

IP1**Shape, Homology, Persistence, and Stability**

My personal journey to the fascinating world of geometric forms started 30 years ago with the invention of alpha shapes in the plane. It took about 10 years before we generalized the concept to higher dimensions, we produced working software with a graphics interface for the 3-dimensional case, and we added homology to the computations. Needless to say that this foreshadowed the inception of persistent homology, because it suggested the study of filtrations to capture the scale of a shape or data set. Importantly, this method has fast algorithms. The arguably most useful result on persistent homology is the stability of its diagrams under perturbations.

Herbert Edelsbrunner

Institute of Science and Technology Austria
edels@ist.ac.at

IP2**Recent Developments in the Sparse Fourier Transform**

The Fast Fourier Transform (FFT) is a widely used numerical algorithm. It computes the Discrete Fourier Transform (DFT) of an n -dimensional signal in $O(n \log n)$ time. It is not known whether its running time can be improved. However, in many applications, most of the Fourier coefficients of a signal are "small" or equal to zero, i.e., the output of the transform is (approximately) sparse. In this case, it is known that one can compute the set of non-zero coefficients much faster, even in sub-linear time. In this talk I will give an overview of recent highly efficient algorithms for computing the Sparse Fourier Transform. I will also give a few examples of applications impacted by these developments.

Piotr Indyk

Massachusetts Institute of Technology
indyk@mit.edu

IP3**Interlacing Families: Mixed Characteristic Polynomials and the Kadison-Singer Problem**

We introduce a new technique for demonstrating the existence of combinatorial objects that we call the "Method of Interlacing Polynomials". We then use this technique to prove a conjecture in discrepancy theory posed by Nik Weaver. Weaver's conjecture implies the truth of the Paving Conjecture of Akemann and Anderson, which in turn provides a solution to the fifty year old problem of Kadison and Singer. Our main technical result is an upper bound on the largest root of the expected characteristic polynomial of a sum of random symmetric rank-1 matrices. We use this bound to prove that there exists a balanced partition of certain sets of vectors. This result can be viewed as a strengthening of the "matrix Chernoff bounds" of Ahlswede and Winter, Rudelson and Vershynin, and Tropp. This is joint work with Adam Marcus and Nikhil Srivastava.

Daniel Spielman
Yale University

spielman@cs.yale.edu

CP0**Distributed Computation of Persistent Homology**

Advances in algorithms for computing persistent homology have reduced the computation time drastically – as long as the algorithm does not exhaust the available memory. Following up on a recently presented parallel method for persistence computation on shared memory systems, we demonstrate that a simple adaptation of the standard reduction algorithm leads to a variant for distributed systems. Our algorithmic design ensures that the data is distributed over the nodes without redundancy; this permits the computation of much larger instances than on a single machine. The parallelism often speeds up the computation compared to sequential and even parallel shared memory algorithms. In our experiments, we were able to compute the persistent homology of filtrations with more than a billion (10^9) elements within seconds on a cluster with 32 nodes using less than 10GB of memory per node.

Ulrich Bauer

Institute of Science and Technology (Austria)
mail@ulrich-bauer.org

Michael Kerber

Stanford University
mkerber@mpi-inf.mpg.de

Jan Reininghaus

IST Austria
jan.reininghaus@ist.ac.at

CP0**On the Asymptotic Number of $BCK(2)$ -Terms**

We investigate the asymptotic number of a particular class of closed lambda-terms. This class is a generalization of a class of terms related to the axiom system BCK which is well known in combinatory logic. We determine the asymptotic number of terms, when their size tends to infinity, up to a constant multiplicative factor and discover a surprising asymptotic behaviour involving an exponential with fractional powers in the exponent.

Olivier Bodini

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

Bernhard Gittenberger

Institute for Discrete Mathematics and Geometry, TU
Wien
gittenberger@dmg.tuwien.ac.at

CP0**On the Scalability of Computing Triplet and Quartet Distances**

We present an experimental evaluation of the algorithms by Brodal *et al.* [SODA 2013] for computing the triplet and quartet distance measures between two leaf labelled rooted and unrooted trees of arbitrary degree, respectively. In our experimental evaluation of the algorithms the typical overhead is about 2 KB and 10 KB per node in the input trees for triplet and quartet computations, respectively. This allows us to compute the distance measures for trees with up

to millions of nodes.

Gerth S. Brodal

MADALGO, Aarhus University
gerth@cs.au.dk

Jens Johansen, Morten Holt
Aarhus University
jensjoha@gmail.com, thehawk@gmail.com

CP0

Enumerating Fundamental Normal Surfaces: Algorithms, Experiments and Invariants

Computational knot theory and 3-manifold topology have seen significant breakthroughs in recent years, despite the fact that many key algorithms have complexity bounds that are exponential or greater. In this setting, experimentation is essential for understanding the limits of practicality, as well as for gauging the relative merits of competing algorithms. Here we focus on normal surface theory, a key tool in low-dimensional computational topology. In particular, we aim to compute fundamental normal surfaces, which form an integral basis for a high-dimensional polyhedral cone. We develop, implement and experimentally compare a primal and a dual algorithm, both of which combine domain-specific techniques with classical Hilbert basis algorithms. Our experiments indicate that we can solve extremely large problems that were once though intractable. As a practical application, we fill gaps from the KnotInfo database by computing 398 previously-unknown crosscap numbers of knots.

Benjamin A. Burton

The University of Queensland
bab@maths.uq.edu.au

CP0

An Exact Approach To Upward Crossing Minimization

The upward crossing number problem asks for a drawing of the graph into the plane with the minimum number of edge crossings where the edges are drawn as monotonously increasing curves w.r.t. the y-axis. While there is a large body of work on solving this central graph drawing problem heuristically, we present the first approach to solve the problem to proven optimality. Our approach is based on a reformulation of the problem as a boolean formula that can be iteratively tightened and resolved. In our experiments, we show the practical applicability and limits of our approach.

Robert Zeranski
Friedrich-Schiller-Universität Jena
robert.zeranski@uni-jena.de

Markus Chimani

Osnabrueck University
markus.chimani@uni-osnabrueck.de

CP0

A Bijection for Plane Graphs and Its Applications

This paper is concerned with the counting and random sampling of plane graphs (simple planar graphs embedded in the plane). Our main result is a bijection between the class of plane graphs with triangular outer face, and a class

of oriented binary trees. The number of edges and vertices of the plane graph can be tracked through the bijection.

Gwendal Collet

LIX, Ecole Polytechnique
gcollet@lix.polytechnique.fr

Olivier Bernardi
Brandeis University, MA
bernardi@brandeis.edu

Eric Fusy

LIX, Ecole Polytechnique, France
fusy@lix.polytechnique.fr

CP0

A Statistical View on Exchanges in Quickselect

In this paper we study the number of key exchanges required by Hoare's FIND algorithm (also called Quickselect) when operating on a uniformly distributed random permutation and selecting an independent uniformly distributed rank. After normalization we give a limit theorem where the limit law is a perpetuity characterized by a recursive distributional equation. To make the limit theorem usable for statistical methods and statistical experiments we provide an explicit rate of convergence in the Kolmogorov–Smirnov metric, a numerical table of the limit law's distribution function and an algorithm for exact simulation from the limit distribution. We also investigate the limit law's density. This case study provides a program applicable to other cost measures, alternative models for the rank selected and more balanced choices of the pivot element such as median-of- $2t + 1$ versions of Quickselect as well as further variations of the algorithm.

Benjamin Dadoun

École Normale Supérieure de Cachan
Département Informatique
benjamin.dadoun@ens-cachan.fr

Ralph Neininger

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

CP0

Small Superpatterns for Dominance Drawing

We exploit the connection between dominance drawings of directed acyclic graphs and permutations, in both directions, to provide improved bounds on the size of universal point sets for certain types of dominance drawing and on superpatterns for certain natural classes of permutations. In particular we show that there exist universal point sets for dominance drawings of the Hasse diagrams of width-two partial orders of size $O(n^{3/2})$, universal point sets for dominance drawings of st -outerplanar graphs of size $O(n \log n)$, and universal point sets for dominance drawings of directed trees of size $O(n^2)$. We show that 321-avoiding permutations have superpatterns of size $O(n^{3/2})$, riffle permutations (321-, 2143-, and 2413-avoiding permutations) have superpatterns of size $O(n)$, and the concatenations of sequences of riffles and their inverses have superpatterns of size $O(n \log n)$. Our analysis includes a calculation of the leading constants in these bounds.

Michael J. Bannister, William Devanny, David Eppstein
University of California, Irvine

Department of Computer Science
mbannist@uci.edu, wdevanny@uci.edu, eppstein@uci.edu

CP0

Permuted Random Walk Exits Typically in Linear Time.

Given a permutation σ of the integers $\{-n, -n+1, \dots, n\}$ we consider the Markov chain X_σ , which jumps from k to $\sigma(k \pm 1)$ equally likely if $k \neq -n, n$. We prove that the expected hitting time of $\{-n, n\}$ starting from any point is $\Theta(n)$ with high probability when σ is a uniformly chosen permutation. We prove this by showing that with high probability, the digraph of allowed transitions is an Eulerian expander; we then utilize general estimates of hitting times in directed Eulerian expanders.

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

Yuval Peres
Microsoft Research, Redmond
peres@microsoft.com

CP0

Practical Experience with Hanani-Tutte for Testing C-Planarity

We propose an algorithm for c-planarity testing which is correct and efficient, but not, in general, complete, i.e., there are input instances on which the algorithm declines to give an answer. At the core of this algorithm is an algebraic criterion with the following properties: (1) The criterion is a necessary condition for c-planarity; (2) for special graph classes, including c-connected graphs, it is also sufficient; and (3) it can be tested efficiently in polynomial time. The algebraic criterion is not sufficient in general; however, we can extend it to a (still efficient) algorithm that verifies the answer of the criterion by building a c-planar embedding of the input graph. Our practical experiments show that this algorithm works well in practice.

Petra Mutzel
Dortmund University
Germany
petra.mutzel@cs.uni-dortmund.de

Carsten Gutwenger
Technische Universität Dortmund
carsten.gutwenger@tu-dortmund.de

Marcus Schaefer
DePaul University
mschaefer@cdm.depaul.edu

CP0

Clump Combinatorics, Automata, and Word Asymptotics

Given a set of words and a probability model for random texts, we are interested in the behavior of occurrences of the words in a random text. Clumps are shown here to play a central role in these problems. They can be used to calculate relevant quantities, such as the probability that a random text contains a given number of pattern word occurrences. We provide combinatorial properties that

greatly simplify the classical enumeration of these texts by inclusion-exclusion approaches. We describe two clump automata that can be used to efficiently calculate generating functions.

Mireille A. Regnier
INRIA
mireille.regnier@inria.fr

Billy Fang
Princeton University
bfang@princeton.edu

Daria Iakovishina
Ecole Polytechnique
daria.iakovishina@inria.fr

CP0

Fast Shortest-Path Distance Queries on Road Networks by Pruned Highway Labeling

We propose a new labeling method for shortest-path and distance queries on road networks. We present a new framework (i.e. data structure and query algorithm) referred to as *highway-based labelings* and a preprocessing algorithm for it named *pruned highway labeling*. Our proposed method has several appealing features from different aspects in the literature. The experimental results show that the proposed method is much faster than the previous state-of-the-art labeling method in the preprocessing time.

Takuya Akiba
The University of Tokyo
JST, ERATO, Kawarabayashi Project
t.akiba@is.s.u-tokyo.ac.jp

Yoichi Iwata
The University of Tokyo
y.iwata@is.s.u-tokyo.ac.jp

Ken-ichi Kawarabayashi
National Institute of Informatics
JST, ERATO, Kawarabayashi Project
k.keniti@nii.ac.jp

Yuki Kawata
The University of Tokyo
JST, ERATO, Kawarabayashi Project
kawatea@lager.is.s.u-tokyo.ac.jp

CP0

Tight Analysis of Randomized Rumor Spreading in Complete Graphs

We present a tight analysis of the basic randomized rumor spreading process in complete graphs introduced by Frieze and Grimmett (1985), where in each round of the process each node knowing the rumor gossips the rumor to a node chosen uniformly at random. The process starts with a single node knowing the rumor. We show that the number S_n of rounds required to spread a rumor in a complete graph with n nodes is very closely described by $\log_2 n$ plus $(1/n)$ times the completion time of the coupon collector process. This in particular gives very precise bounds for the expected runtime of the process, namely $\lceil \log_2 n \rceil +$

$\ln n - 1.116 \leq E[S_n] \leq \lceil \log_2 n \rceil + \ln n + 2.765 + o(1)$.

Marvin Künnemann

Max-Planck-Institut für Informatik, Saarbrücken,
Germany
marvin@mpi-inf.mpg.de

Benjamin Doerr

Max Planck Institute for Informatics
doerr@mpi-inf.mpg.de

CP0

A Back-to-Basics Empirical Study of Priority Queues

The theory community has proposed several new heap variants in the recent past which have remained largely untested experimentally. We take the field back to the drawing board, with straightforward implementations of both classic and novel structures using only standard, well-known optimizations. We study the behavior of each structure on a variety of inputs, including artificial workloads, workloads generated by running algorithms on real map data, and workloads from a discrete event simulator used in recent systems networking research. We provide observations about which characteristics are most correlated to performance. For example, we find that the L1 cache miss rate appears to be strongly correlated with wallclock time. We also provide observations about how the input sequence affects the relative performance of the different heap variants. Overall, our findings suggest that while the conventional wisdom holds in some cases, it is sorely mistaken in others.

Daniel H. Larkin

Princeton University
dhlarkin@cs.princeton.edu

Siddhartha Sen

Microsoft Research
sisen@microsoft.com

Robert Tarjan

Princeton University
Microsoft Research
ret@cs.princeton.edu

CP0

Expected External Profile of Patricia Tries

We analyze the expected number of leaves at level k for PATRICIA tries on n binary strings generated by a memoryless source. We study three natural ranges of k with respect to n . The most interesting range is k growing logarithmically with n , where we apply Mellin transforms, analytic depoissonization, and the saddle point method. The analysis features an intricate Mellin transform that satisfies a recursive functional equation, which presents interesting analytic challenges.

Abram Wagner

Department of Computer Science
Purdue University
anwagner@purdue.edu

Charles Knessl

University of Illinois at Chicago
knessl@uic.edu

Wojciech Szpankowski

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

CP0

Survivors in Leader Election Algorithms

We consider the number of survivors in a broad class of fair leader election algorithms after a number of election rounds. We give sufficient conditions for the number of survivors to converge to a product of independent identically distributed random variables. The number of terms in the product is determined by the round number considered. Each individual term in the product is a limit of a scaled random variable associated with the splitting protocol. The proof is established via convergence (to 0) of the first-order Wasserstein distance from the product limit. In a broader context, the paper is a case study of a class of stochastic recursive equations. We give two illustrative examples, one with binomial splitting protocol (for which we show that a normalized version is asymptotically Gaussian), and one with power distribution splitting protocol.

Ravi Kalpathy, Hosam M. Mahmoud

The George Washington University
kalpathy@gwu.edu, hosam@gwu.edu

Walter Rosenkrantz

University of Massachusetts Amherst
w_rosenkrantz@verizon.net

CP0

On the Average Number of Edges in Theta Graphs

Not available at the time of publication.

Pat Morin, Sander Verdonschot

Carleton University
morin@scs.carleton.ca, sverdons@connect.carleton.ca

CP0

Top-K Substring Matching for Auto-Completion

Given the user's input as a query, auto-completion selects the top- k strings with the highest scores from the strings matching the query in a dictionary. In this paper, we present a novel approach to solve the top- k substring matching problem for auto-completion using a small space. Experimental results show that our algorithm can suggest top- k completion sufficiently fast, while taking much less space than a compressed full-text index of the dictionary.

Yuzuru Okajima

VALWAY Technology Center, NEC Soft, Ltd.
okajima-yuzuru@mx.nes.nec.co.jp

CP0

Triangle Listing Algorithms: Back from the Diver-sion

We present a unifying framework for all common algorithms listing the triangles of a graph. This framework shows that most of them run in $\mathcal{O}(a(G)m)$ time, where $a(G)$ is the arboricity of G – an upper bound that had been proven only for the algorithm of Chiba and Nishizeki (*SIAM J. Computing*, 1985). More importantly, we show

by experimentation that this latter algorithm, if implemented carefully, actually outperforms all subsequently proposed algorithms.

Mark Ortman, Ulrik Brandes
Department of Computer & Information Science
University of Konstanz
mark.ortmann@uni-konstanz.de, ulrik.brandes@uni-konstanz.de

CP0

Multi-Pivot Quicksort: Theory and Experiments

The idea of multi-pivot quicksort has recently received attention after Yaroslavskiy proposed a dual pivot quicksort algorithm that outperforms standard quicksort under the Java JVM. More recently, this algorithm has been analysed in terms of comparisons and swaps by Wild and Nebel. We perform the previous experiments using a native C implementation, removing extraneous effects of the JVM. Additionally, we provide analyses on cache behavior of these algorithms. We then provide strong evidence that cache behavior is causing most of the performance differences. Lastly, we build upon prior work in multi-pivot quicksort and propose a 3-pivot variant that performs very well in theory and practice.

Shrinu Kushagra
University of Waterloo
David R. Cheriton School of Computer Science
skushagr@uwaterloo.ca

Aurick Qiao
University of Waterloo
Department of Computer Science
a2qiao@uwaterloo.ca

Alejandro Lopez-Ortiz
University of Waterloo
David R. Cheriton School of Computer Science
alopez-o@uwaterloo.ca

J. Ian Munro
University of Waterloo
jmunro@uwaterloo.ca

CP0

Precomputation Techniques for the Stochastic on-Time Arrival Problem

We consider the stochastic on-time arrival (SOTA) problem of finding the optimal routing strategy to reach a destination within a pre-specified time budget and provide the first results on preprocessing techniques to speed up the query time. We start by identifying some properties of the SOTA problem that limit the preprocessing techniques that can be used in this setting, and define the stochastic variants of two techniques that can be adapted to the SOTA problem, namely reach and arc-flags. We present the preprocessing and query algorithms for each technique, and also an extension to the standard reach method that provides additional pruning. Finally, we explain the limitations of this approach due to the cost of the preprocessing phase and present a fast heuristic preprocessing scheme. Numerical results for San Francisco, Luxembourg and a synthetic network show up to an order of magnitude improvement in the query time for short queries, with larger gains expected for

longer queries.

