

IP1**Structural and Logical Approaches to the Graph Isomorphism Problem**

The question of whether there is a polynomial time algorithm deciding whether two graphs are isomorphic has been one of the best known open problems in theoretical computer science for more than forty years. Indeed, the graph isomorphism problem is one of the very few natural problems in NP that is neither known to be in P nor known to be NP-complete. The question is still wide open, but a number of deep partial results giving polynomial time algorithms for specific classes of graphs are known. Many of them have been obtained through a group theoretic approach that dominated the research on the graph isomorphism problem since the early 1980s. After an introductory survey on the graph isomorphism problem, in my talk I will focus on approaches to the graph isomorphism problem based on structural graph theory and connections between logical definability and certain combinatorial algorithms for the isomorphism problem. In particular, I will speak about two recent results: The first says that the Weisfeiler Lehman algorithm (a simple combinatorial algorithm) can be used to decide isomorphism on graph classes with excluded minors in polynomial time. The second says that isomorphism can be decided in polynomial time on graph classes with excluded topological subgraphs.

Martin Grohe

Humboldt-Universität zu Berlin
grohe@informatik.hu-berlin.de

IP2**Belief Propagation Algorithms: From Matching to Cancer Genomics**

We review belief propagation algorithms inspired by the study of phase transitions in combinatorial optimization problems. In particular, we present rigorous results on convergence of such algorithms for matching and associated bargaining problems. We also present a belief propagation algorithm for the prize-collecting Steiner tree problem, for which rigorous convergence results are not yet known. Finally, we show how this algorithm can be used to discover associations in cancer genomics, and to suggest drug therapies for cancer.

Jennifer Chayes

Microsoft Research
jchayes@microsoft.com

CP1**The Maximum Number of Faces of the Minkowski Sum of Two Convex Polytopes**

We derive tight bounds for the maximum number of k -faces of the Minkowski sum of two d -dimensional convex polytopes, as a function of their number of vertices. For even dimensions, the maximum values are attained when the two polytopes are cyclic d -polytopes with disjoint vertex sets. For odd dimensions, the maximum values are attained when the two polytopes are $\lfloor \frac{d}{2} \rfloor$ -neighborly d -polytopes, whose vertex sets are chosen appropriately from two distinct d -dimensional moment-like curves.

Menelaos I. Karavelas, Eleni Tzanaki

University of Crete
mkaravel@tem.uoc.gr, etzanaki@tem.uoc.gr

CP1**Computing All Maps into a Sphere**

We present an algorithm for computing $[X, Y]$, i.e., all homotopy classes of continuous maps $X \rightarrow Y$, where X, Y are topological spaces given as finite simplicial complexes, Y is $(d - 1)$ -connected for some $d \geq 2$ (for example, Y can be the d -dimensional sphere S^d), and $\dim X \leq 2d - 2$. These conditions on X, Y guarantee that $[X, Y]$ has a natural structure of a finitely generated Abelian group, and the algorithm finds generators and relations for it. We combine several tools and ideas from homotopy theory (such as *Postnikov systems*, *simplicial sets*, and *obstruction theory*) with algorithmic tools from effective algebraic topology (*objects with effective homology*). We hope that a further extension of the methods developed here will yield an algorithm for computing, in some cases of interest, the \mathbf{Z}_2 -index, which is a quantity playing a prominent role in Borsuk–Ulam style applications of topology in combinatorics and geometry, e.g., in topological lower bounds for the chromatic number of a graph. In a certain range of dimensions, deciding the embeddability of a simplicial complex into \mathbf{R}^d also amounts to a \mathbf{Z}_2 -index computation. This is the main motivation of our work. We believe that investigating the computational complexity of questions in homotopy theory and similar areas presents a fascinating research area, and we hope that our work may help bridge the cultural gap between algebraic topology and theoretical computer science.

Martin Čadek

Department of Mathematics and Statistics
Masaryk University
cadek@math.muni.cz

Marek Krčál

ITI Charles University Prague
krcal@kam.mff.cuni.cz

Jiří Matoušek, Francis Sergeraert

Institute for Theoretical Computer Science - ITI
Charles University
krcal@kam.mff.cuni.cz, matousek@kam.mff.cuni.cz

Lukáš Vokřínek

Department of Mathematics and Statistics
Masaryk University
koren@math.muni.cz

Uli Wagner

Institute of Theoretical Computer Science
ETH Zurich
uli@inf.ethz.ch

CP1**Dimension Reduction for Finite Trees in L_1**

We show that every n -point tree metric admits a $(1 + \varepsilon)$ -embedding into $\ell_1^{C(\varepsilon) \log n}$, for every $\varepsilon > 0$, where $C(\varepsilon) \leq O\left(\left(\frac{1}{\varepsilon}\right)^4 \log \frac{1}{\varepsilon}\right)$. This matches the natural volume lower bound up to a factor depending only on ε . Previously, it was unknown whether even complete binary trees on n nodes could be embedded in $\ell_1^{O(\log n)}$ with $O(1)$ distortion. For complete d -ary trees, our construction achieves $C(\varepsilon) \leq O\left(\frac{1}{\varepsilon^2}\right)$.

James R. Lee, Mohammad Moharrami

University of Washington

jrl@cs.washington.edu, moharrami@cs.washington.edu

Arnaud de Mesmay
Ecole Normale Supérieure
ademesmay@gmail.com

CP1

On Multiplicative λ -Approximations and Some Geometric Applications

Let \mathcal{F} be a set system over an underlying finite set X , and let μ be a non-negative measure over X . A measure μ^* on X is called a *multiplicative λ -approximation* of μ on (\mathcal{F}, X) if for every $S \in \mathcal{F}$ it holds that $a \cdot \mu(S) \leq \mu^*(S) \leq b \cdot \mu(S)$, and $b/a = \lambda \geq 1$. We show that $\text{trk}(\mathcal{F})$, defined as the size of the largest square lower-triangular submatrix (with 1's on the diagonal) of the incidence matrix of (\mathcal{F}, X) , determines the smallest support that a multiplicative approximation of μ may have, in the worst case. For any μ on X and $0 < \epsilon < 1$, we construct μ^* that $\frac{1+\epsilon}{1-\epsilon}$ -approximates μ on (X, \mathcal{F}) , and has support of size $O(\text{trk}(\mathcal{F})^2 \log(\text{trk}(\mathcal{F}))/\text{poly}(\epsilon))$. We also present two alternative constructions which in some cases improve upon this bound. On the other hand, we show that $\text{trk}(\mathcal{F})$ is a lower bound on the support size for some measure.

Ilan I. Newman, Yuri Rabinovich
Haifa University
ilan@cs.haifa.ac.il, yuri@cs.haifa.ac.il

CP1

Polytope Approximation and the Mahler Volume

The problem of approximating convex bodies by polytopes is an important and well studied problem. Given a convex body K in R^d , the objective is to minimize the number of vertices (alternatively, the number of facets) of an approximating polytope for a given Hausdorff error ϵ . Results to date have been of two types. The first type assumes that K is smooth, and bounds hold in the limit as ϵ tends to zero. The second type requires no such assumptions. The latter type includes the well known results of Dudley (1974) and Bronshteyn and Ivanov (1976), which show that in spaces of fixed dimension, $O((\text{diam}(K)/\epsilon)^{(d-1)/2})$ vertices (alt., facets) suffice. Our results are of this latter type. In our first result, under the assumption that the width of the body in any direction is at least ϵ , we strengthen the above bound to $\tilde{O}(\sqrt{\text{diam}(K)}/\epsilon^{(d-1)/2})$. This is never worse than the previous bound (by more than logarithmic factors) and may be significantly better for skinny bodies. Our analysis exploits an interesting analogy with a classical concept from the theory of convexity, called the Mahler volume. This is a dimensionless quantity that involves the product of the volumes of a convex body and its polar dual. In our second result, we apply the same machinery to improve upon the best known bounds for answering ϵ -approximate polytope membership queries. Given a convex polytope P defined as the intersection of halfspaces, such a query determines whether a query point q lies inside or outside P , but may return either answer if q 's distance from P 's boundary is at most ϵ . We show that, without increasing storage, it is possible to dramatically reduce the best known search times for ϵ -approximate polytope membership. This further implies improvements to the best known search times for approximate nearest neighbor searching in spaces of fixed dimension.

Sunil Arya

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

Guilherme D. da Fonseca
Universidade Federal do Estado do Rio de Janeiro
(UNIRIO)
fonseca@uniriotec.br

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

CP2

Kernelization of Packing Problems

Kernelization algorithms are polynomial-time reductions from a problem to itself that guarantee their output to have a size not exceeding some bound. We study the kernelization of set and graph packing problems. For example, d -SET MATCHING for $d \geq 3$ is the problem of finding a matching of size at least k in a given d -uniform hypergraph and has kernels with $O(k^d)$ edges. We prove that this problem does not have kernels with $O(k^{d-\epsilon})$ edges for any $\epsilon > 0$ unless coNP is in NP/poly.

Holger Dell
University of WisconsinMadison
holger@cs.wisc.edu

Dániel Marx
Institut für Informatik, Humboldt-Universität zu Berlin
and
Hungarian Academy of Sciences (MTA SZTAKI),
Budapest, Hungary
dmarx@cs.bme.hu

CP2

Co-Nondeterminism in Compositions: A Kernelization Lower Bound for a Ramsey-Type Problem

In this work we present the first example of how co-nondeterminism can help to make a composition algorithm. We study the existence of polynomial kernels for a Ramsey-type problem: Given a graph G and an integer k , does G contain an independent set or a clique of size at least k ? We provide a co-nondeterministic composition based on embedding t instances into a single host graph H ; co-nondeterminism is used to find an appropriate host graph H .

Stefan Kratsch
Utrecht University
s.kratsch@uu.nl

CP2

Linear Kernels for (Connected) Dominating Set on H -Minor-Free Graphs

We give the first linear kernels for DOMINATING SET and CONNECTED DOMINATING SET problems on graphs excluding a fixed graph H as a minor. In other words, we give polynomial time algorithms that, for a given H -minor free graph G and positive integer k , output an H -minor free graph G' on $O(k)$ vertices such that G has a (connected) dominating set of size k if and only if G' has. Prior to our work, the only polynomial kernel for DOMINATING SET on graphs excluding a fixed graph H as a minor was due to Alon and Gutner [ECCC 2008, IWPEC 2009] and to

Philip, Raman, and Sikdar [ESA 2009] but the size of their kernel is $k^{c(H)}$, where $c(H)$ is a constant depending on the size of H . Alon and Gutner asked explicitly, whether one can obtain a linear kernel for DOMINATING SET on H -minor free graphs. We answer this question in affirmative. For CONNECTED DOMINATING SET no polynomial kernel on H -minor free graphs was known prior to our work. Our results are based on a novel generic reduction rule producing an equivalent instance of the problem with treewidth (\sqrt{k}) . The application of this rule in a divide-and-conquer fashion together with protrusion techniques brings us to linear kernels. As a byproduct of our results we obtain the first subexponential time algorithms for CONNECTED DOMINATING SET, a deterministic algorithm solving the problem on an n -vertex H -minor free graph in time $2^{O(\sqrt{k} \log k)} + n^{O(1)}$ and a Monte Carlo algorithm of running time $2^{O(\sqrt{k})} + n^{O(1)}$. For DOMINATING SET our results implies a significant simplification and refinement of a $2^{O(\sqrt{k})} n^{O(1)}$ algorithm on H minor free graphs due to Demaine et al. [SODA 2003, J. ACM 2005].

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

Daniel Lokshtanov
 UCSD
 daniello@ii.uib.no

Saket Saurabh Last Name
 IMSc
 saket@imsc.res.in

Dimitrios Thilikos
 University of Athens
 sedthilk@math.uoa.gr

CP2

Compression Via Matroids: A Randomized Polynomial Kernel for Odd Cycle Transversal

We give a (randomized) polynomial kernelization for the Odd Cycle Transversal problem (OCT), using a novel kernelization approach based on matroid theory. We encode all relevant information about a problem instance into a matroid with a representation of size polynomial in k . For OCT, we apply this to the iterative compression algorithm of Reed, Smith, and Vetta (ORL 2004), which, combined with an $O(\sqrt{k})$ -approximation, gives a polynomial kernel. This settles an important open problem in parameterized complexity.

Stefan Kratsch
 Utrecht University
 s.kratsch@uu.nl

Magnus Wahlstrom
 Max-Planck-Institute for Informatics
 wahl@mpi-inf.mpg.de

CP2

Weak Compositions and Their Applications to Polynomial Lower Bounds for Kernelization

We use the notion of weak composition to obtain polynomial kernelization lower-bounds for several natural pa-

parameterized problems. Let $d \geq 2$ be some constant and let $L_1, L_2 \subseteq \{0, 1\}^* \times N$ be two parameterized problems where the unparameterized version of L_1 is NP-hard. Assuming $coNP \not\subseteq NP/poly$, our framework essentially states that composing t L_1 -instances each with parameter k , to an L_2 -instance with parameter $k' \leq t^{1/d} k^{O(1)}$, implies that L_2 does not have a kernel of size $O(k^{d-\epsilon})$ for any $\epsilon > 0$. We show two examples of weak composition and derive polynomial kernelization lower bounds for d -Bipartite Regular Perfect Code and d -Dimensional-Matching, parameterized by the solution size k . By reduction using linear parameter transformation, we then derive the following lower-bounds for kernel sizes when the parameter is the solution size k (assuming $coNP \not\subseteq NP/poly$): - d -Set Packing, d -Set Cover, d -Exact Set Cover, Hitting Set with d -Bounded Occurrences, and Exact Hitting Set with d -Bounded Occurrences have no kernels of size $O(k^{d-3-\epsilon})$ for any $\epsilon > 0$. - K_d Packing and Induced $K_{s,d}$ Packing for any constant $s \geq 1$ have no kernels of size $O(k^{d-4-\epsilon})$ for any $\epsilon > 0$. - d -Red-Blue Dominating Set and d -Steiner Tree have no kernels of sizes $O(k^{d-3-\epsilon})$ and $O(k^{d-4-\epsilon})$, respectively, for any $\epsilon > 0$.

XI Wu
 University of Wisconsin-Madison
 wu.andrew.xi@gmail.com

Danny Hermelin
 Max Planck Institute for Informatics
 hermelin@mpi-inf.mpg.de

CP3

A Matroid Approach to Stable Matchings with Lower Quotas

In SODA'10, Huang introduced the laminar classified stable matching problem (**LCSM** for short) that is motivated by academic hiring. This problem is an extension of the well-known hospitals/residents problem in which a hospital has laminar classes of residents and it sets lower and upper bounds on the number of residents that it would hire in that class. Against the intuition that stable matching problems with lower quotas are difficult in general, Huang proved that this problem can be solved in polynomial time. In this paper, we propose a matroid-based approach to this problem and we obtain the following results. (i) We solve a generalization of the **LCSM** problem. (ii) We exhibit a polyhedral description for stable assignments of the **LCSM** problem, which gives a positive answer to Huang's question. (iii) We prove that the set of stable assignments of the **LCSM** problem has a lattice structure similarly to the ordinary stable matching model.

Tamas Fleiner
 Budapest University of Technology and Economics
 fleiner@cs.bme.hu

Naoyuki Kamiyama
 Chuo University
 nkamiyama@gmail.com

CP3

Popularity Vs Maximum Cardinality in the Stable Marriage Setting

We consider the problem of computing matchings with small unpopularity and large size in a stable marriage instance $G = (A \cup B, E)$ with incomplete lists and strict

preferences. For any integer $k \geq 2$, we extend the Gale-Shapley stable matching algorithm to k layers, to show that a matching with unpopularity factor at most $k - 1$ and size at least $\frac{k}{k+1}$ (size of a maximum matching in G) always exists. Such a matching can be computed in $O(km)$ time, where $|E| = m$.

Telikepalli Kavitha
Tata Institute of Fundamental Research, Mumbai, India
kavitha.telikepalli@gmail.com

CP3

On the (In)security of Hash-Based Oblivious Ram and a New Balancing Scheme

A secure Oblivious RAM simulation allows for a client, with small (e.g., constant size) protected memory, to hide not only the data but also the sequence of locations it accesses (both reads and writes) in the unprotected memory of size n . Our main results are as follows:

- We analyze several schemes from the literature, observing a repeated design flaw that leaks information on the memory access pattern.
- On the positive side, we present a new secure oblivious RAM scheme, extending a recent scheme by Goodrich and Mitzenmacher. Our scheme uses only $O(1)$ local memory, and its (amortized) overhead is $O(\log^2 n / \log \log n)$, outperforming the previously-best $O(\log^2 n)$ overhead (among schemes where the client only uses $O(1)$ additional local memory).
- We also present a transformation of our scheme above (whose amortized overhead is $O(\log^2 n / \log \log n)$) into a scheme with worst-case overhead of $O(\log^2 n / \log \log n)$.

Steve Lu
Stealth Software Technologies, Inc.
stevelu8@gmail.com

Eyal Kushilevitz
Technion
eyalk@cs.technion.ac.il

Rafail Ostrovsky
UCLA
rafail@cs.ucla.edu

CP3

Privacy-Preserving Group Data Access Via Stateless Oblivious Ram Simulation

Motivated by cloud computing applications, we study the problem of providing privacy-preserving access to an outsourced honest-but-curious data repository for a group of trusted users. We show how to achieve efficient privacy-preserving data access using a combination of probabilistic encryption, which directly hides data values, and stateless oblivious RAM simulation, which hides the pattern of data accesses. We give a method with $O(\log n)$ amortized access overhead for simulating a RAM algorithm that has a memory of size n , using a scheme that is data-oblivious with very high probability.

Olga Ohrimenko
Brown University
olya@cs.brown.edu

Michael Goodrich

University of California, Irvine
goodrich@ics.uci.edu

Michael Mitzenmacher
Harvard University
michaelm@eecs.harvard.edu

Roberto Tamassia
Brown University
rt@cs.brown.edu

CP3

Private Data Release Via Learning Thresholds

This work considers computationally efficient privacy-preserving data release. Given a set of statistical queries on a database containing sensitive information about individual participants, we want to release approximate answers to queries while guaranteeing differential privacy. Our main contribution is a computationally efficient reduction from differentially private data release for a class of counting queries, to learning thresholded sums of predicates from a related class. We instantiate this general reduction with a variety of algorithms for learning thresholds.

Moritz Hardt
IBM Research
mhardt@cs.princeton.edu

Guy Rothblum
Microsoft Research
rothblum@alum.mit.edu

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

CP4

Near Linear Time $(1 + \epsilon)$ -Approximation for Restricted Shortest Paths in Undirected Graphs

In the *restricted* shortest path problem, each edge in the graph has two weights: a cost-weight, and a delay-weight. The goal is to find the shortest path from s - t that minimizes total cost while having delay-length less than some given threshold T . The problem is NP-complete, but a sequence of papers culminated in an $O(mn)$ time, $(1 + \epsilon)$ approximation. We break through this barrier for undirected graphs, achieving a $(1 + \epsilon)$ approximation in close to linear time.

