

IP1**"If You Can Specify It, You Can Analyze It" —The Lasting Legacy of Philippe Flajolet**

The "Flajolet School" of the analysis of algorithms and combinatorial structures is centered on an effective calculus, known as analytic combinatorics, for the development of mathematical models that are sufficiently accurate and precise that they can be validated through scientific experimentation. It is based on the generating function as the central object of study, first as a formal object that can translate a specification into mathematical equations, then as an analytic object whose properties as a function in the complex plane yield the desired quantitative results. Universal laws of sweeping generality can be proven within the framework, and easily applied. Standing on the shoulders of Cauchy, Polya, de Bruijn, Knuth, and many others, Philippe Flajolet and scores of collaborators developed this theory and demonstrated its effectiveness in a broad range of scientific applications. Flajolet's legacy is a vibrant field of research that holds the key not just to understanding the properties of algorithms and data structures, but also to understanding the properties of discrete structures that arise as models in all fields of science. This talk will survey Flajolet's story and its implications for future research.

Robert Sedgewick
Princeton University
rs@CS.Princeton.EDU

IP2**Submodular Functions and Their Applications**

Submodular functions, a discrete analogue of convex functions, have played a fundamental role in combinatorial optimization since the 1970s. In the last decade, there has been renewed interest in submodular functions due to their interpretation as valuation functions of self-interested agents in algorithmic game theory. These developments have led to new questions as well as new algorithmic techniques. In this talk, we will discuss the concept of submodularity, its motivation and its unifying role in combinatorial optimization, as well as the evolution of the relevant algorithmic techniques. We will survey the state of the art in optimization of submodular functions, as well as selected applications in algorithmic game theory, social networks and machine learning, and some future challenges.

Jan Vondrak
IBM Almaden Research Center
jvondrak@us.ibm.com

IP3**On Graph Property Testing, Arithmetic Progressions and Communication**

Tools from Extremal Graph Theory are helpful in the study of problems in Additive Number Theory, Theoretical Computer Science, and Information Theory. I will illustrate this fact by several closely related examples focusing on a recent one in a joint work with Moitra and Sudakov. The main combinatorial question addressed in this work is that of estimating the maximum possible density of a graph consisting of a union of pairwise edge disjoint large induced matchings. It turns out that variants of this problem appear in the investigation of several seemingly unrelated topics. The lecture will include a brief general introduction, a description of the methods used and results obtained, as

well as a discussion of the remaining challenges.

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

CP0**Inducing Suffix and Lcp Arrays in External Memory**

We consider text index construction in external memory (EM). Our first contribution is an inducing algorithm for suffix arrays in external memory. Practical tests show that this outperforms the previous best EM suffix sorter [Dementiev et al., ALENEX 2005] by a factor of 2-3 in time and I/O-volume. Our second contribution is to augment the first algorithm to also construct the array of longest common prefixes (LCPs). This yields the first EM construction algorithm for LCP arrays. The overhead in time and I/O volume for this extended algorithm over plain suffix array construction is roughly two. Our algorithms scale far beyond problem sizes previously considered in the literature (text size of 80 GiB using only 4 GiB of RAM in our experiments).

Johannes Fischer
KIT
76131 Karlsruhe, Germany
johannes.fischer@kit.edu

Timo Bingmann, Vitaly Osipov
KIT
timo.bingmann@kit.edu, vitaly.osipov@kit.edu

CP0**The Number of Ways to Assemble a Graph**

Motivated by the question of how macromolecules assemble, the notion of an *assembly tree* of a graph is introduced. Given a graph G , the paper is concerned with enumerating the number of assembly trees of G . Explicit formulas or generating functions are provided for the number of assembly trees of several families of graphs. In some special cases, we apply results of Zeilberger and Apagodu on multivariate generating functions, and results of Wimp and Zeilberger to deduce recurrence relations and very precise asymptotic formulas for the number of assembly trees of complete bipartite and tripartite graphs.

Miklos Bona
University of Florida
bona@ufl.edu

Andrew Vince
Department of Mathematics
University of Florida
avince@ufl.edu

CP0**Bootstrap Percolation on Random Geometric Graphs**

Bootstrap percolation has been used effectively to model phenomena as diverse as emergence of magnetism in materials, spread of infection, diffusion of software viruses in computer networks, adoption of new technologies, and emergence of collective action and cultural fads in human societies. It is defined on an (arbitrary) network of inter-

acting agents whose state is determined by the state of their neighbors according to a threshold rule. In a typical setting, bootstrap percolation starts by random and independent ‘activation’ of nodes with a fixed probability p , followed by a deterministic process for additional activations based on the density of active nodes in each neighborhood (θ activated nodes). Here, we study bootstrap percolation on random geometric graphs in the regime when the latter are (almost surely) connected. Random geometric graphs provide an appropriate model in settings where the neighborhood structure of each node is determined by geographical distance, as in wireless *ad hoc* and sensor networks as well as in contagion. We derive bounds on the critical thresholds p'_c, p'_c such that for all $p > p'_c(\theta)$ full percolation takes place, whereas for $p < p'_c(\theta)$ it does not. We conclude with simulations that compare numerical thresholds with those obtained analytically.

Milan Bradonjic
Mathematics of Networks and Communications
Bell Labs, Alcatel-Lucent
bradonjic@gmail.com

Iraj Saniee
Bell Labs, Alcatel-Lucent
iis@research.bell-labs.com

CP0

A Min-Edge Cost Flow Framework for Capacitated Covering Problems

We present the Cov-MECF framework, a special case of minimum-edge cost flow capturing several important covering problems. We also introduce a heuristic LPO, which harnesses the fractional relaxation to greedily build its solution. We empirically establish that LPO returns higher quality solutions than Wolsey’s classical greedy algorithm. We also observe that Leskovec et. al.’s optimization benefits LPO more than Wolsey’s algorithm, and that a freez-ing step heavily reduces the number of subroutine calls for both.

Jessica Chang
University of Washington
jschang@cs.washington.edu

Samir Khuller
University of Maryland, College Park
samir@cs.umd.edu

CP0

Extremal Parameters in Sub-Critical Graph Classes

We analyze several extremal parameters like the diameter or the maximum degree in sub-critical graph classes. Sub-critical graph classes cover several well-known classes of graphs like trees, outerplanar graph or series-parallel graphs which have been intensively studied during the last few years. However, this paper is the first one, where these kinds of parameters are studied from a general point of view.

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

Marc Noy
UPC Barcelona

marc.noy@upc.edu

CP0

Fast Packed String Matching for Short Patterns

Searching for all occurrences of a pattern in a text is a fundamental problem in computer science with applications in many other fields, like natural language processing, information retrieval and computational biology. In the last two decades a general trend has appeared trying to exploit the power of the word RAM model to speed-up the performances of classical string matching algorithms. In this model an algorithm operates on words of length w , grouping blocks of characters, and arithmetic and logic operations on the words take one unit of time. In this paper we use specialized word-size packed string matching instructions, based on the Intel streaming SIMD extensions (SSE) technology, to design very fast string matching algorithms in the case of short patterns. From our experimental results it turns out that, despite their quadratic worst case time complexity, the new presented algorithms become the clear winners on the average for short patterns, when compared against the most effective algorithms known in literature.

Simone Faro
Dipartimento di Matematica e Informatica
Università di Catania
faro@dmi.unict.it

Oguzhan Kulekci
Nation. Res. Institute of Electronics and Cryptology
oguzhan.kulekci@tubitak.gov.tr

CP0

On Unifying the Space of ℓ_0 -Sampling Algorithms

The problem of building an ℓ_0 -sampler is to sample near-uniformly from the support set of a dynamic multiset. This problem has a variety of applications within data analysis, computational geometry and graph algorithms. In this paper, we abstract a set of steps for building an ℓ_0 -sampler, based on sampling, recovery and selection. We analyze the implementation of an ℓ_0 -sampler within this framework, and show how prior constructions of ℓ_0 -samplers can all be expressed in terms of these steps. Our experimental contribution is to provide a first detailed study of the accuracy and computational cost of ℓ_0 -samplers.

Graham Cormode
AT&T Labs-Research
graham@research.att.com

Donatella Firmani
Sapienza University of Rome
firmani@dis.uniroma1.it

CP0

Short and Simple Cycle Separators in Planar Graphs

We provide an implementation of an algorithm that, given a triangulated planar graph with m edges, returns a simple cycle that is a $2/3$ -balanced separator consisting of at most $\sqrt{8m}$ edges. Cycle separators are essential in numerous planar graph algorithms, e.g., for computing shortest paths, minimum cuts, or maximum flows. We demonstrate that our algorithm is competitive with, and sometimes outperforms other planar separator algorithms. Our implemen-

tation scales well to large graphs.

Eli Fox-Epstein
Tufts University
ef@cs.tufts.edu

Shay Mozes, Phitchaya Phothilimthana
Massachusetts Institute of Technology
shaym@mit.edu, mangpo@csail.mit.edu

Christian Sommer
None
csom@csail.mit.edu

CP0

Approximating Fault-Tolerant Domination in General Graphs

We study the NP-complete problem of finding small k -dominating sets in general graphs, which allow $k - 1$ nodes to fail and still dominate the graph. We show that the approach of having at least k dominating nodes in the neighborhood of every node can be give solutions off by a multiplicative factor of nearly k and also show that a greedy algorithm achieves an almost tight approximation ratio of $\ln(\Delta + k) + 1$.

Klaus-Tycho Förster
ETH Zürich, Distributed Computing Group
k-t.foerster@tik.ee.ethz.ch

CP0

The Variance of the Number of 2-Protected Nodes in a Trie

We derive an asymptotic expression for the variance of the number of 2-protected nodes in a binary trie. In an *unbiased* trie on n leaves we find that the variance is approximately $.934n$ plus small fluctuations (of order n); but our result covers the general (biased) case as well. Our proof relies on the asymptotic similarities between a trie and its Poissonized counterpart, using Mellin transform and singularity analysis.

Jeffrey B. Gaither
Purdue University
Department of Mathematics
jgaither@math.purdue.edu

Mark Daniel Ward
Purdue University
Department of Statistics
mdw@purdue.edu

CP0

When Is It Worthwhile to Propagate a Constraint? A Probabilistic Analysis of Alldifferent

We present new work on analyzing the behaviour of a constraint solver, with a view towards optimization. We propose to quantify useless calls to propagator for the ALLDIFFERENT constraint (bound consistency propagator). Our first contribution is the definition of a probabilistic model for the constraint and its variables. This model is used to compute the probability that a call to the propagation algorithm for ALLDIFFERENT modifies its input,

then give an asymptotic approximation of it.

Daniele Gardy
University of Versailles Saint-Quentin (France)
daniele.gardy@prism.uvsq.fr

Jeremie Du Boisberranger
PRISM, University of Versailles St-Quentin
jdb@prism.uvsq.fr

Xavier Lorca
Ecole des Mines, Nantes
xavier.lorca@mines-nantes.fr

Charlotte Truchet
LINA, University of Nantes
charlotte.truchet@univ-nantes.fr

CP0

Exact-Size Sampling for Motzkin Trees in Linear Time Via Boltzmann Samplers and Holonomic Specification

Given a combinatorial class and an algebraic combinatorial specification, one can automatically build a so-called Boltzmann random sampler for that class. We introduce a Boltzmann sampler for Motzkin trees built from a holonomic specification, that is, a specification that uses the pointing operator. This sampler is inspired by Rmy's algorithm on binary trees. We show that our algorithm gives an exact size sampler with a linear time and space complexity in average.

Axel Bacher, Olivier Bodini, Alice Jacquot
LIPN, Université Paris 13
axel.bacher@lipn.univ-paris13.fr, olivier.bodini@lipn.univ-paris13.fr, alice.jacquot@lipn.univ-paris13.fr

CP0

The Cost of Address Translation

Modern computers with memory hierarchy, multiple cores, and virtual memory are not random access machines (RAMs). Hence, the analysis of some algorithms (random scan of an array, binary search, heapsort) in the RAM model does not correctly predict growth rates of actual running times. We propose the VAT model (virtual address translation) that accounts for the cost of address translations. Then, we analyze the VAT-cost of the algorithms mentioned above, cache-oblivious algorithms, and others.

Tomasz Jurkiewicz
Max Planck Institute for Informatics
The Saarbrücken Graduate School of Computer Science
tomasz.jurkiewicz@mpi-inf.mpg.de

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

CP0

3D Kinetic Alpha Complexes and Their Implementation

Motivated by an application in cell biology, we describe an extension of the kinetic data structures framework from Delaunay triangulations to fixed-radius alpha complexes.

Our algorithm is implemented using CGAL, following the exact geometric computation paradigm. We report on several techniques to accelerate the computation that turn our implementation applicable to the underlying biological problem.

Michael Kerber, Herbert Edelsbrunner
Institute of Science and Technology (Austria)
mkerber@mpi-inf.mpg.de, edels@ist.ac.at

CP0

Unlabeled Equivalence for Matroids Representable over Finite Fields

Two matrices representing the same matroid over a finite field are projective equivalent if one can be obtained from the other by elementary row operations and column scaling. Bounds for the number of projective inequivalent representations are known for only a few classes of matroids. We define two matrices to be geometric equivalent if column permutations are also allowed, and show that from an algorithmic perspective this is the better definition of equivalence for matroids.

Sandra R. Kingan
Brooklyn College
CUNY
skingan@brooklyn.cuny.edu

CP0

On Delta-Method of Moments and Probabilistic Sums

Abstract not available at time of publication.

Jacek Cichon, Zbigniew Golebiewski, Marcin Kardas,
Marek Klonowski
Wroclaw University of Technology
jacek.cichon@pwr.wroc.pl, n/a,
zbigniew.golebiewski@pwr.wroc.pl,
marek.klonowski@pwr.wroc.pl

CP0

Analysis of Parameters of Trees Corresponding to Huffman Codes and Sums of Unit Fractions

For fixed $t \geq 2$, we consider the class of representations of 1 as sum of unit fractions whose denominators are powers of t or equivalently the class of canonical compact t -ary Huffman codes or equivalently rooted t -ary plane “canonical” trees. We study the probabilistic behaviour of the height (limit distribution is shown to be normal), the number of distinct summands (normal distribution), the path length (normal distribution), the width (main term of the expectation and concentration property) and the number of leaves at maximum distance from the root (discrete distribution).

Clemens Heuberger
Alpen-Adria-Universität Klagenfurt, Austria
clemens.heuberger@aa.u.at

Daniel Krenn
TU Graz
Austria
krenn@math.tugraz.at

Stephan Wagner

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

CP0

Practical Batch-Updatable External Hashing with Sorting

This paper presents a practical external hashing scheme that supports fast lookup (7 microseconds) for large datasets (millions to billions of items) with a small memory footprint (2.5 bits/item) and fast index construction (151 K items/s for 1-KiB key-value pairs). Our scheme combines three key techniques: (1) a new index data structure (Entropy-Coded Tries); (2) the use of sorting as the main data manipulation method; and (3) support for incremental index construction for dynamic datasets. We evaluate our scheme by building an external dictionary on flash-based drives and demonstrate our scheme’s high performance, compactness, and practicality.

Hyeontaek Lim, David Andersen
Carnegie Mellon University
hl@cs.cmu.edu, dga@cs.cmu.edu

Michael Kaminsky
Intel Labs
michael.e.kaminsky@intel.com

CP0

Perpetuities in Fair Leader Election Algorithms

We consider a broad class of fair leader election algorithms, and study the duration of contestants (the number of rounds a randomly selected contestant stays in the competition) and the overall cost of the algorithm. We give sufficient conditions for the duration to have a geometric limit distribution (a perpetuity built from Bernoulli random variables), and for the limiting distribution of the total cost (after suitable normalization) to be a perpetuity. For the duration, the proof is established via convergence (to 0) of the first-order Wasserstein distance from the geometric limit. For the normalized overall cost, the method of proof is also convergence of the first-order Wasserstein distance, augmented with an argument based on a contraction mapping in the first-order Wasserstein metric space to show that the limit approaches a unique fixed-point solution of a perpetuity distributional equation. (The two steps are commonly referred to as the contraction method.)

Hosam M. Mahmoud, Ravi Kalpathy
The George Washington University
hosam@gwu.edu, kalpathy@gwu.edu

CP0

Robust Gossip-Based Aggregation: A Practical Point of View

Several existing gossip-based aggregation algorithms have the potential for providing resilience in failure-prone distributed systems. In this paper, we evaluate the algorithms with the most promising theoretical fault tolerance properties experimentally in faulty environments. Our experiments reveal several shortcomings of existing approaches in terms of performance as well as in terms of numerical accuracy. We explain this behavior observed in practice and motivate further theoretical research to efficiently meet the

desired robustness requirements.

Gerhard Niederbrucker

University of Vienna
gerhard.niederbrucker@univie.ac.at

Wilfried N. Gansterer

University of Vienna
Department of Computer Science and Business
Informatics
wilfried.gansterer@univie.ac.at

CP0

Lempel-Ziv Factorization: Simple, Fast, Practical

For decades the Lempel-Ziv (LZ77) factorization has been a cornerstone of data compression and string processing, and uses for it are still being uncovered. For example, LZ77 is central to recent index data structures for highly repetitive collections. However, in many applications factorization remains a computational bottleneck. We describe several fast factorization algorithms which consistently outperform all previous approaches by almost a factor of two, use less memory, and still offer strong worstcase performance guarantees.

Simon J. Puglisi, Dominik Kempa

Department of Computer Science
University of Helsinki
simon.j.puglisi@gmail.com, dominik.kempa@gmail.com

CP0

Approximate Counting of Matchings in Sparse Uniform Hypergraphs

In this paper we give a fully polynomial randomized approximation scheme (FPRAS) for the number of matchings in k -uniform hypergraphs whose intersection graphs contain few claws. Our method gives a generalization of the canonical path method of Jerrum and Sinclair to hypergraphs satisfying a local restriction. The proof depends on an application of the Euler tour technique for the canonical paths of the underlying Markov chains. On the other hand, we prove that it is NP-hard to approximate the number of matchings even for the class of 2-regular, linear, k -uniform hypergraphs, for all $k \geq 6$, without the above restriction.

Edyta Szymanska

Adam Mickiewicz University, Poznan, Poland
edka@amu.edu.pl

Andrzej Rucinski

Adam Mickiewicz University, Poznan, Poland
Emory University, Atlanta
rucinski@amu.edu.pl

Marek Karpinski

Department of Computer Science
University of Bonn
marek@cs.uni-bonn.de

CP0

Computational Topology and Normal Surfaces: Theoretical and Experimental Complexity Bounds

In three-dimensional computational topology, the theory of normal surfaces is a tool of great theoretical and practical significance. Although this theory typically leads to ex-

ponential time algorithms, very little is known about how these algorithms perform in “typical” scenarios, or how far the best known theoretical bounds are from the real worst-case scenarios. Here we study the combinatorial and algebraic complexity of normal surfaces from both the theoretical and experimental viewpoints. Theoretically, we obtain new exponential lower bounds on the worst-case complexities in a variety of settings that are important for practical computation. Experimentally, we study the worst-case and average-case complexities over a comprehensive body of roughly three billion input triangulations. Many of our lower bounds are the first known exponential lower bounds in these settings, and experimental evidence suggests that many of our theoretical lower bounds on worst-case growth rates may indeed be asymptotically tight.

Benjamin A. Burton, Jonathan Spreer

The University of Queensland
bab@maths.uq.edu.au, j.spreer@uq.edu.au

Joo Paixo

Pontifcia Universidade Católica, Rio de Janeiro
jpaixao@mat.puc-rio.br

CP0

Polynomial-Time Construction of Contraction Hierarchies for Multi-Criteria Objectives

We consider multi-criteria shortest path problems and show that contraction hierarchies – a speed-up technique originally developed for standard shortest path queries – can be constructed efficiently for the case of linearly weighted combinations of the edge costs. This extends previous results which considered only the bicriteria case. On the theory side we prove a polynomial time bound via some polyhedral considerations; experiments complement these results by showing the practicability of our approach.

Stefan Funke

FMI, University of Stuttgart, Germany
funke@fmi.uni-stuttgart.de

Sabine Storandt

University of Freiburg
storandt@informatik.uni-freiburg.de

CP0

On Parallelizing Matrix Multiplication by the Column-Row Method

We consider the problem of sparse matrix multiplication by the column row method in a distributed setting where the matrix product is not necessarily sparse. We present a surprisingly simple method for “consistent” parallel processing of sparse outer products (column-row vector products) over several processors, in a *communication-avoiding* setting where each processor has a copy of the input. The method is consistent in the sense that a given output entry is always assigned to the same processor independently of the specific structure of the outer product. We show guarantees on the work done by each processor, and achieve linear speedup down to the point where the cost is dominated by reading the input. Our method gives a way of distributing (or parallelizing) matrix product computations in settings where the main bottlenecks are storing the result matrix, and inter-processor communication. Motivated by observations on real data that often the absolute values of the entries in the product adhere to a power law, we combine our approach with frequent items mining algo-

gorithms and show how to obtain a tight approximation of the weight of the heaviest entries in the product matrix. As a case study we present the application of our approach to frequent pair mining in transactional data streams, a problem that can be phrased in terms of sparse $\{0, 1\}$ -integer matrix multiplication by the column-row method. Experimental evaluation of the proposed method on real-life data supports the theoretical findings.