Guillaume Sabran
Ecole Polytechnique, Paris
guillaume.sabran@polytechnique.org

Samitha Samaranyake
University of California, Berkeley
UC Berkeley
samitha@berkeley.edu

Alexandre Bayen
University of California, Berkeley
bayen@berkeley.edu

CP0

On The Average-Case Complexity of the Bottleneck Tower of Hanoi Problem

The Bottleneck Tower of Hanoi (BTH) problem, posed in 1981 by Wood, is a natural generalization of the classic Tower of Hanoi (TH) problem. There, a generalized placement rule allows a larger disk to be placed higher than a smaller one if their size difference is less than a given parameter $k \geq 1$. The objective is to compute a shortest move-sequence transferring a legal (under the above rule) configuration of n disks on three pegs to another legal configuration. In SOFSEM'07, Dinitz and the second author established tight asymptotic bounds for the worst-case complexity of the BTH problem, for all n and k . Moreover, they proved that the average-case complexity is asymptotically the same as the worst-case complexity, for all $n > 3k$ and $n \leq k$, and conjectured that the same phenomenon also occurs in the complementary range $k < n \leq 3k$. In this paper we settle this conjecture of Dinitz and the second author in the affirmative.

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

Noam Solomon
Tel Aviv University
noam.solom@gmail.com

CP0

Flow-Based Guidebook Routing

Using public transit to get from A to B, one is typically confronted with the agony of choice: Considering Pareto-costs (including travel time and number of transfers), there are often several optimal journeys for a given departure time. Over a day or week, the number of possible journeys accumulates, with some routes being useful for long periods of time while others are only valid for a single departure. Such large solution sets (often containing many very similar ones) lead to difficulties to present a reasonable subset to the user and also increase query times and space consumption of routing algorithms. In this paper, we aim for computing so called guidebook routes, which are characterised by being *stable* over long periods. The goal is to detect a small set of guidebook routes, such that for any departure time going by guidebook provides a reasonable journey with the costs exceeding the optimal ones not too much.

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

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

CP0

Connection Scan Accelerated

We study the problem of efficiently computing journeys in timetable networks. Our algorithm optimally answers profile queries, computing all journeys given a time interval. Our study demonstrates that queries can be answered optimally on large country-scale timetable networks within several milliseconds and fast delay integration is possible. Previous work either had to drop optimality or only considered comparatively small timetable networks. Our technique is a combination of the Connection Scan Algorithm and multi-level overlay graphs.

Ben Strasser, Dorothea Wagner
 Karlsruhe Institute of Technology
 Institute of Theoretical Informatics
 strasser@kit.edu, dorothea.wagner@kit.edu

CP0

Typical Depth of a Digital Search Tree Built on a General Source

The digital search tree (dst) plays a central role in compression algorithms. This important structure is a mixing of a digital structure with a binary search tree. Its probabilistic analysis is thus involved, even in the case when the text is produced by a simple source (a memoryless source, or a Markov chain). After the seminal paper of Flajolet and Sedgewick (1986), many papers, published between 1990 and 2005, dealt with general simple sources, and perform the analysis of the main parameters of dst's (namely, path length, profile, typical depth). Here, we are interested in a more realistic analysis, when the words are emitted by a general source. There exist previous analyses that have been performed for general sources, but the case of dst's has not yet been considered. The idea of this study is due to Philippe Flajolet and the first steps of the work were performed with him, during the end of 2010. Our paper is dedicated to his memory.

Brigitte Vallee
 CNRS, GREYC Laboratory
 brigitte.vallee@unicaen.fr

Kanal Hun
 GREYC laboratory
 kanal.hun@unicaen.fr

CP0

Order Constraints for Single Machine Scheduling with Non-Linear Cost

Typically in a scheduling problem we are given jobs of different processing times p_j and different priority weights w_j , and need to schedule them on a single machine in order to minimize a specific cost function. In this paper we consider the non-linear objective function $\sum w_j C_j^\beta$, where C_j is the completion time of job j and $\beta > 0$ is some arbitrary real constant. Except for $\beta = 1$ the complexity status of this problem is open. Past research mainly focused on the quadratic case ($\beta = 2$) and proposed different techniques to speed up exact algorithms. This paper proposes new pruning rules and generalizations of existing rules to non-

integral β . An experimental study evaluates the impact of the proposed rules on the exact algorithm A*.

Christoph Dürr
 CNRS, LIP6
 Université Pierre et Marie Curie
 christoph.durr@lip6.fr

Oscar C. Vásquez

LIP6
 Université Pierre et Marie Curie
 oscar.vasquez@lip6.fr

CP0

Simplifying Massive Planar Subdivisions

We present the first I/O- and practically-efficient algorithm for simplifying a planar subdivision, such that no point is moved more than a given distance ε_{xy} and such that neighbor relations between faces (homotopy) are preserved. Under some practically realistic assumptions, our algorithm uses $\mathcal{O}(\text{SORT}(N))$ I/Os, where N is the size of the decomposition and $\text{SORT}(N)$ is the number of I/Os need to sort in the standard external-memory model of computation.

Jakob Truelsen, Lars Arge, Jungwoo Yang
 Madalgo, Aarhus University
 antialize@gmail.com, large@madalgo.au.dk, jungwoo@madalgo.au.dk

CP1

Better Approximation Bounds for the Joint Replenishment Problem

We provide new approximation results for the Joint Replenishment Problem (JRP). In the offline case, we give an 1.791-approximation algorithm, breaking the current barrier of 1.8, and we show that for the variant with linear costs, the LP integrality gap is at least 1.09. In the online case, we show that the competitive ratio for JRP is at least 2.754 and we prove the optimal ratio of 2 for the version of JRP with deadlines.

Marcin Bienkowski
 University of Wrocław
 Institute of Computer Science
 marcin.bienkowski@cs.uni.wroc.pl

Jaroslaw Byrka
 Institute of Computer Science
 University of Wrocław, Poland
 jby@cs.uni.wroc.pl

Marek Chrobak

Department of Computer Science
 University of California at Riverside, US
 marek@cs.ucr.edu

Lukasz Jez
 Department of Computer, Control, and Management Engineering
 Sapienza University of Rome, Italy
 lukasz.jez@cs.uni.wroc.pl

Dorian Nogneng
 LIX, Ecole Polytechnique, Palaiseau, France
 d.nogneng@free.fr

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

CP1

Improved Approximation Algorithm for Two-Dimensional Bin Packing

We study the two-dimensional bin packing problem with and without rotations. Here we are given a set of two-dimensional rectangular items I and the goal is to pack these into a minimum number of unit square bins. We consider the *orthogonal packing* case where the edges of the items must be aligned parallel to the edges of the bin. Our main result is a 1.405-approximation for two-dimensional bin packing with and without rotation, which improves upon a recent 1.5 approximation due to Jansen and Prädél. We also show that a wide class of rounding based algorithms cannot improve upon the factor of 1.5.

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

Arindam Khan

Georgia Tech
 Georgia Tech
 akhan67@gatech.edu

CP1

A Constant Factor Approximation Algorithm for Fault-Tolerant k -Median

In this paper, we consider the fault-tolerant k -median problem and give the *first* constant factor approximation algorithm for it. In the fault-tolerant generalization of classical k -median problem, each client j needs to be assigned to at least $r_j \geq 1$ distinct open facilities. The service cost of j is the sum of its distances to the r_j facilities, and the k -median constraint restricts the number of open facilities to at most k . Previously, a constant factor was known only for the special case when all r_j s are the same, and a logarithmic approximation ratio was known for the general case. We also consider the fault-tolerant facility location problem, where the service cost of j can be a weighted sum of its distance to the r_j facilities. We give a simple constant factor approximation algorithm, generalizing several previous results which only work for nonincreasing weight vectors.

Mohammadtaghi Hajiaghayi
 University of Maryland, College Park
 hajiaghag@cs.umd.edu

Wei Hu, Jian Li
 Tsinghua University
 huwei9527@gmail.com, lijian83@mail.tsinghua.edu.cn

Shi Li
 Princeton University
 shili@cs.princeton.edu

Barna Saha
 AT&T Research Laboratory

barna@research.att.com

CP1

Better Algorithms and Hardness for Broadcast Scheduling Via a Discrepancy Approach

We study the broadcast scheduling problem with the objective of minimizing the average response time. We give an $\tilde{O}(\log^{1.5} n)$ approximation algorithm for the problem improving upon the previous $\tilde{O}(\log^2 n)$ approximation. We also show an $\Omega(\log^{1/2-\epsilon} n)$ hardness result, and an integrality gap of $\Omega(\log n)$ for the natural LP relaxation for the problem. Prior to our work, only NP-Hardness and a (tiny) constant integrality gap was known.

Nikhil Bansal
 Eindhoven University of Technology
 bansal@gmail.com

Moses Charikar
 Princeton University
 mooses@cs.princeton.edu

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

Shi Li
 Princeton University
 shili@cs.princeton.edu

CP1

A Mazing $2 + \epsilon$ Approximation for Unsplittable Flow on a Path

We study the unsplittable flow on a path problem (UFP), which arises naturally in many applications such as bandwidth allocation, job scheduling, and caching. Here we are given a path with nonnegative edge capacities and a set of tasks, which are characterized by a subpath, a demand, and a profit. The goal is to find the most profitable subset of tasks whose total demand does not violate the edge capacities. In this paper we present a PTAS for δ -large tasks, for any constant $\delta > 0$. Key to this result is a complex geometrically inspired dynamic program. Each task is represented as a segment underneath the capacity curve, and we identify a proper maze-like structure so that each *corridor* of the maze is *crossed* by only $O(1)$ tasks in the optimal solution. The maze has a tree topology, which guides our dynamic program. Our result implies a $2 + \epsilon$ approximation for UFP, for any constant $\epsilon > 0$.

Aris Anagnostopoulos
 Sapienza University of Rome
 aris@dis.uniroma1.it

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

Stefano Leonardi
 Sapienza University of Rome
 leonardi@dis.uniroma1.it

Andreas Wiese
 MPII Saarbrücken

awiese@mpi-inf.mpg.de

CP2

Minimum Common String Partition Parameterized by Partition Size Is Fixed-Parameter Tractable

The NP-hard MINIMUM COMMON STRING PARTITION problem asks whether two strings x and y can each be partitioned into at most k substrings such that both partitions use exactly the same substrings in a different order. We present the first fixed-parameter algorithm for MINIMUM COMMON STRING PARTITION using only parameter k .

Laurent Bulteau, Christian Komusiewicz
Institut für Softwaretechnik und Theoretische Informatik
TU Berlin
l.bulteau@gmail.com, christian.komusiewicz@univ-nantes.fr

CP2

Interval Deletion Is Fixed-Parameter Tractable

We study the minimum *interval deletion* problem, which asks for the removal of a set of at most k vertices to make a graph on n vertices into an interval graph. We present a parameterized algorithm of runtime $10^k \cdot n^{O(1)}$ for this problem, thereby showing its fixed-parameter tractability.

Yixin Cao, Dániel Marx
Institute for Computer Science and Control
Hungarian Academy of Sciences
yixin@sztaki.hu, dmarx@cs.bme.hu

CP2

Finding Small Patterns in Permutations in Linear Time

Given two permutations σ and π , the PERMUTATION PATTERN problem asks if σ is a subpattern of π . We show that the problem can be solved in time $2^{O(\ell^2 \log \ell)} \cdot n$, where $\ell = |\sigma|$ and $n = |\pi|$. In other words, the problem is fixed-parameter tractable parameterized by the size of the subpattern to be found. We introduce a novel type of decompositions for permutations and a corresponding width measure. We present a linear-time algorithm that either finds σ as a subpattern of π , or finds a decomposition of π whose width is bounded by a function of $|\sigma|$. Then we show how to solve the PERMUTATION PATTERN problem in linear time if a bounded-width decomposition is given in the input.

Sylvain Guillemot
Hungarian Academy of Sciences (MTA SZTAKI),
Budapest, Hunga
sguillem@sztaki.hu

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

CP2

An Excluded Grid Theorem for Digraphs with For-

bidden Minors

In 1997, Reed and later Johnson, Robertson, Seymour and Thomas conjectured the existence of a function f of k such that every digraph of directed tree-width at least $f(k)$ contains a directed grid of order k . In this paper we prove the conjecture for the case of digraphs excluding a fixed undirected graph as a minor. For algorithmic applications our theorem is particularly interesting as it covers those classes of digraphs to which, on undirected graphs, theories based on the excluded grid theorem such as bidimensionality theory apply.

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

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

CP2

Efficient Computation of Representative Sets with Applications in Parameterized and Exact Algorithms

We demonstrate how the efficient construction of representative families can be a powerful tool for designing single-exponential parameterized and exact exponential time algorithms.

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

Daniel Lokshtanov
UCSD
daniello@ii.uib.no

Saket Saurabh
IMS
saket@imsc.res.in

CP3

Implicit Manifold Reconstruction

Let P be a dense set of points sampled from an m -dimensional compact smooth manifold Σ in \mathbb{R}^d . We show how to construct an implicit function $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}^{d-m}$ from P so that the zero-set S_φ of φ contains a homeomorphic approximation of Σ . The Hausdorff distance between Σ and this homeomorphic approximation is at most ε^τ for any fixed $\tau < 2$. Moreover, for every point x at distance ε^τ or less from Σ , the normal space of S_φ at x makes an $O(\varepsilon^{(\tau-1)/2})$ angle with the normal space of Σ at the point nearest to x . The function φ has local support, which makes local homeomorphic reconstruction possible without a complete sampling.

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

Man-Kwun Chiu
HKUST

chiumk@cse.ust.hk

CP3

Robust Satisfiability of Systems of Equations

We study the problem of *robust satisfiability* of systems of nonlinear equations, namely, whether for a given continuous function $f : K \rightarrow \mathbb{R}^n$ on a finite simplicial complex K and $\alpha > 0$, it holds that each function $g : K \rightarrow \mathbb{R}^n$ such that $\|g - f\|_\infty \leq \alpha$, has a root in K . Via a reduction to the extension problem of maps into a sphere, we show that this problem is decidable if $\dim K \leq 2n - 3$. This is a substantial extension of previous computational applications of *topological degree* and related concepts in numerical and interval analysis. Via a reverse reduction we prove that the problem is undecidable when $\dim K \geq 2n - 2$, where the threshold comes from the *stable range* in homotopy theory. For the lucidity of our exposition, we focus on the setting when f is piecewise linear. Such functions can approximate general continuous functions, and thus we get approximation schemes and undecidability of the robust satisfiability in other possible settings.

Peter Franek

Institute of Computer Science
Academy of Sciences of the Czech Republic
peter.franek@gmail.com

Marek Krcal
Charles University
krcal@kam.mff.cuni.cz

CP3

Solving 1-Laplacians of Convex Simplicial Complexes in Nearly Linear Time: Collapsing and Expanding a Topological Ball

We present an efficient algorithm for solving a linear system arising from the 1-Laplacian of a collapsible simplicial complex with a known collapsing sequence. When combined with a result of Chillingworth, our algorithm is applicable to convex simplicial complexes embedded in R^3 . The running time of our algorithm has nearly-linear dependency on the size of the complex, and logarithmic dependency on its numerical properties. Our algorithm is based on projection operators for higher dimensional objects, and combinatorial steps for transferring between them. The former relies on decomposing flows into circulations and potential flows using fast solvers for graph Laplacians, and the latter relates Gaussian elimination to topological properties of simplicial complexes.

Michael B. Cohen
MIT
n/a

Brittany Terese Fasy
Tulane
n/a

Gary Miller
Carnegie Mellon University
glmiller@cs.cmu.edu

Amir Nayyeri

Department of Computer Science
Carnegie Mellon University
amirn@cs.cmu.edu

Richard Peng
MIT
n/a

Noel J. Walkington
Department of Mathematical Sciences
Carnegie Mellon University
noelw@andrew.cmu.edu

CP3

On the Computational Complexity of Betti Numbers: Reductions from Matrix Rank

We give evidence for the difficulty of computing Betti numbers of simplicial complexes over a finite field. We do this by reducing the rank computation for sparse matrices with m non-zero entries to computing Betti numbers of simplicial complexes consisting of at most a constant times m simplices. Together with the known reduction in the other direction, this implies that the two problems have the same computational complexity.

Herbert Edelsbrunner
IST Austria
edels@ist.ac.at

Salman Parsa
Duke University
IST Austria
salparsa@cs.duke.edu

CP3

Approximating Local Homology from Samples

Recently, multi-scale notions of local homology (a variant of persistent homology) have been used to study the local structure of spaces around a given point from a point cloud sample. Current reconstruction guarantees rely on constructing embedded complexes which become difficult to construct in higher dimensions. We show that the persistence diagrams used for estimating local homology can be approximated using families of Vietoris-Rips complexes, whose simpler construction are robust in any dimension. To the best of our knowledge, our results, for the first time make applications based on local homology, such as stratification learning, feasible in high dimensions.

Primoz Skraba

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

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

CP4

Near Linear Time Approximation Schemes for Uncapacitated and Capacitated B-Matching Problems in Nonbipartite Graphs

We present the first fully polynomial approximation schemes for the maximum weighted (uncapacitated or capacitated) b -Matching problem for nonbipartite graphs that run in time (near) linear in the number of edges, that is, given any $\delta > 0$ the algorithm produces a $(1 - \delta)$ approximation in $O(m \text{poly}(\delta^{-1}, \log n))$ time. We provide frac-

tional solutions for the standard linear programming formulations for these problems and subsequently also provide fully polynomial (near) linear time approximation schemes for rounding the fractional solutions.

