online Pricing with Impatient Bidders

In this paper we consider the following online pricing problem. An auctioneer is selling identical items in unlimited supply, whereas each bidder from a given set is interested in purchasing a single copy of the item. Each bidder is characterized by a budget and a time interval, in which he is considering to buy the item. Bidders are willing to buy the item at the earliest time provided it is within their time intervals and the price at that time is within their budgets. We call such bidders *impatient bidders*. The problem is considered in the online setting, i.e., each bidder arrives at the start of his time interval, and only then an algorithm learns of his existence and his budget. The goal of the seller is to set the price of the item over time so that the total revenue is maximized. We study two versions of the impatient bidders problem: the one introduced by Bansal et al. [TALG'10], and a more restricted setting in which the deadline of each bidder remains unknown until it is hit. We give tight bounds for both settings. Rather surprisingly, in both cases the optimum competitive ratios are the same. In particular we prove that the competitive ratio of an optimum deterministic algorithm is $\Theta(\log h / \log\log h)$, whereas for randomized algorithms it is $\Theta(\log\log h)$.

Marek Cygan
Institute of Informatics, University of Warsaw
cygan@mimuw.edu.pl

Marcin Mucha
University of Warsaw
Faculty of Mathematics, Informatics and Mechanics
mucha@mimuw.edu.pl

Piotr Sankowski, Qiang Zhang
Institute of Informatics, University of Warsaw
sank@mimuw.edu.pl, qzhang@mimuw.edu.pl

CP3

An $\mathcal{O}(\log m)$-Competitive Algorithm for Online Machine Minimization

We consider the online machine minimization problem in which jobs with hard deadlines arrive online over time at their release dates. The task is to determine a feasible preemptive schedule on a minimum number of machines. Our main result is a general $\mathcal{O}(\log m)$-competitive algorithm for the online problem, where m is the optimal number of machines used in an offline solution. This is the first improvement on an intriguing problem in nearly two decades. To date, the best known result is a $\mathcal{O}(\log(p_{\max}/p_{\min}))$-

competitive algorithm by Phillips et al. (STOC 1997) that depends on the ratio of maximum and minimum job sizes, $p_{\max}$ and $p_{\min}$. Even for $m = 2$ no better algorithm was known. Our algorithm is in this case constant-competitive. When applied to laminar or agreeable instances, our algorithm achieves a competitive ratio of $\mathcal{O}(1)$ even independently of m.

Lin Chen, Nicole Megow
TU München
lchen@ma.tum.de, nmegow@ma.tum.de

Kevin Schewior
TU Berlin
schewior@math.tu-berlin.de

CP3
Online Degree-Bounded Steiner Network Design

We initiate the study of degree-bounded network design problems in the online setting. The degree-bounded Steiner tree problem, which asks for a subgraph with minimum degree that connects a given set of vertices, is perhaps one of the most representative problems in this class. We design an intuitive greedy-like algorithm that achieves a competitive ratio of O(log n) where n is the number of vertices. We show that no (randomized) algorithm can achieve a (multiplicative) competitive ratio o(log n); thus our result is asymptotically tight. Our results carry over to the more general degree-bounded Steiner forest problem as well.

Sina Dehghani, Soheil Ehsani
University of Maryland at College Park
dehghani@umd.edu, ehsani@umd.edu

MohammadTaghi Hajiaghayi
University of Maryland, College Park
hajiagha@cs.umd.edu

Vahid Liaghat
University of Maryland at College Park
vliaghat@cs.umd.edu

Saeed Seddighin
University of Maryland, College Park
saeedreza.seddighin@gmail.com

CP3
Make-to-Order Integrated Scheduling and Distribution

Production and distribution are fundamental operational functions in supply chains. The main challenge is to design algorithms that optimize operational performance by jointly scheduling production and delivery of customer orders. In this paper we study a model of scheduling customer orders on multiple identical machines and their distribution to customers afterwards. The goal is to minimize the total time from release to distribution plus total distribution cost to the customers. We design the first poly-logarithmic competitive algorithm for the problem, improving upon previous algorithms with linear competitive ratios. Our model generalizes two fundamental problems: scheduling of jobs on multiple identical machines (where the goal function is to minimize the total flow time) as well as the TCP Acknowledgment problem.

Yossi Azar
Tel-Aviv University
azar@tau.ac.il

Amir Epstein
IBM Research-Haifa
amirepst@gmail.com

Lukasz Jez
Tel Aviv University
University of Wroclaw, Institute of Computer Science
lje@cs.uni.wroc.pl

Adi Vardi
Tel Aviv University
adi.vardi@gmail.com

CP4
Expanders Via Local Edge Flips

Designing distributed and scalable algorithms to improve network connectivity is a central topic in peer-to-peer networks. In this paper we focus on the following well-known problem: given an n-node d-regular network, we want to design a decentralized, local algorithm that transforms the graph into one that has good connectivity properties (low diameter, expansion, etc.). To this end, Mahlmann and Schindelhauer introduced the random "flip' transformation, where in each time step, a random pair of vertices that have an edge decide to 'swap a neighbor'. They conjectured that performing $poly(d) \times n \log n$ such flips at random would convert any connected d-regular graph into a d-regular expander graph, with high probability. However, the best known upper bound for the number of steps is roughly $O(n^{17}d^{23})$, obtained via a delicate Markov chain comparison argument. Our main result is to prove that a natural instantiation of the random flip produces an expander in at most $O(n^2d^2\sqrt{\log n})$ steps, with high probability. Our argument uses a potential-function analysis based on the matrix exponential, together with the recent beautiful results on the higher-order Cheeger inequality of graphs. We also show that our technique can be used to analyze another well-studied random process known as the 'random switch', and show that it produces an expander in $O(nd)$ steps for $d = \Omega(\log n)$, with high probability.

Zeyuan Allen-Zhu
MIT CSAIL
zeyuan@csail.mit.edu

Aditya Bhaskara, Silvio Lattanzi
Google Research NYC
bhaskaraaditya@gmail.com, silviol@google.com

Vahab Mirrokni
Google Inc.
mirrokni@google.com

Lorenzo Orecchia
Boston University
orecchia@bu.edu

CP4
An Improved Distributed Algorithm for Maximal Independent Set

The Maximal Independent Set (MIS) problem is one of the basics in the study of locality in distributed graph algorithms. This paper presents an extremely simple randomized algorithm providing a near-optimal local complexity

for this problem, which incidentally, when combined with some known techniques, also leads to a near-optimal global complexity. Classical MIS algorithms of Luby [STOC'85]

and Alon, Babai and Itai [JALG'86] provide the global complexity guarantee that, with high probability, all nodes terminate after $O(\log n)$ rounds. In contrast, our initial focus is on the local complexity, and our main contribution is to provide a very simple algorithm guaranteeing that each particular node v terminates after $O(\log \mathsf{deg}(v) + \log 1/\epsilon)$ rounds, with probability at least $1 - \epsilon$. This degree-dependency is optimal, due to a lower bound of Kuhn, Moscibroda, and Wattenhofer [PODC'04]. Interestingly,

this local complexity smoothly transitions to a global complexity: by adding techniques of Barenboim, Elkin, Pettie, and Schneider [FOCS'12; arXiv: 1202.1983v3], we get a randomized MIS algorithm with a high probability global complexity of $O(\log \Delta) + 2^{O(\sqrt{\log \log n})}$, where Δ denotes the maximum degree. This improves over the $O(\log^2 \Delta) + 2^{O(\sqrt{\log \log n})}$ result of Barenboim et al., and gets close to the $\Omega(\min\{\log \Delta, \sqrt{\log n}\})$ lower bound of Kuhn et al. Corollaries include improved algorithms

for MIS in graphs of upper-bounded arboricity, or lower-bounded girth, for Ruling Sets, for MIS in the Local Computation Algorithms (LCA) model, and a faster distributed algorithm for the Lovász Local Lemma.

Mohsen Ghaffari
MIT
ghaffari@mit.edu

CP4

Distributed Algorithms for Planar Networks II: Low-Congestion Shortcuts, MST, and Min-Cut

This paper introduces the concept of low-congestion shortcuts for (near-)planar networks, and demonstrates their power by using them to obtain near-optimal distributed algorithms for problems such as Minimum Spanning Tree (MST) or Minimum Cut, in planar networks.Consider a

graph $G = (V, E)$ and a partitioning of V into subsets of nodes $S_1, \ldots, S_N$, each inducing a connected subgraph $G[S_i]$. We define an α-congestion shortcut with dilation β to be a set of subgraphs $H_1, \ldots, H_N \subseteq G$, one for each subset S_i, such that:

1. For each $i \in [1, N]$, the diameter of the subgraph $G[S_i] + H_i$ is at most β.
2. For each edge $e \in E$, the number of subgraphs $G[S_i] + H_i$ containing e is at most α.

We prove that any partition of a D-diameter planar graph into individually-connected parts admits an $O(D \log D)$-congestion shortcut with dilation $O(D \log D)$, which we prove to be nearly best-possible, and we also present a distributed construction of it in $\tilde{O}(D)$ rounds. Finally, we

use low-congestion shortcuts, and their efficient distributed construction, to derive $\tilde{O}(D)$-round distributed algorithms for MST and Min-Cut, in planar networks. This complexity nearly matches the trivial lower bound of $\Omega(D)$. We remark that this is the first result bypassing the well-known $\tilde{\Omega}(D + \sqrt{n})$ existential lower bound of general graphs (see Peleg and Rubinovich [FOCS'99]; Elkin [STOC'04]; and Das Sarma et al. [STOC'11]) in a family of graphs of interest.

Mohsen Ghaffari
MIT
ghaffari@mit.edu

Bernhard Haupler
CMU
haeupler@cs.cmu.edu

CP4

The Adversarial Noise Threshold for Distributed Protocols

We consider the problem of implementing distributed protocols, despite adversarial channel errors, on synchronous-messaging networks with arbitrary topology. Our main result: Any n-party T-round protocol on an undirected communication network with m edges can be compiled into an $\mathcal{O}\left(\frac{m \log n}{n} T\right)$-round simulation protocol that runs on an $\mathcal{O}(n)$-edge subnetwork and tolerates an (optimal) adversarial error rate of $\Omega\left(\frac{1}{n}\right)$.

William M. Hoza, Leonard Schulman
Caltech
whoza@caltech.edu, schulman@caltech.edu

CP4

Local-on-Average Distributed Tasks

A distributed task is *local* if its time complexity is (nearly) constant, otherwise it is *global*. Unfortunately, local tasks are relatively scarce, and most distributed tasks require time at least logarithmic in the network size (and often higher than that). In a dynamic setting, i.e., when the network undergoes repeated and frequent topological changes, such as vertex and edge insertions and deletions, it is desirable to be able to perform a *local* update procedure around the modified part of the network, rather than running a static *global* algorithm from scratch following each change. This paper makes a step towards establishing the hypothesis that many (statically) *non-local* distributed tasks are *local-on-average* in the dynamic setting, namely, their *amortized* time complexity is $O(\log^* n)$. Towards establishing the plausibility of this hypothesis, we propose a strategy for transforming static $O(\mathrm{polylog}(n))$ time algorithms into dynamic $O(\log^* n)$ amortized time update procedures. We then demonstrate the usefulness of our strategy by applying it to several fundamental problems whose static time complexity is logarithmic, including forest-decomposition, edge-orientation and coloring sparse graphs, and show that their amortized time complexity in the dynamic setting is indeed $O(\log^* n)$.

Merav Parter
MIT
meravparter@gmail.com

David Peleg, Shay Solomon
Weizmann Institute of Science
david.peleg@weizmann.ac.il, solo.shay@gmail.com

CP5

Nearly-Optimal Bounds for Sparse Recovery in Generic Norms, with Applications to k-Median Sketching

We initiate the study of trade-offs between sparsity and the number of measurements in sparse recovery schemes for *generic* norms. Specifically, for a norm $\|\cdot\|$, sparsity parameter k, approximation factor $K > 0$, and probability of failure $P > 0$, we ask: what is the minimal value of m so

that there is a distribution over $m \times n$ matrices A with the property that for any x, given Ax, we can recover a k-sparse approximation to x in the given norm with probability at least $1 - P$? We give a partial answer to this problem, by showing that for norms that admit efficient linear sketches, the optimal number of measurements m is closely related to the *doubling dimension* of the metric induced by the norm $\|\cdot\|$ on the set of all k-sparse vectors. By applying our result to specific norms, we cast known measurement bounds in our general framework (for the ℓ_p norms, $p \in [1, 2]$) as well as provide new, measurement-efficient schemes (for the Earth-Mover Distance norm). The latter result directly implies more succinct linear sketches for the well-studied planar *k-median clustering* problem. Finally, our lower bound for the doubling dimension of the EMD norm enables us to resolve the open question of [Frahling-Sohler, STOC'05] about the space complexity of clustering problems in the dynamic streaming model.

Arturs Backurs
MIT
backurs@mit.edu

Piotr Indyk
Massachusetts Institute of Technology
indyk@mit.edu

Ilya Razenshteyn
MIT
ilyaraz@mit.edu

David Woodruff
IBM Almaden
dpwoodru@us.ibm.com

CP5

Nearly Optimal Deterministic Algorithm for Sparse Walsh-Hadamard Transform

For every fixed constant $\alpha > 0$, we design an algorithm for computing the k-sparse Walsh-Hadamard transform (i.e., Discrete Fourier Transform over the Boolean cube) of an N-dimensional vector $x \in \mathbb{R}^N$ in time $k^{1+\alpha}(\log N)^{O(1)}$. Specifically, the algorithm is given query access to x and computes a k-sparse $\tilde{x} \in \mathbb{R}^N$ satisfying $\|\tilde{x} - \hat{x}\|_1 \leq c\|\hat{x} - H_k(\hat{x})\|_1$, for an absolute constant $c > 0$, where $\hat{x}$ is the transform of x and $H_k(\hat{x})$ is its best k-sparse approximation. Our algorithm is fully deterministic and only uses non-adaptive queries to x (i.e., all queries are determined and performed in parallel when the algorithm starts). An important technical tool that we use is a construction of nearly optimal and linear lossless condensers which is a careful instantiation of the GUV condenser (Guruswami, Umans, Vadhan, JACM 2009). Moreover, we design a deterministic and non-adaptive ℓ_1/ℓ_1 compressed sensing scheme based on general lossless condensers that is equipped with a fast reconstruction algorithm running in time $k^{1+\alpha}(\log N)^{O(1)}$ (for the GUV-based condenser) and is of independent interest. Our scheme significantly simplifies and improves an earlier expander-based construction due to Berinde, Gilbert, Indyk, Karloff, Strauss (Allerton 2008). Our methods use linear lossless condensers in a black box fashion; therefore, any future improvement on explicit constructions of such condensers would immediately translate to improved parameters in our framework (potentially leading to $k(\log N)^{O(1)}$ reconstruction time with a reduced exponent in the poly-logarithmic factor, and eliminating the extra parameter α). By allowing the algorithm to use randomness, while still using non-adaptive queries, the running time of the algorithm can be improved to $\tilde{O}(k \log^3 N)$.

Mahdi Cheraghchi
Carnegie Mellon University
Computer Science Department
m.cheraghchi@imperial.ac.uk

Piotr Indyk
Massachusetts Institute of Technology
indyk@mit.edu

CP5

Nearly Tight Oblivious Subspace Embeddings by Trace Inequalities

We present a new analysis of sparse oblivious subspace embeddings, based on the "matrix Chernoff" technique. These are probability distributions over (relatively) sparse matrices such that for any d-dimensional subspace of R^n, the norms of all vectors in the subspace are simultaneously approximately preserved by the embedding with high probability–typically with parameters depending on d but not on n. The families of embedding matrices considered here are essentially the same as those in [Nelson-Nguyen '13], but with better parameters (sparsity and embedding dimension). Because of this, this analysis essentially serves as a "drop-in replacement' for Nelson-Nguyen's, improving bounds on its many applications to problems such as as least squares regression and low-rank approximation. This new method is based on elementary tail bounds combined with matrix trace inequalities (Golden-Thompson or Lieb's theorem), and does not require combinatorics, unlike the Nelson-Nguyen approach. There are also variants of this method that are even simpler, at the cost of worse parameters. Furthermore, the bounds obtained are much tighter than previous ones, matching known lower bounds up to a single log(d) factor in embedding dimension (previous results had more log factors and also had suboptimal tradeoffs with sparsity).

Michael B. Cohen
MIT
n/a

CP5

The Restricted Isometry Property of Subsampled Fourier Matrices

A matrix $A \in \mathbb{C}^{q \times N}$ satisfies the restricted isometry property of order k with constant ϵ if it preserves the ℓ_2 norm of all k-sparse vectors up to a factor of $1 \pm \epsilon$. We prove that a matrix A obtained by randomly sampling $q = O(k \cdot \log^2 k \cdot \log N)$ rows from an $N \times N$ Fourier matrix satisfies the restricted isometry property of order k with a fixed ϵ with high probability. This improves on Rudelson and Vershynin (Comm. Pure Appl. Math., 2008), its subsequent improvements, and Bourgain (GAFA Seminar Notes, 2014).

Ishay Haviv
The Academic College of Tel Aviv-Yaffo
ishayhav@mta.ac.il

Oded Regev
New York University

regev@cims.nyu.edu

CP6

Simpler, Faster and Shorter Labels for Distances in Graphs

We present a distance labeling scheme for undirected graphs with n nodes having labels of length $\frac{\log 3}{2}n + o(n)$ bits and constant decoding time. This outperforms all existing results with respect to both size and decoding time, including Winkler's (Combinatorica 1983) decade-old result, which uses labels of size $(\log 3)n$ and $O(n/\log n)$ decoding time, and Gavoille et al. (SODA'01), which uses labels of size $11n + o(n)$ and $O(\log \log n)$ decoding time.

Stephen Alstrup
Department of Computer Science
University of Copenhagen
s.alstrup@diku.dk

Cyril Gavoille
Bordeaux University
gavoille@labri.fr

Esben Bistrup Halvorsen
University of Copenhagen
esben@bistruphalvorsen.dk

Holger Petersen
dr.holger.petersen@googlemail.com
n/a

CP6

New Bounds for Approximating Extremal Distances in Undirected Graphs

We provide new approximation bounds for the extremal distances (the diameter, the radius, and the eccentricities of all nodes) of an undirected graph with n nodes and m edges. First we show under the Strong Exponential Time Hypothesis (SETH) of Impagliazzo, Paturi and Zane [JCSS01] that it is impossible to get a $(3/2 - \varepsilon)$-approximation of the diameter and a $(5/3 - \varepsilon)$-approximation of all the eccentricities in $O(m^{2-\delta})$ time for any $\varepsilon, \delta > 0$, even allowing a constant additive term in the approximation. Second, we present an algorithmic scheme that gives a $(2 - 1/2^k)$-approximation of the diameter and the radius and a $(3 - 4/(2^k + 1))$-approximation of all eccentricities in $\tilde{O}(mn^{\frac{1}{k+1}})$ expected time for any $k \geq 0$. For $k \geq 2$, this gives a family of previously unknown bounds, and approaches near-linear running time as k grows. Third, we observe a connection between the approximation of the diameter and the h-dominating sets, which are subsets of nodes at distance $\leq h$ from every other node. We give bounds for the size of these sets, related with the diameter.