Konstantin Kutzkov, Andrea Campagna, Rasmus Pagh, Morten Stöckel
IT University of Copenhagen
konk@itu.dk, acam@itu.dk, pagh@itu.dk, mstc@itu.dk

CP0

Efficient Algorithms for Dualizing Large-Scale Hypergraphs

A hypergraph \mathcal{F} is a set family defined on vertex set V . The dual of \mathcal{F} is the set of minimal subsets H of V such that $F \cap H \neq \emptyset$ for any $F \in \mathcal{F}$. We propose efficient dualization algorithms using a new set system induced by the minimality condition, and efficient methods for the minimality check. The computational experiments show that our algorithms outperform to existing algorithms drastically.

Takeaki Uno
National Institute of Informatics
uno@nii.jp

Keisuke Murakami
Aoyama Gakuin University
murakami@ise.aoyama.ac.jp

CP0

Engineering Java 7's Dual Pivot Quicksort Using Malijan

Recent results on Java 7's dual pivot Quicksort have revealed its highly asymmetric nature. These insights suggest that asymmetric pivot choices are preferable to symmetric ones for this Quicksort variant. From a theoretical point of view, this should allow us to improve on the current implementation in Oracle's Java 7 runtime library. In this paper, we use our new tool MaLiJAn to confirm this for asymptotic combinatorial cost measures such as the total number of executed instructions. However, the observed running times show converse behavior. With the support of data provided by MaLiJAn we are able to identify the profiling capabilities of Oracle's just-in-time compiler to be responsible for this unexpected outcome.

Sebastian Wild, Markus Nebel, Raphael Reitzig, Ulrich Laube
Computer Science Department
University of Kaiserslautern
s.wild@cs.uni-kl.de, nebel@cs.uni-kl.de, r.reitzig@cs.uni-kl.de, laube@cs.uni-kl.de

CP1

Convergence of Multivariate Belief Propagation, with Applications to Cuckoo Hashing and Load Balancing

This paper is motivated by two applications, namely i) generalizations of cuckoo hashing, a computationally simple approach to assigning keys to objects, and ii) load balanc-

ing in content distribution networks, where one is interested in determining the impact of content replication on performance. These two problems admit a common abstraction: in both scenarios, performance is characterized by the maximum weight of a generalization of a matching in a bipartite graph, featuring node and edge capacities. Our main result is a law of large numbers characterizing the asymptotic maximum weight matching in the limit of large bipartite random graphs, when the graphs admit a *local weak limit* that is a tree. This result specializes to the two application scenarios, yielding new results in both contexts. In contrast with previous results, the key novelty is the ability to handle edge capacities with arbitrary integer values. An analysis of belief propagation algorithms (BP) with multivariate belief vectors underlies the proof. In particular, we show convergence of the corresponding BP by exploiting monotonicity of the belief vectors with respect to the so-called *upshifted likelihood ratio* stochastic order. This auxiliary result can be of independent interest, providing a new set of structural conditions which ensure convergence of BP.

Mathieu Leconte
Technicolor - INRIA
mathieu.leconte@inria.fr

Marc Lelarge
INRIA-ENS
marc.lelarge@ens.fr

Laurent Massoulié
Inria, Microsoft research - Inria Joint Center
laurent.massoulie@inria.fr

CP1

Mixing Times of Markov Chains for Self-Organizing Lists and Biased Permutations

We study the mixing time of a Markov chain $\mathcal{M}_{\setminus\setminus}$ on permutations in the non-uniform setting. We are given "positively biased" probabilities $\{p_{i,j} \geq 1/2\}$ for all $i < j$ and let $p_{j,i} = 1 - p_{i,j}$. In each step $\mathcal{M}_{\setminus\setminus}$ chooses two adjacent elements k , and ℓ and exchanges their positions with probability $p_{\ell,k}$. Here we define two general classes and give the first proofs that the chain is rapidly mixing for both. We also demonstrate that the chain can mix slowly.

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

Amanda Streib
National Institute of Standards and Technology
amanda.streib@nist.gov

CP1

Balls into Bins Via Local Search

We propose a natural process for allocating n balls into n bins that are organized as the vertices of an undirected graph G . Each ball first chooses a vertex u in G uniformly at random. Then the ball performs a local search in G starting from u until it reaches a vertex with local minimum load, where the ball is finally placed on. In our main result, we prove that this process yields a maximum load of only $\Theta(\log \log n)$ on expander graphs. In addition, we show that for d -dimensional grids the maximum load

is $\Theta\left(\left(\frac{\log n}{\log \log n}\right)^{\frac{1}{d+1}}\right)$. Finally, for almost regular graphs with minimum degree $\Omega(\log n)$, we prove that the maximum load is constant and also reveal a fundamental difference between random and arbitrary tie-breaking rules.

Paul Bogdan
Carnegie Mellon University
pbogdan@ece.cmu.edu

Thomas Sauerwald
Max Planck Institute for Computer Science
sauerwal@mpi-inf.mpg.de

Alexandre Stauffer
Microsoft Research
alstauff@microsoft.com

He Sun
Max Planck Institute for Informatics
hsun@mpi-inf.mpg.de

CP1

Approximate Counting Via Correlation Decay on Planar Graphs

We show for a broad class of counting problems, correlation decay (strong spatial mixing) implies FPTAS on planar graphs. The framework for the counting problems considered by us is the Holant problems with arbitrary constant-size domain and symmetric constraint functions. We define a notion of regularity on the constraint functions, which covers a wide range of natural and important counting problems, including all multi-state spin systems, counting graph homomorphisms, counting weighted matchings or perfect matchings, the subgraphs world problem transformed from the ferromagnetic Ising model, and all counting CSPs and Holant problems with symmetric constraint functions of constant arity. The core of our algorithm is a fixed-parameter tractable algorithm which computes the exact values of the Holant problems with regular constraint functions on graphs of bounded treewidth. By utilizing the locally tree-like property of apex-minor-free families of graphs, the parameterized exact algorithm implies an FPTAS for the Holant problem on these graph families whenever the Gibbs measure defined by the problem exhibits strong spatial mixing. We further extend the recursive coupling technique to Holant problems and establish strong spatial mixing for the ferromagnetic Potts model and the subgraphs world problem. As consequences, we have new deterministic approximation algorithms on planar graphs and all apex-minor-free graphs for several counting problems.

Yitong Yin
Nanjing University
yinyt@nju.edu.cn

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

CP1

Correlation Decay Up to Uniqueness in Spin Systems

We give a complete characterization of the two-state anti-ferromagnetic spin systems which exhibit strong spatial

mixing on general graphs. We show that a two-state anti-ferromagnetic spin system exhibits strong spatial mixing on all graphs of maximum degree at most Δ if and only if the system has a unique Gibbs measure on infinite regular trees of degree up to Δ , where Δ can be either bounded or unbounded. As a consequence, there exists an FPTAS for the partition function of a two-state anti-ferromagnetic spin system on graphs of maximum degree at most Δ when the uniqueness condition is satisfied on infinite regular trees of degree up to Δ . In particular, an FPTAS exists for arbitrary graphs if the uniqueness is satisfied on all infinite regular trees. This covers as special cases all previous algorithmic results for two-state anti-ferromagnetic systems on general-structure graphs. Combining with the FPRAS for two-state ferromagnetic spin systems of Jerrum-Sinclair and Goldberg-Jerrum-Paterson, and the very recent hardness results of Sly-Sun and independently of Galanis-Štefankovič-Vigoda, this gives a complete classification, except at the phase transition boundary, of the approximability of all two-state spin systems, on either degree-bounded families of graphs or family of all graphs.

Liang Li
Peking University
liang.li@pku.edu.cn

Pinyan Lu
Microsoft Research Asia
pinyanl@microsoft.com

Yitong Yin
Nanjing University
yinyt@nju.edu.cn

CP2

Online Mixed Packing and Covering

Recent advances in online optimization have heavily relied on algorithmic techniques for obtaining near-optimal solutions to LPs that are revealed online. We give the first polylogarithmic-competitive algorithm for obtaining a fractional solution to mixed packing and covering LPs online, where packing constraints are given offline and covering constraints arrive online. Using this generic technique, we give near-optimal algorithms for two canonical problems in this class: unrelated machine scheduling with startup costs and capacity-constrained facility location.

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

Umang Bhaskar
California Institute of Technology
umang@caltech.edu

Lisa Fleischer
Dartmouth College
lkf@cs.dartmouth.edu

Debmalya Panigrahi
Microsoft Research
debmalaya@alum.mit.edu

CP2

Randomized Primal-Dual Analysis of Ranking for

Online Bipartite Matching

We give a simple proof that the algorithm of Karp, Vazirani and Vazirani is $1-1/e$ competitive for the online bipartite matching problem. The proof is via a randomized primal-dual argument. Primal-dual algorithms have been successfully used for many online algorithm problems, but the dual constraints are always satisfied deterministically. This is the first instance of a non-trivial randomized primal-dual algorithm in which the dual constraints only hold in expectation. The approach also generalizes easily to the vertex-weighted version considered by Agarwal et al. Further we show that the proof is very similar to the deterministic primal-dual argument for the online budgeted allocation problem with small bids (also called the AdWords problem) of Mehta et al.

Nikhil R. Devanur

Microsoft Research, Redmond
nikdev@microsoft.com

Kamal Jain

Ebay Research
kamaljain@gmail.com.

Robert Kleinberg

Cornell University
rdk@cs.cornell.edu

CP2

Matroid Secretary for Regular and Decomposable Matroids

In the *matroid secretary problem* we are given a Matroid and a stream of elements. The goal is to choose an independent set with maximum profit. If we choose to accept an element when it is presented by the stream then we can never get rid of it, and if we choose not to accept it then we cannot later add it. We give $O(1)$ -competitive algorithms for *regular* matroids and Min-cut max flow matroids

Michael Dinitz

Computer Science,
michael.dinitz@weizmann.ac.il

Guy Kortsarz

Computer Science
guyk@crab.rutgers.edu

CP2

Weighted Flowtime on Capacitated Machines

Motivated by modern systems which run multiple jobs/VMs at a time subject to overall CPU, memory, network, and disk capacity constraints, we give the first scheduling algorithms to minimize weighted flow time on capacitated machines. We give several positive and negative results, including a near-optimal $(1+\epsilon)$ -speed, $(1+\epsilon)$ -capacity $O(1/\epsilon^3)$ -competitive algorithm using a novel combination of knapsacks, densities, job classification into categories, and potential function methods. Our results also extend to the multiple unrelated machines setting.

Kyle Fox

University of Illinois, Urbana-Champaign
kylefox2@illinois.edu

Madhukar Korupolu

Google Research

mkar@google.com

CP2

A New Approach to Online Scheduling: Approximating the Optimal Competitive Ratio

We propose a new approach to competitive analysis in online scheduling by introducing the novel concept of competitive-ratio approximation schemes. Such a scheme algorithmically constructs an online algorithm with a competitive ratio arbitrarily close to the best possible competitive ratio for *any* online algorithm. We study the problem of scheduling jobs online to minimize the weighted sum of completion times on parallel, related, and unrelated machines, and we derive both deterministic and randomized algorithms which are almost best possible among all online algorithms of the respective settings.

Andreas Wiese

MPII Saarbrücken
awiese@mpi-inf.mpg.de

Elisabeth Günther, Olaf Maurer, Nicole Megow
TU Berlin

eguenth@math.tu-berlin.de, maurer@math.tu-berlin.de,
nmegow@math.tu-berlin.de

CP3

Computing the Discrete Fréchet Distance in Subquadratic Time

The *discrete Fréchet distance* is a similarity measure between two sequences of points A and B , that takes into account the location and ordering of the points along the sequences. We present a subquadratic algorithm for computing the discrete Fréchet distance between two sequences of points in the plane, of respective lengths m and n . The algorithm runs in $O\left(\frac{mn \log \log n}{\log n}\right)$ time and uses $O(n+m)$ storage.

Pankaj Agarwal

Duke University
pankaj@cs.duke.edu

Rinat Ben Avraham

Tel Aviv University
rinatba@gmail.com

Haim Kaplan

Tel-Aviv University
haimk@post.tau.ac.il

Micha Sharir

Tel Aviv University & New York University
michas@post.tau.ac.il

CP3

The Complexity of Detecting Taut Angle Structures on Triangulations

There are many fundamental algorithmic problems on triangulated 3-manifolds whose complexities are unknown. Here we study the problem of finding a taut angle structure on a 3-manifold triangulation, whose existence has implications for both the geometry and combinatorics of the triangulation. We prove that detecting taut angle structures is NP-complete, but also fixed-parameter tractable

in the treewidth of the face pairing graph of the triangulation. These results have deeper implications: the core techniques can serve as a launching point for approaching decision problems such as unknot recognition and prime decomposition of 3-manifolds.

Benjamin A. Burton, Jonathan Spreer
The University of Queensland
bab@maths.uq.edu.au, j.spreer@uq.edu.au

CP3

Approximate Shortest Descending Paths

We present an approximate algorithm for the shortest descending path (SDP) problem. Given a source s and a destination t in a polygonal terrain \mathcal{T} , an SDP from s to t is a path in \mathcal{T} of minimum Euclidean length subject to the constraint that the height decreases monotonically as we traverse that path from s to t . Given any $\varepsilon \in (0, 1)$, our algorithm returns in $O(n^4 \log(n/\varepsilon))$ time a descending path of length at most $1 + \varepsilon$ times the optimum. This is the first algorithm whose running time is polynomial in n and $\log(1/\varepsilon)$ and independent of the terrain geometry.

Siu-Wing Cheng
The Hong Kong University of Science and Technology
scheng@cse.ust.hk

Jiongxin Jin
The Hong Kong University of Science and Technology
jamesjx@cse.ust.hk

CP3

An Infinite Class of Sparse-Yao Spanners

Given a finite set \mathcal{P} of points in the plane and a fixed integer $k > 1$, the Yao graph Y_k for \mathcal{P} is obtained by extending k equally spaced rays from each point $u \in \mathcal{P}$, and connecting u to a closest neighbor in each sector formed by these rays. It has been shown that, Y_k is a t -spanner, for any $k \geq 6$ and some real constant $t > 1$ depending on k , meaning that it has a path between each pair of points in \mathcal{P} no longer than t times the Euclidean distance between the two points. To reduce the potential high in-degree of Y_k , the Sparse-Yao graph YY_k eliminates from Y_k all but a shortest incoming edge in each sector. In this paper we establish that the Sparse-Yao graph defined by integer parameter $6k$, with $k \geq 6$, is an 11.67-spanner.

Mirela Damian, Matthew Bauer
Villanova University
mirela.damian@villanova.edu, mbauer03@villanova.edu

CP3

Weighted Graph Laplace Operator under Topological Noise

Recently, various applications have motivated the study of spectral structures (eigenvalues and eigenfunctions) of the so-called Laplace-Beltrami operator of a manifold and their discrete versions. A popular choice for the discrete version is the so-called Gaussian weighted graph Laplacian which can be applied to point cloud data that samples a manifold. Naturally, the question of stability of the spectrum of this discrete Laplacian under the perturbation of the sampled manifold becomes important for its practical usage. Previous results showed that the spectra of both the manifold Laplacian and discrete Laplacian are stable when

the perturbation is “nice” in the sense that it is restricted to a diffeomorphism with minor area distortion. However, this forbids, for example, small topological changes. We study the stability of the spectrum of the weighted graph Laplacian under more general perturbations. In particular, we allow arbitrary, including topological, changes to the hidden manifold as long as they are localized in the ambient space and the area distortion is small. Manifold Laplacians may change dramatically in this case. Nevertheless, we show that the weighted graph Laplacians computed from two sets of points, uniformly randomly sampled from a manifold and a perturbed version of it, have similar spectra. The distance between the two spectra can be bounded in terms of the size of the perturbation and some intrinsic properties of the original manifold.

Tamal Dey, Pawas Ranjan
The Ohio State University
tamaldehy@cse.ohio-state.edu, ranjan@cse.ohio-state.edu

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

CP4

Compressed Static Functions with Applications

Given a set of integer keys from a bounded universe along with associated data, the dictionary problem asks to answer two queries: membership and retrieval. *Membership* has to tell whether a given element is in the dictionary or not; *Retrieval* has to return the data associated with the searched key. In this paper we provide time and space optimal solutions for three well-established relaxations of this basic problem: *(Compressed) Static functions*, *Approximate membership* and *Relative membership*.

Djamal Belazzougui
LIAFA, Univ. Paris Diderot - Paris 7
dbelaz@liafa.jussieu.fr

Rossano Venturini
University of Pisa
rossano@di.unipi.it

CP4

Near-Optimal Range Reporting Structures for Categorical Data

Range reporting on categorical (or colored) data is a well-studied generalization of the classical range reporting problem in which each of the N input points has an associated color (category). A query then asks to report the set of colors of the points in a given rectangular query range, which may be far smaller than the set of all points in the query range. We study two-dimensional categorical range reporting in both the word-RAM and I/O-model. For the I/O-model, we present two alternative data structures for three-sided queries. The first answers queries in optimal $O(\lg_B N + K/B)$ I/Os using $O(N \lg^* N)$ space, where K is the number of distinct colors in the output, B is the disk block size, and $\lg^* N$ is the iterated logarithm of N . Our second data structure uses linear space and answers queries in $O(\lg_B N + \lg^{(h)} N + K/B)$ I/Os for any constant integer $h \geq 1$. Here $\lg^{(1)} N = \lg N$ and $\lg^{(h)} N = \lg(\lg^{(h-1)} N)$ when $h > 1$. Both solutions use only comparisons on the coordinates. We also show that the $\lg_B N$ terms in the query costs can be reduced to optimal $\lg_B U$ when the

input points lie on a $U \times U$ grid and we allow word-level manipulations of the coordinates. We further reduce the query time to just $O(1)$ if the points are given on an $N \times N$ grid. Both solutions also lead to improved data structures for four-sided queries. For the word-RAM, we obtain optimal data structures for three-sided range reporting, as well as improved upper bounds for four-sided range reporting. Finally, we show a tight lower bound on one-dimensional categorical range *counting* using an elegant reduction from (standard) two-dimensional range counting.

Kasper G. Larsen
MADALGO, Aarhus University
larsen@cs.au.dk

Freek Van Walderveen
MADALGO
freek@vanwal.dk

CP4

Twisted Tabulation Hashing

We introduce a new tabulation-based hashing scheme called ‘twisted tabulation’. It is essentially as simple and fast as simple tabulation, but has some powerful distributional properties illustrating its promise:

1. If we sample keys with arbitrary probabilities, then with high probability, the number of samples inside any subset is concentrated exponentially. With bounded independence we only get polynomial concentration, and with simple tabulation, we have no good bound even in the basic case of tossing an (unbiased) coin for each key.
2. With classic hash tables such as linear probing and collision-chaining, a window of B operations takes $O(B)$ time with high probability, for $B = \Omega(\lg n)$. Good amortized performance over any window of size B is equivalent to guaranteed throughput for an on-line system processing a stream via a buffer of size B (e.g., Internet routers).

Mihai Patrascu
AT&T Labs - Research
Died June 5, 2012
mip@heaven

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

CP4

The Space Complexity of 2-Dimensional Approximate Range Counting

We study the problem of 2-dimensional orthogonal range counting with additive error. Given a set P of n points drawn from an $n \times n$ grid and an error parameter ϵ , the goal is to build a data structure, such that for any orthogonal range R , the data structure can return the number of points in $P \cap R$ with additive error ϵn . A well-known solution for this problem is the ϵ -approximation. Informally speaking, an ϵ -approximation of P is a subset $A \subseteq P$ that allows us to estimate the number of points in $P \cap R$ by counting the number of points in $A \cap R$. It is known that an ϵ -approximation of size $O(\frac{1}{\epsilon} \log^{2.5} \frac{1}{\epsilon})$ exists for any P with respect to orthogonal ranges, and the best lower bound is $\Omega(\frac{1}{\epsilon} \log \frac{1}{\epsilon})$. The ϵ -approximation is a rather restricted data structure, as we are not allowed to store any information

other than the coordinates of a subset of points in P . In this paper, we explore what can be achieved without any restriction on the data structure. We first describe a data structure that uses $O(\frac{1}{\epsilon} \log \frac{1}{\epsilon} \log \log \frac{1}{\epsilon} \log n)$ bits that answers queries with error ϵn . We then prove a lower bound that any data structure that answers queries with error $O(\log n)$ must use $\Omega(n \log n)$ bits. This lower bound has two consequences: 1) answering queries with error $O(\log n)$ is as hard as answering the queries exactly; and 2) our upper bound cannot be improved in general by more than an $O(\log \log \frac{1}{\epsilon})$ factor.

