

IP1**Pursuit on a Graph**

Pursuit games—motivated historically by military tactics—are a natural for graphical settings, and take many forms. We will present some recent results involving (among other things) drunks, Kakeya sets and a “ketchup graph.” Lastly, we discuss what we think is the most important open problem in the field.

Peter Winkler

Dartmouth College
Department of Mathematics
peter.winkler@dartmouth.edu

IP2**Optimization with Random Inputs**

The traditional approach to NP-hard problems in combinatorial optimization aims to design an algorithm with the best possible approximation factor. Many problems, however, are resistant to approximation, even though simple heuristics are said to be unreasonably effective in practice. To circumvent worst-case lower bounds, it is natural to focus on random inputs, but how random should they be? Classical average-case analysis goes all out and assumes fully random, artificially generated inputs, and then the results tend to be applicable within a narrow scope only. Instead, a growing body of work studies inputs with only a little bit of randomness, either on only one aspect of the input or with only a little bit of random noise. I will give examples from combinatorial optimization, online and streaming models, and discuss ways to define “minimally random” input models.

Claire Mathieu

CNRS, Ecole Normale Supérieure, France &
Brown University
cmathieu@di.ens.fr

IP3**New Directions in Learning Theory**

The classic PAC-learning model focuses on the problem of learning a single unknown function, over data drawn from a distribution that is unrelated to what is being learned, and that resides on a single machine. A number of current directions in machine learning differ from this setting in crucial ways. These include distributed learning, multi-task learning, semi-supervised learning, and deep/representation learning. This talk will survey some theoretical results in these directions and a number of open algorithmic problems that these directions bring.

Avrim Blum

Carnegie Mellon University
avrim@cs.cmu.edu

CP0**(Semi-)External Algorithms For Graph Partitioning and Clustering**

In this paper, we develop semi-external and external memory algorithms for graph partitioning and clustering problems. Graph partitioning and clustering are key tools for processing and analyzing large complex networks. Our algorithm is able to partition and cluster huge complex networks with billions of edges on cheap commodity machines. Experiments demonstrate that the semi-external

graph partitioning algorithm is scalable and can compute high quality partitions in time that is comparable to the running time of an efficient internal memory implementation.

Yaroslav Akhremtsev

Karlsruhe Institute of Technology
yaroslav.akhremtsev@kit.edu

Peter Sanders

Karlsruhe Institute of Technology
Faculty of Informatics
sanders@kit.edu

Christian Schulz

Karlsruhe Institute of Technology
Faculty of Informatics
christian.schulz@kit.edu

CP0**Approximating Betweenness Centrality in Large Evolving Networks**

Betweenness centrality ranks the importance of nodes by their participation in *all* shortest paths of the network. Therefore computing exact betweenness values is impractical in large networks. For static networks, approximation based on randomly sampled paths has been shown to be significantly faster in practice. However, for dynamic networks, no approximation algorithm for betweenness centrality is known that improves on static recomputation. We address this deficit by proposing two incremental approximation algorithms (for weighted and unweighted connected graphs) which provide a provable guarantee on the absolute approximation error. Our experimental study shows that our algorithms are the first to make in-memory computation of a betweenness ranking practical for million-edge semi-dynamic networks. Moreover, our results show that the accuracy is even better than the theoretical guarantees in terms of absolute errors and the rank of nodes is well preserved.

Elisabetta Bergamini, Henning Meyerhenke

Institute of Theoretical Informatics
Karlsruhe Institute of Technology (KIT)
elisabetta.bergamini@kit.edu,
henning.meyerhenke@kit.edu

Christian Staudt

Karlsruhe Institute of Technology
christian.staudt@kit.edu

CP0**The 1-2-3-Toolkit for Building Your Own Balls-into-Bins Algorithm**

We examine a generic class of simple distributed balls-into-bins algorithms and compute accurate estimates of the remaining balls and the load distribution after each round. Each algorithm is classified by (i) the load that bins accept in a given round and (ii) the number of messages each ball sends in a given round. Our algorithms employ a novel ranking mechanism resulting in notable improvements. Simulations independently verify our results and their high accuracy.

Pierre Bertrand

Ecole Normale Supérieure Cachan
Département de Maths

pierre.bertrand@ens-cachan.fr

Christoph Lenzen
MPI for Informatics
clenzen@mpi-inf.mpg.de

CP0

Cuts in Increasing Trees

Increasing trees have been extensively studied, since it is a simple model for many natural phenomena. Our paper focuses on sub-families of increasing trees. We measure the number of connected components obtained after having removed the nodes whose labels are smaller than a given value. This measure of cut-length allows, for example, to analyse in average an algorithm for tree-labelling. It is noticeable that we give exact formulae for the distribution of trees according to their size and cut-lengths. Our approach is based on a construction using grafting processes.

Antoine Genitrini
UPMC-LIP6
antoine.genitrini@lip6.fr

Olivier Bodini
LIPN, Université Paris 13
olivier.bodini@lipn.univ-paris13.fr

CP0

Bootstrap Percolation on Periodic Trees

We study bootstrap percolation with the threshold parameter $\theta \geq 2$ and the initial probability p on infinite periodic trees that are defined as follows. Each node of a tree has degree selected from a finite predefined set of non-negative integers and starting from any node, all nodes at the same graph distance from it have the same degree. We show the existence of the critical threshold $p_f(\theta) \in (0, 1)$ such that with high probability, (i) if $p > p_f(\theta)$ then the periodic tree becomes fully active, while (ii) if $p < p_f(\theta)$ then a periodic tree does not become fully active. We also derive a system of recurrence equations for the critical threshold $p_f(\theta)$ and compute these numerically for a collection of periodic trees and various values of θ , thus extending previous results for regular (homogeneous) trees.

Milan Bradonjic
Bell Labs, Alcatel-Lucent
Mathematics of Networks and Systems
milan@research.bell-labs.com

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

CP0

Experimental Evaluation of Multi-Round Matrix Multiplication on MapReduce

A common approach in the MapReduce literature is to design monolithic algorithms, that is algorithms that require just one or two rounds. We claim that this may not be the best approach in cloud systems, where multi-round algorithms may have some advantages. We carry out an experimental study of multi-round MapReduce algorithms aiming at investigating their performance and we use matrix multiplication as a case study. We first propose a scalable Hadoop library for matrix multiplication in the dense and

sparse cases which allows to tradeoff round number with the amount of data shuffled in each round and the amount of memory required by reduce functions. Then, we present an extensive study of this library on an in-house cluster and on Amazon Web Services aiming at showing its performance and at comparing monolithic and multi-round approaches. The experiments show that, even without a low level optimization, it is possible to design multi-round algorithms with a small overhead.

Francesco Silvestri, Matteo Ceccarello
University of Padova
silvest1@dei.unipd.it, ceccarel@dei.unipd.it

CP0

Computing MaxMin Edge Length Triangulations

We consider the problem of computing a triangulation for a given planar set of points with shortest edge as long as possible. We resolve an open problem dating back to 1991 by showing this to be NP-hard. Making use of integer programming formulations, we compute optimal solutions for benchmark instances; applying a number of refinements, we are able to do this for instances with up to 200 points.

Sándor P. Fekete
TU Braunschweig
Department of Computer Science
s.fekete@tu-bs.de

Winfried Hellmann, Michael Hemmer, Arne Schmidt,
Julian Troegel
TU Braunschweig
w.hellmann@tu-bs.de, m.hemmer@tu-bs.de,
arne.schmidt@tu-bs.de, j.troegel@tu-bs.de

CP0

On the Algorithmic Lovász Local Lemma and Acyclic Edge Coloring

The algorithm for Lovász Local Lemma by Moser and Tardos gives a constructive way to prove the existence of combinatorial objects that satisfy a system of constraints. We present an alternative probabilistic analysis of the algorithm that does not involve reconstructing the history of the algorithm from the witness tree. We apply our technique to improve the best known upper bound to acyclic chromatic index. Specifically we show that a graph with maximum degree Δ has an acyclic proper edge coloring with at most $\lceil 3.74(\Delta - 1) \rceil + 1$ colors, whereas the previously known best bound was $4(\Delta - 1)$. The same technique is also applied to improve corresponding bounds for graphs with bounded girth. An interesting aspect of this application is that the probability of the “undesirable” events do not have a uniform upper bound, i.e. it constitutes a case of the *asymmetric* Lovász Local Lemma.

Ioannis Giotis, Lefteris Kirousis, Kostas Psaromiligkos
Department of Mathematics
National Kapodistrian University of Athens, Greece
igiotis@cs.upc.edu, lkirousis@math.uoa.gr,
kostaspsa@gmail.com

Dimitrios Thilikos
CNRS, LIRMM, Montpellier

sedthilk@thilikos.info

CP0

Improved Single-Term Top-K Document Retrieval

On natural language text collections, finding the k documents most relevant to a query is generally solved with inverted indexes. On general string collections, however, more sophisticated data structures are necessary. Navarro and Nekrich [SODA 2012] showed that a linear-space index can solve such top- k queries in optimal time $O(m + k)$, where m is the query length. Konow and Navarro [DCC 2013] implemented the scheme, managing to solve top- k queries within microseconds with an index using 3.34.0 bytes per character. In this paper we introduce a new implementation using significantly less space and retaining similar query times. For short queries, which are the most difficult, our new index actually outperforms the previous one, as well as all the other solutions in the literature. We also show that our index can be built on very large text collections, and that it can handle phrase queries efficiently on natural language text collections.

Simon Gog

Karlsruhe Institute of Technology
gog@kit.edu

Gonzalo Navarro
University of Chile
gnavarro@dcc.uchile.cl

CP0

A Data-Aware Fm-Index

In this paper we present some practical modifications of the higher-order entropy-compressed text indexing method of Foschini et al. [Foschini et al. 2006] based upon the Burrows-Wheeler transform and the FM-index. Our method, called FM-Adaptive, retains the theoretical performance of previous work and introduces some improvements in practice. At the same time, broad experiments indicate that our index achieves superior performance, especially in terms of compression, in comparison to the state-of-the-art indexing techniques. The source code is available online.

Hongwei Huo, Longgang Chen, Heng Zhao
Xidian University
hwhuo@mail.xidian.edu.cn, chenlonggang.love@163.com,
goldbeef@163.com

Jeffrey Scott Vitter
The University of Kansas
jsv@ku.edu

Yakov Nekrich
University of Waterloo
yakov.nekrich@googlemail.com

Qiang Yu
Xidian University
feqond@163.com

CP0

2-Connectivity in Directed Graphs: An Experimental Study

Graph connectivity is a fundamental concept in graph the-

ory with numerous practical applications. In this paper we consider various notions of 2-edge and 2-vertex connectivity in directed graphs (digraphs). In particular, we consider the computation of the 2-vertex- and 2-edge-connected components and blocks in practice. These notions have a much richer and more complicated structure than their counterparts in undirected graphs, which has only been recently investigated. We present efficient implementations of both previously proposed and new algorithms for computing the 2-vertex- and 2-edge-connected components and blocks of a digraph, and evaluate their performance experimentally on large graphs taken from a variety of application areas. We conduct an extensive experimental study that sheds light on the relative difficulty of computing these notions of 2-connectivity in directed graphs in practice.

William Di Luigi
University of Bologna, Italy
williamdilui@gmail.com

Loukas Georgiadis
University of Ioannina, Greece
loukas@cs.uoi.gr

Giuseppe Italiano
Universita di Roma "Tor Vergata"
giuseppe.italiano@uniroma2.it

Luigi Laura
"Sapienza" Universita di Roma
laura@dis.uniroma1.it

Nikos Parotsidis
University of Ioannina, Greece
nparotsi@cs.uoi.gr

CP0

Branch-and-Reduce Exponential/FPT Algorithms in Practice: A Case Study of Vertex Cover

We investigate the gap between theory and practice for exact branching algorithms. In theory, branch-and-reduce algorithms currently have the best time complexity for numerous important problems. On the other hand, in practice, state-of-the-art methods are based on different approaches, and the empirical efficiency of such theoretical algorithms have seldom been investigated probably because they are seemingly inefficient because of the plethora of complex reduction rules. In this paper, we design a branch-and-reduce algorithm for the vertex cover problem using the techniques developed for theoretical algorithms and compare its practical performance with other state-of-the-art empirical methods. The results indicate that branch-and-reduce algorithms are actually quite practical and competitive with other state-of-the-art approaches for several kinds of instances, thus showing the practical impact of theoretical research on branching algorithms.

Takuya Akiba, Yoichi Iwata
The University of Tokyo
t.akiba@is.s.u-tokyo.ac.jp, y.iwata@is.s.u-tokyo.ac.jp

CP0

Engineering Motif Search for Large Graphs

In the *graph motif problem*, we are given as input a vertex-colored graph H and a multiset of colors M . Our task is to decide whether H has a connected set of vertices whose

multiset of colors agrees with M . The graph motif problem is NP-complete but admits parameterized algorithms that run in linear time in the size of H . We demonstrate that algorithms based on *constrained multilinear sieving* are viable in practice, scaling to graphs with hundreds of millions of edges as long as M remains small.

Andreas Björklund
Lund University
andreas.bjorklund@yahoo.se

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

Lukasz Kowalik
University of Warsaw
Institute of Informatics
kowalik@mimuw.edu.pl

Juho Lauri
Aalto University
Department of Information and Computer Science & HIIT
juho.lauri@aalto.fi

CP0 A Bound for the Diameter of Random Hyperbolic Graphs

Random hyperbolic graphs were recently introduced as a model for large networks. Their rigorous study was initiated using the following model: for $\alpha > 1/2$, $C > 0$, set $R = 2 \ln n + C$ and define the graph $G = (V, E)$ on n nodes as follows: For each $v \in V$, generate i.i.d. polar coordinates (r_v, θ_v) using the joint density function $f(r, \theta)$, with θ_v chosen uniformly from $[0, 2\pi)$ and r_v with density $f(r) = \frac{\alpha \sinh(\alpha r)}{\cosh(\alpha R) - 1}$ for $0 \leq r < R$. Then join two vertices by an edge, if their hyperbolic distance is at most R . Our main result establishes that for the relevant α range, a.a.s. for any two vertices of the same component, their graph distance is $O(\log^{C_0+1+o(1)})$, where $C_0 = 2/(\frac{1}{2} - \frac{3}{4}\alpha + \frac{\alpha^2}{4})$, thus answering a question raised by Gugelmann et al. (2012) concerning the diameter of such random graphs.

Marcos Kiwi
U. Chile
mkiwi@dim.uchile.cl

Dieter Mitsche
Université de Nice Sophia-Antipolis
dmitsche@unice.fr

CP0 Analysis and Experimental Evaluation of Time-Dependent Distance Oracles

Urban road networks are represented as directed graphs, with a metric which assigns distance *functions* (rather than scalars) to the arcs. In this work, we present distance oracles for providing time-dependent route plans, and conduct their experimental evaluation on a real-world time-dependent data set (city of Berlin). In particular: (i) We present the oracles, their implementation details concerning the digestion of the raw traffic data, and several heuristic improvements of both the preprocessing phase and the

query algorithms. The common preprocessing phase is based on a novel, quite efficient one-to-all approximation method for creating approximate distance summaries. (ii) We conduct an extensive experimental study with three query algorithms and six landmark sets. Our results are quite encouraging, achieving remarkable speedups (by two orders of magnitude) and quite small approximation guarantees (even PTAS), over the time-dependent variant of Dijkstra's algorithm.

Spyros Kontogiannis
University of Ioannina
kontog@cs.uoi.gr

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

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

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

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

CP0 Cache Replacement with Memory Allocation

We consider a variant of the generalized caching problem in which the cache management policy must not only specify which items to evict to make room for an incoming item, but must also specify a location in memory where each object can be placed contiguously. The problem is motivated by implementations of key-value stores in commercial databases with high read-to-write ratio such as those maintained by Facebook. We propose a simple algorithm and show that if the algorithm is given some additional memory to account for fragmentation, it is competitive against an offline optimal algorithm that does not specify memory layout. On the benchmark traces in the experiments presented here, our algorithm requires approximately 10–15% additional space to be k -competitive against the optimal offline algorithm. Through trace-driven simulations, we demonstrate that the caching performance of our algorithm is close to that of competitive strategies that are not required to manage memory layout.

Shahram Ghandeharizadeh
University of Southern California
shahram@dblab.usc.edu

Sandy Irani, Jenny Lam
University of California, Irvine
irani@ics.uci.edu, jenny.lam@uci.edu

CP0 On Distributed Cardinality Estimation: Random Arcs Recycled

We introduce a distributed cardinality estimation algorithm that is a generalization of the classic 'balls-and-bins' counting algorithm and which refers to the process of cov-

ering the circle with random arcs. We prove the correctness of the algorithm and by the methods of complex analysis we examine the properties of the applied estimator. We also show that the construction of the proposed algorithm is a backbone for a simple distributed summation algorithm.

Marcin Kardas, Mirosław Kutylowski
Wrocław University of Technology
Department of Computer Science
marcin.kardas@pwr.edu.pl,
miroslaw.kutylowski@pwr.edu.pl

Jakub Lemiesz
Wrocław University of Technology
jakub.lemiesz@pwr.edu.pl

CP0

Analytic Samplers and the Combinatorial Rejection Method

Abstract not available at time of publication.

Jérémie Lumbroso
Université Pierre et Marie Curie
jeremie.lumbroso@gmail.com

CP0

Region-Based Approximation Algorithms for Visibility Between Imprecise Locations

In this paper we present new geometric algorithms for approximating the visibility between two imprecise locations amidst a set of obstacles, where the imprecise locations are modeled by continuous probability distributions. Our techniques are based on approximating distributions by a set of regions rather than on approximating by a discrete point sample. In this way we obtain guaranteed error bounds, and the results are more robust than similar results based on discrete point sets. We implemented our techniques and present an experimental evaluation. The experiments show that the actual error of our region-based approximation scheme converges quickly when increasing the complexity of the regions.