Aaron Bernstein
Columbia University
bernstei@gmail.com

CP4

Physarum Can Compute Shortest Paths

Physarum Polycephalum is a slime mold that apparently is able to solve shortest path problems. A mathematical model has been proposed by biologists to describe the feedback mechanism used by the slime mold to adapt its tubular channels while foraging two food sources s_0 and s_1 . We prove that, under this model, the mass of the mold will eventually converge to the shortest s_0 - s_1 path of the network that the mold lies on, independently of the structure of the network or of the initial mass distribution. This matches the experimental observations by the biologists and can be seen as an example of a “natural algo-

rithm', that is, an algorithm developed by evolution over millions of years.

Vincenzo Bonifaci
IASI-CNR, Italy
bonifaci@mpi-inf.mpg.de

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

Girish Varma
Tata Institute for Fundamental Research
Mumbai, India
girish@tcs.tifr.res.in

CP4

Single Source Distance Oracle for Planar Digraphs Avoiding a Failed Node Or Link

Let $G = (V, E)$ be a planar digraph on $n = |V|$ vertices with weights on the edges, and let $s \in V$ be any source vertex. We show that G can be preprocessed in $O(n \text{ polylog } n)$ time to build a data structure of $O(n \text{ polylog } n)$ size which can answer the following query in $O(\log n)$ time for any $u, v \in V$:

report distance from s to v in the graph $G \setminus \{u\}$

We also address all-pairs version of this problem and present a data structure with $O(n\sqrt{n})$ preprocessing time and space which guarantees $O(\sqrt{n} \text{ polylog } n)$ query time.

Surender Baswana, Utkarsh Lath, Anuradha Mehta
IIT Kanpur
sbaswana@gmail.com, utkarsh.lath@gmail.com,
anuradha.s.mehta@gmail.com

CP4

Exact Distance Oracles for Planar Graphs

We present exact distance oracles (data structures that answer node-to-node distance queries) in planar graphs. For any directed planar graph on n nodes, given any space allocation $S \in [n, n^2]$, we construct in $\tilde{O}(S)$ time a data structure of size $O(S)$ that answers queries in $\tilde{O}(n/\sqrt{S})$ time. This was known only for $S = n$ or $S \geq n^{4/3}$. For planar graphs with tree-width $w = o(\sqrt{n})$, we obtain a $\tilde{O}(n)$ -space distance oracle with query time $\tilde{O}(w)$. Building on this, we provide an oracle with query time proportional to the distance between the query nodes. Comparable query performance had been observed experimentally but could not be explained theoretically.

Shay Mozes
Brown University
shay@cs.brown.edu

Christian Sommer
MIT
csom@mit.edu

CP4

Approximate Distance Oracles with Improved Preprocessing Time

Given a weighted undirected graph G with m edges and n vertices and given a $k \geq 1$, we show that for some con-

stant c , a $(2k - 1)$ -approximate distance oracle for G of size $O(kn^{1+1/k})$ and query time $O(k)$ can be constructed in $O(\sqrt{km} + kn^{1+c/\sqrt{k}})$ time. We break the quadratic preprocessing time bound of Baswana and Kavitha for all $k \geq 6$ and improve the $O(kmn^{1/k})$ time bound of Thorup and Zwick except for very sparse graphs and small k .

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

CP5

The Condensation Transition in Random Hypergraph 2-Coloring

For many random CSPs the best current estimates of the threshold for the existence of solutions are based on the first and the second moment method. However, in most cases these techniques do not yield matching upper and lower bounds. Sophisticated but non-rigorous arguments from statistical mechanics have ascribed this discrepancy to the existence of a phase transition called condensation (Krzakala, Montanari, Ricci-Tersenghi, Semerjian, Zdeborova: PNAS 2007). In this paper we prove for the first time that a condensation transition exists in a natural random CSP, namely in random hypergraph 2-coloring.

Amin Coja-Oghlan
University of Warwick
a.coja-oghlan@warwick.ac.uk

Lenka Zdeborova
CEA Saclay
lenka.zdeborova@gmail.com

CP5

A Simple Algorithm for Random Colouring $G(n, D/n)$ Using $(2 + \epsilon)d$ Colours.

We study the problem of approximate random k -colouring $G(n, d/n)$ with fixed d . On input $G(n, d/n)$ and $k \geq (2 + \epsilon)d$, the algorithm returns w.h.p. a k -colouring distributed within total variation distance $n^{-\Omega(1)}$ from the uniform. Our algorithm is neither MCMC nor message-passing based. It is novel, based on a rather simple recursion. The lower bound on k for our algorithm to be polynomial-time is *significantly* smaller than those of previous algorithm.

Charilaos Efthymiou
University of Warwick - DIMAP
c.efthymiou@warwick.ac.uk

CP5

A New Approach to the Orientation of Random Hypergraphs

A h -uniform hypergraph $H = (V, E)$ is called (ℓ, k) -orientable if there exists an assignment of each hyperedge $e \in E$ to exactly ℓ of its vertices $v \in e$ such that no vertex is assigned more than k hyperedges. Let $H_{n,m,h}$ be a hypergraph, drawn uniformly at random from the set of all h -uniform hypergraphs with n vertices and m edges. In this paper, we determine the threshold of the existence of a (ℓ, k) -orientation of $H_{n,m,h}$ for $k \geq 1$ and $h > \ell \geq 1$, extending recent results motivated by applications such as cuckoo hashing or load balancing with guaranteed maxi-

mum load. Our proof combines the local weak convergence of sparse graphs and a careful analysis of a Gibbs measure on spanning subgraphs with degree constraints. It allows us to deal with a much broader class than the uniform hypergraphs.

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

CP5

The Max-Cut of Sparse Random Graphs

MAX-CUT is the problem of partitioning the vertex set of graph into two parts A and B so that the number of edges from A to B is maximized. We consider the MAX-CUT problem on random connected graphs and on Erdős-Rényi random graphs. More specifically, we consider the *distance from bipartiteness* of a graph, the minimum number of edge deletions needed to turn it into a bipartite graph. If we denote this distance $\text{DISTBIP}(G)$, the size of the MAX-CUT of a graph $G = (V, E)$ is clearly given by $|E| - \text{DISTBIP}(G)$. Fix $\epsilon > 0$. For random connected graphs, we prove that asymptotically almost surely (a.a.s) $\text{DISTBIP}(C(n, m)) \sim \frac{m-n}{4}$ whenever $m = n + O(n^{1-\epsilon})$. For sparse random graphs $G(n, m = \frac{n}{2} + \frac{\mu n^{2/3}}{2})$ we show that $\text{DISTBIP}(G(n, m))$ is a.a.s about $\frac{(2m-n)^3}{6n^2} + \frac{1}{12} \log n - \frac{1}{4} \log \mu$ for any $1 \ll \mu \leq O(n^{1/3-\epsilon})$. Thus, our results on $G(n, m)$ contrast with those on dense random graphs ($G(n, m = (1 + \epsilon)n/2)$), where a.a.s $\text{DISTBIP}(G(n, m)) \sim \frac{m}{2}$.

Conrado Martínez
Dept. Software
Universitat Politècnica de Catalunya
conrado@lsi.upc.edu

Hervé Daudé
LATP - UMR CNRS 6622
Université de Provence, France
daude@cmi.univ-mrs.fr

Vonjy Rasendrasana
LIPN - UMR CNRS 7030
Université de Paris Nord (Paris 13), France
vonjy@lipn.univ-paris13.fr

Vlady Ravelomanana
LIAFA - UMR CNRS 7089
Université Denis Diderot (Paris 7), France
vlad@liafa.jussieu.fr

CP5

The Maximum Degree of Random Planar Graphs

Let P_n denote a graph drawn uniformly at random from the class of all simple planar graphs with n vertices. We show that the maximum degree of a vertex in P_n is with probability $1 - o(1)$ asymptotically equal to $c \log n$, where $c \approx 2.529$ is determined explicitly. A similar result is also true for random 2-connected planar graphs. Our analysis combines two orthogonal methods that complement each other. First, in order to obtain the upper bound, we resort to *generating functions* and *analytic combinatorics*. This allows us to obtain precise asymptotic estimates for the expected number of vertices of any given degree in. On the other hand, for the lower bound we use *Boltzmann sampling*. In particular, by tracing the execution of an

adequate algorithm that generates a random planar graph, we are able to explicitly find vertices of sufficiently high degree.

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

Omer Gimenez
Google
omer.gimenez@upc.edu

Marc Noy
UPC Barcelona
marc.noy@upc.edu

Konstantinos Panagiotou
Departement of Computer Science
ETH Zurich
kpanagio@mpi-inf.mpg.de

Konstantinos Panagiotou
Max-Planck-Institute for Informatics
kpanagio@mpi-inf.mpg.de

Angelika Steger
Institute of Theoretical Computer Science
ETH, Zürich
steger@inf.ethz.ch

CP6

Packing Anchored Rectangles

Let S be a set of n points in the unit square $[0, 1]^2$, one of which is the origin. We construct n pairwise interior-disjoint axis-aligned empty rectangles such that the lower left corner of each rectangle is a point in S , and the rectangles jointly cover at least a positive constant area (about 0.09). This is a first step towards the solution of a long-standing conjecture that the rectangles in such a packing can jointly cover an area of at least $1/2$.

Adrian Dumitrescu
Department of Engineering and Computer Science
University of Wisconsin-Milwaukee
dumitres@uwm.edu

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

CP6

Jaywalking Your Dog – Computing the Frechet Distance with Shortcuts

The similarity of two polygonal curves can be measured using the Frechet distance. We introduce the notion of a more robust Frechet distance, where one is allowed to shortcut between vertices of one of the curves. This is a natural approach for handling noise, in particular batched outliers. We compute a constant factor approximation to the minimum Frechet distance over all possible such shortcuts. Our algorithm runs in $O(c^2 kn \log^3 n)$ time if one is allowed to take at most k shortcuts and the input curves are c -packed. For the case where the number of shortcuts is unrestricted, we describe an algorithm which runs in $O(c^2 n \log^3 n)$ time. To facilitate the new algorithm we develop several new data-structures, which we believe to be

of independent interest: (i) for range reporting on a curve, and (ii) for preprocessing a curve to answer queries for the Frechet distance between a subcurve and a line segment.

Sariel Har-Peled

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

Anne Driemel

Department of Information and Computing Sciences;
Utrecht
anne@cs.uu.nl

CP6

Computing the Distance Between Piecewise-Linear Bivariate Functions

We consider the problem of computing the distance between two piecewise linear bivariate functions f and g defined over a common domain M . We focus on the distance induced by the L_2 -norm, that is $\|f - g\|_2 = \sqrt{\int_M (f - g)^2}$. If f is defined by linear interpolation over a triangulation of M with n triangles, while g is defined over another such triangulation, the obvious naïve algorithm requires $\Theta(n^2)$ arithmetic operations to compute this distance. We show that it is possible to compute it in $\mathcal{O}(n \log^4 n)$ arithmetic operations, by reducing the problem to multi-point evaluation of a certain type of polynomials. We also present an application to terrain matching.

Guillaume Moroz

Inria Nancy - Grand Est
guillaume.moroz@inria.fr

Boris Aronov

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

CP6

Submatrix Maximum Queries in Monge Matrices and Monge Partial Matrices, and Their Applications

We describe how to process an $n \times n$ Monge matrix M in $\mathcal{O}(n \log^2 n)$ time into a data structure of size $\mathcal{O}(n \log n)$ that outputs the maximum element in any contiguous submatrix of M in $\mathcal{O}(\log^2 n)$ time. This data structure extends to Monge partial matrices with a slight penalty in the bounds. We use this data structure to (1) preprocess a set of points so that we can find the largest-area empty rectangle containing a query point, and to (2) report distances in a planar graph with dynamically changing real weights.

Haim Kaplan

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

Shay Mozes

Brown University
shay@cs.brown.edu

Yahav Nussbaum

Tel-Aviv University
yahav.nussbaum@cs.tau.ac.il

Micha Sharir

Tel Aviv University
michas@post.tau.ac.il

CP6

Algorithms for the Transportation Problem in Geometric Settings

For $A, B \subset \mathbb{R}^d$, $|A| + |B| = n$, let $a \in A$ have a demand $d_a \in \mathbb{Z}^+$ and $b \in B$ have a supply $s_b \in \mathbb{Z}^+$, $\sum_{a \in A} d_a = \sum_{b \in B} s_b = U$ and let $d(\cdot, \cdot)$ be a distance function. Suppose the diameter of $A \cup B$ is Δ under $d(\cdot, \cdot)$, and $\epsilon > 0$ is a parameter. We present an algorithm that in $\mathcal{O}((n\sqrt{U} \log^2 n + U \log U)\Phi(n) \log(\Delta U/\epsilon))$ time computes a solution to the transportation problem on A, B which is within an additive error ϵ from the optimal solution. Here $\Phi(n)$ is the query and update time of a dynamic weighted nearest-neighbor data structure under distance function $d(\cdot, \cdot)$. Note that the $(1/\epsilon)$ appears only in the log term. As a consequence, we obtain improved exact and approximation algorithms for computing bipartite matchings under L_p and RMS distance. For point sets, $A, B \subset [\Delta]^d$, for the L_p norm and for $0 < \alpha, \beta < 1$, we present a randomized dynamic data structure that maintains a partial solution to the transportation problem under insertions and deletions of points, in which at least $(1 - \alpha)U$ of the demands are satisfied and whose cost is within $(1 + \beta)$ of that of the optimal (complete) solution to the transportation problem with high probability. The insertion, deletion and update times are $\mathcal{O}(\text{poly}(\log(n\Delta)/\alpha\beta))$, provided $U = n^{\mathcal{O}(1)}$.

R Sharathkumar, Pankaj K. Agarwal

Department of Computer Science, Duke University
sharath@cs.duke.edu, pankaj@cs.duke.edu

CP7

Linear Index Coding Via Semidefinite Programming

We show a polynomial time algorithm that, given an n vertex graph G with minrank k , finds a linear index code for G of length $\tilde{\mathcal{O}}(n^{f(k)})$, where $f(k)$ depends only on k . Our algorithm employs a semidefinite program (SDP) introduced by Karger, Motwani and Sudan for graph coloring. At the heart of our analysis we show an exact expression for the maximum possible value of the Lovász ϑ -function of a graph with minrank k .

Ishay Haviv, Eden Chlamtac

Tel Aviv University
havivish@post.tau.ac.il, chlamtac@cs.princeton.edu

CP7

The Entropy Rounding Method in Approximation Algorithms

Let A be a matrix and x be a fractional vector, say an LP solution to some discrete optimization problem. We give a new randomized rounding procedure for obtaining an integral vector y such that $Ay \approx Ax$, provided that A has bounded Δ -approximate entropy. This property means that for uniformly chosen random signs $\chi(j) \in \{1\}$ on any subset of the columns, the outcome $A\chi$ can be approximately described using at most $\frac{m}{5}$ bits in expectation (with m being the number of selected columns). To achieve this result, we modify well-known techniques from the field of discrepancy theory, especially we rely on Beck's entropy method. We demonstrate the versatility of our procedure

by providing an $OPT + O(\log^2 OPT)$ approximation for Bin Packing With Rejection and the first AFPTAS for the Train Delivery problem.

Thomas Rothvoss

Massachusetts Institute of Technology
Department of Mathematics
rothvoss@math.mit.edu

CP7

Concentration and Moment Inequalities for Polynomials of Independent Random Variables

In this work we propose new randomized rounding algorithms for matroid intersection and matroid base polytopes. We prove concentration inequalities for polynomial objective functions and constraints that have numerous applications and can be used in approximation algorithms for Minimum Quadratic Spanning Tree, Unrelated Parallel Machines Scheduling and Scheduling with Time Windows and Nonlinear objectives. We also show applications related to Constraint Satisfaction and Dense Polynomial optimization.

Konstantin Makarychev
IBM
konstantin@us.ibm.com

Warren Schudy
IBM TJ Watson
ws@cs.brown.edu

Maxim Sviridenko
IBM T.J. Watson Research Center
sviri@us.ibm.com

CP7

Concentration Inequalities for Nonlinear Matroid Intersection

In this work we design a general method for proving moment inequalities for polynomials of independent random variables. Our method works for a wide range of random variables including Gaussian, Boolean, exponential, Poisson and many others. We apply our method to derive general concentration inequalities for polynomials of independent random variables. We show that our method implies concentration inequalities for some previously open problems, e.g. permanent of random symmetric matrices. We show that our concentration inequality is stronger than the wellknown concentration inequality due to Kim and Vu [29]. The main advantage of our method in comparison with the existing ones is a wide range of random variables we can handle and bounds for previously intractable regimes of high degree polynomials and small expectations. On the negative side we show that even for boolean random variables each term in our concentration inequality is tight.

Warren Schudy
IBM TJ Watson
ws@cs.brown.edu

Maxim Sviridenko
IBM T.J. Watson Research Center
sviri@us.ibm.com

CP7

Approximating Csps with Global Cardinality Constraints Using Sdp Hierarchies

This work is concerned with approximating constraint satisfaction problems (CSPs) with an additional global cardinality constraints. For example, the Max-Bisection problem is a variant of Max-Cut with an additional global constraint that each side of the cut has exactly half the vertices, i.e., $|S| = |V|/2$. Several other natural optimization problems like Min-Bisection and approximating Graph Expansion can be formulated as CSPs with global constraints. In this work, we formulate a general approach towards approximating CSPs with global constraints using SDP hierarchies. To demonstrate the approach we present the following results:

- Using the Lasserre hierarchy, we present an algorithm that runs in time $O(n^{\text{poly}(1/\epsilon)})$ that given an instance of Max-Bisection with value $1 - \epsilon$, finds a bisection with value $1 - O(\sqrt{\epsilon})$. This approximation is near-optimal (up to constant factors in $O(\cdot)$) under the Unique Games Conjecture.
- By a computer-assisted proof, we show that the same algorithm also achieves a 0.85-approximation for Max-Bisection, improving on the previous bound of 0.70 (note that it is hard to approximate better than a 0.878 factor). The same algorithm also yields a 0.92-approximation for Max-2SAT with cardinality constraints.
- For every CSP with a global cardinality constraints, we present a generic conversion from integrality gap instances for the Lasserre hierarchy to a *dictatorship test* whose soundness is at most integrality gap. Dictatorship testing gadgets are central to hardness results for CSPs, and a generic conversion of the above nature lies at the core of the tight Unique Games based hardness result for CSPs.