Massimo Cairo
Scuola Normale Superiore
cairomassimo@gmail.com

Roberto Grossi
Universita‘ di Pisa
grossi@di.unipi.it

Romeo Rizzi
Universita di Verona
romeo.rizzi@univr.it

CP6

Approximation and Fixed Parameter Subquadratic Algorithms for Radius and Diameter in Sparse Graphs

The radius and diameter are fundamental graph parameters, with several natural definitions for directed graphs. All versions of diameter and radius can be solved via solving all-pairs shortest paths (APSP), followed by a fast postprocessing step. However, solving APSP on n-node graphs requires $\Omega(n^2)$ time even in sparse graphs. We study the question: when can diameter and radius in *sparse* graphs be solved in truly subquadratic time, and when is such an algorithm unlikely? Motivated by our conditional lower bounds on computing these measures exactly in truly subquadratic time, we search for *approximation* and *fixed parameter subquadratic* algorithms, and alternatively, for reasons why they do not exist. We find that: - Most versions of Diameter and Radius can be solved in truly subquadratic time with *optimal* approximation guarantees, under plausible assumptions. For example, there is a 2-approximation algorithm for directed Radius with one-way distances that runs in $\tilde{O}(m\sqrt{n})$ time, while a $(2 - \delta)$-approximation algorithm in $O(n^{2-\epsilon})$ time is considered unlikely. - On graphs with treewidth k, we can solve all versions in $2^{O(k \log k)}n^{1+o(1)}$ time. We show that these algorithms are near *optimal* since even a $(3/2 - \delta)$-approximation algorithm that runs in time $2^{o(k)}n^{2-\epsilon}$ would refute plausible assumptions. Two conceptual contributions of this work that we hope will incite future work are: the introduction of a *Fixed Parameter Tractability in P* framework, and the statement of a differently-quantified variant of the Orthogonal Vectors Conjecture, which we call the *Hitting Set Conjecture*.

Amir Abboud, Virginia Vassilevska Williams,
Joshua R. Wang
Stanford University
abboud@cs.stanford.edu, virgi@cs.stanford.edu, jrwang@stanford.edu

CP6

Approximate Distance Oracles for Planar Graphs with Improved Query Time-Space Tradeoff

We consider approximate distance oracles for edge-weighted n-vertex undirected planar graphs. Given fixed $\epsilon > 0$, we present a $(1+\epsilon)$-approximate distance oracle with $O(n(\log \log n)^2)$ space and $O((\log \log n)^3)$ query time. This improves the previous best product of query time and space of the oracles of Thorup (FOCS 2001, J. ACM 2004) and Klein (SODA 2002) from $O(n \log n)$ to $O(n(\log \log n)^5)$.

Christian Wulff-Nilsen
Department of Computer Science
University of Copenhagen
koolooz@di.ku.dk

CP7

Deterministic Algorithms for Submodular Maximization Problems

Randomization is a fundamental tool used in many theoretical and practical areas of computer science. We study here the role of randomization in the area of submodular function maximization. In this area most algorithms

are randomized, and in almost all cases the approximation ratios obtained by current randomized algorithms are superior to the best results obtained by known deterministic algorithms. Derandomization of algorithms for general submodular function maximization seems hard since the access to the function is done via a value oracle. This makes it hard, for example, to apply standard derandomization techniques such as conditional expectations. Therefore, an interesting fundamental problem in this area is whether randomization is inherently necessary for obtaining good approximation ratios. In this work we give evidence that randomization is not necessary for obtaining good algorithms by presenting a new technique for derandomization of algorithms for submodular function maximization. Our high level idea is to maintain explicitly a (small) distribution over the states of the algorithm, and carefully update it using marginal values obtained from an extreme point solution of a suitable linear formulation. We demonstrate our technique on two recent algorithms for unconstrained submodular maximization and for maximizing submodular function subject to a cardinality constraint. In particular, for unconstrained submodular maximization we obtain an optimal deterministic 1/2-approximation showing that randomization is unnecessary for obtaining optimal results for this setting.

Niv Buchbinder
Statistics and Operations Research Dept.
Tel Aviv University, Israel
niv.buchbinder@gmail.com

Moran Feldman
EPFL
moran.feldman@epfl.ch

CP7

Exact and Approximation Algorithms for Weighted Matroid Intersection

In this paper, we propose new exact and approximation algorithms for the weighted matroid intersection problem. Our exact algorithm is faster than previous algorithms when the largest weight is relatively small. Our approximation algorithm delivers a $(1-\epsilon)$-approximate solution with a running time significantly faster than known exact algorithms. The core of our algorithms is a decomposition technique: we decompose an instance of the weighted matroid intersection problem into a set of instances of the unweighted matroid intersection problem. The computational advantage of this approach is that we can make use of fast unweighted matroid intersection algorithms as a black box for designing algorithms. Precisely speaking, we prove that we can solve the weighted matroid intersection problem via solving W instances of the unweighted matroid intersection problem, where W is the largest given weight. Furthermore, we can find a $(1-\epsilon)$-approximate solution via solving $O(\epsilon^{-1} \log r)$ instances of the unweighted matroid intersection problem, where r is the smallest rank of the given two matroids. Our algorithms are simple and flexible: they can be adapted to special cases of the weighted matroid intersection problem.

Chien-Chung Huang
Chalmers University
villars@gmail.com

Naonori Kakimura
University of Tokyo
kakimura@global.c.u-tokyo.ac.jp

Naoyuki Kamiyama
Kyushu University
kamiyama@imi.kyushu-u.ac.jp

CP7

A Fast Approximation for Maximum Weight Matroid Intersection

Given two matroids defined over a common ground set of n elements, let k be the rank of the matroid intersection and let Q denote the cost of an independence query for either matroid. We present a $(1-\epsilon)$ approximation algorithm for the maximum weight matroid intersection problem with running time $O(nk \log^2(1/\epsilon)Q/\epsilon^2)$.

Chandra Chekuri, Kent Quanrud
University of Illinois at Urbana-Champaign
chekuri@illinois.edu, quanrud2@illinois.edu

CP7

Locally Adaptive Optimization: Adaptive Seeding for Monotone Submodular Functions

The Adaptive Seeding problem is an algorithmic challenge motivated by influence maximization in social networks: One seeks to select among certain accessible nodes in a network, and then select, adaptively, among neighbors of those nodes as they become accessible in order to maximize a global objective function. More generally, adaptive seeding is a stochastic optimization framework where the choices in the first stage affect the realizations in the second stage, over which we aim to optimize. Our main result is a $(1-1/e)^2$-approximation for the adaptive seeding problem for any monotone submodular function. While adaptive policies are often approximated via non-adaptive policies, our algorithm is based on a novel method we call *locally-adaptive* policies. These policies combine a non-adaptive global structure, with local adaptive optimizations. This method enables the $(1-1/e)^2$-approximation for general monotone submodular functions and circumvents some of the impossibilities associated with non-adaptive policies. We also introduce a fundamental problem in submodular optimization that may be of independent interest: given a ground set of elements where every element appears with some small probability, find a set of expected size at most k that has the highest expected value over the realization of the elements. We show a surprising result: there are classes of monotone submodular functions (including coverage) that can be approximated almost optimally as the probability vanishes. For general monotone submodular functions we show via a reduction from Planted-Clique that approximations for this problem are not likely to be obtainable. This optimization problem is an important tool for adaptive seeding via non-adaptive policies, and its hardness motivates the introduction of *locally-adaptive* policies we use in the main result.

Ashwinkumar Badanidiyuru
Google
ashwinkumarbv@gmail.com

Christos Papadimitriou, Aviad Rubinstein
UC Berkeley
christos@cs.berkeley.edu, aviad@cs.berkeley.edu

Yaron Singer, Lior Seeman
Harvard University

yaron@seas.harvard.edu, lior.seeman@gmail.com

CP7

Improved Approximation Algorithms for k-Submodular Function Maximization

This paper presents a polynomial-time 1/2-approximation algorithm for maximizing nonnegative k-submodular functions. This improves upon the previous $\max\{1/3, 1/(1+a)\}$-approximation by Ward and Živný, where $a = \max\{1, \sqrt{(k-1)/4}\}$. We also show that for monotone k-submodular functions there is a polynomial-time $k/(2k-1)$-approximation algorithm while for any $\varepsilon > 0$ a $((k+1)/2k+\varepsilon)$-approximation algorithm for maximizing monotone k-submodular functions would require exponentially many queries. In particular, our hardness result implies that our algorithms are asymptotically tight. We also extend the approach to provide constant factor approximation algorithms for maximizing skew-bisubmodular functions, which were recently introduced as generalizations of bisubmodular functions.

Satoru Iwata
University of Tokyo
iwata@mist.i.u-tokyo.ac.jp

Shin-ichi Tanigawa
Kyoto University
tanigawa@kurims.kyoto-u.ac.jp

Yuichi Yoshida
National Institute of Informatics
Preferred Infrastructure
yyoshida@nii.ac.jp

CP8

Random-Cluster Model in Z^2

The random-cluster model has been widely studied as a unifying framework for random graphs, spin systems and electrical networks, but its dynamics have so far largely resisted analysis. In this paper we analyze the Glauber dynamics of the random-cluster model in the canonical case where the underlying graph is an $n \times n$ box in the Cartesian lattice Z^2. Our main result is a $O(n^2 \log n)$ upper bound for the mixing time at all values of the model parameter p except the critical point $p = p_c(q)$, and for all values of the second model parameter $q \geq 1$. We also provide a matching lower bound proving that our result is tight. Our analysis takes as its starting point the recent breakthrough by Beffara and Duminil-Copin on the location of the random-cluster phase transition in Z^2. It is reminiscent of similar results for spin systems such as the Ising and Potts models, but requires the reworking of several standard tools in the context of the random-cluster model, which is not a spin system in the usual sense.

Antonio Blanca
U.C. Berkeley
ablanca@berkeley.edu

Alistair Sinclair
University of California, Berkeley
sinclair@cs.berkeley.edu

CP8

The Complexity of Approximately Counting in 2-Spin Systems on k-Uniform Bounded-Degree Hypergraphs

An important development in approximate counting is the classification of antiferromagnetic 2-spin systems on bounded-degree graphs, based on a beautiful connection to the uniqueness phase transition. For 2-spin models on *hypergraphs*, the connection between uniqueness and counting breaks down. Nevertheless, we show that for every non-trivial symmetric k-ary Boolean function f, for Δ sufficiently large, it is NP-hard to approximate the partition function of the 2-spin system associated with f on k-uniform hypergraphs of max degree Δ, even within an exponential factor.

Andreas Galanis, Leslie Ann Goldberg
University of Oxford
andreas.galanis@cs.ox.ac.uk, leslie.goldberg@cs.ox.ac.uk

CP8

Bounds for Random Constraint Satisfaction Problems Via Spatial Coupling

We report on a novel technique called *spatial coupling* and its application in the analysis of random constraint satisfaction problems (CSP). Spatial coupling was invented as an engineering construction in the area of error correcting codes where it has resulted in efficient capacity-achieving codes for a wide range of channels. However, this technique is not limited to problems in communications, and can be applied in the much broader context of graphical models. We describe here a general methodology for applying spatial coupling to random constraint satisfaction problems and obtain lower bounds for their rough satisfiability threshold, defined as the largest constraint density for which all but a vanishing fraction of the constraints can be satisfied. The main idea is to construct a distribution of *geometrically structured* random k-SAT instances - namely the spatially coupled ensemble - which has the same rough satisfiability threshold, and is at the same time algorithmically easier to solve. Then by running well-known algorithms on the spatially coupled ensemble we obtain a lower bound on the rough satisfiability threshold of the original ensemble. The method is versatile because one can choose the CSP, there is a certain amount of freedom in the construction of the bigger ensemble, and also in the choice of the algorithm. In this work we focus on random k-SAT but we have also checked that the method is successful for graph coloring, NAE-SAT and XOR-SAT. We choose Unit Clause propagation for the algorithm which is analyzed over the spatially coupled instances. For $k = 3$, for instance, our lower bound is equal to 3.67 which is better than the current bounds in the literature. Similarly, for graph 3-colorability we get a bound of 4.44 which is also better than the current bounds in the literature.

Dimitris Achlioptas
University of California, Santa Cruz
optas@cs.ucsc.edu

Hamed Hassani, Nicolas Macris, Rudiger Urbanke
EPFL
hamed@inf.ethz.ch, nicolas.macris@epfl.ch, rudiger.urbanke@epfl.ch

CP8

Canonical Paths for Mcmc: from Art to Science

Markov Chain Monte Carlo (MCMC) method is a widely

used algorithm design scheme with many applications. To make efficient use of this method, the key step is to prove that the Markov chain is rapid mixing. Canonical paths is one of the two main tools to prove rapid mixing. However, there are much fewer success examples comparing to coupling, the other main tool. The main reason is that there is no systematic approach or general recipe to design canonical paths. Building up on a previous exploration by McQuillan [**?**], we develop a general theory to design canonical paths for MCMC: We reduce the task of designing canonical paths to solving a set of linear equations, which can be automatically done even by a machine. Making use of this general approach, we obtain fully polynomial-time randomized approximation schemes (FPRAS) for counting the number of b-matching with $b \leq 7$ and b-edge-cover with $b \leq 2$. They are natural generalizations of matchings and edge covers for graphs. No polynomial time approximation was previously known for these problems.

Lingxiao Huang
IIIS, Tsinghua University
huanglingxiao1990@126.com

Pinyan Lu
Microsoft Research Asia
pinyanl@microsoft.com

Chihao Zhang
Shanghai Jiaotong University
chihao.zhang@gmail.com

CP8

Evolutionary Dynamics in Finite Populations Mix Rapidly

In this paper we prove that the mixing time of a broad class of evolutionary dynamics in finite, unstructured populations is roughly logarithmic in the size of the state space. An important special case of such a stochastic process is the Wright-Fisher model from evolutionary biology (with selection and mutation) on a population of size N over m genotypes. Our main result implies that the mixing time of this process is $O(\log N)$ for all mutation rates and fitness landscapes, and solves the main open problem from [N. Dixit, P. Srivastava and N. K. Vishnoi, A finite population model of molecular evolution: Theory and computation]. In particular, it significantly extends the main result in [N. K. Vishnoi, The speed of evolution] who proved this for $m = 2$. Biologically, such models have been used to study the evolution of viral populations with applications to drug design strategies countering them. Here the time it takes for the population to reach a steady state is important both for the estimation of the steady-state structure of the population as well in the modeling of the treatment strength and duration. Our result, that such populations exhibit rapid mixing, makes both of these approaches sound. Technically, we make a novel connection between Markov chains arising in evolutionary dynamics and dynamical systems on the probability simplex. This allows us to use the local and global stability properties of the fixed points of such dynamical systems to construct a contractive coupling in a fairly general setting. We expect that our mixing time result would be useful beyond the evolutionary biology setting, and the techniques used here would find applications in bounding the mixing times of Markov chains which have a natural underlying dynamical system.

Ioannis Panageas
Georgia Institute of Technology
panageasj@gmail.com

Piyush Srivastava
California Institute of Technology
piyushs@caltech.edu

Nisheeth K. Vishnoi
EPFL
EPFL
nisheeth.vishnoi@epfl.ch

CP9

Sparse Approximation Via Generating Point Sets

For a set P of n points in the unit ball $\mathsf{b} \subseteq \mathbb{R}^d$, consider the problem of finding a small subset $T \subseteq P$ such that its convex-hull ε-approximates the convex-hull of the original set. Specifically, the Hausdorff distance between the convex hull of T and the convex hull of P should be at most ε. We present an efficient algorithm to compute such an ε'-approximation of size k_{alg}, where ε' is a function of ε, and k_{alg} is a function of the minimum size k_{opt} of such an ε-approximation. Surprisingly, there is no dependence on the dimension d in either of the bounds. Furthermore, every point of P can be ε-approximated by a convex-combination of points of T that is $O(1/\varepsilon^2)$-sparse. Our result can be viewed as a method for sparse, convex autoencoding: approximately representing the data in a compact way using sparse combinations of a small subset T of the original data. The new algorithm can be kernelized, and it preserves sparsity in the original input.

Avrim Blum
Carnegie Mellon University
avrim@cs.cmu.edu

Sariel Har-Peled
Department of Computer Science
University of Illinois, Urbana Champaign
sariel@illinois.edu

Benjamin A. Raichel
University of Texas at Dallas
benjamin.raichel@utdallas.edu

CP9

Towards Optimal Algorithms for Prediction with Expert Advice

We study the classical problem of prediction with expert advice in the adversarial setting with a geometric stopping time. In 1965, Cover gave the optimal algorithm for the case of 2 experts. In this paper, we design the optimal algorithm, adversary and regret for the case of 3 experts. Further, we show that the optimal algorithm for 2 and 3 experts is a probability matching algorithm (analogous to Thompson sampling) against a particular randomized adversary.

Nick Gravin
Microsoft Research
ngravin@gmail.com

Yuval Peres
Microsoft Research, Redmond
peres@microsoft.com

Balasubramanian Sivan
Google Research

balu2901@gmail.com

CP9

Non-Convex Compressed Sensing with the Sum-of-Squares Method

In this paper, we consider stable signal recovery in ℓ_q quasi-norm for $0 < q \leq 1$. In this problem, given a measurement vector $\boldsymbol{y} = A\boldsymbol{x}$ for some unknown signal vector $\boldsymbol{x} \in \mathbb{R}^n$ and a known matrix $A \in \mathbb{R}^{m \times n}$, we want to recover $\boldsymbol{z} \in \mathbb{R}^n$ with $\|\boldsymbol{x}-\boldsymbol{z}\|_q = O(\|\boldsymbol{x}-\boldsymbol{x}^*\|_q)$ from a measurement vector, where $\boldsymbol{x}^*$ is the s-sparse vector closest to $\boldsymbol{x}$ in ℓ_q quasi-norm. Although smaller q is favorable for measuring the distance to sparse vectors, previous methods for $q < 1$ involve ℓ_q quasi-norm minimization and thus are intractable. In this paper, we overcome this issue by using the sum-of-squares method, and give the first polynomial-time stable recovery scheme for a large class of A in ℓ_q quasi-norm for any fixed constant $0 < q \leq 1$.

Tasuku Soma
university of Tokyo
tasuku_soma@mist.i.u-tokyo.ac.jp

Yuichi Yoshida
National Institute of Informatics
Preferred Infrastructure
yyoshida@nii.ac.jp

CP9

Jointly Private Convex Programming

We present a general method for approximately solving convex programs defined by private information from agents, when the solution can be naturally partitioned among the agents. This class of problems includes multi-commodity flow problems, general allocation problems, and multi-dimensional knapsack problems, among other examples. The accuracy of our algorithm depends on the number of *coupling constraints*, which bind multiple agents. On the other hand, our accuracy is nearly independent of the number of variables, and in many cases, actually improves as the number of agents increases. A special case of our result (solving general allocation problems beyond "Gross Substitute' preferences) resolves the main open problem from [Hsu et al. STOC 2014]. We also consider strategic agents who have preferences over their part of the solution. For any convex program in our class that maximizes *social welfare*, we show how to create an *approximately dominant strategy truthful* mechanism, approximately maximizing welfare. The central idea is to charge agents prices based on the approximately optimal dual variables, which are themselves computed under differential privacy. Our results substantially broaden the class of problems that are known to be solvable under privacy and/or incentive constraints.

Justin Hsu
University of Pennsylvania
justhsu@cis.upenn.edu

Zhiyi Huang
The University of Hong Kong
zhiyi@cs.hku.hk

Aaron Roth
University of Pennsylvania, USA
aaroth@cis.upenn.edu

Steven Wu
University of Pennsylvania
wuzhiwei@cis.upenn.edu

CP9

Weighted SGD for ℓ_p Regression with Randomized Preconditioning

In recent years, stochastic gradient descent (SGD) methods and randomized linear algebra (RLA) algorithms have been applied to many large-scale problems in machine learning and data analysis. We aim to bridge the gap between these two classes of methods in solving constrained overdetermined linear regression problems—e.g., ℓ_2 and ℓ_1 regression problems.