Zhewei Wei
Center for Massive Data Algorithmics (MADALGO)
Aarhus University
zhewei@cs.au.dk

Ke Yi
The Hong Kong University of Science and Technology
yike@cse.ust.hk

CP4

Adaptive and Approximate Orthogonal Range Counting

We consider *2-D orthogonal range counting*, a basic problem in geometric data structures, under the w -bit word RAM model. There are linear-space data structures for this problem with optimal worst-case query time $O(\log_w n)$. We give an $O(n \log \log n)$ -space *adaptive* data structure with query time $O(\log \log n + \log_w k)$, where k is the output count. We also give an efficient data structure for *approximate* counting and an optimal adaptive data structure for the closely related 1-D *range selection* problem.

Bryan T. Wilkinson
MADALGO, Aarhus University
btw@cs.au.dk

Timothy Chan
David R. Cheriton School of Computer Science
University of Waterloo
tmchan@uwaterloo.ca

CP5

Better Balance by Being Biased: A 0.8776-Approximation for Max Bisection

We present an approximation algorithm for Max-Bisection that achieves ratio 0.8766. As Max-Bisection is hard to approximate within $\alpha_{GW} + \epsilon \approx 0.8786$ under the Unique Games Conjecture (UGC), our algorithm is nearly optimal. We also obtain an optimal algorithm (assuming the UGC) for the analogous variant of Max-2-Sat. Our approximation ratio for this problem exactly matches the optimal approximation for Max-2-Sat, i.e. $\alpha_{LLZ} + \epsilon \approx 0.9401$.

Konstantinos Georgiou
University of Waterloo
k2georgiou@math.uwaterloo.ca

Per Austrin
KTH Royal Institute of Technology
austrin@cs.toronto.edu

Siavosh Benabbas
University of Toronto
siavosh@cs.toronto.edu

CP5**Local Distribution and the Symmetry Gap: Approximability of Multiway Partitioning Problems**

We study the approximability of multiway partitioning problems, examples of which include Multiway Cut, Node-weighted Multiway Cut, and Hypergraph Multiway Cut. We investigate these problems from the point of view of two possible generalizations: as Min-CSPs, and as Submodular Multiway Partition problems. These two generalizations lead to two natural relaxations that we call respectively the *Local Distribution LP*, and the *Lovász relaxation*. The Local Distribution LP is generally stronger than the Lovász relaxation, but applicable only to Min-CSP with predicates of constant size. The relaxations coincide in some cases such as Multiway Cut where they are both equivalent to the CKR relaxation. We show that the Lovász relaxation gives a $(2 - 2/k)$ -approximation for Submodular Multiway Partition with k terminals, improving a recent 2-approximation. We prove that this factor is optimal in two senses: (1) A $(2 - 2/k - \epsilon)$ -approximation for Submodular Multiway Partition with k terminals would require exponentially many value queries (in the oracle model), or imply $NP = RP$ (for certain explicit submodular functions). (2) For Hypergraph Multiway Cut and Node-weighted Multiway Cut with k terminals, both special cases of Submodular Multiway Partition, we prove that a $(2 - 2/k - \epsilon)$ -approximation is NP-hard, assuming the Unique Games Conjecture. Both our hardness results are more general: (1) We show that the notion of *symmetry gap*, previously used for submodular maximization problems, also implies hardness results for submodular minimization problems. (2) Assuming the Unique Games Conjecture, we show that the Local Distribution LP gives an optimal approximation for every Min-CSP that includes the Not-Equal predicate.

Alina EneUniversity of Illinois at Urbana-Champaign
ene1@illinois.edu

Jan Vondrak

IBM Almaden Research Center
jvondrak@us.ibm.com

Yi Wu

Purdue University
wuyi@purdue.edu**CP5****Poly-Logarithmic Approximation for Maximum Node Disjoint Paths with Constant Congestion**

We consider the Maximum Node Disjoint Paths (MNDP) problem in undirected graphs. The input consists of an undirected graph $G = (V, E)$ and a collection $\{(s_1, t_1), \dots, (s_k, t_k)\}$ of k source-sink pairs. The goal is to select a maximum cardinality subset of pairs that can be routed/connected via node-disjoint paths. A relaxed version of MNDP allows up to c paths to use a node, where c is the congestion parameter. We give a polynomial time algorithm that routes $\Omega(\text{OPT}/\text{polylog}(k))$ pairs with $O(1)$ congestion, where OPT is the value of an optimum fractional solution to a natural multicommodity flow relaxation. Our result builds on the recent breakthrough of Chuzhoy who gave the first poly-logarithmic approximation with constant congestion for the Maximum Edge Disjoint Paths (MEDP) problem.

Alina Ene, Chandra Chekuri

University of Illinois at Urbana-Champaign

ene1@illinois.edu, chekuri@illinois.edu

CP5**How to Sell Hyperedges: The Hypermatching Assignment Problem**

In the Hypermatching Assignment Problem (HAP), we are given a set of clients with budget constraints and a set of indivisible items. Each client is willing to buy one or more bundles of (at most) k items each (bundles can be seen as hyperedges in a k -hypergraph). If client i gets a bundle e , she pays $b_{i,e}$ and yields a net profit $w_{i,e}$. The goal is to assign a set of pairwise disjoint bundles to clients so as to maximize the total profit while respecting the budgets. HAP generalizes both the weighted k -hypergraph matching problem and the Generalized Assignment Problem (GAP). We give a randomized $(k + 1 + \epsilon)$ approximation algorithm for HAP, which is based on rounding the 1-round Lasserre strengthening of a novel LP. Furthermore we obtain better approximation ratios for special cases of the problem.

Fabrizio Grandoni, Marek Cygan, Monaldo Mastrolilli
IDSIA, University of Lugano
fabrizio@idsia.ch, marek@idsia.ch, monaldo@idsia.ch**CP5****Approximating Non-Uniform Sparsest Cut Via Generalized Spectra**

We give an approximation algorithm for non-uniform sparsest cut with the following guarantee: For any $\epsilon, \delta \in (0, 1)$, given cost and demand graphs, we can find a set $T \subseteq V$ whose non-uniform sparsity is at most $\frac{1+\epsilon}{\delta}$ times the optimal non-uniform sparsest cut value, in time $2^{r/(\delta\epsilon)} n^{O(1)}$ provided the r^{th} minimum generalized eigenvalue satisfies $\lambda_r \geq \Phi^*/(1 - \delta)$. In words, we show that the non-uniform sparsest cut problem is easy when the generalized spectrum grows moderately fast.

Venkatesan Guruswami, Ali K. Sinop
Carnegie Mellon University
guruswami@cmu.edu, asinop@cs.cmu.edu**CP6****Shortest Non-Trivial Cycles in Directed and Undirected Surface Graphs**

Let G be a graph embedded on a surface of genus g with b boundary cycles. If G is undirected, then we give algorithms to compute a shortest non-separating, non-contractible, or non-null-homologous cycle in $2^{O(g)} n \log \log n$ time. In the more interesting case where G is directed, we give an algorithm to compute a shortest non-contractible cycle in $O(g^3 n \log n)$ time or a shortest non-null-homologous cycle in $O(g^2 n \log n)$ time.

Kyle FoxUniversity of Illinois, Urbana-Champaign
kylefox2@illinois.edu**CP6****A Simple Algorithm for the Graph Minor Decomposition - Logic Meets Structural Graph Theory -**

A key result of Robertson and Seymour's graph minor theory is a structure theorem stating that all graphs excluding some fixed graph as a minor have a tree decomposition into

pieces that are almost embeddable in a fixed surface. Most algorithmic applications of graph minor theory rely on an algorithmic version of this result. However, the known algorithms for computing such graph minor decompositions heavily rely on the very long and complicated proofs of the existence of such decompositions, essentially they retrace these proofs and show that all steps are algorithmic. In this paper, we give a simple quadratic time algorithm for computing graph minor decompositions. The best previously known algorithm due to Kawarabayashi and Wollan runs in cubic time and is far more complicated. Our algorithm combines techniques from logic and structural graph theory, or more precisely, a variant of Courcelle's Theorem stating that monadic second-order logic formulas can be evaluated in linear time on graphs of bounded tree width and Robertson and Seymour's so called Weak Structure Theorem.

Martin Grohe
Humboldt-Universität zu Berlin
grohe@informatik.rwth-aachen.de

Ken-ichi Kawarabayashi
National Institute of Informatics, Japan
k_keniti@nii.ac.jp

Bruce Reed
McGill University
breed@cs.mcgill.ca

CP6

Packing Directed Cycles Through a Specified Vertex Set

We show that the so-called Erdős-Pósa property does hold for half-integral packings of directed cycles each containing a vertex from a given set S , i.e. where every vertex of the graph is contained in at most 2 cycles. On the other hand, an example shows that the Erdős-Pósa property does not hold without this relaxation.

Ken-ichi Kawarabayashi
National Institute of Informatics, Japan
k_keniti@nii.ac.jp

Daniel Kral, Marek Krcal
Charles University
D.Kral@warwick.ac.uk, krcal@kam.mff.cuni.cz

Stephan Kreutzer
Technical University Berlin
stephan.kreutzer@tu-berlin.de

CP6

4-Connected Projective-Planar Graphs Are Hamiltonian-Connected

We prove that every 4-connected projective planar graph is hamiltonian-connected. Here, *hamiltonian-connected* means that for any two vertices, there is a hamiltonian path between them (and hence this generalizes "hamiltonian"). This proves a conjecture of Dean in 1990. Our proof is constructive in the sense that there is a polynomial time (in fact, $O(n^2)$ time) algorithm to find, given two vertices in a 4-connected projective planar graph, a hamiltonian path between these two vertices.

Kenta Ozeki, Ken-ichi Kawarabayashi
National Institute of Informatics, Japan

ozeki@nii.ac.jp, k_keniti@nii.ac.jp

CP6

Jungles, Bundles, and Fixed-Parameter Tractability

We give a fixed-parameter tractable (FPT) approximation algorithm computing the pathwidth of a tournament, and more generally, of a semi-complete digraph. Based on this result, we prove that the TOPOLOGICAL CONTAINMENT and ROOTED IMMERSION problems are FPT on semi-complete digraphs. Moreover, our results also imply that computing vertex deletion distance to every immersion-closed class of semi-complete digraphs is fixed-parameter tractable.

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

Michal Pilipczuk
University of Bergen
michal.pilipczuk@ii.uib.no

CP7

Fast Algorithms for Interactive Coding

Consider two parties who wish to communicate in order to execute some interactive protocol π . However, the communication channel between them is noisy: An adversary sees everything that is transmitted over the channel and can change a constant fraction of the bits as he pleases, thus interrupting the execution of π (which was designed for an errorless channel). If π only contained one message, then a good error correcting code would have overcome the noise with only a constant overhead in communication, but this solution is not applicable to interactive protocols with many short messages. Schulman (FOCS 92, STOC 93) presented the notion of *interactive coding*: A simulator that, given any protocol π , is able to simulate it (i.e. produce its intended transcript) even with constant rate adversarial channel errors, and with only constant (multiplicative) communication overhead. Until recently, however, the running time of all known simulators was exponential (or sub-exponential) in the communication complexity of π (denoted N in this work). Brakerski and Kalai (FOCS 12) recently presented a simulator that runs in time $\text{poly}(N)$. Their simulator is randomized (each party flips private coins) and has failure probability roughly 2^{-N} . In this work, we improve the computational complexity of interactive coding. While at least N computational steps are required (even just to output the transcript of π), the BK simulator runs in time $\tilde{O}(N^2)$. We present two efficient algorithms for interactive coding: The first with computational complexity $O(N \log N)$ and exponentially small failure probability; and the second with computational complexity $O(N)$, but failure probability $1/\text{poly}(N)$. (Computational complexity is measured in the RAM model.)

Zvika Brakerski
Stanford University
zvika@stanford.edu

Moni Naor
Weizmann Institute of Science
moni.naor@weizmann.ac.il

CP7**Restricted Isometry of Fourier Matrices and List Decodability of Random Linear Codes**

We prove that a random linear code over $\text{GF}(q)$, with probability arbitrarily close to 1, is list decodable at radius $1 - 1/q - \epsilon$ with list size $L = O(1/\epsilon^2)$ and rate $R = \Omega_q(\epsilon^2/(\log^3(1/\epsilon)))$. Up to the polylogarithmic factor in $1/\epsilon$ and constant factors depending on q , this matches the lower bound $L = \Omega_q(1/\epsilon^2)$ for the list size and upper bound $R = O_q(\epsilon^2)$ for the rate. Previously only existence (and not abundance) of such codes was known for the special case $q = 2$ (Guruswami, Håstad, Sudan and Zuckerman, 2002). In order to obtain our result, we employ a relaxed version of the well known Johnson bound on list decoding that translates the *average* Hamming distance between codewords to list decoding guarantees. We furthermore prove that the desired average-distance guarantees hold for a code provided that a natural complex matrix encoding the codewords satisfies the Restricted Isometry Property with respect to the Euclidean norm (RIP-2). For the case of random binary linear codes, this matrix coincides with a random submatrix of the Hadamard-Walsh transform matrix that is well studied in the compressed sensing literature. Finally, we improve the analysis of Rudelson and Vershynin (2008) on the number of random frequency samples required for exact reconstruction of k -sparse signals of length N . Specifically, we improve the number of samples from $O(k \log(N) \log^2(k) (\log k + \log \log N))$ to $O(k \log(N) \log^3(k))$. The proof involves bounding the expected supremum of a related Gaussian process by using an improved analysis of the metric defined by the process. This improvement is crucial for our application in list decoding.

Mahdi Cheraghchi
Carnegie Mellon University
Computer Science Department
cheraghchi@cmu.edu

Venkatesan Guruswami
Carnegie Mellon University
guruswami@cmu.edu

Ameya Velingker
Carnegie Mellon University
Computer Science Department
avelingk@cs.cmu.edu

CP7**Shift Finding in Sub-Linear Time**

We study the following basic pattern matching problem. Consider a “code” sequence \mathbf{c} consisting of n bits chosen uniformly at random, and a “signal” sequence \mathbf{x} obtained by shifting \mathbf{c} (modulo n) and adding noise. The goal is to efficiently recover the shift with high probability. The problem models tasks of interest in several applications, including GPS synchronization and motion estimation. We present an algorithm that solves the problem in time $\tilde{O}(n^{f/(1+f)})$, where $\tilde{O}(N^f)$ is the running time of the best algorithm for finding the closest pair among N “random” sequences of length $O(\log N)$. A trivial bound of $f = 2$ leads to a simple algorithm with a running time of $\tilde{O}(n^{2/3})$. The asymptotic running time can be further improved by plugging in recent more efficient algorithms for the closest pair problem. Our results also yield a sub-linear time algorithm for approximate pattern matching algorithm for a random signal (text), even for the case when the error between the

signal and the code (pattern) is asymptotically as large as the code size. This is the first sublinear time algorithm for such error rates.

Haitham Hassanieh
MIT
haithamh@mit.edu

Alexandr Andoni
Microsoft Research
andoni@microsoft.com

Piotr Indyk
Massachusetts Institute of Technology
indyk@mit.edu

Dina Katabi
MIT
dk@mit.edu

CP7**The Fast Cauchy Transform and Faster Robust Linear Regression**

We provide fast algorithms for overconstrained L_p regression and related problems that improve upon the best previous algorithms. We also provide a suite of results for finding well-conditioned bases via ellipsoidal rounding, illustrating various tradeoffs. To complement this theory, we provide an empirical evaluation. Among other things, this show the theory is a good guide to practical performance. Our algorithms use a fast subspace embedding of independent interest that we call the Fast Cauchy Transform.

Michael Mahoney
Stanford University
Applied and Computational Mathematics
mmahoney@cs.stanford.edu

Kenneth Clarkson
IBM
klclarks@us.ibm.com

Petros Drineas
Computer Science Department
Rensselaer Polytechnic Institute
drinep@cs.rpi.edu

Malik Magdon-Ismail
Rensselaer Polytechnic Institute
magdon@cs.rpi.edu

Xiangrui Meng
Stanford
mengxr@stanford.edu

David Woodruff
IBM Almaden Research Center, USA
dpwoodru@us.ibm.com

CP7**Generalized Perron-Frobenius Theorem for Multiple Choice Matrices, and Applications**

The celebrated Perron-Frobenius (PF) theorem is stated for irreducible nonnegative square matrices, and provides a simple characterization of their eigenvectors and eigenvalues. The importance of this theorem stems from the

fact that eigenvalue problems on such matrices arise in many fields of science and engineering, including dynamical systems theory, economics, statistics and optimization. However, many real-life scenarios give rise to nonsquare matrices. Despite the extensive development of spectral theories for nonnegative matrices, the applicability of such theories to non-convex optimization problems is not clear. In particular, a natural question is whether the PF Theorem (along with its applications) can be generalized to a nonsquare setting. Our paper provides a generalization of the PF Theorem to nonsquare multiple choice matrices. The extension can be interpreted as representing systems with additional degrees of freedom, where each client entity may choose between multiple servers that can cooperate in serving it (while potentially interfering with other clients). This formulation is motivated by applications to power control in wireless networks, economics and others, all of which extend known examples for the use of the original PF Theorem. We show that the option of cooperation does not improve the situation, in the sense that in the optimum solution, no cooperation is needed, and only one server per client entity needs to work. Hence, the additional power of having several potential servers per client translates into choosing the “best” single server and not into sharing the load between the servers in some way, as one might have expected. The two main contributions of the paper are (i) a generalized PF Theorem that characterizes the optimal solution for a non-convex problem, and (ii) an algorithm for finding the optimal solution in polynomial time.

Chen Avin
Ben Gurion University, Beer-Sheva, Israel.
avin@cse.bgu.ac.il

Michael Borokhovich
Ben Gurion University, Beer-Sheva, Israel
borokhom@cse.bgu.ac.il

Yoram Haddad
Jerusalem College of Technology, Jerusalem, Israel
yoram.haddad@gmail.com

Erez Kantor
The Technion, Haifa, Israel
erez.kantor@gmail.com

Zvi Lotker
Ben Gurion University, Beer-Sheva, Israel.
zvilo@cse.bgu.ac.il

Merav Parter
The Weizmann Institute of Science, Rehovot, Israel.
merav.parter@weizmann.ac.il

David Peleg
Weizmann Institute of Science
david.peleg@weizmann.ac.il

CP8 Distance Oracles for Stretch Less Than 2

We present distance oracles for weighted undirected graphs that return distances of stretch less than 2. For the realistic case of sparse graphs, our distance oracles exhibit a smooth three-way trade-off between space, query time and stretch — a phenomenon that does not occur in dense graphs. In particular, for any positive integer t and for any $1 \leq \alpha \leq n$, our distance oracle is of size $O(m + n^2/\alpha)$ and returns

distances of stretch at most $(1 + \frac{2}{t+1})$ in time $O((\alpha\mu)^t)$, where $\mu = 2m/n$ is the average degree of the graph. The query time can be further reduced to $O((\alpha + \mu)^t)$ at the expense of a small additive stretch.

Rachit Agarwal, Philip Godfrey
University of Illinois at Urbana-Champaign
rachit.ee@gmail.com, pbg@illinois.edu

CP8 New Additive Spanners

This paper considers additive and purely additive spanners. We present a new purely additive spanner of size $\tilde{O}(n^{7/5})$ with additive stretch 4. This construction fills in the gap between the two existing constructions for purely additive spanners, one for 2-additive spanner of size $O(n^{3/2})$ and the other for 6-additive spanner of size $O(n^{4/3})$, and thus answers a main open question in this area. In addition, we present a construction for additive spanners with $\tilde{O}(n^{1+\delta})$ edges and additive stretch of $\tilde{O}(n^{1/2-3\delta/2})$ for any $3/17 \leq \delta < 1/3$, improving the stretch of the existing constructions from $O(n^{1-3\delta})$ to $\tilde{O}(\sqrt{n^{1-3\delta}})$. Finally, we show that our $(1, n^{1/2-3\delta/2})$ -spanner construction can be tweaked to give a sublinear additive spanner of size $\tilde{O}(n^{1+3/17})$ with additive stretch $O(\sqrt{\text{distance}})$.