Ning Tan, Prasad Raghavendra
Georgia Tech
ntan3@math.gatech.edu, raghavendra@cc.gatech.edu

CP7

Polynomial Integrality Gaps for Strong Sdp Relaxations of Densest K-Subgraph

We give evidence for the hardness of approximating the Densest k -Subgraph (DkS) problem within polynomial factors. Specifically, we expose the limitations of strong semidefinite programs from SDP hierarchies in solving DkS. Our results include:

- A lower bound of $\Omega(n^{1/4}/\log^3 n)$ on the integrality gap for $\Omega(\frac{\log n}{\log \log n})$ rounds of the Sherali-Adams relaxation for DkS.
- For every $\epsilon > 0$, a lower bound of $n^{\Omega_\epsilon(1)}$ on the integrality gap of $n^{1-\epsilon}$ rounds of the Lasserre SDP relaxation for DkS.

In the absence of inapproximability results for DkS, our results show that beating a factor of $n^{\Omega(1)}$ is a barrier for even the most powerful SDPs, and in fact even beating the best known $n^{1/4}$ factor is a barrier for current techniques.

Yuan Zhou
Carnegie Mellon University
yuanzhou@cs.cmu.edu

Aditya Bhaskara, Moses Charikar
Princeton University
bhaskara@cs.princeton.edu, moses@cs.princeton.edu

Venkatesan Guruswami
Carnegie Mellon University
guruswami@cmu.edu

Aravindan Vijayaraghavan
Princeton University
aravindv@cs.princeton.edu

CP8

Analyzing Graph Structure Via Linear Measurements

We initiate the study of *graph sketching*, i.e., algorithms that use a limited number of linear measurements of a graph to determine the properties of the graph. While a graph on n nodes is essentially $O(n^2)$ -dimensional, we show the existence of a distribution over random projections into d -dimensional “sketch” space ($d \ll n^2$) such that the relevant properties of the original graph can be inferred from the sketch with high probability. Specifically, we show that:

1. $d = O(n \cdot n)$ suffices to evaluate properties including connectivity, k -connectivity, bipartiteness, and to return any constant approximation of the weight of the minimum spanning tree.
2. $d = O(n^{1+\gamma})$ suffices to compute graph sparsifiers, the exact MST, and approximate the maximum weighted matchings if we permit $O(1/\gamma)$ -round adaptive sketches, i.e., a sequence of projections where each projection may be chosen dependent on the outcome of earlier sketches.

Our results have two main applications, both of which have the potential to give rise to fruitful lines of further research. First, our results can be thought of as giving the first *compressed-sensing style algorithms for graph data*. Secondly, our work initiates the study of *dynamic graph streams*. There is already extensive literature on processing massive graphs in the data-stream model. However, the existing work focuses on graphs defined by a sequence of inserted edges and does not consider edge deletions. We think this is a curious omission given the existing work on both dynamic graphs in the non-streaming setting and dynamic geometric streaming. Our results include the first dynamic graph semi-streaming algorithms for connectivity, spanning trees, sparsification, and matching problems.

Kook Jin Ahn
University of Pennsylvania
kookjin@cis.upenn.edu

Sudipto Guha
Computer and Information Science
University of Pennsylvania
sudipto@cis.upenn.edu

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

CP8

On the Communication and Streaming Complexity of Maximum Bipartite Matching

Consider the following communication problem. Alice holds a graph $G_A = (P, Q, E_A)$ and Bob holds a graph

$G_B = (P, Q, E_B)$, where $|P| = |Q| = n$. Alice is allowed to send Bob a message m that depends only on the graph G_A . Bob must then output a matching $M \subseteq E_A \cup E_B$. What is the minimum size of the message m that Alice sends to Bob that allows Bob to recover a matching of size at least $(1 - \epsilon)$ times the maximum matching in $G_A \cup G_B$? The minimum message length is the *one-round communication complexity* of approximating bipartite matching, which also gives a lower bound on the space needed by a one-pass streaming algorithm to compute a $(1 - \epsilon)$ -approximate matching. In this paper we address the question of how well one can approximate maximum bipartite matchings with linear communication and space. In order to study these questions, we introduce the concept of an ϵ -matching cover of a bipartite graph G , which is a sparse subgraph of the original graph that preserves the size of maximum matching between large subsets of vertices. We give a polynomial time construction of a $1/2$ -matching cover of size $O(n)$ with some crucial additional properties, thereby showing that Alice and Bob can achieve a $2/3$ -approximation with a message of size $O(n)$. Further, we establish a connection between ϵ -matching covers and so-called ϵ -Ruzsa-Szemerédi graphs, using it to show that any better approximation requires communication complexity at least $n^{1+\Omega(1/\log \log n)}$. Finally, we build on our techniques to design a one-pass streaming algorithm for the case of vertex arrivals, obtaining the first *deterministic* one-pass streaming $(1 - 1/e)$ -approximation using $O(n)$ space for this setting (*randomized* $1 - 1/e$ approximation can be obtained using the celebrated algorithm of Karp-Vazirani-Vazirani (KVV) for the *online* problem).

Ashish Goel, Michael Kapralov
Stanford University
ashishg@stanford.edu, kapralov@stanford.edu

Sanjeev Khanna
University of Pennsylvania
sanjeev@cis.upenn.edu

CP8

The Shifting Sands Algorithm

We resolve the problem of small-space approximate selection in random-order streams. Specifically, we present an algorithm that reads the n elements of a set in random order and returns an element whose rank differs from the true median by at most $n^{1/3+o(1)}$ while storing a constant number of elements and counters at any one time. This is optimal: it was previously shown that achieving better accuracy required $\text{poly}(n)$ memory. However, it was conjectured that the lower bound was not tight and that a previous algorithm achieving an $n^{1/2+o(1)}$ approximation was optimal. We therefore consider the new result a surprising resolution to a natural and basic question.

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

Paul Valiant
UC Berkeley
pvaliant@gmail.com

CP8

Width of Points in the Streaming Model

We show how to compute the width of a dynamic set of

low-dimensional points in the streaming model. We assume the stream contains both insertions and deletions of points to a set S , and the goal is to compute the width of S , the minimal distance between two parallel lines sandwiching S . Our algorithm is the first to $1+\epsilon$ approximate the width of S using space polylogarithmic in the size and aspect ratio of S .

Alexandr Andoni
Microsoft Research
andoni@microsoft.com

Huy L. Nguyen
Princeton University
hnguyen@princeton.edu

CP8

Lower Bounds for Number-in-Hand Multiparty Communication Complexity, Made Easy

In this paper we prove lower bounds on randomized multiparty communication complexity, both in the *blackboard model* (where each message is written on a blackboard for all players to see) and (mainly) in the *message-passing model*, where messages are sent player-to-player. We introduce a new technique for proving such bounds, called *symmetrization*, which is natural, intuitive, and often easy to use.

Qin Zhang
MADALGO, Aarhus University
qinzhang@cs.au.dk

Jeff Phillips
University of Utah
jeffp@cs.utah.edu

Elad Verbin
Dept. of Computer Science
Aarhus University
elad.verbin@gmail.com

CP9

Gathering Despite Mischief

A team consisting of an unknown number of mobile agents, starting from different nodes of an unknown network, have to meet at the same node. Agents move in synchronous rounds. Each agent has a different label. Up to f of the agents are Byzantine. What is the minimum number of good agents that guarantees deterministic gathering of all of them, with termination? We analyze this problem in our paper and give some tight bounds.

Yoann Dieudonne
Département d'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

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

CP9

Wireless Connectivity and Capacity

Given n wireless transceivers located in a plane, a fundamental problem in wireless communications is to construct a strongly connected digraph on them such that the constituent links can be scheduled in fewest possible time slots, assuming the SINR model of interference. In this paper, we provide an algorithm that connects an arbitrary point set in $O(\log n)$ slots, improving on the previous best bound of $O(\log^2 n)$ due to Moscibroda. This is complemented with a super-constant lower bound on our approach to connectivity. An important feature is that the algorithms allow for bi-directional (half-duplex) communication. One implication of this result is an improved bound of $\Omega(1/\log n)$ on the worst-case capacity of wireless networks, matching the best bound known for the extensively studied average-case. We explore the utility of oblivious power assignments, and show that essentially all such assignments result in a worst case bound of $\Omega(n)$ slots for connectivity. This rules out a recent claim of a $O(\log n)$ bound using oblivious power. On the other hand, using our result we show that $O(\min(\log \Delta, \log n \cdot (\log n + \log \log \Delta)))$ slots suffice, where Δ is the ratio between the largest and the smallest links in a minimum spanning tree of the points. We also initiate the study of network design problems in the SINR model beyond strong connectivity. We show that our approach extends to constructing biconnected, k -edge connected and low-diameter structures with similar bounds.

Magnús M. Halldórsson, Pradipta Mitra
Reykjavik University
mmh@ru.is, ppmitra@gmail.com

CP9

Stochastic Coalescence in Logarithmic Time

We study a distributed coalescence protocol due to Dahlia Malkhi (2006), motivated by social networking applications. Empirical results by Fernandess and Malkhi suggested the protocol concludes in $O(\log n)$ rounds with high probability, whereas numerical estimates by Oded Schramm, based on an ingenious analytic approximation, suggested that the coalescence time should be super-logarithmic. Here we confirm that the above process indeed requires super-logarithmic time w.h.p., and remedy this by showing that a simple modification produces an essentially optimal distributed protocol that *does* terminate in $O(\log n)$ rounds w.h.p. Simulations show that the new protocol readily outperforms the original one. Our upper bound hinges on a potential function involving the logarithm of the number of clusters and the cluster-susceptibility, carefully chosen to form a supermartingale. The analysis of the lower bound builds upon the novel approach of Schramm which may find additional applications: Rather than seeking a single parameter that controls the system behavior, instead one approximates the system by the Laplace transform of the entire cluster-size distribution.

Eyal Lubetzky
Microsoft Research
eyal@microsoft.com

Po-Shen Loh
Department of Mathematical Sciences
Carnegie Mellon University
ploh@cmu.edu

CP9**SINR Diagram with Interference Cancellation**

This paper studies the reception zones of a wireless network in the SINR model with receivers that employ interference cancellation (IC). IC is a recently developed technique that allows a receiver to decode interfering signals, and *cancel* them from the received signal in order to decode its intended message. We first derive the important topological properties of the reception zones and their relation to *high-order Voronoi diagrams* and other geometric objects. We then discuss the computational issues that arise when seeking an efficient description of the zones. Our main fundamental result states that although potentially there are exponentially many possible cancellation orderings, and as a result, reception zones, in fact there are much fewer nonempty such zones. We prove a linear bound (hence tight) on the number of zones and provide a polynomial time algorithm to describe the diagram. Moreover, we introduce a novel parameter, the *Compactness Parameter*, which influences the tightness of our bounds. We then utilize these properties to devise a logarithmic time algorithm to answer point-location queries for networks with IC.

Chen Avin, Asaf Cohen
Ben Gurion University, Beer-Sheva, Israel.
avin@cse.bgu.ac.il, coasaf@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
Weizmann Institute of Science
merav.parter@weizmann.ac.il

David Peleg
The Weizmann Institute of Science, Rehovot, Israel.
davod.peleg@weizmann.ac.il

CP9**Towards Robust and Efficient Computation in Dynamic Peer-to-Peer Networks**

Motivated by the need for robust and fast distributed computation in highly dynamic Peer-to-Peer (P2P) networks, we study algorithms for the fundamental distributed agreement problem. P2P networks are highly dynamic networks that experience heavy node *churn* (i.e., nodes join and leave the network continuously over time). Our goal is to design fast algorithms (running in a small number of rounds) that guarantee, despite high node churn rate, that almost all nodes reach a stable agreement. Our main contributions are randomized distributed algorithms that guarantee *stable almost-everywhere agreement* with high probability even under high adversarial churn in a polylogarithmic number of rounds. In particular, we present the following results:

1. An $O(\log^2 n)$ -round (n is the stable network size) randomized algorithm that achieves almost-everywhere

agreement with high probability under up to *linear churn per round* (i.e., ϵn , for some small constant $\epsilon > 0$), assuming that the churn is controlled by an oblivious adversary.

2. An $O(\log m \log^3 n)$ -round randomized algorithm that achieves almost-everywhere agreement with high probability under up to $\epsilon\sqrt{n}$ churn per round (for some small $\epsilon > 0$), where m is the size of the input value domain, that works even under an adaptive adversary.

Our algorithms are the first-known, fully-distributed, agreement algorithms that work under highly dynamic settings (i.e., high churn rates per step). Furthermore, they are localized (i.e., do not require any global topological knowledge), simple, and easy to implement.

John Augustine
Department of Computer Science and Engineering
Indian Institute of Technology Madras
augustine@cse.iitm.ac.in

Gopal Pandurangan, Peter Robinson
Division of Mathematical Sciences
Nanyang Technological University
gopalpandurangan@gmail.com, peter.robinson@ntu.edu.sg

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

CP10**Fully Persistent B-Trees**

We present I/O-efficient fully persistent B-Trees that support range searches at any version in $O(\log_B n + t/B)$ I/Os and updates at any version in $O(\log_B n + \log_2 B)$ amortized I/Os, using $O(m/B)$ disk blocks. n denotes the size of the accessed version, m the number of updates, t the query's output size, and B the disk block size.

We first present a new B-Tree implementation that supports searches and updates in $O(\log_B n)$ I/Os, using $O(n/B)$ blocks. Every update makes $O(1)$ worst case modifications to the structure. We make these B-Trees fully persistent using an I/O-efficient method for full persistence that is inspired by the *node-splitting* method of Driscoll et al. The method we present is interesting in its own right and can be applied to any external memory pointer based data structure with maximum in-degree d_{in} bounded by a constant and out-degree bounded by $O(B)$, where every node occupies a constant number of blocks on disk. The I/O-overhead per modification to the ephemeral structure is $O(d_{in} \log_2 B)$ amortized I/Os, and the space overhead is $O(d_{in}/B)$ amortized blocks. Access to a field of an ephemeral block is supported in $O(\log_2 d_{in})$ worst case I/Os.

Konstantinos A. Tsakalidis
Aarhus University
tsakalid@madalgo.au.dk

Gerth Brodal
MADALGO, Aarhus University
gerth@madalgo.au.dk

Kostas Tsichlas
Aristotle University of Thessaloniki
tsichlas@csd.auth.gr

Spyros Sioutas
 Ionian University
 sioutas@ionio.gr

CP10

A Little Advice Can Be Very Helpful

Proving superpolylogarithmic lower bounds for dynamic data structures has remained an open problem despite years of research. Recently, Patrascu proposed an exciting new approach for breaking this barrier via a two player communication model in which one player gets private advice at the beginning of the protocol. He gave reductions from the problem of solving an asymmetric version of set-disjointness in his model to a diverse collection of natural dynamic data structure problems in the cell probe model. He also conjectured that, for any hard problem in the standard two-party communication model, the asymmetric version of the problem is hard in his model, provided not too much advice is given. In this paper, we prove several surprising results about his model. We show that there exist Boolean functions requiring linear randomized communication complexity in the two-party model, for which the asymmetric versions in his model have deterministic protocols with exponentially smaller complexity. For set-disjointness, which also requires linear randomized communication complexity in the two-party model, we give a deterministic protocol for the asymmetric version in his model with a quadratic improvement in complexity. These results demonstrate that Patrascu's conjecture, as stated, is false. In addition, we show that the randomized and deterministic communication complexities of problems in his model differ by no more than a logarithmic multiplicative factor. We also prove lower bounds in some restricted versions of this model for natural functions such as set-disjointness and inner product. All of our upper bounds conform to these restrictions.

Arkadev Chattopadhyay
 University of Toronto
 arkadev@cs.toronto.edu

Jeff Edmonds
 York University
 jeff@cs.yorku.ca

Faith Ellen, Toniann Pitassi
 University of Toronto
 faith@cs.toronto.edu, toni@cs.toronto.edu

CP10

I/O-Efficient Data Structures for Colored Range and Prefix Reporting

Motivated by information retrieval applications, we consider the *one-dimensional colored range reporting* problem in rank space. The goal is to build a static data structure for sets $C_1, \dots, C_m \subseteq \{1, \dots, \sigma\}$ that supports queries of the kind: Given indices a, b , report the set $\bigcup_{a \leq i \leq b} C_i$. We study the problem in the I/O model, and show that there exists an optimal linear-space data structure that answers queries in $O(1+k/B)$ I/Os, where k denotes the output size and B the disk block size in words. In fact, we obtain the same bound for the harder problem of *three-sided orthogonal range reporting*. In this problem, we are to preprocess a set of n two-dimensional points in rank space, such that all points inside a query rectangle of the form $[x_1, x_2] \times (-\infty, y]$ can be reported. The best previous bounds for this problem

is either $O(n \lg_B^2 n)$ space and $O(1+k/B)$ query I/Os, or $O(n)$ space and $O(\lg_B^{(h)} n + k/B)$ query I/Os, where $\lg_B^{(h)} n$ is the base B logarithm iterated h times, for any constant integer h . The previous bounds are both achieved under the *indivisibility assumption*, while our solution exploits the full capabilities of the underlying machine. Breaking the indivisibility assumption thus provides us with cleaner and optimal bounds. Our results also imply an optimal solution to the following *colored prefix reporting* problem. Given a set S of strings, each $O(1)$ disk blocks in length, and a function $c : S \rightarrow 2^{\{1, \dots, \sigma\}}$, support queries of the kind: Given a string p , report the set $\bigcup_{x \in S \cap p^*} c(x)$, where p^* denotes the set of strings with prefix p . Finally, we consider the possibility of top- k extensions of this result, and present a simple solution in a model that allows non-blocked I/O.

Kasper Green Larsen
 MADALGO
 Aarhus University
 larsen@cs.au.dk

Rasmus Pagh
 IT University of Copenhagen
 pagh@itu.dk

CP10

Using Hashing to Solve the Dictionary Problem (In External Memory)

We consider the dictionary problem in external memory and improve the update time of the well-known *buffer tree* by roughly a logarithmic factor. For any $\lambda \geq \max\{\lg \lg n, \log_{M/B}(n/B)\}$, we can support updates in time $O(\frac{\lambda}{B})$ and queries in sublogarithmic time, $O(\log_\lambda n)$. We also present a lower bound in the cell-probe model showing that our data structure is optimal. In the RAM, hash tables have been used to solve the dictionary problem faster than binary search for more than half a century. By contrast, our data structure is the first to beat the comparison barrier in external memory. Ours is also the first data structure to depart convincingly from the *indivisibility* paradigm.