- We propose a hybrid algorithm named pwSGD that uses RLA techniques for preconditioning and constructing an importance sampling distribution, and then performs an SGD-like iterative process with weighted sampling on the preconditioned system.
- By rewriting the ℓ_p regression problem into a stochastic optimization problem, we connect pwSGD to several existing ℓ_p solvers including RLA methods with algorithmic leveraging (RLA for short).
- Particularly, when solving ℓ_1 regression with size n by d, pwSGD returns an approximate solution with ϵ relative error on the objective value in $\mathcal{O}(\log n \cdot \text{nnz}(A) + \text{poly}(d)/\epsilon^2)$ time. This complexity is *uniformly* better than that of RLA methods in terms of both ϵ and d when the problem is unconstrained.
- For ℓ_2 regression, pwSGD returns an approximate solution with ϵ relative error on the objective value and solution vector in prediction norm in $\mathcal{O}(\log n \cdot \text{nnz}(A) + \text{poly}(d)\log(1/\epsilon)/\epsilon)$ time. We show that when solving unconstrained ℓ_2 regression, this complexity is comparable to that of RLA and is asymptotically better over several state-of-the-art solvers in the regime where the desired accuracy ϵ, high dimension n and low dimension d satisfy $d \geq 1/\epsilon$ and $n \geq d^2/\epsilon$.

Yinlam Chow
ICME
Stanford University
ychow.stanford.edu

Michael Mahoney
UC Berkeley
mmahoney@stat.berkeley.edu

Christopher Re
Department of Computer Science
Stanford University
chrismre@cs.stanford.edu

Jiyan Yang
ICME
Stanford University
jiyan@stanford.edu

CP10

Time Vs. Information Tradeoffs For Leader Election in Anonymous Trees

In anonymous leader election, every node of the network must output a simple path so that all these paths end at a common node: the leader. We establish tradeoffs between the number of communication rounds and the amount of

information that has to be given *a priori* to the nodes so that they can deterministically solve leader election in anonymous trees. This is achieved using the *algorithms with advice* framework.

Christian Glacet
CNR - IEIIT
christian.glacet@gmail.com

Avery Miller
Tel Aviv University
avery@averymiller.ca

Andrzej Pelc
Département dinformatique
Université du Québec en Outaouais
andrzej.pelc@uqo.ca

CP10

Balanced Allocation: Patience is not a Virtue

We present a load-balancing algorithm, FirstDiff$[d]$, that combines the simplicity of Azar et al.'s Greedy$[d]$ and the improved balance of Vöcking's Left$[d]$. Using the standard balls-and-bins framework, we show that FirstDiff$[d]$ requires d probes on average per ball and closely matches the maximum load ensured by Left$[d]$ in both the standard and heavily-loaded settings. Unlike Left$[d]$, FirstDiff$[d]$ uses no additional structure on the bins. Experimentally, we show that FirstDiff$[d]$ equals Left$[d]$ in practice.

John Augustine
Indian Institute Of Technology Madras
j.e.augustine@gmail.com

William K. Moses Jr.
Department of Computer Science & Engineering
Indian Institute of Technology Madras
wkmjr3@gmail.com

Amanda Redlich
Bowdoin College
aredlich@bowdoin.edu

Eli Upfal
Department of Computer Science
Brown University
eli@cs.brown.edu

CP10

Stabilizing Consensus with Many Opinions

We consider the following distributed consensus problem: Each node in a complete communication network of size n initially holds an *opinion*, which is chosen arbitrarily from a finite set Σ. The system must converge toward a consensus state in which all, or almost all nodes, hold the same opinion. Moreover, this opinion should be *valid*, i.e., it should be one among those initially present in the system. This condition should be met even in the presence of a malicious adversary who can modify the opinions of a bounded subset of nodes, adaptively chosen in every round. We consider the *3-majority dynamics*: At every round, every node pulls the opinion from three random neighbors and sets his new opinion to the majority one (ties are broken arbitrarily). Let k be the number of valid opinions. We show that, if $k \leq n^{\alpha}$, where α is a suitable positive constant, the 3-majority dynamics converges in time polynomial in k and $\log n$ with high probability even in the presence of an adversary who can affect up to $o(\sqrt{n})$ nodes at each round. Previously, the convergence of the 3-majority protocol was known for $|\Sigma| = 2$ only, with an argument that is robust to adversarial errors. On the other hand, no anonymous, uniform-gossip protocol that is robust to adversarial errors was known for $|\Sigma| > 2$.

Luca Becchetti
Sapienza University of Rome
becchett@dis.uniroma1.it

Andrea Clementi
University of Rome Tor Vergata
clementi@mat.uniroma2.it

Emanuele Natale
Sapienza University of Rome
natale@di.uniroma1.it

Francesco Pasquale
Università di Roma Tor Vergata
pasquale@mat.uniroma2.it

Luca Trevisan
UC Berkeley
luca@berkeley.edu

CP10

Markovian Hitters and the Complexity of Blind Rendezvous

We define and construct a novel pseudorandom tool, the *Markovian hitter*. Given an input sequence of n independent random bits, a Markovian hitter produces a sequence of pseudorandom samples in $\{0,1\}^k$, in an online fashion, that hits any subset $W \subset \{0,1\}^k$ of size $\epsilon 2^k$ with probability $\approx 1 - 2^{-(n-k)\epsilon}$. This is comparable to the behavior of truly random samples or classical pseudorandom hitting sets. A Markovian hitter has an additional "Markovian' property of interest: each pseudorandom sample is a function of only the $O(k)$ most recent bits of the input sequence (of random bits). In particular, we apply Markovian hitters to obtain a new algorithm for the well-studied blind rendezvous problem for cognitive radios. This is the problem faced by two parties equipped with radios that can access channels in potentially different subsets, S_1 and S_2, of a universe of n channels. Their challenge is to discover each other (by tuning their radios to the same channel at the same time) as quickly as possible. In prior work [**?**] it was shown that *deterministic* schedules have a lower bound for rendezvous time of $\Omega(|S_1| \cdot |S_2|)$. We beat this quadratic barrier by utilizing a public source of randomness in conjunction with a Markovian hitter to achieve rendezvous in expected time

$$O\left(\log n + \frac{|S_1 \cup S_2|}{|S_1 \cap S_2|}\right).$$

We counterbalance this result by establishing two lower bounds on expected rendezvous time: an

$$\Omega\left(\frac{|S_1 \cup S_2|}{|S_1 \cap S_2|}\right)$$

bound for the setting with public randomness, and an $\Omega(|S_1| \cdot |S_2|)$ bound in the setting with private randomness but no public randomness, which is a strengthening of the result for deterministic schedules.

Sixia Chen

Central Connecticut State University
chensixia09@gmail.com

Matthew Dippel
Northeastern University
mdippel@ccs.neu.edu

Alexander Russell
University of Connecticut
acr@cse.uconn.edu

Abhishek Samanta
Google, Inc
avisheksamantajunior@gmail.com

Ravi Sundaram
Northeastern University
koods@ccs.neu.edu

CP10

How to Scale Exponential Backoff: Constant Throughput, Polylog Access Attempts, and Robustness

Randomized exponential backoff is a widely deployed technique for coordinating access to a shared resource. A good backoff protocol should, arguably, satisfy three natural properties: (i) it should provide constant throughput, wasting as little time as possible; (ii) it should require few failed access attempts, minimizing the amount of wasted effort; and (iii) it should be robust, continuing to work efficiently even if some of the access attempts fail for spurious reasons. Unfortunately, exponential backoff has some well-known limitations in two of these areas: it provides poor (sub-constant) throughput (in the worst case), and is not robust (to adversarial disruption).The goal of this paper is to "fix' exponential backoff by making it scalable, particularly focusing on the case where processes arrive in an on-line, worst-case fashion. We present a relatively simple backoff protocol, Re-Backoff, that has, at its heart, a version of exponential backoff. It guarantees expected constant throughput with dynamic process arrivals and requires only an expected polylogarithmic number of access attempts per process.Re-Backoff is also robust to periods where the shared resource is unavailable for a period of time. If it is unavailable for D time slots, Re-Backoff provides the following guarantees. When the number of packets is a finite n, the average expected number of access attempts for successfully sending a packet is $O(\log^2(n+D))$. In the infinite case, the average expected number of access attempts for successfully sending a packet is $O(\log^2(\eta+D))$ where η is the maximum number of processes that are ever in the system concurrently.

Michael A. Bender
Stony Brook University
bender@cs.stonybrook.edu

Jeremy T. Fineman
Georgetown University
jfineman@cs.georgetown.edu

Seth Gilbert
NUS
seth.gilbert@comp.nus.edu.sg

Maxwell Young
Mississippi State University
myoung@cse.msstate.edu

CP11

On Dynamic Approximate Shortest Paths for Planar Graphs with Worst-Case Costs

Given a base weighted planar graph G_{input} on n nodes and parameters M, ϵ we present a dynamic distance oracle with $1+\epsilon$ stretch and worst case update and query costs of $\epsilon^{-3}M^4 \cdot \text{poly-log}(n)$. We allow arbitrary edge weight updates as long as the shortest path metric induced by the updated graph has stretch of at most M relative to the shortest path metric of the base graph G_{input}. For example, on a planar road network, we can support fast queries and dynamic traffic updates as long as the shortest path from any source to any target (including using arbitrary detours) is between, say, 80 and 3 miles-per-hour. As a warm-up we also prove that graphs of bounded treewidth have exact distance oracles in the dynamic edge model. To the best of our knowledge, this is the first dynamic distance oracle for a non-trivial family of dynamic changes to planar graphs with worst case costs of $o(n^{1/2})$ both for query and for update operations.

Ittai Abraham
Microsoft
iabraham@vmware.com

Shiri Chechik
Tel-Aviv University, Israel
schechik@post.tau.ac.il

Daniel Delling
Unaffiliated
daniel.delling@gmail.com

Andrew Goldberg
Microsoft Research Silicon Valley
andvgol@gmail.com

Renato F. Werneck
Unaffiliated
rwerneck@acm.org

CP11

Faster Fully Dynamic Matchings with Small Approximation Ratios

Maximum cardinality matching is a fundamental algorithmic problem with many algorithms and applications. The fully dynamic version, in which edges are inserted and deleted over time has also been the subject of much attention. Existing algorithms for dynamic matching (in general n-vertex m-edge graphs) fall into two groups: there are fast (mostly randomized) algorithms that achieve a 2-approximation or worse, and there are slow algorithms with $\Omega(\sqrt{m})$ update time that achieve a better-than-2 approximation. Thus the obvious question is whether we can design an algorithm that achieves a tradeoff between these two: a $o(\sqrt{m})$ update time and a better-than-2 approximation simultaneously. We answer this question in the affirmative. Previously, such bounds were only known for the special case of bipartite graphs.
Our main result is a fully dynamic *deterministic* algorithm that maintains a $(3/2+\epsilon)$-approximation in amortized update time $O(m^{1/4}\epsilon^{-2.5})$. In addition to achieving the tradeoff described above, our algorithm manages to be polynomially faster than all existing deterministic algorithms

(excluding an existing $\log n$-approximation of Onak and Rubinfeld), while still maintaining a better-than-2 approximation.
We also give stronger results for graphs whose arboricity is at most α. We show how to maintain a $(1+\epsilon)$-approximate *fractional* matching or a $(3/2+\epsilon)$-approximate *integral* matching in worst-case time $O(\alpha(\alpha+\log n))$ for constant ϵ. When the arboricity is constant, this bound is $O(\log n)$ and when the arboricity is polylogarithmic the update time is also polylogarithmic. Previous results for small arboricity non-bipartite graphs could only maintain a maximal matching (2-approximation).

Aaron Bernstein
Columbia University
bernstei@gmail.com

CP11

Weighted Dynamic Finger in Binary Search Trees

It is shown that the online binary search tree data structure GreedyASS performs asymptotically as well on a sufficiently long sequence of searches as any static binary search tree where each search begins from the previous search (rather than the root). This bound is known to be equivalent to assigning each item i in the search tree a positive weight w_i and bounding the search cost of an item in the search sequence $s_1, \ldots, s_m$ by $O\left(1+\log \frac{\sum_{\min(s_{i-1},s_i)\le x\le \max(s_{i-1},s_i)} w_x}{\min(w_{s_i}, w_{s_{i-1}})}\right)$ amortized. This result is the strongest finger-type bound to be proven for binary search trees. By setting the weights to be equal, one observes that our bound implies the dynamic finger bound. Compared to the previous proof of the dynamic finger bound for Splay trees, our result is significantly shorter, stronger, simpler, and has reasonable constants.

John Iacono
Polytechnic Institute of New York University
jiacono@poly.edu

Stefan Langerman
Université libre de Bruxelles
soda16@slef.org

CP11

Dynamic Dfs in Undirected Graphs: Breaking the $O(m)$ Barrier

Given an undirected graph $G=(V,E)$ on n vertices and m edges, we address the problem of maintaining a DFS tree when the graph is undergoing *updates* (insertion and deletion of vertices or edges). We present the following results for this problem.

1. *Fault tolerant DFS tree*: There exists a data structure of size $\tilde{O}(m)$ such that given any set $\mathcal{F}$ of failed vertices or edges, a DFS tree of the graph $G \setminus \mathcal{F}$ can be reported in $\tilde{O}(n|\mathcal{F}|)$ time.
2. *Fully dynamic DFS tree*: There exists a fully dynamic algorithm for maintaining a DFS tree that takes worst case $\tilde{O}(\sqrt{mn})$ time per update for any arbitrary online sequence of updates.
3. *Incremental DFS tree*: Given any arbitrary online sequence of edge insertions, we can maintain a DFS tree in $\tilde{O}(n)$ worst case time per edge insertion.

These are the first $o(m)$ worst case time results for maintaining a DFS tree in a dynamic environment. Moreover, our fully dynamic algorithm provides, in a seamless manner, the first deterministic algorithm with $O(1)$ query time and $o(m)$ worst case update time for the dynamic subgraph connectivity, biconnectivity, and 2-edge connectivity.

Surender Baswana, Shreejit Chaudhury, Keerti Choudhary, Shahbaz Khan
IIT Kanpur
sbaswana@cse.iitk.ac.in, shreejit.1@gmail.com, keerti@cse.iitk.ac.in, shahbazk@cse.iitk.ac.in

CP11

Dynamic $(1+\epsilon)$-Approximate Matchings: A Density-Sensitive Approach

Approximate matchings in fully dynamic graphs have been intensively studied in recent years. Gupta and Peng [FOCS'13] presented a deterministic algorithm for maintaining fully dynamic $(1+\epsilon)$-approximate maximum cardinality matching (MCM) in general graphs with worst-case update time $O(\sqrt{m}\cdot\epsilon^{-2})$, for any $\epsilon > 0$, where m denotes the current number of edges in the graph. Despite significant research efforts, this $\sqrt{m}$ update time barrier remains the state-of-the-art even if amortized time bounds and randomization are allowed or the approximation factor is allowed to increase from $1+\epsilon$ to $2-\epsilon$, and even in basic graph families such as planar graphs. This paper presents a simple deterministic algorithm whose performance depends on the *density* of the graph. Specifically, we maintain fully dynamic $(1+\epsilon)$-approximate MCM with worst-case update time $O(\alpha\cdot\epsilon^{-2})$ for graphs with *arboricity*[1] bounded by α. Since the arboricity ranges between 1 and $\sqrt{m}$, our density-sensitive bound $O(\alpha\cdot\epsilon^{-2})$ naturally generalizes the $O(\sqrt{m}\cdot\epsilon^{-2})$ bound of Gupta and Peng. For the family of bounded arboricity graphs (which includes forests, planar graphs, and graphs excluding a fixed minor), in the regime $\epsilon = O(1)$ our update time reduces to a constant. This should be contrasted with the previous best 2-approximation results for bounded arboricity graphs, which achieve either an $O(\log n)$ worst-case bound (Kopelowitz et al., ICALP'14) or an $O(\sqrt{\log n})$ amortized bound (He et al., ISAAC'14), where n stands for the number of vertices in the graph.

David Peleg, Shay Solomon
Weizmann Institute of Science
david.peleg@weizmann.ac.il, solo.shay@gmail.com

CP12

Reducing Curse of Dimensionality: Improved PTAS for TSP (with Neighborhoods) in Doubling Metrics

We consider the Traveling Salesman Problem with Neighborhoods (TSPN) in doubling metrics. The goal is to find a shortest tour that visits each of a given collection of subsets (regions or neighborhoods) in the underlying metric space. We give a randomized polynomial time approximation scheme (PTAS) when the regions are fat weakly disjoint. Moreover, more refined procedures are used to improve the dependence of the running time on the doubling dimension.

T-H. Hubert Chan, Shaofeng H.-C. Jiang

[1] The arboricity of a graph is the minimum number of edge-disjoint forests into which it can be partitioned, and it is close to the density of its densest subgraph.

The University of Hong Kong
hubert@cs.hku.hk, sfjiang@cs.hku.hk

CP12

Clustering Time Series under the Frechet Distance

The Fréchet distance is a popular distance measure for curves. We study the problem of clustering time series under the Fréchet distance. In particular, we give $(1+\varepsilon)$-approximation algorithms for variations of the following problem with parameters k and ℓ. Given n univariate time series P, each of complexity at most m, we find k time series, not necessarily from P, which we call cluster centers and which each have complexity at most ℓ, such that (a) the maximum distance of an element of P to its nearest cluster center or (b) the sum of these distances is minimized. Our algorithms have running time near-linear in the input size for constant ε, k and ℓ. To the best of our knowledge, our algorithms are the first clustering algorithms for the Fréchet distance which achieve an approximation factor of $(1+\varepsilon)$ or better.

Anne Driemel
TU Eindhoven, The Netherlands
Department of Mathematics and Computer Science
a.driemel@tue.nl

Amer Krivosija
TU Dortmund, Department of Computer Science
amer.krivosija@tu-dortmund.de

Christian Sohler
Department of Computer Science
Technische Universitat Dortmund
christian.sohler@tu-dortmund.de

CP12

Approximating Capacitated K-Median with $(1+\epsilon)k$ Open Facilities

In this paper, we give the first constant approximation for the capacitated k-median problem, that only violates the cardinality constraint by a factor of 1+eps. This generalizes the result of [Li15], which only works for the case when all capacities are the same. Our algorithm is based on a novel configuration LP relaxation for the problem.

Shi Li
Toyota Technological Institute at Chicago
shil@buffalo.edu

CP12

Simple and Fast Rounding Algorithms for Directed and Node-Weighted Multiway Cut

In Multiway Cut problem, input is an edge/node-weighted graph and a set of terminals $S = \{s_1, \ldots, s_k\}$, and the goal is to remove minimum weight set of edges to disconnect all terminals. We present a very simple 2-approximation for Directed Multiway Cut and $2(1-1/k)$ approximation for Node-weighted Multiway Cut.

Chandra Chekuri
University of Illinois at Urbana-Champaign
chekuri@illinois.edu

Vivek Madan
University of illinois Urbana-Champaign
vmadan2@illinois.edu

CP12

Constant Factor Approximation for Subset Feedback Set Problems Via a New Lp Relaxation

In the Subset Feedback Edge/Vertex Set problem (SFES/SFVS) the input is edge/node-weighted graph $G = (V, E)$ and a set $S = \{s_1, \ldots, s_k\} \subset V$ of k terminals and the goal is to remove a minimum weight set of edges/nodes such that no interesting containing a terminal remains. One reason for the difficulty in addressing these problems has been the lack of LP relaxations with constant factor integrality gaps. We give first LP relaxation for SFVS and SFES with constant integrality gap.