Maarten Löffler
Utrecht University
m.loffler@uu.nl

Irina Kostitsyna
Technical University Eindhoven
i.kostitsyna@tue.nl

Rodrigo Silveira
Universitat Politècnica de Catalunya
rodrigo.silveira@upc.edu

Kevin Buchin
TU Eindhoven
k.a.buchin@tue.nl

CP0

Variance of Size in Regular Graph Tries

Graph tries are tries built from label functions on a graph G . Here we determine asymptotics for the variance of the size of a G -trie on a memoryless source with a uniform alphabet distribution, where G comes from a large class of infinite, M -regular DAGs. We show that the variance is

$\Theta(n^{\rho'})$, with ρ' depending on G minimized when G is a tree. We also give an expression for ρ' for grids.

Philippe Jacquet
Alcatel Lucent
philippe.jacquet@alcatel-lucent.com

Abram Magner
Department of Computer Science
Purdue University
anmagner@purdue.edu

CP0

Faster Linear-Space Orthogonal Range Searching in Arbitrary Dimensions

We consider the problem of orthogonal range searching in linear space for any d dimensions. The kd-tree achieves $O(n^{(d-1)/d})$ query time for range counting, which has been considered to be the best complexity bound in practice, while the non-overlapping k-range achieves $O(n^\epsilon)$ time in theory. In this paper, we propose a new succinct data structure with $O(n^{(d-2)/d} \log n)$ query time for range counting. In experiments, it significantly outperformed the kd-tree for high selectivity queries.

Yuzuru Okajima
VALWAY Technology Center, NEC Solution Innovators,
Ltd.
okajima-yuzuru@mxn.nes.nec.co.jp

Kouichi Maruyama
NEC Solution Innovators, Ltd.
maruyama-k@mxn.nes.nec.co.jp

CP0

Linear-Time Generation of Inhomogeneous Random Directed Walks

Directed two-dimensional random walks modelise diffusion, and are related to Galton–Watson trees. These random objects can be generated in linear time, through an algorithm of Devroye, based on the fact that the steps form an exchangeable sequence of random variables. We consider here the random generation of a more general family of structures, in which the transition rates, instead of being fixed, evolve in time. The steps are not exchangeable anymore. This describes time-dependent diffusion, and, for trees arising from exploration processes on finite random graphs, accounts for effects of excluded volume. Applications include rooted accessible K -maps, which provide the uniform generation of random minimal automata. We present an algorithm for this wider family of problems, for which the achieved (bit-)complexity remains linear.

Andrea Sportiello
LIPN, Université Paris Nord
Villetaneuse, FRANCE
andrea.sportiello@gmail.com

Frédérique Bassino
LIPN, Univ. Paris Nord
Villetaneuse, FRANCE
frederique.bassino@lipn.fr

CP0

An I/O-Efficient Distance Oracle for Evolving

Real-World Graphs

We present an I/O-efficient distance oracle that is able to answer online queries with a constant number of I/O. Furthermore, we developed batched queries that have an amortized I/O-complexity of $O(\frac{1}{B})$ I/Os per query. Online queries can be processed in milliseconds on SSDs and batched queries within microseconds even on HDDs. All results have been achieved on real world graphs. We explain the experimental results and discuss improvements for the future.

David Veith

Goethe University, Frankfurt
dveith@cs.uni-frankfurt.de

Deepak Ajwani
Bell Labs Ireland
deepak.ajwani@gmail.com

Ulrich Meyer
Goethe University, Frankfurt
uli.meyer.cs@googlemail.com

CP0

Repeated Fringe Subtrees in Random Rooted Trees

A fringe subtree of a rooted tree consists of a node and all its descendants. We show that the average number of fringe subtrees that occur at least r times is of asymptotic order $n/(\log n)^{3/2}$ for every $r \geq 2$ if an n -node tree is taken uniformly at random from a simply generated family. Moreover, we also prove a strong concentration result for the size of the smallest tree that does not occur as a fringe subtree.

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

CP0

Lattice Paths of Slope 2/5

We analyze some enumerative and asymptotics properties of Dyck paths under a line of slope 2/5. This answers to Knuth's problem #4 from his "Flajolet lecture" during the conference "Analysis of Algorithms" (AofA'2014). Our approach relies on the work of Banderier and Flajolet for asymptotics and enumeration of directed lattice paths. A key ingredient is the "kernel method", an old trick of Knuth himself. We also show how to obtain similar results for other slopes.

Michael Wallner

TU Wien
michael.wallner@tuwien.ac.at

Cyril Banderier
Université Paris-Nord
cyril.banderier@lipn.univ-paris13.fr

CP0

Analysis of Branch Misses in Quicksort

Analyzing classic elementary operations is not sufficient when features of modern processors like pipelined execution and memory hierarchies influence running time. For Quicksort, for example, it has been demonstrated experi-

mentally that on certain hardware the classically optimal choice of the pivot as median of a sample gets *harmful* because the rollback costs of mispredicted branches become dominating. In this paper, we give the first precise analytical investigation of branch mispredictions in (classic) Quicksort and Yaroslavskiy's dual-pivot Quicksort variant as implemented in Oracle's Java 7 library. It is still not fully understood why the latter is 10% faster in experiments. For different branch prediction strategies, we give the expected number of branch misses caused by the aforementioned Quicksort variants when pivots are chosen from a sample. We conclude that the difference in branch misses is too small to explain the superiority of the dual-pivot algorithm.

Conrado Martínez
Department of Computer Science
Univ. Politècnica de Catalunya
conrado@cs.upc.edu

Markus Nebel, Sebastian Wild
Computer Science Department
University of Kaiserslautern
nebel@cs.uni-kl.de, s_wild@cs.uni-kl.de

CP0

On k -Connectivity and Minimum Vertex Degree in Random s -Intersection Graphs

Random s -intersection graphs have recently received much interest in a wide range of application areas. Broadly speaking, a random s -intersection graph is constructed by first assigning each vertex a set of items in some random manner, and then putting an undirected edge between all pairs of vertices that share at least s items (the graph is called a random intersection graph when $s = 1$). A special case of particular interest is a *uniform* random s -intersection graph, where each vertex independently selects the same number of items uniformly at random from a common item pool. Another important case is a *binomial* random s -intersection graph, where each item from a pool is independently assigned to each vertex with the same probability. Both models have found numerous applications thus far including cryptanalysis, and modeling recommender systems, secure sensor networks, online social networks, trust networks and small-world networks (uniform random s -intersection graphs), as well as clustering analysis, classification, and the design of integrated circuits (binomial random s -intersection graphs). In this paper, for binomial/uniform random s -intersection graphs, we present results related to k -connectivity and minimum vertex degree. Specifically, we derive the asymptotically exact probabilities and zero-one laws for the following three properties: (i) k -vertex-connectivity, (ii) k -edge-connectivity and (iii) the property of minimum vertex degree being at least k .

Jun Zhao, Osman Yagan, Virgil Gligor
Carnegie Mellon University
junzhao@cmu.edu, oyagan@ece.cmu.edu, gligor@cmu.edu

CP1

Approximating Independent Sets in Sparse Graphs

We consider the independent set problem on graphs with maximum degree d . Currently, an $O(d \log \log d / \log d)$ approximation and a $(d / \log^2 d)$ UGC hardness is known. We show the following: (i) The natural LP strengthened by $O(\log^4(d))$ levels of the mixed-hierarchy has integrality gap $\tilde{O}(d / \log^2 d)$. However, our proof is non-constructive and

does not give a good algorithm. (ii) An $O(d/\log d)$ approximation using $\text{polylog}(d)$ -levels of the mixed hierarchy that runs in time $n^{O(1)} \exp(\log^{O(1)} d)$, improving upon the previous approximation by $\log \log d$.

Nikhil Bansal

Eindhoven University of Technology
bansal@gmail.com

CP1

New Approximation Schemes for Unsplittable Flow on a Path

We study the unsplittable flow on a path problem. Given is a path with capacities on its edges and a set of tasks. The goal is to find a subset of the tasks of maximum total profit such that all task demands can be routed simultaneously without violating the capacity constraints. When the task densities—defined as the ratio of a task's profit and demand—lie in a constant range, we obtain a PTAS. We improve the QPTAS of Bansal et al. by removing the assumption that the demands need to lie in a quasi-polynomial range. Our third result is a PTAS for the case where we are allowed to shorten the paths of the tasks by at most an ϵ -fraction. Each of these results uses a sparsification lemma which could be of independent interest. The lemma shows that in any (optimal) solution there exists an $O(\epsilon)$ -fraction (measured by weight) of its tasks whose removal creates, on each edge, a slack which is at least as large as the $(1/\epsilon)$ th largest demand using that edge.

Jatin Batra, Naveen Garg, Amit Kumar
IIT Delhi

jatinbatra@cse.iitd.ac.in, naveen@cse.iitd.ac.in,
amitk@cse.iitd.ac.in

Tobias Mömke
Saarland University
moemke@cs.uni-saarland.de

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

CP1

Spider Covers for Prize-Collecting Network Activation Problem

In the network activation problem, each edge in a graph is associated with an activation function that decides whether the edge is activated from weights assigned to its end nodes. The feasible solutions of the problem are node weights such that the activated edges form graphs of required connectivity, and the objective is to find a feasible solution minimizing its total weight. In this paper, we consider a prize-collecting version of the network activation problem and present the first nontrivial approximation algorithms. Our algorithms are based on a new linear programming relaxation of the problem. They round optimal solutions for the relaxation by repeatedly computing node weights activating subgraphs, called spiders. For the problem with element- and node- connectivity requirements, we also present a new potential function on uncrossable biset families and use it to analyze our algorithms.

Takuro Fukunaga
National Institute of Informatics

takuro@nii.ac.jp

CP1

On Survivable Set Connectivity

In the Set Connectivity problem, we are given an n -node edge-weighted undirected graph and a collection of h set pairs (S_i, T_i) , where S_i and T_i are subsets of nodes. The goal is to compute a min-cost subgraph H so that, for each (S_i, T_i) , H has a path between S_i and T_i .

We initiate the study of the Survivable Set Connectivity problem (SSC) – the generalization of Set Connectivity where we are additionally given an integer requirement $k_i \geq 1$ for each (S_i, T_i) , and we want to find a min-cost subgraph H that has k_i edge-disjoint paths between S_i and T_i . We prove the following results:

- An approximation hardness of $2^{\log^{1-\epsilon} n}$ for SSC, for any $\epsilon > 0$ assuming NP does not have a quasi-polynomial time algorithm.
- A bicriteria approximation algorithm for SSC that computes a solution H such that $\text{cost}(H) \leq \text{opt} * \text{polylog}(nh)$ and connectivity $\geq k_i / \log n$ for all (S_i, T_i) . So, by relaxing connectivity, we obtain a cost-guarantee beyond the lower bound.

Parinya Chalermsook
Max-Planck-Institut für Informatik
parinya@mpi-inf.mpg.de

Fabrizio Grandoni
IDSIA, University of Lugano
fabrizio@idsia.ch

Bundit Laekhanukit
Simons Institute & I-CORE
bundit.laekhanukit@mail.mcgill.ca

CP1

A Note on the Ring Loading Problem

The Ring Loading Problem is an optimal routing problem arising in the planning of optical communication networks which use bidirectional SONET rings. In mathematical terms, it is an unsplittable multicommodity flow problem on undirected ring networks. We prove that any split routing solution can be turned into an unsplittable solution while increasing the load on any edge of the ring by no more than $+\frac{2}{5}D$, where D is the maximum demand value. This improves upon a classical result of Schrijver, Seymour, and Winkler (1998) who obtained a slightly larger bound of $+\frac{3}{2}D$. We also present an improved lower bound of $\frac{11}{10}D$ (previously $\frac{101}{100}D$) on the best possible bound and disprove a famous and long-standing conjecture of Schrijver et al. in this context.

Martin Skutella
Institut für Mathematik, Sekr. MA 5, Tech Uni Berlin,
Strae des 17. Juni 136, 10623 Berlin, Germany
skutella@math.tu-berlin.de

CP2

Pricing Online Decisions: Beyond Auctions

We consider dynamic pricing schemes in online settings where selfish agents generate online events. Previous work on online mechanisms has dealt almost entirely with the

goal of maximizing social welfare or revenue in an auction settings. This paper deals with quite general settings and minimizing social costs. We show that appropriately computed posted prices allow one to achieve essentially the same performance as the best online algorithm. Unlike online algorithms that learn about the event, and then make enforceable decisions, prices are posted without knowing the future events or even the current event, and are thus inherently dominant strategy incentive compatible. In particular we show that one can give efficient posted price mechanisms for metrical task systems, some instances of the k -server problem, and metrical matching problems.

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

Alon Eden, Amos Fiat
Tel Aviv University
alonarden@gmail.com, fiat@tau.ac.il

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

CP2

On the Complexity of Computing An Equilibrium in Combinatorial Auctions

We study the computation of pure Nash equilibria with good social welfare in the game of combinatorial auctions where each item is sold separately but simultaneously via a second-price auction. We show that when bidders' valuations are submodular, in many settings computing equilibria with good welfare is as easy as computing good allocations. For subadditive valuations, computing an equilibrium requires exponential communication, and for XOS valuations, this requires techniques different from the ones currently known.

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

Hu Fu, Robert Kleinberg
Cornell University
hufu@microsoft.com, rdk@cs.cornell.edu

CP2

Combinatorial Auctions Via Posted Prices

We study anonymous posted price mechanisms for Combinatorial Auctions(CA) in a Bayesian framework. In such a mechanism item prices are posted, then the consumers approach the seller sequentially in an arbitrary order, each taking her favorite bundle from among the unsold items. These mechanisms are simple, transparent and Dominant Strategy Incentive Compatible(DSIC). For fractionally subadditive(includes submodular) agents' preferences, we find prices that obtain at least half of the optimal welfare in expectation. Our result is constructive: given black-box access to a CA algorithm A, sample access to the prior, appropriate query access to the sampled valuations, we compute prices that generate at least half of the expected welfare of A. As a corollary, we get the first polytime (in n and m) constant-factor DSIC mechanism for Bayesian submodular CA given access to demand queries.

For valuations with complements our approximations degrade linearly with the level of complementarity.

Michal Feldman
Tel-Aviv University
mfeldman@tau.ac.il

Nick Gravin, Brendan Lucier
Microsoft Research
ngravin@microsoft.com, brlucier@microsoft.com

CP2

Welfare Maximization with Production Costs: A Primal Dual Approach

We study online combinatorial auctions with production costs using the online primal dual framework. For arbitrary (strictly convex and differentiable) production cost functions, we characterize the optimal competitive ratio achievable by online mechanisms/algorithms. We show that online posted pricing mechanisms, which are incentive compatible, can achieve competitive ratios arbitrarily close to the optimal, and construct lower bound instances on which no online algorithms, not necessarily incentive compatible, can do better.

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

Anthony Kim
Stanford University
tonyekim@stanford.edu

CP2

An n -to-1 Bidder Reduction for Multi-Item Auctions and Its Applications

We introduce a novel approach for reducing the maximal revenue k -item n -bidder auction with additive valuation to k -item 1-bidder auctions. This approach leads to a simple deterministic mechanism that achieves a constant fraction of the optimal revenue when items are independent. In addition, it leads to several structural results, including a constant-ratio relationship between the optimal revenue obtainable by dominant strategy mechanisms and that obtainable by Bayesian mechanisms.

Andrew C. Yao
Tsinghua University
andrewcyao@tsinghua.edu.cn

CP3

Efficient and Robust Persistent Homology for Measures

A new paradigm for point cloud data analysis has emerged recently, where point clouds are treated as empirical measures instead of compact sets. A notion of distance to a measure has been defined and shown to be stable with respect to perturbations. It can easily be computed pointwise in the case of a point cloud, but its sublevel-sets, which carry the geometric information, remain hard to compute or approximate. This makes it challenging to adapt many powerful techniques based on the Euclidean distance to a point cloud to the larger setting of the distance to a measure on a metric space. We propose an efficient and reliable scheme to approximate the topological structure of

the family of sublevel-sets of the distance to a measure. We obtain an algorithm for approximating the persistent homology of the distance to an empirical measure working in arbitrary metric spaces. Precise quality and complexity guarantees are given with a discussion on the behavior of our approach in practice.

Mickaël Buchet, Frédéric Chazal, Steve Oudot
INRIA Saclay Ile-de-France
mickael.buchet@m4x.org, frederic.chazal@inria.fr,
steve.oudot@inria.fr

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

CP3

Minimum Forcing Sets for Miura Folding Patterns

We introduce the study of forcing sets in mathematical origami. The origami material folds flat along mountain or valley creases. A subset F of creases is forcing if the mountain/valley assignment can be deduced from F . We develop efficient algorithms for constructing a minimum forcing set of a Miura-ori map, and for deciding whether a given set of creases is forcing or not. We also provide tight bounds on the size of a forcing set.

Brad Ballinger
Humboldt State University
bradley.ballinger@humboldt.edu

Mirela Damian
Villanova University
mirela.damian@villanova.edu

David Eppstein
University of California, Irvine
Department of Computer Science
eppstein@uci.edu

Robin Flatland
Siena College
flatland@siena.edu

Jessica Ginepro
University of Connecticut
jessginepro@yahoo.com

Thomas Hull
Western New England University
thull@wne.edu

CP3

Universal Computation with Arbitrary Polyomino Tiles in Non-Cooperative Self-Assembly

We explore the power of geometry to overcome limitations of non-cooperative self-assembly. We define a generalization of the abstract Tile Assembly Model with tile systems consisting of collections of polyomino tiles, the Polyomino Tile Assembly Model, and investigate the computational powers of polyTAM systems at *temperature-1*, where attachment among tiles occurs without glue cooperation. We prove that for any polyomino P of size ≥ 3 there exists a computationally universal temperature-1 polyTAM system

containing only shape- P tiles.

