Searchable Abstracts Document

ACM-SIAM Symposium on Discrete Algorithms (SODA24)

January 7–10, 2024

SIAM Symposium on Algorithm Engineering and Experiments (ALENEX24)

January 7–8, 2024

SIAM Symposium on Simplicity in Algorithms (SOSA24)

January 8–10, 2024

This document was current as of December 11, 2023. Abstracts appear as submitted.

3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 U.S.
Telephone: 800-447-7426 (U.S. & Canada) +1-215-382-9800 (Worldwide)
meetings@siam.org

IP1

New Frontiers in Structure vs Randomness with Applications to Combinatorics, Complexity, Algorithms.

In 1936, Erdos and Hajnal asked the following: Suppose you have a set S of integers from $\{1, 2, \ldots, N\}$ that contains at least N/C elements. Then, for large enough N, must S have three equally spaced numbers (i.e., a 3-term arithmetic progression)? Behrend in 1946 showed that C can be at most $\exp(\sqrt{(}\log N))$. Since then, the problem has been a cornerstone of the area of additive combinatorics, with the best bound being $C = (\log N)^{(1+c)}$ for some constant $c > 0$. Recent work obtained an exponential improvement showing that C can be as big as $\exp((\log N)^{0.09})$, thus getting closer to Behrend's construction. In this talk, I will describe this result and the main ingredient, a new variant of the "structure vs. randomness" paradigm. The latter is an old technique with many applications in complexity theory, algorithm design, and number theory, and the new variant can potentially lead to further progress. I will highlight two such applications: 1. Communication complexity: explicit separations between randomized and deterministic multiparty protocols. 2. Algorithm design: fast combinatorial algorithms for Boolean matrix multiplication, leading to fast algorithms saving more than poly(log) factors for enumerating triangles in graphs. Based on works with Amir Abboud, Nick Fischer, Zander Kelley, Shachar Lovett.

Raghu Meka
UCLA
raghum@cs.ucla.edu

IP2

How to get from A to B?

Starting with the famous "14-15 puzzle" attributed to Sam Loyd, many geometric, combinatorial, algebraic, and algorithmic questions can be phrased in the following form. Given two configurations, A and B, how can one transform A to B through a sequence of moves of a given type? Is this possible at all? We discuss several problems of this kind, including a basic problem in motion planning (joint work with Herbert Edelsbrunner), and a more recent puzzle by Kovaldzhi and Brunck with linear algebraic background (joint work with Gábor Tardos). We will also state some challenging open problems.

János Pach
Rényi Institute of Mathematics, Budapest
pach@cims.nyu.edu

IP3

Constructing and deConstructing Trust: A Cryptographers Perspective on the ML Pipeline

I will describe how to use cryptographic modeling and tools to achieve privacy, accuracy and robustness throughout the machine learning pipeline in the presence of computationally bounded adversaries.

Shafi Goldwasser
University of California, Berkeley and Simons Institute, U.S
shafi.goldwasser@gmail.com

IP4

Markov Chains for Collective Behaviors

Programmable matter explores how collections of computationally limited agents acting locally and asynchronously can achieve some useful coordinated behavior, spanning particle systems, colonies, and simple robot swarms. We take a stochastic approach using techniques from randomized algorithms and statistical physics to develop distributed algorithms for emergent collective behaviors that give guarantees and are robust to failures.

Dana Randall
Georgia Institute of Technology
randall@cc.gatech.edu

CP1

Breaking the $3/4$ Barrier for Approximate Maximin Share

We study the fundamental problem of fairly allocating a set of indivisible goods among n agents with additive valuations using the desirable fairness notion of maximin share (MMS). MMS is the most popular share-based notion, in which an agent finds an allocation fair to her if she receives goods worth at least her MMS value. An allocation is called MMS if all agents receive at least their MMS value. However, since MMS allocations need not exist when $n > 2$, a series of works showed the existence of approximate MMS allocations with the current best factor of $\frac{3}{4} + O(\frac{1}{n})$. The recent work by Akrami et al. showed the limitations of existing approaches and proved that they cannot improve this factor to $3/4 + \Omega(1)$. In this paper, we bypass these barriers to show the existence of $(\frac{3}{4} + \frac{3}{3836})$-MMS allocations by developing new reduction rules and analysis techniques.

Hannaneh Akrami
Max Planck Institute for Informatics
Graduate School of Computer Science, Saarland University
hakrami@mpi-inf.mpg.de

Jugal Garg
University of Illinois at Urbana-Champaign
jugal@illinois.edu

CP1

On the Tractability Frontier of Combinatorial Contracts

We explore the tractability frontier for combinatorial contracts in the single-agent combinatorial action model of [Duetting, Ezra, Feldman, Kesselheim, FOCS 2021] and the multi-agent model of [Duetting, Ezra, Feldman, Kesselheim, STOC 2023]. For the single-agent combinatorial action setting, we give a poly-time algorithm for finding the optimal contract, provided that the agent's demand problem can be solved efficiently and there are poly-many breakpoints in the agent's demand. This implies an efficient algorithm for supermodular rewards and submodular costs. We also explore a setting with XOS rewards and additive costs, and show that existing approaches provably fail. For the multi-agent setting, our focus is on settings with supermodular rewards. We give an additive PTAS for a natural class of graph-based rewards when the agents have identical costs, and show that even in this special case

it's NP-hard to obtain any finite multiplicative approximation, or an additive FPTAS.

Ramiro Deo-Campo Vuong
University of Southern California
rdeocamp@usc.edu

Paul Duetting
Google Research
duetting@google.com

Shaddin Dughmi
University of Southern California
shaddin@usc.edu

Michal Feldman
Tel Aviv University
Microsoft Research Israel
mfeldman@tauex.tau.ac.il

Yoav Gal Tzur
Tel Aviv University
yoavgaltzur@mail.tau.ac.il

Neel B. Patel, Aditya Prasad
University of Southern California
neelbpat@usc.edu, aprasad4@usc.edu

CP1

Revenue Maximization for Buyers with Costly Participation

We study mechanisms for selling a single item when buyers have private costs for participating in the mechanism. An agent's participation cost can also be interpreted as an outside option value that she must forego to participate. This substantially changes the revenue maximization problem, which becomes non-convex in the presence of participation costs. For multiple buyers, we show how to construct a $(2+\epsilon)$-approximately revenue-optimal mechanism in polynomial time. Our approach makes use of a many-buyers-to-single-buyer reduction, and in the single-buyer case our mechanism improves to an FPTAS. We also bound the menu size and the sample complexity for the optimal single-buyer mechanism. Moreover, we show that posting a single price in the single-buyer case is in fact optimal under the assumption that either (1) the participation cost is independent of the value, and the value distribution has decreasing marginal revenue or monotone hazard rate; or (2) the participation cost is a concave function of the value. When there are multiple buyers, we show that sequential posted pricing guarantees a large fraction of the optimal revenue under similar conditions.

Yannai A. Gonczarowski
Microsoft Research
yannai@gonch.name

Nicole Immorlica
Microsoft Research, USA
nicimm@gmail.com

Yingkai Li
Yale University
yingkai.li@yale.edu

Brendan Lucier
Microsoft Research
brlucier@microsoft.com

CP1

Impossibilities for Obviously-Strategy-Proof Mechanisms

We explore the approximation power of deterministic obviously strategy-proof mechanisms in auctions, where the objective is welfare maximization. A trivial ascending auction on the grand bundle guarantees an approximation of $\min\{m, n\}$ for all valuation classes, where m is the number of items and n is the number of bidders. We focus on two classes of valuations considered simple: additive valuations and unit-demand valuations. For additive valuations, Bade and Gonczarowski [EC'17] have shown that exact welfare maximization is impossible. No impossibilities are known for unit-demand valuations. We show that if bidders' valuations are additive or unit-demand, then no obviously strategy-proof mechanism gives an approximation better than $\min\{m, n\}$. Thus, the aforementioned trivial ascending auction on the grand bundle is the optimal obviously strategy-proof mechanism. These results illustrate a stark separation between the power of dominant-strategy and obviously strategy-proof mechanisms. The reason for it is that for both of these classes the dominant-strategy VCG mechanism does not only optimize the welfare exactly, but is also easy both from a computation and communication perspective. In addition, we prove tight impossibilities for unknown single-minded bidders in a multi-unit auction and in a combinatorial auction. We show that in these environments as well, a trivial ascending auction on the grand bundle is optimal.