Chandra Chekuri
University of Illinois at Urbana-Champaign
chekuri@illinois.edu

Vivek Madan
University of illinois Urbana-Champaign
vmadan2@illinois.edu

CP13

Error Amplification for Pairwise Spanner Lower Bounds

A *pairwise spanner* of a graph $G = (V, E)$ and a "pair set' $P \subseteq V \times V$ is a subgraph H that preserves all pairwise distances in P, up to some additive error term $+\beta$. When $\beta = 0$ the object is called a *pairwise distance preserver*. A large and growing body of work has considered upper bounds for these objects, but lower bounds have been elusive. The only known lower bound results are (1) Coppersmith and Elkin (SODA'05) against preservers, and (2) considerably weaker bounds by Woodruff (FOCS'06) against spanners. Our main result is an amplification theorem: we prove that lower bounds against pairwise *distance preservers* imply lower bounds against pairwise *spanners*. In other words, to prove lower bounds against *any* constant error spanners, it is enough to consider only subgraphs that are not allowed any error at all! We apply this theorem to obtain drastically improved lower bounds. Some of these include:

- Linear size pairwise spanners with up to $+(2k-1)$ error cannot span $|P| = \omega(n^{(1+k)/(3+k)})$ pairs. This is a large improvement over Woodruff's $|P| = \omega(n^{2-2/k})$ ($|P|$ is now linear, rather than quadratic, as k gets large).
- $|E(H)| = \Omega(n^{1+1/k})$ edges are required for a $+(2k-1)$ spanner of $|P| = \Omega(n^{1+1/k})$ pairs - another large improvement over Woodruff's $|P| = \Omega(n^2)$.
- The first tight bounds for pairwise spanners: for $+2$ error and $P = \Theta(n^{3/2})$ we show that $\Theta(n^{3/2})$ edges are necessary and sufficient (this also reflects a new *upper* bound: we construct $+2$ pairwise spanners on $O(n|P|^{1/3})$ edges, removing a log factor from a prior algorithm).

We also show improved lower bounds against subset spanners (where $P = S \times S$ for some node subset S), and lower bounds against D threshold spanners (where P is the set of node pairs at distance at least D).

Amir Abboud, Greg Bodwin
Stanford University

abboud@cs.stanford.edu, gbodwin@cs.stanford.edu

CP13

Better Distance Preservers and Additive Spanners

We make improvements to the upper bounds on several popular types of distance preserving graph sketches. The first part of our paper concerns *pairwise distance preservers*, which are sparse subgraphs that exactly preserve the pairwise distances for a set of given pairs of vertices. Our main result here is that all unweighted, undirected n-node graphs G and all pair sets P have distance preservers on $|H| = O(n^{2/3}|P|^{2/3} + n|P|^{1/3})$ edges. This improves the known bounds whenever $|P| = \omega(n^{3/4})$. We then develop a new graph clustering technique, based on distance preservers, and we apply this technique to show new upper bounds for additive *(standard) spanners*, in which *all* pairwise distances must be preserved up to an additive error function, and for *subset spanners*, in which only distances within a given node subset must be preserved up to an error function. For both of these objects, we obtain the new best tradeoff between spanner sparsity and error allowance in the regime where the error is polynomial in the graph size. We leave open a conjecture that $O(n^{2/3}|P|^{2/3} + n)$ pairwise distance preservers are possible for undirected unweighted graphs. Resolving this conjecture in the affirmative would improve and simplify our upper bounds for all the graph sketches mentioned above.

Greg Bodwin, Virginia Williams
Stanford University
gbodwin@cs.stanford.edu, virgi@cs.stanford.edu

CP13

Near-Optimal Light Spanners

In this work we consider spanners of light weight for general undirected graphs. For any n-vertex graph G and any positive integer k, we present a spanner with $(2k-1)\cdot(1+\epsilon)$ stretch, $O_\epsilon(\omega(MST(G))n^{1/k})$ weight, and $O_\epsilon(n^{1+1/k})$ edges. Up to a $(1+\epsilon)$ factor in the stretch this matches the girth conjecture of Erdős. It improves over the recent work of Elkin, Neiman and Solomon [ICALP 14] by reducing the spanner weight by a factor of $k/\log k$.

Shiri Chechik
Tel-Aviv University, Israel
schechik@post.tau.ac.il

Christian Wulff-Nilsen
Department of Computer Science
University of Copenhagen
koolooz@di.ku.dk

CP13

On Notions of Distortion and An Almost Minimum Spanning Tree with Constant Average Distortion

Minimum Spanning Trees of weighted graphs are fundamental objects in numerous applications. In particular in distributed networks, the minimum spanning tree of the network is often used to route messages between network nodes. Unfortunately, while being most efficient in the total cost of connecting all nodes, minimum spanning trees fail miserably in the desired property of approximately preserving distances between pairs. While known lower bounds exclude the possibility of the worst case distortion of a tree being small, it was shown in [ABN15] that there exists a spanning tree with constant average distortion. Yet, the weight of such a tree may be significantly larger than that of the MST. In this paper, we show that any weighted undirected graph admits a *spanning tree* whose weight is at most $(1+\rho)$ times that of the MST, providing *constant average distortion* $O(1/\rho^2)$. The constant average distortion bound is implied by a stronger property of *scaling distortion*, i.e., improved distortion for smaller fractions of the pairs. The result is achieved by first showing the existence of a low weight *spanner* with small *prioritized distortion*, a property allowing to prioritize the nodes whose associated distortions will be improved. We show that prioritized distortion is essentially equivalent to coarse scaling distortion via a general transformation, which has further implications and may be of independent interest. In particular, we obtain an embedding for arbitrary metrics into Euclidean space with optimal prioritized distortion.

Yair Bartal
Hebrew University
yair@cs.huji.ac.il

Arnold Filtser
Ben gurion university of the negev
arnoldf@cs.bgu.ac.il

Ofer Neiman
Ben-Gurion University of the Negev
neimano@cs.bgu.ac.il

CP13

Approximating Low-Stretch Spanners

Despite significant recent progress on approximating graph spanners (subgraphs which approximately preserve distances), there are still several large gaps in our understanding. We give new results for two of them: approximating basic k-spanner (particularly for small k), and the dependence on f when approximating f-fault tolerant spanners. We first design an $\tilde{O}(n^{1/3})$-approximation for 4-spanner (both basic and directed). This was the last value of k for which only an $O(\sqrt{n})$-approximation was known for basic k-spanner, and thus implies that for any k the approximation ratio is at most $\tilde{O}(n^{1/3})$. For basic k-spanner, we also show an integrality gap for the natural flow-based LP (the main tool in almost all nontrivial spanner approximations) which nearly matches the trivial approximation of $n^{\frac{1}{\lfloor (k+1)/2 \rfloor}}$. For f-fault tolerant spanners, we show that in the small-stretch setting ($k \in \{3,4\}$) it is possible to entirely remove the dependence on f from the approximation ratio, at the cost of moving to bicriteria guarantees. The previous best dependence on f was either almost-linear (in the undirected setting) or exponential (in the directed setting for stretch 4).

Michael Dinitz
Computer Science,
mdinitz@cs.jhu.edu

Zeyu Zhang
Johns Hopkins University

zyzhang92@gmail.com

CP14

Permutation Patterns Are Hard to Count

Abstract not available at time of publication.

Igor Pak, Scott Garrabrant
UCLA
pak@math.ucla.edu, scott.garrabrant@gmail.com

CP14

Efficient Quantum Algorithms for (Gapped) Group Testing and Junta Testing

In the k-junta testing problem, a tester has to efficiently decide whether a given function $f : \{0,1\}^n \to \{0,1\}$ is a k-junta (i.e., depends on at most k of its input bits) or is ε-far from any k-junta. Our main result is a quantum algorithm for this problem with query complexity $\tilde{O}(\sqrt{k/\varepsilon})$ and time complexity $\tilde{O}(n\sqrt{k/\varepsilon})$. This quadratically improves over the query complexity of the previous best quantum junta tester, due to Atıcı and Servedio. Our tester is based on a new quantum algorithm for a gapped version of the combinatorial group testing problem, with an up to quartic improvement over the query complexity of the best classical algorithm. For our upper bound on the *time* complexity we give a near-linear time implementation of a shallow variant of the quantum Fourier transform over the symmetric group, similar to the Schur-Weyl transform. We also prove a lower bound of $\Omega(k^{1/3})$ queries for junta-testing (for constant ε).

Andris Ambainis
Faculty of Computing
University of Latvia
ambainis@lu.lv

Aleksandrs Belovs
CWI
stiboh@inbox.lv

Ronald de Wolf
CWI and University of Amsterdam
rdewolf@cwi.nl

Oded Regev
New York University
regev@cims.nyu.edu

CP14

Computing in Continuous Space with Self-Assembling Polygonal Tiles

In this paper we investigate the computational power of the polygonal tile assembly model (polygonal TAM) at temperature 1, i.e. in non-cooperative systems. The polygonal TAM is an extension of Winfree's abstract tile assembly model (aTAM) which not only allows for square tiles (as in the aTAM) but also allows for tile shapes which are arbitrary polygons. Although a number of self-assembly results have shown computational universality at temperature 1, these are the first results to do so by fundamentally relying on tile placements in continuous, rather than discrete, space. With the square tiles of the aTAM, it is conjectured that the class of temperature 1 systems is not computationally universal. Here we show that for each $n > 6$, the class of systems whose tiles are the shape of the regular polygon P with n sides is computationally universal. On the other hand, we show that the class of systems whose tiles consist of a regular polygon P with $n \leq 6$ sides cannot compute using any known techniques. In addition, we show a number of classes of systems whose tiles consist of a non-regular polygon with $n \geq 3$ sides are computationally universal.

Jacob Hendricks
University of Wisconsin-River Falls
jacob.hendricks@uwrf.edu

Gilbert Oscar, Matthew Patitz, Trent Rogers
University of Arkansas
oogilbert@uark.edu, mpatitz@self-assembly.net, tar003@uark.edu

CP14

Efficient Quantum Algorithms for Computing Class Groups and Solving the Principal Ideal Problem in Arbitrary Degree Number Fields

This paper gives polynomial time quantum algorithms for computing the ideal class group (CGC) under the Generalized Riemann Hypothesis and solving the principal ideal problem (PIP) in number fields of arbitrary degree. These are are fundamental problems in number theory and they are connected to many unproven conjectures in both analytic and algebraic number theory. Previously the best known algorithms by Hallgren only allowed to solve these problems in quantum polynomial time for number fields of constant degree. In a recent breakthrough, Eisentrager et al showed how to compute the unit group in arbitrary fields, thus opening the way to the resolution of CGC and PIP in the general case. The methods we introduce in this paper run in quantum polynomial time in arbitrary classes of number fields. They can be applied to solve other problems in computational number theory as well including computing the ray class group and solving relative norm equations. They are also useful for ongoing cryptanalysis of cryptographic schemes based on ideal lattices.

Jean-Francois Biasse
University of South Florida
biasse@lix.polytechnique.fr

Fang Song
University of Waterloo
University of Waterloo
fang.song@uwaterloo.ca

CP15

Treetopes and Their Graphs

We define treetopes, a generalization of the three-dimensional roofless polyhedra (Halin graphs) to arbitrary dimensions. Like roofless polyhedra, treetopes have a designated base facet such that every face of dimension greater than one intersects the base in more than one point. We prove an equivalent characterization of the 4-treetopes using the concept of clustered planarity from graph drawing, and we use this characterization to recognize the graphs of 4-treetopes in polynomial time. This result provides one of the first classes of 4-polytopes, other than pyramids and stacked polytopes, that can be recognized efficiently from their graphs.

David Eppstein
University of California, Irvine

Department of Computer Science
eppstein@uci.edu

CP15

Multiscale Mapper: Topological Summarization Via Codomain Covers

Summarizing topological information from datasets and maps defined on them is a central theme in topological data analysis. Mapper, a tool for such summarization, takes as input both a possibly high dimensional dataset and a map defined on the data, and produces a summary of the data by using a cover of the codomain of the map. This cover, via a pullback operation to the domain, produces a simplicial complex connecting the data points. The resulting view of the data through a cover of the codomain offers flexibility in analyzing the data. However, it offers only a view at a fixed scale at which the cover is constructed. Inspired by the concept, we explore a notion of a tower of covers which induces a tower of simplicial complexes connected by simplicial maps, which we call *multiscale mapper*. We study the resulting structure, and design practical algorithms to compute its persistence diagrams efficiently. Specifically, when the domain is a simplicial complex and the map is a real-valued piecewise-linear function, the algorithm can compute the exact persistence diagram only from the 1-skeleton of the input complex. For general maps, we present a combinatorial version of the algorithm that acts only on *vertex sets* connected by the 1-skeleton graph, and this algorithm approximates the exact persistence diagram thanks to a stability result that we show to hold.

Tamal Dey
The Ohio State University
tamaldey@cse.ohio-state.edu

Facundo Memoli
Ohio State University
memoli@math.osu.edu

Yusu Wang
The Ohio State University, Columbus, Ohio
yusu@cse.ohio-state.edu

CP15

Beyond the Richter-Thomassen Conjecture

If two closed Jordan curves in the plane have precisely one point in common, then it is called a *touching point*. All other intersection points are called *crossing points*. The main result of this paper is a Crossing Lemma for closed curves: In any family of n pairwise intersecting simple closed curves in the plane, no three of which pass through the same point, the number of crossing points exceeds the number of touching points by a factor of $\Omega((\log\log n)^{1/8})$. As a corollary, we prove the following long-standing conjecture of Richter and Thomassen: The total number of intersection points between any n pairwise intersecting simple closed curves in the plane, no three of which pass through the same point, is at least $(1-o(1))n^2$.

János Pach
EPFL, Lausanne
Rényi Institute, Budapest
pach@cims.nyu.edu

Natan Rubin
Ben Gurion University
rubinnat.ac@gmail.com

Gábor Tardos
Rényi Institute, Budapest
tardos@renyi.hu

CP15

Windrose Planarity: Embedding Graphs with Direction-Constrained Edges

Given a planar graph $G(V,E)$ and an assignment of the neighbours of each vertex $v \in V$ to the four quadrants around v, problem Windrose Planarity asks whether G admits a *windrose-planar drawing*, i.e., a planar drawing in which edges are xy-monotone curves and the neighbours of each vertex v lie in the assigned quadrants. We give a polynomial-time algorithm for Windrose Planarity on embedded graphs, which in the positive case constructs a 1-bend windrose-planar drawing on an $O(n) \times O(n)$ grid.

Patrizio Angelini
Universitaet Tuebingen
angelini@informatik.uni-tuebingen.de

Giordano Da Lozzo
Roma Tre University
dalozzo@dia.uniroma3.it

Giuseppe Di Battista
Dipartimento di Informatica e Automazione,
Roma Tre University
gdb@dia.uniroma3.it

Valentino Di Donato
Roma Tre University
didonato@dia.uniroma3.it

Philipp Kindermann
FernUniversitaet in Hagen
philipp.kindermann@fernuni-hagen.de

Guenter Rote
Freie Universitaet Berlin
rote@inf.fu-berlin.de

Ignaz Rutter
Karlsruhe Institute of Technology
rutter@kit.edu

CP16

An Improved Approximation Guarantee for the Maximum Budgeted Allocation Problem

We study the Maximum Budgeted Allocation problem, which is the problem of assigning indivisible items to players with budget constraints. The best approximation algorithms we know for the MBA problem achieve a 3/4-approximation ratio, and employ the natural Assignment-LP. In this paper, we present a $3/4 + c$-approximation algorithm for MBA, for some constant $c > 0$. This algorithm rounds solutions to the Configuration-LP, therefore also showing that the Configuration-LP is stronger than the Assignment-LP for the MBA problem.

Christos Kalaitzis
EPFL

christos.kalaitzis@epfl.ch

CP16

On the Integrality Gap of Degree-4 Sum of Squares for Planted Clique

The problem of finding large cliques in random graphs and its "planted" variant, where one wants to recover a clique of size $\omega \gg \log(n)$ added to an Erdős-Rényi graph $G \sim G(n, \frac{1}{2})$, have been intensely studied. Nevertheless, existing polynomial time algorithms can only recover planted cliques of size $\omega = \Omega(\sqrt{n})$. By contrast, information theoretically, one can recover planted cliques so long as $\omega \gg \log(n)$. In this work, we continue the investigation of algorithms from the sum of squares hierarchy for solving the planted clique problem begun by Meka, Potechin, and Wigderson [STOC 2015] and Deshpande and Montanari [COLT 2015]. Our main results improve upon both these previous works by showing:

1. Degree four SoS does not recover the planted clique unless $\omega \gg \sqrt{n}/\operatorname{polylog} n$, improving upon the bound $\omega \gg n^{1/3}$ due to Deshpande and Montanari.
2. For $2 < d = o(\sqrt{\log(n)})$, degree $2d$ SoS does not recover the planted clique unless $\omega \gg n^{1/(d+1)}/(2^d \operatorname{polylog} n)$, improving upon the bound due to Meka, Potechin and Wigderson.

Our proof for the second result is based on a fine spectral analysis of the certificate used in the prior works by decomposing it along an appropriately chosen basis. Along the way, we develop combinatorial tools to analyze the spectrum of random matrices with dependent entries and to understand the symmetries in the eigenspaces of the set symmetric matrices inspired by work of Grigoriev. An argument of Kelner shows that the first result cannot be proved using the same certificate. Rather, our proof involves constructing and analyzing a new certificate that yields the nearly tight lower bound by "correcting" the previous certificate.

Samuel Hopkins
Cornell University
samhop@cs.cornell.edu

Pravesh Kothari
Department of Computer Science
University of Texas at Austin
kothari@cs.utexas.edu

Aaron Potechin
Simons Institute
aaronpotechin@gmail.com

Prasad Raghavendra
U C Berkeley
prasad@cs.berkeley.edu

Tselil Schramm
UC Berkeley
tschramm@cs.berkeley.edu

CP16

The Matching Problem Has No Small Symmetric Sdp

Yannakakis showed that the matching problem does not have a small symmetric linear program. Rothvoß recently proved that any, not necessarily symmetric, linear program also has exponential size. It is natural to ask whether the matching problem can be expressed compactly in a framework such as semidefinite programming (SDP) that is more powerful than linear programming but still allows efficient optimization. We answer this question negatively for symmetric SDPs: any symmetric SDP for the matching problem has exponential size. We also show that an $O(k)$-round Lasserre SDP relaxation for the metric traveling salesperson problem yields at least as good an approximation as any symmetric SDP relaxation of size n^k. The key technical ingredient underlying both these results is an upper bound on the degree needed to derive polynomial identities that hold over the space of matchings or traveling salesperson tours.

Gabor Braun
Georgia Tech
gabor.braun@isye.gatech.edu

Jonah Brown-Cohen
UC Berkeley
jonahbc@eecs.berkeley.edu

Arefin Huq, Sebastian Pokutta
Georgia Tech
arefin@gatech.edu, sebastian.pokutta@isye.gatech.edu

Prasad Raghavendra
UC Berkeley
raghavendra@berkeley.edu

Aurko Roy
Georgia Tech
aurko@gatech.edu

Benjamin Weitz
UC Berkeley
bsweitz@cs.berkeley.edu

Daniel Zink
Georgia Tech
daniel.zink@gatech.edu

CP16

Finding a Stable Allocation in Polymatroid Intersection

The stable matching model of Gale and Shapley (1962) has been generalized in various directions such as matroid kernels due to Fleiner (2001) and stable allocations in bipartite networks due to Baïou and Balinski (2002). Unifying these generalizations, we introduce a concept of stable allocations in polymatroid intersection. Our framework includes both integer- and real-variable versions. The integer-variable version corresponds to a special case of the discrete-concave function model due to Eguchi, Fujishige, and Tamura (2003), who established the existence of a stable allocation by showing that a simple extension of the deferred acceptance algorithm of Gale and Shapley finds a stable allocation in pseudo-polynomial time. It has been open to develop a polynomial time algorithm even for our special case. In this paper, we present the first strongly polynomial algorithm for finding a stable allocation in polymatroid intersection. To achieve this, we utilize the augmenting path technique for polymatroid intersection. In each iteration, the algorithm searches for an augmenting path by simulating a chain of proposes and rejects in the deferred acceptance algorithm. The running time of our algorithm is $O(n^3\gamma)$, where n and γ respectively denote the

cardinality of the ground set and the time for computing the saturation and exchange capacities. This is as fast as the best known algorithm for the polymatroid intersection problem.

Satoru Iwata
University of Tokyo
iwata@mist.i.u-tokyo.ac.jp

Yu Yokoi
Department of Mathematical Informatics
University of Tokyo
yu_yokoi@mist.i.u-tokyo.ac.jp

CP16

Online Contention Resolution Schemes

We introduce a new rounding technique designed for online optimization problems, called online contention resolution schemes (OCRSs). OCRSs are applicable to various online selection problems, including Bayesian online selection, oblivious posted pricing mechanisms, and stochastic probing models. Furthermore, they share many properties of offline contention resolution schemes. In particular, OCRSs for different constraint families can be combined to obtain an OCRS for their intersection. Moreover, we can approximately maximize submodular functions in various online settings.