Sándor P. Fekete
TU Braunschweig
Department of Computer Science
s.fekete@tu-bs.de

Jacob Hendricks
University of Arkansas
University of Arkansas
jhendric@uark.edu

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

Robert T. Schweller
University of Texas - Pan American
rtschweller@utpa.edu

CP3

Zigzag Persistence Via Reflections and Transpositions

We introduce a simple algorithm for computing zigzag persistence, designed in the same spirit as the standard persistence algorithm. Our algorithm reduces a single matrix, maintains an explicit set of chains encoding the persistent homology of the current zigzag, and updates it under simplex insertions and removals. The total worst-case running time matches the usual cubic bound. A noticeable difference with the standard persistence algorithm is that we do not insert or remove new simplices "at the end" of the zigzag, but rather "in the middle". To do so, we use arrow reflections and transpositions, in the same spirit as reflection functors in quiver theory. Our analysis introduces two new kinds of reflection called the "injective and surjective diamonds" and the "transposition diamond", for which we are able to predict the changes in the interval decomposition and associated compatible bases.

Clément Maria
INRIA Sophia Antipolis-Méditerranée
clement.maria@inria.fr

Steve Oudot
INRIA Saclay Ile-de-France
steve.oudot@inria.fr

CP3

Improved Bounds for Orthogonal Point Enclosure Query and Point Location in Orthogonal Subdivisions in \mathbb{R}^3

In this paper, new results for two fundamental data structure problems in the field of computational geometry are presented: orthogonal point enclosure query (*OPEQ*) in $3d$ and point location in orthogonal subdivisions in $3d$. All the results are in the pointer machine model of computation. In *OPEQ* the objective is to report those (preprocessed) rectangles which are stabbed by a query point. *OPEQ* has been optimally solved (in terms of space and query time) in $1d$ and $2d$ many decades back. For *OPEQ* in $3d$, optimal solution is known only when rectangles are orthant. The key result of this work is an almost optimal solution for more general rectangles in $3d$. For point location in orthogonal subdivisions (a special case of *OPEQ* where rectangles are disjoint), a data structure is presented which

occupies $O(n)$ space and answers the query in $O(\log^{3/2} n)$ time, improving the previously best known query time by roughly $\sqrt{\log n}$ factor.

Saladi Rahul

University of Minnesota, Twin Cities
sala0198@umn.edu

CP4

Speeding Up the Four Russians Algorithm by About One More Logarithmic Factor

We present a new combinatorial algorithm for Boolean matrix multiplication that runs in $O(n^3(\log \log n)^3/\log^3 n)$ time. This improves the previous combinatorial algorithm by Bansal and Williams [FOCS'09] that runs in $O(n^3(\log \log n)^2/\log^{9/4} n)$ time. Whereas Bansal and Williams' algorithm uses regularity lemmas for graphs, the new algorithm is simple and uses entirely elementary techniques: table lookup, word operations, plus a deceptively straightforward divide-and-conquer.

Timothy M. Chan

University of Waterloo
tmchan@uwaterloo.ca

CP4

Degree-3 Treewidth Sparsifiers

We study treewidth sparsifiers. Informally, given a graph G of treewidth k , a treewidth sparsifier H is a minor of G , whose treewidth is close to k , $|V(H)|$ is small, and the maximum vertex degree in H is bounded. In this paper we describe an algorithm that, given a graph G of treewidth k , computes a topological minor H of G such that (i) the treewidth of H is $\Omega(k/\text{polylog}(k))$; (ii) $|V(H)| = O(k^4)$; and (iii) the maximum vertex degree in H is 3. The running time of the algorithm is polynomial in $|V(G)|$ and k .

Chandra Chekuri

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

Julia Chuzhoy

Toyota Technological Institute at Chicago
cjulia@ttic.edu

CP4

Improved Bounds for the Flat Wall Theorem

The Flat Wall Theorem of Robertson and Seymour states that there is some function f , such that for all integers $w, t > 1$, every graph G containing a wall of size $f(w, t)$, contains either (i) a K_t -minor; or (ii) a small set $A \subset V(G)$ of vertices, and a flat wall of size w in $G \setminus A$. Kawarabayashi, Thomas and Wollan recently showed a proof of this theorem with the following two sets of parameters: (1) $f(w, t) = \Theta(t^{24}(t^2 + w))$ with $|A| = O(t^{24})$, and (2) $f(w, t) = w^{2^{\Theta(t^{24})}}$ with $|A| \leq t - 5$. In this paper we improve their bounds to $f(w, t) = \Theta(t(t + w))$ with $|A| \leq t - 5$. For the special case where the maximum vertex degree in G is bounded by D , we show that, if G contains a wall of size $\Omega(Dt(t + w))$, then either G contains a K_t -minor, or there is a flat wall of size w in G . We complement our result for the low-degree case by proving an almost matching lower bound.

Julia Chuzhoy

Toyota Technological Institute at Chicago
cjulia@ttic.edu

CP4

Beating Exhaustive Search for Quantified Boolean Formulas and Connections to Circuit Complexity

We present new algorithms for the satisfiability problem for quantified Boolean formulas (QBFs), and consequences of faster QBF algorithms for circuit complexity. For example, we show that when the number of quantifier blocks in the QBF is either $\omega(\log n)$ or $o(\log n / \log \log n)$ (where n is the number of variables), the 2^n running time of exhaustive search can be beaten in a significant way. On the other hand, we prove that beating exhaustive search for QBF when the number of quantifier blocks is $O(\log n)$ would imply major circuit complexity lower bounds.

Rahul Santhanam

University of Edinburgh
rsanthan@inf.ed.ac.uk

Ryan Williams

Stanford University
rrwilliams@gmail.com

CP4

More Applications of the Polynomial Method to Algorithm Design

In low-depth circuit complexity, the polynomial method is a way to prove lower bounds by translating weak circuits into low-degree polynomials, then analyzing their properties. Recently, this method found an application to algorithm design: Williams (STOC 2014) used it to compute all-pairs shortest paths in $n^3/2^{\Omega(\sqrt{\log n})}$ time on dense n -node graphs. In this paper, we extend this methodology to solve a number of problems faster than previously known methods. We first give an algorithm for BOOLEAN ORTHOGONAL DETECTION, which is to detect among two sets $A, B \subseteq \{0, 1\}^d$ of size n if there is an $x \in A$ and $y \in B$ such that $\langle x, y \rangle = 0$. For vectors of dimension $d = c(n) \log n$, we solve the problem in $n^{2-1/O(\log c(n))}$ time. We apply this as a subroutine in several other new algorithms: Batch Partial Match, DNF evaluation, LCS with don't cares, Longest Substrings with small edit distance, Symmetric Boolean CSPs.

Amir Abboud, Ryan Williams, Huacheng Yu

Stanford University
amir.abboud@gmail.com, rrwilliams@gmail.com,
yuhch123@gmail.com

CP5

Short Paths on the Voronoi Graph and Closest Vector Problem with Preprocessing

We give a $\tilde{O}(2^n)$ expected time and space algorithm for CVPP (the preprocessing version of the Closest Vector Problem). This improves on the $\tilde{O}(4^n)$ deterministic runtime of the Micciancio Voulgaris algorithm for CVPP, based on Voronoi cell techniques, at the cost of a polynomial amount of randomness (which only affects runtime, not correctness). As a byproduct of our work, we also give an optimal relationship between geometric and path distance on the Voronoi graph.

Nicolas Bonifas

Ecole Polytechnique & IBM
nicolas.bonifas@polytechnique.edu

Daniel Dadush
New York University
dadush@cs.nyu.edu

CP5

On Largest Volume Simplices and Sub-Determinants

We show that the problem of finding the simplex of largest volume in the convex hull of n points in \mathbb{Q}^d can be approximated with a factor of $O(\log d)^{d/2}$ in polynomial time. This improves upon the previously best known approximation guarantee of $d^{(d-1)/2}$ by Khachiyan. On the other hand, we show that there exists a constant $c > 1$ such that this problem cannot be approximated with a factor of c^d , unless $P=NP$. Our hardness result holds even if $n = O(d)$, in which case there exists a \bar{c}^d -approximation algorithm that relies on recent sampling techniques, where \bar{c} is again a constant. We show that similar results hold for the problem of finding the largest absolute value of a subdeterminant of a $d \times n$ matrix.

Marco Di Summa
Universita degli Studi di Padova
disumma@math.unipd.it

Friedrich Eisenbrand, Yuri Faenza, Carsten Moldenhauer
Ecole Polytechnique Federale de Lausanne
friedrich.eisenbrand@epfl.ch, yuri.faenza@epfl.ch,
carsten.moldenhauer@epfl.ch

CP5

Approximate Nearest Line Search in High Dimensions

We consider the Approximate Nearest Line Search (NLS) problem. Given a set L of N lines in high dimensional Euclidean space R^d , the goal is to build a data structure that, given a query point $q \in R^d$, reports a line $\ell \in L$ s.t. its distance to q is within $(1 + \epsilon)$ factor of the distance of the closest line to q . This is a natural generalization of the Approximate Nearest Neighbor problem for point sets (ANN), and is a natural first step to build nearest-neighbor data structures for objects other than points. We present a data structure that, for any fixed $\epsilon > 0$, reports the approximate nearest line in time $(d + \log N + 1/\epsilon)^{O(1)}$ using $O(N + d)^{O(1/\epsilon^2)}$ space. This is the first high-dimensional data structure for this problem with poly-log query time and polynomial space. The best previous, due to Magen, required quasi-polynomial space. The bounds we achieve match the performance of the best algorithm for the ANN problem.

Sepideh Mahabadi
MIT
mahabadi@mit.edu

CP5

Approximating Hereditary Discrepancy Via Small Width Ellipsoids

The *Discrepancy* of a hypergraph is the minimum attainable value, over two-colorings of its vertices, of the maximum absolute imbalance of any hyperedge. The *Hered-*

itary Discrepancy of a hypergraph, defined as the maximum discrepancy of a restriction of the hypergraph to a subset of its vertices, is a measure of its complexity. Recent work by Nikolov, Talwar and Zhang (2013) showed a polynomial time $\tilde{O}(\log^3 n)$ -approximation to hereditary discrepancy, as a by-product of their work in differential privacy. In this paper, we give a direct simple $O(\log^{3/2} n)$ -approximation algorithm for this problem. We show that up to this approximation factor, the hereditary discrepancy is characterized by the optimal value of a geometric convex program that seeks to minimize the largest ℓ_∞ norm of any point in an ellipsoid containing the columns of the incidence matrix of the hypergraph. This characterization promises to be a useful tool in discrepancy theory.

Aleksandar Nikolov
Rutgers University
anikolov@cs.rutgers.edu

Kunal Talwar
No affiliation
k@kunalatalwar.org

CP5

Fast Lattice Point Enumeration with Minimal Overhead

Enumeration algorithms are arguably the best currently known methods to solve lattice problems. However, there is an uncomfortable gap between our theoretical understanding and practical performance of enumeration algorithms. We introduce a new class of lattice enumeration algorithms that simultaneously achieve asymptotic efficiency and practicality, matching or surpassing the performance of practical algorithms already in moderately low dimension.

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

Michael Walter
University of California, San Diego
miwalter@eng.ucsd.edu

CP6

Gossip vs. Markov Chains, and Randomness-Efficient Rumor Spreading

We study gossip algorithms for the rumor spreading problem which asks one node to deliver a rumor to all nodes in an unknown network, and every node is only allowed to call one neighbor in each round. In this work we introduce two fundamentally new techniques in studying the rumor spreading problem: (1) We establish a new connection between the rumor spreading process in an *arbitrary* graph and certain Markov chains, and show that the mixing time of a certain Markov chain suffices to bound the rumor spreading time in an arbitrary graph. (2) We construct a reduction from rumor spreading processes to branching programs. Based on this reduction, we show that, for *any* n -vertex expander graph, there is a protocol which informs every node in $O(\log n)$ rounds with high probability, and uses $O(\log n \cdot \log \log n)$ random bits in total.

Zeyu Guo
Caltech
zguo@caltech.edu

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

CP6

Plurality Consensus in the Gossip Model

We study Plurality Consensus in the Gossip Model over a network of n anonymous agents. Each agent supports an initial opinion or color. The goal is to provide a protocol that brings the system into the configuration in which all agents support the (initial) plurality. We consider the Undecided-State Dynamics, a well-known protocol which uses just one more state than those necessary to store colors. We show that the speed of convergence of this protocol depends on the initial color configuration as a whole. We capture this dependence by introducing the notion of monochromatic distance $md(\bar{c})$ which measures the distance of the initial color configuration \bar{c} from the closest monochromatic one. In the complete graph, we prove that this dynamics converges within $O(md(\bar{c}) \log n)$ rounds. We prove that this upper bound is almost tight and we adapt the Undecided-State Dynamics to obtain a fast, random walk-based protocol for plurality consensus on regular expanders.

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

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

Emanuele Natale, Francesco Pasquale, Riccardo Silvestri
Sapienza University of Rome
natale@di.uniroma1.it, pasquale@di.uniroma1.it, silvestri@di.uniroma1.it

CP6

Distributed Computation of Large-Scale Graph Problems

Motivated by the increasing need for fast distributed processing of large-scale graphs such as the Web graph and various social networks, we study a number of fundamental graph problems in the message-passing model, where we have k machines that jointly perform computation on an arbitrary n -node (typically, $n \gg k$) input graph. The graph is assumed to be *randomly* partitioned among the $k \geq 2$ machines (a common implementation in many real world systems). The communication is point-to-point, and the goal is to minimize the time complexity, i.e., the number of communication rounds, of solving various fundamental graph problems.

Hartmut Klauck
Division of Mathematical Sciences
Nanyang Technological University
hklauck@gmail.com

Danupon Nanongkai
Nanyang Technological University, Singapore
danupon@gmail.com

Gopal Pandurangan
Division of Mathematical Sciences
Nanyang Technological University

gopalpandurangan@gmail.com

Peter Robinson
Department of Computer Science
National University of Singapore
peter.robinson@monoid.at

CP6

Sequential Random Permutation, List Contraction and Tree Contraction Are Highly Parallel

We show that simple sequential randomized iterative algorithms for random permutation, list contraction, and tree contraction are highly parallel. In particular, if iterations of the algorithms are run as soon as all of their dependencies have been resolved, the resulting computations have logarithmic depth (parallel time) with high probability. Our proofs make an interesting connection between the dependence structure of two of the problems and random binary trees. Building upon this analysis, we describe linear-work, polylogarithmic-depth algorithms for the three problems. Although asymptotically no better than the many prior parallel algorithms for the given problems, their advantages include very simple and fast implementations, and returning the same result as the sequential algorithm. Experiments on a 40-core machine show reasonably good performance relative to the sequential algorithms.

Guy Blelloch
Computer Science Department
Carnegie Mellon University
guyb@cs.cmu.edu

Jeremy Fineman
Georgetown University
jfineman@cs.georgetown.edu

Phillip Gibbons
Intel Labs and Carnegie Mellon University
gibbons@cs.cmu.edu

Yan Gu, Julian Shun
Carnegie Mellon University
yan.gu@cs.cmu.edu, jshun@cs.cmu.edu

CP6

$(2\Delta - 1)$ -Edge-Coloring Is Much Easier Than Maximal Matching in the Distributed Setting

Graph coloring is a central problem in distributed computing. Both vertex- and edge-coloring problems have been extensively studied in this context. In this paper we show that a $(2\Delta - 1)$ -edge-coloring can be computed in time smaller than $\log^\epsilon n$ for any $\epsilon > 0$, specifically, in $e^{O(\sqrt{\log \log n})}$ rounds. This establishes a separation between the $(2\Delta - 1)$ -edge-coloring and Maximal Matching problems, as the latter is known to require $\Omega(\sqrt{\log n})$ time. We devise a $(1 + \epsilon)\Delta$ -edge-coloring algorithm for any constant $\epsilon > 0$. The running time of this algorithm is $O(\log^* \Delta + \frac{\log n}{\Delta^{1-\epsilon(1)}})$. For $\Delta = (\log n)^{1+\Omega(1)}$ the running time of our algorithm is only $O(\log^* n)$. This constitutes a drastic improvement of the previous logarithmic bound. Our result for $(2\Delta - 1)$ -edge-coloring also follows from our more general results concerning $(1 - \epsilon)$ -locally sparse graphs.

Michael Elkin
Ben-Gurion University

elkinm@cs.bgu.ac.il

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

Hsin-Hao Su
University of Michigan
EECS Department
hsinhao@umich.edu

CP7

Learning from Satisfying Assignments

This paper studies the problem of learning “low-complexity” probability distributions over the Boolean hypercube $\{-1, 1\}^n$. In our model the learning algorithm is given uniform random satisfying assignments of an unknown $f \in \mathcal{C}$ and its goal is to output a high-accuracy approximation of the uniform distribution over $f^{-1}(1)$. We show that the two most widely studied classes of Boolean functions in learning theory — LTFs and DNFs — have efficient distribution learning algorithms in our model.

Ilias Diakonikolas
University of Edinburgh
iliadiakonikolas@gmail.com

Anindya De
DIMACS/IAS
anindya@math.ias.edu

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

CP7

Approximate Resilience, Monotonicity, and the Complexity of Agnostic Learning

A function f is d -resilient if all its Fourier coefficients of degree at most d are zero, i.e. f is uncorrelated with all low-degree parities. We study the notion of *approximate resilience* of Boolean functions, where we say that f is α -approximately d -resilient if f is α -close to a $[-1, 1]$ -valued d -resilient function in ℓ_1 distance. We show that approximate resilience essentially characterizes the complexity of agnostic learning of a concept class \mathcal{C} over the uniform distribution. Focusing on monotone Boolean functions, we exhibit the existence of near-optimal α -approximately $\tilde{\Omega}(\alpha\sqrt{n})$ -resilient monotone functions for all $\alpha > 0$. Our constructions are based on general resilience analysis and amplification techniques we introduce. These structural results, together with the characterization, imply nearly optimal lower bounds for agnostic learning of monotone juntas.