Kook Jin Ahn

University of Pennsylvania
kookjin@google.com

Sudipto Guha

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

CP4

Towards $(1 + \epsilon)$ -Approximate Flow Sparsifiers

A useful approach to compress a large network G is to represent it with a flow-sparsifier, ie, a small network H that supports the same flows as G , up to a factor $q \geq 1$ called the quality of sparsifier. Specifically, we assume the network G contains a set of k terminals T , shared with the network H , i.e., $T \subseteq V(G) \cap V(H)$, and we want H to preserve all multicommodity flows that can be routed between the terminals T . The challenge is to construct H that is small. These questions have received a lot of attention in recent years, leading to some known tradeoffs between the sparsifier quality q and its size $|V(H)|$. Nevertheless, it remains an outstanding question whether every G admits a flow-sparsifier H with quality $q = 1 + \epsilon$, or even $q = O(1)$, and size $|V(H)| \leq f(k, \epsilon)$ (in particular, independent of $|V(G)|$ and edge capacities). Making a first step in this direction, we present new constructions for several scenarios:

Alexandr Andoni

Microsoft Research
andoni@microsoft.com

Anupam Gupta

Carnegie Mellon University
anupamg@cs.cmu.edu

Robert Krauthgamer

Weizmann Institute of Science
robert.krauthgamer@weizmann.ac.il

CP4

Improved Bounds and Algorithms for Graph Cuts and Network Reliability

Karger (*SIAM Journal on Computing*, 1999) developed the first fully-polynomial approximation scheme to estimate the probability that a graph G becomes disconnected, given that its edges are removed independently with probability p . This algorithm runs in $n^{5+o(1)}\epsilon^{-3}$ time. We improve this in two key ways. First, there is a certain key sub-problem encountered by Karger, for which a generic estimation procedure is employed. We show that a more efficient algorithm can be used. Second, we show better bounds on the number of edge cuts which are likely to fail. Karger's analysis depends on bounds for various graph parameters; we show that these bounds cannot be simultaneously tight. These techniques improve the runtime to $n^{3+o(1)}\epsilon^{-2}$. A key driver of Karger's approach is bounding the number of small cuts: we also show how to improve this when the min-cut size is "small" and odd, augmenting, in part, a

result of Bixby (*Bull. AMS*, 1974).

David Harris, Aravind Srinivasan

University of Maryland
davidgharris29@hotmail.com, srin@cs.umd.edu

CP4

An Almost-Linear-Time Algorithm for Approximate Max Flow in Undirected Graphs, and Its Multicommodity Generalizations

In this paper, we present an almost linear time algorithm for solving approximate maximum flow in undirected graphs. In particular, given a graph with m edges we show how to produce a $1 - \epsilon$ approximate maximum flow in time $O(m^{1+o(1)}\epsilon^{-2})$. Furthermore, we present this algorithm as part of a general framework that also allows us to achieve a running time of $O(m^{1+o(1)}k^2\epsilon^{-2})$ for the maximum concurrent k -commodity flow problem, the first such algorithm with an almost linear dependence on m .

Jonathan Kelner

MIT
kelner@mit.edu

Yin Tat Lee

Massachusetts Institute of Technology
yintat@mit.edu

Lorenzo Orecchia, Aaron Sidford

MIT
orecchia@mit.edu, sidford@mit.edu

CP4

Computing Cut-Based Hierarchical Decompositions in Almost Linear Time

We present a fast construction for the hierarchical tree decompositions that lie at the heart of oblivious routing strategies and that form the basis for approximation and online algorithms for various cut problems in graphs. Given an undirected graph $G = (V, E, c)$, we compute a single tree $T = (V_T, E_T, c_T)$ s.t. T approximates the cut-structure of G up to a factor of $O(\log^4 n)$. The best existing construction by Harrelson et al. just guarantees a polynomial running time but offers a better approximation guarantee of $O(\log^2 n \log \log n)$. In other words, for a graph $G = (V, E)$ with a subset S of terminals, we compute a tree T with at most $2|S|$ vertices such that T is a flow-sparsifier for S in G with quality $O(\log^2 |V| \log^2 |S|)$. Our algorithm runs in time $O(\text{polylog } n \cdot T(m, 1/\log^3 |V|))$ where $T(m, \epsilon)$ is the time for computing an approximate maxflow. The latter is almost linear due to the recent results of Sherman and Kelner et al.

Harald Räcke, Chintan Shah, Hanjo Täubig

Technische Universität München
raecke@in.tum.de, shahc@in.tum.de, taebig@in.tum.de

CP5

Clustering and Mixing Times for Segregation Models on \mathbb{Z}^2

The Schelling segregation model attempts to explain racial segregation in cities. Schelling considered residents of two types, where one prefers neighbors of the same type. He showed through simulations that even mild preferences of this type can lead to segregation if residents move when-

ever they are unhappy. We generalize the Schelling model to include a broad class of influence dynamics, called the General Influence Model. We show that in this model, the dynamics will be rapidly mixing and cities will be integrated if the racial bias is sufficiently low. We show complementary results for two broad classes of influence functions: Increasing Bias Functions, where one's likelihood of moving increases each time someone of the same color leaves, and Threshold Bias Functions, reminiscent of Schelling's original model. For both classes, we show that when the bias is sufficiently high, the dynamics take exponential time to mix, a large "ghetto" will form, and we will have segregation.

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

CP5

Mcmc Sampling Colourings and Independent Sets of $G(n, d/n)$ Near Uniqueness Threshold

Sampling from the Gibbs distribution is a well-studied problem in CS. We focus on the k -colouring model and the hard-core model with fugacity λ , when the underlying graph is a random graph $G(n, p)$, for $p = d/n$ and fixed d . Our approach is based on the Markov Chain Monte Carlo method, i.e. Glauber (block) dynamics. We show a dramatic improvement on the bounds for rapid mixing in terms of colours and the fugacity for the corresponding models.

Charilaos Efthymiou
University of Warwick - DIMAP
efthymiou@math.uni-frankfurt.de

CP5

A Simple FPTAS for Counting Edge Covers

An edge cover of a graph is a set of edges such that every vertex has at least an adjacent edge in it. Previously, approximation algorithm for counting edge covers is only known for 3 regular graphs and it is randomized. We design a very simple deterministic fully polynomial-time approximation scheme (FPTAS) for counting the number of edge covers for any graph. Our main technique is correlation decay, which is a powerful tool to design FPTAS for counting problems. In order to get FPTAS for general graphs without degree bound, we make use of a stronger notion called computationally efficient correlation decay, which is introduced in [Li, Lu, Yin SODA 2012].

Chengyu Lin, Jingcheng Liu
Shanghai Jiao Tong University
linmrain@gmail.com, liuexp@gmail.com

Pinyan Lu
Microsoft Research Asia
pinyanl@microsoft.com

CP5

Arboricity and Spanning-Tree Packing in Random Graphs with an Application to Load Balancing

We study the arboricity A and the maximum number T of edge-disjoint spanning trees of the classical random graph $G(n, p)$. For all $p(n) \in [0, 1]$, we show that, with high

probability, T is precisely the minimum between δ and $\lfloor m/(n-1) \rfloor$, where δ is the smallest degree of the graph and m denotes the number of edges. Moreover, we explicitly determine a sharp threshold value for p such that: above this threshold, T equals $\lfloor m/(n-1) \rfloor$ and A equals $\lfloor m/(n-1) \rfloor$; and below this threshold, T equals δ , and we give a two-value concentration result for the arboricity A in that range. Finally, we include a stronger version of these results in the context of the random graph process where the edges are sequentially added one by one. A direct application of our result gives a sharp threshold for the maximum load being at most k in the two-choice load balancing problem, where $k \rightarrow \infty$.

Pu Gao
University of Toronto
pu.gao@utoronto.ca

Xavier Pérez Giménez, Cristiane M. Sato
University of Waterloo
xperez@uwaterloo.ca, cmsato@uwaterloo.ca

CP5

Uniform Random Sampling of Simple Branched Coverings of the Sphere by Itself

We present the first polynomial uniform random sampling algorithm for simple branched coverings of degree n of the sphere by itself. More precisely, our algorithm generates in linear time *increasing quadrangulations*, which are equivalent combinatorial structures. Our result is based on the identification of some canonical *labelled* spanning trees, and yields a constructive proof of a celebrated formula of Hurwitz for the number of some factorizations of permutations in transpositions. The previous approaches were either non constructive or lead to exponential time algorithms for the sampling problem.

Enrica Duchi, Dominique Poulalhon
Université Paris Diderot
enrica.duchi@liafa.univ-paris-diderot.fr,
dominique.poulalhon@liafa.univ-paris-diderot.fr

Gilles Schaeffer
CNRS / École Polytechnique
Department LABRI
gilles.schaeffer@lix.polytechnique.fr

CP6

Polynomial Solvability of Variants of the Trust-Region Subproblem

We consider an optimization problem of the form

$$\begin{aligned} \min \quad & x^T Q x + c^T x \\ \text{s.t.} \quad & \|x - \mu_h\| \leq r_h, \quad h \in S, \\ & \|x - \mu_h\| \geq r_h, \quad h \in K, \\ & x \in P, \end{aligned}$$

where $P \subseteq \mathbb{R}^n$ is a polyhedron defined by m inequalities and Q is general and the $\mu_h \in \mathbb{R}^n$ and the r_h quantities are given. In the case $|S| = 1$, $|K| = 0$ and $m = 0$ one obtains the classical trust-region subproblem; a strongly NP-hard problem which has been the focus of much interest because of applications to combinatorial optimization and nonlinear programming. We prove that for each fixed pair $|S|$ and $|K|$ our problem can be solved in polynomial time provided that either (1) $|K| > 0$ and the number of faces of P that

intersect $\bigcap_h \{x \in \mathbb{R}^n : \|x - \mu_h\| \leq r_h, 1 \leq j \leq p\}$ is polynomially bounded, or (2) $|K| = 0$ and m is bounded.

Daniel Bienstock

Columbia University IEOR and APAM Departments
IEOR Department
dano@columbia.edu

Alexander Michalka
Columbia University
admichalka@gmail.com

CP6

On the Lattice Isomorphism Problem

We study the Lattice Isomorphism Problem (LIP) and present an algorithm solving it in time $n^{O(n)}$ times a polynomial in the input size, where n is the rank of the input lattices. A crucial component is a new generalized isolation lemma, which can isolate n linearly independent vectors in a given subset of Z^n and might be useful elsewhere. We also prove that LIP lies in the complexity class SZK.

Ishay Haviv

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

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

CP6

The Complexity of Order Type Isomorphism

The order type of a point set in R^d maps each $(d+1)$ -tuple of points to its orientation (e.g., clockwise or counterclockwise in R^2). Two point sets X and Y have the same order type if there exists a mapping f from X to Y for which every $(d+1)$ -tuple $(a_1, a_2, \dots, a_{d+1})$ of X and the corresponding tuple $(f(a_1), f(a_2), \dots, f(a_{d+1}))$ in Y have the same orientation. In this paper we investigate the complexity of determining whether two point sets have the same order type. We provide an $O(n^d)$ algorithm for this task, thereby improving upon the $O(n^{\lfloor 3d/2 \rfloor})$ algorithm of Goodman and Pollack (1983). The algorithm uses only order type queries and also works for abstract order types (or acyclic oriented matroids). Our algorithm is optimal, both in the abstract setting and for realizable points sets if the algorithm only uses order type queries.

Greg Aloupis
Universite Libre de Bruxelles
aloupis.greg@gmail.com

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

Stefan Langerman
Universite Libre de Bruxelles
stefan.langerman@ulb.ac.be

Ozgur Ozkan
Polytechnic Institute of New York University
ozgurozkan@gmail.com

Stefanie Wuhler
Saarland University

swuhrer@mmci.uni-saarland.de

CP6

Positivity Problems for Low-Order Linear Recurrence Sequences

We consider two decision problems for linear recurrence sequences (LRS) over the integers, namely the *Positivity Problem* (are all terms of a given LRS positive?) and the *Ultimate Positivity Problem* (are all but finitely many terms of a given LRS positive?). We show decidability of both problems for LRS of order 5 or less, with complexity in the Counting Hierarchy for Positivity, and in polynomial time for Ultimate Positivity. Moreover, we show by way of hardness that extending the decidability of either problem to LRS of order 6 would entail major breakthroughs in analytic number theory, more precisely in the field of Diophantine approximation of transcendental numbers.

Joel Ouaknine, James Worrell

Department of Computer Science
Oxford University
joel@cs.ox.ac.uk, jbw@cs.ox.ac.uk

CP6

Space Complexity of List H-Coloring: a Dichotomy

The Dichotomy Conjecture for constraint satisfaction problems (CSPs) states that every CSP is in P or is NP-complete (Feder-Vardi, 1993). It has been verified for conservative problems (also known as list homomorphism problems) by A. Bulatov (2003). We augment this result by showing that for digraph templates H , every conservative CSP, denoted LHOM(H), is solvable in logspace or is hard for NL. More precisely, we introduce a digraph structure we call a circular N, and prove the following dichotomy: if H contains no circular N then LHOM(H) admits a logspace algorithm, and otherwise LHOM(H) is hard for NL. Our algorithm operates by reducing the lists in a complex manner based on a novel decomposition of an auxiliary digraph, combined with repeated applications of Reingold's algorithm for undirected reachability (2005). Moreover, we show that the presence of a circular N can be decided in time polynomial in the size of H . We also prove an algebraic version of this dichotomy.

Laszlo Egri
Academy of Science, Budapest, Hungary
laszlo.egri@mail.mcgill.ca

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

Benoit Larose
Department of Mathematics, Concordia University,
Canada
benoit.larose@concordia.ca

Arash Rafiey
Simon Fraser University
arashr@sfu.ca

CP7

Competitive Analysis Via Regularization

We provide a framework for designing competitive online

algorithms using *regularization*, a widely used technique in online learning, particularly in online convex optimization. An online algorithm that uses regularization serves requests by computing a solution, in each step, to an objective function involving a smooth convex regularization function. Applying the technique of regularization allows us to obtain new results in the domain of competitive analysis. In our new framework we exhibit a general $O(\log m)$ -competitive deterministic algorithm for generating a fractional solution that satisfies a time-varying set of online covering and precedence constraints, where m is the number of variables. This framework allows to incorporate both service costs (over time) and setup costs into a host of applications. We then provide an $O(\log m \log n)$ -competitive randomized algorithm for the online set cover problem with service cost.

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

Shahar Chen, Joseph (Seffi) Naor
Technion
shaharch@cs.technion.ac.il, naor@cs.technion.ac.il

CP7

Dynamic Task Allocation in Asynchronous Shared Memory

Task allocation is a classic distributed problem in which a set of p potentially faulty processes must cooperate to perform a set of tasks. This paper considers a new *dynamic* version of the problem, in which tasks are injected adversarially during an asynchronous execution. We give the first asynchronous shared-memory algorithm for dynamic task allocation, and we prove that our solution is optimal within logarithmic factors. The main algorithmic idea is a randomized concurrent data structure called a *dynamic to-do tree*, which allows processes to pick new tasks to perform at random from the set of available tasks, and to insert tasks at random empty locations in the data structure.

Dan Alistarh
MIT
d.alistarh@gmail.com

James Aspnes
Yale
aspnes@cs.yale.edu

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

Rati Gelashvili
MIT
gelash@mit.edu

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

CP7

Online Steiner Tree with Deletions

In the online Steiner tree problem, the input is a set of vertices that appear one-by-one (with a metric between them),

and we have to maintain a low-cost Steiner tree on the current set of vertices. What if the set of vertices sees both additions and deletions, and we are allowed to change a small number of existing edges at each point in time? In this paper we improve on prior results, giving online algorithms that maintains a constant-competitive Steiner tree under only deletions (where we change only a constant number of edges upon each request). We also give an algorithm that maintains a Steiner tree in the fully-dynamic model (both vertex insertions and deletions), where we make a constant number of changes per request in an amortized sense.

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

Amit Kumar
IIT Delhi
amitk@cse.iitd.ac.in

CP7

Maintaining Assignments Online: Matching, Scheduling, and Flows

We consider matching and scheduling problems, where objects can be reassigned, and we want to minimize the number of reassignments done. For online matching, where the left vertices arrive online and must be matched to the right vertices, we give an algorithm that reassigns the left vertices an (amortized) constant number of times, and maintains a constant factor to the optimal load on the right vertices. For restricted machine scheduling with arbitrary sized jobs, we give an algorithm that maintains load which is $O(\log \log mn)$ times the optimum, and reassigns each job an (amortized) constant number of times. Finally, in a digraph with a single source, where sinks arrive online and want to send unit flow to the source with minimum the congestion on the edges. Suppose there is an offline flow such that the total length of the flow paths is F^* . We give an algorithm that reroutes flow along $O(F^*)$ edges and achieves a $O(1)$ -approximation to the congestion.

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

Amit Kumar
IIT Delhi
amitk@cse.iitd.ac.in

Clifford Stein
Columbia University
cliff@ieor.columbia.edu

CP7

First Come First Served for Online Slot Allocation and Huffman Coding

Can one choose a good Huffman code without knowing the distribution? Online Slot Allocation (OSA) models this and similar problems: There are n slots with known costs. Requests for items are drawn i.i.d. from a hidden distribution p . After the first request to each item i , the algorithm (knowing the slot costs and previous requests, but not p) must place i in a vacant slot j_i . The goal is to minimize $\sum_{i=1}^n p_i c(j_i)$. The optimal offline algorithm puts the j th most probable item in the j th cheapest slot. The optimal online algorithm, First Come First Served (FCFS), puts the j th distinct requested item in the j th cheapest slot.

The optimal competitive ratio for any online algorithm is $1 + H_{n-1} \sim \ln n$ for general costs and 2 for concave costs. For logarithmic cost, FCFS gives cost $\text{opt} + O(\log \text{opt})$. For Huffman coding, FCFS allocates codewords on the fly while guaranteeing cost $\text{opt} + 2 \log_2(1 + \text{opt}) + 2$.

Monik Khare
yellowpages.com
monikkhare@gmail.com

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

Neal E. Young
University of California, Riverside
neal.young@ucr.edu

CP8

Learning Sparse Polynomial Functions

We study the question of learning a sparse multi-variate polynomial over the real domain. In particular, for some unknown polynomial $f(\vec{x})$ of degree- d and k monomials, we show how to reconstruct f , within error ϵ , given only a set of examples \vec{x}_i drawn uniformly from the n -dimensional cube (or an n -dimensional Gaussian distribution), together with evaluations $f(\vec{x}_i)$ on them. The result holds even in the “noisy setting”, where we have only values $f(\vec{x}_i) + g$ where g is noise (say, modeled as a Gaussian random variable). The runtime of our algorithm is polynomial in $n, k, 1/\epsilon$ and C_d where C_d depends only on d . Note that, in contrast, in the “boolean version” of this problem, where \vec{x} is drawn from the hypercube, the problem is at least as hard as the “noisy parity problem”, where we do not know how to break the $n^{\Omega(d)}$ time barrier, even for $k = 1$, and some believe it may be impossible to do so.