Shiri Chechik
Microsoft Research SVC
shiri.chechik@gmail.com

CP8 Fast Constructions of Light-Weight Spanners for General Graphs

Since the pioneering works of Peleg and Schäffer [PS89], Althöfer et al. [ADDJS93], and Chandra et al. [CDNS92], it is known that for every *weighted* undirected n -vertex m -edge graph $G = (V, E)$, and every integer $k \geq 1$, there exists a $((2k - 1) \cdot (1 + \epsilon))$ -spanner with $O(n^{1+1/k})$ edges and weight $O(k \cdot n^{1/k}) \cdot \omega(MST(G))$. Nearly linear time algorithms for constructing $(2k - 1)$ -spanners with nearly $O(n^{1+1/k})$ edges were devised in [Baswana-Sen 2003, Roditty-Zwick 2004, Roditty-Thorup-Zwick 2005]. However, these algorithms fail to guarantee any meaningful upper bound on the weight of the constructed spanners. To our knowledge, there are only two known algorithms for constructing sparse and light spanners for general graphs. One of them is the greedy algorithm of [ADDJS93], analyzed in [CDNS92]. The drawback of the greedy algorithm is that it requires $O(m \cdot (n^{1+1/k} + n \cdot \log n))$ time. The other algorithm is due to Awerbuch et al. [ABP91], from 1991. It constructs $O(k)$ -spanners with $O(k \cdot n^{1+1/k} \cdot \Lambda)$ edges, weight $O(k^2 \cdot n^{1/k} \cdot \Lambda) \cdot \omega(MST(G))$, within time $O(m \cdot k \cdot n^{1/k} \cdot \Lambda)$, where Λ is the logarithm of the aspect ratio of the graph. In this paper we devise an efficient algorithm for constructing sparse and light spanners. Specifically, our algorithm constructs $((2k - 1) \cdot (1 + \epsilon))$ -spanners with $O(k \cdot n^{1+1/k})$ edges and weight $O(k \cdot n^{1/k}) \cdot \omega(MST(G))$, where $\epsilon > 0$ is an arbitrarily small constant. The running time of our algorithm is $O(k \cdot m + \min\{n \cdot \log n, m \cdot \alpha(n)\})$.

Michael Elkin
Ben-Gurion University
elkinm@cs.bgu.ac.il

Shay Solomon
Weizmann Institute of Science
solo.shay@gmail.com

CP8

More Compact Oracles for Approximate Distances in Undirected Planar Graphs

Distance oracles are data structures that provide fast (possibly approximate) answers to shortest-path and distance queries in graphs. The tradeoff between the space requirements and the query time of distance oracles is of particular interest and the main focus of this paper. Unless stated otherwise, we assume all graphs to be planar and undirected. In FOCS 2001 (J. ACM 2004), Thorup introduced approximate distance oracles for planar graphs (concurrent with Klein, SODA 2002). Thorup proved that, for any $\epsilon > 0$ and for any undirected planar graph on n nodes, there exists a $(1 + \epsilon)$ -approximate distance oracle using space $O((n \log n)/\epsilon)$ such that approximate distance queries can be answered in time $O(1/\epsilon)$. In this paper, we aim at reducing the polynomial dependency on $1/\epsilon$ and $\log n$, getting the first improvement in the query time-space tradeoff. To simplify the statement of our bounds, we define $\hat{O}()$ to hide $\log \log n$ and $\log(1/\epsilon)$ factors. We provide the first oracle with a time-space product that is subquadratic in $1/\epsilon$. We obtain an oracle with space $\hat{O}(n \log n)$ and query time $\hat{O}(1/\epsilon)$. For unweighted graphs we show how the logarithmic dependency on n can be removed. We obtain an oracle with space $\hat{O}(n)$ and query time $\hat{O}(1/\epsilon)$. This bound also holds for graphs with polylogarithmic average edge length, which may be a quite reasonable assumption, e.g., for road networks.

Christian Sommer
MIT
csom@mit.edu

Ken-ichi Kawarabayashi
National Institute of Informatics, Japan
k_keniti@nii.ac.jp

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

CP8

Approximate Distance Oracles with Improved Query Time

For a graph with m edges and n vertices and for an integer $k \geq 2$, we give a $(2k - 1)$ -approximate distance oracle of size $O(kn^{1+1/k})$ and with $O(\log k)$ query time, improving the $O(k)$ query time of Thorup and Zwick. Furthermore, for any $0 < \epsilon \leq 1$, we give an oracle of size $O(kn^{1+1/k})$ that answers $((2 + \epsilon)k)$ -approximate distance queries in $O(1/\epsilon)$ time. This improves a $128k$ approximation of Mendel and Naor and approaches the best possible tradeoff between size and stretch.

Christian Wulff-Nilsen
Department of Mathematics and Computer Science
University of Southern Denmark
koolooz@diku.dk

Christian Wulff-Nilsen
Department of Computer Science
University of Copenhagen

koolooz@diku.dk

CP9

Optimal and Efficient Parametric Auctions

Consider a seller who seeks to provide service to a collection of interested parties, subject to feasibility constraints on which parties may be simultaneously served. Assuming that a distribution is known on the value of each party for service—arguably a strong assumption—Myerson’s seminal work provides revenue optimizing auctions. We show instead that, for very general feasibility constraints, only knowledge of the median of each party’s value distribution, or any other quantile of these distributions, or approximations thereof, suffice for designing simple auctions that simultaneously approximate both the optimal revenue and the optimal welfare. Our results apply to all downward-closed feasibility constraints under the assumption that the underlying, unknown value distributions are monotone hazard rate, and to all matroid feasibility constraints under the weaker assumption of regularity of the underlying distributions. Our results jointly generalize the single-item results obtained by Azar and Micali on parametric auctions, and Daskalakis and Pierrakos for simultaneously approximately optimal and efficient auctions.

Pablo Azar, Constantinos Daskalakis
MIT
pazar@mit.edu, costis@csail.mit.edu

Silvio Micali
Massachusetts Institute of Technology
silvio@csail.mit.edu

Matt Weinberg
MIT
smweinberg@csail.mit.edu

CP9

Simple and Nearly Optimal Multi-Item Auctions

We provide a Polynomial Time Approximation Scheme (PTAS) for the Bayesian optimal multi-item multi-bidder auction problem under two conditions. First, bidders are independent, have additive valuations and are from the same population. Second, every bidder’s value distributions of items are independent but not necessarily identical monotone hazard rate (MHR) distributions. For non-i.i.d. bidders, we also provide a PTAS when the number of bidders is small. Prior to our work, even for a single bidder, only constant factor approximations are known. Another appealing feature of our mechanism is the simple allocation rule. Indeed, the mechanism we use is either the second-price auction with reserve price on every item individually, or VCG allocation with a few outlying items that requires additional treatments. It is surprising that such simple allocation rules suffice to obtain nearly optimal revenue.

Yang Cai
MIT
MIT
ycai@mit.edu

Zhiyi Huang
University of Pennsylvania
hzhiyi@cis.upenn.edu

CP9**Clinching Auctions with Online Supply**

Auctions for perishable goods such as internet ad inventory need to make real-time allocation and pricing decisions as the supply of the good arrives in an online manner, without knowing the entire supply in advance. These allocation and pricing decisions get complicated when buyers have some global constraints. In this work, we consider a multi-unit model where buyers have global *budget* constraints, and the supply arrives in an online manner. Our main contribution is to show that for this setting there is an individually-rational, incentive-compatible and Pareto-optimal auction that allocates these units and calculates prices on the fly, without knowledge of the total supply. We do so by showing that the Adaptive Clinching Auction satisfies a *supply-monotonicity* property. We also analyze and discuss, using examples, how the insights gained by the allocation and payment rule can be applied to design better ad allocation heuristics in practice. Finally, while our main technical result concerns multi-unit supply, we propose a formal model of online supply that captures scenarios beyond multi-unit supply and has applications to sponsored search. We conjecture that our results for multi-unit auctions can be extended to these more general models.

Renato Paes Leme

Cornell University
renatopl@cs.cornell.edu

Gagan Goel

Google Research, New York
gagangoel@google.com

Vahab Mirrokni

Google Inc.
mirrokni@google.com

CP9**Ironing in Dynamic Revenue Management: Posted Prices & Biased Auctions**

We consider the design of the revenue maximizing mechanism for a seller with a fixed capacity selling to buyers who arrive over time. The buyers have single unit demand and both their value for the object and the deadline by which they must make a purchase are unknown to the seller. In previous work where buyers' deadlines are publicly observed and only values are private, the optimal mechanism can be computed by running a dynamic stochastic knapsack algorithm. However, these mechanisms are only feasible with private deadlines when the calculated allocation rule is appropriately monotone. In the classic static environment of Myerson (1981) monotonicity is only violated for 'irregular' value distributions. Here, he characterizes the optimal mechanism by a procedure he calls 'ironing.' We characterize the optimal mechanism in our general dynamic environment by providing the dynamic counterpart of ironing. We show that only a subset of the monotonicity constraints can bind in a solution of the seller's dynamic programming problem. The optimal mechanism can be characterized by relaxing these constraints with their appropriate dual multiplier. Further, the optimal mechanism can be implemented by a series of posted prices followed by a 'biased' auction in the final period where buyers have the auction biased in their favor depending on their arrival time. Our theoretical characterization complements the existing computational approaches for ironing in these

settings (e.g. Parkes et al. (2007)).

Malleesh M. Pai

Department of Economics
University of Pennsylvania
malleesh.pai@gmail.com

Rahul Deb

Department of Economics
University of Toronto
rahul.deb@utoronto.ca

CP9**Reducing Revenue to Welfare Maximization: Approximation Algorithms and Other Generalizations**

It was recently shown in [CDW12b] that revenue optimization can be computationally efficiently reduced to welfare optimization in all multi-dimensional Bayesian auction problems with arbitrary (possibly combinatorial) feasibility constraints and independent additive bidders with arbitrary (possibly combinatorial) demand constraints. This reduction provides a poly-time solution to the optimal mechanism design problem in all auction settings where welfare optimization can be solved efficiently, but it is fragile to approximation and cannot provide solutions to settings where welfare maximization can only be tractably approximated. In this paper, we extend the reduction to accommodate approximation algorithms, providing an approximation preserving reduction from (truthful) revenue maximization to (not necessarily truthful) welfare maximization. The mechanisms output by our reduction choose allocations via black-box calls to welfare approximation on randomly selected inputs, thereby generalizing also our earlier structural results on optimal multi-dimensional mechanisms to approximately optimal mechanisms. Unlike [CDW12b], our results here are obtained through novel uses of the Ellipsoid algorithm and other optimization techniques over *non-convex regions*.

Yang Cai, Constantinos Daskalakis, Matt Weinberg
MIT

ycai@csail.mit.edu,
smweinberg@csail.mit.edu

costis@csail.mit.edu,

CP10**On the Number of Matroids**

We consider the problem of determining m_n , the number of matroids on n elements. Knuth (1974) showed that $\log \log m_n \geq n - \frac{3}{2} \log n - O(1)$, while Piff (1973) showed that $\log \log m_n \leq n - \log n + \log \log n + O(1)$, and it has been conjectured that the right answer is perhaps closer to Knuth's bound. We show that this is indeed the case, and prove an upper bound on $\log \log m_n$ that is within an additive $1 + o(1)$ term of Knuth's lower bound.

Nikhil Bansal

Eindhoven University of Technology
bansal@gmail.com

Rudi Pendavingh, Jorn van Der Pol
TU Eindhoven

rudi.pendavingh@gmail.com, j.g.v.d.pol@student.tue.nl

CP10**Playing Mastermind with Many Colors**

We analyze the general version of the classic guessing game Mastermind with n positions and k colors. Since the case $k \leq n^{1-\epsilon}$, $\epsilon > 0$ constant, is well understood, we concentrate on larger numbers of colors. For the most prominent case $k = n$, our results imply that Codebreaker can find the secret code with $O(n \log \log n)$ guesses. This bound is valid also when only black answer-pegs are used. It improves the $O(n \log n)$ bound first proven by Chvátal (Combinatorica 3 (1983), 325–329). We also show that if both black and white answer-pegs are used, then the $O(n \log \log n)$ bound holds for up to $n^2 \log \log n$ colors. These bounds are almost tight as the known lower bound of $\Omega(n)$ shows. Unlike for $k \leq n^{1-\epsilon}$, simply guessing at random until the secret code is determined is not sufficient. In fact, we show that any non-adaptive strategy needs an expected number of $\Omega(n \log n)$ guesses.

Benjamin Doerr, Reto Spöhel
Max Planck Institute for Informatics
doerr@mpi-inf.mpg.de, rspoehel@mpi-inf.mpg.de

Henning Thomas
Institute of Theoretical Computer Science ETH Zurich
8092 Zurich, Switzerland
hthomas@inf.ethz.ch

Carola Winzen
Max-Planck-Institut für Informatik
winzen@mpi-inf.mpg.de

CP10**Tight Cell-Probe Bounds for Online Hamming Distance Computation**

We show tight bounds for online Hamming distance computation in the cell-probe model with word size w . The task is to output the Hamming distance between a fixed string of length n and the last n symbols of a stream. We give a lower bound of $\Omega\left(\frac{\delta}{w} \log n\right)$ time on average per output, where δ is the number of bits needed to represent an input symbol. We argue that this bound is tight within the model. The lower bound holds under randomisation and amortisation.

Markus Jalsenius, Raphael Clifford
University of Bristol
m.jalsenius@bristol.ac.uk, clifford@compsci.bristol.ac.uk

Benjamin Sach
University of Warwick
sach@dcs.warwick.ac.uk

CP10**Lower Bounds for Adaptive Sparse Recovery**

We give lower bounds for the problem of stable sparse recovery from *adaptive* linear measurements. In this problem, one would like to estimate a vector $x \in \mathbf{R}^n$ from m linear measurements A_1x, \dots, A_mx . One may choose each vector A_i based on $A_1x, \dots, A_{i-1}x$, and must output \hat{x} satisfying

$$\|\hat{x} - x\|_p \leq (1 + \epsilon) \min_{k\text{-sparse } x'} \|x - x'\|_p$$

with probability at least $1 - \delta > 2/3$, for some $p \in \{1, 2\}$. For $p = 2$, it was recently shown that this is possible with

$m = O\left(\frac{1}{\epsilon} k \log \log(n/k)\right)$, while nonadaptively it requires $\Theta\left(\frac{1}{\epsilon} k \log(n/k)\right)$. It is also known that even adaptively, it takes $m = \Omega(k/\epsilon)$ for $p = 2$. For $p = 1$, there is a non-adaptive upper bound of $\tilde{O}\left(\frac{1}{\sqrt{\epsilon}} k \log n\right)$. We show:

- For $p = 2$, $m = \Omega(\log \log n)$. This is tight for $k = O(1)$ and constant ϵ , and shows that the $\log \log n$ dependence is correct.
- If the measurement vectors are chosen in R “rounds”, then $m = \Omega(R \log^{1/R} n)$. For constant ϵ , this matches the previously known upper bound up to an $O(1)$ factor in R .
- For $p = 1$, $m = \Omega(k/(\sqrt{\epsilon} \cdot \log k/\epsilon))$. This shows that adaptivity cannot improve more than logarithmic factors, providing the analogue of the $m = \Omega(k/\epsilon)$ bound for $p = 2$.

Eric Price
MIT
ecprice@mit.edu

David Woodruff
IBM Almaden Research Center, USA
dpwoodru@us.ibm.com

CP10**The Communication Complexity of Addition**

Abstract not available at time of publication.

Emanuele Viola
Northeastern University
viola@ccs.neu.edu

CP11**Finding Endogenously Formed Communities**

We define a natural notion of (potentially overlapping) communities, in which the members of a given community collectively prefer each other to anyone else outside the community. We provide a tight polynomial bound on the number of such communities as a function of the robustness of the community. Moreover, we present polynomial-time algorithms for finding these communities and a local algorithm that can find a community in time nearly linear in its size.

Maria-Florina Balcan
Georgia Institute of Technology
ninamf@cc.gatech.edu

Christian Borgs
Microsoft Research
borgs@microsoft.com

Mark Braverman
Princeton University
mbraverm@gmail.com

Jennifer Chayes
Microsoft Research
jchayes@microsoft.com

Shang-Hua Teng
University of Southern California
shanghua.teng@gmail.com

CP11**Anonymous Meeting in Networks**

A team of an unknown number of mobile agents, in an unknown network, have to meet at the same node. Agents are anonymous, execute the same deterministic algorithm and move synchronously. An initial configuration is gatherable if there exists a deterministic algorithm that achieves meeting of all agents in one node. Which configurations are gatherable and how to gather the agents deterministically by the same algorithm? We give a complete solution of this gathering problem in arbitrary networks.

Yoann Dieudonne

Departement Informatique
Université du Québec en Outaouais
yoann.dieudonne@u-picardie.fr

Andrzej Pelc

Département d'informatique
Université du Québec en Outaouais
andrzej.pelc@uqo.ca

CP11**On the Complexity of Information Spreading in Dynamic Networks**

We study how to spread k tokens of information to every node on an n -node dynamic network, the edges of which change at each round. This basic *gossip problem* can be completed in $O(n+k)$ rounds in any static network, and determining its complexity in dynamic networks is central to understanding the algorithmic limits and capabilities of various dynamic network models. We provide several almost tight upper and lower bounds in some of these models.

Chinmoy Dutta

Northeastern University
chinmoy@ccs.neu.edu

Gopal Pandurangan

Division of Mathematical Sciences
Nanyang Technological University
gopalpandurangan@gmail.com

Rajmohan Rajaraman, Zhifeng Sun

Northeastern University
USA
rraj@ccs.neu.edu, austin@ccs.neu.edu

Emanuele Viola

Northeastern University
viola@ccs.neu.edu

CP11**Near Optimal Leader Election in Multi-Hop Radio Networks**

We design leader election protocols for multi-hop radio networks that elect a leader in almost the same time T_{BC} that it takes for broadcasting one message (one ID). For the setting without collision detection our algorithm runs with high probability in $O(D \log \frac{n}{D} + \log^3 n) \cdot \min\{\log \log n, \log \frac{n}{D}\}$ rounds on any n -node network with diameter D . Since $T_{BC} = \Theta(D \log \frac{n}{D} + \log^2 n)$ is a lower bound, our upper bound is optimal up to a factor of at most $\log \log n$ and the extra $\log n$ factor on the additive

term. Our algorithm is furthermore the first $O(n)$ time algorithm for this setting. This algorithm improves over a 23 year old simulation approach of Bar-Yehuda, Goldreich and Itai with a $O(T_{BC} \log n)$ running time. We also give an $O(D + \log n \log \log n) \cdot \min\{\log \log n, \log \frac{n}{D}\} = O(D + \log n) \cdot O(\log \log n)^2$ leader election algorithm for the setting with collision detection (even with single-bit messages). This time bound — which is also in $O(n)$ — is optimal up to $\log \log n$ factors and improves over a deterministic algorithm that requires $\Theta(n)$ rounds independently of D .

Mohsen Ghaffari, Bernhard Haeupler

MIT
ghaffari@mit.edu, haeupler@mit.edu

CP11**Simple, Fast and Deterministic Gossip and Rumor Spreading**

Abstract not available at time of publication.

Bernhard Haeupler

MIT
haeupler@mit.edu

CP12**Reporting Neighbors in High-Dimensional Euclidean Space**

We consider the following problem, which arises in many database and web-based applications: Given a set P of n points in a high-dimensional space R^d and a distance r , we want to report all pairs of points of P at Euclidean distance at most r . We present two randomized algorithms, one based on randomly shifted grids, and the other on randomly shifted and rotated grids. The running time of both algorithms is of the form $C(d)(n+k) \log n$, where k is the output size and $C(d)$ is a constant that depends on the dimension d . The $\log n$ factor is needed to guarantee, with high probability, that all neighbor pairs are reported, and can be dropped if it suffices to report, in expectation, an arbitrarily large fraction of the pairs. When only translations are used, $C(d)$ is of the form $(a\sqrt{d})^d$, for some (small) absolute constant $a \approx 0.484$; this bound is worst-case tight, up to an exponential factor of about 2^d . When both rotations and translations are used, $C(d)$ can be improved to roughly 6.74^d , getting rid of the super-exponential factor \sqrt{d}^d . When the input set (lies in a subset of d -space that) has low *doubling dimension* δ , the performance of the first algorithm improves to $C(d, \delta)(n+k) \log n$ (or to $C(d, \delta)(n+k)$), where $C(d, \delta) = O((ed/\delta)^\delta)$, for $\delta \leq \sqrt{d}$. Otherwise, $C(d, \delta) = O\left(e^{\sqrt{d}} \sqrt{d}^\delta\right)$. We also present experimental results on several large datasets, demonstrating that our algorithms run significantly faster than all the leading existing algorithms for reporting neighbors.

Dror Aiger

Google Inc.
aigerd@google.com

Haim Kaplan, Micha Sharir

Tel Aviv University
haimk@tau.ac.il, michas@tau.ac.il

CP12**The Traveling Salesman Problem for Lines, Balls and Planes**

We revisit the traveling salesman problem with neighborhoods (TSPN) and obtain several approximation algorithms. These constitute either improvements over previously best approximations achievable in comparable times (for unit disks in the plane), or first approximations ever (for planes, lines and unit balls in 3-space). (I) Given a set of n planes in 3-space, a TSP tour that is at most 2.31 times longer than the optimal can be computed in $O(n)$ time. (II) Given a set of n lines in 3-space, a TSP tour that is at most $O(\log^3 n)$ times longer than the optimal can be computed in polynomial time. (III) Given a set of n unit disks in the plane (resp., unit balls in 3-space), we improve the approximation ratio using a black box that computes an approximate tour for a set of points (the centers of a subset of the disks or the balls).