John Iacono
 The Polytechnic Institute of New York University
 johniacono@gmail.com

Mihai Pătraşcu
 AT&T
 mip@research.att.com

CP10

Confluent Persistence Revisited

It is shown how to enhance any data structure in the pointer model to make it confluent persistent, with efficient query and update times and limited space overhead. Updates are performed in $O(\log n)$ amortized time, and following a pointer takes $O(\log c \log n)$ time where c is the in-degree of a node in the data structure. In particular, this proves that confluent persistence can be achieved at a logarithmic cost in the bounded in-degree model used widely in previous work. This is a $O(n/\log n)$ -factor improvement over the previous known transform to make a data structure confluent persistent.

John Iacono

The Polytechnic Institute of New York University
johniacono@gmail.com

Sébastien Collette, Stefan Langerman
Universite Libre de Bruxelles
secollet@ulb.ac.be, stefanlangerman@ulb.ac.be

CP11

A Polynomial-Time Approximation Scheme for Planar Multiway Cut

Given an undirected graph with edge lengths and a subset of nodes (called the terminals), the multiway cut (also called the multi-terminal cut) problem asks for a subset of edges, with minimum total length, whose removal disconnects each terminal from all others. The problem generalizes minimum s-t cut, but is NP-hard for planar graphs and APX-hard for general graphs [12]. In this paper, we present a PTAS for multiway cut on planar graphs.

Mohammadhossein Bateni
Google Research
New York, NY
bateni@gmail.com

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

Philip Klein, Claire Mathieu
Brown University
klein@cs.brown.edu, claire@cs.brown.edu

CP11

Finding An Induced Path of Given Parity in Planar Graphs in Polynomial Time

The problem of deciding, given a graph G and two vertices s and t , whether there exists an induced path of given parity between s and t in G is known to be NP-complete. We show how to solve the problem in $\mathcal{O}(|V(G)|^7)$ time, when the input graph is planar. We use techniques from the theory of graph minors as well as the theory of perfect graphs.

Marcin Kaminski
Université Libre de Bruxelles
Belgium
Marcin.Kaminski@ulb.ac.be

Naomi Nishimura
University of Waterloo
nishi@uwaterloo.ca

CP11

Approximate Tree Decompositions of Planar Graphs in Linear Time

Many algorithms have been developed for NP-hard problems on graphs with small treewidth k . For example, all problems that are expressible in linear extended monadic second order can be solved in linear time on graphs of bounded treewidth. It turns out that the bottleneck of many algorithms for NP-hard problems is the computation of a tree decomposition of width $O(k)$. In particular, by the bidimensional theory, there are many linear extended monadic second order problems that can be solved on n -vertex planar graphs with treewidth k in a time linear in

n and subexponential in k if a tree decomposition of width $O(k)$ can be found in such a time. We present the first algorithm that, on n -vertex planar graphs with treewidth k , finds a tree decomposition of width $O(k)$ in such a time. In more detail, our algorithm has a running time of $O(nk^3 \log k)$. The previous best algorithm with a running time subexponential in k was the algorithm of Gu and Tamaki (ISAAC 2009) with a running time of $O(n^{1+\epsilon} \log n)$ and an approximation ratio $1.5 + 1/\epsilon$ for any $\epsilon > 0$. The running time of our algorithm is also better than the running time of $O(f(k) \cdot n \log n)$ of Reed's algorithm (STOC 1992) for general graphs, where f is a function exponential in k .

Frank Kammer
University of Augsburg
kammer@informatik.uni-augsburg.de

Torsten Tholey
Universität Augsburg
tholey@informatik.uni-augsburg.de

CP11

An Efficient Polynomial-Time Approximation Scheme for Steiner Forest in Planar Graphs

We give an $O(n \log^3 n)$ approximation scheme for Steiner forest in planar graphs, improving on the previous approximation scheme for this problem, which runs in $O(n^{f(\epsilon)})$ time.

David Eisenstat, Philip Klein, Claire Mathieu
Brown University
eisenstatdavid@gmail.com, klein@cs.brown.edu,
claire@cs.brown.edu

CP11

Spanning Closed Walks and Tsp in 3-Connected Planar Graphs

We consider the following problem which is motivated by two different contexts independently, namely graph theory and combinatorial optimization. Given a 3-connected planar graph G with n vertices, is there a spanning closed walk W with at most $4n/3$ edges? We give a positive answer to the above problem. Moreover, our proof is constructive (and purely combinatorial) in a sense that there is an $O(n^2)$ algorithm to construct, given a 3-connected planar graph, such a walk.

Ken-ichi Kawarabayashi, Kenta Ozeki
National Institute of Informatics, Japan
k_keniti@nii.ac.jp, ozeki@comb.math.keio.ac.jp

CP12

The Set of Solutions of Random Xorsat Formulae

We consider random k -XORSAT instances, drawn uniformly at random from the ensemble of Boolean formulae containing n variables and m XOR clauses of size k each. As m/n grows, the set of solutions shatters into an exponential number of well-separated components. This phenomenon appears to be related to hardness in random CSPs like k -SAT. We prove a complete characterization of this clustering phase transition for random k -XORSAT. In particular, we prove that the clustering threshold is sharp, and each of the clusters is well connected above this thresh-

old.

Yashodhan Kanoria, Morteza Ibrahimi, Matt Kranning,
Andrea Montanari
Stanford University
ykanoria@stanford.edu, ibrahimi@stanford.edu, mkran-
ing@stanford.edu, montanari@stanford.edu

CP12

The Complexity of Conservative Valued CSPs

We study the complexity of valued constraint satisfaction problems (VCSP). A problem from VCSP is characterised by a *constraint language*, a fixed set of cost functions over a finite domain. An instance of the problem is specified by a sum of cost functions from the language and the goal is to minimise the sum. We consider the case of languages containing all possible unary cost functions. In the case of languages consisting of only $\{0, \infty\}$ -valued cost functions (i.e. relations), such languages have been called *conservative* and studied by Bulatov [LICS'03] and recently by Barto [LICS'11]. Since we study valued languages, we call a language *conservative* if it contains all finite-valued unary cost functions. The computational complexity of conservative valued languages has been studied by Cohen et al. [AIJ'06] for languages over Boolean domains, by Deineko et al. [JACM'08] for $\{0, 1\}$ -valued languages (a.k.a Max-CSP), and by Takhanov [STACS'10] for $\{0, \infty\}$ -valued languages containing all finite-valued unary cost functions (a.k.a. Min-Cost-Hom). We prove a Schaefer-like dichotomy theorem for conservative valued languages: if all cost functions in the language satisfy a certain condition (specified by a complementary combination of *STP* and *MJN multimorphisms*), then any instance can be solved in polynomial time (via a new algorithm developed in this paper), otherwise the language is NP-hard. This is the *first* complete complexity classification of *general-valued constraint languages* over non-Boolean domains. Our results generalise previous results by Takhanov [STACS'10] and (a subset of results) by Cohen et al. [AIJ'06] and Deineko et al. [JACM'08]. Moreover, our results do not rely on any computer-assisted search as in Deineko et al. [JACM'08], and provide a powerful tool for proving hardness of finite-valued and general-valued languages.

Vladimir Kolmogorov
Institute of Science and Technology (IST), Austria
vnk@ist.ac.at

Stanislav Zivny
Oxford University
standa.zivny@cs.ox.ac.uk

CP12

Inapproximability Results for the Multi-Level Uncapacitated Facility Location Problem

In this paper, we present improved inapproximability results for the k -level uncapacitated facility location problem. In particular, we show that there is no polynomial time approximation algorithm with performance guarantee better than 1.539 unless NP is contained in $DTIME(n^{O(\log \log n)})$ for the case when $k = 2$. For the case of general k (tending to infinity) we obtain a better hardness factor of 1.61. Interestingly, our results show that the two-level problem is *computationally harder* than the well known uncapacitated facility location problem ($k = 1$) since the best known approximation guarantee for the latter problem is 1.488 due to Li, and our inapproximability is a factor of 1.539 for

the two-level problem. The only inapproximability result known before for this class of metric facility location problems is the bound of 1.463 due to Guha and Khuller, which holds even for the case of $k = 1$.

Ravishankar Krishnaswamy
Computer Science Department
Carnegie Mellon University
ravishan@cs.cmu.edu

Maxim Sviridenko
IBM T.J. Watson Research Center
sviri@us.ibm.com

CP12

Bypassing UGC from some Optimal Geometric Inapproximability Results

The Unique Games Conjecture (UGC) asserts that a certain two variable constraint satisfaction problem with bijective constraints is NP-hard to approximate. While the conjecture remains an important open question, it has yielded several optimal inapproximability results - many of which depend critically on the assumption of bijective constraints. In this work we observe that some geometric inapproximability results - previously only known based on the UGC - can be shown without appealing to the conjecture. Specifically, we obtain optimal NP-hardness of approximation results for two problems: the L_p Subspace Approximation Problem and the L_p Quadratic Grothendieck Maximization Problem for constant $p > 2$. These results are based on the so called "smooth" Label Cover instances which have "locally bijective" constraints. We also obtain an optimal approximation for the L_p Quadratic Grothendieck Maximization problem using a convex relaxation of the problem. A similar result was independently obtained by Naor and Schechtman.

Venkatesan Guruswami
Carnegie Mellon University
guruswami@cmu.edu

Prasad Raghavendra
Georgia Tech
raghavendra@cc.gatech.edu

Rishi Saket
Princeton University
rsaket@cs.princeton.edu

Yi Wu
IBM Almaden Research
wuyi@us.ibm.com

CP12

On the Hardness of Pricing Loss Leaders

Consider the problem of pricing n items under an unlimited supply with m buyers. Each buyer is interested in a bundle of at most k of the items. These buyers are single minded, which means each of them has a budget and they will either buy all the items if the total price is within their budget or they will buy none of the items. The goal is to price each item with profit margin p_1, p_2, \dots, p_n so as to maximize the overall profit. When $k = 2$, such a problem is called the graph vertex pricing problem. Another special case of the problem is the highway problem when the items (toll-booths) are arranged linearly on a line and each buyer

(as a driver) is interested in paying for a path that consists of consecutive items. The goal again is to price the items (tolls) so as to maximize the total profits. In this paper, we obtain strong hardness of approximation result for the problem of pricing loss leaders. First we show that it is NP-hard to get better than $O(\log \log \log n)$ -approximation when $k \geq 3$. This improves a previous super-constant hardness result assuming the Unique Games Conjecture [Wu11]. In addition, we show a super-constant unique games hardness for the highway problem as well as the graph vertex pricing problem.

Yi Wu

IBM Almaden Research
wuyi@us.ibm.com

CP13

Approximate Duality of Multicommodity Multi-route Flows and Cuts: Single Source Case

Given an integer h , a graph $G = (V, E)$ with edge capacities and k pairs of vertices $(s_1, t_1), (s_2, t_2), \dots, (s_k, t_k)$, called terminals, an h -route cut is a set $F \subseteq E$ of edges such that after the removal of the edges in F no pair $s_i - t_i$ is connected by h edge-disjoint paths (i.e., the connectivity of every $s_i - t_i$ pair is at most $h - 1$ in $(V, E \setminus F)$). The h -route cut is a natural generalization of the classical cut problem for multicommodity flows (take $h = 1$). The main result of this paper is an $O(h^7 2^{2h} \log^2 k)$ -approximation algorithm for the minimum h -route cut problem in the case that $s_1 = s_2 = \dots = s_k$, called the single source case. As a corollary of it we obtain an approximate duality theorem for multiroute multicommodity flows and cuts with a single source. This partially answers an open question posted in several previous papers dealing with cuts for multicommodity multiroute problems.

Petr Kolman

Charles University, Faculty of Mathematics and Physics
Department of Applied Mathematics
kolman@kam.mff.cuni.cz

Christian Scheideler

University of Paderborn
scheidel@mail.uni-paderborn.de

CP13

Constant Factor Approximation Algorithm for the Knapsack Median Problem

We give a constant factor approximation algorithm for the following generalization of the k -median problem. We are given a set of clients and facilities in a metric space. Each facility has a facility opening cost, and we are also given a budget B . The objective is to open a subset of facilities of total cost at most B , and minimize the total connection cost of the clients. This settles an open problem of Krishnaswamy-Kumar-Nagarajan-Sabharwal-Saha. The natural linear programming relaxation for this problem has unbounded integrality gap. Our algorithm strengthens this relaxation by adding constraints which stipulate which facilities a client can get assigned to.

Amit Kumar

IIT Delhi
amitk@cse.iitd.ac.in

CP13

Approximation Algorithms and Hardness of the k -Route Cut Problem

We study the k -route cut problem: given an undirected edge-weighted graph $G = (V, E)$, a collection $\{(s_1, t_1), (s_2, t_2), \dots, (s_r, t_r)\}$ of source-sink pairs, and an integer connectivity requirement k , the goal is to find a minimum-weight subset E' of edges to remove, such that the connectivity of every pair (s_i, t_i) falls below k . Specifically, in the edge-connectivity version, EC- k RRC, the requirement is that there are at most $(k - 1)$ edge-disjoint paths connecting s_i to t_i in $G \setminus E'$, while in the vertex-connectivity version, VC- k RRC, the same requirement is for vertex-disjoint paths. Prior to our work, poly-logarithmic approximation algorithms have been known for the special case where $k \leq 3$, but no non-trivial approximation algorithms were known for any value $k > 3$, except in the single-source setting. We show an $O(k \log^{3/2} r)$ -approximation algorithm for EC- k RRC with uniform edge weights, and several polylogarithmic bi-criteria approximation algorithms for EC- k RRC and VC- k RRC, where the connectivity requirement k is violated by a constant factor. We complement these upper bounds by proving that VC- k RRC is hard to approximate to within a factor of k^ϵ for some fixed $\epsilon > 0$. We then turn to study a simpler version of VC- k RRC, where only one source-sink pair is present. We give a simple bi-criteria approximation algorithm for this case, and show evidence that even this restricted version of the problem may be hard to approximate. For example, we prove that the single source-sink pair version of VC- k RRC has no constant-factor approximation, assuming Feige's Random κ -AND assumption.

Julia Chuzhoy, Yury Makarychev

Toyota Technological Institute at Chicago
cjulia@ttic.edu, ymakarychev@gmail.com

Aravindan Vijayaraghavan

Princeton University
aravindv@cs.princeton.edu

Yuan Zhou

Carnegie Mellon University
yuanzhou@cs.cmu.edu

CP13

Subquadratic Time Approximation Algorithms for the Girth

We study the problem of computing the girth of an unweighted undirected graph. We obtain several new efficient approximation algorithms for graphs with n nodes and m edges and unknown girth g . We consider additive and multiplicative approximations.

Additive Approximations. We present:

- an $\tilde{O}(n^3/m)$ -time algorithm which returns a cycle of length at most $g + 2$ if g is even and $g + 3$ if g is odd. This complements the seminal work of Itai and Rodeh [SIAM J. Computing'78] who gave an algorithm that in $O(n^2)$ time finds a cycle of length g if g is even, and $g + 1$ if g is odd.
- an $\tilde{O}(n^3/m)$ -time algorithm which returns a cycle of length at most $g' + 2$, where g' is the length of the shortest *even* cycle in G . This result complements the work of Yuster and Zwick [SIAM J. Discrete Math'97] who showed how to compute g' in $O(n^2)$ time.

Multiplicative Approximations. We present:

- an $\tilde{O}(n^{5/3})$ -time algorithm which returns a cycle of length at most $3g/2 + z/2$ when g is even and $3g/2 + z/2 + 1$ when g is odd, where $z = -g \bmod 4$, $z \in \{0, 1, 2, 3\}$. This gives an $\tilde{O}(n^{5/3})$ -time 2-approximation for the girth, the first subquadratic time 2-approximation algorithm, resolving an open question of Lingas and Lundell [IPL'09].
- an $O(n^{1.968})$ -time $(8/5)$ -approximation algorithm for the girth in graphs with girth at least 4 (i.e., triangle-free graphs). This is the first subquadratic time $(2 - \varepsilon)$ -approximation algorithm for the girth for triangle-free graphs, for any $\varepsilon > 0$. We prove that a deterministic algorithm of this kind is not possible for directed graphs, thus showing a strong separation between undirected and directed graphs for girth approximation.

Virginia V. Williams
 UC Berkeley
 virgito@gmail.com

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

CP13

On a Linear Program for Minimum-Weight Triangulation

Minimum-weight triangulation (MWT) is NP-hard. It has a polynomial-time constant-factor approximation algorithm, and a variety of effective polynomial-time heuristics that, for many instances, can find the exact MWT. Linear programs (LPs) for MWT are well-studied, but previously no connection was known between any LP and any approximation algorithm or heuristic for MWT. Here we show the first such connections, for an LP formulation due to Dantzig et al. (1985): (i) the integrality gap is constant; (ii) given any instance, if the aforementioned heuristics find the MWT, then so does the LP.

Arman Yousefi, Neal E. Young
 University of California, Riverside
 yousefia@cs.ucr.edu, neal@cs.ucr.edu

CP14

Mechanism Design Via Consensus Estimates, Cross Checking, and Profit Extraction

There is only one technique for *prior-free optimal mechanism design* that generalizes beyond the structurally benevolent setting of digital goods. This technique uses random sampling to estimate the distribution of agent values and then employs the *Bayesian optimal mechanism* for this estimated distribution on the remaining players. Though quite general, even for digital goods, this random sampling auction has a complicated analysis and is known to be suboptimal. To overcome these issues we generalize the profit extraction and consensus techniques from Goldberg and Hartline (2003) to structurally rich environments that include, e.g., single-minded combinatorial auctions.

Bach Q. Ha, Jason D. Hartline
 Northwestern University
 bach@u.northwestern.edu, hartline@eecs.northwestern.edu

CP14

Optimal Crowdsourcing Contests

We study the design and approximation of optimal crowdsourcing contests. Crowdsourcing contests can be modeled as all-pay auctions because entrants must exert effort upfront to enter. Unlike all-pay auctions where a usual design objective would be to maximize revenue, in crowdsourcing contests, the principal only benefits from the submission with the highest quality. We give a theory for optimal crowdsourcing contests that mirrors the theory of optimal auction design: the optimal crowdsourcing contest is a virtual valuation optimizer (the virtual valuation function depends on the distribution of contestant skills and the number of contestants). We also compare crowdsourcing contests with more conventional means of procurement. In this comparison, crowdsourcing contests are relatively disadvantaged because the effort of losing contestants is wasted. Nonetheless, we show that crowdsourcing contests are 2-approximations to conventional methods for a large family of “regular” distributions, and 4-approximations, otherwise.