Alexandr Andoni, Rina Panigrahy
Microsoft Research
andoni@microsoft.com, rina@microsoft.com

Gregory Valiant
Microsoft
gregory.valiant@gmail.com

Li Zhang
Microsoft Research
lzha@microsoft.com

CP8

Exploiting Metric Structure for Efficient Private Query Release

We consider the problem of privately answering queries defined on databases which are collections of points belonging to some metric space. We give simple, computationally efficient algorithms for answering *distance queries* defined over an arbitrary metric. Distance queries are specified by points in the metric space, and ask for the average distance from the query point to the points contained in the database, according to the specified metric. Our algorithms run efficiently in the database size and the dimension of the space, and operate in both the online query release setting, and the offline setting in which they must in polynomial time generate a fixed data structure which can answer *all* queries of interest. This represents one of the

first subclasses of linear queries for which *efficient* algorithms are known for the private query release problem, circumventing known hardness results for generic linear queries.

Zhiyi Huang
Stanford University
hzhiyi@stanford.edu

Aaron Roth
University of Pennsylvania, USA
n/a

CP8

(Nearly) Sample-Optimal Sparse Fourier Transform

We consider the problem of computing a k -sparse approximation to the discrete Fourier transform of an n -dimensional signal. Our main result is a randomized algorithm that computes such an approximation using $O(k \log n (\log \log n)^{O(1)})$ signal samples in time $O(k \log^2 n (\log \log n)^{O(1)})$, assuming that the entries of the signal are polynomially bounded. The sampling complexity improves over the recent bound of $O(k \log n \log(n/k))$ given in [HIKP12], and matches the lower bound of $\Omega(k \log(n/k) / \log \log n)$ from the same paper up to $\text{poly}(\log \log n)$ factors when $k = O(n^{1-\delta})$ for a constant $\delta > 0$.

Piotr Indyk
Massachusetts Institute of Technology
indyk@mit.edu

Michael Kapralov
Stanford University
kapralov@mit.edu

Eric Price
MIT
ecprice@mit.edu

CP8

Learning Entangled Single-Sample Gaussians

We introduce a new model of Gaussian mixtures, motivated by the setting where the data points correspond to ratings on a set of items provided by users who have widely varying expertise, and each user can rate an item at most once. In this mixture model, each item i has a true quality μ_i , each user has a variance (lack of expertise) σ_j^2 , and the rating of a user j on an item i consists of a *single* sample independently drawn from the Normal distribution $N(\mu_i, \sigma_j^2)$. The aim is to learn the unknown item qualities μ_i ’s as precisely as possible. We study the single item case and obtain efficient algorithms for the problem, complemented by near-matching lower bounds; we also obtain preliminary results for the multiple items case.

Flavio Chierichetti
Cornell University
Cornell University
flavio@chierichetti.name

Anirban Dasgupta
Yahoo! Research
anirban.dasgupta@gmail.com

Ravi Kumar, [Silvio Lattanzi](#)
 Google Inc.
 ravi.k53@gmail.com, silviolat@gmail.com

CP8

On the Compatibility of Quartet Trees

Phylogenetic tree reconstruction is a fundamental biological problem. Quartet trees are the minimal informational unit for phylogenetic classification. Here we focus on the compatibility of quartet sets. We provide several results addressing the question of what can be inferred about the compatibility of a set from its subsets. In particular we show that there are quartet sets Q of size $m = cn \log n$ in which every subset of cardinality $c'n/\log n$ is consistent, and yet no fraction of more than $1/3 + \epsilon$ of Q is consistent.

Noga Alon
 Schools of Mathematics and Computer Science
 Tel Aviv University
 nogaa@post.tau.ac.il

Sagi Snir
 University of Haifa
 ssagi@research.haifa.ac.il

Raphael Yuster
 Department of Mathematics
 University of Haifa
 raphy@math.haifa.ac.il

CP9

A New Perspective on Vertex Connectivity

Edge connectivity and vertex connectivity are two fundamental concepts in graph theory. Although by now there is a good understanding of the structure of graphs based on their edge connectivity, our knowledge in the case of vertex connectivity is much more limited. An essential tool in capturing edge connectivity are the classical results of Tutte and Nash-Williams about edge-disjoint spanning trees. We argue that connected dominating set partitions and packings are the natural analogues of edge-disjoint spanning trees in the context of vertex connectivity and we use them to obtain structural results about vertex connectivity in the spirit of those for the edge connectivity. Using this new perspective, we also prove results about vertex connectivity under vertex sampling, similar to the results of Karger about edge-connectivity under edge sampling. Our results also find applications in networking and yield routing-based broadcast algorithms with optimal throughput.

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

[Mohsen Ghaffari](#)
 MIT
 ghaffari@mit.edu

Fabian Kuhn
 University of Freiburg, Germany

kuhn@cs.uni-freiburg.de

CP9

Counting Thin Subgraphs Via Packings Faster Than Meet-in-the-Middle Time

Vassilevska and Williams (STOC 2009) showed how to count simple paths on k vertices and matchings on $k/2$ edges in an n -vertex graph in time $n^{k/2+O(1)}$. Two different algorithms with the same runtime were given by Koutis and Williams (ICALP 2009), and Björklund *et al.* (ESA 2009), via $n^{st/2+O(1)}$ -time algorithms for counting t -tuples of pairwise disjoint sets drawn from a given family of s -sized subsets of an n -element universe. Alon and Gutner (TALG 2010) showed that these problems have $\Omega(n^{\lfloor st/2 \rfloor})$ and $\Omega(n^{\lfloor k/2 \rfloor})$ lower bounds when counting by color coding. We show that the “meet-in-the-middle” exponent $st/2$ can be beaten and give an algorithm that counts in time $n^{0.4547st+O(1)}$ for t a multiple of three. This implies algorithms for counting occurrences of a fixed subgraph on k vertices and pathwidth $p \ll k$ in an n -vertex graph in $n^{0.4547k+2p+O(1)}$ time.

Andreas Björklund
 Lund University
 andreas.bjorklund@yahoo.se

[Petteri Kaski](#)
 Aalto University
 Helsinki Institute for Information Technology HIIT
 petteri.kaski@aalto.fi

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

CP9

Large Induced Subgraphs Via Triangulations and CMSO

Consider the following optimization problem. Let φ be a Counting Monadic Second Order Logic formula and t be an integer. Given a graph $G = (V, E)$, the task is to find a maximum induced subgraph $G[F]$ of treewidth at most t , that models φ . Using the theory of potential maximal cliques, we show that the problem can be solved in polynomial time for classes of graphs with polynomially many minimal separators, and in time $O(1.7347^n)$ for arbitrary graphs.

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

[Ioan Todinca](#)
 Univ. Orleans
 ioan.todinca@univ-orleans.fr

Yngve Villanger
 University of Bergen
 yngve.villanger@gmail.com

CP9

Independent Set in P_5 -Free Graphs in Polynomial

Time

We give the first polynomial time algorithm for Independent Set on P_3 -free graphs. Our algorithm also works for the Weighted Independent Set problem.

Daniel Lokshantanov
UCSD
daniello@ii.uib.no

Martin Vatshelle, Yngve Villanger
University of Bergen
martin.vatshelle@uib.no, yngve.villanger@gmail.com

CP9**Packing A -Paths in Group-Labelled Graphs via Linear Matroid Parity**

Some frameworks of packing A -paths in group-labelled graphs include various interesting problems such as Mader's disjoint \mathcal{S} -paths problem and packing odd-length A -paths. We discuss the fast solvability of these frameworks like that of Mader's problem via a reduction to the linear matroid parity problem due to Schrijver (2003). This paper characterizes the groups in question, which is a complete answer to the question, "when can we extend such a reduction technique to the generalized frameworks?"

Yutaro Yamaguchi
Department of Mathematical Informatics
The University of Tokyo
yutaro_yamaguchi@mist.i.u-tokyo.ac.jp

CP10**A QPTAS for Maximum Weight Independent Set of Polygons with Polylogarithmically Many Vertices**

The Maximum Weight Independent Set of Polygons (MWISP) problem is a fundamental problem in computational geometry. Given a set of weighted polygons in the two-dimensional plane, the goal is to find a set of pairwise non-overlapping polygons with maximum total weight. Despite a lot of research, the general case of the problem is not well-understood yet. Currently the best known polynomial time algorithm achieves an approximation ratio of n^ϵ , and it is not even clear whether the problem is APX-hard. We present a $(1 + \epsilon)$ -approximation algorithm, assuming that each polygon in the input has at most a polylogarithmic number of vertices. Our algorithm has quasi-polynomial running time, i.e., it runs in time $2^{\text{poly}(\log n, 1/\epsilon)}$.

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

CP10**On the Optimality of Approximation Schemes for the Classical Scheduling Problem**

We consider the classical scheduling problem on parallel identical machines to minimize the makespan, and achieve the following results under the Exponential Time Hypothesis (ETH).

- The scheduling problem on a constant number m of identical machines, denoted as $Pm||C_{max}$, is known to admit an FPTAS of running time $O(n) + (1/\epsilon)^{O(m)}$. We prove it is essentially the best possible in the sense

that a $(1/\epsilon)^{O(m^{1-\delta})} + n^{O(1)}$ time FPTAS for any $\delta > 0$ implies that ETH fails.

- The scheduling problem on an arbitrary number of identical machines, denoted as $P||C_{max}$, is known to admit a PTAS of running time $2^{O(1/\epsilon^2 \log^3(1/\epsilon))} + n^{O(1)}$. We prove it is nearly optimal in the sense that a $2^{O((1/\epsilon)^{1-\delta})} + n^{O(1)}$ time PTAS for any $\delta > 0$ implies that ETH fails.

We also obtain lower bounds of exact algorithms for the scheduling problem that almost matches upper bounds.

Lin Chen

Department of Mathematics
Technical University of Berlin
chenlin198662@gmail.com

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

Guochuan Zhang
Zhejiang University
zgc@zju.edu.cn

CP10**A Polynomial-Time Approximation Scheme for Fault-Tolerant Distributed Storage**

We consider a problem which has received considerable attention in systems literature because of its applications to routing in delay tolerant networks and replica placement in distributed storage systems. In abstract terms the problem can be stated as follows: Given a random variable X generated by a known product distribution over $\{0, 1\}^n$ and a target value $0 \leq \theta \leq 1$, output a non-negative vector w , with $\|w\|_1 \leq 1$, which maximizes the probability of the event $w \cdot X \geq \theta$. We provide an additive EPTAS for this problem which, for constant-bounded product distributions, runs in $\text{poly}(n) \cdot 2^{\text{poly}(1/\epsilon)}$ time and outputs an ϵ -approximately optimal solution vector w for this problem. Our approach is inspired by recent results from the complexity-theoretic study of linear threshold functions.

Constantinos Daskalakis
MIT
costis@csail.mit.edu

Anindya De
Simons Institute, Berkeley
anindya@cs.berkeley.edu

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

Ankur Moitra
MIT
moitra@mit.edu

Rocco A. Servedio
Columbia University

rocco@cs.columbia.edu

CP10

Approximating k -Center in Planar Graphs

We consider variants of the metric k -center problem. Imagine you must choose locations for k firehouses in a city so as to minimize the maximum distance of any house from the nearest firehouse. An instance is specified by a graph with arbitrary nonnegative edge-lengths, a set of vertices that can serve as firehouses (i.e. centers) and a set of vertices that represent houses. For general graphs, this problem is exactly equivalent to the metric k -center problem, which is APX-hard. We give a polynomial-time bicriteria approximation scheme when the input graph is a planar graph. We also give polynomial-time bicriteria approximation schemes for several generalizations: if, instead of all houses, we wish to cover a specified proportion of the houses; if the candidate locations for firehouses have rental costs and we wish to minimize, not the number of firehouses but the sum of their rental costs; and if the input graph is not planar but of bounded genus.

David Eisenstat, Philip Klein
Brown University
david@cs.brown.edu, klein@brown.edu

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

CP10

Polynomial Time Approximation Schemes for the Traveling Repairman and Other Minimum Latency Problems.

We give a polynomial time, $(1+\epsilon)$ -approximation algorithm (PTAS) for the traveling repairman problem in the Euclidean plane, on weighted planar graphs, and on weighted trees. This improves on the known quasi-polynomial time approximation schemes for these problems. Further, we give a PTAS for the scheduling problem of minimizing total weighted completion time on a single machine under interval order precedence constraints. This improves on the known $3/2$ -approximation for this problem.

Rene A. Sitters
Vrije Universiteit, Amsterdam
r.a.sitters@vu.nl

CP11

Relative Errors for Deterministic Low-Rank Matrix Approximations

We consider processing an $n \times d$ matrix A in a stream with row-wise updates according to a recent algorithm called Frequent Directions (Liberty, KDD 2013). This algorithm maintains an $\ell \times d$ matrix Q deterministically. We show that if one sets $\ell = \lceil k + k/\epsilon \rceil$ and returns Q_k , we achieve:

$$\|A - \pi_{Q_k}(A)\|_F^2 \leq (1 + \epsilon) \|A - A_k\|_F^2$$

where $\pi_{Q_k}(A)$ is the projection of A onto rowspace of Q_k .

Mina Ghashami
University of Utah
University of Utah

ghashami@cs.utah.edu

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

CP11

Approximating Matching Size from Random Streams

We present a streaming algorithm that makes one pass over the edges of an unweighted graph presented in *random* order, and produces a polylogarithmic approximation to the size of the maximum matching in the graph, while using only polylogarithmic space. Prior to this work the only approximations known were a folklore $\tilde{O}(\sqrt{n})$ approximation with polylogarithmic space in an n vertex graph and a constant approximation with $\Omega(n)$ space. Our work thus gives the first algorithm where both the space and approximation factors are smaller than any polynomial in n .

Michael Kapralov
Stanford University
kapralov@mit.edu

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

Madhu Sudan
Microsoft Research New England
madhu@mit.edu

CP11

Improved Concentration Bounds for Count-Sketch

We present a refined analysis of the classic Count-Sketch streaming heavy hitters algorithm [CCF02]. Count-Sketch uses $O(k \log n)$ linear measurements to estimate a vector $x \in \mathbb{R}^n$. Our main result is that the average error is substantially smaller than the maximum error. We also give high probability results for the recovery of distributions with suitable decay, such as a power law or lognormal. We complement our results with simulations of Count-Sketch on a power law distribution. The empirical evidence indicates that our theoretical bounds give a precise characterization of the algorithm's performance: the asymptotics are correct and the associated constants are small. Our proof shows that any symmetric random variable with finite variance and positive Fourier transform concentrates at least as well as a Gaussian. This result, which may be of independent interest, gives good concentration even when the noise does not converge to a Gaussian.

Gregory Minton, Eric Price
MIT
gminton@math.mit.edu, ecprice@mit.edu

CP11

Annotations for Sparse Data Streams

We continue the study of annotated data streams. In this setting, a verifier (cloud user), lacking the resources to store and manipulate his massive input locally, accesses a powerful but untrusted prover (cloud service). The verifier must work within the restrictive data streaming paradigm. The prover, who can annotate the data stream, must both

supply the final answer and convince the verifier of its correctness. While optimal protocols are now known for several basic problems, such optimality holds only for streams whose length is commensurate with the size of the data universe. In contrast, many real-world data sets are relatively *sparse* (e.g. graphs that contain only $o(n^2)$ edges). We design the first protocols that allow both the annotation and the space usage to be sublinear in the number of stream updates rather than the size of the data universe. We solve significant problems, including variations of Frequency Moments and several natural problems on graphs.

Amit Chakrabarti
Dartmouth
ac@cs.dartmouth.edu

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

Navin Goyal
Microsoft Research
navingo@microsoft.com

Justin Thaler
Simons Institute for the Theory of Computing, UC
Berkeley
jthaler@fas.harvard.edu

CP11

An Optimal Lower Bound for Distinct Elements in the Message Passing Model

In the message-passing model of communication, there are k players each with their own private input, who try to compute or approximate a function of their inputs by sending messages to one another over private channels. We consider the setting in which each player holds a subset S_i of elements of a universe of size n , and their goal is to output a $(1 + \varepsilon)$ -approximation to the total number of distinct elements in the union of the sets S_i with constant probability, which can be amplified by independent repetition. This problem has applications in data mining, sensor networks, and network monitoring.

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

Qin Zhang
Indiana University Bloomington
qzhangcs@indiana.edu

CP12

Tight Lower Bounds for Greedy Routing in Higher-Dimensional Small-World Grids

We consider Kleinberg's small world graph model [STOC 2000]: a D -dimensional grid $\{0, \dots, n-1\}^D$ augmented by a constant number of additional unidirectional edges per node, chosen independently at random according to the same probability distribution. Kleinberg showed that if $D = 2$ and v is the long range contact of u with a probability proportional to $\|u - v\|_1^{-2}$, the "greedy" routing algorithm leads to an expected number of hops of $O((\log n)^2)$. We prove that greedy routing does not perform asymptotically better for any distribution in which the probability that node u has long range contact v is only a function of

$\|u - v\|_1$, as long as D is constant. For this, we introduce a novel proof technique, involving a so-called *budget game*, which is a simple process in one dimension. We prove a tight lower bound for such a budget game and then obtain a lower bound for greedy routing by a reduction.

Martin Dietzfelbinger
Technische Universität Ilmenau
martin.dietzfelbinger@tu-ilmenau.de

Philipp Woelfel
University of Calgary
woelfel@ucalgary.ca

CP12

Timing in Chemical Reaction Networks

Chemical reaction networks (CRNs) formally model chemistry in a well-mixed solution. CRNs are formally equivalent to a model of distributed computing known as population protocols, results transfer readily between the two models. We show that if a CRN respects **finite density** (at most $O(n)$ additional molecules can be produced from n initial molecules), then starting from any **dense** initial state (all molecular species initially present have count $\Omega(n)$, where n is the initial molecular count and volume), **every** producible species is produced in constant time with high probability. This implies that no such CRN can be a "timer", able to produce a molecule, but doing so only after a time that is an unbounded function of the input size. This has consequences regarding an open question of Angluin, Aspnes, and Eisenstat concerning the ability of population protocols to perform fast, reliable leader election and to simulate arbitrary algorithms from a uniform initial state.