Dana Dachman-Soled
University of Maryland
danadach@ece.umd.edu

Vitaly Feldman
IBM Research - Almaden
vitaly.edu@gmail.com

Li-Yang Tan
Columbia University
liyong@cs.columbia.edu

Andrew Wan
Simons Institute, UC Berkeley
atw12@columbia.edu

Karl Wimmer
Duquesne University
wimmerk@duq.edu

CP7

Robust Probabilistic Inference

Robust probabilistic inference is an extension of probabilistic inference, where some of the observations are adversarially corrupted. We model it as a zero-sum game between the adversary, who can select a modification rule, and the predictor, who wants to accurately predict the state of nature. Given a black-box access to a Bayesian inference in the classic (adversary-free) setting, our near optimal policy runs in polynomial time in the number of observations and the number of possible modification rules.

Yishay Mansour
Tel Aviv University
Microsoft Research, Hertzelia
mansour@tau.ac.il

Aviad Rubinfeld
UC Berkeley
aviad@eecs.berkeley.edu

Moshe Tennenholtz
Microsoft Research
Technion - Israel Institute of Technology
moshet@ie.technion.ac.il

CP7

Learning Privately with Labeled and Unlabeled Examples

Inspired by the (non-private) models of semi-supervised and active learning, where the focus is on the sample complexity of *labeled* examples, we consider *private* semi-supervised learners that operate on a random sample where only a (hopefully small) portion of this sample is labeled. The learners have no control over which of the sample elements are labeled. Our main result is that the labeled sample complexity of *private* learners is characterized by the VC dimension.

Amos Beimel, Kobbi Nissim, Uri Stemmer
Ben Gurion University
beimel@cs.bgu.ac.il, kobbi@cs.bgu.ac.il, stemmer@cs.bgu.ac.il

CP7

Linear-Programming Based Approximation Algorithms for Multi-Vehicle Minimum Latency Problems

We consider various *multi-vehicle versions of the minimum latency problem*. There is a fleet of k vehicles located at one or more depot nodes, and we seek a collection of routes for these vehicles that visit all nodes so as to minimize the total latency incurred, which is the sum of the client waiting times. We obtain an 8.497-approximation for the version where vehicles may be located at multiple depots and a 7.183-approximation for the version where all vehicles are located at the same depot, both of which are the first im-

improvements on this problem in a decade. Perhaps more significantly, our algorithms exploit various LP-relaxations for minimum-latency problems. We show how to effectively leverage two classes of LPs—*configuration LPs* and *bidirected LP-relaxations*—that are often believed to be quite powerful but have only sporadically been effectively leveraged for network-design and vehicle-routing problems.

Ian Post, [Chaitanya Swamy](#)
University of Waterloo
ian@ianpost.org, cswamy@uwaterloo.ca

CP8

A New Characterization of Maximal Repetitions by Lyndon Trees

We give a new characterization of maximal repetitions (or runs) in strings, using a tree defined on recursive standard factorizations of Lyndon words, called the Lyndon tree. This leads to a very simple proof of $\rho(n) < 1.5n$, where $\rho(n)$ is the maximum number of runs in a string of length n . We also present a new, conceptually simple linear-time algorithm for computing all the runs in a string, not relying on the Lempel-Ziv factorization.

[Hideo Bannai](#)
Department of Informatics
Kyushu University
bannai@inf.kyushu-u.ac.jp

Tomohiro I
Department of Computer Science
TU Dortmund
tomohiro.i@cs.tu-dortmund.de

Shunsuke Inenaga, Yuto Nakashima, Masayuki Takeda,
Kazuya Tsuruta
Department of Informatics
Kyushu University
inenaga@inf.kyushu-u.ac.jp,
yuto.nakashima@inf.kyushu-u.ac.jp,
takeda@inf.kyushu-u.ac.jp,

CP8

Internal Pattern Matching Queries in a Text and Applications

We develop a data structure for internal pattern matching in a text, which finds in constant time the occurrences of one subword x in another subword y provided that $|y| = \mathcal{O}(|x|)$. The data structure has linear size and admits a linear-time construction algorithm. We then use it to answer queries about primitivity and periods of subwords, general substring compression, and cyclic equivalence of subwords. Efficient construction allows to improve the algorithm finding δ -subrepetitions (a generalized version of runs). Recently, our data structure has been applied for further queries about subwords of a given text.

[Tomasz Kociumaka](#), Jakub Radoszewski
University of Warsaw
kociumaka@mimuw.edu.pl, jrad@mimuw.edu.pl

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

Tomasz Walen

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

CP8

Sharp Bounds on Formation-Free Sequences

An (r, s) -formation is the concatenation of s permutations over an r -letter alphabet. Formation-free sequences are a generalization of standard Davenport-Schinzel sequences (where $r = 2$) and can be used to obtain good bounds on the extremal function of any forbidden subsequence. More recently, formation-free sequences have been applied to bounding the size of sets of permutations with fixed VC-dimension. Improving on earlier work by Klazar, Nivasch, and Cibulka and Kynčl, we prove sharp bounds on the length of (r, s) -formation-free sequences, for every r and s , and on sequences avoiding “doubled” (r, s) -formations.

[Seth Pettie](#)
University of Michigan
seth@pettie.net

CP8

Cell-Probe Bounds for Online Edit Distance and Other Pattern Matching Problems

We give new cell-probe lower bounds for Hamming distance and convolution in a stream. In contrast to previous results, these bounds are non-trivial even for single bit cells. Using these new ideas, we then show the first known time lower bounds for both online edit distance and LCS. Finally, for the online edit distance problem we give a cell-probe algorithm which is exponentially faster than the best known RAM algorithm.

Raphael Clifford, Markus Jalsenius, [Benjamin G. Sach](#)
University of Bristol, UK
clifford@cs.bris.ac.uk, markus.jalsenius@gmail.com,
sach@cs.bris.ac.uk

CP8

Wavelet Trees Meet Suffix Trees

We start by presenting a faster construction algorithm for the famous wavelet trees — a data structure with “myriad virtues”. We then combine wavelet and suffix trees in a clever way to develop a new data structure, a wavelet suffix tree, which operates on substrings of a string rather than single characters. Sample applications of wavelet suffix trees are: substring suffix rank, substring suffix selection, and computation of a run-length-encoded Burrows-Wheeler transform of a substring.

Maxim Babenko
National Research University Higher School of Economics
mbabenko@hse.ru

Paweł Gawrychowski
Max-Planck-Institut f
gawry1@gmail.com

Tomasz Kociumaka
University of Warsaw
kociumaka@mimuw.edu.pl

[Tatiana Starikovskaya](#)
National Research University Higher School of Economics

tstarikovskaya@hse.ru

CP9

Characterizing the Easy-to-Find Subgraphs from the Viewpoint of Polynomial-Time Algorithms, Kernels, and Turing Kernels

We study two fundamental problems related to finding subgraphs: given graphs G and H , test whether H is isomorphic to a subgraph of G , or determine the maximum number of vertex-disjoint H -subgraphs that can be packed in G . We investigate these problems when the graph H belongs to a fixed hereditary family \mathcal{F} . Our goal is to study which classes \mathcal{F} make the two problems tractable in one of the following senses: (a) (randomized) polynomial-time solvable, (b) admits a polynomial (many-one) kernel, or (c) admits a polynomial Turing kernel.

Bart Jansen

Eindhoven University of Technology
bjansen@gmail.com

Daniel Marx

Computer and Automation Research Institute, Hungarian Academy of Sciences (MTA SZTAKI)
dmarx@cs.bme.hu

CP9

The Parameterized Complexity of K-Biclique

Given a graph G and a parameter k , the k -Biclique problem asks whether G contains a complete bipartite subgraph $K_{k,k}$. We prove that k -Biclique is $W[1]$ -hard by giving an fpt-reduction from k -Clique to k -Biclique, thus solving this longstanding open problem. We also provide a probabilistic reduction, which indicates that k -Biclique has no $f(k) \cdot n^{o(\sqrt{k})}$ -time algorithm for any computable function f unless a randomized version of Exponential Time Hypothesis fails. Our result implies the dichotomy classification of the parameterized complexity of cardinality constrain satisfaction problem and the inapproximability of maximum k -intersection problem.

Bingkai Lin

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

CP9

Solving D-Sat Via Backdoors to Small Treewidth

For a CNF formula ϕ and integer t , a *strong backdoor set* to treewidth t is a set of variables such that *each* possible partial assignment τ to this set reduces ϕ to a formula whose incidence graph is of treewidth at most t . A *weak backdoor set* to treewidth t is a set of variables such that there is a partial assignment to this set that reduces ϕ to a *satisfiable* formula of treewidth at most t . Our main contribution is an algorithm that, given a d -CNF formula ϕ and an integer k , in time $2^{O(k)}|\phi|$, either finds a satisfying assignment of ϕ , or reports correctly that ϕ is not satisfiable, or concludes correctly that ϕ has no weak or strong backdoor set to treewidth t of size at most k . Prior to our work, such results were known only for the very special case of $t = 1$ (Gaspers and Szeider, ICALP 2012).

Fedor Fomin

Dep. of Informatics
University of Bergen

fomin@ii.uib.no

Daniel Lokshantov

University of Bergen
daniello@ii.uib.no

Neeldhara Misra

Indian Institute of Science
neeldhara@csa.iisc.ernet.in

Ramanujan M. Sridharan

University of Bergen
ramanujan.sridharan@ii.uib.no

Saket Saurabh

IMS
saket@imsc.res.in

CP9

An Exact Characterization of Tractable Demand Patterns for Maximum Disjoint Path Problems

We study the following general disjoint paths problem: given a supply graph G , a set $T \subseteq V(G)$ of terminals, a demand graph H on the vertices T , and an integer k , the task is to find a set of k pairwise vertex-disjoint valid paths, where we say that a path of the supply graph G is valid if its endpoints are in T and adjacent in the demand graph H . For a class \mathcal{H} of graphs, we denote by MAXIMUM DISJOINT \mathcal{H} -PATHS the restriction of this problem when the demand graph H is assumed to be a member of \mathcal{H} . We study the fixed-parameter tractability of this family of problems, parameterized by k . Our main result is a complete characterization of the FPT cases of MAXIMUM DISJOINT \mathcal{H} -PATHS for every hereditary class \mathcal{H} of graphs: it turns out that complexity depends on the existence of large induced matchings and large induced skew bicliques in the demand graph H .

Dániel Marx

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

Paul Wollan

University of Rome "La Sapienza"
paul.wollan@gmail.com

CP9

Tight Lower Bound for the Channel Assignment Problem

We study the complexity of the Channel Assignment problem. An open problem asks whether Channel Assignment admits an $O(c^n)$ -time algorithm, for a constant c independent of the weights on the edges. We answer this question in the negative i.e. we show that there is no $2^{o(n \log n)}$ -time algorithm solving Channel Assignment unless the Exponential Time Hypothesis fails. Note that the currently best known algorithm works in time $O^*(n!) = 2^{O(n \log n)}$ so our lower bound is tight.

Arkadiusz Socala

University of Warsaw

as277575@students.mimuw.edu.pl

khoa@cs.umd.edu

CP10**On Uniform Capacitated K-Median Beyond the Natural Lp Relaxation**

In this paper, we study the uniform capacitated k -median problem. Our result is an $\exp(O(1/\epsilon^2))$ -approximation algorithm for the problem that violates the cardinality constraint by a factor of $1 + \epsilon$, by considering a novel LP relaxation for the problem.

Shi LiToyota Technological Institute at Chicago
shili@ttic.edu**CP10****Dynamic Facility Location via Exponential Clocks**

We present a new LP-rounding algorithm for facility location problems, which yields the first constant approximation algorithm for the dynamic facility location problem. This problem is a generalization of the classic facility location problem, proposed by Eisenstat, Mathieu, and Schabanel to model the dynamics of evolving social/infrastructure networks. Our algorithm exhibits several properties distinguishing it from previous LP-roundings for facility location problems, which are crucial for its application to the dynamic problem.

Hyung-Chan An, Ashkan Norouzi-Fard, Ola Svensson
EPFLhyung-chan.an@epfl.ch, ashkan.norouzifard@epfl.ch,
ola.svensson@epfl.ch**CP10****An Improved Approximation for k -Median, and Positive Correlation in Budgeted Optimization**

k -median is a classic problem in combinatorial optimization. We improve upon the best known approximation ratio of Li and Svensson from $2.732 + \epsilon$ to $2.611 + \epsilon$ by improving two separate stages of their algorithm. We do this by exploiting the structure of their worst case instance. We also show a general dependent rounding technique which bounds positive as well as negative correlation, and allows us to decrease the run-time dependence on ϵ .

Jaroslaw Byrka
University of Wrocław
jby@ii.uni.wroc.plThomas Pensyl
University of Maryland, College Park
tpensyl@cs.umd.eduBartosz Rybicki
University of Wrocław
bry@ii.uni.wroc.plAravind Srinivasan
University of Maryland
srin@cs.umd.eduKhoa Trinh
University of Maryland, College Park**CP10****Bi-Factor Approximation Algorithms for Hard Capacitated K-Median Problems**

In the classical k -median problem the goal is to select a subset of at most k facilities in order to minimize the total cost of opened facilities and established connections between clients and opened facilities. We consider the capacitated version of the problem, where a single facility may only serve a limited number of clients. It is known that violating capacities by a factor of $2 + \epsilon$ is sufficient to obtain constant factor approximation of the connection cost. We substantially extend this result in the following two directions. First, we extend the $2 + \epsilon$ capacity violating algorithm to the more general k -facility location problem with uniform capacities, where opening facilities incurs a location specific opening cost. Second, we obtain a constant factor approximation algorithm for the k -median problem also in the case of non-uniform capacities with a slightly bigger violation factor of $3 + \epsilon$.

Jaroslaw Byrka
University of Wrocław
jby@ii.uni.wroc.plKrzysztof Fleszar
University of Würzburg
krzysztof.fleszar@uni-wuerzburg.deBartosz Rybicki
University of Wrocław
bry@ii.uni.wroc.plJoachim Spoerhase
University of Würzburg
joachim.spoerhase@uni-wuerzburg.de**CP10****Improved Region-Growing and Combinatorial Algorithms for k -Route Cut Problems**

We study the k -route generalizations of various cut problems, the most general of which is k -route multicut, wherein we have r source-sink pairs and the goal is to delete a minimum-cost set of edges to reduce the edge-connectivity of every pair to below k . We present various approximation and hardness results that improve the state-of-the-art for these problems in several cases. Our algorithms are based on combinatorial techniques and on a new powerful region-growing approach.

Guru Guruganesh
Carnegie Mellon University
prashanth818@gmail.comLaura Sanita, Chaitanya Swamy
University of Waterloo
lsanita@uwaterloo.ca, cswamy@uwaterloo.ca**CP11****A Linear-Size Logarithmic Stretch Path-Reporting Distance Oracle for General Graphs**

For any n -vertex undirected graph $G = (V, E)$ and $k = 1, 2, \dots$, Thorup-Zwick's distance oracle has size

$O(kn^{1+1/k})$. It produces paths with stretch at most $2k - 1$, and its query time is, as was recently shown by Chechik, $O(1)$. A major drawback of the Thorup-Zwick oracle is that its size is $\Omega(n \cdot \log n)$. Mendel and Naor devised an oracle with space $O(n^{1+1/k})$ and stretch $O(k)$, but their oracle is not path-reporting. In this paper we devise a path-reporting distance oracle with size $O(n^{1+1/k})$, stretch $O(k)$ and query time $O(n^\epsilon)$, for any constant $\epsilon > 0$.

Michael Elkin

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

Seth Pettie

University of Michigan, Ann Arbor
pettie@gmail.com

CP11

Set Membership with a Few Bit Probes

We consider the bit-probe complexity of the set membership problem, where a set S of size at most n from a universe of size m is to be represented as a short bit vector in order to answer membership queries of the form “Is x in S ?” by adaptively probing the bit vector at t places. Let $s(m, n, t)$ be the minimum number of bits of storage needed for such a scheme. Improving the results of Alon and Feige (2009), we obtain better bounds for $s(m, n, 2)$ and $s(m, n, 3)$.

Mohit Garg, Jaikumar Radhakrishnan

Tata Institute of Fundamental Research
bootup.c@gmail.com, jaikumar@tifr.res.in

CP11

Approximate Range Emptiness in Constant Time and Optimal Space

We study a generalization of Bloom Filters. The ARE problem asks to represent elements from a universe in order to answer emptiness queries on a range, with a probability of false positives allowed. We provide space-error tradeoffs, and then show how one can achieve constant query time while matching the space lower bound up to a lower order additive term. This result is achieved through a novel succinct data structure for (exact Id) range emptiness/reporting.

Mayank Goswami

Max-Planck Institute for Informatics
Algorithms and Complexity.
gmayank@mpi-inf.mpg.de

CP11

Deterministic Fully Dynamic Data Structures for Vertex Cover and Matching

We present the first deterministic data structures for maintaining approximate minimum vertex cover and maximum matching in a fully dynamic graph in $o(\sqrt{m})$ time per update. In particular, for minimum vertex cover we provide deterministic data structures for maintaining a $(2 + \epsilon)$ approximation in $O(\log n / \epsilon^2)$ amortized time per update. For maximum matching, we show how to maintain a $(3 + \epsilon)$ approximation in $O(m^{1/3} / \epsilon^2)$ amortized time per update, and a $(4 + \epsilon)$ approximation in $O(m^{1/3} / \epsilon^2)$ worst-case time

per update.