Adrian Dumitrescu

Department of Computer Science
University of Wisconsin-Milwaukee
dumitres@uwm.edu

Csaba D. Toth
University of Calgary
cdtoth@ucalgary.ca

CP12**Clustering Affine Subspaces: Hardness and Algorithms**

We study a generalization of the famous k -center problem where each object is an affine subspace of dimension Δ , and give either the first or significantly improved algorithms and hardness results for many combinations of parameters. This generalization from points ($\Delta = 0$) is motivated by the analysis of incomplete data, a pervasive challenge in statistics: incomplete data objects in \mathbf{R}^d can be modeled as affine subspaces. We give three algorithmic results for different values of k , under the assumption that all subspaces are axis-parallel, the main case of interest because of the correspondence to missing entries in data tables. 1) $k = 1$: Two polynomial time approximation schemes which runs in $\text{poly}(\Delta, 1/\epsilon)nd$. 2) $k = 2$: $O(\Delta^{1/4})$ -approximation algorithm which runs in $\text{poly}(n, d, \Delta)$ 3) General k : Polynomial time approximation scheme which runs in $2^{O(\Delta k \log k(1+1/\epsilon^2))}nd$ We also prove nearly matching hardness results; in both the general (not necessarily axis-parallel) case (for $k \geq 2$) and in the axis-parallel case (for $k \geq 3$), the running time of an approximation algorithm with any approximation ratio cannot be polynomial in even one of k and Δ , unless $P = NP$. Furthermore, assuming that the 3-SAT problem cannot be solved subexponentially, the dependence on both k and Δ must be exponential in the general case (in the axis-parallel case, only the dependence on k drops to $2^{\Omega(\sqrt{k})}$). The simplicity of the first and the third algorithm suggests that they might be actually used in statistical applications. The second algorithm, which demonstrates a theoretical gap between the axis-parallel and general case for $k = 2$, displays a strong connection between geometric clustering and classical coloring problems on graphs and hypergraphs, via a new Helly-type theorem.

Euiwoong Lee
Carnegie Mellon University

euiwoonl@cs.cmu.edu

Leonard Schulman
Caltech
schulman@caltech.edu

CP12**Approximating Watchman Routes**

Given a connected polygonal domain P , the watchman route problem is to compute a shortest path or tour for a mobile guard (the “watchman”) that is required to see every point of P . While the watchman route problem is polynomially solvable in simple polygons, it is known to be NP-hard in polygons with holes. We present the first polynomial-time approximation algorithm for the watchman route problem in polygonal domains. Our algorithm has an approximation factor $O(\log^2 n)$. Further, we prove that the problem cannot be approximated in polynomial time to within a factor of $c \log n$, for a constant $c > 0$, assuming that $P \neq NP$.

Joseph Mitchell
Stony Brook University
jsbm@ams.sunysb.edu

CP12**Euclidean Spanners in High Dimensions**

A classical result in metric geometry asserts that any n -point metric admits a linear-size spanner of dilation $O(\log n)$ [Peleg, and Schäffer, 1989]. More generally, for any $c > 1$, any metric space admits a spanner of size $O(n^{1+1/c})$, and dilation at most c . This bound is tight assuming the well-known girth conjecture of Erdős [Erdős, 1963]. We show that for a metric induced by a set of n points in high-dimensional Euclidean space, it is possible to obtain improved dilation/size trade-offs. More specifically, we show that any n -point Euclidean metric admits a near-linear size spanner of dilation $O(\sqrt{\log n})$. Using the LSH scheme of Andoni and Indyk [Andoni, and Indyk, 2006] we further show that for any $c > 1$, there exist spanners of size roughly $O(n^{1+1/c^2})$ and dilation $O(c)$. Finally, we also exhibit super-linear lower bounds on the size of spanners with constant dilation.

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

Piotr Indyk
MIT
indyk@theory.csail.mit.edu

Anastasios Sidiropoulos
Toyota Technological Institute at Chicago, USA
sidiropo@gmail.com

CP13**Windows into Relational Events: Data Structures for Contiguous Subsequences of Edges**

We consider the problem of analyzing social network data sets in which the edges of the network have timestamps, and we wish to analyze the subgraphs formed from edges in contiguous subintervals of these timestamps. We pro-

vide data structures for these problems that use near-linear preprocessing time, linear space, and sublogarithmic query time to handle queries that ask for the number of connected components, number of components that contain cycles, number of vertices whose degree equals or is at most some predetermined value, number of vertices that can be reached from a starting set of vertices by time-increasing paths, and related queries.

Michael J. Bannister

University of California, Irvine
Department of Computer Science
mbannist@uci.edu

Christopher DuBois

University of California, Irvine
Department of Computer Statistics
duboisc@ics.uci.edu

David Eppstein

Department of Information and Computer Science
University of California, Irvine
eppstein@ics.uci.edu

Padhraic Smyth

UC Irvine
smyth@ics.uci.edu

CP13

Towards More Realistic Probabilistic Models for Data Structures: The External Path Length in Tries under the Markov Model

Tries are among the most versatile and widely used data structures on words. Its performance analysis, even for simplest memoryless sources has proved difficult and rigorous findings about inherently complex parameters under more realistic models of string generations were rarely analyzed. Here, we consider general Markov sources and prove a central limit theorem for the external path length of a trie.

Ralph Neininger

J.W. Goethe University, Frankfurt
neinigr@math.uni-frankfurt.de

Kevin Leckey

Goethe University Frankfurt am Main, Germany
leckey@math.uni-frankfurt.de

Wojciech Szpankowski

Purdue University
Department of Computer Science
spa@cs.purdue.edu

CP13

Optimal Dynamic Sequence Representations

We describe a data structure that supports access, rank and select queries, as well as symbol insertions and deletions, on a string $S[1, n]$ over alphabet $[1..σ]$ in time $O(\lg n / \lg \lg n)$, which is optimal. The time is worst-case for the queries and amortized for the updates. This complexity is better than the best previous ones by a $\Theta(1 + \lg \sigma / \lg \lg n)$ factor. Our structure uses $nH_0(S) + O(n + \sigma(\lg \sigma + \lg^{1+\varepsilon} n))$ bits, where $H_0(S)$ is the zero-order entropy of S and $0 < \varepsilon < 1$ is any constant. This space redundancy over $nH_0(S)$ is also better, almost al-

ways, than that of the best previous dynamic structures, $o(n \lg \sigma) + O(\sigma(\lg \sigma + \lg n))$. We can also handle general alphabets in optimal time, which has been an open problem in dynamic sequence representations.

Gonzalo Navarro, Yakov Nekrich

University of Chile
gnavarro@dcc.uchile.cl, yakov.nekrich@googlemail.com

CP13

Output-Sensitive Skyline Algorithms in External Memory

This paper presents new results in external memory for finding the skyline (a.k.a. maxima) of N points in d -dimensional space. The state of the art uses $O((N/B) \log_{M/B}^{d-2}(N/B))$ I/Os for fixed $d \geq 3$, and $O((N/B) \log_{M/B}(N/B))$ I/Os for $d = 2$, where M and B are the sizes (in words) of memory and a disk block, respectively. We give algorithms whose running time depends on the number K of points in the skyline. Specifically, we achieve $O((N/B) \log_{M/B}^{d-2}(K/B))$ expected cost for fixed $d \geq 3$, and $O((N/B) \log_{M/B}(K/B))$ worst-case cost for $d = 2$. As a side product, we solve two problems both of independent interest. The first one, the *M-skyline problem*, aims at reporting M arbitrary skyline points, or the entire skyline if its size is at most M . We settle this problem in $O(N/B)$ expected time in any fixed dimensionality d . The second one, the *M-pivot problem*, is more fundamental: given a set S of N elements drawn from an ordered domain, it outputs M evenly scattered elements (called pivots) from S , namely, S has asymptotically the same number of elements between each pair of consecutive pivots. We give a deterministic algorithm for solving the problem in $O(N/B)$ I/Os.

Yufei Tao, Xiaocheng Hu, Cheng Sheng

Chinese University of Hong Kong
taoyf@cse.cuhk.edu.hk, xchu@cse.cuhk.edu.hk,
csheng@cse.cuhk.edu.hk

Yi Yang, Shuigeng Zhou

Fudan University
yyang1@fudan.edu.cn, sgzhou@fudan.edu.cn

CP13

Multway Simple Cycle Separators and I/O-Efficient Algorithms for Planar Graphs

We revisit I/O-efficient solutions to a number of fundamental problems on planar graphs: single-source shortest paths, topological sorting, and computing strongly connected components. Existing I/O-efficient solutions to these problems pay for I/O efficiency using excessive computation time in internal memory, thereby completely negating the performance gain achieved by minimizing the number of disk accesses. In this paper, we show how to make these algorithms simultaneously efficient in internal and external memory so they achieve I/O complexity $O(\text{sort}(N))$ and take $O(N \log N)$ time in internal memory, where $\text{sort}(N)$ is the number of I/Os needed to sort N items in external memory. The key, and the main technical contribution of this paper, is a multway version of Miller's simple cycle separator theorem. We show how to compute these separators in linear time in internal memory, and using $O(\text{sort}(N))$ I/Os and $O(N \log N)$ (internal-memory computation) time in external memory.

Lars Arge, Freek Van Walderveen

MADALGO, Department of Computer Science
Aarhus University
large@cs.au.dk, freek_siam@vanwal.nl

Norbert Zeh
Dalhousie University
nzeh@cs.dal.ca

CP14

A Constant Factor Approximation Algorithm for Reordering Buffer Management

In the reordering buffer management problem (RBM) a sequence of n colored items enters a buffer with limited capacity k . When the buffer is full, one item is removed to the output sequence, making room for the next input item. This step is repeated until the input sequence is exhausted and the buffer is empty. The objective is to find a sequence of removals that minimizes the total number of color changes in the output sequence. The problem formalizes numerous applications in computer and production systems, and is known to be NP-hard. We give the first constant factor approximation guarantee for RBM. Our algorithm is based on an intricate “rounding” of the solution to an LP relaxation for RBM, so it also establishes a constant upper bound on the integrality gap of this relaxation. Our results improve upon the best previous bound of $O(\sqrt{\log k})$ of Adamaszek et al. (STOC 2011) that used different methods and gave an online algorithm. Our constant factor approximation beats the super-constant lower bounds on the competitive ratio given by Adamaszek et al. This is the first demonstration of a polynomial time offline algorithm for RBM that is provably better than any online algorithm.

Noa Avigdor-Elgrabli
Technion
noaelg@cs.technion.ac.il

Yuval Rabani
The Hebrew University of Jerusalem
yrabani@cs.huji.ac.il

CP14

Energy Efficient Scheduling of Parallelizable Jobs

We consider scheduling parallelizable jobs in the non-clairvoyant speed scaling setting to minimize the objective of weighted flow time plus energy. Previously, strong lower bounds for certain classes of algorithms were given even assuming constant resource augmentation and unweighted jobs. We overcome these lower bounds and give a *scalable* algorithm under the natural assumption that algorithms know the current parallelizability of alive jobs. Our algorithm is also the first positive result in this setting when considering general power functions.

Kyle Fox
University of Illinois, Urbana-Champaign
kylefox2@illinois.edu

Sungjin Im
Duke University
sungjin@cs.duke.edu

Benjamin Moseley
Toyota Technological Institute at Chicago
moseley@ttic.edu

CP14

Lyndon Words and Short Superstrings

In the Shortest Superstring Problem (SSP), we are given a set of strings $\{s_1, \dots, s_k\}$ and want to find a string that contains all s_i as substrings and has minimum length. This is a classical problem in approximation and the best known approximation factor is $2\frac{1}{2}$, given by Sweedyk in 1999. Since then no improvement has been made, however two other approaches yielding a $2\frac{1}{2}$ -approximation algorithms have been proposed by Kaplan et al. and recently by Paluch et al. — both based on a reduction to maximum asymmetric TSP path (MAX-ATSP) and structural results of Breslauer et al. In this paper we give an algorithm that achieves an approximation ratio of $2\frac{11}{23}$, breaking through the long-standing bound of $2\frac{1}{2}$. We use the standard reduction of SSP to MAX-ATSP. The new, somewhat surprising, algorithmic idea is to take the better of the two solutions obtained by using: (a) the currently best $\frac{2}{3}$ -approximation algorithm for MAX-ATSP, and (b) a naïve cycle-cover based $\frac{1}{2}$ -approximation algorithm. To prove that this indeed results in an improvement, we further develop a theory of string overlaps, extending the results of Breslauer et al. This theory is based on the novel use of Lyndon words, as a substitute for generic unbordered rotations and critical factorizations, as used by Breslauer et al.

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

CP14

New Approximability Results for Two-Dimensional Bin Packing

We study the two-dimensional bin packing problem: Given a list of n rectangles the objective is to find a feasible packing of all rectangles into the minimum number of unit sized bins. Rotation of the rectangles by 90° may be allowed or forbidden. We give a new asymptotic upper bound of our problem: For any fixed ε and any instance our algorithm computes a packing into $(1.5 + \varepsilon) \cdot \text{OPT} + 69$ bins with polynomial running time.

Lars Praedel
University of Kiel
lap@informatik.uni-kiel.de

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

CP14

Minimum Makespan Scheduling with Low Rank Processing Times

We investigate approximation algorithms for the classical minimum makespan scheduling problem, focusing on instances where the rank of the matrix describing the processing times of the jobs is bounded. A bounded rank matrix arises naturally when the processing time of a job on machine depends upon a bounded set of resources. A bounded rank matrix also shows up when jobs have varying degrees of parallelizability, and the machines have multiple cores. We are interested in studying the tractability of the problem as a function of the (positive) rank of

the processing-time matrix. At one extreme is the case of unit rank, also known as *related machines*, which admits a PTAS, and at the other extreme is the full rank case (*unrelated machines*), which is NP-hard to approximate within a factor better than $3/2$. Our main technical contribution is in showing that the approximability of the problem is *not* smooth with the rank of the matrix. From the inapproximability side, we show that the problem becomes APX-hard, *even for rank four matrices*. For rank seven matrices, we prove that it is hard to approximate to a factor $3/2$, matching the inapproximability result for general unrelated machines. From the algorithmic side, we obtain a quasi-polynomial approximation scheme (i.e., a $(1+\epsilon)$ approximation in time $n^{\text{poly}(1/\epsilon, \log n)}$) for the rank two case. This implies that the problem is not APX-hard in this case, unless NP has quasi-polynomial algorithms. Our algorithm is a subtle dynamic program which runs in polynomial time in some interesting special cases. The classification of the three dimensional problem remains open.

Aditya Bhaskara
EPFL
bhaskara@cs.princeton.edu

Ravishankar Krishnaswami
Princeton
ravishan@cs.cmu.edu

Kunal Talwar
Microsoft Research Silicon Valley
kunal@microsoft.com

Udi Wieder
Microsoft Research
uwieder@microsoft.com

CP15

Simultaneous PQ-Ordering with Applications to Constrained Embedding Problems

In this paper, we define and study the new problem SIMULTANEOUS PQ-ORDERING. Its input consists of a set of PQ-trees, which represent sets of circular orders of their leaves, together with a set of child-parent relations between these PQ-trees, such that the leaves of the child form a subset of the leaves of the parent. SIMULTANEOUS PQ-ORDERING asks whether orders of the leaves of each of the trees can be chosen *simultaneously*, that is, for every child-parent relation the order chosen for the parent is an extension of the order chosen for the child. We show that SIMULTANEOUS PQ-ORDERING is \mathcal{NP} -complete in general and we identify a family of instances that can be solved efficiently, the *2-fixed instances*. We show that this result serves as a framework for several other problems that can be formulated as instances of SIMULTANEOUS PQ-ORDERING.

Ignaz Rutter
Karlsruhe Institute of Technology
rutter@kit.edu

Thomas Bläsius
Karlsruhe Institute of Technology (KIT)
blaesius@kit.edu

CP15

List-Coloring Embedded Graphs

For any fixed surface Σ of genus g , we give an algorithm to

decide whether a graph G of girth at least five embedded in Σ is colorable from an assignment of lists of size three in time $O(|V(G)|)$. Furthermore, we can allow a subgraph (of any size) with at most s components to be precolored, at the expense of increasing the time complexity of the algorithm to $O(|V(G)|^{K(g+s)+1})$ for some absolute constant K ; in both cases, the multiplicative constant hidden in the O -notation depends on g and s . This also enables us to find such a coloring when it exists. The idea of the algorithm can be applied to other similar problems, e.g., 5-list-coloring of graphs on surfaces.

Zdenek Dvorak
Computer Science Institute
Charles University, Prague
rakdver@kam.mff.cuni.cz

Ken-ichi Kawarabayashi
National Institute of Informatics, Japan
k_keniti@nii.ac.jp

CP15

Totally Odd Subdivisions and Parity Subdivisions: Structures and Coloring

A *totally odd H -subdivision* means a subdivision of a graph H in which each edge of H corresponds to a path of odd length. Thus this concept is a generalization of a subdivision of H . In this paper, we give a structure theorem for graphs without a fixed graph H as a totally odd subdivision. Namely, every graph with no totally odd H -subdivision has a tree-decomposition such that each piece is either

1. after deleting bounded number of vertices, an “almost” embedded graph into a bounded-genus surface, or
2. after deleting bounded number of vertices, a bipartite graph, or
3. after deleting bounded number of vertices, a graph with maximum degree at most $f|H|$ for some function f of $|H|$ or a $6|H|$ -degenerate graph.

We also prove that any graph with no totally odd K_k -subdivision is $79k^2/4$ -colorable. The bound on the chromatic number is essentially best possible since the correct order of the magnitude of the chromatic number even for graphs with no K_k -subdivision is Θk^2 .

Ken-ichi Kawarabayashi
National Institute of Informatics, Japan
k_keniti@nii.ac.jp

CP15

5-Coloring $K_{3,k}$ -Minor-Free Graphs

We first show that if G is 6-color-critical 4-connected $K_{3,k}$ -minor-free graphs, then G has tree-width at most $g(k)$ for some function g of k . This allows us to show the following algorithmic result, which is of independent interest: For a 4-connected graph G and any k , there is an $O(n^3)$ algorithm to test whether or not G is 5-colorable. We also show the following algorithmic result. For a graph G and any k , there is an $O(n^3)$ algorithm to output one of the following:

1. a 5-coloring of G , or
2. a $K_{3,k}$ -minor of G , or
3. a minor R of G of order at most $f(k)$ ($f(k)$ comes

from the above theorem) such that R does not have a $K_{3,k}$ -minor nor is 5-colorable.

Ken-ichi Kawarabayashi
National Institute of Informatics, Japan
k_keniti@nii.ac.jp

CP15

Known Algorithms for Edge Clique Cover are Probably Optimal

In the Edge Clique Cover (ECC) problem, given a graph G and an integer k , we ask whether the edges of G can be covered with k complete subgraphs of G . Gramm et al. [JEA 2008] have shown a set of simple rules that reduce the number of vertices of G to 2^k , and no algorithm is known with significantly better running time bound than a brute-force search on this reduced instance. In this paper we show that the approach of Gramm et al. is essentially optimal: we present a polynomial time algorithm that reduces an arbitrary 3-CNF-SAT formula with n variables and m clauses to an equivalent ECC instance (G, k) with $k = O(\log n)$ and $|V(G)| = O(n + m)$. Consequently, there is no $2^{2^{o(k)}}$ poly(n) time algorithm for the ECC problem, unless the Exponential Time Hypothesis fails. To the best of our knowledge, these are the first results for a natural, fixed-parameter tractable problem, and proving that a doubly-exponential dependency on the parameter is essentially necessary.

Marek Cygan
IDSIA, University of Lugano
marek@idsia.ch

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

Michał Pilipczuk
University of Bergen
michal.pilipczuk@ii.uib.no

CP16

Fast Matrix Multiplication Using Coherent Configurations

We introduce a relaxation of the notion of tensor rank, called s -rank, and show that upper bounds on the s -rank of the matrix multiplication tensor imply upper bounds on the ordinary rank. In particular, if the “ s -rank exponent of matrix multiplication” equals 2, then $\omega = 2$. Using this machinery, we generalize the group-theoretic approach of Cohn and Umans from group algebras to general algebras. We identify *adjacency algebras of coherent configurations* as a promising family of algebras in the generalized framework. Finally, we prove a closure property involving symmetric powers of adjacency algebras, which enables us to prove nontrivial bounds on ω using *commutative coherent configurations* and suggests that commutative coherent configurations may be sufficient to prove $\omega = 2$.