David Doty
California Institute of Technology
pexatus@gmail.com

CP12

Tight Bounds for Rumor Spreading with Vertex Expansion

We establish a bound for the classic PUSH-PULL rumor spreading protocol on general graphs, in terms of the vertex expansion of the graph. We show that $O(\log^2(n)/\alpha)$ rounds suffice with high probability to spread a rumor from any single node to all n nodes, in any graph with vertex expansion at least α . This bound matches a known lower bound, and settles the natural question on the relationship between rumor spreading and vertex expansion asked by Chierichetti, Lattanzi, and Panconesi (SODA 2010). Further, some of the arguments used in the proof may be of independent interest, as they give new insights, for example, on how to choose a small set of nodes in which to plant the rumor initially, to guarantee fast rumor spreading.

George Giakkoupis
INRIA Rennes
ggiakk@gmail.com

CP12

Faster Agreement Via a Spectral Method for Detecting Malicious Behavior

We address the problem of Byzantine agreement, to bring processors to agreement on a bit in the presence of a it

strong adversary. This adversary has full information of the state of all processors, the ability to control message scheduling in an asynchronous model, and the ability to control the behavior of a constant fraction of processors which it may choose adaptively. We use a novel technique for detecting malicious behavior via spectral analysis. In particular, our algorithm uses coin flips from individual processors to repeatedly try to generate a fair global coin. The corrupted processors can bias this global coin by generating biased individual coin flips. However, we can detect which processors generate biased coin flips by analyzing the top right singular vector of a matrix containing the sums of coin flips generated by each processor. Entries in this singular vector with high absolute value correspond to processors that are trying to bias the global coin.

Valerie King
University of Victoria
val@uvic.ca

Jared Saia
Department of Computer Science
University of New Mexico
saia@cs.unm.edu

CP12

Intrinsic Universality in Tile Self-Assembly Requires Cooperation

We prove a negative result on the power of a model of algorithmic self-assembly for which finding general techniques and results has been notoriously difficult. Specifically, we prove that Winfree's abstract Tile Assembly Model is not intrinsically universal when restricted to use noncooperative tile binding. Noncooperative self-assembly, also known as "temperature 1", is where all tiles bind if they match on at least one side. Cooperative self-assembly, on the other hand, requires that some tiles bind on at least two sides. Our result shows that the change from non-cooperative to cooperative binding qualitatively improves the dynamics and behaviors found in these models of nanoscale self-assembly. In addition to the negative result, we exhibit a 3D noncooperative self-assembly tile set capable of simulating any 2D noncooperative self-assembly system. This implies that, in a restricted sense, non-cooperative self-assembly is intrinsically universal for itself.

Pierre-Etienne Meunier
Caltech
meunier@dna.caltech.edu

Matthew Patitz
University of Arkansas
patitz@uark.edu

Scott Summers
University of Wisconsin - Oshkosh
summerss@uwosh.edu

Guillaume Theyssier
LAMA, CNRS, Université de Savoie
guillaume.theyssier@univ-savoie.fr

Andrew Winslow
Tufts University
awinslow@cs.tufts.edu

Damien Woods
Caltech

woods@caltech.edu

CP13

Integer Quadratic Programming in the Plane

We show that the problem of minimizing a quadratic polynomial with integer coefficients over the integer points in a general two-dimensional rational polyhedron is solvable in time bounded by a polynomial in the input size.

Alberto Del Pia
ETH
adelpia@us.ibm.com

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

CP13

Improved Upper Bounds for Random-Edge and Random-Jump on Abstract Cubes

Upper bounds are given for the complexity of two very natural randomized algorithms for finding the *sink* of an *Acyclic Unique Sink Orientation* (AUSO) of the n -cube. For Random-Edge, we obtain an upper bound of about 1.80^n , improving upon the previous upper bound of about $2^n/n^{\log n}$ obtained by Gärtner and Kaibel. For Random-Jump, we obtain an upper bound of about $(3/2)^n$, improving upon the previous upper bound of about 1.72^n obtained by Mansour and Singh. AUSOs provide an appealing combinatorial abstraction of linear programming and other computational problems such as finding optimal strategies for turn-based *Stochastic Games*.

Thomas D. Hansen
Stanford University
tdh@cs.au.dk

Mike Paterson
University of Warwick
msp@dcs.warwick.ac.uk

Uri Zwick
Tel Aviv University
zwick@tau.ac.il

CP13

Dantzig's Pivoting Rule for Shortest Paths, Deterministic MDPs, and Minimum Cost to Time Ratio Cycles

Dantzig's pivoting rule is one of the most studied pivoting rules for the simplex algorithm. Orlin showed that $O(mn^2 \log n)$ Dantzig's pivoting steps suffice to solve *shortest paths* problems. Post and Ye recently showed that $O(m^2 n^3 \log^2 n)$ Dantzig's steps suffice to solve *deterministic MDPs* with a uniform discount factor. We improve both these bounds by a factor of n . Finally, we show how to reduce the problem of finding a *minimum cost to time ratio cycle* to the problem of finding an optimal policy for a discounted deterministic MDP with varying discount factors that tend to 1. This gives a strongly polynomial time algorithm for the problem that does not use Megiddo's parametric search technique.

Thomas D. Hansen
Stanford University

tdh@cs.au.dk

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

Uri Zwick
Tel Aviv University
zwick@tau.ac.il

CP13

Optimization Despite Chaos: Convex Relaxations to Complex Limit Sets Via Poincaré Recurrence

Decentralized systems, such as evolutionary dynamics in network extensions of zero-sum games, can give rise to complex behavior patterns. The challenge of any analytic investigation is to identify and characterize persistent properties despite the irregularities of such systems. We combine ideas from dynamical systems and game theory to produce topological characterizations of system trajectories. These trajectories are complex and do not necessarily converge to limit points or even limit cycles. We provide relaxed convex descriptions that can be computed in polynomial time. This allows us to compute extremal values of system features (e.g. expected utility of an agent) within these relaxations. Finally, we use information theoretic conservation laws along with Poincaré recurrence theory to argue about the optimality of our techniques.

Georgios Piliouras, Jeff S. Shamma
Georgia Institute of Technology
georgios@gatech.edu, shamma@gatech.edu

CP13

Polynomiality for Bin Packing with a Constant Number of Item Types

We consider the bin packing problem with d different item sizes s_i and item multiplicities a_i , where all numbers are given in binary encoding. This problem formulation is also known as the *1-dimensional cutting stock problem*. In this work, we provide an algorithm which, for constant d , solves bin packing in polynomial time. This was an open problem for all $d \geq 3$. In fact, for constant d our algorithm solves the following problem in polynomial time: given two d -dimensional polytopes P and Q , find the smallest number of integer points in P whose sum lies in Q .

Michel Goemans
Massachusetts Institute of Technology
goemans@math.mit.edu

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

CP14

Maximizing Social Influence in Nearly Optimal Time

Diffusion is a fundamental graph process, underpinning such phenomena as epidemic disease contagion and the spread of innovation by word-of-mouth. We address the algorithmic problem of finding a set of k initial seed nodes in a network so that the expected size of the resulting cascade is maximized, under the standard *independent cascade*

model of network diffusion. Runtime is a primary consideration for this problem due to the massive size of the relevant input networks. We provide a fast algorithm for the influence maximization problem, obtaining the near-optimal approximation factor of $(1 - \frac{1}{e} - \epsilon)$, for any $\epsilon > 0$, in time $O((m+n)\epsilon^{-3} \log n)$. Our algorithm is runtime-optimal (up to a logarithmic factor) and substantially improves upon the previously best-known algorithms which run in time $\Omega(mnk \cdot \text{POLY}(\epsilon^{-1}))$. Furthermore, our algorithm can be modified to allow early termination: if it is terminated after $O(\beta(m+n) \log n)$ steps for some $\beta < 1$ (which can depend on n), then it returns a solution with approximation factor $O(\beta)$. Finally, we show that this runtime is optimal (up to logarithmic factors) for any β and fixed seed size k .

Christian Borgs
Microsoft Research
borgs@microsoft.com

Michael Brautbar
MIT
brautbar@mit.edu

Jennifer Chayes
Microsoft Research
jchayes@microsoft.com

Brendan Lucier
Microsoft Research New England
brlucier@microsoft.com

CP14

Smoothed Analysis of Local Search for the Maximum-Cut Problem

Although the Maximum-Cut Problem is PLS-complete, local search heuristics for this problem perform well in practice. We study local search in the framework of smoothed analysis in which inputs are subject to a small amount of random noise. We show that the smoothed number of iterations is bounded by a polynomial in $n^{\log n}$ and ϕ where ϕ denotes the perturbation parameter. This shows that worst-case instances are fragile and unlikely to occur in practice.

Michael Etscheid, Heiko Röglin
Department of Computer Science
University of Bonn, Germany
etscheid@cs.uni-bonn.de, roeglin@cs.uni-bonn.de

CP14

A Constructive Algorithm for the Lovász Local Lemma on Permutations

While there has been significant progress on algorithmic aspects of the Lovász Local Lemma (LLL) in recent years, a noteworthy exception is when the LLL is used in the context of random permutations: the “lopsided” version of the LLL is usually at play here, and we do not yet have subexponential-time algorithms. We resolve this by developing a randomized polynomial-time algorithm for such applications. A noteworthy application is for Latin Transversals: the best-known general result here (Bissacot et al., improving on Erdős and Spencer), states that any $n \times n$ matrix in which each entry appears at most $(27/256)n$ times, has a Latin transversal. We present the first polynomial-time algorithm to construct such a transversal. Our ap-

proach also yields RNC algorithms: for Latin transversals, as well as the first efficient ones for the strong chromatic number and (special cases of) acyclic edge-coloring.

David Harris, Aravind Srinivasan
University of Maryland
davidgharris29@hotmail.com, srin@cs.umd.edu

CP14

Pipage Rounding, Pessimistic Estimators and Matrix Concentration

Pipage rounding is a dependent random sampling technique that has several interesting properties. One property that has been useful in applications is negative correlation of the resulting vector. There are some further properties that would be interesting to derive, but do not follow from negative correlation. In particular, recent concentration results for sums of independent random matrices are not known to extend to a negatively dependent setting. We introduce a technique called concavity of pessimistic estimators. This technique allows us to show concentration of submodular functions and concentration of matrix sums under pipage rounding. We provide an application to spectrally-thin trees, a spectral analog of the thin trees that played a crucial role in the recent breakthrough on ATSP. We show an algorithm that, given a graph where every edge has effective conductance at least κ , returns an $O(\kappa^{-1} \cdot \log n / \log \log n)$ -spectrally-thin tree.

Nicholas Harvey
University of Waterloo
nickhar@cs.ubc.ca

Neil Olver
VU Amsterdam
n.olver@vu.nl

CP14

Bilu-Linial Stable Instances of Max Cut and Minimum Multiway Cut

We investigate the notion of stability proposed by Bilu and Linial. We obtain an exact robust algorithm for γ -stable Max Cut instances with $\gamma \geq c\sqrt{\log n} \log \log n$ for some $c > 0$. We also prove that there is no robust polynomial-time algorithm for instances when γ is less than the best approximation factor for Sparsest Cut. That suggests that solving γ -stable instances with $\gamma = o(\sqrt{\log n})$ might be difficult or even impossible. Additionally, we present an algorithm for 4-stable instances of Minimum Multiway Cut. We also study a relaxed notion of weak stability.

Konstantin Makarychev
Microsoft Research
komakary@microsoft.com

Yury Makarychev
Toyota Technological Institute at Chicago
yury@ttic.edu

Aravindan Vijayaraghavan
CMU

v.aravindan@gmail.com

CP15

Cache-Adaptive Algorithms

We introduce the cache-adaptive model, which generalizes the external-memory model to apply to environments in which the amount of memory available to an algorithm changes over time. We prove that if an optimal cache-oblivious algorithm has a particular recursive structure, then it is also an optimal cache-adaptive algorithm. Cache-oblivious algorithms having this form include Floyd-Warshall all pairs shortest paths, naïve recursive matrix multiplication, matrix transpose, and Gaussian elimination. While the cache-oblivious sorting algorithm Lazy Funnel Sort does not have this recursive structure, we prove that it is nonetheless optimally cache-adaptive. We give paging algorithms for the case where the cache size changes dynamically. LRU with 4-memory and 4-speed augmentation is competitive with optimal. Moreover, Belady's algorithm remains optimal even when the cache size changes.

Roohbeh Ebrahimi
Stony Brook University
rebrahimi@cs.stonybrook.edu

CP15

Near-Optimal Labeling Schemes for Nearest Common Ancestors

We consider NCA labeling schemes: given a rooted tree T , label the nodes of T with binary strings such that, given the labels of any two nodes, one can determine, by looking only at the labels, the label of their nearest common ancestor. For trees with n nodes we present upper and lower bounds establishing that labels of size $(2 \pm \epsilon) \log_2 n$, $\epsilon < 1$ are both sufficient and necessary. Our lower bound separates NCA from ancestor labeling schemes.

Stephen Alstrup, Esbén B. Halvorsen
Department of Computer Science
University of Copenhagen
s.alstrup@diku.dk, esbenbh@diku.dk

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

CP15

Selection and Sorting in the 'Restore' Model

We consider the classical selection and sorting problems in a model where the initial permutation of the input has to be *restored* after completing the computation. While the requirement of the restoration is stringent compared to the classical versions of the problems, this model is more relaxed than a read-only memory where the input elements are not allowed to be moved within the input array. We show that for a sequence of n integers, selection can be done in $O(n)$ time using $O(\lg n)$ words. For sorting n integers in this model, we present an $O(n \lg n)$ -time algorithm using $O(\lg n)$ words of extra space, and we show how to match the time bound of any word-RAM integer-sorting algorithms using $O(n^\epsilon)$ words of extra space. Such time-space tradeoffs are provably not possible for general 'invisible' input elements.

Timothy M. Chan

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

J. Ian Munro
University of Waterloo
imunro@uwaterloo.ca

Venkatesh Raman
Institute of Mathematical Sciences, Chennai
vraman@imsc.res.in

CP15

Disjoint Set Union with Randomized Linking

A classic result in the analysis of data structures is that path compression with linking by rank solves the disjoint set union problem in almost-constant amortized time per operation. Recent experiments suggest that in practice, a naïve linking method works just as well if not better than linking by rank, in spite of being theoretically inferior. How can this be? We prove that *randomized* linking is asymptotically as efficient as linking by rank. This result provides theory that matches the experiments, which implicitly do randomized linking as a result of the way the input instances are generated.

Ashish Goel
Stanford University
ashishg@stanford.edu

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

Daniel H. Larkin
Princeton University
dhlarkin@cs.princeton.edu

Robert E. Tarjan
Princeton and Microsoft Research
prof.tarjan@gmail.com

CP15

Concurrent Range Reporting in Two-Dimensional Space

In the *concurrent range reporting* problem, the input is L disjoint sets S_1, \dots, S_L of points with a total of N points. The goal is to preprocess the sets into a structure such that, given a query range r and an arbitrary set $Q \subseteq \{1, \dots, L\}$, we can efficiently report all the points in $S_i \cap r$ for each $i \in Q$. We focus on the one- and two-dimensional cases of the problem. We prove that in the pointer-machine model (as well as comparison models such as the real RAM model), answering queries requires $\Omega(|Q| \log(L/|Q|) + \log N + K)$ time in the worst case, where K is the number of output points. In one dimension, we achieve this query time with a linear-space dynamic data structure that requires optimal $O(\log N)$ time to update. In the static case, for three-sided two-dimensional ranges, we get close to within an inverse Ackermann ($\alpha(\cdot)$) factor: we answer queries in $O(|Q| \log(L/|Q|)\alpha(L) + \log N + K)$ time.

Peyman Afshani
MADALGO, Aarhus University
peyman@madalgo.au.dk

Cheng Sheng, Yufei Tao
Chinese University of Hong Kong
csheng@cse.cuhk.edu.hk, taoyf@cse.cuhk.edu.hk

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

CP16

Better Approximation Algorithms for the Graph Diameter

The diameter of the graph is a fundamental graph parameter. In a seminal result Aingworth, Chekuri, Indyk and Motwani [SICOMP'99] designed an algorithm that computes in $\tilde{O}(n^2 + m\sqrt{n})$ time an estimate \tilde{D} for the diameter D , such that $\lfloor 2D/3 \rfloor - (M - 1) \leq \tilde{D} \leq D$, where M is the maximum edge weight in the graph. In recent work, Roditty and Vassilevska W. [STOC 13] gave an algorithm that has the same approximation guarantee but improves the runtime to $\tilde{O}(m\sqrt{n})$. For weighted graphs however, this approximation guarantee can be meaningless, as M can be arbitrarily large. In this paper we exhibit two algorithms that achieve a genuine $3/2$ -approximation for the diameter, one running in $\tilde{O}(m^{3/2})$ time, and one running in $\tilde{O}(mn^{2/3})$ time. In addition, we address the question of obtaining an additive c -approximation for the diameter. We present essentially tight lower bounds assuming the Strong Exponential Time Hypothesis.

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

Daniel H. Larkin
Princeton University
dhlarkin@cs.princeton.edu

Liam Roditty
Bar Ilan University
liam.roditty@gmail.com

Grant Schoenebeck
University of Michigan
schoenebeck@gmail.com

Robert E. Tarjan
Princeton and Microsoft Research
prof.tarjan@gmail.com

Virginia Vassilevska Williams
Stanford University
virgito@gmail.com

CP16

Cutting Corners Cheaply, Or How to Remove Steiner Points.