Shiri Ron
Weizmann Institute of Science
shiri.ron@mail.huji.ac.il

CP1

Prior-Independent Auctions for Heterogeneous Bidders

We study the design of prior-independent auctions in a setting with heterogeneous bidders. In particular, we consider the setting of selling to n bidders whose values are drawn from n independent but not necessarily identical distributions. We work in the robust auction design regime, where we assume the seller has no knowledge of the bidders' value distributions and must design a mechanism that is prior-independent. While there have been many strong results on prior-independent auction design in the i.i.d. setting, not much is known for the heterogeneous setting, even though the latter is of significant practical importance. Unfortunately, no prior-independent mechanism can hope to always guarantee any approximation to Myerson's revenue in the heterogeneous setting; similarly, no prior-independent mechanism can consistently do better than the second-price auction. In light of this, we design a family of (parametrized) randomized auctions which approximates at least one of these benchmarks: For heterogeneous bidders with regular value distributions, our mechanisms either achieve a good approximation of the expected revenue of an optimal mechanism (which knows the bidders' distributions) or exceeds that of the second-price auction by a certain multiplicative factor.

Guru Guruganesh
Google Research
gurug@google.com

Aranyak Mehta, Di Wang
Google
aranyak@google.com, wadi@google.com

Kangning Wang
Stanford University
knwang@stanford.edu

CP2

Dynamically Maintaining the Persistent Homology of Time Series

We present a dynamic data structure for maintaining the persistent homology of a time series of real numbers. The data structure supports local operations, including the insertion and deletion of an item and the cutting and concatenating of lists, each in time $O(\log? n + k)$, in which n is the number of critical items and k is the number of changes in the augmented persistence diagram. To achieve this, we design a tailor-made tree structure with an unconventional representation, referred to as banana tree, which may be useful in its own right.

Sebastiano Cultrera Di Montesano,
Herbert Edelsbrunner, Monika Henzinger
Institute of Science and Technolology Austria
sebastiano.cultrera@ist.ac.at,
herbert.edelsbrunner@ist.ac.at,
monika.henzinger@ist.ac.at

Lara Ost
Department of Computer Science, University of Vienna
lara.ost@univie.ac.at

CP2

Vertical Decomposition in 3D and 4D with Applications to Line Nearest-Neighbor Searching in 3D

Vertical decomposition is a widely used general technique for decomposing the cells of arrangements of semi-algebraic sets in $\mathbb{R}^d$ into constant-complexity subcells. In this paper, we settle in the affirmative a few long-standing open problems involving the vertical decomposition of substructures of arrangements for $d = 3, 4$: (i) Let $\mathcal{S}$ be a collection of n semi-algebraic sets of constant complexity in $\mathbb{R}^3$, and let $U(m)$ be an upper bound on the complexity of the union $\mathcal{U}(\mathcal{S}')$ of any subset $\mathcal{S}' \subseteq \mathcal{S}$ of size at most m. We prove that the complexity of the vertical decomposition of the complement of $\mathcal{U}(\mathcal{S})$ is $O^*(n^2 + U(n))$ (where the $O^*(\cdot)$ notation hides subpolynomial factors). We also show that the complexity of the vertical decomposition of the entire arrangement $\mathcal{AS})$ is $O^*(n^2 + X)$, where X is the number of vertices in $\mathcal{A}(\mathcal{S})$. (ii) Let $\mathcal{F}$ be a collection of n trivariate functions whose graphs are semi-algebraic sets of constant complexity. We show that the complexity of the vertical decomposition of the portion of the arrangement $\mathcal{A}(\mathcal{F})$ in $\mathbb{R}^4$ lying below the lower envelope of $\mathcal{F}$ is $O^*(n^3)$. These results lead to efficient algorithms for a variety of problems involving these decompositions. In addition, as a main domain of applications, we study various proximity problems involving points and lines in $\mathbb{R}^3$.

Esther Ezra
Bar-Ilan University
ezraest@cs.biu.ac.il

Pankaj K. Agarwal
Duke University
pankaj@duke.edu

MICHA Sharir
Tel-Aviv University
s.micha7777@gmail.com

CP2

Dynamic Dynamic Time Warping

The Dynamic Time Warping (DTW) distance is a popular similarity measure for polygonal curves (i.e., sequences of points). It finds many theoretical and practical applications, especially for temporal data, and is known to be a robust, outlier-insensitive alternative to the Frechet distance. For static curves of at most n points, the DTW distance can be computed in $O(n^2)$ time in constant dimension. This tightly matches a SETH-based lower bound, even for curves in $\mathbb{R}^1$. In this work, we study dynamic algorithms for the DTW distance. Here, the goal is to design a data structure that can be efficiently updated to accommodate local changes to one or both curves, such as inserting or deleting vertices and, after each operation, reports the updated DTW distance. We give such a data structure with update and query time $O(n^{1.5} \log n)$, where n is the maximum length of the curves. As our main result, we prove that our data structure is conditionally optimal, up to subpolynomial factors. More precisely, we prove that, already for curves in $\mathbb{R}^1$, there is no dynamic algorithm to maintain the DTW distance with update and query time $O(n^{1.5-\delta})$ for any constant $\delta > 0$, unless the Negative-k-Clique Hypothesis fails. In fact, we give matching upper and lower bounds for various trade-offs between update and query time, even in cases where the lengths of the curves differ.

Karl Bringmann
Saarland University and Max Planck Institute for Informatics
Saarland Informatics Campus, Germany
bringmann@cs.uni-saarland.de

Nick Fischer
Weizmann Institute of Science
nick.fischer@weizmann.ac.il

Ivor Hoog
Technical University of Denmark
vanderhoog@gmail.com

Evangelos Kipouridis
BARC (Basic Algorithms Research Copenhagen)
University of Copenhagen
kipouridis@cs.uni-saarland.de

Tomasz Kociumaka
Max Planck Institute for Informatics
tomasz.kociumaka@mpi-inf.mpg.de

Eva Rotenberg
Technical University of Denmark
erot@dtu.dk

CP2

Sorting Pattern-Avoiding Permutations via 0-1 Matrices Forbidding Product Patterns

We consider the problem of comparison-sorting an n-permutation S that avoids some k-permutation π. Chalermsook, Goswami, Kozma, Mehlhorn, and Saranurak prove that when S is sorted by inserting the elements into the GreedyFuture binary search tree, the running time

is linear in the extremal function $\mathrm{Ex}(P \otimes \mathrm{hat}, n)$. This is the maximum number of 1s in an $n \times n$ 0–1 matrix avoiding $P \otimes \mathrm{hat}$, where $P \otimes \mathrm{hat}$ is the Kronecker product of the $k \times k$ permutation matrix P representing π and a 2x3 "hat' pattern. Applying off-the-shelf results on the extremal functions of 0–1 matrices, it was known that $\mathrm{Ex}(P \otimes \mathrm{hat}, n) = \Omega(n\alpha(n))$ and $O(n \cdot 2^{(\alpha(n))^{3k/2-O(1)}})$, where $\alpha(n)$ is the inverse-Ackermann function. In this paper we give nearly tight upper and lower bounds on the density of $P \otimes$ hat-free matrices in terms of "n', and improve the dependence on "k' from doubly exponential to singly exponential. $\mathrm{Ex}(P \otimes \mathrm{hat}, n) = \Omega(n \cdot 2^{\alpha(n)})$ and $O(n \cdot 2^{O(k^2)+(1+o(1))\alpha(n)})$. As a consequence, sorting π-free sequences can be performed in $O(n2^{(1+o(1))\alpha(n)})$ time. For many corollaries of the dynamic optimality conjecture, the best analysis uses forbidden 0–1 matrix theory. Our analysis may be useful in analyzing other classes of access sequences on binary search trees.