Moran Feldman
The Open University of Isael
moranfe3@gmail.com

Ola Svensson
EPFL
ola.svensson@epfl.ch

Rico Zenklusen
ETH Zurich
ricoz@math.ethz.ch

CP17

Lower Bounds for the Parameterized Complexity of Minimum Fill-In and Other Completion Problems

In this work, we focus on several completion problems for subclasses of chordal graphs: Minimum Fill-In, Interval Completion, Proper Interval Completion, Threshold Completion, and Trivially Perfect Completion. In these problems, the task is to add at most k edges to a given graph in order to obtain a chordal, interval, proper interval, threshold, or trivially perfect graph, respectively. We prove the following lower bounds for all these problems, as well as for the related Chain Completion problem:

- Assuming the Exponential Time Hypothesis, none of these problems can be solved in time $2^{\mathcal{O}(n^{1/2}/\log^c n)}$ or $2^{\mathcal{O}(k^{1/4}/\log^c k)} \cdot n^{\mathcal{O}(1)}$, for some integer c.
- Assuming the non-existence of a subexponential-time approximation scheme for Min Bisection on d-regular graphs, for some constant d, none of these problems can be solved in time $2^{o(n)}$ or $2^{o(\sqrt{k})} \cdot n^{\mathcal{O}(1)}$.

The second result is an evidence, that a significant improvement of the best known algorithms for parameterized completion problems would lead to a surprising breakthrough in the design of approximation algorithms for Min Bisection.

Ivan Bliznets
Institute of Mathematic, St.Petersburg Department of Steklov
iabliznets@gmail.com

Marek Cygan, Pawel Komosa
Institute of Informatics, University of Warsaw
cygan@mimuw.edu.pl, p.komosa@mimuw.edu.pl

Lukas Mach
Computer Science Institute of Charles University (IUUK)
Lesser Town Square, no. 25, Prague
lukas.mach@gmail.com

Michal Pilipczuk
Institute of Informatics, University of Warsaw
michal.pilipczuk@mimuw.edu.pl

CP17

Linear Recognition of Almost Interval Graphs

Let interval $+ kv$, interval $+ ke$, and interval $- ke$ denote the classes of graphs that can be obtained from some interval graph by adding k vertices, adding k edges, and deleting k edges, respectively. When k is small, these graph classes are called almost interval graphs. They are well motivated from computational biology, where the data ought to be represented by an interval graph while we can only expect an almost interval graph for the best. For any fixed k, we give linear-time algorithms for recognizing all these classes, and in the case of membership, our algorithms provide also a specific interval graph as evidence. When k is part of the input, these problems are also known as graph modification problems, all NP-complete. Our results imply that they are fixed-parameter tractable parameterized by k, thereby resolving the long-standing open problem on the parameterized complexity of recognizing interval$+ke$, first asked by Bodlaender et al. [Bioinformatics, 11:49–57, 1995]. Moreover, our algorithms for recognizing interval $+ kv$ and interval $- ke$ run in times $O(6^k \cdot (n+m))$ and $O(8^k \cdot (n+m))$, (where n and m stand for the numbers of vertices and edges respectively in the input graph,) significantly improving the $O(k^{2k} \cdot n^3 m)$-time algorithm of Heggernes et al. [STOC 2007; SICOMP 2009] and the $O(10^k \cdot n^9)$-time algorithm of Cao and Marx [SODA 2014; TALG 2015] respectively.

Yixin Cao
Institute for Computer Science and Control
Hungarian Academy of Sciences
yixin.cao@polyu.edu.hk

CP17

Raising The Bar For Vertex Cover: Fixed-Parameter Tractability Above A Higher Guarantee

We investigate the following above-guarantee parameterization of the classical Vertex Cover problem: Given a graph G and an integer k as input, does G have a vertex cover of size at most $(2LP - MM) + k$? Here MM is the size of a maximum matching of G, LP is the value of an optimum solution to the relaxed (standard) LP for Vertex Cover on G, and k is the parameter. It can be shown that $(2LP - MM)$ is a lower bound on vertex cover size, and since $(2LP - MM) \geq LP \geq MM$, this is a stronger pa-

rameterization than those—namely, above MM, and above LP—which have been studied so far. We prove that Vertex Cover is fixed-parameter tractable for this stronger parameter k: We derive an algorithm which solves Vertex Cover in time $\mathcal{O}^*(3^k)$, and thus push the envelope further on the parameterized tractability of Vertex Cover.

Shivam Garg
IIT Bombay, Mumbai, India
shivamgarg@iitb.ac.in

Geevarghese Philip
Max-Planck-Institute for Informatics,
Saarbruecken, Germany
gphilip@cmi.ac.in

CP17

Directed Multicut Is $W[1]$-Hard, Even for Four Terminal Pairs

We prove that Multicut in directed graphs, parameterized by the size of the cutset, is $W[1]$-hard and hence unlikely to be fixed-parameter tractable even if restricted to instances with only four terminal pairs. This negative result almost completely resolves one of the central open problems in the area of parameterized complexity of graph separation problems, posted originally by Marx and Razgon [SIAM J. Comput. 43(2):355–388 (2014)], leaving only the case of three terminal pairs open. The case of two terminal pairs was shown to be FPT by Chitnis et al. [SIAM J. Comput. 42(4):1674–1696 (2013)]. Our gadget methodology also allows us to prove $W[1]$-hardness of the Steiner Orientation problem parameterized by the number of terminal pairs, resolving an open problem of Cygan, Kortsarz, and Nutov [SIAM J. Discrete Math. 27(3):1503-1513 (2013)].

Marcin Pilipczuk
University of Warsaw
malcin@mimuw.edu.pl

Magnus Wahlström
Royal Holloway, University of London
Magnus.Wahlstrom@rhul.ac.uk

CP17

Subexponential Parameterized Algorithm for Interval Completion

In the Interval Completion problem we are given an n-vertex graph G and an integer k, and the task is to transform G by making use of at most k edge additions into an interval graph. This is a fundamental graph modification problem with applications in sparse matrix multiplication and molecular biology. The question about fixed-parameter tractability of Interval Completion was asked by Kaplan, Shamir and Tarjan [FOCS 1994; SIAM J. Comput. 1999] and was answered affirmatively more than a decade later by Villanger at el. [STOC 2007; SIAM J. Comput. 2009], who presented an algorithm with running time $O(k^{2k}n^3m)$. We give the first subexponential parameterized algorithm solving Interval Completion in time $k^{O(\sqrt{k})}n^{O(1)}$. This adds Interval Completion to a very small list of parameterized graph modification problems solvable in subexponential time.

Ivan Bliznets
Institute of Mathematic, St.Petersburg Department of Steklov
iabliznets@gmail.com

Fedor Fomin
Dep. of Informatics
University of Bergen
fomin@ii.uib.no

Marcin Pilipczuk, Michal Pilipczuk
University of Warsaw
malcin@mimuw.edu.pl, michal.pilipczuk@mimuw.edu.pl

CP18

Improved Deterministic Algorithms for Linear Programming in Low Dimensions

At SODA'93, Chazelle and Matoušek presented a derandomization of Clarkson's sampling-based algorithm [FOCS'88] for solving linear programs with n constraints and d variables in $d^{(7+o(1))d}n$ deterministic time. The time bound can be improved to $d^{(5+o(1))d}n$ with subsequent work by Brönnimann, Chazelle, and Matoušek [FOCS'93]. We first point out a much simpler derandomization of Clarkson's algorithm that avoids ε-approximations and runs in $d^{(3+o(1))d}n$ time. We then describe a few additional ideas that eventually improve the deterministic time bound to $d^{(1/2+o(1))d}n$.

Timothy M. Chan
University of Waterloo
tmchan@uwaterloo.ca

CP18

An Efficient Algorithm for Computing High-Quality Paths Amid Polygonal Obstacles

We study a path-planning problem amid a set of obstacles in $\mathbb{R}^2$. Specifically, the problem asks for a path minimizing the reciprocal of the clearance integrated over the length of the path. We present the first polynomial-time approximation scheme for this problem. Let n be the total number of obstacle vertices and let $\varepsilon \in (0, 1]$. Our algorithm computes in time $O(\frac{n^2}{\varepsilon^2} \log \frac{n}{\varepsilon})$ a path of total cost at most $(1+\varepsilon)$ times the cost of the optimal path.

Pankaj Agarwal
Duke University
pankaj@cs.duke.edu

Kyle Fox
University of Illinois, Urbana-Champaign
kylefox@cs.duke.edu

Oren Salzman
Tel Aviv University
orenzalz@post.tau.ac.il

CP18

Approximating the K-Level in Three-Dimensional Plane Arrangements

For a set of n non-vertical planes in three dimensions, and a parameter $r < n$, we give a simple alternative proof of the xistence of $O(1/r)$-cutting of the first n/r levels of $A(H)$ via a terrain consisting of $O(r/\epsilon^3)$ triangular faces. The proof avoids sampling, and exploits techniques based on planar separators and various structural properties of levels in three-dimensional arrangements and planar maps.

Sariel Har-Peled

Department of Computer Science
University of Illinios, Urbana-Champaign
sariel@uiuc.edu

CP18

A Fast and Simple Algorithm for Computing Approximate Euclidean Minimum Spanning Trees

The Euclidean minimum spanning tree (EMST) is a fundamental and widely studied structure. In the approximate version we are given an n-element point set P in R^d and an error parameter $\epsilon > 0$, and the objective is to compute a spanning tree over P whose weight is at most $(1+\epsilon)$ times that of the true minimum spanning tree. Assuming that d is a fixed constant, existing algorithms have running times that (up to logarithmic factors) grow as $O(n/\epsilon^{\Omega(d)})$. We present an algorithm whose running time is $O(n \log n + (\epsilon^{-2} \log^2 \frac{1}{\epsilon})n)$. Thus, this is the first algorithm for approximate EMSTs that eliminates the exponential ϵ dependence on dimension. (Note that the O-notation conceals a constant factor of the form $O(1)^d$.) The algorithm is deterministic and very simple.

Sunil Arya
The Hong Kong University of Science and Technology
arya@cse.ust.hk

David M. Mount
University of Maryland
mount@cs.umd.edu

CP18

Persistent Homology and Nested Dissection

Nested dissection exploits the underlying topology to do matrix reductions while persistent homology exploits matrix reductions to the reveal underlying topology. It seems natural that one should be able to combine these techniques to beat the currently best bound of matrix multiplication time for computing persistent homology. However, nested dissection works by fixing a reduction order, whereas persistent homology generally constrains the ordering according to an input filtration. Despite this obstruction, we show that it is possible to combine these two theories. This shows that one can improve the computation of persistent homology if the underlying space has some additional structure. We give reasonable geometric conditions under which one can beat the matrix multiplication bound for persistent homology.

Michael Kerber
University of Technology, Graz, Austria
kerber@tugraz.at

Donald Sheehy
University of Connecticut
don.r.sheehy@gmail.com

Primoz Skraba
Jozef Stefan Institute
Ljubljana, Slovenia
primoz.skraba@ijs.si

CP19

Deterministic APSP, Orthogonal Vectors, and More: Quickly Derandomizing Razborov-Smolensky

We show how to solve all-pairs shortest paths on n nodes in *deterministic* $n^3/2^{\Omega(\sqrt{\log n})}$ time, and how to deterministically count the pairs of orthogonal vectors among n 0-1 vectors in $d = c \log n$ dimensions in $n^{2-1/O(\log c)}$ time. These running times essentially match the best known randomized algorithms of (Williams, STOC'14) and (Abboud, Williams, and Yu, SODA 2015) respectively, and the ability to count was open even for randomized algorithms. By reductions, these two results yield faster deterministic algorithms for many other problems. Our techniques can also be used to deterministically count k-SAT assignments on n variable formulas in $2^{n-n/O(k)}$ time, roughly matching the best known running times for detecting satisfiability and resolving an open problem of Santhanam (2013). A key to our constructions is an efficient way to deterministically simulate certain probabilistic polynomials critical to the algorithms of prior work, carefully applying small-biased sets and modulus-amplifying polynomials.

Timothy Chan
University of Waterloo
tmchan@cs.uwaterloo.ca

Ryan Williams
Stanford University
rrwilliams@gmail.com

CP19

A Faster Subquadratic Algorithm for Finding Outlier Correlations

We study the problem of detecting *outlier pairs* of strongly correlated variables among a collection of n variables with otherwise weak pairwise correlations. After normalization, this task amounts to the geometric task where we are given as input a set of n vectors with unit Euclidean norm and dimension d, and we are asked to find all the outlier pairs of vectors whose inner product is at least ρ in absolute value, subject to the promise that all but at most q pairs of vectors have inner product at most τ in absolute value for some constants $0 < \tau < \rho < 1$. Improving on an algorithm of G. Valiant [FOCS 2012; J. ACM 2015], we present a randomized algorithm that for Boolean inputs ($\{-1,1\}$-valued data normalized to unit Euclidean length) runs in time

$$\tilde{O}\big(n^{\max\{1-\gamma+M(\Delta\gamma,\gamma),\, M(1-\gamma,2\Delta\gamma)\}} + qdn^{2\gamma}\big),$$

where $0 < \gamma < 1$ is a constant tradeoff parameter and $M(\mu,\nu)$ is the exponent to multiply an $\lfloor n^\mu \rfloor \times \lfloor n^\nu \rfloor$ matrix with an $\lfloor n^\nu \rfloor \times \lfloor n^\mu \rfloor$ matrix and $\Delta = 1/(1-\log_\tau \rho)$. As corollaries we obtain randomized algorithms that run in time

$$\tilde{O}\big(n^{\frac{2\omega}{3-\log_\tau \rho}} + qdn^{\frac{2(1-\log_\tau \rho)}{3-\log_\tau \rho}}\big)$$

and in time

$$\tilde{O}\big(n^{\frac{4}{2+\alpha(1-\log_\tau \rho)}} + qdn^{\frac{2\alpha(1-\log_\tau \rho)}{2+\alpha(1-\log_\tau \rho)}}\big),$$

where $2 \le \omega < 2.38$ is the exponent for square matrix multiplication and $0.3 < \alpha \le 1$ is the exponent for rectangular matrix multiplication. We present further corollaries for the light bulb problem and for learning sparse Boolean functions. (The notation $\tilde{O}(\cdot)$ hides polylogarithmic factors in n and d whose degree may depend on ρ and τ.)

Matti Karppa
Aalto University

matti.karppa@aalto.fi

Petteri Kaski
Aalto University
Department of Information and Computer Science & HIIT
petteri.kaski@aalto.fi

Jukka Kohonen
Aalto University
jukka.kohonen@aalto.fi

CP19

Higher Lower Bounds from the 3sum Conjecture

Abstract not available at time of publication.

Tsvi Kopelowitz
University of Michigan
kopelot@gmail.com

Ely Porat
Bar Ilan University
porately@cs.biu.ac.il

Seth Pettie
University of Michigan, Ann Arbor
pettie@gmail.com

CP19

Algorithmic Complexity of Power Law Networks

We define a deterministic condition for checking whether a graph has a power law degree distribution and experimentally validate it on real-world networks. This definition allows us to derive interesting properties of power law networks. We observe that for exponents of the degree distribution in the range $[1, 2]$ such networks exhibit double power law phenomenon that was observed for several real-world networks. Moreover, we give a novel theoretical explanation why many algorithms run faster on real-world data than what is predicted by algorithmic worst-case analysis. We show how to exploit the power law degree distribution to design faster algorithms for a number of classic problems including transitive closure, maximum matching, determinant, PageRank, matrix inverse, counting triangles and finding maximum clique. In contrast to previously done average-case analyses, we believe that this is the first "waterproof' argument that explains why many real-world networks are easier.

Pawel Brach
University of Warsaw
pawel.brach@mimuw.edu.pl

Marek Cygan
Institute of Informatics, University of Warsaw
cygan@mimuw.edu.pl

Jakub Lacki
Sapienza University of Rome
j.lacki@mimuw.edu.pl

Piotr Sankowski
Institute of Informatics, University of Warsaw
sank@mimuw.edu.pl

CP19

Subtree Isomorphism Revisited

The *Subtree Isomorphism* problem asks whether a given tree is contained in another given tree. The problem is of fundamental importance and has been studied since the 1960s. For some variants, e.g., *ordered trees*, near-linear time algorithms are known, but for the general case truly subquadratic algorithms remain elusive. Our first result is a reduction from the Orthogonal Vectors problem to Subtree Isomorphism, showing that a truly subquadratic algorithm for the latter refutes the Strong Exponential Time Hypothesis (SETH). In light of this conditional lower bound, we focus on natural special cases for which no truly subquadratic algorithms are known. We classify these cases against the quadratic barrier, showing in particular that:

- Even for binary, rooted trees, a truly subquadratic algorithm refutes SETH.
- Even for rooted trees of depth $O(\log \log n)$, where n is the total number of vertices, a truly subquadratic algorithm refutes SETH.
- For every constant d, there is a constant $\epsilon_d > 0$ and a randomized, truly subquadratic algorithm for degree-d rooted trees of depth at most $(1 + \epsilon_d) \log_d n$. In particular, there is an $O(\min\{2.85^h, n^2\})$ algorithm for binary trees of depth h.

Our reductions utilize new "tree gadgets" that are likely useful for future SETH-based lower bounds for problems on trees. Our upper bounds apply a folklore result from randomized decision tree complexity.

Amir Abboud
Stanford University
abboud@cs.stanford.edu

Arturs Backurs
MIT
backurs@mit.edu

Thomas Dueholm Hansen
Aarhus University
tdh@cs.au.dk

Virginia Vassilevska Williams
Stanford University
virgi@cs.stanford.edu

Or Zamir
Tel Aviv University
orzamir@mail.tau.ac.il

CP20

Maximum Matchings in Dynamic Graph Streams and the Simultaneous Communication Model

We study the problem of finding an approximate maximum matching in two closely related computational models, namely, the dynamic graph streaming model and the simultaneous multi-party communication model. We resolve the space complexity of single-pass turnstile streaming algorithms for approximating matchings. Our results for dynamic graph streams also resolve the simultaneous communication complexity of approximating matchings in the edge partition model. We further design new random-

ized and deterministic protocols for the vertex partition model.

Sepehr Assadi, Sanjeev Khanna, Yang Li, Grigory Yaroslavtsev
University of Pennsylvania
sassadi@cis.upenn.edu, sanjeev@cis.upenn.edu, yangli2@cis.upenn.edu, grigory@grigory.us

CP20

Incidence Geometries and the Pass Complexity Of Semi-Streaming Set Cover

Set cover, over a universe of size n, may be modelled as a data-streaming problem, where the m sets that comprise the instance are to be read one by one. A semi-streaming algorithm is allowed only $O(n \operatorname{poly}\{\log n, \log m\})$ space to process this stream. For each $p \geq 1$, we give a very simple deterministic algorithm that makes p passes over the input stream and returns an appropriately certified $(p+1)n^{1/(p+1)}$-approximation to the optimum set cover. More importantly, we proceed to show that this approximation factor is essentially tight, by showing that a factor better than $0.99\, n^{1/(p+1)}/(p+1)^2$ is unachievable for a p-pass semi-streaming algorithm, even allowing randomisation. In particular, this implies that achieving a $\Theta(\log n)$-approximation requires $\Omega(\log n/\log\log n)$ passes, which is tight up to the $\log\log n$ factor. These results extend to a relaxation of the set cover problem where we are allowed to leave an ϵ fraction of the universe uncovered: the tight bounds on the best approximation factor achievable in p passes turn out to be $\Theta_p(\min\{n^{1/(p+1)}, \epsilon^{-1/p}\})$. Our lower bounds are based on a construction of a family of high-rank incidence geometries, which may be thought of as vast generalisations of affine planes. This construction, based on algebraic techniques, appears flexible enough to find other applications and is therefore interesting in its own right.