Sayan Bhattacharya

University of Vienna
Vienna, Austria
jucse.sayan@gmail.com

Monika Henzinger

University of Vienna
mhenzinger@gmail.com

Giuseppe Italiano

University of Rome, Tor Vergata
Rome, Italy
italiano@disp.uniroma2.it

CP11

The Amortized Cost of Finding the Minimum

We obtain an essentially optimal tradeoff between the *amortized* cost of the three basic priority queue operations insert, delete and find-min in the comparison model. More specifically, we show that

$$A(\text{find-min}) = \Omega\left(\frac{n}{(2 + \epsilon)^{A(\text{insert}) + A(\text{delete})}}\right),$$

$$A(\text{find-min}) = O\left(\frac{n}{2^{A(\text{insert}) + A(\text{delete})}} + \log n\right),$$

for any fixed $\epsilon > 0$, where n is the number of items in the priority queue and $A(\text{insert})$, $A(\text{delete})$ and $A(\text{find-min})$ are the *amortized* costs of the insert, delete and find-min operations, respectively. Our lower bound holds even if *randomization* is allowed. Surprisingly, such fundamental bounds on the amortized cost of the operations were not known before. Brodal, Chaudhuri and Radhakrishnan, obtained similar bounds for the *worst-case* complexity of find-min

Haim Kaplan, Or Zamir, Uri Zwick

Tel Aviv University
haimk@tau.ac.il, orzamir@mail.tau.ac.il, zwick@tau.ac.il

CP12

The Simplex Algorithm Is NP-Mighty

We propose to classify the power of algorithms by the complexity of the problems that they can be used to solve. Instead of restricting to the problem an algorithm was designed to solve, we include problems that are *implicitly* during the algorithm’s execution. We show that the (Network) Simplex Method and the Successive Shortest Path Algorithm can be used to solve any problem in NP. As a consequence, we obtain several novel hardness results.

Yann Disser, Martin Skutella

TU Berlin
disser@math.tu-berlin.de, martin.skutella@math.tu-berlin.de

CP12

Towards a Characterization of Constant-Factor Approximable Min Csps

We study the approximability of Minimum Constraint Satisfaction Problems, where the goal is to minimize the number of unsatisfied constraints in a given CSP instance. The basic LP relaxation is known to be optimal for constant-factor approximating Min CSP unless the Unique Games

Conjecture fails. Using the algebraic approach to the CSP, we show that a new natural algebraic condition is necessary and sufficient for the finiteness of the integrality gap of this LP.

Victor Dalmau
University Pompeu Fabra, Barcelona
vdalmau@gmail.com

Andrei Krokhin
Durham University
andrei.krokhin@durham.ac.uk

Rajsekar Manokaran
KTH Stockholm
rajsekar@gmail.com

CP12

Strong Inapproximability Results on Balanced Rainbow-Colorable Hypergraphs

Consider a K -uniform hypergraph $H = (V, E)$. A coloring $c : V \rightarrow \{1, 2, \dots, k\}$ with k colors is *rainbow* if every hyperedge e contains at least one vertex from each color, and is called *perfectly balanced* when each color appears the same number of times. A simple polynomial-time algorithm finds a 2-coloring if H admits a perfectly balanced rainbow k -coloring. For a hypergraph that admits an *almost balanced rainbow* coloring, we prove that it is NP-hard to find an independent set of size ϵ , for any $\epsilon > 0$. Consequently, we cannot *weakly color* (avoiding monochromatic hyperedges) it with $O(1)$ colors. With $k = 2$, it implies strong hardness for discrepancy minimization of systems of bounded set-size. We give a *recipe* for converting a promising test distribution and a suitable choice of an outer PCP to a hardness result, and use this recipe to prove additional results almost in a black-box manner.

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

CP12

The Matching Polytope Does Not Admit Fully-Polynomial Size Relaxation Schemes

Rothvo established that every linear program expressing the matching polytope has exponential many inequalities. We generalize this to sharp bounds on the polyhedral inapproximability of the matching polytope: for $0 < \epsilon < 1$, every polyhedral $(1 + \epsilon/n)$ -approximation requires exponential many inequalities, where n is the number of vertices. Thus matching is the first problem in P, whose natural linear encoding does not admit a fully polynomial-size relaxation scheme (the polyhedral version of FPTAS).

Gábor Braun
ISYE Georgia Institute of Technology
gabor.braun@isye.gatech.edu

Sebastian Pokutta
Friedrich-Alexander-Universität Erlangen-Nürnberg
sebastianpokutta@me.com

CP12

Sperner's Colorings, Hypergraph Labeling Prob-

lems and Fair Division

We prove three results about colorings of the simplex reminiscent of Sperner's Lemma, with applications in hardness of approximation and fair division. First, we prove a coloring lemma which implies an optimal Unique-Games hardness of $(k - 1 - \epsilon)$ -approximation for the Hypergraph Labeling problem with label set $[k]$. We also show that a $(k - 1)$ -approximation can be achieved. Second, we show a Sperner-admissible labeling of the simplex with an application in the context of fair division. Third, we prove that there are subdivisions of the simplex with a fractional labeling such that every hyperedge in the subdivision uses only labelings with 1 or 2 colors. This means that a natural LP cannot distinguish instances of Hypergraph Labeling that can be labeled so that every hyperedge uses at most 2 colors, and instances that must have a rainbow hyperedge. We prove that this problem is indeed NP-hard for $k = 3$.

Maryam Mirzakhani
Stanford University
mmirzakh@stanford.edu

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

CP13

Tighter Low-Rank Approximation Via Sampling the Leveraged Element

We propose a new randomized algorithm for computing a low-rank approximation to a given matrix. Taking an approach different from existing literature, our method first involves a specific biased sampling, with an element being chosen based on the leverage scores of its row and column, and then involves weighted alternating minimization over the factored form, to minimize error only on these samples. Our method can leverage input sparsity, yet produce approximations in *spectral* norm.

Srinadh Bhojanapalli
UT Austin
bsrinadh@utexas.edu

Prateek Jain
Microsoft Research
prajain@microsoft.com

Sujay Sanghavi
Electrical and Computer Engineering
University of Texas at Austin
sanghavi@mail.utexas.edu

CP13

Online Principal Component Analysis

We consider the online version of the Principal Component Analysis problem: given a set of vectors $x_1, \dots, x_n \in R^{d \times n}$ and a target dimension $k < d$, return a set of vectors $y_1, \dots, y_n \in R^{k \times n}$ that minimize a certain objective. In the online setting, the vectors x_t are presented to the algorithm one by one; for every x_t the algorithm must output a y_t before receiving x_{t+1} . In this work, we present new algorithms for online PCA.

Christos Boutsidis
Yahoo! Labs
boutsidis@yahoo-inc.com

Dan Garber
Technion
dangar@tx.technion.ac.il

Zohar Karnin
Yahoo Labs
zkarnin@gmail.com
karnin

Edo Liberty
Yahoo Labs
edo.liberty@gmail.com

CP13

The Polyhedron-Hitting Problem

We consider polyhedral versions of Kannan and Lipton's Orbit Problem (STOC '80 and JACM '86)—determining whether a target polyhedron V may be reached from a starting point x under repeated applications of a linear transformation A in an ambient vector space Q^m . In the context of program verification, very similar reachability questions were also considered and left open by Lee and Yannakakis in (STOC '92). We present what amounts to a complete characterisation of the decidability landscape for the Polyhedron-Hitting Problem, expressed as a function of the dimension m of the ambient space, together with the dimension of the polyhedral target V : more precisely, for each pair of dimensions, we either establish decidability, or show hardness for longstanding number-theoretic open problems.

Ventsislav Chonev
Department of Computer Science
University of Oxford
chonev@cs.ox.ac.uk

Joel Ouaknine, James Worrell
Department of Computer Science
Oxford University
joel@cs.ox.ac.uk, jbw@cs.ox.ac.uk

CP13

On Termination of Integer Linear Loops

A fundamental problem in program verification concerns the termination of simple linear loops of the form:

$$x \leftarrow u; \quad \text{while } Bx \geq c \text{ do } x \leftarrow Ax + a$$

where x is a vector of variables, u , a , and c are integer vectors, and A and B are integer matrices. Assuming the matrix A is diagonalisable, we give a decision procedure for the problem of whether, for all initial integer vectors u , such a loop terminates. The correctness of our algorithm relies on sophisticated tools from algebraic and analytic number theory, Diophantine geometry, and real algebraic geometry. To the best of our knowledge, this is the first substantial advance on a 10-year-old open problem of Tiwari (2004) and Braverman (2006).

Joel Ouaknine
Department of Computer Science
Oxford University
joel@cs.ox.ac.uk

Joao Sousa Pinto
University of Oxford
jspinto@cs.ox.ac.uk

James Worrell
Department of Computer Science
Oxford University
jbw@cs.ox.ac.uk

CP13

Sketching for M-Estimators: A Unified Approach to Robust Regression

We give algorithms for M -estimators $\min_x \|Ax - b\|_G$, where $A \in \mathbf{R}^{n \times d}$, $b \in \mathbf{R}^n$, and $\|y\|_G$ for $y \in \mathbf{R}^n$ is specified by a non-negative function G with $\|y\|_G \equiv \sum_i G(y_i)$. We give a $(1 + \epsilon)$ -approximation to the Huber measure in $O(\text{nnz}(A) \log n + \text{poly}(d(\log n)/\epsilon))$ time, where $\text{nnz}(A)$ is the number of non-zero entries of A . We then solve the M -estimation problem in $\text{nnz}(A) + \text{poly}(d)$ time for any convex G with at most quadratic growth, finding a constant factor approximation with high probability.

Kenneth Clarkson
IBM
klclarks@us.ibm.com

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

CP14

A Stable Marriage Requires Communication

The Gale-Shapley algorithm for the Stable Marriage Problem is known to take $\Theta(n^2)$ steps to find a stable marriage in the worst case (with n women and n men). Using a reduction to the communication complexity of the disjointness problem, we show that $\Omega(n^2)$ Boolean queries of any type are indeed required to find a stable marriage. This lower bound also generalizes to randomized algorithms, finding approximately stable marriages, and several related problems.

Yannai Gonczarowski
Departments of Mathematics and Computer Science
The Hebrew University of Jerusalem
yannai@gonch.name

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

Rafail Ostrovsky
UCLA
rafail@cs.ucla.edu

Will Rosenbaum
Department of Mathematics
University of California, Los Angeles
will.rosenbaum@gmail.com

CP14

The Value 1 Problem Under Finite-Memory Strategies for Concurrent Mean-Payoff Games

The talk will be about two-player, zero-sum concurrent mean-payoff games with payoffs in $[0,1]$. We present an algorithm for finding the set of states, such that for play starting from any state in that set, for each $\epsilon > 0$, player 1 has a finite memory strategy ensuring mean-payoff $1 - \epsilon$,

against all strategies for player 2. We also argue that in such states player 1 has a strategy that does not use memory.

Krishnendu Chatterjee, [Rasmus Ibsen-Jensen](#)

IST Austria

krishnendu.chatterjee@ist.ac.at,
jensen@gmail.com

rasmusibsen-

CP14

Approximating the Best Nash Equilibrium in $n^{o(\log N)}$ -Time Breaks the Exponential Time Hypothesis

We study the computational complexity of finding an ϵ -approximate Nash equilibrium with good social welfare. Hazan and Krauthgamer and subsequent improvements showed that finding an ϵ -approximate Nash equilibrium with good social welfare in a two player game and many variants of this problem is at least as hard as finding a planted clique of size $O(\log n)$ in the random graph $\mathcal{G}(n, 1/2)$. We show that any polynomial time algorithm that finds an ϵ -approximate Nash equilibrium with good social welfare refutes (the worst-case) Exponential Time Hypothesis by Impagliazzo and Paturi, confirming the recent conjecture by Aaronson, Impagliazzo and Moshkovitz. Specifically, it would imply a $2^{\tilde{O}(n^{1/2})}$ algorithm for SAT. Our lower bound matches the quasi-polynomial time algorithm by Lipton, Markakis and Mehta for solving the problem.

Mark Braverman, [Young Kun Ko](#), Omri Weinstein
Princeton University

mbraverm@cs.princeton.edu, yko@cs.princeton.edu, oweinste@cs.princeton.edu

CP14

Robust Price of Anarchy Bounds Via Lp and Fenchel Duality

In this paper, we develop a new framework based on LP and Fenchel duality for bounding the robust price of anarchy for a large class of games. We use our framework to give to the *first* PoA bounds for temporal routing games on graphs and energy minimization games in machine scheduling. We demonstrate the wide applicability of our framework by giving new proofs of the PoA bounds for several classical games, such as congestion games.

[Janardhan Kulkarni](#)

Department of Computer Science, Duke University
Durham, NC, 27708.
kulkarni@cs.duke.edu

Vahab Mirrokni

Google Research
mirrokni@gmail.com

CP14

Perfect Bayesian Equilibria in Repeated Sales

A special case of Myerson's classic result describes the revenue-optimal equilibrium when a seller offers a single item to a buyer. We study a natural repeated sales extension of this model: a seller offers to sell a single fresh copy of an item to the same buyer every day via a posted price. The buyer's value for the item is unknown to the seller but is drawn initially from a publicly known distribution F

and remains the same throughout. One key aspect of this game is revelation of the buyer's type through his actions: while the seller might try to learn this value to extract more revenue, the buyer is motivated to hide it to induce lower prices. If the seller is able to commit to future prices, it is known that the best he can do is extract the Myerson optimal revenue each day. In a more realistic scenario, the seller is unable to commit and must play a perfect Bayesian equilibrium. We explore this setting without commitment and find several surprises.

Nikhil R. Devanur, Yuval Peres

Microsoft Research, Redmond
nikdev@microsoft.com, peres@microsoft.com

[Balasubramanian Sivan](#)

University of Wisconsin-Madison
bsivan@microsoft.com

CP15

A Dynamic Programming Framework for Non-Preemptive Scheduling Problems on Multiple Machines

In this paper, we consider a variety of scheduling problems where n jobs with release times are to be scheduled *non-preemptively* on a set of m identical machines. The problems considered are machine minimization, (weighted) throughput maximization and min-sum objectives such as (weighted) flow time and (weighted) tardiness. We develop a novel quasi-polynomial time dynamic programming framework that gives $O(1)$ -speed $O(1)$ -approximation algorithms for these problems.

[Sungjin Im](#)

University of California, Merced
sim3@ucmerced.edu

Shi Li

Toyota Technological Institute at Chicago
shili@ttic.edu

Benjamin Moseley

Washington University in St. Louis
bmoseley@wustl.edu

Eric Torng

Michigan State University
torng@msu.edu

CP15

New Approximations for Broadcast Scheduling Via Variants of α -Point Rounding

In the pull-based broadcast scheduling, there are n unit-sized pages of information available at the server. When the server transmits a page, all outstanding requests for the page are simultaneously satisfied. In this paper, we give new approximation results for two popular objectives, minimizing the maximum flow time and maximizing the throughput. In particular, we obtain the first PTAS for the former objective. Our key techniques for these improvements are novel variants of α -point rounding that can effectively reduce congestion in schedule.

[Sungjin Im](#)

University of California, Merced
sim3@ucmerced.edu

Maxim Sviridenko
 Yahoo! Labs
 sviri@yahoo-inc.com

CP15

On $(1, \epsilon)$ -Restricted Assignment Makespan Minimization

Makespan minimization on unrelated machines is a classic problem in approximation algorithms. In this paper we consider the $(1, \epsilon)$ -restricted assignment problem where each job is either heavy ($p_j = 1$) or light ($p_j = \epsilon$), for some parameter $\epsilon > 0$. Each job can be assigned to a specified set of machines. Our main result is a $(2 - \delta)$ -approximate *polynomial time* algorithm for the $(1, \epsilon)$ -restricted assignment problem for a fixed constant $\delta > 0$.

Deeparnab Chakrabarty
 Microsoft Research India
 dechakr@microsoft.com

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

Shi Li

Toyota Technological Institute at Chicago
 shili@ttic.edu

CP15

A Fully Polynomial-Time Approximation Scheme for Speed Scaling with Sleep State

We study deadline-based preemptive scheduling of jobs in a computing environment equipped with both dynamic speed scaling and sleep state capabilities, in order to minimize the total energy consumption. We present a fully polynomial-time approximation scheme for this problem. Previously, only a $4/3$ -approximation was known (Albers and Antoniadis, SODA 2012).

Antonios Antoniadis
 Max-Planck-Institut für Informatik
 aantonia@mpi-inf.mpg.de

Chien-Chung Huang
 Chalmers University
 villars@gmail.com

Sebastian Ott

Max-Planck-Institut für Informatik
 ott@mpi-inf.mpg.de

CP15

Rejecting Jobs to Minimize Load and Maximum Flow-Time

Online algorithms are usually analyzed using the notion of competitive ratio which compares the algorithm solution to that obtained by an online adversary for the worst possible input sequence. Often this measure turns out to be too pessimistic, and one popular approach especially for scheduling problems has been that of “resource augmentation” which was first proposed by Kalyanasundaram and Pruhs. Although resource augmentation has been very successful in dealing with a variety of objective functions, there are problems for which even a (arbitrary) constant speedup cannot lead to a constant competitive algorithm.

In this paper we propose a “rejection model” which requires no resource augmentation but which permits the online algorithm to not serve an epsilon-fraction of the requests. We consider the problems of minimizing the maximum load on any machine and minimizing (weighted) flow time in the restricted assignment setting.