Parinya Chalermsook
Aalto University, Espoo, Finland
parinya.chalermsook@aalto.fi

Seth Pettie
University of Michigan, Ann Arbor
seth@pettie.net

Sorrachai Yingchareonthawornchai
Simons Institute, UC Berkeley
sorrachai.yingchareonthawornchai@aalto.fi

CP2

Dynamic Dictionary with Subconstant Wasted Bits Per Key

Dictionaries have been one of the central questions in data structures. A dictionary data structure maintains a set of key-value pairs under insertions and deletions such that given a query key, the data structure efficiently returns its value. The state-of-the-art dictionaries [Bender, Farach-Colton, Kuszmaul, Kuszmaul, Liu 2022] store n key-value pairs with only $O(n \log^{(k)} n)$ bits of redundancy, and support all operations in $O(k)$ time, for $k \leq \log^* n$. It was recently shown to be optimal [Li, Liang, Yu, Zhou 2023b]. In this paper, we study the regime where the redundant bits is $R = o(n)$, and show that when R is at least $n/\mathrm{poly}\log n$, all operations can be supported in $O(\log^* n + \log(n/R))$ time, matching the lower bound in this regime [Li, Liang, Yu, Zhou 2023b]. We present two data structures based on which range R is in. The data structure for $R < n/\log^{0.1} n$ utilizes a generalization of adapters studied in [Berger, Kuszmaul, Polak, Tidor, Wein 2022] and [Li, Liang, Yu, Zhou 2023a]. The data structure for $R \geq n/\log^{0.1} n$ is based on recursively hashing into buckets with logarithmic sizes.

Tianxiao Li, Jingxun Liang
Tsinghua University
litx20@mails.tsinghua.edu.cn,
liangjx20@mails.tsinghua.edu.cn

Huacheng Yu
Princeton University
yuhch123@gmail.com

Renfei Zhou
Tsinghua University
zhourf20@mails.tsinghua.edu.cn

CP3

Optimally Repurposing Existing Algorithms to Obtain Exponential-Time Approximations

The goal of this paper is to understand how exponential-time approximation algorithms can be obtained from existing polynomial-time approximation algorithms, existing parameterized exact algorithms, and existing parameterized approximation algorithms. More formally, we consider a monotone subset minimization problem over a universe of size n (e.g., Vertex Cover or Feedback Vertex Set). We have access to an algorithm that finds an α-approximate solution in time $c^k \cdot n^{O(1)}$ if a solution of size k exists (and more generally, an extension algorithm that can approximate in a similar way if a set can be extended to a solution with k further elements). Our goal is to obtain a $d^n \cdot n^{O(1)}$ time β-approximation algorithm for the problem with d as small as possible. That is, for every fixed $\alpha, c, \beta \geq 1$, we would like to determine the smallest possible d that can be achieved in a model where our problem-specific knowledge is limited to checking the feasibility of a solution and invoking the α-approximate extension algorithm. Our results completely resolve this question: - For every fixed $\alpha, c, \beta \geq 1$, a simple algorithm ("approximate monotone local search") achieves the optimum value of d. - Given $\alpha, c, \beta \geq 1$, we can efficiently compute the optimum d up to any precision $\epsilon > 0$.

Baris Can Esmer
CISPA Helmholtz Center for Information Security
Saarland University
baris-can.esmer@cispa.de

Ariel Kulik
CISPA Helmholtz Center for Information Security
, Saarland Informatics Campus, Germany
ariel.kulik@cispa.de

Dániel Marx
CISPA Helmholtz Center for Information Security
marx@cispa.de

Daniel Neuen
University of Bremen
dneuen@uni-bremen.de

Roohani Sharma
Max Planck Institute of Informatics
roohani.sharma90@gmail.com

CP3

Tree Containment Above Minimum Degree is FPT

According to the classic Chvatal's Lemma from 1977, a graph of minimum degree $\delta(G)$ contains every tree on $\delta(G)+1$ vertices. Our main result is the following algorithmic "extension' of Chvl's Lemma: For any n-vertex graph G, integer k, and a tree T on at most $\delta(G) + k$ vertices, deciding whether G contains a subgraph isomorphic to T, can be done in time $f(k) \cdot n^{O(1)}$ for some function f of k only. The proof of our main result is based on an interplay between extremal graph theory and parameterized algorithms.

Fedor Fomin
Dep. of Informatics

University of Bergen
fedor.fomin@uib.no

Petr Golovach
University of Bergen, Norway
petr.golovach@uib.no

Danil Sagunov
St. Petersburg Department of V.A. Steklov Institute of Mathematics
danilka.pro@gmail.com

Kirill Simonov
TU Wien
kirillsimonov@gmail.com

CP3

Determinantal Sieving

We introduce determinantal sieving, a new, remarkably powerful tool in the toolbox of algebraic FPT algorithms. Given a polynomial $P(X)$ on a set of variables $X = \{x_1, \ldots, x_n\}$ and a linear matroid $M = (X, \mathcal{I})$ of rank k, both over a field $\mathbb{F}$ of characteristic 2, in 2^k evaluations we can sieve for multilinear terms in the monomial expansion of P whose support is a basis for M. Alternatively, using 2^k evaluations of P we can sieve for those monomials whose odd support spans M. Applying this framework, we improve on a range of algebraic FPT algorithms, such as: 1. q-Matroid Intersection in time $O^*(2^{(q-2)k})$ and q-Matroid Parity in time $O^*(2^{qk})$, improving on $O^*(4^{qk})$ (Brand and Pratt, ICALP 2021) 2. T-Cycle, Colourful (s,t)-Path, Colourful (S,T)-Linkage in undirected graphs, all in $O^*(2^k)$ time, improving on $O^*(2^{k+|S|})$ (Fomin et al., SODA 2023) 3. Instances of the Diverse X paradigm, such as k-Distinct Branchings improving from time $O^*(2^{O(k \log k)})$ to $O^*(2^k)$ (Bang-Jensen et al., ESA 2021), and Diverse Perfect Matchings from $O^*(2^{2^{O(rd)}})$ to $O^*(2^{r(r-1)/2})$ (Fomin et al., STACS 2021) Over general fields, we achieve similar results using exponential space via the exterior algebra. For a certain class of monotone arithmetic circuits, we achieve the same running time. However, the odd support sieving result appears to be specific to characteristic 2.

Eduard Eiben
Royal Holloway University of London, Egham, UK
eduard.eiben@rhul.ac.uk

Tomohiro Koana
TU Berlin
tomohiro.koana@tu-berlin.de

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

CP3

Shortest Disjoint Paths on a Grid

The well-known k-disjoint paths problem involves finding pairwise vertex-disjoint paths between k specified pairs of vertices within a given graph if they exist. In the shortest k-disjoint paths problem one looks for such paths of minimum total length. Despite nearly 50 years of active research on the k-disjoint paths problem, many open problems and complexity gaps still persist. A particularly well-defined scenario, inspired by VLSI design, focuses on infinite rectangular grids where the terminals are placed at arbitrary grid points. While the decision problem in this context remains NP-hard, no prior research has provided any positive results for the optimization version. The main result of this paper is a fixed-parameter tractable (FPT) algorithm for this scenario. It is important to stress that this is the first result achieving the FPT complexity of the shortest disjoint paths problem in any, even very restricted classes of graphs where we do not put any restriction on the placements of the terminals.

Mathieu Mari
Ecole Normale Superieure, Paris
mathieu.mari@lirmm.fr

Anish Mukherjee
University of Warwick
anish.mukherjee@warwick.ac.uk

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

CP3

Fully Dynamic Approximation Schemes on Planar and Apex-Minor-Free Graphs

The classic technique of Baker [J. ACM '94] is the most fundamental approach for designing approximation schemes on planar, or more generally topologically-constrained graphs, and it has been applied in a myriad of different variants and settings throughout the last 30 years. In this work we propose a dynamic variant of Baker's technique, where instead of finding an approximate solution in a given static graph, the task is to design a data structure for maintaining an approximate solution in a fully dynamic graph a graph changing over time by edge deletions and edge insertions. We address the two most basic problems Maximum Weight Independent Set and Minimum Weight Dominating Set and we prove that for a fully dynamic n-vertex planar graph G, one can: maintain a $(1-\varepsilon)$-approximation of the maximum weight of an independent set in G with amortized update time $f(\varepsilon) \cdot n^{o(1)}$; and, under the additional assumption that the maximum degree of the graph is bounded at all times by a constant, also maintain a $(1+\varepsilon)$-approximation of the minimum weight of a dominating set in G with amortized update time $f(\varepsilon) \cdot n^{o(1)}$. In both cases, $f(\varepsilon)$ is doubly-exponential in $\text{poly}(1/\varepsilon)$ and the data structure can be initialized in time $f(\varepsilon) \cdot n^{1+o(1)}$. All our results in fact hold in the larger generality of any graph class that excludes a fixed apex-graph as a minor.