Our main result is that the Steiner Point Removal (SPR) problem can always be solved with polylogarithmic distortion, which resolves in the affirmative a question posed by Chan, Xia, Konjevod, and Richa (2006). Specifically, we prove that for every edge-weighted graph $G = (V, E, w)$ and a subset of terminals $T \subseteq V$, there is a graph $G' = (T, E', w')$ that is isomorphic to a minor of G , such that for every two terminals $u, v \in T$, the shortest-path distances between them in G and in G' satisfy $d_{G,w}(u, v) \leq$

$d_{G',w'}(u,v) \leq O(\log^6 |T|) \cdot d_{G,w}(u,v)$. Our proof features a new variant of metric decomposition.

Lior Kamma, Robert Krauthgamer
Weizmann Institute of Science
lior.kamma@weizmann.ac.il,
robert.krauthgamer@weizmann.ac.il

Huy Nguyen
Princeton University
hnguyen@princeton.edu

CP16

A Subquadratic-Time Algorithm for Decremental Single-Source Shortest Paths

We study dynamic $(1 + \epsilon)$ -approximation algorithms for the single-source shortest paths problem in an unweighted undirected n -node m -edge graph under edge deletions. The fastest algorithm for this problem is an algorithm with $O(n^{2+o(1)})$ total update time and constant query time by Bernstein and Roditty (SODA 2011). In this paper, we improve the total update time to $O(n^{1.8+o(1)} + m^{1+o(1)})$ while keeping the query time constant. This running time is essentially tight when $m = \Omega(n^{1.8})$ since we need $\Omega(m)$ time even in the static setting. For smaller values of m , the running time of our algorithm is *subquadratic*, and is the first that breaks through the quadratic time barrier.

Monika Henzinger, Sebastian Krinninger
University of Vienna
mhenzinger@gmail.com, sebastian.krinninger@univie.ac

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

CP16

Fault Tolerant Approximate BFS Structures

This paper addresses the problem of designing a *fault-tolerant* (α, β) approximate BFS structure, namely, a subgraph H of the network G satisfying $(s, v, H \setminus F) \leq \alpha \cdot (s, v, G \setminus F) + \beta$ for every $v \in V$. We first consider *multiplicative* $(\alpha, 0)$ FT-ABFS structures and present an algorithm that given an n -vertex unweighted undirected graph G and a source s constructs a $(3, 0)$ FT-ABFS structure rooted at s with at most $3n$ edges. We then consider *additive* $(1, \beta)$ FT-ABFS structures. In contrast to the linear size of $(\alpha, 0)$ FT-ABFS structures, we show that for every $\beta \in [1, O(\log n)]$ there exists an n -vertex graph G with a source s for which any $(1, \beta)$ FT-ABFS structure rooted at s has $\Omega(n^{1+\epsilon(\beta)})$ edges, for some function $\epsilon(\beta) \in (0, 1)$. Our lower bounds are complemented by an upper bound.

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

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

CP16

Beyond Locality-Sensitive Hashing

We present a new data structure for the c -approximate

near neighbor problem (ANN) in the Euclidean space. For n points in \mathbb{R}^d , our algorithm achieves $O(n^\rho + d \log n)$ query time and $O(n^{1+\rho} + d \log n)$ space, where $\rho \leq 7/(8c^2) + O(1/c^3) + o(1)$. This is the first improvement over the result by Andoni and Indyk (FOCS 2006) and the first data structure that bypasses a locality-sensitive hashing lower bound proved by O'Donnell, Wu and Zhou (ICS 2011). By a standard reduction we obtain a data structure for the Hamming space and ℓ_1 norm with $\rho \leq 7/(8c) + O(1/c^{3/2}) + o(1)$, which is the first improvement over the result of Indyk and Motwani (STOC 1998).

Alexandr Andoni
Microsoft Research
andoni@microsoft.com

Piotr Indyk
Massachusetts Institute of Technology
indyk@mit.edu

Huy Nguyen
Princeton University
hnguyen@princeton.edu

Ilya Razenshteyn
Lomonosov Moscow State University
ilyaraz@mit.edu

CP17

Ranking on Arbitrary Graphs: Rematch via Continuous LP with Monotone and Boundary Condition Constraints

We revisit the well-known oblivious matching problem on general graphs, in which an adversary fixes a graph whose edges are oblivious to the algorithm. We analyze the famous Ranking algorithm, and form an LP with monotone and boundary constraints. We solve its continuous LP relaxation to give the currently best performance ratio 0.523.

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

Xiaowei Wu
the University of Hong Kong
xwwu@cs.hku.hk

Zhichao Zhao
The University of Hong Kong
zczhao@cs.hku.hk

CP17

Primal Dual Gives Almost Optimal Energy Efficient Online Algorithms

We consider the problem of online scheduling of jobs on unrelated machines with dynamic speed scaling to minimize the sum of energy and weighted flow time. We give an algorithm with an almost optimal competitive ratio for arbitrary power functions. For power functions of the form $f(s) = s^\alpha$ for some constant $\alpha > 1$, we get a competitive ratio of $O(\frac{\alpha}{\log \alpha})$, improving upon a previous competitive ratio of $O(\alpha^2)$ by Anand et al., along with a matching lower bound of $\Omega(\frac{\alpha}{\log \alpha})$. Further, in the resource augmentation model, with a $1 + \epsilon$ speed up, we give a $2(\frac{1}{\epsilon} + 1)$ competitive algorithm, improving the bound of $1 + O(\frac{1}{\epsilon^2})$ by Gupta

et al. Unlike the previous results most of which used an amortized local competitiveness argument or dual fitting methods, we use a primal-dual method, which is useful not only to analyze the algorithms but also to design the algorithm itself.

Nikhil R. Devanur
Microsoft Research, Redmond
nikdev@microsoft.com

Zhiyi Huang
Stanford University
hzhyyi@stanford.edu

CP17

Hallucination Helps: Energy Efficient Virtual Circuit Routing

We consider virtual circuit routing protocols, with an objective of minimizing energy, in a network of components that are speed scalable, and that may be shutdown when idle. We give a polynomial-time offline algorithm for multicommodity routing, that has approximation ratio $O(\log^\alpha k)$, where k is the number of demand pairs. The key step of the algorithm design is a random sampling technique that we call *hallucination*. The algorithm extends rather naturally to an online algorithm. The analysis of the online algorithm introduces a natural ‘priority’ multicommodity flow problem, and bounds the priority multicommodity flow-cut gap. We also explain how our hallucination technique can be used to achieve an $(O(\log km), O(\log km))$ bicriteria approximation result for the problem of buying a minimum cost collection of unit-capacitated edges to support a concurrent multicommodity flow, where m is the number of links in the network.

Antonios Antoniadis
Humboldt-Universität zu Berlin
antoniosantoniadis@gmail.com

Sungjin Im
Duke University
sungjin@cs.duke.edu

Ravishankar Krishnaswamy
Princeton University
ravishankar.k@gmail.com

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

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

Kirk Pruhs
University of Pittsburgh
krp2@pitt.edu

Cliff Stein
Columbia University
cliff@ieor.columbia.edu

CP17

Improvements and Generalizations of Stochastic Knapsack and Multi-Armed Bandit Approximation

Algorithms

In 2011, A. Gupta et al. obtained an LP-based $\frac{1}{48}$ -approximation algorithm for the multi-armed bandit (MAB) problem without the martingale assumption. We improve the algorithm to a $\frac{4}{27}$ -approximation, with simpler analysis. Our algorithm also generalizes to the case of MAB superprocesses with (stochastic) multi-period actions, yielding new results in the budgeted learning framework of Guha and Munagala. Also, we obtain a $(\frac{1}{2} - \epsilon)$ -approximation for the variant of MAB where preemption (playing an arm, switching to another arm, then coming back to the first arm) is not allowed. This contains the stochastic knapsack problem of Dean, Goemans, and Vondrák with correlated rewards, improving the $\frac{1}{16}$ and $\frac{1}{8}$ approximations of A. Gupta et al. for the cases with and without cancellation, respectively, providing to our knowledge the first tight algorithm for these problems that matches the integrality gap of 2.

Will Ma
Massachusetts Institute of Technology
willma353@gmail.com

CP17

New Approximations for Reordering Buffer Management

In this paper we consider the buffer reordering management problem. In this model there are n elements that arrive over time with different colors. There is a buffer that can store up to k elements and when the buffer becomes full an element must be output. If an element is output that has a color different from the previous element, a cost depending on the color must be paid. This cost could be uniform or non-uniform over colors; these are called unweighted and weighted cases, respectively. The goal is to reorder elements within the buffer before outputting them to minimize the total cost incurred. Our main result is a randomized $O(\log \log k\gamma)$ -approximation for the weighted case. Here γ is the ratio of the maximum to minimum weight. We also revisit the unweighted case and give an improved randomized -approximation which improves (modestly) upon the approximation guarantee given previously.

Sungjin Im
University of California, Merced
sim3@ucmerced.edu

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

CP18

Testing Equivalence Between Distributions Using Conditional Samples

We study a recently introduced framework for property testing of probability distributions, by considering distribution testing algorithms that have access to a *conditional sampling oracle*. This is an oracle that takes as input a subset S of the domain $[N]$ of the unknown probability distribution D and returns a draw from the conditional distribution D restricted to S . This model allows considerable flexibility in the design of distribution testing algorithms; in particular, testers in this model can be adaptive. In this paper we focus on algorithms for two fundamental distribution testing problems: testing whether $D = D^*$

for an explicitly provided D^* , and testing whether two unknown D_1 and D_2 are equivalent. While these problems have $\Omega(\sqrt{N})$ sample complexity in the standard model we give $\text{poly}(\log N, 1/\epsilon)$ -query algorithms (and in some cases $\text{poly}(1/\epsilon)$ independent of N) for both in our conditional sampling setting.

Clément Canonne
Columbia University
ccanonne@cs.columbia.edu

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

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

CP18

A Cubic Algorithm for Computing Gaussian Volume

We present randomized algorithms for sampling the standard Gaussian distribution restricted to a convex set and for estimating the Gaussian measure of a convex set, in the general membership oracle model. The complexity of the integration algorithm is $O^*(n^3)$ while the complexity of the sampling algorithm is $O^*(n^3)$ for the first sample and $O^*(n^2)$ for every subsequent sample. These bounds improve on the corresponding state-of-the-art by a factor of n . Our improvement comes from several aspects: better isoperimetry, smoother annealing, avoiding transformation to isotropic position and the use of the “speedy walk” in the analysis.

Ben Cousins, Santosh Vempala
Georgia Institute of Technology
bcousins3@gatech.edu, vempala@gatech.edu

CP18

Optimal Algorithms for Testing Closeness of Discrete Distributions

We study the question of closeness testing for two discrete distributions. More precisely, given samples from two distributions p and q over an n -element set, we wish to distinguish whether $p = q$ versus p is at least ϵ -far from q , in either ℓ_1 or ℓ_2 distance. Batu et al [?, ?] gave the first sub-linear time algorithms for these problems, which matched the lower bounds of [?] up to a logarithmic factor in n , and a polynomial factor of ϵ . In this work, we present simple (and new) testers for both the ℓ_1 and ℓ_2 settings, with sample complexity that is information-theoretically optimal, to constant factors, both in the dependence on n , and the dependence on ϵ ; for the ℓ_1 testing problem we establish that the sample complexity is $\Theta(\max\{n^{2/3}/\epsilon^{4/3}, n^{1/2}/\epsilon^2\})$.

Siu On Chan
Microsoft Research New England
siuon@cs.berkeley.edu

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

Gregory Valiant
Stanford

valiant@stanford.edu

Paul Valiant
UC Berkeley
pvaliant@gmail.com

CP18

Hereditary Properties of Permutations Are Strongly Testable

We show that for every hereditary permutation property P and every $\epsilon > 0$, there exists an integer M such that if a permutation π is ϵ -far from P in the Kendall’s tau distance, then a random subpermutation of π of order M has the property P with probability at most ϵ . This settles an open problem whether hereditary permutation properties are strongly testable, i.e., testable with respect to the Kendall’s tau distance, which is considered to be the edit distance for permutations.

Tereza Klimosova, Daniel Kral
University of Warwick
t.klimosova@warwick.ac.uk, D.Kral@warwick.ac.uk

CP18

Testing Surface Area

We give an $O(1/\epsilon)$ -query property testing algorithm which distinguishes whether an unknown set has surface area at most A or (is ϵ -far from) surface area at least $(4/\pi)A$. Our result works under n -dimensional Lebesgue measure or n -dimensional Gaussian measure. Previous work only treated the 1-dimensional case. We improve previous results for the 1-dimensional case and give the first results for higher n and evades the “curse of dimensionality”.

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

Amir Nayyeri
Department of Computer Science
Carnegie Mellon University
amirn@cs.cmu.edu

Ryan O’Donnell
Carnegie Mellon University, USA
odonnell@cs.cmu.edu

Chenggang Wu
IIIS, Tsinghua University
wuchenggang0316@gmail.com

CP19

Approximation Algorithm for Sparsest K-Partitioning

Given a graph G , the sparsest-cut problem asks to find the set of vertices S which has the least expansion defined as

$$\phi_G(S) := \frac{w(E(S, \bar{S}))}{\min\{w(S), w(\bar{S})\}},$$

where w is the total edge weight of a subset. Here we study the natural generalization of this problem: given an integer k , compute a k -partition $\{P_1, \dots, P_k\}$ of the vertex set so

as to minimize

$$\phi_k(\{P_1, \dots, P_k\}) := \max_i \phi_G(P_i).$$

Our main result is a polynomial time bi-criteria approximation algorithm which outputs a $(1 - \varepsilon)k$ -partition of the vertex set such that each piece has expansion at most $O_\varepsilon(\sqrt{\log n \log k})$ times OPT . We also study balanced versions of this problem.

Anand Louis
Georgia Tech
anandl@gatech.edu

Konstantin Makarychev
Microsoft Research
komakary@microsoft.com

CP19

Flow-Based Algorithms for Local Graph Clustering

Given a subset A of vertices of an undirected graph G , the cut-improvement problem asks us to find a subset S similar to A but with smaller conductance. An algorithm for this problem was given by Andersen and Lang in 2008 and requires maximum flow computations over the whole graph G . In this paper, we introduce LocalImprove, the first cut-improvement algorithm that is LOCAL, i.e., that runs in time dependent on the size of the input set A rather than that of G . LocalImprove closely matches the same theoretical guarantee as Andersen-Lang. The main application of LocalImprove is to design better local-graph-partitioning algorithms. Zhu, Lattanzi and Mirrokni showed in 2013 that random-walk based local algorithms can obtain an improved-Cheeger's like conductance performance, while we show in the regime when their improved Cheeger's outperforms Cheeger's, we can use LocalImprove to obtain an $O(1)$ approximation. This yields the first flow-based local algorithm for graph partitioning.

Lorenzo Orecchia, Zeyuan Allen Zhu
MIT
orecchia@mit.edu, zeyuan@csail.mit.edu

CP19

Partitioning into Expanders

There is a basic fact in algebraic graph theory that for any graph $G = (V, E)$, $\lambda_k > 0$ if and only if G has at most $k - 1$ connected components. We prove a robust version of this fact. For some $1 \leq l \leq k - 1$, V can be partitioned into l sets P_1, \dots, P_l such that each P_i is a low-conductance set in G and induces a high conductance induced subgraph.

Shayan Oveis Gharan
Stanford University
oveisgharan@berkeley.edu

Luca Trevisan
Stanford
n/a

CP19

Non-Uniform Graph Partitioning

We consider the problem of Non-Uniform Graph Partitioning, where the input is an edge-weighted undirected graph

$G = (V, E)$ and k capacities n_1, \dots, n_k , and the goal is to find a partition $\{S_1, S_2, \dots, S_k\}$ of V satisfying $|S_j| \leq n_j$ for all $1 \leq j \leq k$, that minimizes the total weight of edges crossing between different parts. We present a bicriteria approximation algorithm for Non-Uniform Graph Partitioning that approximates the objective within $O(\log n)$ factor while deviating from the required capacities by at most a constant factor. Our approach is to apply stopping-time based concentration results to a simple randomized rounding of a configuration LP. These concentration bounds are needed as the commonly used techniques of bounded differences and bounded conditioned variances do not suffice.

Robert Krauthgamer
Weizmann Institute of Science
robert.krauthgamer@weizmann.ac.il

Seffi Naor
Technion - Israel Institute of Technology
naor@cs.technion.ac.il

Roy Schwartz
Microsoft Research, Redmond
roysch@microsoft.com

Kunal Talwar
Microsoft Research Silicon Valley
kunal@microsoft.com

CP19

Streaming Balanced Graph Partitioning Algorithms for Random Graphs

With recent advances in storage technology, it is now possible to store the vast amounts of data generated by cloud computing applications. The sheer size of big data motivates the need for streaming algorithms that can compute approximate solutions without full random access to all of the data. In this paper, we consider the problem of loading a graph onto a distributed cluster with the goal of optimizing later computation. We model this as computing an approximately balanced k -partitioning of a graph in a streaming fashion with only one pass over the data. We give lower bounds on this problem, showing that no algorithm can obtain an $o(n)$ approximation with a random or adversarial stream ordering.

Isabelle Stanton
Google Inc
stanton@google.com

CP20

Constrained Signaling in Auction Design

We consider the problem of an auctioneer who faces the task of selling a good to a set of buyers, when the auctioneer does not have the capacity to describe to the buyers the exact identity of the good that he is selling. Instead, he must come up with a constrained signalling scheme: a mapping from goods to signals, that satisfies the constraints of his setting. For example, the auctioneer may be able to communicate only a bounded length message for each good, or he might be legally constrained in how he can advertise the item being sold. Each candidate signaling scheme induces an incomplete-information game among the buyers, and the goal of the auctioneer is to choose the signaling scheme and accompanying auction format that optimizes welfare. In this paper, we use techniques from submodular

function maximization and no-regret learning to give algorithms for computing constrained signaling schemes for a variety of constrained signaling problems.