Anamitra Roychoudhury

IBM Research-India
 anamchou@in.ibm.com

Syamantak Das
 Indian Institute of Technology Delhi
 sdas@cse.iitd.ac.in

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

CP16

Revealing Optimal Thresholds for Generalized Secretary Problem Via Continuous Lp: Impacts on Online K -Item Auction and Bipartite K -Matching with Random Arrival Order

We consider the general (J, K) -secretary problem. An algorithm observes the relative merits of arriving items and is allowed to make J selections. The objective is to maximize the expected number of items selected among the K best items. We give a construction procedure for an optimal algorithm involving JK thresholds via a continuous linear programming method. We further provide new results on online K -item auction and bipartite K -matching with random arrival order.

T-H. Hubert Chan, Fei Chen, Shaofeng H.-C. Jiang
 The University of Hong Kong
 hubert@cs.hku.hk, fchen@cs.hku.hk, sfjiang@cs.hku.hk

CP16

Comparing Apples and Oranges: Query Tradeoff in Submodular Maximization

Fast algorithms for submodular maximization problems have a vast potential use. Our main result is a new algorithm for the general problem of maximizing a submodular function subject to a matroid independence constraint. Our algorithm is faster than the previous algorithm of Badanidiyuru and Vondrák (2014) when the rank is sub-linear in the size of the ground set. Furthermore, our algorithm establishes a surprising *tradeoff* between the number of value and independence oracle queries.

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

Moran Feldman
 EPFL
 moran.feldman@epfl.ch

Roy Schwartz
 Princeton University
 roysch@cs.princeton.edu

CP16

Online Submodular Maximization with Preemp-

tion

We study a natural online variant of submodular function maximization in which the algorithm has to maintain an online solution obeying certain constraints. Upon arrival of an element, the algorithm has to decide whether to accept the element and may preempt previously chosen elements. The goal is to maximize a submodular function over the maintained solution. We derive upper and lower bounds on the competitive ratio for the unconstrained case and for cardinality constraints.

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

Moran Feldman
 EPFL
 moran.feldman@epfl.ch

Roy Schwartz
 Princeton University
 roysch@cs.princeton.edu

CP16

A Simple $O(\log \log(\text{rank}))$ -Competitive Algorithm for the Matroid Secretary Problem

Only recently progress has been made in obtaining $o(\log(\text{rank}))$ -competitive algorithms for the matroid secretary problem. Chakraborty and Lachish (2012) presented a $O(\sqrt{\log(\text{rank})})$ -competitive procedure, and Lachish (2014) recently presented a $O(\log \log(\text{rank}))$ -competitive algorithm. Both algorithms are involved with complex analyses. Using different tools, we present a considerably simpler $O(\log \log(\text{rank}))$ -competitive algorithm. Our algorithm can be interpreted as a distribution over a simple type of matroid secretary algorithms which are easy to analyze.

Moran Feldman, [Ola Svensson](#)
 EPFL
 moran.feldman@epfl.ch, ola.svensson@epfl.ch

Rico Zenklusen
 ETH Zurich
 ricoz@math.ethz.ch

CP16

Optimal Approximation for Submodular and Supermodular Optimization with Bounded Curvature

We design new approximation algorithms for the problems of optimizing submodular and supermodular functions with bounded curvature subject to a single matroid constraint. Our approach is based on modifications of the continuous greedy algorithm and non-oblivious local search, and allows us to approximately maximize the sum of a nonnegative, nondecreasing submodular function and a (possibly negative) linear function. We show how to reduce both submodular maximization and supermodular minimization to this general problem when the objective function has bounded total curvature. We prove that the approximation results we obtain are the best possible in the value oracle model, even in the case of a cardinality constraint. Finally, we give two concrete applications of our results in the settings of maximum entropy sampling,

and the column-subset selection problem.

Maxim Sviridenko
 Yahoo! Labs
 sviri@yahoo-inc.com

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

[Justin Ward](#)
 University of Warwick
 J.D.Ward@warwick.ac.uk

CP17

Streaming Algorithms for Estimating the Matching Size in Planar Graphs and Beyond

We consider the problem of estimating the size of a maximum matching when the edges are revealed in a streaming fashion. When the input graph is planar, we present a simple and elegant streaming algorithm that with high probability estimates the size of a maximum matching within a constant factor using $O(n^{2/3})$ space, where n is the number of vertices. The approach generalizes to the family of graphs that have bounded arboricity, which include graphs with an excluded constant-size minor. We furthermore design a reduction from the Boolean Hidden Matching Problem to show that there is no randomized streaming algorithm that estimates the size of the maximum matching to within a factor better than $3/2$ and uses only $o(n^{1/2})$ bits of space.

[Hossein Esfandiari](#)
 University of Maryland at College Park
 hossein@cs.umd.edu

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

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

Morteza Monemizadeh
 University of Frankfurt
 monemizadeh@em.uni-frankfurt.de

Krzysztof Onak
 IBM TJ Watson Research Center
 krzysztof@onak.pl

CP17

Streaming Lower Bounds for Approximating Max-Cut

We consider the problem of estimating the value of max cut in a graph in the streaming model of computation. At one extreme, there is a trivial 2-approximation for this problem that uses only $O(\log n)$ space, namely, count the number of edges and output half of this value as the estimate for max cut value. On the other extreme, a near-optimal solution to the max cut value can be obtained in $\tilde{O}(n)$ space by storing a cut sparsifier. An intriguing question is if poly-logarithmic space suffices to obtain a non-trivial approximation to the max-cut value (that is, beating the factor 2). Our main result is that any streaming algorithm that

breaks the 2-approximation barrier requires $\tilde{\Omega}(\sqrt{n})$ space even if the edges of the input graph are presented in random order. We also show that for any $\epsilon > 0$, any streaming algorithm that obtains a $(1 + \epsilon)$ -approximation when edges arrive in adversarial order requires $n^{1-O(\epsilon)}$ space.

Michael Kapralov
MIT

michael.kapralov@gmail.com

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

Madhu Sudan
Microsoft Research New England
madhu@mit.edu

CP17

Parameterized Streaming: Maximal Matching and Vertex Cover

In this paper, we introduce a new approach to handling graph streams, by instead seeking solutions for the parameterized versions of these problems. Here, we are given a parameter k and the objective is to decide whether there is a solution bounded by k . By combining kernelization techniques with randomized sketch structures, we obtain the first streaming algorithms for the parameterized versions of Maximal Matching and Vertex Cover.

Rajesh Chitnis
Department of Computer Science
University of Maryland at College Park, USA
rchitnis@cs.umd.edu

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

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

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

CP17

Decomposing a Graph Into Expanding Subgraphs

A paradigm that was successfully applied in the study of both pure and algorithmic problems in graph theory can be colloquially summarized as stating that any graph is close to being the disjoint union of expanders. Our goal in this paper is to show that in several of the instantiations of the above approach, the quantitative bounds that were obtained are essentially best possible. These instantiations include the Unique Games algorithm of Trevisan [FOCS '05], the eigenspace enumeration method of Arora, Barak and Steurer [FOCS '10], and a classical result of Lipton, Rose and Tarjan on graph separators.

Guy Moshkovitz, Asaf Shapira
Tel Aviv University

guymosko@post.tau.ac.il, asafico@tau.ac.il

CP17

A Polylogarithmic Space Deterministic Streaming Algorithm for Approximating Distance to Monotonicity

We give the first deterministic streaming algorithm that approximates the distance to monotonicity within a $1 + \epsilon$ factor for any fixed $\epsilon > 0$ and runs in space polylogarithmic in the length of the sequence and the range of the numbers. The best previous deterministic algorithm achieving the same approximation factor required space $\Omega(\sqrt{n})$. Previous polylogarithmic space algorithms were either randomized, or had approximation factor no better than 2.

Timothy R. Naumovitz, Michael Saks
Rutgers University
tbaumovi@math.rutgers.edu, saks@math.rutgers.edu

CP18

Combinatorial Algorithm for Restricted Max-Min Fair Allocation

We study the basic allocation problem of assigning resources to players so as to maximize fairness. This is one of the few natural problems that enjoys the intriguing status of having a better estimation algorithm than approximation algorithm. Indeed, a certain configuration-LP can be used to estimate the value of the optimal allocation to within a factor of $4 + \epsilon$. In contrast, however, the best known approximation algorithm for the problem has an unspecified large constant guarantee. In this paper we significantly narrow this gap by giving a 13-approximation algorithm for the problem.

Chidambaram Annamalai, Christos Kalaitzis, Ola Svensson
EPFL
chidambaram.annamalai@epfl.ch, christos.kalaitzis@epfl.ch, ola.svensson@epfl.ch

CP18

Capacity of Interactive Communication over Erasure Channels and Channels with Feedback

We consider interactive communication performed over two simple types of noisy channels: binary error channels with noiseless feedback and binary erasure channels. In both cases, the noise model is adversarial. Assuming at most ϵ -fraction of the bits can be corrupted, we show coding schemes that simulate any alternating interactive protocol with rate $1 - \Theta(H(\epsilon))$. All our simulations are simple, randomized, and computationally efficient. The rates of our coding schemes stand in contrast to the interactive communication rates supported by random or adversarial error channels without feedback. Such a gap has no equivalent in the standard one-way communication setting.

Ran Gelles
Princeton University
rgelles@cs.princeton.edu

Bernhard Haeupler
Carnegie Mellon U.

haeupler@cs.cmu.edu

CP18

Lp/sdp Hierarchy Lower Bounds for Decoding Random Ldpc Codes

Random (d_v, d_c) -regular LDPC codes are well-known to achieve the Shannon capacity of the binary symmetric channel (for sufficiently large d_v and d_c) under exponential time decoding. However, polynomial time algorithms are only known to correct a much smaller fraction of errors. One of the most powerful polynomial-time algorithms with a formal analysis is the LP decoding algorithm of Feldman et al. which is known to correct an $\Omega(1/d_c)$ fraction of errors. In this work, we show that fairly powerful extensions of LP decoding, based on the Sherali-Adams and Lasserre hierarchies, fail to correct much more errors than the basic LP-decoder. In particular, we show that for any (resp. infinitely many) values of d_v and d_c , a linear number of rounds of the Sherali-Adams LP (resp. Lasserre SDP) hierarchy cannot correct more than an $O(1/d_c)$ fraction of errors on a random (d_v, d_c) -regular LDPC code.

Badih Ghazi

MIT

badih@mit.edu

Euiwoong Lee

Carnegie Mellon University

euiwoonl@cs.cmu.edu

CP18

On the Quickest Flow Problem in Dynamic Networks – A Parametric Min-Cost Flow Approach

We consider the Quickest Flow (QF) problem in dynamic networks with a single source s and a single sink t : given an amount of flow F , find the minimum time needed to send it from s to t , and the corresponding optimal flow over time. We introduce new mathematical formulations and derive optimality conditions for QF. Based on the optimality conditions, we develop a cost-scaling algorithm that leverages the parametric nature of the problem. The algorithm solves QF with integer arc costs in $O(nm \log(n^2/m) \log(nC))$ time, where n , m , and C are the number of nodes, arcs, and the maximum arc cost, respectively. Our algorithm runs in the same time bound as the cost-scaling algorithm by Goldberg and Tarjan for solving the min-cost flow problem. This result shows for the first time that QF can be solved within the same time bound as one of the fastest algorithms for the min-cost flow problem.

Patrick Jaillet

MIT

jaillet@mit.edu

Maokai Lin

Operations Research Center, MIT

lmk@csail.mit.edu

CP18

Limitations on Testable Affine-Invariant Codes in the High-Rate Regime

Locally testable codes (LTCs) of constant minimum (absolute) distance that allow the tester to make a nearly linear number of queries have become the focus of attention recently. The binary Reed-Muller code of block length N

and absolute distance d is testable with $O(N/d)$ queries, and has a dimension of $\approx N - (\log N)^{\log d}$. On the other hand, the smallest possible co-dimension for a distance d code (without any testability requirement) is $\approx \frac{d}{2} \log N$, achieved by BCH codes. This raises the natural question of understanding the optimal co-dimension of a distance d LTC. One promising approach for constructing LTCs is to focus on affine-invariant codes. In this work, we show that a recent family of affine-invariant constructions is essentially optimal among linear affine-invariant LTCs that contain the Reed-Muller code of the appropriate degree.

Venkatesan Guruswami

Carnegie Mellon University

guruswami@cmu.edu

Madhu Sudan

Microsoft Research New England

madhu@mit.edu

Ameya Velingker

Carnegie Mellon University

Computer Science Department

avelingk@cs.cmu.edu

Carol Wang

Carnegie Mellon University

wangc@cs.cmu.edu

CP19

Fast Algorithms for Online Stochastic Convex Programming

We introduce the online stochastic Convex Programming(CP) problem, a very general version of stochastic online problems which allows arbitrary concave objectives and convex feasibility constraints. Online stochastic packing and covering form special cases of online stochastic CP. We present fast algorithms for these problems, which achieve near-optimal regret guarantees for both the i.i.d. and the random permutation models of stochastic inputs. Our techniques make explicit the connection of primal-dual paradigm and online learning to online stochastic CP

Shipra Agrawal

Microsoft

shipra@microsoft.com

Nikhil R. Devanur

Microsoft Research

nikdev@microsoft.com

CP19

Using Optimization to Break the Epsilon Barrier: A Faster and Simpler Width-Independent Algorithm for Solving Positive Linear Programs in Parallel

We study the design of nearly-linear-time algorithms for approximately solving positive linear programs. Both the parallel and the sequential deterministic versions of these algorithms require $\tilde{O}(\varepsilon^{-4})$ iterations (i.e., have a $O(T^{-1/4})$ convergence rate), a dependence that has not been improved since the introduction of these methods in 1993 by Luby and Nisan. Moreover, previous algorithms and their analyses rely on update steps and convergence arguments that are combinatorial in nature, but do not seem to arise naturally from an optimization viewpoint. In this paper,

we leverage insights from optimization theory to construct a novel algorithm that breaks the $\tilde{O}(\varepsilon^{-4})$ barrier. Our algorithm has a simple analysis and a clear motivation. Our work introduces a number of techniques, such as the combined application of gradient and mirror descent, and a truncated version of the multiplicative weight update, which may be of independent interest.

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

Lorenzo Orecchia
MIT
orecchia@mit.edu

CP19

Online Network Design Algorithms Via Hierarchical Decompositions

We develop a new approach for online network design and obtain improved competitive ratios for several problems. At the heart of our work is a novel application of embeddings into hierarchically well-separated trees (HSTs) to the analysis of online network design algorithms. Using this approach, we obtain simple greedy-style online algorithms and analyses. In this work, we apply it to Steiner network (with edge-duplication), rent-or-buy, connected facility location and prize-collecting Steiner forest.

Seeun W. Umboh
University of Wisconsin-Madison
seeun.umbogh@gmail.com

CP19

Online Stochastic Matching with Unequal Probabilities

We study a variant of online bipartite matching in which edges are labeled with the probability with which a match along this edge ‘succeeds’. Prior work gave a competitive ratio of 0.567 for vanishing probabilities when all probabilities in the graph are equal. We consider general, unequal probabilities, and show a ratio of 0.534 for the vanishing case with a different algorithmic approach: maximizing marginal contributions while carefully controlling the level of adaptiveness.

Aranyak Mehta
Google
aranyak@google.com

Bo Waggoner
Harvard
bwaggoner@fas.harvard.edu

Morteza Zadimoghaddam
Google
zadim@google.com

CP19

The Optimal Absolute Ratio for Online Bin Packing

We present an online bin packing algorithm with absolute competitive ratio $5/3$, which is optimal.

Janos Balogh, Jozsef Bekesi

University of Szeged
Hungary
balogh@jgypk.u-szeged.hu, bekesi@jgypk.u-szeged.hu

Gyorgy Dosa
University of Pannonia
Hungary
dosagy@almos.vein.hu

Jiri Sgall
Computer Science Institute
Charles University, Czech Republic
sgall@iuuk.mff.cuni.cz

Rob van Stee
Max Planck Institute for Informatics
rvs4@le.ac.uk

CP20

A Unified Framework for Clustering Constrained Data Without Locality Property

We consider a class of constrained clustering problems in R^d , where d could be rather high. A common feature of these problems is that their optimal clusterings no longer have the locality property (due to the additional constraints), which is a key property required by many existing clustering algorithms. To overcome this difficulty, we present a unified framework, and improve considerably the best known results.

Hu Ding, Jinhui Xu
Department of Computer Science and Engineering
State University of New York at Buffalo
huding@buffalo.edu, jinhui@buffalo.edu

CP20

On the Richter-Thomassen Conjecture About Pairwise Intersecting Closed Curves

A long standing conjecture of Richter and Thomassen states that the total number of intersection points between any n pairwise intersecting simple closed Jordan curves in the plane is at least $(1 - o(1))n^2$. We confirm the above conjecture in several important cases, including the case (1) when all curves are convex, and (2) when the family of curves can be partitioned into two equal classes such that each curve from the first class is touching every curve from the second class. An important ingredient of our proofs is the following statement: Let S be a family of the graphs of n continuous real functions defined on \mathbb{R} , no three of which pass through the same point. If there are nt pairs of touching curves in S , then the number of crossing points is $\Omega(nt\sqrt{\log t / \log \log t})$.

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

Natan Rubin
Department of Computer Science
Ben Gurion University
rubinnat.ac@gmail.com

Gábor Tardos
Rényi Institute, Budapest

tardos@renyi.hu

CP20

Density and Regularity Theorems for Semi-Algebraic Hypergraphs

A k -uniform semi-algebraic hypergraph H is a pair (P, E) , where P is a subset of R^d and E is a collection of k -tuples $\{p_1, \dots, p_k\} \subset P$ such that $(p_1, \dots, p_k) \in E$ if and only if the kd coordinates of the p_i -s satisfy a boolean combination of a finite number of polynomial inequalities. Improving a theorem of Fox, Gromov, Lafforgue, Naor, and Pach, we establish a polynomial regularity lemma: For any $\epsilon > 0$, the vertex set of every k -uniform semi-algebraic hypergraph can be partitioned into at most $(1/\epsilon)^c$ parts such that all but an ϵ -fraction of the k -tuples of parts are complete or empty. We also establish an improved lower bound on the best constant $\delta > 0$ such that the vertex classes of every dense k -partite k -uniform semi-algebraic hypergraph $H = (P_1 \cup \dots \cup P_k, E)$ have $\delta|P_i|$ -element subsets that induces a complete k -partite hypergraph. We give 3 geometric applications of our results.