Marek Sokolowski
University of Warsaw
marek.sokolowski@mimuw.edu.pl

Tuukka Korhonen
University of Bergen
tuukka.korhonen@uib.no

Wojciech Nadara
University of Warsaw
w.nadara@mimuw.edu.pl

Michal Pilipczuk
Institute of Informatics, University of Warsaw

michal.pilipczuk@mimuw.edu.pl

CP4

Rationality-Robust Information Design: Bayesian Persuasion under Quantal Response

Classic mechanism/information design imposes the assumption that agents are *fully rational*, meaning each of them always selects the action that maximizes her expected utility. Yet many empirical evidence suggests that human decisions may deviate from this full rationality assumption. In this work, we attempt to relax the full rationality assumption with *bounded rationality*. Specifically, we formulate the bounded rationality of an agent by adopting the quantal response model [**?**]. We develop a theory of rationality-robust information design in the canonical setting of Bayesian persuasion [**?**] with binary receiver action. We first identify conditions under which the optimal signaling scheme structure for a fully rational receiver remains optimal or approximately optimal for a boundedly rational receiver. In practice, it might be costly for the designer to estimate the degree of the receivers bounded rationality level. Motivated by this practical consideration, we then study the existence and construction of *robust* signaling schemes when there is uncertainty about the receiver's bounded rationality level.

Yiding Feng
Northwestern University
Department of Electrical Engineering & Computer Science
yidingfeng2021@u.northwestern.edu

Chien-Ju Ho
Washington University in St. Louis
chienju.ho@wustl.edu

Wei Tang
Columbia University
wt2359@columbia.edu

CP4

Minimization Is Harder in the Prophet World

We study I.I.D. prophet inequalities for cost minimization, where the problem is to pick a cost from a sequence $X_1, \ldots, X_n$ drawn independently from a known distribution in an online manner, and compete against the prophet who can see all the realizations upfront and select the minimum. In contrast to the well-studied rewards maximization setting where a simple threshold strategy achieves a competitive ratio of ≈ 0.745 for all distributions, the cost minimization setting turns out to be much more complex. In our main result, we obtain a complete and nuanced characterization of the I.I.D. cost prophet inequality: if the expected value of the given distribution is infinite, then the competitive ratio is also infinite. On the other hand, when the expected value is finite, we show that the competitive ratio of the optimal stopping strategy is a (distribution-dependent) constant, which we characterize precisely as the solution to a simple inequality. Furthermore, we obtain a closed form for this constant for a broad class of distributions we call entire distributions, we show that the constant is 2 for MHR distributions and obtain matching lower bounds for all our results. We then focus on single-threshold strategies and design a single threshold that achieves a tight $\mathrm{O}(\mathrm{polylog}(n))$-factor approximation.

Vasilis Livanos
University of Illinois Urbana-Champaign
livanos3@illinois.edu

Ruta Mehta
University of Illinois at Urbana-Champaign
rutameht@illinois.edu

CP4

Equilibrium Dynamics in Market Games with Exchangeable and Divisible Resources

We study a market game with $n \geq 2$ players competing over $m \geq 1$ divisible resources of different finite capacities. Resources are traded via the proportional sharing mechanism, where players are price-anticipating, meaning that they can influence the prices with their bids. Additionally, each player has an initial endowment of the resources which are sold at market prices. Although the players' total profit functions may be discontinuous in the bids, we prove existence and uniqueness of pure Nash equilibria of the resulting market game. Then, we study a discrete dynamic arising from repeatedly taking the (unique) equilibrium resource allocation as initial endowments for the next market game. We prove that the total utility value of the dynamic converges to either an optimal allocation value (maximizing total utility over the allocation space) or to a restricted optimal allocation value, where the restriction is defined by fixing some tight resources which are exclusively allocated to a single player. As a corollary, it follows that for strictly concave utility functions, the aggregated allocation vector of the dynamic converges to the unique (possibly restricted) optimal aggregated allocation, and for linear utility functions, we even get convergence of the dynamic to a (possibly restricted) optimal solution in the (non-aggregated) original allocation space.

José Correa
Universidad de Chile
correa@uchile.cl

Tobias Harks
Faculty of Computer Science and Mathematics, University of Passau, Germany
tobias.harks@uni-passau.de

Anja Schedel
University of Passau
anja.schedel@uni-passau.de

José Verschae
Institute for Mathematical and Computational Engineering,
Pontificia Universidad Católica de Chile, Chile
jverschae@uc.cl

CP4

Simple Delegated Choice

This paper studies delegation in a model of discrete choice. In the delegation problem, an uninformed principal must consult an informed agent to make a decision. Both the agent and principal have preferences over the decided-upon action which vary based on the state of the world, and which may not be aligned. The principal designs a menu of actions, from which the agent chooses upon observing the state. We consider a setting where the decision being delegated is a choice of a utility-maximizing action from a set of several options. We assume the shared portion of the

agent's and principal's utilities is drawn from a distribution known to the principal, and that utility misalignment takes the form of a known bias for or against each action. We provide tight approximation analyses for simple threshold policies under three increasingly general sets of assumptions. With independently-distributed utilities, we prove a 3-approximation. When the agent has an outside option the principal cannot rule out, the constant-approximation fails, but we prove a $\log \rho / \log \log \rho$-approximation, where ρ is the ratio of the maximum value to the optimal utility. We also give a weaker but tight bound that holds for correlated values, and complement our upper bounds with hardness results. One special case of our model is utility-based assortment optimization, for which our results are new.

Samuel Taggart
Department of Mathematics
Oberlin College
samuel.taggart@oberlin.edu

Ali Ali Khodabakhsh
University of Texas, Austin
ali.kh@utexas.edu

Emmanouil Pountourakis
Drexel University
manolis@drexel.edu

CP4

Bandit Algorithms for Prophet Inequality and Pandora's Box

The Prophet Inequality and Pandora's Box problems are fundamental stochastic problem with applications in Mechanism Design, Online Algorithms, Stochastic Optimization, Optimal Stopping, and Operations Research. A usual assumption in these works is that the probability distributions of the n underlying random variables are given as input to the algorithm. Since in practice these distributions need to be learned under limited feedback, we initiate the study of such stochastic problems in the Multi-Armed Bandits model. In the Multi-Armed Bandits model we interact with n unknown distributions over T rounds: in round t we play a policy $x^{(t)}$ and only receive the value of $x^{(t)}$ as feedback. The goal is to minimize the regret, which is the difference over T rounds in the total value of the optimal algorithm that knows the distributions vs. the total value of our algorithm that learns the distributions from the limited feedback. Our main results give near-optimal $\widetilde{O}\big(\mathsf{poly}(n)\sqrt{T}\big)$ total regret algorithms for both Prophet Inequality and Pandora's Box. Our proofs proceed by maintaining confidence intervals on the unknown indices of the optimal policy. The exploration-exploitation tradeoff prevents us from directly refining these confidence intervals, so the main technique is to design a regret upper bound function that is learnable while playing low-regret Bandit policies.

Yifan Wang
Georgia Tech
ywang3782@gatech.edu

Khashayar Gatmiry
Massachusetts Institute of Technology
gatmiry@mit.edu

Thomas Kesselheim
University of Bonn
thomas.kesselheim@uni-bonn

Sahil Singla
Georgia Tech
ssingla@gatech.edu

CP5

An $\tilde{\Omega}\big(\sqrt{\log |T|}\big)$ Lower Bound for Steiner Point Removal

In the Steiner point removal (SPR) problem, we are given a (weighted) graph G and a subset T of its vertices called terminals, and the goal is to compute a (weighted) graph H on T that is a minor of G, such that the distance between every pair of terminals is preserved to within some small multiplicative factor, that is called the *stretch* of H. It has been shown that on general graphs we can achieve stretch $O(\log |T|)$ [Filtser, 2018]. On the other hand, the best-known stretch lower bound is 8 [Chan-Xia-Konjevod-Richa, 2006], which holds even for trees. In this work, we show an improved lower bound of $\tilde{\Omega}\big(\sqrt{\log |T|}\big)$.