Henry Cohn
Microsoft Research New England
cohn@microsoft.com

Christopher Umans
Computing + Mathematical Sciences
California Institute of Technology

umans@cms.caltech.edu

CP16

Lattice Sparsification and the Approximate Closest Vector Problem

We give a deterministic algorithm for solving the $(1 + \epsilon)$ approximate Closest Vector Problem (CVP) on any n dimensional lattice and any norm in $2^{O(n)}(1 + 1/\epsilon)^n$ time and $2^n \text{poly}(n)$ space. Our algorithm builds on the lattice point enumeration techniques of Micciancio and Voulgaris (STOC 2010) and Dadush, Peikert and Vempala (FOCS 2011), and gives an elegant, deterministic alternative to the “AKS Sieve” based algorithms for $(1 + \epsilon)$ -CVP (Ajtai, Kumar, and Sivakumar; STOC 2001 and CCC 2002). Furthermore, assuming the existence of a poly(n)-space and $2^{O(n)}$ time algorithm for exact CVP in the l_2 norm, the space complexity of our algorithm can be reduced to polynomial. Our main technical contribution is a method for “sparsifying” any input lattice while approximately maintaining its metric structure. To this end, we employ the idea of random sublattice restrictions, which was first employed by Khot (FOCS 2003) for the purpose of proving hardness for Shortest Vector Problem (SVP) under l_p norms.

Daniel Dadush
Georgia Institute of Technology
dndadush@gatech.edu

Gabor Kun
Courant Institute for Mathematical Sciences
kungabor@cs.elte.hu

CP16

Minimizing the Number of Lattice Points in a Translated Polygon

The parametric lattice-point counting problem is as follows: Given an integer matrix $A \in \mathbb{Z}^{m \times n}$, compute an explicit formula parameterized by $b \in \mathbb{R}^m$ that determines the number of integer points in the polyhedron $\{x \in \mathbb{R}^n : Ax \leq b\}$. In the last decade, this counting problem has received considerable attention in the literature. Several variants of Barvinok’s algorithm have been shown to solve this problem in polynomial time if the number n of columns of A is fixed. Central to our investigation is the following question:

Can one also efficiently determine a parameter b such that the number of integer points in $\{x \in \mathbb{R}^n : Ax \leq b\}$ is minimized?

Here, the parameter b can be chosen from a given polyhedron $Q \subseteq \mathbb{R}^m$. Our main result is a proof that finding such a minimizing parameter is NP -hard, even in dimension 2 and even if the parametrization reflects a translation of a 2-dimensional convex polygon. This result is established via a relationship of this problem to arithmetic progressions and simultaneous Diophantine approximation. On the positive side we show that in dimension 2 there exists a polynomial time algorithm for each fixed k that either determines a minimizing translation or asserts that any translation contains at most $1 + 1/k$ times the minimal number of lattice points.

Nicolai Hähnle
Technische Universität Berlin
nhaehnle@gmail.com

Friedrich Eisenbrand

Ecole Polytechnique Federale de Lausanne
friedrich.eisenbrand@epfl.ch

CP16**Algorithms for the Densest Sub-Lattice Problem**

We give algorithms for computing the densest k -dimensional sublattice of an arbitrary lattice, and related problems. This is an important problem in the algorithmic geometry of numbers that includes as special cases Rankin's problem (which corresponds to the densest sublattice problem with respect to the Euclidean norm, and has applications to the design of lattice reduction algorithms), and the shortest vector problem for arbitrary norms (which corresponds to setting $k = 1$) and its dual ($k = n - 1$). Our algorithm works for any norm and has running time $k^{O(k \cdot n)}$ and uses $2^n \text{poly}(n)$ space. In particular, the algorithm runs in single exponential time $2^{O(n)}$ for any constant $k = O(1)$.

Daniel Dadush
Georgia Institute of Technology
dndadush@gatech.edu

Daniele Micciancio
University of California - San Diego
daniele@eng.ucsd.edu

CP16**Breaking the $n^{\log n}$ Barrier for Solvable-Group Isomorphism**

The $n^{\log n}$ barrier for group isomorphism has withstood all attacks ever since the $n^{\log n + O(1)}$ generator-enumeration algorithm. We present the first significant improvement over $n^{\log n}$ by showing an $n^{(1/2) \log n + O(1)}$ algorithm for p -group isomorphism. We then generalize our techniques from p -groups and derive an $n^{(1/2) \log n + O(\log n / \log \log n)}$ algorithm for solvable-group isomorphism. Finally, we relate group isomorphism to the collision problem which allows us replace the $1/2$ in the exponents with $1/4$ using randomized algorithms and $1/6$ using quantum algorithms.

David J. Rosenbaum
University of Washington
djr@cs.washington.edu

CP17**Smoothed Analysis of the Successive Shortest Path Algorithm**

Several algorithms for the minimum-cost flow problem have been developed, and it seems that both the problem and the algorithms are well understood. However, some of the algorithms' running times observed in empirical studies contrast the running times obtained by worst-case analysis. We explain the good performance of the Successive Shortest Path (SSP) algorithm in empirical studies by studying the SSP algorithm in the framework of smoothed analysis.

Tobias Brunsch
University of Bonn
brunsch@cs.uni-bonn.de

Kamiel Cornelissen, Bodo Manthey
University of Twente

k.cornelissen@utwente.nl, b.manthey@utwente.nl

Heiko Röglin
University of Bonn
heiko@roeglin.org

CP17**Breaking the $O(n^{2.5})$ Deterministic Time Barrier for Undirected Unit-Capacity Maximum Flow**

This paper gives the first $o(n^{2.5})$ deterministic algorithm for the maximum flow problem in any undirected unit-capacity graph with no parallel edges. In an n -vertex, m -edge graph with maximum flow value v , our running time is $\tilde{O}(n^{9/4} v^{1/8}) = \tilde{O}(n^{2.375})$. Note that $v \leq n$ for simple unit-capacity graphs. The previous deterministic algorithms [Karger and Levine 1998] achieve $O(m + nv^{3/2})$ and $O(nm^{2/3} v^{1/6})$ time bound, which are both $O(n^{2.5})$ for dense simple graphs and $v = \Theta(n)$.

Ran Duan
University of Michigan
EECS Department
duanran@mpi-inf.mpg.de

CP17**Dynamic Graph Connectivity in Polylogarithmic Worst Case Time**

The dynamic graph connectivity problem is the following: given a graph on a fixed set of n nodes which is undergoing a sequence of edge insertions and deletions, answer queries of the form $q(a, b)$: "Is there a path between nodes a and b ?" While data structures for this problem with polylogarithmic *amortized* time per operation have been known since the mid-1990's, these data structures have $\Theta(n)$ worst case time. In fact, no previously known solution has worst case time per operation which is $o(\sqrt{n})$. We present a solution with worst case times $O(\log^4 n)$ per edge insertion, $O(\log^5 n)$ per edge deletion, and $O(\log n / \log \log n)$ per query. The answer to each query is correct if the answer is "yes" and is correct with high probability if the answer is "no". The data structure is based on a simple novel idea which can be used to quickly identify an edge in a cutset. Our technique can be used to simplify and significantly speed up the preprocessing time for the emergency planning problem while matching previous bounds for an update, and to approximate the sizes of cutsets of dynamic graphs in time $\tilde{O}(\min\{|S|, |V \setminus S|\})$ for an oblivious adversary.

Bruce M. Kapron
University of Victoria
Computer Science Department
bmkapron@uvic.ca

Valerie King, Ben Mountjoy
University of Victoria
val@uvic.ca, mountjoy@uvic.ca

CP17**Approximate Maximum Flow on Separable Undirected Graphs**

We present faster algorithms for approximate maximum flow in undirected graphs with good separator structures, such as bounded genus, minor free, and geometric

graphs. Given such a graph with n vertices, m edges along with a recursive \sqrt{n} -vertex separator structure, our algorithm finds an $1 - \epsilon$ approximate maximum flow in time $\tilde{O}(m^{6/5} \text{poly}(\epsilon^{-1}))$, ignoring poly-logarithmic terms. Similar speedups are also achieved for separable graphs with larger size separators albeit with larger run times. These bounds also apply to image problems in two and three dimensions. Key to our algorithm is an intermediate problem that we term grouped L_2 flow, which exists between maximum flows and electrical flows. Our algorithm also makes use of spectral vertex sparsifiers in order to remove vertices while preserving the energy dissipation of electrical flows. We also give faster spectral vertex sparsification algorithms on well separated graphs, which may be of independent interest.

Gary Miller, Richard Peng
Carnegie Mellon University
gmiller@cs.cmu.edu, yangp@cs.cmu.edu

CP17

Decremental Maintenance of Strongly Connected Components

We consider the problem of maintaining the strongly connected components (SCCs) of an n -nodes and m -edges directed graph that undergoes a sequence of edge deletions. Recently, in SODA 2011, Lacki presented a deterministic algorithm that preprocesses the graph in $O(mn)$ time and creates a data structure that maintains the SCCs of a graph under edge deletions with a total update time of $O(mn)$. The data structure answers strong connectivity queries in $O(1)$ time. The worst case update time after a single edge deletion might be as large as $O(mn)$. In this paper we reduce the preprocessing time and the worst case update time of Lacki's data structure from $O(mn)$ to $O(m \log n)$. The query time and the total update time remain unchanged.

Liam Roditty
Bar-Ilan University
liamr@macs.biu.ac.il

CP18

Regret Minimization for Reserve Prices in Second-Price Auctions

We show a regret minimization algorithm for setting the reserve price in second-price auctions. We make the assumption that all bidders draw their bids from the same unknown and arbitrary distribution. Our algorithm is computationally efficient, and achieves a regret of $O(T^{1/2})$, even when the number of bidders is stochastic with a known distribution.

Yishay Mansour
Tel Aviv University
mansour@cs.tau.ac.il

Nicolo' Cesa-Bianchi
Universita' degli Studi di Milano, Italy
nicolo.cesa-bianchi@unimi.it

Claudio Gentile
University of Insubria
claudio.gentile@uninsubria.it

CP18

A Unified Approach to Truthful Scheduling on Re-

lated Machines

We present a unified framework for designing deterministic monotone polynomial time approximation schemes (PTAS's) for a wide class of scheduling problems on uniformly related machines. This class includes (among others) minimizing the makespan, maximizing the minimum load, and minimizing the ℓ_p norm of the machine loads vector. Previously, this kind of result was only known for the makespan objective.

Leah Epstein
University of Haifa, Israel
lea@math.haifa.ac.il

Asaf Levin
The Technion
levinas@ie.technion.ac.il

Rob van Stee
Max Planck Institute for Informatics
vanstee@mpi-inf.mpg.de

CP18

Online Submodular Welfare Maximization: Greedy Is Optimal

We prove that no online algorithm (even randomized, against an oblivious adversary) is better than $1/2$ -competitive for welfare maximization with coverage valuations, unless $NP = RP$. Since the Greedy algorithm is known to be $1/2$ -competitive for monotone submodular valuations, of which coverage is a special case, this proves that Greedy provides the optimal competitive ratio. On the other hand, we prove that Greedy in a stochastic setting with i.i.d. items and valuations satisfying diminishing returns is $(1 - 1/e)$ -competitive, which is optimal even for coverage valuations, unless $NP = RP$. For online budget-additive allocation, we prove that no algorithm can be 0.612 -competitive with respect to a natural LP which has been used previously for this problem.

Ian Post, Michael Kapralov
Stanford University
itp@stanford.edu, kapralov@stanford.edu

Jan Vondrak
IBM Almaden Research Center
jvondrak@us.ibm.com

CP18

Towards Polynomial Simplex-Like Algorithms for Market Equilibria

In this paper we consider the problem of computing market equilibria in the Fisher setting for utility models such as spending constraint and perfect, price-discrimination. These models were inspired from modern e-commerce settings and attempt to bridge the gap between the computationally hard but realistic separable, piecewise-linear and concave utility model and, the tractable but less relevant linear utility case. While there are polynomial time algorithms known for these problems, the question of whether there exist polynomial time Simplex-like algorithms has remained elusive, even for linear markets. Such algorithms are desirable due to their conceptual simplicity, ease of implementation and practicality. This paper takes a significant step towards this goal by presenting the first Simplex-like algorithms for these markets assuming a positive res-

olution of an algebraic problem of Cucker, Koiran and Smale. Unconditionally, our algorithms are FPTASs; they compute prices and allocations such that each buyer derives at least a $\frac{1}{1+\epsilon}$ -fraction of the utility at a true market equilibrium, and their running times are polynomial in the input length and $\frac{1}{\epsilon}$. We start with convex programs which capture market equilibria in each setting and, in a systematic way, convert them into linear complementarity problem (LCP) formulations. Then, departing from previous approaches which try to pivot on a single polyhedron associated to the LCP obtained, we carefully construct a polynomial-length sequence of polyhedra, one containing the other, such that starting from an optimal solution to one allows us to obtain an optimal solution to the next in the sequence in a polynomial number of complementary pivot steps. Our framework to convert a convex program into an LCP and then come up with a Simplex-like algorithm that moves on a sequence of connected polyhedra may be of independent interest.

Jugal Garg, Ruta Mehta
Georgia Tech
jugal.garg@gmail.com, mehta.ruta@gmail.com

Milind Sohoni
IIT-Bombay
sohoni@cse.iitb.ac.in

Nisheeth K. Vishnoi
Microsoft Research India
nisheeth.vishnoi@gmail.com

CP18

Communication Complexity of Combinatorial Auctions with Submodular Valuations

We prove the first communication complexity lower bound for constant-factor approximation of the submodular welfare problem. More precisely, we show that a $(1 - \frac{1}{2e} + \epsilon)$ -approximation ($\simeq 0.816$) for welfare maximization in combinatorial auctions with submodular valuations would require exponential communication. We also show NP-hardness of $(1 - \frac{1}{2e} + \epsilon)$ -approximation in a computational model where each valuation is given explicitly by a table of constant size. Both results rule out better than $(1 - \frac{1}{2e})$ -approximations in every oracle model with a separate oracle for each player, such as the demand oracle model. Our main tool is a new construction of monotone submodular functions that we call *multi-peak submodular functions*. Roughly speaking, given a family of sets \mathcal{S} , we construct a monotone submodular function f with a high value $f(S)$ for every set $S \in \mathcal{S}$ (a “peak”), and a low value on every set that does not intersect significantly any set in \mathcal{S} .

Shahar Dobzinski
The Weizmann Institute of Science, Israel
dobzin@gmail.com

Jan Vondrak
IBM Almaden Research Center
jvondrak@us.ibm.com

CP19

Dichotomy for Holant* Problems with a Function on Domain Size 3

Holant problems are a general framework to study the algorithmic complexity of counting problems. Both count-

ing constraint satisfaction problems and graph homomorphisms are special cases. All previous results of Holant problems are over the Boolean domain. In this paper, we give the first dichotomy theorem for Holant problems for domain size > 2 . We discover unexpected tractable families of counting problems, by giving new polynomial time algorithms. This paper also initiates holographic reductions in domains of size > 2 . This is our main algorithmic technique, and is used for both tractable families and hardness reductions. The dichotomy theorem is the following: For any complex-valued symmetric function \mathbf{F} with arity 3 on domain size 3, we give an explicit criterion on \mathbf{F} , such that if \mathbf{F} satisfies the criterion then the problem $\text{Holant}^*(\mathbf{F})$ is computable in polynomial time, otherwise $\text{Holant}^*(\mathbf{F})$ is $\#\text{P-hard}$.

Jin-Yi Cai
University of Wisconsin, Madison
jyc@cs.wisc.edu

Pinyan Lu
Microsoft Research Asia
pinyanl@microsoft.com

Mingji Xia
Max Planck Institute for Informatics
xmjljx@gmail.com

CP19

Skew Bisubmodularity and Valued CSPs

The Valued Constraint Satisfaction Problem (VCSP) amounts to minimising a rational-valued function (of many variables) on a finite set, given as a sum of bounded-arity “basic” functions depending on subsets of the variables. We show that, when the set is small, any VCSP with a fixed set of “basic” functions either satisfies a submodularity-like condition and can be solved by LP, or else can express Max Cut in a specific way, and so is NP-hard.

Andrei Krokhin, Anna Huber, Robert Powell
Durham University
andrei.krokhin@durham.ac.uk,
anna.huber@durham.ac.uk, robert.powell@durham.ac.uk

CP19

Complexity of Sat Problems, Clone Theory and the Exponential Time Hypothesis

The construction of exact exponential-time algorithms for NP-complete problems has for some time been a very active research area. Unfortunately, there is a lack of general methods for studying and comparing the time complexity of algorithms for such problems. We propose such a method based on *clone theory* and demonstrate it on the SAT problem. Schaefer has completely classified the complexity of SAT with respect to the set of allowed relations and proved that this parameterized problem exhibits a dichotomy: it is either in P or is NP-complete. We show that there is a certain partial order on the NP-complete SAT problems with a close connection to their worst-case time complexities; if a problem $\text{SAT}(S)$ is below a problem $\text{SAT}(S')$ in this partial order, then $\text{SAT}(S')$ cannot be solved strictly faster than $\text{SAT}(S)$. By using this order, we identify a relation R such that $\text{SAT}(\{R\})$ is the *computationally easiest* NP-complete $\text{SAT}(S)$ problem. This result may be interesting when investigating the borderline between P and NP since one appealing way of studying this borderline is to identify problems that, in some

sense, are situated close to it. We strengthen the result by showing that $\text{SAT}(\{R\})\text{-}2$ (i.e. $\text{SAT}(\{R\})$ restricted to instances where no variable appears more than twice) is NP-complete, too. This is in contrast to, for example, 1-in-3-SAT, which is in P under the same restriction. We then relate $\text{SAT}(\{R\})\text{-}2$ to the exponential-time hypothesis (ETH) and show that ETH holds if and only if $\text{SAT}(\{R\})\text{-}2$ is not sub-exponential. This constitutes a strong connection between ETH and the SAT problem under both severe relational and severe structural restrictions. In the process, we also prove a stronger version of Impagliazzo et al.'s sparsification lemma for k -SAT; namely that all finite Boolean constraint languages S and S' such that $\text{SAT}(\cdot)$ is NP-complete can be sparsified into each other.

Victor Lagerkvist, Peter Jonsson, Gustav Nordh
Linköping University
victor.lagerkvist@liu.se, peter.jonsson@liu.se,
gustav.nordh@liu.se

Bruno Zanuttini
Université de Caen Basse-Normandie
bruno.zanuttini@unicaen.fr

CP19

Frozen Variables in Random Boolean Constraint Satisfaction Problems

We determine the exact *freezing threshold*, r^f , for a family of models of random boolean constraint satisfaction problems, including NAE-SAT and hypergraph 2-colouring, when the constraint size is sufficiently large. If the constraint-density of a random CSP, F , in our family is greater than r^f then for almost every solution of F , a linear number of variables are *frozen*, meaning that their colours cannot be changed by a sequence of alterations in which we change $o(n)$ variables at a time, always switching to another solution. If the constraint-density is less than r^f , then almost every solution has $o(n)$ frozen variables.

Michael Molloy
University of Toronto
molloy@cs.toronto.edu

Ricardo Restrepo
Universidad di Antioquia
ricrest@gmail.com

CP19

Exponential Lower Bounds for the PPSZ k -Sat Algorithm

In 1998, Paturi, Pudlák, Saks, and Zane presented PPSZ, an elegant randomized algorithm for k -SAT. Fourteen years on, this algorithm is still the fastest known worst-case algorithm. They proved that its expected running time on k -CNF formulas with n variables is at most $2^{(1-\epsilon_k)n}$, where $\epsilon_k \in \Omega(1/k)$. So far, no exponential lower bounds at all have been known. In this paper, we construct hard instances for PPSZ. That is, we construct satisfiable k -CNF formulas over n variables on which the expected running time is at least $2^{(1-\epsilon_k)n}$, for $\epsilon_k \in O(\log^2 k/k)$.

Dominik A. Scheder
Aarhus Universitet
dominik.scheder@gmail.com

Shiteng Chen, Bangsheng Tang

Tsinghua University
shitengchen@gmail.com, bangsheng.tang@gmail.com

Navid Talebanfard
Aarhus University
navid@cs.au.dk

CP20

Testing Low-Complexity Affine-Invariant Properties

Invariance with respect to linear/affine transformations of the domain is arguably the most common symmetry exhibited by natural algebraic properties. In this work, we show that any *low complexity* affine-invariant property of multivariate functions over finite fields is testable with a constant number of queries. This immediately reproves, for instance, that the Reed-Muller code over F_p of degree $d < p$ is testable, via an argument that uses no detailed algebraic information about polynomials.