Jacob Fox
MIT
fox@math.mit.edu

Janos Pach
Courant Institute
pach@cims.nyu.edu

Andrew Suk
University of Illinois at Chicago
suk@math.uic.edu

CP20

Approximation Schemes for Partitioning: Convex Decomposition and Surface Approximation

Building on the discovery of new types of separators by Adamaszek and Wiese, there has recently been significant progress on approximability of planar packing problems as well as certain covering problems. Following these developments, we revisit two NP-hard geometric partitioning problems - convex decomposition and surface approximation. These problems bring up additional issues that are worth examining in the context of the wider applicability of the separator methodology; we show how these issues can be addressed.

Sayan Bandyapadhyay
Department of Computer Science
The University of Iowa
sayan-bandyapadhyay@uiowa.edu

Santanu Bhowmick
University of Iowa
santanu-bhowmick@uiowa.edu

Kasturi Varadarajan
Department of Computer Science
The University of Iowa
kasturi-varadarajan@uiowa.edu

CP20

A Quasi-Ptas for the Two-Dimensional Geometric

Knapsack Problem

We consider the two-dimensional geometric knapsack problem. Given a collection of rectangular axis-parallel items with weights, we want to find a maximum weight subset of the items that can be packed into a rectangular knapsack. The goal is to compute the optimal collection of items together with a feasible packing. For the general case the best known result is a $(2 + \epsilon)$ -approximation algorithm, while the only hardness result is NP-hardness. Our main result is a $(1 + \epsilon)$ -approximation algorithm that runs in quasipolynomial time, provided that the input data consists of (quasi-)polynomially bounded integers.

Anna Adamaszek, Andreas Wiese
MPII Saarbrücken
anna@mpi-inf.mpg.de, awiese@mpi-inf.mpg.de

CP21

The Switch Markov Chain for Sampling Irregular Graphs

The problem of efficiently sampling from a set of (undirected) graphs with a given degree sequence has many applications. One approach to this problem uses a simple Markov chain, which we call the switch chain, to perform the sampling. The switch chain is known to be rapidly mixing for regular degree sequences. We prove that the switch chain is rapidly mixing for any degree sequence with minimum degree at least 1 and with maximum degree d_{\max} which satisfies $3 \leq d_{\max} \leq \frac{1}{4}\sqrt{M}$, where M is the sum of the degrees. The mixing time bound obtained is only an order of n larger than that established in the regular case, where n is the number of vertices.

Catherine Greenhill
University of New South Wales
School of Mathematics and Statistics
c.greenhill@unsw.edu.au

CP21

FPTAS for Counting Monotone CNF

A monotone CNF formula is a Boolean formula in conjunctive normal form where each variable appears positively. We design a deterministic fully polynomial-time approximation scheme (FPTAS) for counting the number of satisfying assignments for a given monotone CNF formula when each variable appears in at most 5 clauses. Equivalently, this is also an FPTAS for counting set covers where each set contains at most 5 elements. If we allow variables to appear in a maximum of 6 clauses (or sets to contain 6 elements), it is NP-hard to approximate it. Thus, this gives a complete understanding of the approximability of counting for monotone CNF formulas. It is also an important step towards a complete characterization of the approximability for all bounded degree Boolean #CSP problems. In addition, we study the hypergraph matching problem, which arises naturally towards a complete classification, and show an FPTAS for counting 3D matchings of hypergraphs with maximum degree 4.

Jingcheng Liu
UC Berkeley
liuexp@berkeley.edu

Pinyan Lu
Microsoft Research Asia

pinyanl@microsoft.com

yinyt@nju.edu.cn

CP21**Phase Transitions in Random Dyadic Tilings and Rectangular Dissections**

We study rectangular dissections of an $n \times n$ lattice region into rectangles of area n , where $n = 2^k$ for an even integer k . We show there is a natural edge-flipping Markov chain that connects the state space. A similar chain is known to connect the state space when restricted to dyadic tilings, where each rectangle has the form $R = [s2^u, (s+1)2^u] \times [t2^v, (t+1)2^v]$, where s, t, u and v are nonnegative integers. We consider a weighted version where, given a parameter $\lambda > 0$, we generate each rectangular dissection (or dyadic tiling) σ with probability proportional to $\lambda^{|\sigma|}$, where $|\sigma|$ is the total edge length. We show there is a phase transition in the dyadic setting: when $\lambda < 1$, the edge-flipping chain mixes in time $O(n^2 \log n)$, and when $\lambda > 1$, the mixing time is $\exp(\Omega(n^2))$. In the general setting, we show the chain requires exponential time when $\lambda > 1$ and when $\lambda < 1$.

Sarah Cannon, [Sarah Miracle](#), Dana Randall
Georgia Institute of Technology
sarah.cannon@gatech.edu, sarah.miracle@gatech.edu, randall@cc.gatech.edu

CP21**Spatial Mixing and the Connective Constant: Optimal Bounds**

We study the problem of deterministic approximate counting of matchings and independent sets in graphs of bounded *connective constant*. More generally, we consider the problem of evaluating the partition functions of the monomer-dimer model and the hard core model on such graphs. The *connective constant* is a natural notion of average degree of a graph which has been studied extensively in combinatorics and mathematical physics, and can be bounded by a constant even for certain unbounded degree graphs such as those sampled from the sparse Erdős–Rényi model $\mathcal{G}(n, d/n)$. For both models, we prove the best possible rates of decay of correlations in graphs with a given bound on the connective constant. We then use these optimal decay of correlations results to obtain FPTASs for the two problems on graphs of bounded connective constant. Our results improve upon previous results which were often applicable only to bounded degree graphs.

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

Piyush Srivastava
California Institute of Technology
piyushs@caltech.edu

Daniel tefankovic
University of Rochester
stefanko@cs.rochester.edu

Yitong Yin
Nanjing University

CP21**The Speed of Evolution**

In this paper we study the mixing times of a Wright-Fisher model for an asexual population. Despite their importance, and having been widely studied, there has been a lack of rigorous results establishing fast mixing times for many relevant ranges of model parameters *even for the case of two genotypes* (single locus). The main result of this paper is a tight analytical bound on the mixing time of a Wright-Fisher model for two genotypes when there is no restriction on the mutation rate or the fitness.

[Nisheeth K. Vishnoi](#)
EPFL
EPFL
nisheeth.vishnoi@epfl.ch

CP22**Optimal Detection of Intersections Between Convex Polyhedra**

For a polyhedron P in \mathbb{R}^d , denote by $|P|$ its combinatorial complexity, i.e., the number of faces of all dimensions of the polyhedra. In this paper, we revisit the classic problem of preprocessing polyhedra independently so that given two preprocessed polyhedra P and Q in \mathbb{R}^d , each translated and rotated, their intersection can be tested rapidly. For $d = 3$ we show how to perform such a test in $O(\log |P| + \log |Q|)$ time after linear preprocessing time and space. This running time is the best possible and improves upon the last best known query time of $O(\log |P| \log |Q|)$ by Dobkin and Kirkpatrick (1990). We then generalize our method to any constant dimension d , achieving the same optimal $O(\log |P| + \log |Q|)$ query time using a representation of size $O(|P|^{d/2+\epsilon})$ for any $\epsilon > 0$ arbitrarily small. This answers an even older question posed by Dobkin and Kirkpatrick 30 years ago.

[Luis Barba](#)
Carleton University
lfbalba@gmail.com

Stefan Langerman
Université Libre de Bruxelles
slanger@ulb.ac.be

CP22**Compatible Connectivity-Augmentation of Planar Disconnected Graphs**

Motivated by applications to graph morphing, we consider the following *compatible connectivity-augmentation problem*: We are given a labelled n -vertex planar graph, \mathcal{G} , that has $r \geq 2$ connected components, and $k \geq 2$ isomorphic planar straight-line drawings, G_1, \dots, G_k , of \mathcal{G} . We wish to augment \mathcal{G} by adding vertices and edges to make it connected in such a way that these vertices and edges can be added to G_1, \dots, G_k as points and straight-line segments, respectively, to obtain k planar straight-line drawings isomorphic to the augmentation of \mathcal{G} . We show that adding $\Theta(nr^{1-1/k})$ edges and vertices to \mathcal{G} is always sufficient and sometimes necessary to achieve this goal.

Greg Aloupis
Université Libre de Bruxelles

aloupis.greg@gmail.com

Luis Barba

Carleton University
lbarba@gmail.com

Paz Carmi

Ben-Gurion University of the Negev
carmip@cs.bgu.ac.il

Vida Dujmovic

University of Ottawa
vida.dujmovic@uottawa.ca

Fabrizio Frati

The University of Sydney
fabrizio.frati@sydney.edu.au

Pat Morin

Carleton University
morin@scs.carleton.ca

CP22

Geometric k Shortest Paths

We consider computing k homotopically distinct shortest paths in a polygonal domain. We give algorithms for finding the paths, and present the k th shortest path map – a structure allowing one to report efficiently the length of the k shortest paths to a query point.

Sylvester Eriksson-Bique

Courant Institute, NYU
ebs@cims.nyu.edu

John Hershberger

Mentor Graphics Corporation
john_hershberger@mentor.com

Valentin Polishchuk

Communications and Transport Systems
ITN, Linköping University
valentin.polishchuk@liu.se

Bettina Speckmann

Dept. of Mathematics and Computer Science
TU Eindhoven
b.speckmann@tue.nl

Subhash Suri

University of California, Santa Barbara
suri@cs.ucsb.edu

Topi L. Talvitie

University of Helsinki
topi.talvitie@helsinki.fi

Kevin Verbeek, Hakan Yildiz

Computer Science
University of California Santa Barbara
kverbeek@cs.ucsb.edu, hakan@cs.ucsb.edu

CP22

Triangulation Refinement and Approximate Shortest Paths in Weighted Regions

We consider the problem of computing a shortest path in a

weighted planar subdivision, where the cost of a path within a face is the product of its length and the face weight. We present a fully polynomial-time approximation scheme for this problem. Our algorithm relies on a new triangulation refinement method that produces a triangulation such that no triangle has two small angles.

Siu-Wing Cheng

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

Jiongxin Jin

Google Inc.
jamesjx@google.com

Antoine Vigneron

King Abdullah University of Science and Technology
(KAUST)
antoine.vigneron@kaust.edu.sa

CP22

Detecting Weakly Simple Polygons

A closed curve in the plane is weakly simple if it is the limit (in the Frchet metric) of a sequence of simple closed curves. We describe an algorithm to determine whether a closed walk of length n in a simple plane graph is weakly simple in $O(n \log n)$ time, improving an earlier $O(n^3)$ -time algorithm of Cortese et al. [*Discrete Math.* 2009]. As an immediate corollary, we obtain the first efficient algorithm to determine whether an arbitrary n -vertex polygon is weakly simple; our algorithm runs in $O(n^2 \log n)$ time. We also describe algorithms that detect weak simplicity in $O(n \log n)$ time for two interesting classes of polygons. Finally, we discuss subtle errors in several previously published definitions of weak simplicity.

Hsien-Chih Chang, Jeff Erickson, Chao Xu

University of Illinois, Urbana-Champaign
hchang17@illinois.edu, jeffe@cs.uiuc.edu,
chaoxu3@illinois.edu

CP23

Subcubic Equivalences Between Graph Centrality Problems, ApSP and Diameter

In the analysis of social networks there are many important centrality parameters such as radius, diameter, betweenness centrality, and reach centrality. The known algorithms for computing these measures run in time that is roughly cubic in the number of nodes. We present a set of reductions showing that these algorithms cannot be improved without improving the best known algorithms for the classical problems All Pairs Shortest Paths (APSP) and graph diameter. Moreover, we prove that most graph centrality measures are equivalent to either APSP or diameter.

Amir Abboud

Stanford University
abboud@cs.stanford.edu

Fabrizio Grandoni

IDSIA, University of Lugano
fabrizio@idsia.ch

Virginia Vassilevska Williams
Stanford University

virgi@cs.stanford.edu

CP23

Forbidden Structure Characterization of Circular-Arc Graphs and a Certifying Recognition Algorithm

A circular-arc graph is the intersection graph of arcs of a circle. It is a well-studied graph model with numerous natural applications. A certifying algorithm is an algorithm that outputs a certificate, along with its answer (be it positive or negative), where the certificate can be used to easily justify the given answer. We present the first forbidden structure characterization of circular-arc graphs and the first polynomial-time certifying algorithm for the recognition of circular-arc graphs.

Mathew Francis
Indian Statistical Institute, Chennai Centre
Computer Science Unit
mathewcfrancis@gmail.com

Pavol Hell
School of Computer Science
Simon Fraser University
pavol@sfu.ca

Juraj Stacho
Columbia University in the City of New York
Department of Industrial Engineering & Operations
Research
stacho@cs.toronto.edu

CP23

Four Terminal Planar Delta-Wye Reducibility Via Rooted $K_{2,4}$ Minors

A graph with four special vertices (terminals) is wye-delta reducible if we can obtain a graph on 4 vertices by a sequence of wye-delta operations and series-parallel reductions, none of which removes a terminal. A good characterization of wye-delta reducible 3-connected planar graphs with 4 terminals is given. The proofs yield an $O(n^2)$ algorithm that either finds an obstruction to reducibility or returns a sequence of reductions. The main ingredient in the proofs is a good characterization of planar graphs with 4 terminals that do not admit a rooted $K_{2,4}$ minor with the terminals corresponding to the roots on the large side of the bipartition of $K_{2,4}$. Up to small connectivity reductions, cases without the rooted minor fall into five structural cases that lead to a polynomial-time algorithm for recognition of these graphs and construction of rooted $K_{2,4}$ minors. This result is of independent interest in structural graph theory.

Lino Demasi, Bojan Mohar
Simon Fraser University
lmd14@sfu.ca, mohar@sfu.ca

CP23

Minors and Dimension

Streib and Trotter proved in 2012 that posets with bounded height and with planar cover graphs have bounded dimension. Recently, Joret et al. proved that the dimension is bounded for posets with bounded height whose cover graphs have bounded tree-width. In this paper, it is proved that posets of bounded height whose cover graphs exclude a fixed (topological) minor have bounded dimension. This

generalizes both the aforementioned results and verifies a conjecture of Joret et al. The proof relies on the Robertson-Seymour and Grohe-Marx structural decomposition theorems.

Bartosz Walczak
Jagiellonian University
walczak@tcs.uj.edu.pl

CP23

Finding Four-Node Subgraphs in Triangle Time

We present new algorithms for finding induced four-node subgraphs in a given graph, which run in time roughly that of detecting a clique on three nodes.

- The best known algorithms for triangle finding in an n -node graph take $O(n^\omega)$ time, where $\omega < 2.373$ is the matrix multiplication exponent. We give a general randomized technique for finding any induced four-node subgraph, except for the clique or independent set on 4 nodes, in $\tilde{O}(n^\omega)$ time with high probability. The algorithm can be derandomized in some cases.
- For sparse graphs with m edges, the best known triangle finding algorithm runs in $O(m^{2\omega/(\omega+1)}) \leq O(m^{1.41})$ time. We give a randomized $\tilde{O}(m^{2\omega/(\omega+1)})$ time algorithm for finding any induced four-node subgraph other than C_4 , K_4 and their complements. Some cases can be derandomized. For C_4 or its complement, we get slightly slower algorithms.

Virginia Vassilevska Williams, Joshua R. Wang, Ryan Williams, Huacheng Yu
Stanford University
virgito@gmail.com, jrwang@stanford.edu,
rrwilliams@gmail.com, yuhch123@gmail.com

CP24

Lopsidedependency in the Moser-Tardos Framework: Beyond the Lopsided Lovász Local Lemma

The Lopsided Lovász Local Lemma (LLLL) is a powerful principle in probability, along with the efficient resampling algorithm Moser & Tardos. We show a new criterion for the Moser-Tardos algorithm to converge. This criterion is stronger than the LLLL criterion, and in fact can yield better results even than the full Shearer criterion. A noteworthy application is for k -SAT. As shown in Gebauer, Szábo, and Tardos, a k -SAT instance in which every variable appears $L \leq \frac{2^{k+1}}{e^{(k+1)}}$ times, is satisfiable. We improve it to $L \leq \frac{2^{k+1}(1-1/k)^k}{k-1} - \frac{2}{k}$ which can be significantly stronger when k is small. We introduce a new parallel algorithm for the LLLL, which applies in nearly all settings as the sequential algorithm.

David Harris
University of Maryland
davidgharris29@hotmail.com

CP24

Characterization of Cutoff for Reversible Markov Chains

A sequence of Markov chains is said to exhibit (total variation) cutoff if the convergence to stationarity in total variation distance is abrupt. We consider reversible lazy chains.

We prove a necessary and sufficient condition for the occurrence of the cutoff phenomena in terms of concentration of hitting time of “worst” (in some sense) sets of stationary measure at least α , for some $\alpha \in (0, 1)$. As an application of our techniques we show that a sequence of lazy Markov chains on finite trees exhibits a cutoff iff the product of their spectral gaps and their (lazy) mixing-times tends to ∞ .