Balasubramanian Sivan
 University of Wisconsin-Madison
 balu2901@cs.wisc.edu

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

Jason D. Hartline
 Northwestern University
 hartline@eecs.northwestern.edu

CP14

Black-Box Reductions for Cost-Sharing Mechanism Design

We consider the design of *strategyproof* cost-sharing mechanisms. We give two simple, but extremely versatile, *black-box reductions*, that in combination reduce the cost-sharing mechanism-design problem to the *algorithmic* problem of finding a minimum-cost solution for a set of players. Our first reduction shows that any α -approximation, truthful mechanism for the social-cost-minimization (SCM) problem that satisfies a technical no-bossiness condition can be morphed into a truthful mechanism that achieves an $\alpha \log n$ -approximation where the prices recover the cost incurred. This disconnects the task of truthfully computing an outcome with near-optimal social cost from the cost-sharing problem. Complementing this, our second reduction shows that any LP-based ρ -approximation for the problem of finding a min-cost solution for a set of players can be used to obtain a truthful, no-bossy, $(\rho + 1)$ -approximation for the SCM problem (and hence, a truthful $(\rho + 1) \log n$ -approximation cost-sharing mechanism).

Konstantinos Georgiou, Chaitanya Swamy
 University of Waterloo
 k2georgiou@math.uwaterloo.ca,
 cswamy@math.uwaterloo.ca

CP14

Sequential Auctions and Externalities

In many settings agents participate in multiple different auctions that are not necessarily implemented simultane-

ously. Future opportunities affect strategic considerations of the players in each auction, introducing externalities. Motivated by this consideration, we study a setting of a market of buyers and sellers, where each seller holds one item, bidders have combinatorial valuations and sellers hold item auctions sequentially. Our results are qualitatively different from those of simultaneous auctions, proving that simultaneity is a crucial aspect of previous work. We prove that if sellers hold sequential first price auctions then for unit-demand bidders (matching market) every subgame perfect equilibrium achieves at least half of the optimal social welfare, while for submodular bidders or when second price auctions are used, the social welfare can be arbitrarily worse than the optimal. We also show that a first price sequential auction for buying or selling a base of a matroid is always efficient, and implements the VCG outcome. An important tool in our analysis is studying first and second price auctions with externalities (bidders have valuations for each possible winner outcome), which can be of independent interest. We show that a Pure Nash Equilibrium always exists in a first price auction with externalities.

Renato Paes Leme, Vasilis Syrgkanis
Cornell University
renatopl@cs.cornell.edu, vasilis@cs.cornell.edu

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

CP14

A Universally-Truthful Approximation Scheme for Multi-Unit Auctions

We present a randomized, polynomial-time approximation scheme for multi-unit auctions. Our mechanism is truthful in the universal sense, i.e., a distribution over deterministically truthful mechanisms. Previously known approximation schemes were truthful in expectation which is a weaker notion of truthfulness assuming risk neutral bidders. The existence of a universally truthful approximation scheme was questioned by previous work showing that multi-unit auctions with certain technical restrictions on their output do not admit a polynomial-time, universally truthful mechanism with approximation factor better than two. Our new mechanism employs VCG payments in a non-standard way: The deterministic mechanisms underlying our universally truthful approximation scheme are not maximal in range and do not belong to the class of affine maximizers which, on a first view, seems to contradict previous characterizations of VCG-based mechanisms. Instead, each of these deterministic mechanisms is composed of a collection of affine maximizers, one for each bidder. This yields a subjective variant of VCG in which payments for different bidders are defined on the basis of possibly different affine maximizers.

Berthold Vöcking
RWTH Aachen
voecking@cs.rwth-aachen.de

CP15

Counting Perfect Matchings As Fast As Ryser

We present a polynomial space algorithm that counts the number of perfect matchings in an n -vertex graph in $O(\text{poly}(n)2^{n/2}) \subset O(1.415^n)$ time. The previously fastest algorithms was the exponential space $O(\text{poly}(n)((1 +$

$\sqrt{5})/2)^n) \subset O(1.619^n)$ time algorithm by Koivisto, and for polynomial space, the $O(1.942^n)$ time algorithm by Nederlof. Our new algorithm's runtime matches up to polynomial factors that of Ryser's 1963 algorithm for bipartite graphs. We present our algorithm in the general setting of computing the hafnian over an arbitrary ring.

Andreas Björklund
Lund University
andreas.bjorklund@yahoo.se

CP15

A Satisfiability Algorithm for AC^0

We consider the problem of efficiently enumerating the satisfying assignments to AC^0 circuits. We give a zero-error randomized algorithm which takes an AC^0 circuit as input and constructs a set of restrictions which partitions $\{0, 1\}^n$ so that under each restriction the value of the circuit is constant. Let d denote the depth of the circuit and cn denote the number of gates. This algorithm runs in time $|C|2^{n(1-\mu_{c,d})}$ where $|C|$ is the size of the circuit for $\mu_{c,d} \geq 1/O[\lg c + d \lg d]^{d-1}$ with probability at least $1-2^{-n}$. As a result, we get improved exponential time algorithms for AC^0 circuit satisfiability and for counting solutions. In addition, we get an improved bound on the correlation of AC^0 circuits with parity. As an important component of our analysis, we extend the Håstad Switching Lemma to handle multiple k -CNFs and k -DNFs.

Ramamohan Paturi, Russell Impagliazzo
University of California, San Diego
paturi@cs.ucsd.edu, russell@cs.ucsd.edu

William Matthews
Google Inc.
wgmatthews@eng.ucsd.edu

CP15

Approximating Fixation Probabilities in the Generalized Moran Process

In the Moran process, as generalized by Lieberman, Hauert and Nowak (*Nature*, 433:312–316, 2005), a population including one mutant of relative “fitness” r resides on the vertices of a connected, undirected graph. At each step, an individual is chosen to reproduce with probability proportional to its fitness, replacing one of its neighbours with a copy of itself. We give an FPRAS for the probabilities that the mutant's descendants take over the graph (when $r \geq 1$) and that they become extinct (for all $r > 0$).

George B. Mertzios
School of Engineering and Computing Sciences
Durham University
george.mertzios@durham.ac.uk

Josep Diaz
Departament de Llenguatges i Sistemes Informàtics
Universitat Politècnica de Catalunya, Spain
diaz@lsi.upc.edu

Leslie Ann Goldberg, David Richerby
Department of Computer Science
University of Liverpool, UK
l.a.goldberg@liverpool.ac.uk,
david.richerby@liverpool.ac.uk

Maria Serna
 Departament de Llenguatges i Sistemes Informàtics
 Universitat Politècnica de Catalunya, Spain
 serna@lsi.upc.edu

Paul Spirakis
 Department of Computer Engineering and Informatics
 University of Patras, Greece
 spirakis@cri.gr

CP15

Approximation Algorithms for Two-State Anti-Ferromagnetic Spin Systems on Bounded Degree Graphs

We show that for anti-ferromagnetic Ising models on the d -regular tree, weak spatial mixing implies strong spatial mixing. Our proof uses a message-decay argument inspired by a similar approach proposed recently for the hard-core model by Restrepo *et al* [2011]. Using a reduction of Weitz [2006], we obtain as a corollary a deterministic FPTAS for the partition function of the anti-ferromagnetic Ising model with arbitrary field on graphs of degree at most d , throughout the uniqueness region of the d -regular tree. By a standard correspondence, these results translate to arbitrary two-state anti-ferromagnetic spin systems.

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

Piyush Srivastava
 University of California Berkeley
 piyushsr@eecs.berkeley.edu

Marc Thurley
 Centre de Recerca Matemàtica, Bellaterra, Spain
 marc.thurley@googlemail.com

CP15

Approximate Counting Via Correlation Decay in Spin Systems

We give the first deterministic fully polynomial-time approximation scheme (FPTAS) for computing the partition function of a two-state spin system on an arbitrary graph, when the parameters of the system satisfy the uniqueness condition on infinite regular trees. This condition is of physical significance and is believed to be the right boundary between approximable and inapproximable. The FPTAS is based on the correlation decay technique introduced by Bandyopadhyay and Gamarnik [SODA 06] and Weitz [STOC 06]. The classic correlation decay is defined with respect to graph distance. Although this definition has natural physical meanings, it does not directly support an FPTAS for systems on arbitrary graphs, because for graphs with unbounded degrees, the local computation that provides a desirable precision by correlation decay may take super-polynomial time. We introduce a notion of *computationally efficient correlation decay*, in which the correlation decay is measured in a refined metric instead of graph distance. We use a potential method to analyze the amortized behavior of this correlation decay and establish a correlation decay that guarantees an inverse-polynomial precision by polynomial-time local computation. This gives us an FPTAS for spin systems on arbitrary graphs. This new notion of correlation decay properly reflects the algorithmic aspect of the spin systems, and may be used for designing

FPTAS for other counting problems.

Yitong Yin
 Nanjing University
 yinyt@nju.edu.cn

CP16

Sketching Valuation Functions

Motivated by the problem of querying and communicating bidders' valuations in combinatorial auctions, we study how well different classes of set functions can be sketched. More formally, let f be a function mapping subsets of some ground set $[n]$ to the non-negative real numbers. We say that f' is an α -sketch of f if for every set S , the value $f'(S)$ lies between $f(S)/\alpha$ and $f(S)$, and f' can be specified by $\text{poly}(n)$ bits. We show that for every subadditive function f there exists an α -sketch where $\alpha = n^{1/2} \cdot O(\text{polylog}(n))$. Furthermore, we provide an algorithm that finds these sketches with a polynomial number of demand queries. This is essentially the best we can hope for since:

1. We show that there exist subadditive functions (in fact, XOS functions) that do not admit an $o(n^{1/2})$ sketch. (Balcan and Harvey [?] previously showed that there exist functions belonging to the class of substitutes valuations that do not admit an $O(n^{1/3})$ sketch.)
2. We prove that every deterministic algorithm that accesses the function via *value* queries only cannot guarantee a sketching ratio better than $n^{1-\epsilon}$.

We also show that coverage functions, an interesting subclass of submodular functions, admit arbitrarily good sketches. Finally, we show an interesting connection between sketching and learning. We show that for every class of valuations, if the class admits an α -sketch, then it can be α -approximately learned in the PMAC model of Balcan and Harvey. The bounds we prove are only information-theoretic and do not imply the existence of computationally efficient learning algorithms in general.

Ashwinkumar Badanidiyuru, Shahar Dobzinski, Hu Fu,
 Robert Kleinberg
 Cornell University
 ashwin85@cs.cornell.edu, dobzin@gmail.com,
 hufu@cs.cornell.edu, rdk@cs.cornell.edu

Noam Nisan
 Hebrew University of Jerusalem
 noam@cs.huji.ac.il

Tim Roughgarden
 Stanford University
 Computer Science Department
 tim@cs.stanford.edu

CP16

Voting with Limited Information and Many Alternatives

The traditional axiomatic approach to voting is motivated by the problem of reconciling differences in subjective preferences. In contrast, a dominant line of work in the theory of voting over the past 15 years has considered a different kind of scenario, also fundamental to voting, in which there is a genuinely "best" outcome that voters would agree on if they only had enough information. This type of scenario has its roots in the classical Condorcet Jury Theorem; it includes cases such as jurors in a criminal trial who all

want to reach the correct verdict but disagree in their inferences from the available evidence, or a corporate board of directors who all want to improve the company's revenue, but who have different information that favors different options. This style of voting leads to a natural set of questions: each voter has a *private signal* that provides probabilistic information about which option is best, and a central question is whether a simple plurality voting system, which tabulates votes for different options, can cause the group decision to arrive at the correct option. We show that plurality voting is powerful enough to achieve this: there is a way for voters to map their signals into votes for options in such a way that — with sufficiently many voters — the correct option receives the greatest number of votes with high probability. We show further, however, that any process for achieving this is inherently expensive in the number of voters it requires: succeeding in identifying the correct option with probability at least $1 - \eta$ requires $\Omega(n^3 \epsilon^{-2} \log \eta^{-1})$ voters, where n is the number of options and ϵ is a distributional measure of the minimum difference between the options.

Jon M. Kleinberg
Cornell University
Dept of Computer Science
kleinber@cs.cornell.edu

Flavio Chierichetti
Cornell University
Cornell University
flavio@chierichetti.name

CP16 Metastability of Logit Dynamics for Coordination Games

Logit Dynamics [Blume, Games and Economic Behavior, 1993] is a randomized best response dynamics for strategic games: at every time step a player is selected uniformly at random and she chooses a new strategy according to a probability distribution biased toward strategies promising higher payoffs. This process defines an ergodic Markov chain, over the set of strategy profiles of the game, whose unique stationary distribution is the long-term equilibrium concept for the game. However, when the mixing time of the chain is large (e.g. exponential in the number of players), the stationary distribution loses its appeal as equilibrium concept, and the transient phase of the Markov chain becomes important. In several cases it happens that on a time-scale shorter than mixing time the chain is “quasi-stationary”, meaning that it stays close to some small set of the state space, while in a time-scale multiple of the mixing time it jumps from one quasi-stationary configuration to another; this phenomenon is usually called “metastability”. In this paper we give a quantitative definition of “metastable probability distributions” for a Markov chain and we study the metastability of the Logit dynamics for some classes of coordination games. In particular, we study no-risk-dominant coordination games on the clique (which is equivalent to the well-known Glauber dynamics for the Ising model) and coordination games on a ring (both the risk-dominant and no-risk-dominant case). We also describe a simple “artificial” game that highlights the distinctive features of our metastability notion based on distributions.

Francesco Pasquale
Dipartimento di Informatica - Università di Salerno
pasquale@dia.unisa.it

Vincenzo Auletta, Diodato Ferraioli, Giuseppe Persiano
Università di Salerno
auletta@dia.unisa.it, ferraioli@dia.unisa.it,
giuper@dia.unisa.it

CP16 Beyond Myopic Best Response in Cournot Competition

A Nash Equilibrium is a joint strategy profile at which each agent myopically plays a best response to the other agents' strategies, ignoring the possibility that deviating from the equilibrium could lead to an avalanche of successive changes by other agents. However, such changes could potentially be beneficial to the agent, creating incentive to act non-myopically, so as to take advantage of others' responses. To study this phenomenon, we consider a non-myopic Cournot competition, where each firm selects whether it wants to maximize profit (as in the classical Cournot competition) or to maximize revenue (by masquerading as a firm with zero production costs). We study the properties of Nash Equilibria of *non-myopic* Cournot competition with linear demand functions and show existence of pure Nash Equilibria, that simple best response dynamics will produce such an equilibrium, and that for some natural dynamics this convergence is within linear time. This is in contrast to the well known fact that best response dynamics need not converge in the standard myopic Cournot competition. Furthermore, we compare the outcome of the non-myopic Cournot competition with that of the standard myopic Cournot competition. Not surprisingly, perhaps, prices in the non-myopic game are lower and the firms, in total, produce more and have a lower aggregate utility.

Svetlana Olonetsky, Amos Fiat
Tel Aviv University
olonetsk@post.tau.ac.il, fiat@tau.ac.il

Elias Koutsoupias
University of Athens
elias@di.uoa.gr

Katrina Ligett
Caltech
katrina@caltech.edu

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

CP16 The Notion of a Rational Convex Program, and An Algorithm for the Arrow-Debreu Nash Bargaining Game

We introduce the notion of a *rational convex program (RCP)* and we classify the known RCPs into two classes: quadratic and logarithmic. The importance of rationality is that it opens up the possibility of computing an optimal solution to the program via an algorithm that is either combinatorial or uses an LP-oracle. Next, from the linear case of the Arrow-Debreu market model, we define a new Nash bargaining game, which we call ADNB. We show that the convex program for ADNB is a logarithmic RCP, but unlike other known members of this class, it is non-total. Our main result is a combinatorial, polynomial time algorithm for ADNB. It turns out that the reason for in-

feasibility of logarithmic RCPs is quite different from that for LPs and quadratic RCPs. Finally, we present a number of interesting questions that the new notion of RCP raises.

Vijay V. Vazirani
Georgia Institute of Technology
vazirani@cc.gatech.edu

CP17

Partial Match Queries in Random Quadrees

We consider the problem of recovering items matching a partially specified pattern in multidimensional trees (quad trees and k -d trees). We assume the traditional model where the data consist of independent and uniform points in the unit square. For this model, in a structure on n points, it is known that the number of nodes $C_n(\xi)$ to visit in order to report the items matching an independent and uniformly on $[0, 1]$ random query ξ satisfies $\mathbf{E}C_n(\xi) \sim \kappa n^\beta$, where κ and β are explicit constants. We develop an approach based on the analysis of the cost $C_n(x)$ of any fixed query $x \in [0, 1]$, and give precise estimates for the variance and limit distribution of the cost $C_n(x)$. Our results permit to describe a limit process for the costs $C_n(x)$ as x varies in $[0, 1]$; one of the consequences is that $\mathbf{E}\max_{x \in [0, 1]} C_n(x) \sim \gamma n^\beta$; this settles a question of Devroye [Pers. Comm., 2000].

Nicolas Broutin
Inria
nicolas.broutin@inria.fr

Ralph Neininger, Henning Sulzbach
J.W. Goethe University, Frankfurt
neinigr@math.uni-frankfurt.de,
sulzbach@math.uni-frankfurt.de

CP17

Lsh-Preserving Functions and Their Applications

Locality sensitive hashing (LSH) is a key algorithmic tool that is widely used both in theory and practice. An important goal in the study of LSH is to understand which similarity functions admit an LSH, i.e., are LSHable. In this paper we focus on the class of transformations such that given any similarity that is LSHable, the transformed similarity will continue to be LSHable. We show a tight characterization of all such LSH-preserving transformations: they are precisely the probability generating functions, up to scaling. As a concrete application of this result, we study which set similarity measures are LSHable. We obtain a complete characterization of similarity measures between two sets A and B that are ratios of two linear functions of $|A \cap B|, |A \Delta B|, |A \cup B|$: such a measure is LSHable if and only if its corresponding distance is a metric. This result generalizes the well-known LSH for the Jaccard set similarity, namely, the minwise-independent permutations, and obtains LSHs for many set similarity measures that are used in practice. Using our main result, we obtain a similar characterization for set similarities involving radicals.

Flavio Chierichetti
Cornell University
Cornell University
flavio@chierichetti.name

Ravi Kumar
Yahoo! Research
ravikuma@yahoo-inc.com

CP17

Tight Bounds on the Maximum Size of a Set of Permutations with Bounded VC-Dimension

The *VC-dimension* of a family \mathcal{P} of n -permutations is the largest integer k such that the set of restrictions of the permutations in \mathcal{P} on some k -tuple of positions is the set of all $k!$ permutation patterns. Let $r_k(n)$ be the maximum size of a set of n -permutations with VC-dimension k . We show slightly superexponential upper and lower bounds on $r_k(n)$, in particular, the bounds involve the inverse Ackermann function. The proofs use several results from the area of Davenport–Schinzel sequences and their generalizations.