Shaddin Dughmi

University of Southern California
shaddin@usc.edu

Nicole Immorlica

Northwestern U
nicimm@gmail.com

Aaron Roth

University of Pennsylvania, USA
n/a

CP20

On Computability of Equilibria in Markets with Production

Although production is an integral part of the Arrow-Debreu market model, and of most economies, almost all of the work in theoretical computer science has so far concentrated on markets without production. Our paper takes a significant step towards understanding computation of equilibria in markets with production. For markets with separable, piecewise-linear concave (SPLC) utilities and SPLC production, we obtain an LCP formulation that captures exactly the set of equilibria and we give a complementary pivot algorithm for finding one. This settles a question asked by Eaves in 1975. Since this is a path-following algorithm, we obtain a proof of membership of this problem in PPA, using Todd, 1976. We further give a proof of PPA-hardness for this problem. Experiments show that our algorithm is practical and does not suffer from issues of numerical instability. Also, it is strongly polynomial when the number of goods or the number of agents and firms is constant.

Jugal Garg

Georgia Tech
jugal.garg@gmail.com

Vijay V. Vazirani

Georgia Institute of Technology
vazirani@cc.gatech.edu

CP20

The Complexity of Optimal Multidimensional Pricing

We resolve the complexity of revenue-optimal deterministic auctions in the unit-demand single-buyer Bayesian setting, i.e., the optimal item pricing problem, when the buyer's values for the items are independent. We show that the problem of computing a revenue-optimal pricing can be solved in polynomial time for distributions of support size 2 and its decision version is NP-complete for distributions of support size 3 as well for identical distributions.

Xi Chen

Columbia University
xichen@cs.columbia.edu

Ilias Diakonikolas

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

Dimitris Paparas, Xiaorui Sun, Mihalis Yannakakis

Columbia University

paparas@cs.columbia.edu, xiaoruisun@cs.columbia.edu,
mihalis@cs.columbia.edu

CP20

The Complexity of Optimal Mechanism Design

Myerson's seminal work provides a computationally efficient revenue-optimal auction for selling one item to multiple bidders. Generalizing this work to selling multiple items at once has been a central question in economics and algorithmic game theory, but its complexity has remained poorly understood. We answer this question by showing that a revenue-optimal auction in multi-item settings cannot be found and implemented computationally efficiently, unless $ZPP \supseteq P^{\#P}$. This is true even for a single additive bidder whose values for the items are independently distributed on two rational numbers with rational probabilities. Our result is very general: we show that it is hard to compute any encoding of an optimal auction of any format that can be implemented in expected polynomial time. In particular, under well-believed complexity-theoretic assumptions, revenue-optimization in very simple multi-item settings can only be tractably approximated.

Costis Daskalakis, Alan Deckelbaum, Christos Tzamos

Massachusetts Institute of Technology
costis@csail.mit.edu, deckel@mit.edu, tzamos@mit.edu

CP20

Prophet Inequalities with Limited Information

In the classical prophet inequality, a gambler observes a sequence of stochastic rewards and must decide, for each reward, whether to keep it and stop the game or to forfeit the reward forever and reveal the next value. The gambler's goal is to obtain a constant fraction of the expected reward that the optimal offline algorithm would get. Recently, prophet inequalities have been generalized to settings where the gambler can choose k elements, and, more generally, where he can choose any independent set in a matroid. However, all the existing algorithms require the gambler to know the distribution from which the rewards are drawn. We construct the first single-sample prophet inequalities for many settings of interest, whose guarantees all match the best possible asymptotically even with full knowledge of the distribution. In addition, we apply these results to design the first posted-price and multi-dimensional auction mechanisms with limited information in certain settings.

Pablo Azar

MIT
pazar@mit.edu

Robert Kleinberg

Cornell University
rdk@cs.cornell.edu

Matt Weinberg

MIT
smweinberg@csail.mit.edu

CP21

Fast Computation of Output-Sensitive Maxima in

a Word Ram

In this paper, we study the problem of computing the maxima of a set of points in three dimensions and show that in a word RAM and given a set of n points with integer coordinates, the maxima can be found in $O\left(n \log \log_{n/h} n\right)$ deterministic time in which h is the output size. For $h = n^{1-\alpha}$ this is $O(n \log(1/\alpha))$. This improves the previous $O(n \log \log h)$ time algorithm and can be considered surprising since it gives a linear time algorithm when $\alpha > 0$ is a constant, which is faster than the current best deterministic and randomized integer sorting algorithms. We also observe that improving this running time at least requires faster integer sorting algorithms. Additionally, we show that the same deterministic running time could be achieved for performing n point location queries in an arrangement of size h . Finally, our maxima result can be extended to higher dimensions by paying a $\log_{n/h} n$ factor penalty per dimension.

Peyman Afshani
MADALGO, Aarhus University
peyman@madalgo.au.dk

CP21

Making Octants Colorful and Related Covering Decomposition Problems

We give new positive results on the long-standing open problem of geometric covering decomposition for homothetic polygons. In particular, we prove that for any positive integer k , every finite set of points in \mathbb{R}^3 can be colored with k colors so that every translate of the negative octant containing at least k^6 points contains at least one of each color. The best previously known bound was doubly exponential in k . This yields, among other corollaries, the first polynomial bound for the decomposability of multiple coverings by homothetic triangles. We also prove that no algorithm can dynamically maintain a decomposition of a multiple covering by intervals under insertion of new intervals, even in a semi-online model.

Jean Cardinal
Université libre de Bruxelles (ULB)
jcardin@ulb.ac.be

Kolja Knauer
Université Montpellier 2
kolja.knauer@math.univ-montp2.fr

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

Torsten Ueckerdt
Karlsruhe Institute of Technology
torsten.ueckerdt@kit.edu

CP21

A Size-Sensitive Discrepancy Bound for Set Systems of Bounded Primal Shatter Dimension

Let (X, ξ) be a set system on an n -point set X . We consider the scenario where, for any subset $X' \subseteq X$ of size $m \leq n$ and for any parameter $1 \leq k \leq m$, the number of restrictions of the sets of ξ to X' of size at most k is only $O(m^{d_1} k^{d-d_1})$, for fixed integers $d > 0$ and $1 \leq d_1 \leq d$. In this case we show that (X, ξ) admits a discrepancy bound of

$O^*(|S|^{1/2-d_1/(2d)} n^{(d_1-1)/(2d)})$, for each $S \in \xi$, where $O^*(\cdot)$ hides a polylogarithmic factor in n . This bound is tight up to a polylogarithmic factor and the corresponding coloring can be computed in expected polynomial time using the very recent machinery of Lovett and Meka for constructive discrepancy minimization. Our bound improves and generalizes the bounds obtained from the machinery of Har-Peled and Sharir (and the follow-up work of Sharir and Zaban) for points and halfspaces in d -space for $d \geq 3$.

Esther Ezra
Courant Institute of Mathematical Sciences
New York University
esther@cims.nyu.edu

CP21

Four Soviets Walk the Dog—with An Application to Alt's Conjecture

Given two polygonal curves in the plane, there are many ways to define a notion of similarity between them. One measure that is extremely popular is the Fréchet distance. Since it has been proposed by Alt and Godau in 1992, many variants and extensions have been studied. Nonetheless, even more than 20 years later, the original $O(n^2 \log n)$ algorithm by Alt and Godau for computing the Fréchet distance remains the state of the art (here n denotes the number of vertices on each curve). This has led Helmut Alt to conjecture that the associated decision problem is 3SUM-hard. Building on recent work by Agarwal et al, we give a randomized algorithm to compute the Fréchet distance between two polygonal curves in time $O(n^2 \sqrt{\log n} (\log \log n)^{3/2})$ on a pointer machine and in time $O(n^2 (\log \log n)^2)$ on a word RAM. Furthermore, we provide evidence that the decision problem may not be 3SUM-hard after all.

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

Maike Buchin
Ruhr Universitaet Bochum
maike.buchin@ruhr-uni-bochum.de

Wouter Meulemans
TU Eindhoven
w.meulemans@tue.nl

Wolfgang J. Mulzer
Freie Universitaet Berlin
mulzer@inf.fu-berlin.de

CP21

Optimal Deterministic Shallow Cuttings for 3D Dominance Ranges

Many problems in range searching admit efficient static data structures by the applying shallow cuttings. We present efficient deterministic algorithms that construct optimal size (single and multiple) shallow cuttings for orthogonal dominance ranges on a set of n 3d points. A single shallow cutting takes $O(n \log n)$ worst case time and $O(n)$ space. We construct a logarithmic number of shallow cuttings of the input simultaneously, in the same complexity. This is optimal in the comparison and the algebraic comparison models, and an important step forward, since only polynomial guarantees were previously achieved for

the deterministic construction of shallow cuttings, even in three dimensions. Our methods yield the first worst case efficient preprocessing algorithms for a series of important orthogonal range searching problems in the pointer machine and the word-RAM models, where such shallow cuttings are utilised to support the queries efficiently.

Peyman Afshani
MADALGO, Aarhus University
peyman@madalgo.au.dk

Konstantinos A. Tsakalidis
Aarhus University
tsakalid@cse.ust.hk

CP22

Fast Algorithms for Maximizing Submodular Functions

In this paper we develop algorithms that match the best known approximation guarantees, but with significantly improved running times, for maximizing a monotone submodular function subject to various constraints. Our first result is a simple algorithm that gives a $(1 - 1/e - \epsilon)$ -approximation for a cardinality constraint using $O(\frac{n}{\epsilon} \log \frac{n}{\epsilon})$ queries, and a $1/(p + 2\ell + 1 + \epsilon)$ -approximation for the intersection of a p -system and ℓ knapsack constraints using $O(\frac{n}{\epsilon^2} \log^2 \frac{n}{\epsilon})$ queries. The main idea behind these algorithms serves as a building block in our more sophisticated algorithms. Our main result is a new variant of the continuous greedy algorithm, which interpolates between the classical greedy algorithm and a truly continuous algorithm. We show how this algorithm can be implemented for matroid and knapsack constraints using $\tilde{O}(n^2)$ oracle calls to the objective function.

Ashwinkumar Badanidiyuru
Cornell University
ashwin85@cs.cornell.edu

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

CP22

Submodular Maximization with Cardinality Constraints

We consider the problem of maximizing a (non-monotone) submodular function subject to a cardinality constraint. When at most k elements can be chosen, we improve the current best $1/e - o(1)$ approximation to a factor that is in the range $[1/e + 0.004, 1/2]$. When exactly k elements must be chosen, we improve the current best $1/4 - o(1)$ approximation to a factor that is in the range $[0.356, 1/2]$. Additionally, some of our algorithms have a time complexity of only $O(nk)$.

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

Moran Feldman
EPFL
moran.feldman@epfl.ch

Joseph Naor

Technion
naor@cs.technion.ac.il

Roy Schwartz
Microsoft Research, Redmond
roysch@microsoft.com

CP22

Approximation Algorithms for Stochastic Boolean Function Evaluation and Stochastic Submodular Set Cover

We present approximation algorithms for two problems: Stochastic Boolean Function Evaluation (SBFE) and Stochastic Submodular Set Cover (SSSC). Our results for SBFE problems are obtained by reducing them to SSSC problems through the construction of appropriate utility functions. We give a new algorithm for the SSSC problem that we call Adaptive Dual Greedy. We use this algorithm to obtain a 3-approximation algorithm solving the SBFE problem for linear threshold functions. We also get a 3-approximation algorithm for the closely related Stochastic Min-Knapsack problem, and a 2-approximation for a natural special case of that problem. In addition, we prove a new approximation bound for a previous algorithm for the SSSC problem, Adaptive Greedy. We consider an approach to SBFE using existing techniques, which we call the Q -value approach. We present some uses of the Q -approach but prove its limitations as well.

Amol Deshpande
University of Maryland
amol@cs.umd.edu

Lisa Hellerstein, Devorah Kletenik
Polytechnic Institute of NYU
hstein@poly.edu, dkletenik@gmail.com

CP22

Influence Maximization in Undirected Networks

We consider the problem of finding k vertices of maximum total influence in an undirected network, under the independent cascade model of influence. It is known to be NP-hard to achieve an approximation factor better than $(11/e)$ in directed networks. We show that this barrier can be overcome in undirected networks: the greedy algorithm obtains a $(1 - 1/e + c)$ approximation to the set of optimal influence, for some constant $c > 0$.

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

Brendan Lucier
Microsoft Research New England
brlucier@microsoft.com

CP22

Maximizing Bisubmodular and k -Submodular Functions

We consider the problem of maximizing bisubmodular and, more generally, k -submodular functions in the value oracle model. We provide the first approximation guarantees for maximizing a general bisubmodular or k -submodular function. We give an analysis of the naive random algorithm as well as a randomized greedy algorithm inspired

by the recent randomized greedy algorithm of Buchbinder et al. [FOCS'12] for unconstrained submodular maximization. We show that this algorithm approximates any k -submodular function to a factor of $1/(1 + \sqrt{k/2})$. In the case of bisubmodular functions, we achieve an approximation guarantee $1/2$ and show that it is the best possible.

Justin Ward

University of Warwick
J.D.Ward@warwick.ac.uk

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

CP23

Model-Based Sketching and Recovery with Expanders.

Linear sketching and recovery of sparse vectors with randomly constructed sparse matrices has numerous applications in several areas, including compressive sensing, data stream computing, graph sketching, and combinatorial group testing. This paper considers the same problem with the added twist that the sparse coefficients of the unknown vector exhibit further correlations as determined by a known sparsity model. We prove that exploiting model-based sparsity in sketching provably reduces the sketch size without sacrificing recovery quality. In this context, we present the first polynomial time algorithm for recovering model sparse signals from linear sketches obtained via sparse adjacency matrices of expander graphs with rigorous performance guarantees. The computational cost of our algorithm depends on the difficulty of projecting onto the model-sparse set. For the sparsity models we describe in this paper, such projections can be obtained in linear time. Finally, we provide numerical experiments to illustrate the theoretical results in action.

Bubacarr Bah, Luca Baldassarre
École Polytechnique Fédérale de Lausanne
bubacarr.bah@epfl.ch, luca.baldassarre@epfl.ch

Volkan Cevher
EPFL
volkan.cevher@epfl.ch

CP23

Bicriteria Data Compression

We introduce the Bicriteria LZ77-Parsing problem. The goal is to determine an LZ77 parsing which minimizes the compressed space occupancy, provided that the decompression time is bounded by T (or vice-versa). We solve this problem efficiently, up to a negligible additive constant, in $O(n \log^2 n)$ time and optimal $O(n)$ words of working space. Experiments show that our compressor is very competitive w.r.t. highly engineered competitors such as Snappy, lzma, bzip2.

Andrea Farruggia, Paolo Ferragina, Antonio Frangioni
Dipartimento di Informatica
University of Pisa
farruggi@di.unipi.it, ferragina@di.unipi.it,
frangio@di.unipi.it

Rossano Venturini
Dept. of Computer Science, University of Pisa

rossano@di.unipi.it

CP23

On Sketching Matrix Norms and the Top Singular Vector

In this paper, we study the problem of sketching matrix norms. We consider two sketching models. The first is bilinear sketching, in which there is a distribution over pairs of $r \times n$ matrices S and $n \times s$ matrices T such that for any fixed $n \times n$ matrix A , from $S \cdot A \cdot T$ one can approximate $\|A\|_p$ up to an approximation factor $\alpha \geq 1$ with constant probability, where $\|A\|_p$ is a matrix norm. The second is general linear sketching, in which there is a distribution over linear maps $L : \mathbf{R}^{n^2} \rightarrow \mathbf{R}^k$, such that for any fixed $n \times n$ matrix A , from $L(A)$ one can approximate $\|A\|_p$ up to a factor α . We study some of the most frequently occurring Schatten p -norms for $p \in \{0, 1, 2, \infty\}$, which correspond to the rank of A , the trace, Frobenius, and the operator norm, respectively.

Yi Li

University of Michigan
leeyi@umich.edu

Huy Nguyen
Princeton University
hnguyen@princeton.edu

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

CP23

Approximation-Tolerant Model-Based Compressive Sensing

Results in compressive sensing show that one can recover a k -sparse signal from $m = O(k \log(n/k))$ measurements. The framework of model-based compressive sensing reduces the number of measurements even further to $O(k)$ by limiting the supports of x to a subset of all possible supports. Unfortunately, for the framework to apply, one needs an algorithm that given a signal x , solves the following optimization problem exactly: $\arg \min_{\Omega \in \mathcal{M}} \|x_{[n] \setminus \Omega}\|_2$. This requirement poses an obstacle for extending the framework to a richer class of models. In this paper, we remove this obstacle and show how to extend the model-based compressive sensing framework so that it requires only *approximate* solutions to the aforementioned optimization problem. Further, we apply our framework to the *Constrained Earth Mover's Distance (CEMD)* model, obtaining a sparse recovery scheme that uses significantly less than $O(k \log(n/k))$ measurements.

Chinmay Hegde
MIT
chinmay@csail.mit.edu

Piotr Indyk
Massachusetts Institute of Technology
indyk@mit.edu

Ludwig Schmidt
MIT

ludwigs@mit.edu

CP23

New Constructions of RIP Matrices with Fast Multiplication and Fewer Rows

In this paper, we present novel constructions of matrices with the restricted isometry property (RIP) that support fast matrix-vector multiplication. Our guarantees are the best known, and can also be used to obtain the best known guarantees for fast Johnson-Lindenstrauss transforms.

Jelani Nelson
Harvard
minilek@seas.harvard.edu

Eric Price
MIT
ecprice@mit.edu

Mary Wootters
University of Michigan
wootters@umich.edu

CP24

Point Line Cover: The Easy Kernel Is Essentially Tight

In the NP-hard Point Line Cover problem the task is to cover a given set of n points in the plane with at most k lines. There is a straightforward P-time reduction to k^2 points. Settling an open problem of Lokshtanov we prove that this cannot be significantly improved unless the polynomial hierarchy collapses. A key feature is the use of an active oracle protocol to efficiently transfer an instance with n points.