Yu Chen
University of Pennsylvania
yu.chen@epfl.ch

Zihan Tan
Rutgers University
zihantan1993@gmail.com

CP5

2-Approximation for Prize-Collecting Steiner Forest

Approximation algorithms for the prize-collecting Steiner forest problem (PCSF) have been a subject of research for over three decades, starting with the seminal works of Agrawal, Klein, and Ravi [STOC'91] and Goemans and Williamson [SODA'92] on Steiner forest and prize-collecting problems. In this paper, we propose and analyze a natural deterministic algorithm for PCSF that achieves a 2-approximate solution in polynomial time. This represents a significant improvement compared to the previously best known algorithm with a 2.54-approximation factor developed by Hajiaghayi and Jain [SODA'06]. Furthermore, Könemann, Olver, Pashkovich, Ravi, Swamy, and Vygen [APPROX/RANDOM'17] have established an integrality gap of at least 9/4 for the natural LP relaxation for PCSF. However, we surpass this gap through the utilization of a combinatorial algorithm and a novel analysis technique. Since 2 is the best known approximation guarantee for Steiner forest problem [AKR, SIAM J. Comput.'95] (see also [GW, SIAM J. Comput.'95]), which is a special case of PCSF, our result matches this factor and closes the gap between the Steiner forest problem and its generalized version, PCSF.

Ali Ahmadi, Iman Gholami, MohammadTaghi Hajiaghayi, Peyman Jabbarzade, Mohammad Mahdavi
University of Maryland
ahmadia@umd.edu, igholami@umd.edu, hajiagha@umd.edu, peymanj@umd.edu, mahdavi@umd.edu

CP5

A Polynomial-Time opt$^{\epsilon}$-Approximation Algorithm for Maximum Independent Set of Connected

Subgraphs in a Planar Graph

In the Maximum Independent Set of Objects problem, we are given an n-vertex planar graph G and a family D of N objects, where each object is a connected subgraph of G. The task is to find a subfamily $F \subseteq D$ of maximum cardinality that consists of pairwise disjoint objects. This problem is NP-hard and is equivalent to the problem of finding the maximum number of pairwise disjoint polygons in a given family of polygons in the plane. As shown by Adamaszek et al. (J. ACM'19), the problem admits a quasi-polynomial time approximation scheme (QPTAS): a $(1-\epsilon)$-approximation algorithm whose running time is bounded by $2^{\mathrm{poly}(\log(N),1/\epsilon)}n^{O(1)}$. Nevertheless, to the best of our knowledge, in the polynomial-time regime only the trivial $O(N)$-approximation is known for the problem in full generality. In the restricted setting where the objects are pseudolines in the plane, Fox and Pach (SODA'11) gave an N^{ϵ}-approximation algorithm in $N^{2^{\tilde{O}(1/\epsilon)}}$ time, for any $\epsilon > 0$. In this work, we present an OPT^{ϵ}-approximation algorithm for the problem that runs in time $N^{\tilde{O}(1/\epsilon^2)}n^{O(1)}$, for any $\epsilon > 0$, thus improving upon the result of Fox and Pach both in terms of generality and in terms of the running time. Our approach combines the methodology of Voronoi separators, introduced by Marx and Pilipczuk (TALG'22), with a new analysis of the approximation factor.

Wegrzycki Karol
Saarland University and Max Planck Institute for Informatics
kwegrzyc@mpi-inf.mpg.de

Jana Cslovjecsek
EPFL
Switzerland
jana.cslovjecsek@epfl.ch

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

CP5

Single-Source Unsplittable Flows in Planar Graphs

The single-source unsplittable flow (SSUF) problem asks to send flow from a common source to different terminals with unrelated demands, each terminal being served through a single path. One of the most heavily studied SSUF objectives is to minimize the violation of some given arc capacities. A seminal result of Dinitz, Garg, and Goemans showed that, whenever a fractional flow exists respecting the capacities, then there is an unsplittable one violating the capacities by at most the maximum demand. Goemans conjectured a very natural cost version of the same result, where the unsplittable flow is required to be no more expensive than the fractional one. This intriguing conjecture remains open. More so, there are arguably no non-trivial graph classes for which it is known to hold. We show that a slight weakening of it (with at most twice as large violations) holds for planar graphs. Our result is based on a connection to a highly structured discrepancy problem, whose repeated resolution allows us to successively reduce the number of paths used for each terminal, until we obtain an unsplittable flow. Moreover, our techniques also extend to simultaneous upper and lower bounds on the flow values. This also affirmatively answers a conjecture of Morell and Skutella for planar SSUF.

Vera Traub
University of Bonn
traub@or.uni-bonn.de

Laura Vargas Koch, Rico Zenklusen
ETH Zurich
laura.vargas@ifor.math.ethz.ch, ricoz@ethz.ch

CP5

Fast Algorithms for Directed Graph Partitioning Using Flows and Reweighted Eigenvalues

We consider a new semidefinite programming relaxation for directed edge expansion, which is obtained by adding triangle inequalities to the reweighted eigenvalue formulation. Applying the matrix multiplicative weight update method on this relaxation, we derive almost linear-time algorithms to achieve $O(\sqrt{\log n})$-approximation and Cheeger-type guarantee for directed edge expansion, as well as an improved cut-matching game for directed graphs. This provides a primal-dual flow-based framework to obtain the best known algorithms for directed graph partitioning. The same approach also works for vertex expansion and for hypergraphs, providing a simple and unified approach to achieve the best known results for different expansion problems and different algorithmic techniques.

Lap Chi Lau, Kam Chuen Tung, Robert Wang
University of Waterloo
lapchi@uwaterloo.ca, kctung@uwaterloo.ca, r585wang@uwaterloo.ca

CP6

Factoring Pattern-Free Permutations into Separable Ones

We show that for any permutation π there exists an integer k_π such that every permutation avoiding π as a pattern is a product of at most k_π separable permutations. In other words, every strict class $\mathcal{C}$ of permutations is contained in a bounded power of the class of separable permutations. This factorisation can be computed in linear time, for any fixed π. The central tool for our result is a notion of width of permutations, introduced by Guillemot and Marx [SODA '14] to efficiently detect patterns, and later generalised to graphs and matrices under the name of twin-width. Specifically, our factorisation is inspired by the decomposition used in the recent result that graphs with bounded twin-width are polynomially χ-bounded. As an application, we show that there is a fixed class $\mathcal{C}$ of graphs of bounded twin-width such that every class of bounded twin-width is a first-order transduction of $\mathcal{C}$.

Édouard Bonnet
Univ Lyon, CNRS, ENS de Lyon, France
edouard.bonnet@ens-lyon.fr

Romain Bourneuf
LIP, ENS de Lyon
France
romain.bourneuf@ens-lyon.fr

Colin Geniet
ENS Paris-Saclay
France
colin.geniet@ens-lyon.fr

Stéphan Thomassé
ENS de Lyon
stephan.thomasse@ens-lyon.fr

CP6

Parameterized Algorithms for Block-Structured Integer Programs with Large Entries

We study two classic variants of block-structured integer programming. Two-stage stochastic programs are integer programs of the form $\{A_i\mathbf{x} + D_i\mathbf{y}_i = \mathbf{b}_i$ for all $i = 1, \ldots, n\}$, where A_i and D_i are bounded-size matrices. Next, n-fold programs are integer programs of the form $\{\sum_{i=1}^{n} C_i\mathbf{y}_i = \mathbf{a}$ and $D_i\mathbf{y}_i = \mathbf{b}_i$ for all $i = 1, \ldots, n\}$, where again C_i and D_i are bounded-size matrices. We prove that the parameterized tractability results for two-stage stochastic and n-fold programs persist even when one allows large entries in the global part of the program: - The feasibility problem for two-stage stochastic programs is fixed-parameter tractable when parameterized by the dimensions of matrices A_i, D_i and by the maximum absolute value of the entries of matrices D_i. That is, we allow matrices A_i to have arbitrarily large entries. - The linear optimization problem for n-fold integer programs that are uniform – all matrices C_i are equal – is fixed-parameter tractable when parameterized by the dimensions of matrices C_i and D_i and by the maximum absolute value of the entries of matrices D_i. That is, we require that $C_i = C$ for all $i = 1, \ldots, n$, but we allow C to have arbitrarily large entries.