Amit Chakrabarti
Dartmouth
ac@cs.dartmouth.edu

Anthony Wirth
The University of Melbourne
awirth@unimelb.edu.au

CP20

Clustering Problems on Sliding Windows

We present the first polylogarithmic space $O(1)$-approximation to the metric k-median and metric k-means problems in the sliding window model. We show that using only polylogarithmic space we can maintain a summary of the current window from which we can construct an $O(1)$-approximate clustering solution. Additionally, we give the first algorithm that, given an insertion-only streaming coreset of space s (using merge-and-reduce), maintains a coreset in the sliding window model using $O(s^2\epsilon^{-2}\log W)$ space.

Vladimir Braverman, Harry Lang, Keith Levin
Johns Hopkins University
vova@cs.jhu.edu, hlang@math.jhu.edu, klevin@jhu.edu

Morteza Monemizadeh
Institute for Computer Science
Goethe-University Frankfurt am Main
m.monemizadeh@gmail.com

CP20

Kernelization via Sampling with Applications to Finding Matchings and Related Problems in Dynamic Graph Streams

We present a simple but powerful subgraph sampling primitive that is applicable in a variety of computational models including dynamic graph streams (where the input graph is defined by a sequence of edge/hyperedge insertions and deletions) and distributed systems such as MapReduce. In the case of dynamic graph streams, we use this primitive to prove the following results:

- *Matching:* Our main result for matchings is that there exists an $\tilde{O}(k^2)$ space algorithm that returns the edges of a maximum matching on the assumption the cardinality is at most k. The best previous algorithm used $\tilde{O}(kn)$ space where n is the number of vertices in the graph and we prove our result is optimal up to logarithmic factors. Our algorithm has $\tilde{O}(1)$ update time. We also show that there exists an $\tilde{O}(n^2/\alpha^3)$ space algorithm that returns an α-approximation for matchings of arbitrary size. In independent work, Assadi et al. (SODA 2016) proved this approximation algorithm is optimal and provided an alternative algorithm. We generalize our exact and algorithms to weighted matching. For graphs with low arboricity such as planar graphs, the space required for constant approximation can be further reduced.
- *Vertex Cover and Hitting Set:* There exists an $\tilde{O}(k^d)$ space algorithm that solves the minimum hitting set problem where d is the cardinality of the input sets and k is an upper bound on the size of the minimum hitting set. We prove this is optimal up to logarithmic factors. Our algorithm has $\tilde{O}(1)$ update time.

Finally, we present fast, small-space dynamic graph stream algorithms for a larger family of parameterized problems and show lower bounds for natural problems outside this family.

Rajesh Chitnis
The Weizmann Institute of Science
rajesh.chitnis@weizmann.ac.il

Graham Cormode
University of Warwick
g.cormode@warwick.ac.uk

Hossein Esfandiari
University of Maryland at College Park
hossein@cs.umd.edu

MohammadTaghi Hajiaghayi
University of Maryland, College Park
hajiagha@cs.umd.edu

Andrew McGregor
University of Massachusetts, Amherst
mcgregor@cs.umass.edu

Morteza Monemizadeh
Computer Science Institute of Charles University
monemi@iuuk.mff.cuni.cz

Sofya Vorotnikova
University of Massachusetts Amherst

svorotni@cs.umass.edu

CP20

Gowers Norm, Function Limits, and Parameter Estimation

Let $\{f_i : F_p^i \to \{0,1\}\}$ be a sequence of functions, where p is a fixed prime and F_p is the finite field of order p. The limit of the sequence can be syntactically defined using the notion of ultralimit. Inspired by the Gowers norm, we introduce a metric over limits of function sequences, and study properties of it. One application of this metric is that it provides a simpler characterization of affine-invariant parameters of functions that are constant-query estimable than the previous one obtained by Yoshida (STOC'14). Using this characterization, we show that the property of being a function of a constant number of low-degree polynomials and a constant number of factored polynomials (of arbitrary degrees) is constant-query testable if it is closed under blowing-up. Examples of this property include the property of having a constant spectral norm and degree-structural properties with rank conditions.

Yuichi Yoshida
National Institute of Informatics
Preferred Infrastructure
yyoshida@nii.ac.jp

CP21

On the Economic Efficiency of the Combinatorial Clock Auction

Since the 1990s spectrum auctions have been implemented world-wide. This has provided for a practical examination of an assortment of auction mechanisms and, amongst these, two simultaneous ascending price auctions have proved to be extremely successful. These are the simultaneous multiround ascending auction (SMRA) and the combinatorial clock auction (CCA). It has long been known that, for certain classes of valuation functions, the SMRA provides good theoretical guarantees on social welfare. However, no such guarantees were known for the CCA. In this paper, we show that CCA does provide strong guarantees on social welfare provided the price increment and stopping rule are well-chosen. This is very surprising in that the choice of price increment has been used primarily to adjust auction duration and the stopping rule has attracted little attention. The main result is a polylogarithmic approximation guarantee for social welfare when the maximum number of items demanded $\mathcal{C}$ by a bidder is fixed. Specifically, we show that either the revenue of the CCA is at least an $\Omega\Big(\frac{1}{\mathcal{C}^2 \log n \log^2 m}\Big)$-fraction of the optimal welfare or the welfare of the CCA is at least an $\Omega\Big(\frac{1}{\log n}\Big)$-fraction of the optimal welfare, where n is the number of bidders and m is the number of items. As a corollary, the welfare ratio – the worst case ratio between the social welfare of the optimum allocation and the social welfare of the CCA allocation – is at most $O(\mathcal{C}^2 \cdot \log n \cdot \log^2 m)$. We emphasize that this latter result requires no assumption on bidders valuation functions. Finally, we prove that such a dependence on $\mathcal{C}$ is necessary. In particular, we show that the welfare ratio of the CCA is at least $\Omega\Big(\mathcal{C} \cdot \frac{\log m}{\log \log m}\Big)$.

Nicolas Bousquet, Yang Cai
McGill University
nicolas.bousquet2@mail.mcgill.ca, cai@cs.mcgill.ca

Christoph Hunkenschröder
University of Bonn
chr.hunkenschroeder@gmail.com

Adrian Vetta
Department of Math. & Stat. and School of Comp. Sci.
McGill University
vetta@math.mcgill.ca

CP21

On the Complexity of Dynamic Mechanism Design

We introduce a dynamic mechanism design problem in which the designer wants to offer two items for sale to the same agent, one now and one in the future. We show that finding the revenue maximizing deterministic mechanism, subject to ex-post individual rationality and truthfulness, is NP-hard. We also prove several positive results and we show that, in an environment in which contracts cannot be enforced, the optimum mechanism requires multiple rounds of communication.

Christos Papadimitriou
UC Berkeley
christos@cs.berkeley.edu

George Pierrakos, Alexandros Psomas
University of California, Berkeley
geopier@gmail.com, alexpsomi@cs.berkeley.edu

Aviad Rubinstein
UC Berkeley
aviad@cs.berkeley.edu

CP21

Simple Pricing Schemes for Consumers With Evolving Values

We consider a pricing problem where a buyer is interested in purchasing/using a good, such as an app or music or software, repeatedly over time. The consumer discovers his value for the good only as he uses it, and the value evolves with each use as a martingale. We provide a simple pricing scheme and show that its revenue is a constant fraction of the buyers expected cumulative value.

Shuchi Chawla
Department of Computer Sciences
University of Wisconsin-Madison
shuchi@cs.wisc.edu

Nikhil R. Devanur
Microsoft Research
nikdev@microsoft.com

Anna R. Karlin
University of Washington
karlin@cs.washington.edu

Balasubramanian Sivan
Google Research
balu2901@gmail.com

CP21

Approximately Efficient Double Auctions with

Strong Budget Balance

This paper focuses on two-sided algorithmic mechanism design - a research area that has not been thoroughly studied as its one-sided counterpart. In the light of [Myerson and Satterwhite, 1983]'s impossibility result, which prevents any individually rational, dominant strategy incentive compatible, and strongly budget-balanced mechanism maximizing the social welfare, we devise a two-sided version of sequential posted price mechanisms to obtain a constant approximation to the social welfare while satisfying the above constraints.

Riccardo Colini-Baldeschi, Bart De Keijzer, Stefano Leonardi
Sapienza University of Rome
colini@dis.uniroma1.it, bdekeijzer@gmail.com, leonardi@dis.uniroma1.it

Stefano Turchetta
University of Oxford
stefano.turchetta@gmail.com

CP21

Interpolating Between Truthful and Non-Truthful Mechanisms for Combinatorial Auctions

We study the communication complexity of combinatorial auctions via *interpolation mechanisms* that interpolate between non-truthful and truthful protocols. Specifically, an interpolation mechanism has two phases. In the first phase, the bidders participate in some non-truthful protocol *whose output is itself a truthful protocol.* In the second phase, the bidders participate in the truthful protocol selected during phase one. Note that virtually all existing auctions have either a non-existent first phase (and are therefore truthful mechanisms), or a non-existent second phase (and are therefore just traditional protocols, analyzed via the Price of Anarchy/Stability). The goal of this paper is to understand the benefits of interpolation mechanisms versus truthful mechanisms or traditional protocols, and develop the necessary tools to formally study them. Interestingly, we exhibit settings where interpolation mechanisms greatly outperform the optimal traditional and truthful protocols. Yet, we also exhibit settings where interpolation mechanisms are provably no better than truthful ones. Finally, we apply our new machinery to prove that the recent single-bid mechanism of Devanur et. al. (the only pre-existing interpolation mechanism in the literature) achieves the optimal price of anarchy among a wide class of protocols, a claim that simply can't be addressed by appealing just to machinery from communication complexity or the study of truthful mechanisms.

Mark Braverman, Jieming Mao
Princeton University
mbraverm@cs.princeton.edu, jiemingm@cs.princeton.edu

Seth M. Weinberg
Princet University
smweinberg@csail.mit.edu

CP22

Integrality Gaps and Approximation Algorithms for Dispersers and Bipartite Expanders

We study the problem of approximating the quality of a disperser. A bipartite graph G on $([N],[M])$ is a $(\rho N,(1-\delta)M)$-disperser if for any subset $S\subseteq[N]$ of size ρN, the neighbor set $\Gamma(S)$ contains at least $(1-\delta)M$ distinct vertices. Our main results are strong integrality gaps in the Lasserre hierarchy and an approximation algorithm for dispersers.

1. For any $\alpha>0$, $\delta>0$, and a random bipartite graph G with left degree $D=O(\log N)$, we prove that the Lasserre hierarchy cannot distinguish whether G is an $(N^{\alpha},(1-\delta)M)$-disperser or not an $(N^{1-\alpha},\delta M)$-disperser.
2. For any $\rho>0$, we prove that there exist infinitely many constants d such that the Lasserre hierarchy cannot distinguish whether a random bipartite graph G with right degree d is a $(\rho N,(1-(1-\rho)^d)M)$-disperser or not a $(\rho N,(1-\Omega(\frac{1-\rho}{\rho d+1-\rho}))M)$-disperser. We also provide an efficient algorithm to find a subset of size exact ρN that has an approximation ratio matching the integrality gap within an extra loss of $\frac{\min\{\frac{\rho}{1-\rho},\frac{1-\rho}{\rho}\}}{\log d}$.

Our method gives an integrality gap in the Lasserre hierarchy for bipartite expanders with left degree D. G on $([N],[M])$ is a $(\rho N,a)$-expander if for any subset $S\subseteq[N]$ of size ρN, the neighbor set $\Gamma(S)$ contains at least $a\cdot\rho N$ distinct vertices. We prove that for any constant $\epsilon>0$, there exist constants $\epsilon'<\epsilon,\rho$, and D such that the Lasserre hierarchy cannot distinguish whether a bipartite graph on $([N],[M])$ with left degree D is a $(\rho N,(1-\epsilon')D)$-expander or not a $(\rho N,(1-\epsilon)D)$-expander.

Xue Chen
Department of Computer Science
University of Texas at Austin
xchen@cs.utexas.edu

CP22

Packing Small Vectors

We study the online d-dimensional Vector Bin Packing problem for small vectors (relative to the size of a bin). We give a constant competitive ratio, improving the best known algorithm from $O(\log d)$ to $\approx e$, for arbitrary d. For two dimensions, we present a First Fit variant with a competitive ratio ≈ 1.48. Additionally, we give an essentially tight algorithm (not via a First Fit variant) with a competitive ratio arbitrarily close to 4/3.

Yossi Azar
Tel-Aviv University
azar@tau.ac.il

Ilan R. Cohen
Tel Aviv University
Tel Aviv University
ilanrcohen@gmail.com

Amos Fiat, Alan Roytman
Tel Aviv University
fiat@tau.ac.il, alan.roytman@cs.tau.ac.il

CP22

Improved Approximation for Vector Bin Packing

We study the d-dimensional vector bin packing problem, a well-studied generalization of bin packing arising in resource allocation and scheduling problems. Here we are given a set of d-dimensional vectors $v_1,\ldots,v_n$ in $[0,1]^d$, and the goal is to pack them into the least number of bins so that for each bin B, the sum of the vectors in it is at

most 1 in every dimension, i.e., $||\sum_{v_i \in B} v_i||_\infty \leq 1$. For the 2-dimensional case we give an asymptotic approximation guarantee of $1 + \ln(1.5) + \epsilon \approx (1.405 + \epsilon)$, improving upon the previous bound of $1 + \ln 2 + \epsilon \approx (1.693 + \epsilon)$. We also give an almost tight $(1.5 + \epsilon)$ absolute approximation guarantee, improving upon the previous bound of 2. For the d-dimensional case, we get a $1.5 + \ln(\frac{d+1}{2}) + \epsilon \approx 0.807 + \ln(d + 1) + \epsilon$ guarantee, improving upon the previous $(1 + \ln d + \epsilon)$ guarantee based on the Round & Approx framework. Here $(1 + \ln d)$ was a natural barrier as rounding-based algorithms can not achieve better than d approximation. We get around this by exploiting various structural properties of (near)-optimal packings, and using multi-objective multi-budget matching based techniques and expanding the R & A framework to go beyond rounding-based algorithms. Along the way we also prove several results that could be of independent interest.

Nikhil Bansal
Department of Mathematics and Computer Science
Eindhoven University of Technology, Eindhoven, Netherlands
n.bansal@tue.nl

Marek Elias
Eindhoven University of Technology, Netherlands
m.elias@tue.nl

Arindam Khan
School of Computer Science,
Georgia Tech
akhan67@gatech.edu

CP22

Approximation Schemes for Machine Scheduling with Resource (in-)dependent Processing Times

We consider two related scheduling problems: resource constrained scheduling on identical parallel machines and a generalization with resource dependent processing times. In both problems, jobs require a certain amount of an additional resource and have to be scheduled minimizing the makespan, while at every point in time a given resource capacity is not exceeded. We present a method to obtain asymptotic fully polynomial approximation schemes (AFPTAS) for the problems.

Klaus Jansen
University of Kiel
Institute for Computer Science
kj@informatik.uni-kiel.de

Marten Maack, Malin Rau
University of Kiel
mmaa@informatik.uni-kiel.de, mra@informatik.uni-kiel.de

CP22

On Approximating Strip Packing with a Better Ratio Than 3/2

In the strip packing problem we are given a set of rectangular items that we want to place in a strip of given width such that we minimize the height of the obtained packing. It is a very classical two-dimensional packing problem that has received a lot of attention and it has applications in many settings such as stock-cutting and scheduling. A straight-forward reduction from Partition shows that the problem cannot be approximated with a better absolute factor than 3/2. However, this reduction requires the numeric values to be exponentially large. In this paper, we present a $(1.4 + \epsilon)$-approximation algorithm with pseudo-polynomial running time. This implies that for polynomially bounded input data the problem can be approximated with a strictly better ratio than for exponential input which is a very rare phenomenon in combinatorial optimization.

Giorgi Nadiradze
Saarland University
gionadira@gmail.com

Andreas Wiese
MPII Saarbrücken
awiese@mpi-inf.mpg.de

CP23

Recovery and Rigidity in a Regular Stochastic Block Model

The stochastic block model is a natural model for studying community detection in random networks. Its clustering properties have been extensively studied in the statistics, physics and computer science literature. Recently this area has experienced major mathematical breakthroughs, particularly for the binary (two-community) version, see [Mossel, Neeman, Sly (2012, 2013)] and [Massoulie (2013)]. In this paper, we introduce a variant of the binary model which we call the regular stochastic block model (RSBM). We prove rigidity by showing that with high probability an exact recovery of the community structure is possible. Spectral methods exhibit a regime where this can be done efficiently. Moreover we also prove that, in this setting, any suitably good partial recovery can be bootstrapped to obtain a full recovery of the communities.

Gerandy Brito
Department of Mathematics
University of Washington
gerandy@uw.edu

Ioana Dumitriu
University of Washington, Seattle
dumitriu@math.washington.edu

Shirshendu Ganguly
Department of Mathematics
University of Washington
sganguly@math.washington.edu

Christopher Hoffman
University of Washington
hoffman@math.washington.edu

Linh V. Tran
International University
National University Hochiminh City
tvlinh@hcmiu.edu.vn

CP23

Connectivity in Bridge-Addable Graph Classes: the McDiarmid-Steger-Welsh Conjecture

We prove a conjecture of McDiarmid, Steger and Welsh that says that if a class of graphs is *bridge-addable* (*i.e.* stable by adding edges between different components) then the random graph from this class is connected with proba-

bility at least $e^{-1/2} + o(1)$, when its number of vertices tends to infinity. This lower bound is tight since it is reached for forests. The best previously known constants where e^{-1}, $e^{-0.7983}$ and $e^{-2/3}$ proved respectively by McDiarmid, Steger and Welsh, by Balister, Bollobás and Gerke, and by Norin.

Guillaume Chapuy
CNRS, Université de Montréal and Université Paris Diderot
guillaume.chapuy@liafa.univ-paris-diderot.fr

Guillem Perarnau
McGill University
p.melliug@gmail.com

CP23

The Power of Two Choices with Simple Tabulation

The power of two choices is a classic paradigm for load balancing when assigning m balls to n bins. When placing a ball, we pick two bins according to two hash functions h_0 and h_1, and place the ball in the least loaded bin. Assuming fully random hash functions, when $m = O(n)$, Azar et al. [STOC'94] proved that the maximum load is $\lg \lg n + O(1)$ with high probability. No such bound was known with a hash function implementable in constant time. In this paper, we investigate the power of two choices when the hash functions h_0 and h_1 are implemented with simple tabulation, which is a very efficient hash function evaluated in constant time. Following their analysis of Cuckoo hashing [J.ACM'12], Pătraşcu and Thorup claimed that the expected maximum load with simple tabulation is $O(\lg \lg n)$. This did not include any high probability guarantee, so the load balancing was not yet to be trusted. Here, we show that with simple tabulation, the maximum load is $O(\lg \lg n)$ with high probability, giving the first constant time hash function with this guarantee. We also give a concrete example where, unlike with fully random hashing, the maximum load is not bounded by $\lg \lg n + O(1)$, or even $(1 + o(1)) \lg \lg n$ with high probability. Finally, we show that the expected maximum load is $\lg \lg n + O(1)$, just like with fully random hashing.

Søren Dahlgaard, Mathias Knudsen
University of Copenhagen
soerend@di.ku.dk, knudsen@di.ku.dk

Eva Rotenberg
Department of Computer Science,
University of Copenhagen
eva@rotenberg.dk

Mikkel Thorup
AT&T Labs - Research and University of Copenhagen
mikkel2thorup@gmail.com

CP23

Distance in the Forest Fire Model How Far Are You from Eve?