Josef Cibulka
Department of Applied Mathematics
Charles University
cibulka@kam.mff.cuni.cz

Jan Kyncl
Department of Applied Mathematics
Charles University, Prague
kyncl@kam.mff.cuni.cz

CP17

A Linear Time Algorithm for Seeds Computation

A seed in a word is a relaxed version of a period. We show a linear time algorithm computing a compact representation of all the seeds of a word, in particular, the shortest seed. Thus, we solve an open problem stated in the survey by Smyth (2000) and improve upon a previous over 15-year old $O(n \log n)$ algorithm by Iliopoulos, Moore and Park (1996). Our approach is based on combinatorial relations between seeds and a variant of the LZ-factorization (used here for the first time in context of seeds).

Tomasz Kociumaka, Marcin Kubica, Jakub Radoszewski
University of Warsaw
kociumaka@mimuw.edu.pl, kubica@mimuw.edu.pl,
jrad@mimuw.edu.pl

Wojciech Rytter
University of Warsaw
Nicolaus Copernicus University in Toruń
rytter@mimuw.edu.pl

Tomasz Waleń
University of Warsaw
International Institute of Molecular and Cell Biology
walen@mimuw.edu.pl

CP17

Top-K Document Retrieval in Optimal Time and Linear Space

We describe a data structure that uses $O(n)$ -word space and reports k most relevant documents that contain a query pattern P in optimal $O(|P| + k)$ time. Our construction supports an ample set of important relevance measures, such as the frequency of P in a document and the minimal distance between two occurrences of P in a document. We show how to reduce the space of the data structure from $O(n \log n)$ bits to $O(n(\log \sigma + \log D + \log \log n))$ bits, where σ is the alphabet size and D is the total number of documents.

Gonzalo Navarro, Yakov Nekrich

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

CP18 Parallelism and Time in Hierarchical Self-Assembly

We show three results in the hierarchical tile assembly model. First, it is possible to assemble an $n \times n$ square from the optimal $O(\log n / \log \log n)$ tile types using nearly optimal hierarchical parallelism, in the sense that each producible assembly requires at most $O(\log^2 n)$ attachment events to become the final $n \times n$ square. Second, we develop a model of time complexity of hierarchical systems based on stochastic chemical kinetics and show that, surprisingly, the highly parallel construction requires n^2 time. In fact, all “partial order” systems (encompassing most known constructions) require time $\Omega(n)$ to assemble any shape of diameter n . Finally, we show that a non-partial order system is able to assemble a rectangle of diameter n in time $O(n^{4/5} \log n)$.

David Doty
 California Institute of Technology
 pexatus@gmail.com

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

CP18 Testing Odd-Cycle-Freeness in Boolean Functions

A function $f : F_2^n \rightarrow \{0, 1\}$ is *odd-cycle-free* if there are no $x_1, \dots, x_k \in F_2^n$ with k an odd integer such that $f(x_1) = \dots = f(x_k) = 1$ and $x_1 + \dots + x_k = 0$. We show that one can distinguish odd-cycle-free functions from those ϵ -far from being odd-cycle-free by making $\text{poly}(1/\epsilon)$ queries to an evaluation oracle. We give two proofs of this result, each shedding light on a different connection between testability of properties of Boolean functions and of dense graphs.

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

Elena Grigorescu, Prasad Raghavendra, Asaf Shapira
 Georgia Tech
 elena@cc.gatech.edu, raghavendra@cc.gatech.edu,
 asafico@math.gatech.edu

CP18 Networks Cannot Compute Their Diameter in Sub-linear Time

We study the problem of computing the diameter of a network in a distributed way. Here, in each synchronous round, each node can transmit a different (but short) message to each neighbor. We provide an $\tilde{\Omega}(n)$ lower bound for the number of communication rounds needed in a n -node network even if the diameter is a small constant. We show that a $(3/2 - \epsilon)$ -approximation of the diameter requires $\tilde{\Omega}(\sqrt{n} + D)$ rounds. Furthermore we prove a girth-approximation lower bound.

Silvio Frischknecht, Stephan Holzer, Roger Wattenhofer
 ETH Zurich
 fsilvio@ethz.ch, stholzer@ethz.ch, wattenhofer@ethz.ch

CP18 A Near-Optimal Sublinear-Time Algorithm for Approximating the Minimum Vertex Cover Size

We give a nearly optimal sublinear-time algorithm for approximating the size of a minimum vertex cover in a graph G . The algorithm may query the degree $\deg(v)$ of any vertex v of its choice, and for each $1 \leq i \leq \deg(v)$, it may ask for the i^{th} neighbor of v . Letting $\text{VC}_{\text{opt}}(G)$ denote the minimum size of vertex cover in G , the algorithm outputs, with high constant success probability, an estimate $\widehat{\text{VC}}(G)$ such that $\text{VC}_{\text{opt}}(G) \leq \widehat{\text{VC}}(G) \leq 2\text{VC}_{\text{opt}}(G) + \epsilon n$, where ϵ is a given additive approximation parameter. We refer to such an estimate as a $(2, \epsilon)$ -estimate. The query complexity and running time of the algorithm are $\tilde{O}(\bar{d} \cdot \text{poly}(1/\epsilon))$, where \bar{d} denotes the average vertex degree in the graph. The best previously known sublinear algorithm, of Yoshida et al. (*STOC 2009*), has query complexity and running time $O(d^4/\epsilon^2)$, where d is the maximum degree in the graph. Given the lower bound of $\Omega(\bar{d})$ (for constant ϵ) for obtaining such an estimate (with any constant multiplicative factor) due to Parnas and Ron (*TCS 2007*), our result is nearly optimal. In the case that the graph is dense, that is, the number of edges is $\Theta(n^2)$, we consider another model, in which the algorithm may ask, for any pair of vertices u and v , whether there is an edge between u and v . We show how to adapt the algorithm that uses neighbor queries to this model and obtain an algorithm that outputs a $(2, \epsilon)$ -estimate of the size of a minimum vertex cover whose query complexity and running time are $\tilde{O}(n) \cdot \text{poly}(1/\epsilon)$.

Krzysztof Onak
 Carnegie Mellon University
 konak@cs.cmu.edu

Dana Ron, Michal Rosen
 Tel Aviv University
 danar@eng.tau.ac.il, micalros@post.tau.ac.il

Ronitt Rubinfeld
 CSAIL, MIT and Tel Aviv University
 ronitt@csail.mit.edu

CP18 Space-Efficient Local Computation Algorithms

Recently Rubinfeld et al. (*ICS 2011*, pp. 223–238) proposed a new model of sublinear algorithms called *local computation algorithms*. In this model, a computation problem F may have more than one legal solution and each of them consists of many bits. The local computation algorithm for F should answer in an online fashion, for any index i , the i^{th} bit of some legal solution of F . Further, all the answers given by the algorithm should be consistent with at least one solution of F . In this work, we continue the study of local computation algorithms. In particular, we develop a technique which under certain conditions can be applied to construct local computation algorithms that run not only in polylogarithmic time but also in polylogarithmic *space*. Moreover, these local computation algorithms are easily parallelizable and can answer all parallel queries consistently. Our main technical tools are pseudorandom numbers with bounded independence and the theory of branching processes.

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

Ronitt Rubinfeld
CSAIL, MIT and Tel Aviv University
ronitt@csail.mit.edu

Shai Vardi
Tel Aviv University
shaivardi@gmail.com

Ning Xie
CSAIL, MIT
ningxie@csail.mit.edu

CP19 Simple and Practical Algorithm for Sparse Fourier Transform

We consider the sparse Fourier transform problem: given a complex vector x of length n , and a parameter k , estimate the k largest (in magnitude) coefficients of the Fourier transform of x . The problem is of key interest in several areas, including signal processing, audio/image/video compression, and learning theory. We propose a new algorithm for this problem. The algorithm leverages techniques from digital signal processing, notably Gaussian and Dolph-Chebyshev filters. Unlike the typical approach to this problem, our algorithm is not *iterative*. That is, instead of estimating “large” coefficients, subtracting them and recursing on the remainder, it identifies and estimates the k largest coefficients in “one shot”, in a manner akin to sketching/streaming algorithms. The resulting algorithm is structurally simpler than its predecessors. As a consequence, we are able to extend considerably the range of sparsity, k , for which the algorithm is faster than FFT, both in theory and practice.

Piotr Indyk
Massachusetts Institute of Technology
indyk@mit.edu

Haitham Hassani, Dina Katabi, Eric Price
MIT
haithamh@mit.edu, dina@csail.mit.edu, eprice@mit.edu

CP19 Sparsers Johnson-Lindenstrauss Transforms

The Johnson-Lindenstrauss lemma states that for any $0 < \epsilon, \delta < 1/2$ and for any positive integer d , there exists a distribution $\mathcal{D}_{\epsilon, \delta}$ over linear mappings from R^d into R^k for $k = O(\epsilon^{-2} \log(1/\delta))$ so that for all $x \in R^d$ with $\|x\|_2 = 1$, $\Pr_{S \sim \mathcal{D}_{\epsilon, \delta}}[|\|Sx\|_2^2 - 1| > \epsilon] < \delta$. We give two different con-

structions of such a distribution, both of which are supported only on sparse embeddings: for every embedding matrix in the supports of our distributions, each column has only $s = O(\epsilon k)$ non-zero entries. These are the first distributions which provide column sparsity asymptotically less than k regardless of how ϵ and δ are related.

Jelani Nelson
Princeton University
minilek@princeton.edu

Daniel Kane
Stanford University
dankane@stanford.edu

CP19 Optimal Column-Based Low-Rank Matrix Reconstruction

We prove that for any real-valued matrix $X \in R^{m \times n}$, and positive integers $r \geq k$, there is a subset of r columns of X such that projecting X onto their span gives a $\sqrt{\frac{r+1}{r-k+1}}$ -approximation to best rank- k approximation of X in Frobenius norm. We show that the trade-off we achieve between the number of columns and the approximation ratio is optimal up to lower order terms. Furthermore, there is a deterministic algorithm to find such a subset of columns that runs in $O(rnm^\omega \log m)$ arithmetic operations where ω is the exponent of matrix multiplication. We also give a faster randomized algorithm that runs in $O(rnm^2)$ arithmetic operations.

Ali K. Sinop, Venkatesan Guruswami
Carnegie Mellon University
asinop@cs.cmu.edu, guruswami@cmu.edu

CP19 Sublinear Time, Measurement-Optimal, Sparse Recovery For All

We give an algorithm for the approximate sparse recovery problem in ℓ_1 norm. Our algorithm makes a small number of measurements of a noisy vector of length N with at most k large entries and recovers those heavy hitters up to the factor $(1 + \epsilon)$ in ℓ_1 norm. For any positive integer ℓ , our approach uses time $O(\ell^5 \epsilon^{-3} k (N/k)^{1/\ell})$ and uses $O(\ell^8 \epsilon^{-3} k \log(N/k))$ measurements, with access to a data structure requiring space and preprocessing time $O(\ell N k^{0.2} / \epsilon)$.

Martin Strauss
Mathematics and EECS
University of Michigan
martinjs@umich.edu

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

CP20 Race to Idle: New Algorithms for Speed Scaling with a Sleep State

We study an energy conservation problem where a variable-speed processor is equipped with a sleep state. Executing jobs at high speeds and then setting the processor asleep is an approach that can lead to further energy savings compared to standard dynamic speed scaling. We consider classical deadline-based scheduling, i.e. each job is specified by a release time, a deadline and a processing volume. For general convex power functions, an offline 2-approximation algorithm was devised by [Irani, Shukla and Gupta, Algorithms for Power Savings, ACM TALG, 2007]. Roughly speaking, the algorithm schedules jobs at a critical speed s_{crit} that yields the smallest energy consumption while jobs are processed. We investigate the offline setting of speed scaling with a sleep state. First we prove NP-hardness of the optimization problem. Additionally, we develop lower bounds, for general convex power functions: No algorithm that constructs s_{crit} -schedules, which execute jobs at speeds of at least s_{crit} , can achieve an approximation factor smaller than 2. Furthermore, no algorithm that minimizes the energy expended for process-

ing jobs can attain an approximation ratio smaller than 2. We then present an algorithmic framework for designing good approximation algorithms. For general convex power functions, we derive an approximation factor of $4/3$. For power functions $P(s) = \beta s^\alpha + \gamma$, where s is the processor speed, we obtain an approximation of $137/117 < 1.171$. We finally show that our framework yields the best approximation guarantees for the class of s_{crit} -schedules. For general convex power functions, we give another 2-approximation algorithm. For functions $P(s) = \beta s^\alpha + \gamma$, we present tight upper and lower bounds on the best possible approximation factor. The ratio is exactly $eW_{-1}(-e^{-1-1/e})/(eW_{-1}(-e^{-1-1/e}) + 1) < 1.211$, where W_{-1} is the lower branch of the Lambert W function.

Susanne Albers
Humboldt-Universität zu Berlin
albers@informatik.hu-berlin.de

Antonios Antoniadis
Humboldt-Universität zu Berlin
antoniad@informatik.hu-berlin.de

CP20

Resource Augmentation for Weighted Flow-Time Explained by Dual Fitting

We propose a general dual-fitting technique for analyzing on-line scheduling algorithms in the unrelated machines setting where the objective function involves weighted flow-time, and we allow the machines of the on-line algorithm to have $(1 + \varepsilon)$ -extra speed than the offline optimum (the so-called speed augmentation model). We propose that one can often analyze such algorithms by looking at the dual (or Lagrangian dual) of the linear (or convex) program for the corresponding scheduling problem, and finding a feasible dual solution as the on-line algorithm proceeds. We apply these ideas to the problems of minimizing (i) weighted flow-time, (ii) weighted ℓ_k norm of weighted flow-time, and (iii) sum of energy and weighted flow-time.

Amit Kumar, S Anand, Naveen Garg
IIT Delhi
amitk@cse.iitd.ac.in, anand@cse.iitd.ac.in,
naveen@cse.iitd.ac.in

CP20

Online Scheduling with General Cost Functions

We consider a general online scheduling problem with the objective of minimizing $\sum_j w_j g(F_j)$, where w_j is the weight/importance of job J_j , F_j is the flow time of the job in the schedule, and g is an arbitrary non-decreasing cost function. We show that the scheduling algorithm Highest Density First (HDF) is $(2 + \varepsilon)$ -speed $O(1)$ -competitive for all cost functions g *simultaneously*. We give lower bounds that show the HDF algorithm and this analysis are essentially optimal.

Sungjin Im, Benjamin Moseley
University of Illinois
im3@illinois.edu, bmosele2@illinois.edu

Kirk Pruhs
University of Pittsburgh
kirk@cs.pitt.edu

CP20

Scheduling Heterogeneous Processors Isn't As Easy As You Think

We consider preemptive online scheduling algorithms to minimize the total weighted/unweighted flow time plus energy for speed-scalable heterogeneous multiprocessors. We show that the well-known priority scheduling algorithms are not $O(1)$ -speed $O(1)$ -competitive for the objective of weighted flow. We then show that a variation of the non-clairvoyant algorithm *Late Arrival Processor Sharing* coupled with a non-obvious speed scaling algorithm is scalable for the objective of unweighted flow plus energy on speed-scalable multiprocessors.

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

Sungjin Im
University of Illinois
im3@illinois.edu

Ravishankar Krishnaswamy
Computer Science Department
Carnegie Mellon University
ravishan@cs.cmu.edu

Benjamin Moseley
University of Illinois
bmosele2@illinois.edu

Kirk Pruhs
University of Pittsburgh
kirk@cs.pitt.edu

CP21

A Faster Algorithm to Recognize Even-Hole-Free Graphs

We study the problem of determining whether an n -node m -edge graph has an *even hole*, i.e., an induced simple cycle consisting of an even number of nodes. Conforti, Cornuéjols, Kapoor, and Vušković gave the first polynomial-time algorithm for the problem, which runs in $O(n^{40})$ time. Later, Chudnovsky, Kawarabayashi, and Seymour reduced the running time to $O(n^{31})$. The best previously known algorithm for the problem, due to da Silva and Vušković, runs in $O(n^{19})$ time. In this paper, we solve the problem in time $O(n^{11})$.

Hsien-Chih Chang, Hsueh-I Lu
National Taiwan University
hcsoso@gmail.com, hil@csie.ntu.edu.tw

CP21

Separating Stable Sets in Claw-Free Graphs Via Padberg-Rao and Compact Linear Programs

In this paper, we provide the first linear programming formulations for the stable set problem in claw-free graphs. We then exploit one of those extended formulations and propose a new polytime algorithm for solving the separation problem for the stable set polytope of claw-free graphs in the original space. This routine combines a separation algorithm for the matching polytope due to Padberg and Rao and the solution of (moderate size) compact linear

programs.

Yuri Faenza

Università degli Studi di Padova
yuri.faenza@gmail.com

Gianpaolo Oriolo
Università di Tor Vergata
oriolo@disp.uniroma2.it

Gautier Stauffer
Université de Bordeaux - INRIA
gautier.stauffer@gmail.com

CP21

Global Minimum Cuts in Surface Embedded Graphs

We give a deterministic algorithm to find the minimum cut in a surface-embedded graph in near-linear time. Given an undirected graph embedded on an orientable surface of genus g , our algorithm computes the minimum cut in $g^{O(g)}n \log \log n$ time, matching the running time of the fastest algorithm known for planar graphs, due to Lacki and Sankowski, for any constant g . Indeed, our algorithm calls Lacki and Sankowski's recent $O(n \log \log n)$ time planar algorithm as a subroutine.

Previously, the best time bounds known for this problem followed from two algorithms for general sparse graphs: a *randomized* algorithm of Karger that runs in $O(n \log^3 n)$ time and succeeds with high probability, and a deterministic algorithm of Nagamochi and Ibaraki that runs in $O(n^2 \log n)$ time. We can also achieve a deterministic $g^{O(g)}n^2 \log \log n$ time bound by repeatedly applying the best known algorithm for minimum (s, t) -cuts in surface graphs. The bulk of our work focuses on the case where the dual of the minimum cut splits the underlying surface into multiple components with positive genus.