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

Jana Cslovjecsek
EPFL, Lausanne, Switzerland
cslovj@gmail.com

Martin Kouteck
Charles University, Prague, Czech Republic
koutecky@iuuk.mff.cuni.cz

Alexandra Lassota
TU Eindhoven, Netherlands
alexandra.lassota@web.de

Adam Polak
Bocconi University, Milan, Italy
pl.adam.polak@gmail.com

CP6

Meta-Theorems for Parameterized Streaming Algorithms

In this paper, we study semi-streaming algorithms for parameterized graph problems, and present the first systematic study of this topic. Crucially, we aim to construct succinct representations of the input on which optimal post-processing time complexity can be achieved. 1. We devise meta-theorems specifically designed for parameterized streaming and demonstrate their applicability by obtaining the first $k^{O(1)}n \cdot \mathit{polylog}(n)$-space streaming algorithms for well-studied problems such as Feedback Vertex Set on Tournaments, Cluster Vertex Deletion, Proper Interval Vertex Deletion and Block Vertex Deletion. In the process, we demonstrate a fundamental connection between semi-streaming algorithms for recognizing graphs in a graph class $\mathcal{H}$ and semi-streaming algorithms for the problem of vertex deletion into $\mathcal{H}$. 2. We present an algorithmic machinery for obtaining streaming algorithms for cut problems and exemplify this by giving the first $k^{O(1)}n \cdot \mathit{polylog}(n)$-space streaming algorithms for Graph Bipartitization, Multiway Cut and Subset Feedback Vertex Set.

Ramanujan M. Sridharan
TU Wien
R.Maadapuzhi-Sridharan@warwick.ac.uk

Daniel Lokshtanov
UCSB
daniello@ucsb.edu

Pranabendu Misra
Chennai Mathematical Institute
pranabendu@cmi.ac.in

Fahad Panolan
University of Leeds
f.panolan@leeds.ac.uk

Saket Saurabh
IMSc +UiB
saket@imsc.res.in

Meirav Zehavi
BGU
meiravze@bgu.ac.il

CP6

Representative Set Statements for Delta-Matroids and the Mader Delta-Matroid

We show representative sets statements for linear delta-matroids, generalizing the representative sets lemma for linear matroids. Delta-matroids are generalizations of matroids with important connections to matching theory and graph embeddings. For applications, we show a linear representation for the class of Mader delta-matroids, which should be of independent interest, and use this to construct $O(k^3)$-vertex "Mader-mimicking networks" for graphs with k terminals. This generalizes the powerful cut-covering lemma to path-packing properties for terminal partitions with more than two parts. The representative sets lemma for matroids has powerful applications in parameterized complexity, including improved FPT dynamic programming algorithms, and polynomial kernels for Almost 2-SAT and restrictions of Multiway Cut. However, its usage is sporadic, since it only applies to properties encoded into bounded-rank linear matroids. Thus, despite its success, many problems are left open including polynomial kernels for the general case of Multiway Cut. Our proof uses a new approach of sieving polynomial families, generalizing the linear algebra approach of previous work to a setting of bounded-degree polynomials. The results for linear matroids and linear delta-matroids then follow by applying this to the determinant, respectively the Pfaffian of the corresponding matrix representation. This significantly increases the toolbox for polynomial kernelization.

Magnus Wahlström
Royal Holloway, University of London

Magnus.Wahlstrom@rhul.ac.uk

CP6

Euclidean Bottleneck Steiner Tree is Fixed-Parameter Tractable

In the Euclidean Bottleneck Steiner Tree problem, the input consists of a set of n points in $\mathbb{R}^2$ called terminals and a parameter k, and the goal is to compute a Steiner tree that spans all the terminals and contains at most k points of $\mathbb{R}^2$ as Steiner points such that the maximum edge-length of the Steiner tree is minimized, where the length of a tree edge is the Euclidean distance between its two endpoints. The problem is well-studied and is known to be NP-hard. In this paper, we give a $k^{O(k)}n^{O(1)}$-time algorithm for Euclidean Bottleneck Steiner Tree, which implies that the problem is fixed-parameter tractable (FPT). This settles an open question explicitly asked by Bae et al. [Algorithmica, 2011], who showed that the ℓ_1 and ℓ_∞ variants of the problem are FPT. Our approach can be generalized to the problem with ℓ_p metric for any $1 \leq p \leq \infty$, or even other metrics on $\mathbb{R}^2$.

Jie Xue
New York University Shanghai
jiexue@nyu.edu

Sayan Bandyapadhyay
Portland State University
sayanb@pdx.edu

William Lochet
University of Bergen
william.lochet@gmail.com

Daniel Lokshtanov
UCSB
daniello@ucsb.edu

Saket Saurabh
IMSc +UiB
saket@imsc.res.in

CP7

Matrix Perturbation: Davis-Kahan in the Infinity Norm

Perturbation theory is developed to analyze the impact of noise on data and has been an essential part of numerical analysis. Recently, it has played an important role in designing and analyzing matrix algorithms. One of the most useful tools in this subject, the Davis-Kahan sine theorem, provides an ℓ_2 error bound on the perturbation of the leading singular vectors (and spaces). We focus on the case when the signal matrix has low rank and the perturbation is random, which occurs often in practice. In an earlier paper, O'Rourke, Wang, and the second author showed that in this case, the Davis-Kahan bound can be improved significantly. In particular, the noise-to-gap ratio condition in the original setting can be weakened considerably. In the current paper, we develop an infinity norm version of the O'Rourke-Vu-Wang result. The key ideas in the proof are a new bootstrapping argument and the so-called iterative leave-one-out method, which may be of independent interest. Applying the new bounds, we develop new and very simple algorithms for several well-known problems, such as finding hidden partitions and matrix completion. The core of these new algorithms is the fact that one is now able to quickly approximate certain key objects in the infinity norm, which has critical advantages over approximations in the ℓ_2 norm, Frobenius norm, or spectral norm.

Abhinav Bhardwaj, Van Vu
Yale University
abhinav.bhardwaj@yale.edu, van.vu@yale.edu

CP7

Positivity Certificates for Linear Recurrences

We consider linear recurrences with polynomial coefficients of Poincarpe and with a unique simple dominant eigenvalue. We give an algorithm that proves or disproves positivity of solutions provided the initial conditions satisfy a precisely defined genericity condition. For positive sequences, the algorithm produces a certificate of positivity that is a data-structure for a proof by induction. This induction works by showing that an explicitly computed cone is contracted by the iteration of the recurrence.

Bruno Salvy
INRIA, Paris-Rocquencourt
bruno.salvy@inria.fr

Alaa Ibrahim
ENS de Lyon
France
alaa.ibrahim@inria.fr

CP7

A Ptas for ℓ_0-Low Rank Approximation: Solving Dense CSPS over Reals

We consider the Low Rank Approximation problem, where the input consists of a matrix $A \in R^{n_R \times n_C}$ and an integer k, and the goal is to find a matrix B of rank at most k that minimizes $\|A - B\|_0$, which is the number of entries where A and B differ. For any constant k and $\epsilon > 0$, we present a polynomial time $(1+\epsilon)$-approximation time for this problem, which significantly improves the previous best $poly(k)$-approximation. Our algorithm is obtained by viewing the problem as a Constraint Satisfaction Problem (CSP) where each row and column becomes a variable that can have a value from R^k. In this view, we have a constraint between each row and column, which results in a *dense* CSP, a well-studied topic in approximation algorithms. While most of previous algorithms focus on finite-size (or constant-size) domains and involve an exhaustive enumeration over the entire domain, we present a new framework that bypasses such an enumeration in R^k. On the hardness side, when k is part of the input, we prove that Low Rank Approximation is NP-hard to approximate within a factor of $\Omega(\log n)$. This is the first superconstant NP-hardness of approximation for any $p \in [0, \infty]$ that does not rely on stronger conjectures (e.g., the Small Set Expansion Hypothesis).