Leskovec, Kleinberg and Faloutsos (2005) observed that many social networks exhibit properties such as shrinking (i.e. bounded) diameter, densification, and (power-law) heavy tail degree distributions. To explain these phenomena, they introduced a generative model, called the Forest Fire model, and using simulations showed that this model indeed exhibited these properties; however, proving this rigorously was left as an open problem. In this paper, we analyse one of these properties, shrinking diameter. We define a restricted version of their model that incorporates the main features that seem to contribute towards this property, and prove that the graphs generated by this model exhibit shrinking distance to the seed graph. We prove that an even simpler model, the random walk model, already exhibits this phenomenon.

Frederik Mallmann-Trenn
École normale supérieure
mallmann@di.ens.fr

CP23

Species Trees from Gene Trees Despite a High Rate of Lateral Genetic Transfer: A Tight Bound

Reconstructing the tree of life from molecular sequences is a fundamental problem in computational biology. Modern data sets often contain a large number of genes which can complicate the reconstruction problem due to the fact that different genes may undergo different evolutionary histories. This is the case in particular in the presence of lateral genetic transfer (LGT), whereby a gene is inherited from a distant species rather than an immediate ancestor. Such an event produces a gene tree which is distinct from (but related to) the species phylogeny. In previous work, a stochastic model of LGT was introduced and it was shown that the species phylogeny can be reconstructed from gene trees despite surprisingly high rates of LGT. Both lower and upper bounds on this rate were obtained, but a large gap remained. Here we close this gap, up to a constant.

Sebastien Roch
Department of Mathematics
UW Madison
roch@math.wisc.edu

CP24

Tight Conditional Lower Bounds for Counting Perfect Matchings on Graphs of Bounded Treewidth, Cliquewidth, and Genus

Assuming the Strong Exponential-Time Hypothesis, we show that the problem of counting perfect matchings has no $(2-\epsilon)^k n^{O(1)}$ time algorithm for any $\epsilon > 0$ on graphs of treewidth k, whereas it is known to be solvable in time $2^k n^{O(1)}$ if a tree decomposition of width k is given. We also show that the problem has no $O(n^{(1-\epsilon)k})$ time algorithm for any $\epsilon > 0$ on graphs of cliquewidth k, but it can be solved in time $O(n^{k+1})$ if a k-expression is given. Furthermore, assuming the Exponential-Time Hypothesis, we show that there is no $2^{o(k)} n^{O(1)}$ time algorithm for the problem on graphs of genus k.

Radu Curticapean
Simons Institute, UC Berkeley
radu.curticapean@gmail.com

Dániel Marx
Hungarian Academy of Sciences (MTA SZTAKI),
Budapest,Hungary
dmarx@cs.bme.hu

CP24

Tight Bounds for Graph Homomorphism and Sub-

graph Isomorphism

We prove that unless Exponential Time Hypothesis (ETH) fails, deciding if there is a homomorphism from graph G to graph H cannot be done in time $|V(H)|^{o(|V(G)|)}$. We also show an exponential-time reduction from Graph Homomorphism to Subgraph Isomorphism. This rules out (subject to ETH) a possibility of $|V(H)|^{o(|V(H)|)}$-time algorithm deciding if graph G is a subgraph of H. For both problems our lower bounds asymptotically match the running time of brute-force algorithms trying all possible mappings of one graph into another. Thus, our work closes the gap in the known complexity of these fundamental problems.

Marek Cygan
Institute of Informatics, University of Warsaw
cygan@mimuw.edu.pl

Fedor Fomin
Dep. of Informatics
University of Bergen
fomin@ii.uib.no

Aleksandr Golovnev
New York University
alexgolovnev@gmail.com

Alexander Kulikov
St Petersburg Department of Steklov Institute of Mathematics
kulikov@logic.pdmi.ras.ru

Ivan Mihajlin
University of CaliforniaSan Diego
ivmihajlin@gmail.com

Jakub Pachocki
Carnegie Mellon University
pachocki@cs.cmu.edu

Arkadiusz Socala
Institute of Informatics, University of Warsaw
a.socala@mimuw.edu.pl

CP24

Constructive Algorithm for Path-Width of Matroids

We present a fixed-parameter tractable algorithm to construct a linear layout $V_1, V_2, \ldots, V_n$ of the subspaces such that $\dim((V_1+V_2+\cdots+V_i)\cap(V_{i+1}+\cdots+V_n)) \leq k$ for all i, if it exists, for input subspaces of a finite-dimensional vector space over F. When restricted to 1-dimensional subspaces, this problem is equivalent to computing the path-width of an F-represented matroid in matroid theory and computing the trellis-width of a linear code in coding theory.

Jisu Jeong
KAIST
jjisu@kaist.ac.kr

Eun Jung Kim
CNRS, LAMSADE-Universite Paris Dauphine
eunjungkim78@gmail.com

Sang-Il Oum
KAIST
sangil@kaist.edu

CP24

Blocking Optimal k-Arborescences

Given a digraph $D = (V, A)$ and a positive integer k, an arc set $F \subseteq A$ is called a *k-arborescence* if it is the disjoint union of k spanning arborescences. The problem of finding a minimum cost k-arborescence is known to be polynomial-time solvable using matroid intersection. In this paper we study the following problem: find a minimum cardinality subset of arcs that contains at least one arc from every minimum cost k-arborescence. For $k = 1$, the problem was solved in [A. Bernáth, G. Pap, Blocking optimal arborescences, IPCO 2013]. In this paper we give an algorithm for general k that has polynomial running time if k is fixed.

Attila Bernáth
MTA-ELTE Egerváry Research Group
Eötvös University, Budapest
bernath@cs.elte.hu

Tamás Király
Department of Operations Research
Eötvös University, Budapest
tkiraly@cs.elte.hu

CP24

Discovering Archipelagos of Tractability for Constraint Satisfaction and Counting

This paper addresses the general limit of standard tractability results for CSP and #CSP, that they only apply to instances where all constraints belong to a single tractable language. For this purpose we utilize the notion of a *strong backdoor* of a CSP instance, as introduced by Williams et al. (IJCAI 2003). We consider strong backdoors into *scattered classes*, consisting of CSP instances where each connected component belongs entirely to some class from a list of tractable classes and show that one can solve a given CSP instance in time fixed-parameter tractable in the size of the strong backdoor.

Robert Ganian
TU Wien
rganian@gmail.com

Ramanujan M. Sridharan
University of Bergen
ramanujan.sridharan@ii.uib.no

Stefan Szeider
TU Wien
stefan@szeider.net

CP25

Nearly Optimal Np-Hardness of Unique Coverage

The *Unique Coverage* problem, given a universe V of elements and a collection E of subsets of V, asks to find $S \subseteq V$ to maximize the number of $e \in E$ that intersects S in *exactly one* element. When each $e \in E$ has cardinality at most k, it is also known as 1-*in*-k *Hitting Set*, and admits a simple $\Omega(\frac{1}{\log k})$-approximation algorithm. For constant k, we prove that 1-in-k Hitting Set is NP-hard to approximate within a factor $O(\frac{1}{\log k})$. This improves the result of Guruswami and Zhou [SODA'11, ToC'12], who proved the

same result assuming the Unique Games Conjecture. For Unique Coverage, we prove that it is hard to approximate within a factor $O(\frac{1}{\log^{1-\epsilon} n})$ for any $\epsilon > 0$, unless NP admits quasipolynomial time algorithms. This improves the results of Demaine et al. [SODA'06, SICOMP'08], including their $\approx 1/\log^{1/3} n$ inapproximability factor which was proven under the Random 3SAT Hypothesis. Our simple proof combines ideas from two classical inapproximability results for Set Cover and Constraint Satisfaction Problem, made efficient by various derandomization methods based on bounded independence.

Venkatesan Guruswami, Euiwoong Lee
Carnegie Mellon University
guruswami@cmu.edu, euiwoonl@cs.cmu.edu

CP25

Algorithms and Adaptivity Gaps for Stochastic Probing

A stochastic probing problem consists of a set of elements whose values are independent random variables. The algorithm knows the distributions of these variables, but not the actual outcomes. The only way to learn the actual outcomes is to probe these elements. However, there are constraints on which set of elements may be probed. (E.g., we may have to travel in some metric to probe elements but have limited time.) These constraints are called outer constraints. We want to develop an algorithm that picks some set of elements to maximize the (expected) value, subject to the picked subset of elements satisfying some other set of constraints, called the inner constraints. In the past, probing problems were studied for the case when both inner and outer constraints were intersections of matroids; these modeled kidney matching and Bayesian auctions applications. One limitation of past work was their reliance on linear-programming-like techniques, which made going beyond matroid-like structures difficult. In this work, we give a very general adaptivity gap result that holds for all prefix-closed outer constraints, as long as the inner constraints are intersections of matroids. The adaptivity gap is $O(\log n)$ for any constant number of inner matroid constraints. The prefix-closedness captures most "reasonable' outer constraints, like orienteering, connectivity, and precedence. Based on this we obtain the first approximation algorithms for a number of stochastic probing problems, which have applications, e.g., to path-planning and precedence-constrained scheduling.

Anupam Gupta
Carnegie Mellon University
anupamg@cs.cmu.edu

Viswanath Nagarajan
IBM T. J. Watson Research
viswanath@us.ibm.com

Sahil Singla
CMU
ssingla@cmu.edu

CP25

Discrete Gaussian Sampling Reduces to Cvp and Svp

The discrete Gaussian $D_{L-t,s}$ is the distribution that assigns to each vector x in a shifted lattice $L - t$ probability proportional to $e^{-\pi\|x\|^2/s^2}$. It has long been an important tool in the study of lattices. More recently, algorithms for discrete Gaussian sampling (DGS) have found many applications in computer science. In particular, polynomial-time algorithms for DGS with very high parameters s have found many uses in cryptography and in reductions between lattice problems. And, in the past year, Aggarwal, Dadush, Regev, and Stephens-Davidowitz showed $2^{n+o(n)}$-time algorithms for DGS with a much wider range of parameters and used them to obtain the current fastest known algorithms for the two most important lattice problems, the Shortest Vector Problem (SVP) and the Closest Vector Problem (CVP). Motivated by its increasing importance, we investigate the complexity of DGS itself and its relationship to CVP and SVP. Our first result is a polynomial-time dimension-preserving reduction from DGS to CVP. There is a simple reduction from CVP to DGS, so this shows that DGS is equivalent to CVP. Our second result, which we find to be more surprising, is a polynomial-time dimension-preserving reduction from *centered* DGS (the important special case when $t = 0$) to SVP. In the other direction, there is a simple reduction from γ-approximate SVP for any $\gamma = \Omega(\sqrt{n/\log n})$, and we present some (relatively weak) evidence to suggest that this might be the best achievable approximation factor.

Noah Stephens-Davidowitz
New York University
noahsd@gmail.com

CP25

Approximation of Nonboolean 2csp

We develop a polynomial time $\Omega\left(\frac{1}{R}\log R\right)$ approximate algorithm for Max 2CSP-R, the problem where we are given a collection of constraints, each involving two variables, where each variable ranges over a set of size R, and we want to find an assignment to the variables that maximizes the number of satisfied constraints. Assuming the Unique Games Conjecture, this is the best possible approximation up to constant factors.

Guy Kindler
Hebrew University
gkindler@cs.huji.ac.il

Alexandra Kolla
UIUC
alexkolla@gmail.com

Luca Trevisan
UC Berkeley
luca@berkeley.edu

CP25

An FPTAS for Minimizing Indefinite Quadratic Forms over Integers in Polyhedra

We present a generic approach that allows us to develop a fully polynomial-time approximation scheme (FTPAS) for minimizing nonlinear functions over the integer points in a rational polyhedron in fixed dimension. The approach combines the subdivision strategy of Papadimitriou and Yannakakis (2000) with ideas similar to those commonly used to derive real algebraic certificates of positivity for polynomials. Our general approach is widely applicable. We apply it, for instance, to the Motzkin polynomial and to indefinite quadratic forms x^TQx in a fixed number of variables, where Q has at most one positive, or at most

one negative eigenvalue. In dimension three, this leads to an FPTAS for general Q.

Robert Hildebrand
ETH Zurich
robert.hildebrand@ifor.math.ethz.ch

Robert Weismantel
ETH Zuerich
robert.weismantel@ifor.math.ethz.ch

Kevin Zemmer
ETH Zurich
kevin.zemmer@ifor.math.ethz.ch

CP26

Finding Perfect Matchings in Bipartite Hypergraphs

Haxell's condition is a natural hypergraph analog of Hall's condition, which is a well-known necessary and sufficient condition for a bipartite graph to admit a perfect matching. That is, when Haxell's condition holds it forces the existence of a perfect matching in the bipartite hypergraph. Unlike in graphs, however, there is no known polynomial time algorithm to find the hypergraph perfect matching that is guaranteed to exist when Haxell's condition is satisfied. We prove the existence of an efficient algorithm to find perfect matchings in bipartite hypergraphs whenever a stronger version of Haxell's condition holds. Our algorithm can be seen as a generalization of the classical Hungarian algorithm for finding perfect matchings in bipartite graphs. The techniques we use to achieve this result could be of use more generally in other combinatorial problems on hypergraphs where disjointness structure is crucial, e.g. Set Packing.

Chidambaram Annamalai
EPFL
cannamalai@inf.ethz.ch

CP26

Obstructions for Three-Coloring Graphs with One Forbidden Induced Subgraph

The complexity of coloring graphs without long induced paths is a notorious problem in algorithmic graph theory, an especially intruiging case being that of 3-colorability. So far, not much was known about certification in this context. We prove that there are only finitely many 4-critical P_6-free graphs, and give the complete list that consists of 24 graphs. In particular, we obtain a certifying algorithm for 3-coloring P_6-free graphs, which solves an open problem posed by Golovach et al. Here, P_6 denotes the induced path on six vertices. Our result leads to the following dichotomy theorem: if H is a connected graph, then there are finitely many 4-critical H-free graphs if and only if H is a subgraph of P_6. This answers a question of Seymour. The proof of our main result involves two distinct automatic proofs, and an extensive structural analysis by hand.

Maria Chudnovsky
Princeton University
mchudnov@math.princeton.edu

Jan Goedgebeur
Ghent University
jan.goedgebeur@ugent.be

Oliver Schaudt
University of Cologne
Institute of Computer Science
schaudto@uni-koeln.de

Mingxian Zhong
Columbia University
mingxian.zhong@gmail.com

CP26

An Algorithmic Hypergraph Regularity Lemma

Szemerédi's Regularity Lemma is a powerful tool in graph theory. It asserts that all large graphs G admit a bounded partition of $E(G)$, most classes of which are bipartite subgraphs with uniformly distributed edges. The original proof of this result was non-constructive. A constructive proof was given by Alon, Duke, Lefmann, Rödl and Yuster, which allows one to efficiently construct a regular partition for any large graph. Szemerédi's Regularity Lemma was extended to hypergraphs by various authors. Frankl and Rödl gave one such extension to 3-uniform hypergraphs, and Rödl and Skokan extended this result to k-uniform hypergraphs. W.T. Gowers gave another such extension. Similarly to the graph case, all of these proofs are non-constructive. We present an efficient algorithmic version of the Hypergraph Regularity Lemma for k-uniform hypergraphs.

Brendan Nagle
Department of Mathematics and Statistics
University of South Florida
bnagle@usf.edu

Vojtech Rodl
Emory University
rodl@mathcs.emory.edu

Mathias Schacht
Universität Hamburg
Fachbereich Mathematik
schacht@math.uni-hamburg.de

CP26

Independence and Efficient Domination on P_6-Free Graphs

In the Max Independent Set problem, the input is a graph G, every vertex has a non-negative integer weight, and the task is to find a set S of pairwise non-adjacent vertices, maximizing the total weight of the vertices in S. We give an $n^{O(\log^2 n)}$ time algorithm for this problem on graphs excluding the path P_6 on 6 vertices as an induced subgraph. Currently, there is no constant k known for which Max Independent Set on P_k-free graphs becomes NP-complete, and our result implies that if such a k exists, then $k > 6$ unless all problems in NP can be decided in (quasi)polynomial time. Using the combinatorial tools that we develop for the above algorithm, we also give a polynomial-time algorithm for Efficient Domination on P_6-free graphs. In this problem, the input is a graph G, every vertex has an integer weight, and the objective is to find a set S of maximum weight such that every vertex in G has exactly one vertex in S in its closed neighborhood, or to determine that no such set exists. Prior to our work, the class of P_6-free graphs was the only class of graphs defined by a single forbidden induced subgraph on which the computational

complexity of Efficient Domination was unknown.

Daniel Lokshtanov
University of Bergen
daniello@ii.uib.no

Marcin Pilipczuk
University of Warsaw
malcin@mimuw.edu.pl

Erik Jan van Leeuwen
Max-Planck-Institut für Informatik
erikjan@mpi-inf.mpg.de

CP26

Sparsity and Dimension

We prove that posets of bounded height whose cover graphs belong to a fixed class with bounded expansion have bounded dimension. Bounded expansion, introduced by Nešetřil and Ossona de Mendez as a model for sparsity in graphs, is a property that is naturally satisfied by a wide range of graph classes, from graph structure theory (graphs excluding a minor or a topological minor) to graph drawing (e.g. graphs with constant book thickness). Therefore, our theorem generalizes a number of results including the most recent one for posets of bounded height with cover graphs excluding a fixed graph as a topological minor (Walczak, SODA 2015). We also show that the result is in a sense best possible, as it does not extend to nowhere dense classes; in fact, it already fails for cover graphs with locally bounded treewidth.

Gwenael Joret
The University of Melbourne
gjoret@ulb.ac.be

Piotr Micek
Jagellionian University
piotr.Micek@tcs.uj.edu.pl

Veit Wiechert
TU - Berlin
veit416@hotmail.com

CP27

Using Optimization to Obtain a Width-Independent, Parallel, Simpler, and Faster Positive Sdp Solver

We study the design of polylogarithmic depth algorithms for approximately solving packing and covering semidefinite programs (or positive SDPs for short). This is a natural SDP generalization of the well-studied positive LP problem. Although positive LPs can be solved in polylogarithmic depth while using only $log^2 n/eps^3$ parallelizable iterations [3], the best known positive SDP solvers due to Jain and Yao [18] require $log^{14}(n)/eps^{13}$ parallelizable iterations. Several alternative solvers have been proposed to reduce the exponents in the number of iterations [19, 29]. However, the correctness of the convergence analyses in these works has been called into question [29], as they both rely on algebraic monotonicity properties that do not generalize to matrix algebra. In this paper, we propose a very simple algorithm based on the optimization framework proposed in [3] for LP solvers. Our algorithm only needs $log^2 n/eps^3$ iterations, matching that of the best LP solver. To surmount the obstacles encountered by previous approaches, our analysis requires a new matrix inequality that extends Lieb-Thirring's inequality, and a sign-consistent, randomized variant of the gradient truncation technique proposed in [2, 3].

Zeyuan Allen-Zhu
MIT CSAIL
zeyuan@csail.mit.edu

Yin Tat Lee
MIT
yintat@mit.edu

Lorenzo Orecchia
Boston University
orecchia@bu.edu

CP27

Improved Cheeger's Inequality Using Vertex Expansion and Expansion Profile

We prove two generalizations of the Cheeger's inequality. The first generalization relates the second eigenvalue to the edge expansion and the vertex expansion of the graph G,

$$\lambda_2 = \Omega(\phi^V(G) \cdot \phi(G)),$$

where $\phi^V(G)$ denotes the robust vertex expansion of G and $\phi(G)$ denotes the edge expansion of G. The second generalization relates the second eigenvalue to the edge expansion and the expansion profile of G, for all $k \geq 2$,

$$\lambda_2 = \Omega(\frac{1}{k} \cdot \phi_k(G) \cdot \phi(G)),$$

where $\phi_k(G)$ denotes the k-way expansion of G. These show that the spectral partitioning algorithm has better performance guarantees when $\phi^V(G)$ is large (e.g. planted random instances) or $\phi_k(G)$ is large (instances with few disjoint non-expanding sets). Both bounds are tight up to a constant factor.