Stefan Kratsch
Utrecht University
stefan.kratsch@tu-berlin.de

Geevarghese Philip
Max-Planck-Institute for Informatics,
Saarbruecken, Germany
gphilip@mpi-inf.mpg.de

Saurabh Ray
Ben-Gurion University, Israel
saurabh.ray@gmail.com

CP24

Parameters of Two-Prover-One-Round Game and The Hardness of Connectivity Problems

Optimizing parameters of Two-Prover-One-Round Game (2P1R) is an important task in PCPs literature as it would imply a smaller PCP with the same or stronger soundness. While this is a basic question in PCPs community, the connection between the parameters of PCPs and hardness of approximations is sometimes obscure to approximation algorithm community. In this paper, we investigate the connection between the parameters of 2P1R and the hardness of approximating the class of so-called connectivity problems. Based on recent development on 2P1R by Chan (STOC 2013) and techniques in PCPs literature, we improve hardness results of connectivity problems that are in the form k^σ , for some constant $\sigma > 0$, to hardness of

the form k^c for some explicit constant c , where k is a connectivity parameter. In addition, we show how to convert these hardness into hardness of the form $\mathcal{D}^{c'}$, where \mathcal{D} is the number of demand pairs (or the number of terminals).

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

CP24

Hardness of Finding Independent Sets in 2-Colorable and Almost 2-Colorable Hypergraphs

This work studies the hardness of finding independent sets in hypergraphs which are either 2-colorable, or $(1 - \epsilon)$ -almost 2-colorable, i.e. removing an ϵ fraction of its vertices and all incident hyperedges makes the hypergraph 2-colorable. Our results are:

- For any constant $\gamma > 0$, there is $\xi > 0$, such that, given a 4-uniform hypergraph on n vertices which is $(1 - \epsilon)$ -almost 2-colorable for $\epsilon = 2^{-(\log n)^\xi}$, it is quasi-NP-hard to find an independent set of $n / \left(2^{(\log n)^{1-\gamma}}\right)$ vertices.
- For any constants $\epsilon, \delta > 0$, given a $(1 - \epsilon)$ -almost 2-colorable 3-uniform hypergraph on n vertices, it is NP-hard to find an independent set of δn vertices. Assuming the d -to-1 Games Conjecture, this holds for $\epsilon = 0$, i.e. 2-colorable 3-uniform hypergraphs.

Rishi Saket
IBM T. J. Watson Research Center
ris.saket@gmail.com

Subhash Khot
New York University
khot@cs.nyu.edu

CP24

Hardness of Robust Graph Isomorphism, Lasserre Gaps, and Asymmetry of Random Graphs

Given two graphs which are almost isomorphic, is it possible to find a bijection which preserves most of the edges between the two? This is the problem of Robust Graph Isomorphism, an approximation variation of the Graph Isomorphism problem. We show that no polynomial-time algorithm solves this problem, conditioned on Feige's Random 3XOR Hypothesis. In addition, we show that the Lasserre SDP hierarchy needs a linear number of rounds to distinguish two isomorphic graphs from two far-from-isomorphic graphs.

Ryan O'Donnell
Carnegie Mellon University, USA
odonnell@cs.cmu.edu

John Wright
Department of Computer Science
Carnegie Mellon University
jswright@cs.cmu.edu

Chenggang Wu
IIIS, Tsinghua University
wuchenggang0316@gmail.com

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

CP24

Hypercontractive Inequalities Via SOS, and the Frankl-Rödl Graph

Our main result is a formulation and proof of the reverse hypercontractive inequality in the sum-of-squares (SOS) proof system. As a consequence we show that for any constant $0 < \gamma \leq 1/4$, the SOS/Lasserre SDP hierarchy at degree $4\lceil \frac{1}{4\gamma} \rceil$ certifies the statement “the maximum independent set in the Frankl–Rödl graph FR_n^γ has fractional size $o(1)$ ”. Here $\text{FR}_n^\gamma = (V, E)$ is the graph with $V = \{0, 1\}^n$ and $(x, y) \in E$ whenever $\text{dist}(x, y) = (1 - \gamma)n$ (an even integer). Finally, we also give an SOS proof of (a generalization of) the sharp $(2, q)$ -hypercontractive inequality for any even integer q .

Manuel Kauers
RISC Linz
manuel@kauers.de

Ryan O’Donnell
Carnegie Mellon University, USA
odonnell@cs.cmu.edu

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

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

CP25

The Generalized Terminal Backup Problem

We consider the following network design problem, that we call the Generalized Terminal Backup Problem. Given a graph (or a hypergraph) $G_0 = (V, E_0)$, a set of (at least 2) terminals $T \subseteq V$, costs $c(uv) \geq 0$ for every pair $u, v \in V$, and a requirement $r(t)$ for every $t \in T$, find a minimum cost multigraph $G = (V, E)$ such that $\lambda_{G_0+G}(t, T-t) \geq r(t)$ for any $t \in T$. We give a polynomial time algorithm for the case when c is node-induced. We also solve the problem when G_0 is the empty graph.

Attila Bernath
University of Warsaw
athos@cs.elte.hu

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

CP25

Analyzing the Optimal Neighborhood: Algorithms for Budgeted and Partial Connected Dominating Set Problems

We study natural generalizations of the classical connected dominating set (CDS) problem to the partial and budgeted versions. We develop surprisingly simple approximation algorithms for both these versions. For partial CDS, we are required to dominate a specified number of vertices, and

we obtain the first $O(\log \Delta)$ approximation algorithm. For budgeted CDS, where the goal is to maximize coverage subject to a budget constraint, we obtain a $\frac{1}{13}(1 - \frac{1}{e})$ approximation algorithm. Finally, we extend our results to a more general setting that has capacitated and weighted profit versions of both partial and budgeted CDS as special cases. While our algorithms themselves are simple, the results make a surprising use of the greedy set cover framework in defining a useful profit function.

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

Manish Purohit
Department of Computer Science
University of Maryland, College Park.
manishp@cs.umd.edu

Kantheni K. Sarpatwar
Department of Computer Science
University of Maryland, College Park
kasarpa@cs.umd.edu

CP25

Minimum d -Dimensional Arrangement with Fixed Points

In the Minimum d -Dimensional Arrangement Problem (d -DIMAP) we are given a weighted graph, and the goal is to find a 1-1 map of the vertices into \mathbb{Z}^d minimizing the total weighted stretch of the edges. This problem arises in VLSI placement and chip design. Motivated by these applications, we consider a generalization of d -DIMAP, where the positions of some k of the vertices (pins) is *fixed* and specified as part of the input. We are asked to extend this partial map to all the vertices, again minimizing the weighted stretch of edges. This generalization, which we refer to as d -DIMAP+, arises naturally in these application domains. We give approximation algorithms for this problem, based on a strengthening of the *spreading-metric* LP for d -DIMAP. We complement our upper bounds by integrality gaps, and inapproximability results.

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

Anastasios Sidiropoulos
University of Illinois at Urbana Champaign
sidiropo@gmail.com

CP25

Approximating Minimum Cost Connectivity Orientation and Augmentation

We investigate problems addressing combined connectivity augmentation and orientations settings. We give a polynomial time 6-approximation algorithm for finding a minimum cost subgraph of an undirected graph G that admits an orientation covering a non-negative crossing G -supermodular demand function, as defined by Frank. An important example is (k, l) -edge-connectivity, a common generalization of global and rooted edge-connectivity. Our algorithm is based on a non-standard application of the iterative rounding method. We observe that the standard linear program with cut constraints is not amenable and use an alternative linear program with partition and

co-partition constraints instead. The proof requires a new type of uncrossing technique on partitions and co-partitions.

Mohit Singh
Microsoft Research
mohits@microsoft.com

Laszlo Vegh
London School of Economics
l.vegh@lse.ac.uk

CP25

Improved Algorithms for Vertex Cover with Hard Capacities on Multigraphs and Hypergraphs

We present improved approximation algorithms for the minimum unweighted Vertex Cover problem with Hard Capacity constraints. On multigraphs, we improve the approximation ratio from 38 to 2.155, and on f -hypergraphs from $\max\{6f, 65\}$ to $2f$. Our algorithms consist of a two-step process, each based on rounding an appropriate LP. For multigraphs, the analysis in the second step relies on identifying a *matching structure* within any extreme point solution.

Wang Chi Cheung, Michel Goemans
Massachusetts Institute of Technology
Wangchi@mit.edu, goemans@math.mit.edu

Sam Chiu-Wai Wong
MIT
Samwong@mit.edu

CP26

Tight Bounds for Planar Strongly Connected Steiner Subgraph with Fixed Number of Terminals (and Extensions)

Feldman and Ruhl (FOCS '99; SICOMP '06) gave a novel $n^{O(k)}$ algorithm for the Strongly Connected Steiner Subgraph (SCSS) problem, where n is the number of vertices and k is the number of terminals. We explore how much easier the problem becomes on planar directed graphs (or more generally graphs excluding a fixed minor). In particular, we show

- A $2^{O(k \log k)} \cdot n^{O(\sqrt{k})}$ algorithm for planar SCSS.
- A matching lower bound by showing that planar SCSS does not have an $f(k) \cdot n^{o(\sqrt{k})}$ algorithm for any function f , unless the Exponential Time Hypothesis (ETH) fails.
- Assuming ETH, SCSS in general graphs does not have an $f(k) \cdot n^{o(k/\log k)}$ algorithm for any function f .
- Assuming ETH, there is no $f(k) \cdot n^{o(k)}$ time algorithm for the Directed Steiner Forest problem on acyclic planar graphs. This matches the upper bound of $n^{O(k)}$ for DSF due to Feldman and Ruhl

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

Daniel Marx
Computer and Automation Research Institute, Hungarian Academy of Sciences (MTA SZTAKI)

dmarx@cs.bme.hu

CP26

Linear-Time FPT Algorithms via Network Flow

In the area of parameterized complexity, to cope with NP-Hard problems, we introduce a parameter k besides the input size n , and we aim to design algorithms (called FPT algorithms) that run in $O(f(k)n^d)$ time for some function $f(k)$ and constant d . Though FPT algorithms have been successfully designed for many problems, typically they are not sufficiently fast because of huge $f(k)$ and d . In this paper, we give FPT algorithms with small $f(k)$ and d for many important problems including Odd Cycle Transversal and Almost 2-SAT by exploiting network flows. More specifically, we can choose $f(k)$ as a single exponential (4^k) and d as one, that is, linear in the input size. To the best of our knowledge, our algorithms achieve linear time complexity for the first time for these problems.

Yoichi Iwata, Keigo Oka
The University of Tokyo
y.iwata@is.s.u-tokyo.ac.jp, ogiekako@is.s.u-tokyo.ac.jp

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

CP26

A Near-Optimal Planarization Algorithm

The k -VERTEX PLANARIZATION problem asks whether a graph G can be made planar by deleting k vertices. We show that the problem can be solved in time $2^{O(k \log k)} \cdot n$, thereby improving on an earlier algorithm by Kawarabayashi (FOCS 2009) with running time $\Omega(2^{k^{\Omega(k^3)}}) \cdot n$. While earlier planarization algorithms relied heavily on deep results from the graph minors project, our techniques are elementary and practically self-contained.

Bart M. P. Jansen
University of Bergen
bart.jansen@ii.uib.no

Daniel Lokshtanov
UCSD
daniello@ii.uib.no

Saket Saurabh
IMS
saket@imsc.res.in

CP26

Linear Time Parameterized Algorithms Via Skew-Symmetric Multicuts

A skew-symmetric graph $(D = (V, A), \sigma)$ is a directed graph D with an involution σ on the set of vertices and arcs. In the d -skew symmetric multicut problem where we are given a skew-symmetric graph D , a family \mathcal{T} of d -sized subsets of vertices and an integer k , the objective is to decide if there is a set $X \subseteq A$ of k arcs such that every set J in the family has a vertex v such that v and $\sigma(v)$ are in different strongly connected components of $(V, A \setminus (X \cup \sigma(X)))$. This problem generalizes numerous well studied classical problems including Odd Cycle Transversal and 2-SAT dele-

tion. We give an algorithm for this problem which runs in time $O((4d)^k(m+n+\ell))$, where m is the number of arcs in the graph, n the number of vertices and ℓ the length of the family given in the input. As corollaries, we obtain linear time FPT algorithms for OCT and 2-SAT deletion. This is joint work with Saket Saurabh.

Ramanujan M. S.
University of Bergen
Ramanujan.Sridharan@ii.uib.no

Saket Saurabh
IMS
saket@imsc.res.in

CP26

A Subexponential Parameterized Algorithm for Subset TSP on Planar Graphs

Given a graph G and a subset S of vertices, the SUBSET TSP problem asks for a shortest closed walk in G visiting all vertices of S . The problem can be solved in time $2^k \cdot n^{O(1)}$ using the classical dynamic programming algorithms of Bellman and of Held and Karp, where $k = |S|$ and $n = |V(G)|$. Our main result is showing that the problem can be solved in time $(2^{O(\sqrt{k} \log k)} + W) \cdot n^{O(1)}$ if G is a planar graph with weights that are integers no greater than W . Our algorithm consists of two steps: (1) find a locally optimal solution, and (2) use it to guide a dynamic program. The proof of correctness of the algorithm depends on a treewidth bound on a graph obtained by combining an optimal solution with a locally optimal solution.

Philip Klein
Brown University
klein@brown.edu

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

CP26

Half-Integrality, LP-Branching and FPT Algorithms

We give a new, CSP-based approach to constructing LP-branching FPT algorithms. Instead of directly deriving Nemhauser-Trotter-style half-integrality, as previous work, we view half-integrality as a *discrete* relaxation of the search space (e.g., from $\{0, 1\}^n$ to $\{0, 1/2, 1\}^n$) with certain properties, which can be found with CSP tools. As an application of the framework, we give the first single-exponential algorithm for GROUP FEEDBACK VERTEX SET and SUBSET FEEDBACK VERTEX SET, and significantly improved algorithms for UNIQUE LABEL COVER.

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

CP27

Causal Erasure Channels

We introduce the causal erasure model and provide new upper bounds (impossibility results) and lower bounds (anal-

yses of codes) on the achievable rate. Our bounds separate the achievable rate in the causal erasures setting from the rates achievable in two related models: random erasure channels (strictly weaker) and fully adversarial erasure channels (strictly stronger). Our results contrast with existing results on correcting causal bit-flip errors (as opposed to erasures). For the separations we provide, the analogous separations for bit-flip models are either not known at all or much weaker.

Raef Bassily, Adam Smith
Pennsylvania State University
bassily@psu.edu, asmith@cse.psu.edu

CP27

Broadcast Throughput in Radio Networks: Routing Vs. Network Coding

The *broadcast throughput* in a network is defined as the average number of messages that can be transmitted per unit time from a given source to all other nodes. It has been shown that certain wired networks have asymptotically higher broadcast throughput when using network coding, where messages can be coded together in intermediate nodes, compared to the classical approach of routing, where messages are regarded as atomic tokens. Whether such a gap exists for wireless networks has been an open question of great interest. We show that there is a family of n -nodes radio networks in which the asymptotic throughput achievable via routing is a $\Theta(\log \log n)$ factor smaller than that of the optimal network coding algorithm. We also provide new tight upper and lower bounds showing that the asymptotic worst-case broadcast throughput is $\Theta(\frac{1}{\log n})$ messages-per-round for both routing and network coding.

Noga Alon
Schools of Mathematics and Computer Science
Tel Aviv University
nogaa@post.tau.ac.il

Mohsen Ghaffari, Bernhard Haeupler
MIT
ghaffari@mit.edu, haeupler@mit.edu

Majid Khabbazzian
University of Alberta
mkhabbazzian@ualberta.ca

CP27

Optimal Rate List Decoding of Folded Algebraic-Geometric Codes over Constant-Sized Alphabets

We construct a new list-decodable family of asymptotically good algebraic-geometric codes over fixed alphabets. The function fields underlying these codes are constructed using class field theory, and designed to have an automorphism of large order that is used to “fold” the AG code. This generalizes earlier work by the first author on folded AG codes based on cyclotomic function fields. The recent linear-algebraic approach to list decoding can be applied to our new codes, and crucially, we use the Chebotarev density theorem to establish a polynomial upper bound on the list-size. Our construction yields algebraic codes over constant-sized alphabets that can be list decoded up to the Singleton bound. Previous such results over constant-sized alphabets were either based on concatenation or involved taking a subcode of AG codes. In contrast, our result shows that these folded AG codes themselves have the claimed list

decoding property.

Venkatesan Guruswami
Carnegie Mellon University
guruswami@cmu.edu

Chaoping Xing
NTU Singapore
xingcp@ntu.edu.sg

CP27

Finding Orthogonal Vectors in Discrete Structures

Hopcroft's problem in d dimensions asks: given two sets of n vectors each in \mathbb{R}^{d+1} , is there a pair of vectors (one from each set) that are orthogonal? This problem has a long history and many applications. Known algorithms need at least $f(d) \cdot n^{2-1/O(d)}$ time for fast-growing f , and $n^{2-\epsilon} \cdot \text{poly}(d)$ time looks currently out of reach. We give randomized algorithms and almost matching lower bounds (modulo a breakthrough in SAT algorithms) for Hopcroft's problem over the ring of integers modulo m , and over finite fields. The algorithms arise from studying the communication problem of whether two lists of vectors (one held by Alice, one by Bob) contain an orthogonal pair of vectors (one from each list). We show the randomized communication complexity of the problem is related to the sizes of matching vector families. Building on these ideas, we give a very simple and efficient output-sensitive algorithm for matrix multiplication.

Ryan Williams
IBM Almaden Research Center, USA
rrwilliams@gmail.com

Huacheng Yu
Stanford University
yuhch123@gmail.com

CP27

Efficient Quantum Protocols for XOR Functions

We show that for any Boolean function f , the bounded-error quantum communication complexity $Q_\epsilon(f \circ \oplus)$ of XOR functions $f(x \oplus y)$ satisfies that

$$Q_\epsilon(f \circ \oplus) = O(2^d (\log \|\hat{f}\|_{1,\epsilon} + \log \frac{n}{\epsilon}) \log(1/\epsilon)),$$

where d is the GF[2]-degree of f , and $\|\hat{f}\|_{1,\epsilon} = \min_{g: \|f-g\|_\infty \leq \epsilon} \|\hat{g}\|_1$. This matches a previous lower bound by Lee and Shraibman for f with low GF[2]-degree. The result also confirms the quantum version of the Log-rank Conjecture for low-degree XOR functions. In addition, we show that the exact quantum communication complexity

$$Q_E(f) = O(2^d \log \|\hat{f}\|_0),$$

where $\|\hat{f}\|_0$ is the number of nonzero Fourier coefficients of f . This matches a previous lower bound by Buhrman and de Wolf for low-degree XOR functions.

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