Euiwoong Lee
University of Michigan
euiwoong@umich.edu

Vincent Cohen-Addad
Google Research
cohenaddad@google.com

Chenglin Fan
Sorbonne Université
fanchenglin@gmail.com

Suprovat Ghoshal
Northwestern University and TTIC at Chicago, U.S.
suprovat.ghoshal@gmail.com

Arnaud de Mesmay
LIGM, CNRS, Univ. Gustave Eiffel, ESIEE
Paris, F-77454, Marne-la-Vallée, France
arnaud.de-mesmay@univ-eiffel.fr

Alantha Newman
Laboratoire G-SCOP (CNRS, Grenoble-INP)
alantha.newman@gmail.com

Tony Wang
University of Wisconsin-Madison
tonycw@umich.edu

CP7

On the Unreasonable Effectiveness of Single Vector Krylov Methods for Low-Rank Approximation

Krylov subspace methods are a ubiquitous tool for computing near-optimal rank k approximations of large matrices. While "large block" Krylov methods with block size at least k give the best known theoretical guarantees, block size one or a small constant is often preferred in practice. Despite their popularity, we lack theoretical bounds on the performance of such "small block" methods for low-rank approximation. We address this gap between theory and practice by proving that small block methods essentially match all known low-rank approximation guarantees for large block methods. Via a black-box reduction we show that the single vector method run for t iterations obtains the same spectral and Frobenius norm error bounds as a Krylov method with block size $\ell \geq k$ run for $O(t/\ell)$ iterations, up to a logarithmic dependence on the smallest gap between sequential singular values. For a given number of matrix-vector products, single vector methods are essentially as effective as any choice of large block size. By combining our result with tail-bounds on eigenvalue gaps in random matrices, we prove that the dependence on the smallest singular value gap can be eliminated if the input matrix is perturbed by a small random matrix. Further, we show that single vector methods match the more complex algorithm of [Bakshi et al. '22], which combines the results of multiple block sizes to achieve an improved algorithm for Schatten norm low-rank approximation.

Raphael A. Meyer
New York University
ram900@nyu.edu

Cameron Musco
University of Massachusetts Amherst
cmusco@cs.umass.edu

Christopher Musco
New York University
cmusco@nyu.edu

CP7

Detecting Hidden Communities by Power Iterations with Connections to Vanilla Spectral Algorithms

Community detection in the stochastic block model is one of the central problems of graph clustering. Since its introduction by Holland, Laskey, and Leinhardt (Social Networks, 1983), many subsequent papers have made great strides in solving and understanding this model. However, despite the long history of study, there are still unsolved challenges. In this direction, two primary open problems are: how to recover large clusters in the presence of small clusters, and how to analyze simple and practical spectral algorithms (also known as vanilla spectral algorithms), especially when the number of communities is large. In this paper, we use a power iteration approach to make progress in both these directions. To this end, we design the first parameter-free community recovery algorithm that recovers large clusters in the presence of small clusters. Our algorithm only compares the rows of the powered adjacency matrix and has a recovery guarantee poly-logarithmically close to that of the state-of-the-art algorithms that need model parameter knowledge. Then based on a connection between the powered adjacency matrix and eigenvectors, we provide a "vanilla' spectral algorithm in the balanced case when the number of communities is large. This answers an open question by Van Vu (Combinatorics Probability and Computing, 2018) in the balanced case. Our methods also partially solve technical barriers discussed by Abbe, Fan, Wang, and Zhong (Annals of Statistics, 2020).

Chandra Sekhar Mukherjee, Jiapeng Zhang
University of Southern California
chandrasekhar.mukherjee07@gmail.com, jiapengz@usc.edu

CP7

Strongly Polynomial Frame Scaling to High Precision

The frame scaling problem is: given vectors $U := \{u_1, ..., u_n\} \subseteq \mathbb{R}^d$, marginals $c \in \mathbb{R}^n_{++}$, and precision $\epsilon > 0$, find scalings $L \in \mathbb{R}^{d\times d}, r \in \mathbb{R}^n_{++}$ such that $(v_1, \ldots, v_n) := (Lu_1r_1, \ldots, Lu_nr_n)$ simultaneously satisfies $\sum_{i=1}^n v_i v_i^T = I_d$ and $\|v_j\|_2^2 = c_j, \forall j \in [n]$, up to error ϵ. In this work, we give a strongly polynomial algorithm for frame scaling with $\log(1/\epsilon)$ convergence. This answers a question of Diakonikolas, Tzamos and Kane (STOC 2023), who gave the first strongly polynomial randomized algorithm with $\text{poly}(1/\epsilon)$ convergence for Forster transformation ($c = \frac{d}{n}1_n$). Our algorithm is deterministic, applies for all $c \in \mathbb{R}^n_{++}$, and requires $O(n^3 \log(n/\epsilon))$ iterations as compared to the $O(n^5 d^{11}/\epsilon^5)$ iterations of DTK. Our main technical contribution is generalizing the potential analysis of Linial, Samorodnitsky and Wigderson (Combinatorica 2000) to frames. In fact, we can adapt our analysis to reduce the $O(n^5 \log(n/\epsilon))$ iterations of LSW to $O(n^3 \log(n/\epsilon))$. We also prove a new bound on the size of approximate solutions, involving condition measure $\bar{\chi}$ studied in the linear programming literature.

Daniel Dadush
Centrum Wiskunde & Informatica
dadush@cwi.nl

Akshay Ramachandran
Centrum Wiskunde & Informatica (CWI)
Netherlands
akshay.ramachandran@cwi.nl

CP8

Deterministic Byzantine Agreement with Adaptive $O(n \cdot F)$ Communication

In this paper, we present a deterministic protocol for bi-

nary Byzantine Agreement (BA) with adaptive $O(n \cdot f)$ communication complexity against a corrupt minority in the synchronous communication model, where n is the total number of parties and f is the exact number (not an upper bound) of corruptions. Our protocol represents a significant improvement over the previous best-known deterministic BA protocol developed by Momose and Ren (DISC 2021), whose communication complexity is non-adaptive $O(n^2)$. Our proposed approach combines two distinct primitives, Reliable Voting (which we introduce) and Weak Byzantine Agreement (recently introduced by Cohen et al. in OPODIS 2022). In Reliable Voting, all honest parties agree on the same value only if all honest parties start with that value, except in the scenario where all parties are honest. Additionally, we put forth a novel protocol for Weak Byzantine Agreement with $O(n \cdot f)$ communication complexity. While we use the Cohen et al. protocol as our starting point for our new Weak Byzantine Agreement protocol, we identify and address various issues to create our new protocol.

Fatima Elsheimy
Yale
Fatima.Elsheimy@yale.edu

Giorgos Tsimos
University of Maryland
tsimos@umd.edu

Charalampos Papamanthou
Yale University
charalampos.papamanthou@yale.edu

CP8

Shannon Meets Gray: Noise-Robust, Low-Sensitivity Codes with Applications in Differential Privacy

Integer data is typically made differentially private by adding noise from a Discrete Laplace (or Discrete Gaussian) distribution. We study the setting where differential privacy of a counting query is achieved using bit-wise randomized response, i.e., independent, random bit flips on the encoding of the query answer. Binary *error-correcting codes* transmitted through noisy channels with independent bit flips are well-studied in information theory. However, such codes are unsuitable for differential privacy since they have (by design) high sensitivity, i.e., neighboring integers have encodings with a large Hamming distance. *Gray codes* show that it is possible to create an efficient sensitivity 1 encoding, but are also not suitable for differential privacy due to lack of noise-robustness. Our main result is that it is possible, with a constant rate code, to *simultaneously* achieve the sensitivity of Gray codes *and* the noise-robustness of error-correcting codes (down to the noise level required for differential privacy). An application of this new encoding of the integers is a faster, space-optimal differentially private data structure for histograms.