Jeff Erickson
Computer Science Department
University of Illinois, Urbana-Champaign
jeffe@cs.uiuc.edu

Kyle Fox, Amir Nayyeri
University of Illinois, Urbana-Champaign
kylefox2@illinois.edu, nayyeri2@illinois.edu

CP21

Competitive Routing in the Half- θ_6 -Graph

We present a deterministic local routing scheme that is guaranteed to find a path between any pair of vertices in a half- θ_6 -graph whose length is at most $5/\sqrt{3} = 2.886\dots$ times the Euclidean distance between the pair of vertices. The half- θ_6 -graph is identical to the Delaunay triangulation where the empty region is an equilateral triangle. Moreover, we show that no local routing scheme can achieve a better competitive spanning ratio thereby implying that our routing scheme is optimal. This is somewhat surprising because the spanning ratio of the half- θ_6 -graph is 2. Since every triangulation can be embedded in the plane as a half- θ_6 -graph using $O(\log n)$ bits per vertex coordinate via Schnyder's embedding scheme (SODA 1990), our result provides a competitive local routing scheme for

every such embedded triangulation.

Prosenjit Bose
Carleton University
jit@scs.carleton.ca

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

André Van Renssen, Sander Verdonschot
Carleton University
avrensse@connect.carleton.ca,
sverdons@connect.carleton.ca

CP22

Learning K-Modal Distributions Via Testing

A k -modal probability distribution over $\{1, \dots, n\}$ is one whose histogram has at most k "peaks" and "valleys." We give an efficient algorithm for learning such a distribution to variation distance ϵ , using independent random samples, that runs in time $\text{poly}(k, \log(n), 1/\epsilon)$. For $k \leq \tilde{O}(\sqrt{\log n})$, the number of samples we use is very close to optimal. A novel feature of our learning algorithm is that it uses a new *property testing* algorithm as a key subroutine.

Constantinos Daskalakis
MIT
costis@csail.mit.edu

Ilias Diakonikolas
UC Berkeley
ilias@cs.berkeley.edu

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

CP22

Data Reduction for Weighted and Outlier-Resistant Clustering

Statistical data frequently includes outliers; these can distort the results of estimation procedures and optimization problems. However, there are few algorithmic results about outlier-resistant clustering. We consider several formulations including "k-median with outliers", k-clustering w.r.t. M -estimators such as the Tukey or Huber loss functions, and "k-line-median". The common challenge in these optimizations is data reduction for "weighted k-median". We solve this task, which was previously solved only in one dimension.

Dan Feldman
California Institute of Technology
dannfyf@caltech.edu

Leonard Schulman
Caltech
schulman@caltech.edu

CP22

Local Homology Transfer and Stratification Learning

The objective of this paper is to show that point cloud data

can under certain circumstances be clustered by strata in a plausible way. For our purposes, we consider a stratified space to be a collection of manifolds of different dimensions which are glued together in a locally trivial manner inside some Euclidean space. To adapt this abstract definition to the world of noise, we first define a multi-scale notion of stratified spaces, providing a stratification at different scales which are indexed by a radius parameter. We then use methods derived from kernel and cokernel persistent homology to cluster the data points into different strata. We prove a correctness guarantee for this clustering method under certain topological conditions. We then provide a probabilistic guarantee for the clustering for the point sample setting we provide bounds on the minimum number of sample points required to state with high probability which points belong to the same strata. Finally, we give an explicit algorithm for the clustering.

Paul Bendich
Duke University and IST Austria
bendich@math.duke.edu

Bei Wang
Scientific Computing and Imaging Institute
University of Utah
beiwang@sci.utah.edu

Sayan Mukherjee
Duke University
sayan@stat.duke.edu

CP22

A Near-Linear Algorithm for Projective Clustering Integer Points

We consider the problem of projective clustering in Euclidean spaces of non-fixed dimension. Here, we are given a set P of n points in \mathbf{R}^m and integers $j \geq 1$, $k \geq 0$, and the goal is to find j k -subspaces so that the sum of the distances of each point in P to the nearest subspace is minimized. Observe that this is a shape fitting problem where we wish to find the best fit in the L_1 sense. Here we will treat the number j of subspaces we want to fit and the dimension k of each of them as constants. We consider instances of projective clustering where the point coordinates are integers of magnitude polynomial in m and n . Our main result is a randomized algorithm that for any $\epsilon > 0$ runs in time $O(mn \text{ polylog}(mn))$ and outputs a solution that with high probability is within $(1 + \epsilon)$ of the optimal solution. To obtain this result, we show that the fixed dimensional version of the above projective clustering problem has a small *coreset*. We do that by observing that in a fairly general sense, shape fitting problems that have small coresets in the L_∞ setting also have small coresets in the L_1 setting, and then exploiting an existing construction for the L_∞ setting. This observation seems to be quite useful for other shape fitting problems as well, as we demonstrate by constructing the first ‘regular’ coreset for the circle fitting problem in the plane.

Kasturi Varadarajan, Xin Xiao
The University of Iowa
kasturi-varadarajan@uiowa.edu, xin-xiao@uiowa.edu

CP23

An $O(n^2)$ Time Algorithm for Alternating Büchi Games

Computing the winning set for Büchi objectives in alter-

nating graph games is a central problem in verification. The long standing best known bound for the problem is $\tilde{O}(n \cdot m)$, where n and m are the number of vertices and edges. We are the first to break the $\tilde{O}(n \cdot m)$ bound by presenting a new technique that solves the problem in $O(n^2)$ time. We show how to maintain the winning set under a sequence of edge insertions or deletions in $O(n)$ amortized time.

Krishnendu Chatterjee
IST Austria
krish.chat@gmail.com

Monika Henzinger
University of Vienna, Austria
monika.henzinger@univie.ac.at

CP23

Efficient Algorithms for Maximum Weight Matchings in General Graphs with Small Edge Weights

Let $G = (V, E)$ be a graph with positive integral edge weights. Our problem is to find a matching of maximum weight in G . We present a simple iterative algorithm for this problem that uses a maximum cardinality matching algorithm as a subroutine. Using the current fastest maximum cardinality matching algorithms, we solve the maximum weight matching problem in $O(W\sqrt{nm} \log_n(n^2/m))$ time, or in $O(Wn)$ time with high probability, where $n = |V|$, $m = |E|$, W is the largest edge weight, and < 2.376 is the exponent of matrix multiplication. In relatively dense graphs, our algorithm performs better than all existing algorithms with $W = o(\log^{1.5} n)$. Our technique hinges on exploiting Edmonds’ matching polytope and its dual.

Chien Chung Huang
Humboldt Universität zu Berlin
villars@informatik.hu-berlin.de

Telikepalli Kavitha
Tata Institute of Fundamental Research, Mumbai, India
kavitha.telikepalli@gmail.com

CP23

List-Coloring Graphs Without Subdivisions and Without Immersions

We give additive approximation algorithms for list-coloring within $3.5(k + 1)$ of the list-chromatic number for graphs without K_k as a subdivision, and within $1.5(k - 1)$ of the list-chromatic number for graphs without K_k as an immersion. Clearly our results give rise to additive approximation algorithms for graph-coloring of graphs without K_k as a subdivision (in fact, we shall give an additive approximation algorithm within $2.5(k + 1)$ of the chromatic number) and K_k as an immersion, too. These are the first results in this direction (in fact, these are the first results concerning list-coloring graphs without fixed graph as a subdivision or as an immersion, except for the known upper bound results) and extend the result by Kawarabayashi, Demaine and Hajiaghayi (SODA’09) concerning the additive approximation algorithm for list-coloring graphs without K_k as a minor. We also discuss how our results are related to the famous Hájos’ conjecture and Hadwiger’s conjecture.

Ken-ichi Kawarabayashi
National Institute of Informatics, Japan

k_keniti@nii.ac.jp

Yusuke Kobayashi
University of Tokyo
kobayashi@mist.i.u-tokyo.ac.jp

CP23

A Scaling Algorithm for Maximum Weight Matching in Bipartite Graphs

Given a weighted bipartite graph, the maximum weight matching (MWM) problem is to find a set of vertex-disjoint edges with maximum weight. We present a new scaling algorithm that runs in $O(m\sqrt{n} \log N)$ time, when the weights are integers within the range of $[0, N]$. The result improves the previous bounds of $O(Nm\sqrt{n})$ by Gabow and $O(m\sqrt{n} \log(nN))$ by Gabow and Tarjan over 20 years ago. Our improvement draws ideas from a not widely known result, the primal method by Balinski and Gomory.

Ran Duan, Hsin-Hao Su
University of Michigan
EECS Department
duanran@umich.edu, hsinhao@umich.edu

CP24

Constructing High Order Elements Through Subspace Polynomials

Every finite field has many multiplicative generators. However, finding one in polynomial time is an important open problem. In fact, even finding elements of high order has not been solved satisfactorily. In this paper, we present an algorithm that for any positive integer c and prime power q , finding an element of order $\exp(\Omega(\sqrt{q^c}))$ in the finite field $\mathbf{F}_{q^{(q^c-1)/(q-1)}}$ in deterministic time $(q^c)^{O(1)}$. We also show that there are $\exp(\Omega(\sqrt{q^c}))$ many weak keys for the discrete logarithm problems in those fields with respect to certain bases.

Qi Cheng
University of Oklahoma
qcheng@ou.edu

Shuhong Gao
Department of Mathematical Sciences
Clemson University, Clemson, SC 29634, USA
sgao@clemson.edu

Daqing Wan
University of California, Irvine
dwan@math.uci.edu

CP24

Deterministic Construction of An Approximate M-Ellipsoid and Its Applications to Derandomizing Lattice Algorithms

We give a deterministic $O(\log n)^n$ -time and space algorithm for the Shortest Vector Problem (SVP) under any norm, improving on the previous best deterministic $n^{O(n)}$ -time algorithms for general norms. This approaches the $2^{O(n)}$ -time and space complexity of the randomized sieve based SVP algorithms (Arvind and Joglekar, FSTTCS 2008), first introduced by Ajtai, Kumar and Sivakumar (STOC 2001) for l_2 -SVP, and the M-ellipsoid covering based SVP algorithm of Dadush et al. (FOCS 2011). Here

we continue with the covering approach of Dadush et al., and our main technical contribution is a deterministic approximation of an M-ellipsoid for any convex body. To achieve this we exchange the M-position of a convex body by a related position, known as the minimal mean width position of the polar body. We reduce the task of computing this position to solving a semi-definite program whose objective is a certain Gaussian expectation, which we show can be approximated deterministically.

Daniel Dadush, Santosh Vempala
Georgia Institute of Technology
dndadush@gatech.edu, vempala@gatech.edu

CP24

Fast Zeta Transforms for Lattices with Few Irreducibles

We investigate fast algorithms for changing between the standard basis and an orthogonal basis of idempotents for Möbius algebras of finite lattices. We show that every lattice with v elements, n of which are nonzero and join-irreducible (or, by a dual result, nonzero and meet-irreducible), has arithmetic circuits of size $O(vn)$ for computing the zeta transform and its inverse, thus enabling fast multiplication in the Möbius algebra.

Petteri Kaski
Aalto University
petteri.kaski@aalto.fi

Andreas Björklund
Lund University
andreas.bjorklund@yahoo.se

Thore Husfeldt
IT University of Copenhagen
thore@itu.dk

Mikko Koivisto
University of Helsinki
Helsinki Institute for Information Technology HIIT
mikko.koivisto@cs.helsinki.fi

Jesper Nederlof
University of Bergen
jesper.nederlof@ii.uib.no

Pekka Parviainen
University of Helsinki
Helsinki Institute for Information Technology HIIT
pekka.parviainen@cs.helsinki.fi

CP24

Improved Output-Sensitive Quantum Algorithms for Boolean Matrix Multiplication

We present new quantum algorithms for Boolean Matrix Multiplication in both the time complexity and the query complexity settings. As far as time complexity is concerned, our results show that the product of two $n \times n$ Boolean matrices can be computed on a quantum computer in time $\tilde{O}(n^{3/2} + n\ell^{3/4})$, where ℓ is the number of non-zero entries in the product, improving over the output-sensitive quantum algorithm by Buhrman and Špalek that runs in $\tilde{O}(n^{3/2}\sqrt{\ell})$ time. This is done by constructing a quantum version of a recent algorithm by Lingas, using quantum techniques such as quantum counting to exploit the spar-

sity of the output matrix. As far as query complexity is concerned, our results improve over the quantum algorithm by Vassilevska Williams and Williams based on a reduction to the triangle finding problem. One of the main contributions leading to this improvement is the construction of a triangle finding quantum algorithm tailored especially for the tripartite graphs appearing in the reduction.

Francois Le Gall

The University of Tokyo
legall@is.s.u-tokyo.ac.jp

CP25

Traffic-Redundancy Aware Network Design

We consider network design problems for information networks where routers can replicate data but cannot alter it. This functionality allows the network to eliminate data-redundancy in traffic, thereby saving on routing costs. We consider two problems within this framework and design approximation algorithms. The first problem we study is the traffic-redundancy aware network design (RAND) problem. We are given a weighted graph over a single server and many clients. The server owns a number of different data packets and each client desires a subset of the packets; the client demand sets form a laminar set system. Our goal is to connect every client to the source via a single path, such that the collective cost of the resulting network is minimized. Here the transportation cost over an edge is its weight times the number of *distinct* packets that it carries. The second problem is a facility location problem that we call RAFL. Here the goal is to find an assignment from clients to facilities such that the total cost of routing packets from the facilities to clients (along unshared paths), plus the total cost of “producing” one copy of each desired packet at each facility is minimized. We present a constant factor approximation for the RAFL and an $O(\log P)$ approximation for RAND, where P is the total number of distinct packets. We remark that P is always at most the number of different demand sets desired or the number of clients, and is generally much smaller.

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

CP25

Approximation Algorithms for Stochastic Orienteering

In the Stochastic Orienteering problem, we are given a metric, where each node also has a job located there with some deterministic reward and a *random* size. (Think of the jobs as being chores one needs to run, and the sizes as the amount of time it takes to do the chore.) The goal is to adaptively decide which nodes to visit to maximize total expected reward, subject to the constraint that the total distance traveled plus the total size of jobs processed is at most a given budget of B . (I.e., we get reward for all those chores we finish by the end of the day). The (random) size of a job is not known until it is completely processed. Hence the problem combines aspects of both the stochastic knapsack problem with uncertain item sizes and the deterministic orienteering problem of using a limited travel time to maximize gathered rewards located at nodes. In this paper, we present a constant-factor approximation algorithm for the best non-adaptive policy for the Stochastic Orienteering problem. We also show a small

adaptivity gap—i.e., the existence of a non-adaptive policy whose reward is at least an $\Omega(1/\log \log B)$ fraction of the optimal expected reward—and hence we also get an $O(\log \log B)$ -approximation algorithm for the adaptive problem. Finally we address the case when the node rewards are also random and could be correlated with the waiting time, and give a non-adaptive policy which is an $O(\log n \log B)$ -approximation to the best adaptive policy on n -node metrics with budget B .

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

Ravishankar Krishnaswamy
Computer Science Department
Carnegie Mellon University
ravishan@cs.cmu.edu

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

R Ravi
Carnegie Mellon University
ravi@cmu.edu

CP25

Approximating Rooted Steiner Networks

The Directed Steiner Tree (DST) problem is a cornerstone problem in network design. We focus on the generalization of the problem with higher connectivity requirements. The problem with one root and two sinks is APX-hard. The problem with one root and many sinks is as hard to approximate as the directed Steiner forest problem, and the latter is well known to be as hard to approximate as the label cover problem. Utilizing previous techniques (due to others), we strengthen these results and extend them to undirected graphs. Specifically, we give an $\Omega(k^\epsilon)$ hardness bound for the rooted k -connectivity problem in undirected graphs; this addresses a recent open question of Khanna. As a consequence, we also obtain the $\Omega(k^\epsilon)$ hardness of the undirected subset k -connectivity problem. Additionally, we give a result on the integrality ratio of the natural linear programming relaxation of the directed rooted k -connectivity problem.

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

Joseph Cheriyan
Department of Combinatorics and Optimization
University of Waterloo
jcheriyan@math.uwaterloo.ca

Guylain Naves
Department of Mathematics & Statistics
McGill University
naves@math.mcgill.ca

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

CP25**A Proof of the Boyd-Carr Conjecture**

Determining the precise integrality gap for the subtour LP relaxation of the traveling salesman problem is a significant open question, with little progress made in thirty years in the general case of symmetric costs that obey triangle inequality. Boyd and Carr observe that we do not even know the worst-case upper bound on the ratio of the optimal 2-matching to the subtour LP; they conjecture the ratio is at most $10/9$. In this paper, we prove the Boyd-Carr conjecture. In the case that a fractional 2-matching has no cut edge, we can further prove that an optimal 2-matching is at most $10/9$ times the cost of the fractional 2-matching.

Anke Van Zuylen
Max Planck Institute for Informatics
anke@mpi-inf.mpg.de

Frans Schalekamp
Unaffiliated
fms9@cornell.edu

David Williamson
Cornell University
dpw@cs.cornell.edu

CP25**Matroidal Degree-Bounded Minimum Spanning Trees**

We consider the minimum spanning tree (MST) problem under the restriction that for every vertex v , the edges of the tree that are adjacent to v satisfy a given family of constraints. A famous example thereof is the classical degree-bounded MST problem, where for every vertex v , a simple upper bound on the degree is imposed. Iterative rounding/relaxation algorithms became the tool of choice for degree-constrained network design problems. A cornerstone for this development was the work of Singh and Lau, who showed that for the degree-bounded MST problem, one can find a spanning tree violating each degree bound by at most one unit and with cost at most the cost of an optimal solution that respects the degree bounds. However, current iterative rounding approaches face several limits when dealing with more general degree constraints, where several linear constraints are imposed on the edges adjacent to a vertex v . For example, when a partition of the edges adjacent to v is given and only a fixed number of elements can be chosen out of each set of the partition, current approaches might violate each of the constraints at v by a constant, instead of violating the whole family of constraints at v by at most a constant number of edges. Furthermore, previous iterative rounding approaches are not suited for degree constraints where some edges are in a super-constant number of constraints. We extend iterative rounding/relaxation approaches, both conceptually as well as in their analysis, to address these limitations. Based on these extensions, we present an algorithm for the degree-constrained MST problem where for every vertex v , the edges adjacent to v have to be independent in a given matroid. The algorithm returns a spanning tree of cost at most OPT , such that for every vertex v , it suffices to remove at most 8 edges from the spanning tree to satisfy the matroidal degree constraint at v .

Rico Zenklusen
EPFL Lausanne
ricoz@math.mit.edu

CP26**Weighted Capacitated, Priority, and Geometric Set Cover Via Improved Quasi-Uniform Sampling**

The minimum-weight set cover problem is well-known to be $O(\log n)$ -approximable, with no improvement possible in the general case. We provide a geometry-inspired algorithm whose approximation guarantee depends solely on an instance-specific combinatorial property known as shallow cell complexity (SCC). Roughly speaking, a set cover instance has low SCC if any column-induced submatrix of the corresponding element-set incidence matrix has few distinct rows. We adapt and improve Varadarajan's recent quasi-uniform random sampling method for weighted geometric covering problems.