Riddhipratim Basu
Department of Statistics, UC Berkeley, California, USA
rpbasu.riddhi@gmail.com

Jonathan Hermon
367 Evans Hall, University of California
Berkeley, CA 94720-3860
jonathan.hermon@gmail.com

Yuval Peres
Microsoft Research, Redmond
peres@microsoft.com

CP24

An Algorithmic Framework for Obtaining Lower Bounds for Random Ramsey Problems

In this paper we introduce a general framework for proving lower bounds for various Ramsey type problems within random settings. The main idea is to provide an algorithm that finds the desired colouring w.h.p. Our framework allows to reduce the probabilistic problem of whether the Ramsey property holds for random (hyper)graphs to a deterministic question of whether there exists a finite graph that forms an obstruction. We apply this framework to solve various open problems. We provide a matching lower bound for the result of Friedgut, Rödl and Schacht (2010) and, independently, Conlon and Gowers (2014+) for the classical Ramsey problem for hypergraphs in the case of cliques. We also improve a result of Bohman, Frieze, Pikhurko and Smyth (2010) for bounded anti-Ramsey problems in random graphs and extend it to hypergraphs. Finally, we provide matching lower bounds for a proper-colouring version of anti-Ramsey problems introduced by Kohayakawa, Konstantinidis and Mota (2014).

Rajko Nenadov
ETH Zürich
rnenadov@inf.ethz.ch

Nemanja Skoric
ETH Zurich
nskoric@inf.ethz.ch

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

CP24

Robust Hamiltonicity of Random Directed Graphs

In 1952 Dirac showed that K_n remains Hamiltonian even if we allow an adversary to remove $\lfloor n/2 \rfloor$ edges touching each vertex. In 1960 Ghouila-Houri obtained an analogue statement for digraphs. A natural way to generalize such results to arbitrary graphs (digraphs) is using the notion of local resilience. The local resilience of a graph (digraph) G with respect to a property P is the maximum number r

such that G has the property P even if we allow an adversary to remove an r -fraction of (in- and out-going) edges touching each vertex. Lee and Sudakov (2012) proved that the local resilience of a random graph with edge probability $p = \omega(\log n/n)$ with respect to Hamiltonicity is $1/2 \pm o(1)$. For random directed graphs, Hefetz, Steger and Sudakov (2014+) proved an analogue statement for $p = \omega(\log n/\sqrt{n})$. In this paper we improve their result to $p = \omega(\log^8 n/n)$, which is optimal up to the polylogarithmic factor.

Asaf Ferber
ETH Zurich
asaf.ferber@inf.ethz.ch

Rajko Nenadov
ETH Zürich
rnenadov@inf.ethz.ch

Andreas Noever, Ueli Peter, Nemanja Skoric
ETH Zurich
anoever@inf.ethz.ch, upeter@inf.ethz.ch, nskoric@inf.ethz.ch

CP24

Surprise Probabilities in Markov Chains

In a Markov chain started at a state x , the hitting time $\tau(y)$ is the first time that the chain reaches another state y . We study the probability $\mathbf{P}_x(\tau(y) = t)$ that the first visit to y occurs precisely at a given time t . Informally speaking, the event that a new state is visited at a large time t may be considered a “surprise”. We prove:

- In any Markov chain with n states, $\mathbf{P}_x(\tau(y) = t) \leq \frac{n}{t}$.
- In a reversible chain with n states, $\mathbf{P}_x(\tau(y) = t) \leq \frac{\sqrt{2n}}{t}$ for $t \geq 4n + 4$.
- For random walk on a simple graph with $n \geq 2$ vertices, $\mathbf{P}_x(\tau(y) = t) \leq \frac{4e \log n}{t}$.

We construct examples showing that these bounds are close to optimal. The main feature of our bounds is that they require very little knowledge of the structure of the Markov chain.

James Norris
University of Cambridge
james@statslab.cam.ac.uk

Yuval Peres
Microsoft Research, Redmond
peres@microsoft.com

Alex Zhai
Stanford University
alexlinzhai@gmail.com

CP25

Testing Poisson Binomial Distributions

A Poisson Binomial distribution over n variables is the distribution of the sum of n independent Bernoullis. We provide a sample near-optimal algorithm for testing whether a distribution P supported on $\{0, \dots, n\}$ to which we have sample access is a Poisson Binomial distribution, or far from all Poisson Binomial distributions. The sample complexity of our algorithm is $O(n^{1/4})$ to which we provide a

matching lower bound. We note that our sample complexity improves quadratically upon that of the naive “learn followed by tolerant-test” approach, while instance optimal identity testing [VV14] is not applicable since we are looking to simultaneously test against a whole family of distributions.

Jayadev Acharya, Constantinos Daskalakis
MIT
jayadev@csail.mit.edu, costis@csail.mit.edu

CP25

Testing Identity of Structured Distributions

We study the question of identity testing for structured distributions. More precisely, given samples from a *structured* distribution q over $[n]$ and an explicit distribution p over $[n]$, we wish to distinguish whether $q = p$ versus q is at least ϵ -far from p , in L_1 distance. In this work, we present a unified approach that yields new, simple testers, with sample complexity that is information-theoretically optimal, for broad classes of structured distributions.

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

Daniel Kane
UC San Diego
dakane@cs.ucsd.edu

Vladimir Nikishkin
University of Edinburgh
v.nikishkin@sms.ed.ac.uk

CP25

Property Testing on Product Distributions: Optimal Testers for Bounded Derivative Properties

We study algorithms that, given access to a small number of samples from a function over large discrete domain, approximately determine whether the dataset satisfies a desired property. The distance to having the property is measured with respect to a known or an unknown distribution on the datasets. We focus on functions over domains of the form $[n]^d = 1, 2, \dots, n^d$ (d-dimensional hypergrids). We give an optimal tester for a general class of properties of such functions, called bounded derivative properties (BDP), for the case when the distance to the property is measured with respect to a product distribution. The BDPs include monotonicity and the Lipschitz property. We use our novel and optimum (upto constant factors) dimension reduction, which reduces testing properties of functions over $[n]^d$ to testing functions over $[n]$. Our work resolves two open problems given by Ailon and Chazelle (Information and Computation, 2006) and improves their line-tester for monotonicity.

Deeparnab Chakrabarty
Microsoft Research India
dechakr@microsoft.com

Kashyap Dixit
Penn State university
kashyap@cse.psu.edu

Madhav Jha, C. Seshadhri
Sandia National Labs

mjha@sandia.gov, scomand@sandia.gov

CP25

Algorithmic Regularity for Polynomials and Applications

We discuss several notions of regularity for collections of polynomials, which lead to algorithmic versions of various polynomial regularity lemmas such as the ones by Green and Tao [Contrib. Discrete Math. 2009] and by Kaufman and Lovett [FOCS 2008]. We show that our algorithmic regularity lemmas for polynomials also lead to algorithmic versions of many results relying on regularity, such as decoding Reed-Muller codes beyond the list decoding radius (for certain structured errors) and worst-case to average-case reductions for polynomials. We also provide algorithmic versions of some of the inverse theorems for Gowers norm.

Arnab Bhattacharyya
Indian Institute of Science
arnabb@csa.iisc.ernet.in

Pooya Hatami
University of Chicago
pooya@cs.uchicago.edu

Madhur Tulsiani
TTI Chicago
madhurt@ttic.edu

CP25

The Complexity of Estimating Rényi Entropy

It was recently shown that estimating the Shannon entropy $H(p)$ of a discrete k -symbol distribution p requires $\Theta(k/\log k)$ samples, a number that grows near-linearly in the support size. In many applications $H(p)$ can be replaced by the more general Rényi entropy of order α , $H_\alpha(p)$. We determine the number of samples needed to estimate $H_\alpha(p)$ for all α , showing that $\alpha < 1$ requires super-linear, roughly $k^{1/\alpha}$ samples, noninteger $\alpha > 1$ requires near-linear, roughly k samples, but integer $\alpha > 1$ requires only $\Theta(k^{1-1/\alpha})$ samples. In particular, estimating $H_2(p)$, which arises in security, DNA reconstruction, closeness testing, and other applications, requires only $\Theta(\sqrt{k})$ samples. The estimators achieving these bounds are simple and run in time linear in the number of samples.

Jayadev Acharya
MIT
jayadev@csail.mit.edu

Alon Orlitsky, Ananda Theertha Suresh
University of California, San Diego
alon@ucsd.edu, asuresh@ucsd.edu

Himanshu Tyagi
UC San Diego
htyagi@eng.ucsd.edu

CP26

A Dynamic Model of Barter Exchange

We consider the problem of efficient operation of a barter exchange platform for indivisible goods. We introduce a dynamic model where in each period one agent arrives

with a single item she wants to exchange for a different item. An agent is interested in the item possessed by another agent with probability p , independently for all pairs of agents. We consider three settings with respect to the types of allowed exchanges: a) Only two-way cycles, b) Two or three-way cycles, c) (unbounded) chains initiated by altruistic donors. The goal of the platform is to minimize the average waiting time of an agent. We find that in each of these settings, a policy that conducts exchanges in a greedy fashion is near optimal, among a large class of policies that includes batching policies. Further, for small p , allowing three-cycles can greatly improve the waiting time over the two-cycles only setting, and the presence of altruistic donors can lead to a further large improvement.

Ross Anderson
MIT ORC
rma350@gmail.com

Itai Ashlagi
MIT Sloan School of Management
iashlagi@mit.edu

David Gamarnik
MIT
gamarnik@MIT.EDU

Yash Kanoria
Columbia Business School
ykanoria@gmail.com

CP26

Approximately Stable, School Optimal, and Student-Truthful Many-to-One Matchings (via Differential Privacy)

We present a mechanism for computing asymptotically stable school optimal matchings, while guaranteeing that it is an asymptotic dominant strategy for every student to report their true preferences to the mechanism. Our main tool in this endeavor is differential privacy: we give an algorithm that coordinates a stable matching using differentially private signals, which lead to our truthfulness guarantee. This is the first setting in which it is known how to achieve nontrivial truthfulness guarantees for students when computing school optimal matchings, assuming worst-case preferences (for schools and students) in large markets.

Sampath Kannan
University of Pennsylvania
sampath@cis.penn.edu

Jamie Morgenstern
Carnegie Mellon university
jamiemmt@cs.cmu.edu

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

Zhiwei Steven Wu
University of Pennsylvania

wuzhiwei@cis.penn.edu

CP26

The Structure of the Core in Assignment Markets

Assignment markets (Shapley & Shubik 1971) involve matching with transfers, as in labor and housing markets. We consider a two-sided assignment market with agent types and stochastic structure similar to models used in empirical studies. Each agent has a randomly drawn "productivity" associated with each type on the other side. We characterize how the structure of the core, i.e., the set of stable outcomes, is determined by market characteristics.

Yash Kanoria
Columbia Business School
ykanoria@gmail.com

Daniela Saban, Jay Sethuraman
Columbia University
dhs2131@columbia.edu, jay@ieor.columbia.edu

CP26

Robust Randomized Matchings

The following zero-sum game is played on a weighted graph G : Alice selects a matching M in G and Bob selects a number k . Alice receives a payoff equal to the ratio of the weight of the top k edges of M to the maximum weight of a matching of size at most k in G . If M guarantees a payoff of at least α then it is called α -robust. In 2002, Hassin and Rubinstein gave an algorithm that returns a $\sqrt{1/2}$ -robust matching, which is best possible for this setting. We show that Alice can improve on the guarantee of $\sqrt{1/2}$ when allowing her to play a randomized strategy. We devise a simple algorithm that returns a $1/\ln(4)$ -robust randomized matching. We further show that our robustness results for randomized matchings translate to an asymptotic robustness guarantee for deterministic matchings and we give a new simple LP-based proof of Hassin and Rubinstein's original result.

Jannik Matuschke, Martin Skutella
TU Berlin
matuschke@math.tu-berlin.de,
martin.skutella@tu-berlin.de

José Soto
Universidad de Chile
jsoto@dim.uchile.cl

CP26

Bayesian Truthful Mechanisms for Job Scheduling from Bi-Criterion Approximation Algorithms

We provide polynomial-time approximately optimal Bayesian mechanisms for makespan minimization on unrelated machines as well as for max-min fair allocations of indivisible goods, with approximation factors of 2 and $\min\{m - k + 1, \tilde{O}(\sqrt{k})\}$ respectively, matching the approximation ratios of best known polynomial-time *algorithms*. Our mechanisms are obtained by establishing a polynomial-time approximation-sensitive reduction from the problem of designing approximately optimal *mechanisms* for some arbitrary objective \mathcal{O} to that of designing bi-criterion approximation *algorithms* for the same objective \mathcal{O} plus a linear allocation cost term. Our reduction is itself enabled by extending the celebrated "equivalence of

separation and optimization'[GLS81,KP80] to also accommodate bi-criterion approximations.

Constantinos Daskalakis, Matt Weinberg
MIT
costis@csail.mit.edu, smweinberg@csail.mit.edu

CP27

Tight Bounds on Vertex Connectivity Under Vertex Sampling

A fundamental result by Karger (SODA 1994) states that for any λ -edge-connected graph with n nodes, independently sampling each edge with probability $p = \Omega(\log n/\lambda)$ results in a graph that has edge connectivity $\Omega(\lambda p)$, with high probability. This paper proves the analogous result for vertex connectivity, when sampling vertices. We show that for any k -vertex-connected graph G with n nodes, if each node is independently sampled with probability $p = \Omega(\sqrt{\log n/k})$, then the subgraph induced by the sampled nodes has vertex connectivity $\Omega(kp^2)$, with high probability. This bound improves upon the recent results of Censor-Hillel et al. (SODA 2014), and is existentially optimal.

Keren Censor-Hillel
Technion, Israel
ckeren@cs.technion.ac.il

Mohsen Ghaffari
MIT
ghaffari@mit.edu

George Giakkoupis
INRIA Rennes
george.giakkoupis@inria.fr

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

Fabian Kuhn
University of Freiburg, Germany
kuhn@cs.uni-freiburg.de

CP27

2-Edge Connectivity in Directed Graphs

Edge and vertex connectivity are fundamental concepts in graph theory. While they have been thoroughly studied in the case of undirected graphs, surprisingly not much has been investigated for directed graphs. In this paper we study 2-edge connectivity problems in directed graphs and, in particular, we consider the computation of the following natural relation: We say that two vertices v and w are 2-edge-connected if there are two edge-disjoint paths from v to w and two edge-disjoint paths from w to v . This relation partitions the vertices into blocks such that all vertices in the same block are 2-edge-connected. Differently from the undirected case, those blocks do not correspond to the 2-edge-connected components of the graph. Our main result is an algorithm for computing the 2-edge-connected blocks of a directed graph in linear time. Additionally, we also show how to compute in linear time a sparse certificate for this relation.

Loukas Georgiadis
University of Ioannina, Greece

loukas@gmail.com

Giuseppe Italiano
Universita di Roma "Tor Vergata"
giuseppe.italiano@uniroma2.it

Luigi Laura
"Sapienza" Universita di Roma
laura@dis.uniroma1.it

Nikos Parotsidis
University of Ioannina, Greece
nparotsi@cs.uoi.gr

CP27

Contagious Sets in Expanders

We consider the following activation process in undirected graphs, often referred to as *bootstrap percolation*: a vertex is active either if it belongs to a set of initially activated vertices or if at some point it has at least r active neighbors, where $r > 1$ is the activation threshold. A *contagious set* is a set whose activation results with the entire graph being active. Given a graph G , let $m(G, r)$ be the minimal size of a contagious set. We consider graphs with expansion properties, parameterized by the spectral gap and/or the girth. We prove that sufficiently strong expansion properties imply that $m(G, 2) \leq O(\frac{n}{d^2})$ (and more generally, $m(G, r) \leq O(\frac{n}{d^{r/(r-1)}})$). In addition, we demonstrate that weaker assumptions suffice in order to imply that $m(G, 2) \leq O(\frac{n \log d}{d^2})$. For example, we show this for graphs of girth at least 7, and for graphs with $\lambda(G) < (1 - \epsilon)d$, provided the graph has no 4-cycles.

Amin Coja-Oghlan
Goethe University, Germany
acoghlan@math.uni-frankfurt.de

Uriel Feige
Weizmann Institute, Israel
uriel.feige@weizmann.ac.il

Michael Krivelevich
Tel Aviv University
krivelev@post.tau.ac.il

Daniel Reichman
Cornell University
daniel.reichman@gmail.com

CP27

Fast Generation of Random Spanning Trees and the Effective Resistance Metric

We present a new algorithm for generating a uniformly random spanning tree in an undirected graph. Our algorithm samples such a tree in expected $O(m^{\frac{4}{3}+o(1)})$ time. This improves over the best previously known bound of $\min(O(m\sqrt{n}), O(n^\omega))$ – that follows from the work of Kellner and Madry [FOCS'09] and of Colbourn et al. [J. Algorithms'96] – whenever the input graph is sufficiently sparse. At a high level, our result stems from carefully exploiting the interplay of random spanning trees, random walks, and the notion of effective resistance, as well as from devising a way to algorithmically relate these concepts to the combinatorial structure of the graph. This involves, in particular, establishing a new connection between the effective resis-

tance metric and the cut structure of the underlying graph.

Aleksander Madry
Microsoft Research New England
aleksander.madry@epfl.ch

Damian Straszak, Jakub Tarnawski
University of Wrocław
damian.straszak@gmail.com, jakub.tarnawski@gmail.com

CP27

Connectivity in Random Forests and Credit Networks

In this paper, we study uniformly random forests of a graph. We define the RF-connectivity between a pair of nodes in a graph G as the probability that the two nodes belong to the same connected component in a uniformly random forest of G . Our main result is that for an arbitrary subset S of nodes in G , the average RF-connectivity between pairs of nodes in S is at least $1 - 2/h(G_S)$, where $h(G_S)$ is the edge expansion of the subgraph G_S induced by S . This work is motivated by the study of credit networks (Dandekar et al. '11) and the negative correlation conjecture in the study of uniformly random forests.

Ashish Goel
Stanford University
ashishg@stanford.edu

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

Sharath Raghvendra
Virginia Tech
r.sharath@gmail.com

Hongyang Zhang
Stanford University
hongyz@stanford.edu