Tsz Chiu Kwok
École polytechnique fédérale de Lausanne
tckwok0@gmail.com

Lap Chi Lau
University of Waterloo
lapchi@uwaterloo.ca

Yin Tat Lee
MIT
yintat@mit.edu

CP27

Approximate Undirected Maximum Flows in O(m Polylog(n)) Time

We give the first O(mpolylog(n)) time algorithms for approximating maximum flows in undirected graphs and constructing polylog(n)-quality cut-approximating hierarchical tree decompositions. Our algorithm invokes existing algorithms for these two problems recursively while gradually incorporating size reductions. These size reductions are in turn obtained via ultra-sparsifiers, which are key tools in solvers for symmetric diagonally dominant (SDD) linear systems.

Richard Peng

MIT
rpeng@cc.gatech.edu

CP27

How to Round Subspaces: A New Spectral Clustering Algorithm

We propose a new spectral clustering algorithm, which can recover a k-partition whose span is $O(\sqrt{OPT})$ close to the input in spectral norm. This implies an algorithm for approximating k-expansion in a graph where each cluster have expansion $O(\lambda_{k+1})$ (λ_{k+1} being $(k+1)^{st}$ eigenvalue of Laplacian). This significantly improves upon the previous algorithms, which required $O(\lambda_{k+1}/k)$.

Ali K. Sinop
Princeton University / IAS
asinop@cs.cmu.edu

CP27

Natural Algorithms for Flow Problems

In the last few years, there has been a significant interest in the computational abilities of *Physarum polycephalum* (a slime mold). This arose from a remarkable experiment which showed that this organism can compute shortest paths in a maze [NTY '00]. Subsequently, the workings of Physarum were mathematically modeled as a dynamical system and algorithms inspired by this model were proposed to solve several graph problems: shortest paths, flows, and linear programs to name a few. Indeed, computer scientists have initiated a rigorous study of these dynamics and a first step towards this was taken by BMV '12 and BBDKM '13 who proved that the Physarum dynamics for the shortest path problem are efficient (when edge-lengths are polynomially bounded). In this paper, we take this further: we prove that the discrete time Physarum-dynamics can also efficiently solve the uncapacitated min-cost flow problems on undirected and directed graphs; problems that are non-trivial generalizations of shortest path. This raises the tantalizing possibility that nature, via evolution, developed algorithms that efficiently solve some of the most complex computational problems, about a billion years before we did.

Damian Straszak
EPFL
damian.straszak@epfl.ch

Nisheeth K. Vishnoi
EPFL
EPFL
nisheeth.vishnoi@epfl.ch

CP28

Robust Positioning Patterns

In this paper, we construct large sequences and matrices with the property that the contents of any small window determine the location of the window, robustly. Such objects have found many applications in practical settings, from positioning of wireless devices to smart pens, and have recently gained some theoretical interest. In this context, we give the first explicit constructions of sequences and matrices with high rate and constant relative distance. Accompanying these efficient constructions, we also give efficient decoding algorithms, which can determine the position of the window given its contents, even if a constant fraction of the contents have been corrupted.

Ross Berkowitz, Swastik Kopparty
Rutgers University
rkb73@math.rutgers.edu, swastik.kopparty@gmail.com

CP28

Towards Optimal Deterministic Coding for Interactive Communication

We study *efficient, deterministic* interactive coding schemes that simulate any interactive protocol both under random and adversarial errors, and can achieve a constant communication rate independent of the protocol length. For channels that flip bits independently with probability $\epsilon < 1/2$, our coding scheme achieves a communication rate of $1-O(\sqrt{H(\epsilon)})$ and a failure probability of $\exp(-n/\log n)$ in length n protocols. Prior to our work, all nontrivial deterministic schemes (either efficient or not) had a rate bounded away from 1. Furthermore, the best failure probability achievable by an efficient deterministic coding scheme with constant rate was only quasi-polynomial, i.e., of the form $\exp(-\log^{O(1)} n)$ (Braverman, ITCS 2012). For channels in which an adversary controls the noise pattern our coding scheme can tolerate $\Omega(1/\log n)$ fraction of errors with rate approaching 1. Once more, all previously known nontrivial deterministic schemes (either efficient or not) in the adversarial setting had a rate bounded away from 1, and no nontrivial efficient deterministic coding schemes were known with any constant rate. Essential to both results is an explicit, efficiently encodable and decodable *systematic tree code* of length n that has relative distance $\Omega(1/\log n)$ and rate approaching 1, defined over an $O(\log n)$-bit alphabet. No nontrivial tree code (either efficient or not) was known to approach rate 1, and no nontrivial distance bound was known for any efficient constant rate tree code.

Ran Gelles
Princeton University
rgelles@cs.princeton.edu

Bernhard Haeupler
Carnegie Mellon U.
haeupler@cs.cmu.edu

Gillat Kol
IAS
gillat.kol@gmail.com

Noga Ron-Zewi
IAS and Rutgers University
nogazewi@ias.edu

Avi Wigderson
Institute for Advanced Studies
avi@ias.edu

CP28

An Improved Bound on the Fraction of Correctable Deletions

We consider codes over fixed alphabets against worst-case symbol deletions. For any fixed $k \geq 2$, we construct a family of codes over alphabet of size k with positive rate, which allow efficient recovery from a worst-case deletion fraction approaching $1 - \frac{2}{k+1}$. In particular, for binary codes, we are able to recover a fraction of deletions approaching

1/3. Previously, even non-constructively the largest deletion fraction known to be correctable with positive rate was $1 - \Theta(1/\sqrt{k})$, and around 0.17 for the binary case. Our result pins down the largest fraction of correctable deletions for k-ary codes as $1 - \Theta(1/k)$, since $1 - 1/k$ is an upper bound even for the simpler model of erasures where the locations of the missing symbols are known. Closing the gap between 1/3 and 1/2 for the limit of worst-case deletions correctable by binary codes remains a tantalizing open question.

Boris Bukh, Venkatesan Guruswami
Carnegie Mellon University
bbukh@math.cmu.edu, guruswami@cmu.edu

CP28

Efficient Low-Redundancy Codes for Correcting Multiple Deletions

We consider the problem of constructing binary codes to recover from k–bit deletions with efficient encoding/decoding, for a fixed k. The single deletion case is well understood, with the Varshamov-Tenengolts-Levenshtein code from 1965 giving an asymptotically optimal construction with $\approx 2^n/n$ codewords of length n, i.e., at most $\log n$ bits of redundancy. However, even for the case of two deletions, there was no known explicit construction with redundancy less than $n^{\Omega(1)}$. For any fixed k, we construct a binary code with $c_k \log n$ redundancy that can be decoded from k deletions in $O_k(n \log^4 n)$ time. The coefficient c_k can be taken to be $O(k^2 \log k)$, which is only quadratically worse than the optimal, non-constructive bound of $O(k)$. We also indicate how to modify this code to allow for a combination of up to k insertions and deletions. We also note that among *linear* codes capable of correcting k deletions, the $(k+1)$-fold repetition code is essentially the best possible.

Joshua Brakensiek, Venkatesan Guruswami, Samuel Zbarsky
Carnegie Mellon University
jbrakens@andrew.cmu.edu, guruswami@cmu.edu, szbarsky@andrew.cmu.edu

CP28

Communication Complexity of Permutation-Invariant Functions

Motivated by the quest for a broader understanding of upper bounds in communication complexity, at least for simple functions, we introduce the class of "permutation-invariant' functions. A partial function $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1,?\}$ is permutation-invariant if for every bijection $\pi : \{1,\ldots,n\} \to \{1,\ldots,n\}$ and every $\mathbf{x}, \mathbf{y} \in \{0,1\}^n$, it is the case that $f(\mathbf{x},\mathbf{y}) = f(\mathbf{x}^\pi, \mathbf{y}^\pi)$. Most of the commonly studied functions in communication complexity are permutation-invariant. For such functions, we present a simple complexity measure (computable in time polynomial in n given an implicit description of f) that describes their communication complexity up to polynomial factors and up to an additive error that is logarithmic in the input size. This gives a coarse taxonomy of the communication complexity of simple functions. Our work highlights the role of the well-known lower bounds of functions such as 'Set-Disjointness' and 'Indexing', while complementing them with the relatively lesser-known upper bounds for 'Gap-Inner-Product' (from the sketching literature) and 'Sparse-Gap-Inner-Product' (from the recent work of Canonne et al. [ITCS 2015]). We also present consequences to the study of communication complexity with imperfectly shared randomness where we show that for total permutation-invariant functions, imperfectly shared randomness results in only a polynomial blow-up in communication complexity after an additive $O(\log \log n)$ overhead.

Badih Ghazi
MIT
badih@mit.edu

Pritish Kamath
Massachusetts Institute of Technology
pritish@mit.edu

Madhu Sudan
Harvard
madhu@cs.harvard.edu

CP29

Sampling on Lattices with Free Boundary Conditions Using Randomized Extensions

Many statistical physics models are defined on an infinite lattice by taking appropriate limits of finite lattice regions, where a key consideration is how the boundaries are defined. For several models on planar lattices, such as 3-colorings and lozenge tilings, efficient sampling algorithms are known for regions with *fixed boundary conditions*, where the colors or tiles around the boundary are pre-specified (Luby et al., 2002), but much less is known about how to sample when these regions have *free boundaries*, where we want to include all configurations one could see within a finite window. We introduce a method using *randomized extensions* of a lattice region to relate sampling problems on regions with free boundaries to a constant number of sampling problems on larger regions with fixed boundaries. We demonstrate this principled approach to sample 3-colorings of regions of $\mathbb{Z}^2$ and lozenge tilings of regions of the triangular lattice, building on arguments for the fixed boundary cases due to Luby et al. Our approach also yields an efficient algorithm for sampling 3-colorings with free boundary conditions on regions with one reflex corner, the first such result for a nonconvex region. Our approach can also be generalized to a broad class of mixed boundary conditions. Sampling for these families of regions is significant because it allows us to establish self-reducibility, giving the first algorithm to approximately count the total number of 3-colorings of rectangular lattice regions.

Sarah Cannon, Dana Randall
Georgia Institute of Technology
sarah.cannon@gatech.edu, randall@cc.gatech.edu

CP29

On the Switch Markov Chain for Perfect Matchings

We study a simple Markov chain, the switch chain, on the set of all perfect matchings in a bipartite graph. This Markov chain was proposed by Diaconis, Graham and Holmes as a possible approach to a sampling problem arising in Statistics. They considered several classes of graphs, and conjectured that the switch chain would mix rapidly for graphs in these classes. Here we settle their conjecture almost completely. We ask: for which graph classes is the Markov chain ergodic and for which is it rapidly mixing? We provide a precise answer to the ergodicity question and close bounds on the mixing question. We show for the first time that the mixing time of the switch chain is polynomial

in the class of *monotone* graphs. This class was identified by Diaconis, Graham and Holmes as being of particular interest in the statistical setting.

Martin E. Dyer
University of Leeds
m.e.dyer@leeds.ac.uk

Mark Jerrum
Queen Mary University of London
m.jerrum@qmul.ac.uk

Haiko Muller
University of Leeds
h.muller@leeds.ac.uk

CP29

Algorithmic and Enumerative Aspects of the Moser-Tardos Distribution

Moser & Tardos have developed a powerful algorithmic approach (henceforth "MT") to the Lovasz Local Lemma (LLL); the basic operation done in MT and its variants is a search for "bad" events in a current configuration. In the initial stage of MT, the variables are set independently. We examine the distributions on these variables which arise during intermediate stages of MT. We show that these configurations have a more or less "random' form, building further on the "MT-distribution" concept of Haeupler et al. in understanding the (intermediate and) output distribution of MT. This has a variety of algorithmic applications; the most important is that bad events can be found relatively quickly, improving upon MT across the complexity spectrum: it makes some polynomial-time algorithms sublinear (e.g., for Latin transversals, which are of basic combinatorial interest), gives lower-degree polynomial run-times in some settings, transforms certain super-polynomial-time algorithms into polynomial-time ones, and leads to Las Vegas algorithms for some coloring problems for which only Monte Carlo algorithms were known. We show that in certain conditions when the LLL condition is violated, a variant of the MT algorithm can still produce a distribution which avoids most of the bad events. We show in some cases this MT variant can run faster than the original MT algorithm itself, and develop the first-known criterion for the case of the asymmetric LLL. This can be used to find partial Latin transversals – improving upon earlier bounds of Stein (1975) – among other applications. We furthermore give applications in enumeration, showing that most applications (where we aim for all or most of the bad events to be avoided) have many more solutions than known before by proving that the MT-distribution has "large" Renyi entropy and hence that its support-size is large.

David Harris
University of Maryland
davidgharris29@hotmail.com

CP29

Partial Resampling to Approximate Covering Integer Programs

We consider positive covering integer programs, which generalize set cover and which have attracted a long line of research developing (randomized) approximation algorithms. Srinivasan (2006) gave a rounding algorithm based on the FKG inequality for systems which are "column-sparse." This algorithm may return an integer solution in which the variables get assigned large (integral) values; Kolliopoulos & Young (2005) modified this algorithm to limit the solution size, at the cost of a worse approximation ratio. We develop a new rounding scheme based on the Partial Resampling variant of the Lovasz Local Lemma developed by Harris & Srinivasan (2013). This achieves an approximation ratio of $1 + \frac{\ln(\Delta_1+1)}{a_{\min}} + O(\sqrt{\frac{\log(\Delta_1+1)}{a_{\min}}})$, where $a_{\min}$ is the minimum covering constraint and Δ_1 is the maximum ℓ_1-norm of any column of the covering matrix (whose entries are scaled to lie in $[0, 1]$); we also show nearly-matching inapproximability and integrality-gap lower bounds.

Our approach improves asymptotically, in several different ways, over known results. First, it replaces Δ_0, the maximum number of nonzeroes in any column (from the result of Srinivasan) by Δ_1 which is always – and can be much – smaller than Δ_0; this is the first such result in this context. Second, our algorithm automatically handles multi-criteria programs; we achieve improved approximation ratios compared to the algorithm of Srinivasan, and give, for the first time when the number of objective functions is large, polynomial-time algorithms with good multi-criteria approximations. We also significantly improve upon the upper-bounds of Kolliopoulos & Young when the integer variables are required to be within $(1+\epsilon)$ of some given upper-bounds, and show nearly-matching inapproximability.

David Harris
University of Maryland
davidgharris29@hotmail.com

CP29

Focused Stochastic Local Search and the Lovász Local Lemma

We develop tools for analyzing focused stochastic local search algorithms. These are algorithms which search a state space probabilistically by repeatedly selecting a constraint that is violated in the current state and moving to a random nearby state which, hopefully, addresses the violation without introducing many new ones. A large class of such algorithms arise from the algorithmization of the Lovász Local Lemma, a non-constructive tool for proving the existence of satisfying states. Here we give tools that provide a unified analysis of such algorithms and of many more, expressing them as instances of a general framework.

Dimitris Achlioptas
University of California, Santa Cruz
optas@cs.ucsc.edu

Fotis Iliopoulos
University of California, Berkeley
fotis.iliopoulos@berkeley.edu

CP30

Weak Duality for Packing Edge-Disjoint Odd (u, v)-Trails

There is no packing-covering duality for odd edge-disjoint (u, v)-paths: a graph with no two edge-disjoint (u, v)-paths may need an arbitrarily large number of edges to cover all such paths. In contrast, we show that the relaxed problem of packing odd trails does have an approximate duality. The proof leads to a polynomial-time algorithm to find, for any given k, either k edge-disjoint odd (u, v)-trails or a

set of fewer than $8k$ edges intersecting all odd (u, v)-trails.

Ross Churchley, Bojan Mohar
Simon Fraser University
rchurchl@sfu.ca, mohar@sfu.ca

Hehui Wu
University of Mississippi
hhwu@olemiss.edu

CP30
The k-Mismatch Problem Revisited

We revisit the complexity of one of the most basic problems in pattern matching. In the k-mismatch problem we must compute the Hamming distance between a pattern and every pattern-length substring of a text, as long as that Hamming distance is at most k. Whenever the Hamming distance is greater than k, we simply output "No'. We give improved algorithms for this problem in both the standard offline setting and also in the streaming model.

Raphael Clifford, Allyx Fontaine
University of Bristol, UK
clifford@cs.bris.ac.uk, allyx.fontaine@bristol.ac.uk

Ely Porat
Bar Ilan University
porately@cs.biu.ac.il

Benjamin G. Sach, Tatiana Starikovskaya
University of Bristol, UK
sach@cs.bris.ac.uk, tat.starikovskaya@gmail.com

CP30
Communication with Contextual Uncertainty

We introduce a simple model illustrating the role of context in communication and the challenge posed by uncertainty of knowledge of context. We consider a variant of distributional communication complexity where Alice gets some information x and Bob gets y, where (x, y) is drawn from a known distribution, and Bob wishes to compute some function $g(x, y)$ (with high probability over (x, y)). In our variant, Alice does not know g, but only knows some function f which is an approximation of g. Thus, the function being computed forms the context for the communication, and knowing it imperfectly models (mild) uncertainty in this context. A naive solution would be for Alice and Bob to first agree on some common function h that is close to both f and g and then use a protocol for h to compute $h(x, y)$. We show that any such agreement leads to a large overhead in communication ruling out such a universal solution. In contrast, we show that if g has a one-way communication protocol with complexity k in the standard setting, then it has a communication protocol with complexity $O(k \cdot (1 + I))$ in the uncertain setting, where I denotes the mutual information between x and y. In the particular case where the input distribution is a product distribution, the protocol in the uncertain setting only incurs a constant factor blow-up in communication and error. Furthermore, we show that the dependence on the mutual information I is required. Namely, we construct a class of functions along with a non-product distribution over (x, y) for which the communication complexity is a single bit in the standard setting but at least $\Omega(\sqrt{n})$ bits in the uncertain setting.

Badih Ghazi
MIT
badih@mit.edu

Ilan Komargodski
Weizmann Institute of Science
ilan.komargodski@weizmann.ac.il

Pravesh Kothari
Department of Computer Science
University of Texas at Austin
kothari@cs.utexas.edu

Madhu Sudan
Harvard
madhu@cs.harvard.edu

CP30
On the Maximum Quartet Distance Between Phylogenetic Trees

A conjecture of Bandelt and Dress states that the maximum quartet distance between any two phylogenetic trees on n leaves is at most $(\frac{2}{3} + o(1))\binom{n}{4}$. Using the machinery of flag algebras we improve the currently known bounds regarding this conjecture, in particular we show that the maximum is at most $(0.69 + o(1))\binom{n}{4}$. We also give further evidence that the conjecture is true by proving that the maximum distance between caterpillar trees is at most $(\frac{2}{3} + o(1))\binom{n}{4}$.

Noga Alon
Tel-Aviv University
nogaa@post.tau.ac.il

Humberto Naves, Benjamin Sudakov
ETH Zurich
hnaves@ima.umn.edu, benjamin.sudakov@math.ethz.ch

CP30
Range Predecessor and Lempel-Ziv Parsing

The Lempel-Ziv (LZ77) parsing of a string is one of the most important algorithmic tools in data compression and string processing. We show that the Lempel-Ziv parsing of a string of length n on an alphabet of size σ can be computed in $O(n \log\log \sigma)$ time using $O(n \log \sigma)$ bits of working space. The previous fastest algorithm using $O(n \log \sigma)$ space takes $O(n(\log \sigma + \log\log n))$ time. We also consider the important *rightmost* variant of the problem, where the goal is to associate with each phrase of the parsing its *most recent* occurrence in the input string. We solve this problem in $O(n(1 + (\log \sigma/\sqrt{\log n}))$ time, using the same working space as above. The previous best solution for rightmost parsing uses $O(n(1 + \log \sigma/\log\log n))$ time and $O(n \log n)$ space. As a bonus, we provide a faster construction method for efficient *2D orthogonal range reporting*, which is of independent interest.

Djamal Belazzougui, Simon J. Puglisi
Department of Computer Science
University of Helsinki
djamal.belazzougui@gmail.com, simon.j.puglisi@gmail.com