Jochen Koenemann
University of Waterloo
jochen@math.uwaterloo.ca

Elyot Grant
MIT
elyotgrant@gmail.com

Malcolm Sharpe, Timothy Chan
University of Waterloo
sharpe.malcolm@gmail.com,
tmchan@lucky.math.uwaterloo.ca

timothy

CP26**Expanders Are Universal for the Class of All Spanning Trees**

We show that every n -vertex graph satisfying certain natural expansion properties is universal for the class of all bounded-degree spanning trees, i.e., contains a copy of each n -vertex tree with maximum degree bounded by some function of n . Since random graphs are known to be good expanders, this implies that binomial random graphs and random regular graphs with sufficiently large (average) degree are asymptotically almost surely universal for the class of all bounded-degree spanning trees.

Daniel Johannsen, Michael Krivelevich
Tel Aviv University
johannse@tau.ac.il, krivelev@post.tau.ac.il

Wojciech Samotij
Tel Aviv University / Trinity College
samotij@post.tau.ac.il

CP26**Directed Nowhere Dense Classes of Graphs**

Many natural computational problems on graphs such as finding dominating or independent sets of a certain size are well known to be intractable, both in the classical sense as well as in the framework of parameterized complexity. Much work therefore has focussed on exhibiting restricted classes of graphs on which these problems become tractable. While in the case of undirected graphs, there is a rich structure theory which can be used to develop tractable algorithms for these problems on large classes of undirected graphs, such a theory is much less developed for directed graphs. Many attempts to identify structure properties of directed graphs tailored towards algorithmic applications have focussed on a directed analogue of undirected tree-width. These attempts have proved to be successful in the development of algorithms for linkage prob-

lems but none of the existing width-measures allows for tractable solutions to important problems such as dominating sets and many other related problems. In this paper we take a radically different approach to identifying classes of directed graphs where domination and other problems become tractable. More specifically, we introduce the concept of shallow directed minors and based on this a new classification of classes of directed graphs which is diametric to existing directed graph decompositions and directed width measures proposed in the literature. We then study in depth one type of classes of directed graphs which we call *nowhere crownful*. The classes are very general and yet we are able to show that problems such as directed dominating set and many others become tractable on nowhere crownful classes of directed graphs. This is of particular interest as these problems are not tractable on any existing digraph measure for sparse classes.

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

Siamak Tazari
 MIT
 stazari@mit.edu

CP26 Submodular Functions Are Noise Stable

We show that all non-negative submodular functions have high *noise-stability*. As a consequence, we obtain a polynomial-time learning algorithm for this class with respect to any product distribution on $\{-1, 1\}^n$ (for any constant accuracy parameter ϵ). Our algorithm also succeeds in the agnostic setting. Previous work on learning submodular functions required either query access or strong assumptions about the types of submodular functions to be learned (and did not hold in the agnostic setting). Additionally we give simple algorithms that efficiently release differentially private answers to all Boolean conjunctions and to all halfspaces with constant average error, subsuming and improving recent work due to Gupta, Hardt, Roth and Ullman (STOC 2011).

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

Adam Klivans, Pravesh Kothari, Homin Lee
 Department of Computer Science
 University of Texas at Austin
 klivans@cs.utexas.edu, kothari@cs.utexas.edu,
 homin@cs.utexas.edu

CP26 Bidimensionality and Geometric Graphs

Bidimensionality theory was introduced by Demaine et al. [*JACM 2005*] as a framework to obtain algorithmic results for hard problems on minor closed graph classes. The theory has been successfully applied to yield subexponential time parameterized algorithms, EPTASs and linear kernels for many problems on families of graphs excluding a fixed graph H as a minor. classes of graphs which are not minor closed, but instead exhibit a geometric structure. In this paper we use several of the key ideas from Bidimensionality to give a new generic approach to design (E)PTASs and subexponential time parameterized algorithms for prob-

lems on classes of graphs which are not minor closed, but instead exhibit a geometric structure. In particular we present (E)PTASs and subexponential time parameterized algorithms for FEEDBACK VERTEX SET, VERTEX COVER, CONNECTED VERTEX COVER, DIAMOND HITTING SET, on map graphs and unit disk graphs, and for CYCLE PACKING and MINIMUM-VERTEX FEEDBACK EDGE SET on unit disk graphs. Our results are based on the recent decomposition theorems proved by Fomin et al. in [*SODA 2011*] and novel grid-excluding theorems in unit disc and map graphs without large cliques. We also show that our approach can not be extended in its full generality to more general classes of geometric graphs, such as intersection graphs of unit balls in \mathbf{R}^d , $d \geq 3$. Additionally, we show that the decomposition theorems which our approach is based on fail for disk graphs and that therefore any extension of our results to disk graphs would require new algorithmic ideas. On the other hand, we prove that our EPTASs and subexponential time algorithms for VERTEX COVER and CONNECTED VERTEX COVER carry over both to disk graphs and to unit-ball graphs in \mathbf{R}^d for every fixed d .

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

Daniel Lokshantov
 UCSD
 daniello@ii.uib.no

Saket Saurabh
 IMS
 saket@imsc.res.in

CP27 The Mixing Time of the Newman-Watts Small World

“Small worlds’ are large systems in which any given node has only a few connections to other points, but possessing the property that all pairs of points are connected by a short path, typically logarithmic in the number of nodes. The use of random walks for sampling a uniform element from a large state space is by now a classical technique; to prove that such a technique works for a given network, a bound on the mixing time is required. However, little detailed information is known about the behaviour of random walks on small-world networks, though many predictions can be found in the physics literature. The principal contribution of this paper is to show that for a famous small-world random graph model known as the Newman–Watts small world, the mixing time is of order $\log^2 n$. This confirms a prediction of Richard Durrett, who proved a lower bound of order $\log^2 n$ and an upper bound of order $\log^3 n$.

Louigi Addario-Berry, Tao Lei
 McGill University
 louigi@gmail.com, tao.lei@mail.mcgill.ca

CP27 Random Walks, Electric Networks and The Transience Class Problem of Sandpiles

The Abelian Sandpile Model is a discrete diffusion process defined on graphs (Dhar [?], Dhar et al. [?]) which serves as the standard model of *self-organized criticality*. The transience class of a sandpile is defined as the maximum

number of particles that can be added without making the system recurrent ([?]). We develop the theory of discrete diffusions in contrast to continuous harmonic functions on graphs and establish deep connections between standard results in the study of random walks on graphs and sandpiles on graphs. Using this connection and building other necessary machinery we improve the main result of Babai and Gorodetzky (SODA 2007,[?]) of the bound on the transience class of an $n \times n$ grid, from $O(n^{30})$ to $O(n^7)$. Proving that the transience class is small validates the general notion that for most natural phenomenon, the time during which the system is transient is small. In addition, we use the machinery developed to prove a number of auxiliary results. We exhibit an equivalence between two other tessellations of plane, the honeycomb and triangular lattices. We give general upper bounds on the transience class as a function of the number of edges to the sink. Further, for planar sandpiles we derive an explicit algebraic expression which provably approximates the transience class of G to within $O(|E(G)|)$. This expression is based on the spectrum of the Laplacian of the dual of the graph G . We also show a lower bound of $\Omega(n^3)$ on the transience class on the grid improving the obvious bound of $\Omega(n^2)$.

Ayush Choure, Sundar Vishwanathan
IIT Bombay
ayush@cse.iitb.ac.in, sundar@cse.iitb.ac.in

CP27 Rumor Spreading and Vertex Expansion

We study the relation between the rate at which rumors spread throughout a graph and the vertex expansion of the graph. We consider the standard rumor spreading protocol where every node chooses a random neighbor in each round and the two nodes exchange the rumors they know. For any n -node graph with vertex expansion α , we show that this protocol spreads a rumor from a single node to all other nodes in $O(\alpha^{-1} \log^2 n \sqrt{\log n})$ rounds with high probability. Further, we construct graphs for which $\Omega(\alpha^{-1} \log^2 n)$ rounds are needed. Our results complement a long series of works that relate rumor spreading to edge-based notions of expansion, resolving one of the most natural questions on the connection between rumor spreading and expansion.

George Giakkoupis
University of Calgary
ggiakk@gmail.com

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

CP27 Information Dissemination Via Random Walks in D-Dimensional Space

We study a natural information dissemination problem for multiple mobile agents in a bounded Euclidean space. Agents are placed uniformly at random in the d -dimensional space $\{-n, \dots, n\}^d$ at time zero, and one of the agents holds a piece of information to be disseminated. All the agents then perform independent random walks over the space, and the information is transmitted from one agent to another if the two agents are sufficiently close. We wish to bound the total time before all agents receive the information (with high probability). Our work extends Pettarin et al's work (Infectious random walks, arXiv:1007.1604v2, 2011), which solved the problem for

$d \leq 2$. We present tight bounds up to polylogarithmic factors for the case $d = 3$. (While our results extend to higher dimensions, for space and readability considerations we provide only the case $d = 3$ here.) Our results show the behavior when $d \geq 3$ is qualitatively different from the case $d \leq 2$. In particular, as the ratio between the volume of the space and the number of agents varies, we show an interesting phase transition for three dimensions that does not occur in one or two dimensions.

Henry Lam
Boston University
khlam@math.bu.edu

Zhenming Liu, Michael Mitzenmacher
Harvard University
zliu@fas.harvard.edu, michaelm@eecs.harvard.edu

Xiaorui Sun
Columbia University
xiaoruisun@cs.columbia.edu

Yajun Wang
Microsoft Research Asia
yajunw@microsoft.com

CP27 Ultra-Fast Rumor Spreading in Social Networks

We analyze the popular push-pull protocol for spreading a rumor in networks. Initially, a single node knows of a rumor. In each succeeding round, every node chooses a random neighbor, and the two nodes share the rumor if one of them is already aware of it. We present the first theoretical analysis of this protocol on random graphs that have a power law degree distribution. Our main findings reveal a striking dichotomy in the performance of the protocol that depends on the exponent of the power law. More specifically, we show that if $2 < \beta < 3$, then the rumor spreads to almost all nodes in $\Theta(\log \log n)$ rounds with high probability. On the other hand, if $\beta > 3$, then $\Omega(\log n)$ rounds are necessary. We also investigate the asynchronous version of the push-pull protocol, where the nodes do not operate in rounds, but exchange information according to a Poisson process with rate 1. Surprisingly, we are able to show that, if $2 < \beta < 3$, the rumor spreads even in *constant* time, which is much smaller than the typical distance of two nodes. To the best of our knowledge, this is the first result that establishes a gap between the synchronous and the asynchronous protocol.

Konstantinos Panagiotou
Max-Planck-Institute for Informatics
kpanagio@mpi-inf.mpg.de

Konstantinos Panagiotou
Department of Computer Science
ETH Zurich
kpanagio@mpi-inf.mpg.de

Nikolaos Fountoulakis
University of Birmingham
n.fountoulakis@bham.ac.uk

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

CP28**An $O(\log K)$ -Competitive Algorithm for Generalized Caching**

In the generalized caching problem, we have a set of pages and a cache of size k . Each page p has a size $w_p \geq 1$ and fetching cost c_p . The input consists of a sequence of page requests. If a page is not present in the cache when requested, it has to be loaded at a cost of c_p . We give a randomized $O(\log k)$ -competitive online algorithm for this problem.

Artur Czumaj, Anna Adamaszek, Matthias Englert
 University of Warwick
 A.Czumaj@warwick.ac.uk,
 a.m.adamaszek@warwick.ac.uk, m.englert@warwick.ac.uk

Harald Raecke
 TU Munich
 raecke@informatik.tu-muenchen.de

CP28**Improved Competitive Ratio for the Matroid Secretary Problem**

The Matroid Secretary Problem, introduced by Babaioff *et al.* (2007), is a generalization of the Classical Secretary Problem. In this problem, elements from a matroid are presented to an on-line algorithm in a random order. Each element has a weight associated with it, which is revealed to the algorithm along with the element. After each element is revealed the algorithm must make an irrevocable decision on whether or not to select it. The goal is to pick an independent set with the sum of the weights of the selected elements as large as possible. Babaioff *et al* gave an algorithm for the Matroid Secretary Problem with a *competitive ratio* of $O(\log \rho)$, where ρ is the rank of the matroid. It has been conjectured that a constant competitive ratio is achievable for this problem. In this paper we give an algorithm that has a competitive ratio of $O(\sqrt{\log \rho})$.

Oded Lachish
 Birkbeck, University of London
 oded@dcs.bbk.ac.uk

Sourav Chakraborty
 Chennai Mathematical Institute
 sourav@cmi.ac.in

CP28**Outperforming LRU via Competitive Analysis on Parametrized Inputs for Paging**

We propose a new measure of the "adversarial evilness" for the paging problem, namely the attack rate r , which is shown to be a tight bound for the competitive ratio of deterministic algorithms. Parameter r is usually small on real-world inputs. Using a priority-based framework, always yielding r -competitive algorithms, we propose a new paging strategy which outperforms LRU and other practical algorithms on many real-world traces.

Andrei Negoescu, Gabriel Moruz
 Goethe University, Frankfurt
 negoescu@cs.uni-frankfurt.de, gabi@cs.uni-frankfurt.de

CP28**Simultaneous Approximations for Adversarial and****Stochastic Online Budgeted Allocation**

We study the problem of simultaneous approximations for the adversarial and stochastic online budgeted allocation problem. For unweighted graphs, under some mild assumptions, we show that Balance [MSV07] achieves a competitive ratio of $1 - \epsilon$ in a random permutation model. For weighted graphs, however, we prove this is not possible. To this end, we show Weighted-Balance [MSV07] achieves a competitive ratio of 0.76 for the random arrival model, while having a $1 - \frac{1}{e}$ ratio in the worst case.

Shayan Oveis Gharan
 Stanford University
 shayan@stanford.edu

Vahab Mirrokni
 Google Inc.
 mirrokni@google.com

Morteza Zadimoghaddam
 MIT
 morteza@mit.edu

CP29**Fixed-Parameter Tractability of Directed Multiway Cut Parameterized by the Size of the Cutset**

Given a directed graph G , a set of k terminals and an integer p , the DIRECTED VERTEX MULTIWAY CUT problem asks if there is a set S of at most p (nonterminal) vertices whose removal disconnects each terminal from all other terminals. DIRECTED EDGE MULTIWAY CUT is the analogous problem where S is a set of at most p edges. These two problems indeed are known to be equivalent. A natural generalization of the multiway cut is the *multi-cut* problem, in which we want to disconnect only a set of k given pairs instead of all pairs. Marx (Theor. Comp. Sci. 2006) showed that in undirected graphs multiway cut is fixed-parameter tractable (FPT) parameterized by p . Marx and Razgon (STOC 2011) showed that undirected multicut is FPT and directed multicut is W[1]-hard parameterized by p . We complete the picture here by our main result which is that both DIRECTED VERTEX MULTIWAY CUT and DIRECTED EDGE MULTIWAY CUT can be solved in time $2^{2^{O(p)}} n^{O(1)}$, i.e., FPT parameterized by size p of the cutset of the solution. This answers an open question raised by Marx (Theor. Comp. Sci. 2006) and Marx and Razgon (STOC 2011). It follows from our result that DIRECTED MULTICUT is FPT for the case of $k = 2$ terminal pairs, which answers another open problem raised in Marx and Razgon (STOC 2011).

Rajesh H. Chitnis, MohammadTaghi Hajiaghayi
 University of Maryland, College Park
 rajeshchitnis@gmail.com, hajiagha@cs.umd.edu

Dániel Marx
 Institut für Informatik, Humboldt-Universität zu Berlin
 and
 Hungarian Academy of Sciences (MTA SZTAKI),
 Budapest, Hungary
 dmarx@cs.bme.hu

CP29**Erdős-Pósa Property and Its Algorithmic Applications — Parity Constraints, Subset Feedback Set,**

and Subset Packing

The well-known Erdős-Pósa theorem says that for any integer k and any graph G , either G contains k vertex-disjoint cycles or a vertex set X of order at most $c \cdot k \log k$ (for some constant c) such that $G - X$ is a forest. Thomassen extended this result to the even cycles, but on the other hand, this theorem is no longer true for the odd cycles. However, Reed proved this theorem still holds if we relax k vertex-disjoint odd cycles to k odd cycles with each vertex in at most two of them. These theorems initiate many researches in both graph theory and theoretical computer science. In the graph theory side, our problem setting is that we are given a graph and a vertex set S , and we want to extend all the above results to cycles each of which contains at least one vertex in S (such a cycle is called an S -cycle). It was recently shown that the above Erdős-Pósa theorem still holds for this subset version. In this paper, we extend both Thomassen's result and Reed's result in this way. In the theoretical computer science side, we investigate the following related problems in the framework of parameterized complexity: the feedback set problem with respect to the S -cycles, and the S -cycle packing problem. We give the first fixed parameter algorithms for the two problems (one for the feedback set problem was independently obtained in [Cygan, *et al.*, ICALP '11]). In addition, we extend the above results to those with the parity constraints.

Naonori Kakimura

University of Tokyo
kakimura@mist.i.u-tokyo.ac.jp

Ken-ichi Kawarabayashi

National Institute of Informatics, Japan
k.keniti@nii.ac.jp

Yusuke Kobayashi

University of Tokyo
kobayashi@mist.i.u-tokyo.ac.jp

CP29

Shortest Cycle Through Specified Elements

We give a randomized algorithm that finds a shortest simple cycle through a given set of k vertices or edges in an n -vertex undirected graph in time $2^k n^{O(1)}$.

Nina Taslamán

IT University of Copenhagen
nina.taslaman@gmail.com

Andreas Bjorklund

Lund University
andreas.bjorklund@yahoo.se

Thore Husfeldt

IT University of Copenhagen, Lund University
thore.husfeldt@gmail.com

CP29

Subexponential Parameterized Algorithm for Minimum Fill-in

The MINIMUM FILL-IN problem is to decide if a graph can be triangulated by adding at most k edges. Kaplan, Shamir, and Tarjan [FOCS 1994] have shown that the problem is solvable in time $(2^{(k)} + k^2 nm)$ on graphs with n vertices and m edges and thus is fixed parameter tractable. Here, we give the first subexponential parameterized algorithm solving MINIMUM FILL-IN in time $(2^{(\sqrt{k} \log k)} + k^2 nm)$. This substantially lowers the complexity of the problem.

ter tractable. Here, we give the first subexponential parameterized algorithm solving MINIMUM FILL-IN in time $(2^{(\sqrt{k} \log k)} + k^2 nm)$. This substantially lowers the complexity of the problem.

Yngve Villanger, Fedor V. Fomin

University of Bergen

yngve.villanger@ii.uib.no, fedor.fomin@ii.uib.no