Arnab Bhattacharyya
Center for Computational Intractability, Princeton, NJ,
abhattach@mit.edu

Eldar Fischer
Technion Israel Institute of Technology,
Haifa, Israel
eldar@cs.technion.ac.il

Shachar Lovett
IAS
slovett@math.ias.edu

CP20

Learning Mixtures of Structured Distributions over Discrete Domains

Let C be a class of probability distributions over the discrete domain $[n] = \{1, \dots, n\}$. We show that if C satisfies a rather general condition – essentially, that each distribution in C can be well-approximated by a variable-width histogram with few bins – then there is a highly efficient (both in terms of running time and sample complexity) algorithm that can learn any mixture of k unknown distributions from C . We analyze several natural types of distributions over $[n]$, including log-concave, monotone hazard rate and unimodal distributions, and show that they have the required structural property of being well-approximated by a histogram with few bins. Applying our general algorithm, we obtain near-optimally efficient algorithms for all these mixture learning problems.

Ilias Diakonikolas
University of Edinburgh
ilias.d@ed.ac.uk

Siu-On Chan
UC Berkeley
siuon@cs.berkeley.edu

Rocco A. Servedio, Xiaorui Sun
Columbia University
rocco@cs.columbia.edu, xiaoruisun@cs.columbia.edu

CP20

Exponentially Improved Algorithms and Lower

Bounds for Testing Signed Majorities

A signed majority function is a linear threshold function $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ of the form $f(x) = \text{sign}(\sum_{i=1}^n \sigma_i x_i)$ where each $\sigma_i \in \{-1, 1\}$. Signed majority functions are a highly symmetrical subclass of the class of all linear threshold functions. We study the query complexity of testing whether an unknown $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ is a signed majority function versus ϵ -far from every signed majority function. As our main results we exponentially improve prior upper and lower bounds on the number of queries needed to test signed majority functions.

Rocco A. Servedio
Columbia University
rocco@cs.columbia.edu

Dana Ron
Tel Aviv University
danaron@post.tau.ac.il

CP20

Learning and Testing Integer Submodular Functions

We prove that any submodular function $f : \{0, 1\}^n \rightarrow \{0, 1, \dots, k\}$ can be represented as a pseudo-Boolean $2k$ -DNF formula. Pseudo-Boolean DNFs are a natural generalization of DNF representation for functions with integer range. Each term in such a formula has an associated integral constant. We show that an analog of switching lemma holds for pseudo-Boolean k -DNFs if all constants associated with the terms of the formula are bounded. Our characterization implies efficient learning and testing algorithms for integer submodular functions.

Grigory Yaroslavtsev
Penn State University
grigory@cse.psu.edu

Sofya Raskdhonikova
Pennsylvania State University
sofya@cse.psu.edu

CP20

Learning Disjunctions: Near-Optimal Trade-off Between Mistakes and “I Don’t Know”

We develop polynomial-time online algorithms for learning disjunctions. In this model, we are given an online adversarial sequence of inputs for an unknown function of the form $f(x_1, x_2, \dots, x_n) = \bigvee_{i \in S} x_i$, and for each such input, we must guess “true”, “false”, or “I don’t know”. On the algorithm side, we show how to make at most ϵn mistakes while answering “I don’t know” at most $(1/\epsilon)^{2^{O(1/\epsilon)}}$ n times. Furthermore, we show how to make $O\left(n \frac{\log \log n}{\log n}\right)$ mistakes while answering “I don’t know” $O(n^2 \log \log n)$ times. We also show that any algorithm making $o(n/\log n)$ mistakes must answer “I don’t know” a superpolynomial number of times.

Morteza Zadimoghaddam
MIT
MIT
morteza@mit.edu

Erik Demaine
MIT

edemaine@mit.edu

CP21

On Differentially Private Low Rank Approximation

Low rank approximation is a fundamental computational primitive widely used in data analysis. In many applications the dataset that the algorithm operates on may contain sensitive information about contributing individuals (e.g. user/movie ratings in the Netflix challenge), motivating the need to design low rank approximation algorithms that preserve privacy of individual entries of the input matrix. In this paper, we give a polynomial time algorithm that, given a privacy parameter $\epsilon > 0$, for a symmetric matrix A , outputs an ϵ -differentially approximation to the principal eigenvector of A , and then show how this algorithm can be used to obtain a differentially private rank- k approximation. We also provide lower bounds showing that our utility/privacy tradeoff is close to best possible. While there has been significant progress on this problem recently for a weaker notion of privacy, namely (ϵ, δ) -differential privacy [HR12, BBDS12], our result is the first to achieve $(\epsilon, 0)$ -differential privacy guarantees with a near-optimal utility/privacy tradeoff in polynomial time.

Michael Kapralov
Stanford University
kapralov@stanford.edu

Kunal Talwar
Microsoft Research Silicon Valley
kunal@microsoft.com

CP21

The Power of Linear Reconstruction Attacks

We consider the power of linear reconstruction attacks in statistical data privacy by mounting attacks on any release which gives:

1. the fraction of records that satisfy a given non-degenerate boolean function. Such releases include contingency tables as well as more complex outputs like the error rate of certain classifiers such as decision trees;
2. any one of a large class of M-estimators including the standard estimators for linear and logistic regression.

Shiva Prasad Kasiviswanathan
GE Global Research
kasivisw@gmail.com

Mark Rudelson
University of Michigan
rudelson@umich.edu

Adam Smith
Pennsylvania State University
asmith@cse.psu.edu

CP21

An Almost Optimal Algorithm for Computing Nonnegative Rank

Here, we give an algorithm for deciding if the nonnegative rank of a matrix M of dimension $m \times n$ is at most r which runs in time $(nm)^{O(r^2)}$. This is the first exact algorithm that runs in time singly-exponential in r . This algorithm

(and earlier algorithms) are built on methods for finding a solution to a system of polynomial inequalities (if one exists). Notably, the best algorithms for this task run in time exponential in the number of variables but polynomial in all of the other parameters (the number of inequalities and the maximum degree). Hence these algorithms motivate natural *algebraic* questions whose solution have immediate *algorithmic* implications: How many variables do we need to represent the decision problem, does M have nonnegative rank at most r ? A naive formulation uses $nr + mr$ variables and yields an algorithm that is exponential in n and m even for constant r . (Arora, Ge, Kannan, Moitra, STOC 2012) recently reduced the number of variables to $2r^2 2^r$, and here we *exponentially* reduce the number of variables to $2r^2$ and this yields our main algorithm. In fact, the algorithm that we obtain is nearly-optimal (under the Exponential Time Hypothesis) since an algorithm that runs in time $(nm)^{o(r)}$ would yield a subexponential algorithm for 3-SAT. Our main result is based on establishing a normal form for nonnegative matrix factorization – which in turn allows us to exploit algebraic dependence among a large collection of linear transformations with variable entries. Additionally, we also demonstrate that nonnegative rank cannot be certified by even a very large submatrix of M , and this property also follows from the intuition gained from viewing nonnegative rank through the lens of systems of polynomial inequalities.

Ankur Moitra
IAS
moitra@ias.edu

CP21

The Simplex Method Is Strongly Polynomial for Deterministic Markov Decision Processes

We prove that the simplex method with the highest gain/most-negative-reduced cost pivoting rule converges in strongly polynomial time for deterministic Markov decision processes (MDPs) regardless of the discount factor. For a deterministic MDP with n states and m actions, we prove the simplex method runs in $O(n^3 m^2 \log^2 n)$ iterations if the discount factor is uniform and $O(n^5 m^3 \log^2 n)$ iterations if each action has a distinct discount factor. Previously the simplex method was known to run in polynomial time only for discounted MDPs where the discount was bounded away from 1 [Ye 2011].

Ian Post, Yinyu Ye
Stanford University
itp@stanford.edu, yinyu-ye@stanford.edu

CP21

Turning Big Data into Tiny Data: Constant-Size Coresets for K-Means, PCA and Projective Clustering

Projective Clustering is the task to find k subspaces of dimension j minimizing the sum of the squared distances to a set of points. It includes popular problems as the k -means problem and principal component analysis. We observe that projecting points to a $O(\|\cdot\|/\epsilon)$ -dimensional best fit subspace preserves the cost function up to a factor of $(1 + \epsilon)$. Based on this, we provide coresets for projective clustering problems containing a number of coreset points independent of the number of input points and their dimension, leading to fast streaming and parallel algorithms.

Melanie Schmidt
TU Dortmund
melanie.schmidt@tu-dortmund.de

Dan Feldman
MIT
dannypf.post@gmail.com

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

CP22

Active Self-Assembly of Simple Units Using An Insertion Primitive

We describe a formal model for studying the complexity of self-assembled structures with active molecular components that utilize an insertion primitive. We can directly map our model onto a molecular implementation using DNA. The expressive power of this language is stronger than regular languages, but at most as strong as context free grammars. We explore the trade-off between the complexity of the system and the behavior of the system and speed of its assembly.

Nadine Dabby
California Institute of Technology
ndabby@caltech.edu

Ho-Lin Chen
National Taiwan University
holinc@gmail.com

CP22

Nested Quantum Walks with Quantum Data Structures

We develop a new framework that extends the quantum walk framework of Magniez, Nayak, Roland, and Santha, by utilizing the idea of quantum data structures to construct an efficient method of nesting quantum walks. Surprisingly, only classical data structures were considered before for searching via quantum walks. The recently proposed learning graph framework of Belovs has yielded improved upper bounds for several problems, including triangle finding and more general subgraph detection. We exhibit the power of our framework by giving a simple explicit constructions that reproduce both the $O(n^{35/27})$ and $O(n^{9/7})$ learning graph upper bounds (up to logarithmic factors) for triangle finding, and discuss how other known upper bounds in the original learning graph framework can be converted to algorithms in our framework. We hope that the ease of use of this framework will lead to the discovery of new upper bounds.

Stacey Jeffery, Robin Kothari
Institute for Quantum Computing
University of Waterloo
sjeffery@uwaterloo.ca, rkothari@cs.uwaterloo.ca

Frederic Magniez
CNRS, LIAFA, Univ Paris Diderot
frederic.magniez@univ-paris-diderot.fr

CP22**Improved Quantum Query Algorithms for Triangle Finding and Associativity Testing**

We show that the quantum query complexity of detecting if an n -vertex graph contains a triangle is $O(n^{9/7})$. This improves the previous best algorithm of Belovs making $O(n^{35/27})$ queries. For the problem of determining if an operation $\circ : S \times S \rightarrow S$ is associative, we give an algorithm making $O(|S|^{10/7})$ queries, the first improvement to the trivial $O(|S|^{3/2})$ application of Grover search. Our algorithms are designed using the learning graph framework of Belovs. We give a family of algorithms for detecting constant-sized subgraphs, which can possibly be directed and colored. These algorithms are designed in a simple high-level language; our main theorem shows how this high-level language can be compiled as a learning graph and gives the resulting complexity. The key idea to our improvements is to allow more freedom in the parameters of the database kept by the algorithm. As in our previous work, the edge slots maintained in the database are specified by a graph whose edges are the union of regular bipartite graphs, the overall structure of which mimics that of the graph of the certificate. By allowing these bipartite graphs to be unbalanced and of variable degree we obtain better algorithms.

Frederic Magniez
CNRS, LIAFA, Univ Paris Diderot
frederic.magniez@univ-paris-diderot.fr

Troy Lee
Centre for Quantum Technologies
National University of Singapore
troylee@gmail.com

Miklos Santha
CNRS, LIAFA, Univ Paris Diderot
CQT, National University of Singapore
miklos.santha@liafa.univ-paris-diderot.fr

CP22**Fuel Efficient Computation in Passive Self-Assembly**

In this paper we show that passive self-assembly in the context of the tile self-assembly model is capable of performing fuel efficient, universal computation. A computationally universal tile system is said to be fuel efficient if the number of tiles used up per computation step is bounded by a constant.

Robert T. Schweller, Michael Sherman
University of Texas - Pan American
rtschweller@utpa.edu, msherman1225@gmail.com

CP22**Efficient Protocols for Generating Bipartite Classical Distributions and Quantum States**

We investigate the fundamental problem of generating bipartite classical distributions or quantum states. By designing efficient communication protocols and proving their optimality, we establish a number of intriguing connections to fundamental measures in optimization, convex geometry, and information theory.

1. To generate a classical distribution $P(x, y)$, we tightly characterize the minimum amount of quantum com-

munication needed by the *psd-rank* of P (as a matrix), a measure recently proposed by Fiorini, Massar, Pokutta, Tiwary and de Wolf (*Proceedings of the 44th ACM Symposium on Theory of Computing*, pages 95-106, 2012) in studies of the minimum size of extended formulations of optimization problems such as TSP. This echoes the previous characterization for the optimal classical communication cost by the *nonnegative rank* of P . The result is obtained via investigating the more general case of bipartite *quantum* state generation and designing an optimal protocol for it.

2. When an approximation of ϵ is allowed to generate a distribution $(X, Y) \sim P$, we present a classical protocol of the communication cost $O((C(X, Y) + 1)/\epsilon)$, where $C(X, Y)$ is *common information*, a well-studied measure in information theory introduced by Wyner (*IEEE Transactions on Information Theory*, 21(2):163-179, 1975). This also links nonnegative rank and common information, two seemingly unrelated quantities in different fields.
3. For approximately generating a quantum pure state $|\psi\rangle$, we completely characterize the minimum cost by a corresponding approximate rank, closing a possibly exponential gap left in Ambainis, Schulman, Ta-Shma, Vazirani and Wigderson (*SIAM Journal on Computing*, 32(6):1570-1585, 2003).

Shengyu Zhang
The Chinese University of Hong Kong
syzhang@cse.cuhk.edu.hk

Rahul Jain
National University of Singapore
rahul@comp.nus.edu.sg

Yaoyun Shi
University of Michigan
shiyy@eecs.umich.edu

Zhaohui Wei
National University of Singapore
cqtzw@nus.edu.sg

CP23**Local-Search Based Approximation Algorithms for Mobile Facility Location Problems**

We consider the *mobile facility location* (MFL) problem. We are given a set of facilities and clients located in a common metric space $G = (V, c)$. The goal is to move each facility from its initial location to a destination (in V) and assign each client to the destination of some facility so as to minimize the sum of the movement-costs of the facilities and the client-assignment costs. We give the first *local-search based* approximation algorithm for this problem and achieve the best-known approximation guarantee. Our main result is $(3 + \epsilon)$ -approximation for this problem for any constant $\epsilon > 0$ using local search. The previous best guarantee for MFL was an 8-approximation algorithm due to Friggstad and Salavatipour based on LP-rounding. Our guarantee *matches* the best-known approximation guarantee for the k -median problem. Since there is an approximation-preserving reduction from the k -median problem to MFL, any improvement of our result would imply an analogous improvement for the k -median problem. Furthermore, *our analysis is tight* (up to $o(1)$ factors) since the tight example for the local-search based 3-approximation algorithm for k -median can be easily adapted to show that our local-search algorithm has a

tight approximation ratio of 3. One of the chief novelties in the analysis is that in order to generate a suitable collection of local-search moves whose resulting inequalities yield the desired bound on the cost of a local-optimum, we define a ‘tree-like structure that (loosely speaking) functions as a ‘recursion tree’, using which we spawn off local-search moves by exploring this tree to a constant depth. Our results extend to the weighted generalization wherein each facility i has a non-negative weight w_i and the movement cost for i is w_i times the distance traveled by i .

Chaitanya Swamy, Sara Ahmadian, Zachary Friggstad
University of Waterloo
cswamy@math.uwaterloo.ca, sahma-
dian@math.uwaterloo.ca, zfriggstad@math.uwaterloo.ca

CP23
The Power of Non-Uniform Wireless Power

We study a fundamental measure for wireless interference in the SINR model, known as (weighted) inductive independence. This measure characterizes the effectiveness of using *oblivious* power — when the power used by a transmitter only depends on the distance to the receiver — as a mechanism for improving wireless capacity. We prove optimal bounds for inductive independence, implying a number of algorithmic applications. An algorithm is provided that achieves — due to existing lower bounds — capacity that is asymptotically best possible using oblivious power assignments. Improved approximation algorithms are provided for a number of problems for oblivious power and for power control, including distributed scheduling, connectivity, secondary spectrum auctions, and dynamic packet scheduling.

Magnus M. Halldorsson
Reykjavik University
mmh@ru.is

Stephan Holzer
ETH Zurich
stholzer@ethz.ch

Pradipta Mitra
Reykjavik University
ppmitra@gmail.com

Roger Wattenhofer
ETH Zurich
wattenhofer@ethz.ch

CP23
The Diffusion of Networking Technologies

There has been significant interest in the networking community on the impact of cascade effects on the diffusion of networking technology upgrades in the Internet. Thinking of the global Internet as a graph, where each node represents an economically-motivated Internet Service Provider (ISP), a key problem is to determine the smallest set of nodes that can trigger a cascade that causes every other node in the graph to adopt the protocol. We design the first approximation algorithm with a provable performance guarantee for this problem, in a model that captures the following key issue: a node’s decision to upgrade should be influenced by the decisions of the remote nodes it wishes to communicate with. Given an internet network $G(V, E)$ and threshold function θ , we assume that node u activates (upgrades to the new technology) when it is adjacent to a

connected component of active nodes in G of size exceeding node u ’s threshold $\theta(u)$. Our objective is to choose the smallest set of nodes that can cause the rest of the graph to activate. Our main contribution is an approximation algorithm based on linear programming, which we complement with computational hardness results and a near-optimum integrality gap. Our algorithm, which does not rely on sub-modular optimization techniques, also highlights the substantial algorithmic difference between our problem and similar questions studied in the context of social networks.

Zhenming Liu
Harvard University
zhenming@cs.princeton.edu

Sharon Goldberg
Boston University
goldbe@cs.bu.edu

CP23
Graph Products Revisited: Tight Approximation Hardness of Induced Matching, Poset Dimension and More

Graph product is a fundamental tool with rich applications in both graph theory and theoretical computer science. It is usually studied in the form $f(G * H)$ where G and H are graphs, $*$ is a graph product and f is a graph property. For example, if f is the *independence number* and $*$ is the *disjunctive product*, then the product is known to be *multiplicative*: $f(G * H) = f(G)f(H)$. In this paper, we study graph products in the following non-standard form: $f((G \oplus H) * J)$ where G, H and J are graphs, \oplus and $*$ are two different graph products and f is a graph property. We show that if f is the *induced and semi-induced matching number*, then for some products \oplus and $*$, it is *subadditive* in the sense that $f((G \oplus H) * J) \leq f(G * J) + f(H * J)$. Moreover, when f is the *poset dimension number*, it is *almost* subadditive. As applications of this result (we only need $J = K_2$ here), we obtain tight hardness of approximation for various problems in discrete mathematics and computer science: bipartite induced and semi-induced matching (a.k.a. maximum expanding sequences), poset dimension, maximum feasible subsystem with 0/1 coefficients, unit-demand min-buying and single-minded pricing, donation center location, boxicity, cubicity, threshold dimension and independent packing.

Parinya Chalermsook
IDSIA
parinya@cs.uchicago.edu

Danupon Nanongkai
University of Vienna
danupon@gmail.com

Bundit Laekhanukit
School of Computer Science
McGill University
blaekh@cs.mcgill.ca

CP23
Approximability and Proof Complexity

This work is concerned with the proof-complexity of certifying that optimization problems do *not* have good solutions. Specifically we consider bounded-degree ‘Sum of Squares’ (SOS) proofs, a powerful algebraic proof system

introduced in 1999 by Grigoriev and Vorobjov. Work of Shor, Lasserre, and Parrilo shows that this proof is automatizable using semidefinite programming (SDP), meaning that any n -variable degree- d proof can be found in time $n^{O(d)}$. Furthermore, the SDP is dual to the well-known Lasserre SDP hierarchy, meaning that the “ $d/2$ -round Lasserre value’ of an optimization problem is equal to the best bound provable using a degree- d SOS proof. These ideas were exploited in a recent paper by Barak et al. (STOC 2012) which shows that the known “hard instances’ for the Unique-Games problem are in fact optimally solved a constant level of the Lasserre SDP hierarchy. We continue the study of the power of SOS proofs in the context of difficult optimization problems. In particular, we show that the Balanced-Separator integrality gap instances proposed by Devanur et al. can have their optimal value certified by a degree-4 SOS proof. The key ingredient is an SOS proof of the KKL Theorem. We also investigate the extent to which the Khot–Vishnoi Max-Cut integrality gap instances can have their optimum value certified by an SOS proof. We show they can be certified to within a factor .952 ($> .878$) using a constant-degree proof. These investigations also raise an interesting mathematical question: is there a constant-degree SOS proof of the Central Limit Theorem?

Yuan Zhou

Carnegie Mellon University
yuanzhou@cs.cmu.edu

Ryan O’Donnell

Carnegie Mellon University, USA
odonnell@cs.cmu.edu

CP24

Transforming Curves on Surfaces Redux

Almost exactly 100 years ago, Max Dehn described the first combinatorial algorithm to determine whether two given cycles on a compact surface are homotopic, meaning one cycle can be continuously deformed into the other without leaving the surface. We describe a simple variant of Dehn’s algorithm that runs in linear time, with no hidden dependence on the genus of the surface. Specifically, given two closed vertex-edge walks of length at most ℓ in a combinatorial surface of complexity n , our algorithm determines whether the walks are homotopic in $O(n + \ell)$ time. Our algorithm simplifies and corrects a similar algorithm of Dey and Guha [JCSS 1999] and simplifies the more recent algorithm of Lazarus and Rivaud [FOCS 2012], who identified a subtle flaw in Dey and Guha’s results. Our algorithm combines components of these earlier algorithms, classical results in small cancellation theory by Gersten and Short [Inventiones 1990], and simple run-length encoding.

Jeff Erickson, Kim Whittlesey

University of Illinois, Urbana-Champaign
jeffe@cs.uiuc.edu, kwhittle@illinois.edu

CP24

Higher-Order Geodesic Voronoi Diagrams in a Polygonal Domain with Holes

Given a set of n point sites in a simple polygon with h polygonal holes and c corners, the k^{th} -order geodesic Voronoi diagram partitions the free space into several regions such that all points in a region share the same k nearest sites. We prove that the total complexity of such a diagram is $\Theta(k(n - k) + kc)$ and the number of faces is

$\Theta(k(n - k) + kh)$, and also show that it can be computed in $O(k^2(n + c)\log(n + c))$ time.

Chih-Hung Liu

Research Center for Information Technology Innovation
Academia Sinica
chliu_10@citi.sinica.edu.tw

D. T. Lee

National Chung Hsing University
dtlee@nchu.edu.tw

CP24

Eps-Samples for Kernels

We study the worst case error of kernel density estimates via subset approximation. A kernel density estimate of a distribution is the convolution of that distribution with a fixed kernel (e.g. Gaussian kernel). Given a subset (i.e. a point set) of the input distribution, we can compare the kernel density estimates of the input distribution with that of the subset and bound the worst case error. If the maximum error is ε , then this subset can be thought of as an ε -sample (aka an ε -approximation) of the range space defined with the input distribution as the ground set and the fixed kernel representing the family of ranges. Interestingly, in this case the ranges are not binary, but have a continuous range (for simplicity we focus on kernels with range of $[0, 1]$); these allow for smoother notions of range spaces. It turns out, the use of this smoother family of range spaces has an added benefit of greatly decreasing the size required for ε -samples. For instance, in the plane the size is $O((1/\varepsilon^{4/3})\log^{2/3}(1/\varepsilon))$ for disks (based on VC-dimension arguments) but is only $O((1/\varepsilon)\sqrt{\log(1/\varepsilon)})$ for Gaussian kernels and for kernels with bounded slope that only affect a bounded domain. These bounds are accomplished by studying the discrepancy of these “kernel’ range spaces, and here the improvement in bounds are even more pronounced. In the plane, we show the discrepancy is $O(\sqrt{\log n})$ for these kernels, whereas for balls there is a lower bound of $\Omega(n^{1/4})$.

Jeff Phillips

University of Utah
jeffp@cs.utah.edu

CP24

Morphing Planar Graph Drawings with a Polynomial Number of Steps

In 1944, Cairns proved the following theorem: given any two straight-line planar drawings of a triangulation with the same outer face, there exists a morph (i.e., a continuous transformation) between the two drawings so that the drawing remains straight-line planar at all times. Cairns’s original proof required exponentially many morphing steps. We prove that there is a morph that consists of $O(n^2)$ steps, where each step is a linear morph that moves each vertex at constant speed along a straight line. Using a known result on compatible triangulations this implies that for a general planar graph G and any two straight-line planar drawings of G with the same embedding, there is a morph between the two drawings that preserves straight-line planarity and consists of $O(n^4)$ steps.

Anna Lubiw

University of Waterloo, Canada
alubiw@uwaterloo.ca

Soroush Alamdari
David R. Cheriton School of Computer Science,
University of Waterloo
s26hosseinialamdari@uwaterloo.ca

Patrizio Angelini
Dipartimento di Informatica e Automazione,
Roma Tre University
angelini@dia.uniroma3.it

Timothy M. Chan
David R. Cheriton School of Computer Science,
University of Waterloo
tmchan@uwaterloo.ca

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

Fabrizio Frati
School of Information Technologies,
The University of Sydney
brillo@it.usyd.edu.au

Maurizio Patrignani, Vincenzo Roselli
Dipartimento di Informatica e Automazione,
Roma Tre University
patrigna@dia.uniroma3.it, roselli@dia.uniroma3.it

Sahil Singla
David R. Cheriton School of Computer Science,
University of Waterloo
s2singla@uwaterloo.ca

Bryan T. Wilkinson
MADALGO, Aarhus University
btw@cs.au.dk

CP24

Combinatorial and Geometric Properties of Planar Laman Graphs

Laman graphs naturally arise in structural mechanics and rigidity theory. Specifically, they characterize minimally rigid planar bar-and-joint systems which are frequently needed in robotics, as well as in molecular chemistry and polymer physics. We introduce three new combinatorial structures for planar Laman graphs: angular structures, angle labelings, and edge labelings. The latter two structures are related to Schnyder realizers for maximally planar graphs. We prove that planar Laman graphs are exactly the class of graphs that have an angular structure that is a tree, called angular tree, and that every angular tree has a corresponding angle labeling and edge labeling. Using a combination of these powerful combinatorial structures, we show that every planar Laman graph has an L-contact representation, that is, planar Laman graphs are contact graphs of axis-aligned L-shapes. Moreover, we show that planar Laman graphs and their subgraphs are the only graphs that can be represented this way. We present efficient algorithms that compute, for every planar Laman graph G , an angular tree, angle labeling, edge labeling, and finally an L-contact representation of G . The overall running time is $O(n^2)$, where n is the number of vertices of G , and the L-contact representation is realized on the

$n \times n$ grid.

Stephen Kobourov
Department of Computer Science
University of Arizona
kobourov@cs.arizona.edu

Torsten Ueckerdt
Charles University
torsten@kam.mff.cuni.cz

Kevin Verbeek
TU Eindhoven
k.a.b.verbeek@tue.nl

CP25

Better Bounds for Matchings in the Streaming Model

In this paper we present improved bounds for approximating maximum matchings in bipartite graphs in the streaming model. First, we consider the question of how well maximum matching can be approximated in a single pass over the input when $\tilde{O}(n)$ space is allowed, where n is the number of vertices in the input graph. Two natural variants of this problem have been considered in the literature: (1) the edge arrival setting, where edges arrive in the stream and (2) the vertex arrival setting, where vertices on one side of the graph arrive in the stream together with all their incident edges. The latter setting has also been studied extensively in the context of *online algorithms*, where each arriving vertex has to either be matched irrevocably or discarded upon arrival. In the online setting, the celebrated algorithm of Karp-Vazirani-Vazirani achieves a $1 - 1/e$ approximation by crucially using randomization (and using $\tilde{O}(n)$ space). Despite the fact that the streaming model is less restrictive in that the algorithm is not constrained to match vertices irrevocably upon arrival, the best known approximation in the streaming model with vertex arrivals and $\tilde{O}(n)$ space is the same factor of $1 - 1/e$. We show that no single pass streaming algorithm constrained to use $\tilde{O}(n)$ space can achieve a better than $1 - 1/e$ approximation to maximum matching, even in the vertex arrival setting. This leads to the striking conclusion that no single pass streaming algorithm can get any advantage over online algorithms unless it uses significantly more than $\tilde{O}(n)$ space. Second, we consider the problem of approximating matchings in multiple passes in the vertex arrival setting. We show that a simple fractional load balancing approach achieves approximation ratio $1 - e^{-k}k^{k-1}/(k-1)! = 1 - \frac{1}{\sqrt{2\pi k}} + o(1/k)$ in k passes using linear space, improving upon the state of the art for this setting.

Michael Kapralov
Stanford University
kapralov@stanford.edu

CP25

Beating the Direct Sum Theorem in Communication Complexity with Implications for Sketching

A direct sum theorem for two parties and a function f states that the communication cost of solving k copies of f simultaneously with error probability $1/3$ is at least $k \cdot R_{1/3}(f)$, where $R_{1/3}(f)$ is the communication required to solve a single copy of f with error probability $1/3$. We improve this for a natural family of functions f , showing that

the 1-way communication required to solve k copies of f simultaneously with probability $2/3$ is $\Omega(k \cdot R_{1/k}(f))$. Since $R_{1/k}(f)$ may be as large as $\Omega(R_{1/3}(f) \cdot \log k)$, we asymptotically beat the direct sum bound for such functions, showing that the trivial upper bound of solving each of the k copies of f with probability $1 - O(1/k)$ and taking a union bound is optimal! In order to achieve this, our direct sum involves a novel measure of information cost which allows a protocol to abort with constant probability, and otherwise must be correct with very high probability. Moreover, for the functions considered, we show strong lower bounds on the communication cost of protocols with these relaxed guarantees; indeed, our lower bounds match those for protocols that are not allowed to abort. In the distributed and streaming models, where one wants to be correct not only on a single query, but simultaneously on a sequence of n queries, we obtain optimal lower bounds on the communication or space complexity. Lower bounds obtained from our direct sum result show that a number of techniques in the sketching literature are optimal, including the following: dimension of (oblivious) Johnson-Lindenstrauss transforms, multiple ℓ_1 or ℓ_2 -norm estimation, sketch size for approximate matrix multiplication and database join size estimation.

Marco Molinaro
Carnegie Mellon University
molinaro@cmu.edu

David Woodruff
IBM Almaden Research Center, USA
dpwoodru@us.ibm.com

Grigory Yaroslavtsev
Penn State University
grigory@cse.psu.edu

CP25

$(1 + \epsilon)$ -Approximation for Facility Location in Data Streams

We consider the *Euclidean facility location problem* with uniform opening cost. We obtain two main results:

- A $(1 + \epsilon)$ -approximation algorithm with running time $O(n \log^2 n \log \log n)$ for constant ϵ ,
- The first $(1 + \epsilon)$ -approximation algorithm for the *cost of the facility location problem for dynamic geometric data streams*, i.e., when the stream consists of insert and delete operations of points from a discrete space $\{1, \dots, \Delta\}^2$. The streaming algorithm uses $\left(\frac{\log \Delta}{\epsilon}\right)^{O(1)}$ space.

Morteza Monemizadeh
Institute of Theoretical Informatics
Karlsruhe Institute of Technology
m.monemizadeh@gmail.com

Artur Czumaj
Department of Computer Science and DIMAP
University of Warwick
A.Czumaj@warwick.ac.uk

Christiane Lammersen
School of Computing Science
Simon Fraser University
clammers@cs.sfu.ca

Christian Sohler

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

CP25

Eigenvalues of a Matrix in the Streaming Model

We study the question of estimating the eigenvalues of a matrix in the streaming model, addressing a question posed in [Muthukrishnan '05]. We show that the eigenvalue “heavy hitters” of a matrix can be computed in a single pass. In particular, we show that the ϕ -heavy hitters (in the ℓ_1 or ℓ_2 norms) can be estimated in space proportional to $1/\phi^2$. Such a dependence on ϕ is optimal. We also show how the same techniques may give an estimate of the residual error tail of a rank- k approximation of the matrix (in the Frobenius norm), in space proportional to k^2 . All our algorithms are linear and hence can support arbitrary updates to the matrix in the stream. In fact, what we show can be seen as a form of a bi-linear dimensionality reduction: if we multiply an input matrix with projection matrices on both sides, the resulting matrix preserves the top eigenvalues and the residual Frobenius norm.

Huy Nguyen
Princeton University
hnguyen@princeton.edu

Alexandr Andoni
MSR Silicon Valley
andoni@mit.edu

CP25

Space Efficient Streaming Algorithms for the Distance to Monotonicity and Asymmetric Edit Distance

We study streaming algorithms for approximating the length of the longest increasing sequence (LIS) of an array. We present a 1-pass algorithm that, for any $\delta > 0$, provides a $(1 + \delta)$ -multiplicative approximation to the *distance to monotonicity* of a data stream of length n (n minus the length of the LIS), and uses $O((\log^2 n)/\delta)$ space. Our algorithm is simple and just 3 lines of pseudocode. The previous best known approximation using polylogarithmic space was a multiplicative 2-factor. Our algorithm estimates the length of the LIS to within an additive δn for any $\delta > 0$ while previous algorithms could only achieve additive error $n(1/2 - o(1))$.

C. Seshadhri
Sandia National Labs
scomand@sandia.gov

Michael Saks
Rutgers University
saks@math.rutgers.edu

CP26

Efficient Algorithms for Computing the Triplet and Quartet Distance Between Trees of Arbitrary Degree

The triplet and quartet distances are distance measures to compare two rooted and two unrooted trees, respectively. The leaves of the two trees should have the same set of labels. The distances are defined by enumerating all subsets of three labels (triplets) and four labels (quartets), respec-

tively, and counting how often the induced topologies in the two input trees are different. In this paper we present efficient algorithms for computing these distances.

Gerth S. Brodal
MADALGO, Aarhus University
gerth@cs.au.dk

Rolf Fagerberg
University Of Southern Denmark
rolf@imada.sdu.dk

Thomas Mailund
Aarhus University
Bioinformatics Research Center
mailund@birc.au.dk

Christian N.S. Pedersen
Aarhus University
Department of Computer Science
cstorm@birc.au.dk

Andreas Sand
Aarhus University
Bioinformatics Research Center
asand@birc.au.dk

CP26
Discrete Convexity and Polynomial Solvability in Minimum 0-Extension Problems

The minimum 0-extension problem **0-Ext**[Γ] on a graph Γ includes a number of basic combinatorial optimization problems, such as minimum (s, t) -cut problem and multi-way cut problem. Karzanov proved that **0-Ext**[Γ] is NP-hard if Γ is not modular or not orientable (in a certain sense). In this paper, we prove the converse: if Γ is orientable and modular, then **0-Ext**[Γ] can be solved in polynomial time.

Hiroshi Hirai
Graduate School of Information Science and Technology
University of Tokyo
hirai@mist.i.u-tokyo.ac.jp

CP26
Improved Algorithms for Constructing Consensus Trees

A *consensus tree* is a single phylogenetic tree that summarizes the branching structure in a given set of conflicting phylogenetic trees. Many different types of consensus trees have been proposed in the literature; three of the most well-known and widely used ones are *the majority rule consensus tree*, *the loose consensus tree*, and *the greedy consensus tree*. This paper presents new deterministic algorithms for constructing them that are faster than all the previously known ones. Given k phylogenetic trees with n leaves each and with identical leaf label sets, our algorithms run in $O(nk \log k)$ time (majority rule consensus tree), $O(nk)$ time (loose consensus tree), and $O(n^2k)$ time (greedy consensus tree).

Jesper Jansson
Kyoto University
jj@kuicr.kyoto-u.ac.jp

Chuanqi Shen

Stanford University & Raffles Institution
scq1993@gmail.com

Wing-Kin Sung
National University of Singapore
ksung@comp.nus.edu.sg

CP26
Mimicking Networks and Succinct Representations of Terminal Cuts

Given a large edge-weighted network G with k terminal vertices, we wish to compress and store all the minimum terminal-cut values. One appealing methodology is to construct a *mimicking network*: a small network G' with the same k terminals, where the minimum terminal-cut values are the same as in G . We provide several new bounds on the mimicking network's size, which together narrow the previously known gap from doubly-exponential to only singly-exponential.

Inbal Rika, Robert Krauthgamer
Weizmann Institute of Science
inbal.rika@weizmann.ac.il,
robert.krauthgamer@weizmann.ac.il

CP26
Faster Deterministic Fully-Dynamic Graph Connectivity

We give new deterministic bounds for fully-dynamic graph connectivity. Our data structure supports updates (edge insertions/deletions) in $O(\log^2 n / \log \log n)$ amortized time and connectivity queries in $O(\log n / \log \log n)$ worst-case time, where n is the number of vertices of the graph. This improves the deterministic data structures of Holm, de Lichtenberg, and Thorup (STOC 1998, J.ACM 2001) and Thorup (STOC 2000) which both have $O(\log^2 n)$ amortized update time and $O(\log n / \log \log n)$ worst-case query time.

Christian Wulff-Nilsen
Department of Mathematics and Computer Science
University of Southern Denmark
koolooz@diku.dk

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

CP27
Testing K-Modal Distributions: Optimal Algorithms Via Reductions

We give highly efficient algorithms, and almost matching lower bounds, for a range of basic statistical problems that involve testing and estimating the L_1 (total variation) distance between two k -modal distributions p and q over the discrete domain $\{1, \dots, n\}$. As our main conceptual contribution, we introduce a new reduction-based approach for distribution-testing problems that lets us obtain our results in a unified way. Roughly speaking, this approach enables us to transform various distribution testing problems for k -modal distributions over $\{1, \dots, n\}$ to the corresponding distribution testing problems for unrestricted distributions over a much smaller domain $\{1, \dots, \ell\}$ where

$\ell = O(k \log n)$.

Constantinos Daskalakis
MIT
costis@csail.mit.edu

Ilias Diakonikolas
University of Edinburgh
ilias.d@ed.ac.uk

Rocco Servedio
Columbia University
ras2105@columbia.edu

Gregory Valiant
Microsoft
gregory.valiant@gmail.com

Paul Valiant
UC Berkeley
pvaliant@gmail.com

CP27

Optimal Listing of Cycles and st-Paths in Undirected Graphs

We present the first optimal algorithm for the classical problem of listing all the cycles in an undirected graph. The time complexity is proportional to the time taken to read the input graph plus the time to list the output, namely, the edges in each of the cycles. The algorithm uses a reduction to the problem of listing all the simple paths from a vertex s to a vertex t which we also solve optimally.

Etienne Birmele
Universite dEvr
etienne.birmele@genopole.cnrs.fr

Rui Ferreira
University of Pisa
Microsoft
ferreira@di.unipi.it

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

Andrea Marino
Universita di Firenze
andrea.marino@unifi.it

Nadia Pisanti
Universita di Pisa
pisanti@di.unipi.it

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

Gustavo Sacomoto
INRIA Grenoble Rhone-Alpes
Universite Lyon 1
sacomoto@gmail.com

CP27

Faster Walks in Graphs: A $\tilde{O}(n^2)$ Time-Space

Trade-off for Undirected s-t Connectivity

In this paper, we make use of the Metropolis-type walks due to Nonaka et al. (2010) to provide a faster solution to the s - t -connectivity problem in undirected graphs (USTCON). As our main result, we propose a family of randomized algorithms for USTCON which achieves a time-space product of $S \cdot T = \tilde{O}(n^2)$ in graphs with n nodes and m edges (where the \tilde{O} -notation disregards poly-logarithmic terms). This improves the previously best trade-off of $\tilde{O}(nm)$, due to Feige (1995). Our algorithm consists in deploying several short Metropolis-type walks, starting from landmark nodes distributed using the scheme of Broder et al. (1994) on a modified input graph. In particular, we obtain an algorithm running in time $\tilde{O}(n + m)$ which is, in general, more space-efficient than both BFS and DFS. We close the paper by showing how to fine-tune the Metropolis-type walk so as to match the performance parameters (e.g., average hitting time) of the unbiased random walk for any graph, while preserving a worst-case bound of $\tilde{O}(n^2)$ on cover time.

Adrian Kosowski
INRIA Bordeaux Sud-Ouest
kosowski@labri.fr

CP27

Low-Distortion Inference of Latent Similarities from a Multiplex Social Network

It is commonly assumed that individuals tend to be more similar to their friends than to strangers. Thus, we can view an observed social network as a noisy signal about the latent underlying "social space": the way in which individuals are (dis)similar. We present near-linear time algorithms which - under reasonably standard models of social network generation - can infer the similarities from the observed network with provable guarantees.

Ittai Abraham
Microsoft
ittai@microsoft.com

Shiri Chechik
Weizmann Institute of Science
shiri.chechik@weizmann.ac.il

David Kempe
USC
dkempe@usc.edu

Aleksandrs Slivkins
Microsoft Research
slivkins@microsoft.com

CP27

Segmentation of Trajectories on Non-Monotone Criteria

We study the trajectory segmentation problem, and suggest a general framework for solving it, based on the *start-stop diagram*: a 2-dimensional diagram that represents all valid and invalid segments of a given trajectory. This yields two subproblems: (i) computing the start-stop diagram, and (ii) finding the optimal segmentation for a given diagram. We show that (ii) is NP-hard in general. However, we identify properties of the start-stop diagram that make the problem tractable, and give polynomial-time algorithm

for this case.

Boris Aronov
Department of Computer Science and Engineering,
Polytechnic Institute of NYU
aronov@poly.edu

Anne Driemel, Marc van Kreveld, Maarten Löffler,
Frank Staals
Utrecht University
a.driemel@uu.nl, m.j.vankreveld@uu.nl, m.loffler@uu.nl,
f.staals@uu.nl