David R. Lolck, Rasmus Pagh
University of Copenhagen
david@lolck.com, pagh@di.ku.dk

CP8

Adjacency Sketches in Adversarial Environments

An adjacency sketching or implicit labeling scheme for a family of graphs is a method that defines for any n-vertex G in the family an assignment of labels to each vertex in G, so that the labels of two vertices tell you whether or not they are adjacent. The goal is to have schemes that use as few bits as possible to represent the labels. By using randomness when assigning labels it is sometimes possible to produce adjacency sketches with much smaller label sizes, at the cost of introducing some probability of error. Both deterministic and randomized labeling schemes have applications for distributed data structures and deeper connections to universal graphs and communication complexity. The main question of interest is which graph families have schemes using short labels, usually $O(\log n)$ in the deterministic case or constant for randomized sketches. We consider the resilience of probabilistic adjacency sketches against an adversary making adaptive queries to the labels. We show that in the adaptive adversarial case the size of the labels is tightly related to the maximal degree of the graph. This results in a stronger characterization compared to what is known in the non-adversarial setting. In more detail, we construct sketches that fail with probability ε for graphs with maximal degree d using $2d \log(1/\varepsilon)$ bit labels and show that this is roughly the best possible for any specific graph of maximal degree d, e.g. a d-ary tree.

Moni Naor, Eugene Pekel
Weizmann Institute of Science
moni.naor@weizmann.ac.il, eugune.pekel@weizmann.ac.il

CP8

Sorting and Selection in Rounds with Adversarial Comparisons

We continue the study of selection and sorting of n numbers under the adversarial comparator model, where comparisons can be adversarially tampered with if the arguments are sufficiently close. We derive a randomized sorting algorithm that does $O(n \log^2 n)$ comparisons and gives a correct answer with high probability, addressing an open problem of Ajtai, Feldman, Hassadim, and Nelson. Our algorithm also implies a selection algorithm that does $O(n \log n)$ comparisons and gives a correct answer with high probability. Both of these results are a log factor away from the naive lower bound. In work by Ajtai, Feldman, Hassidim, and Nelson (2015), they show an $\Omega(n^{1+\varepsilon})$ lower bound for both sorting and selection in the deterministic case, so our results also prove a discrepancy between what is possible with deterministic and randomized algorithms in this setting. We also consider both sorting and selection in rounds, exploring the tradeoff between accuracy, number of comparisons, and number of rounds.

Christopher D. Trevisan
University of Waterloo
chris.trevisan@uwaterloo.ca

CP8

A Unifying Framework for Differentially Private Sums under Continual Observation

We study the problem of maintaining a differentially private decaying sum under continual observation. We give a unifying framework and an efficient algorithm for this problem for *any sufficiently smooth* function. Our algorithm is the first differentially private algorithm that does not have a multiplicative error for polynomially decaying weights. Our algorithm improves on all prior works on differentially private decaying sums under continual observation and recovers exactly the additive error for the special case of con-

tinual counting from Henzinger et al. (SODA 2023) as a corollary. Our algorithm is a variant of the matrix mechanism whose error depends on the γ_2 and γ_F norm of the underlying matrix. We give a constructive proof for an almost exact upper bound on the γ_2 and γ_F norm and an almost tight lower bound on the γ_2 norm for a large class of lower-triangular matrices. This is the first non-trivial lower bound for lower-triangular matrices whose non-zero entries are not all the same. It includes matrices for all continual decaying sums problems, resulting in an upper bound on the additive error of any differentially private decaying sums algorithm under continual observation. We also explore some implications of our result in discrepancy theory and operator algebra.

Monika Henzinger
Institute of Science and Technolology Austria
monika.henzinger@ist.ac.at

Jalaj Upadhyay
Rutgers University
jalaj.upadhyay@rutgers.edu

Sarvagya Upadhyay
Fujitsu Research of America
sarvagya@gmail.com

CP8

Optimal Bounds on Private Graph Approximation

We propose an efficient ε-differentially private (DP) algorithm, that given a simple weighted n-vertex, m-edge graph G with a maximum unweighted degree $\Delta(G)$, outputs a synthetic graph which approximates the spectrum with $\widetilde{O}(\min\{\Delta(G), \sqrt{n}\})$ bound on the purely additive error. To the best of our knowledge, this is the first ε-DP algorithm with a non-trivial additive error for approximating the spectrum of the graph. One of our subroutines also precisely simulates the exponential mechanism over a non-convex set, which could be of independent interest given the recent interest in sampling from a log-concave distribution defined over a convex set. As a direct application of our result, we give the first non-trivial bound on approximating all-pairs effective resistances by a synthetic graph, which also implies approximating hitting/commute time and cover time of random walks on the graph. We further show that using our sampler, we can also output a synthetic graph that approximates the sizes of all (S, T)-cuts on n vertices weighted graph G with m edges while preserving (ε, δ)-DP and an optimal additive error. We also give a matching lower bound (with respect to all the parameters) on the private cut approximation for weighted graphs. This removes the gap in the upper and lower bound with respect to the average edge weight in Eli Kapralov, Kulkarni, and Lee (SODA 2020).

Zongrui Zou, Jingcheng Liu
Nanjing University
zou.zongrui@smail.nju.edu.cn, liu@nju.edu.cn

Jalaj Upadhyay
Rutgers University
jalaj.upadhyay@rutgers.edu

CP9

Partial Coloring Complex, Vertex Decomposability and Tverberg's Theorem with Constraints

We present a novel family of simplicial complexes associated with the graph coloring problem. They include many well-known simplicial complexes such as chessboard complexes and crosspolytopes. We then study conditions under which these complexes become vertex decomposable and hence shellable. The connectivity of these complexes is also investigated. We apply these results to Tverberg's theorem with constraints and also to the chromatic number of certain Kneser-type hypergraphs and improve upon existing facts. Notably, we prove a conjecture of Engstrom and Noren on Tverberg graphs.

Sharareh Alipour
Tehran institute for advanced studies, khatam university
sharareh.alipour@gmail.com

Amir Jafari, Mohammad Hassan Mazidi, Seyed Abolfazl Najafian
Sharif university of technology
Department of mathematical sciences
amirjafa@gmail.com, mhmazidi4@gmail.com, a.najafian74@gmail.com

CP9

Small But Unwieldy: A Lower Bound on Adjacency Labels for Small Classes

We show that for any natural number s, there is a constant γ and a subgraph-closed class having, for any natural n, at most γ^n graphs on n vertices up to isomorphism, but no adjacency labeling scheme with labels of size at most $s \log n$. In other words, for every s, there is a small monotone class without universal graphs of size n^s. Prior to this result, it was not excluded that every small class has an almost linear universal graph, or equivalently a labeling scheme with labels of size $(1 + o(1)) \log n$. The existence of such a labeling scheme, a scaled-down version of the recently disproved Implicit Graph Conjecture, was repeatedly raised [Gavoille and Labourel, ESA '07; Dujmović et al., JACM '21; Bonamy et al., SIDMA '22; Bonnet et al., Comb. Theory '22]. Furthermore, our small monotone classes have unbounded twin-width, thus simultaneously disprove the already-refuted Small conjecture; but this time with a self-contained proof, not relying on elaborate group-theoretic constructions. As our main ingredient, we show that with high probability an Erdős–Rényi random graph $G(n, p)$ with $p = O(1/n)$ has, for every kn, at most $2^{O(k)}$ subgraphs on k vertices, up to isomorphism. As a barrier to our general method of producing even more complex tiny classes, we show that when $p = \omega(1/n)$, the latter no longer holds.

Édouard Bonnet
Univ Lyon, CNRS, ENS de Lyon, France
edouard.bonnet@ens-lyon.fr

Julien Duron
LIP - ENS de Lyon
julien.duron@ens-lyon.fr

John Sylvester, Viktor Zamaraev
Department of Computer Science, University of Liverpoo
john.sylvester, viktor.zamaraev@liverpool.ac.uk

Maksim Zhukovskii
Department of Computer Science, University of Sheffield

m.zhukovskii@sheffield.ac.uk

CP9

Random Embeddings of Graphs: The Expected Number of Faces in Most Graphs Is Logarithmic

A random 2-cell embedding of a connected graph G in some orientable surface is obtained by choosing a random local rotation around each vertex. Under this setup, the number of faces or the genus of the corresponding 2-cell embedding becomes a random variable. Random embeddings of two particular graph classes, those of a bouquet of n loops and those of n parallel edges connecting two vertices, have been extensively studied and are well-understood. However, little is known about more general graphs. The results of this paper explain why Monte Carlo methods cannot work for approximating the minimum genus of graphs. In his breakthrough work [Permutation-partition pairs, JCTB 1991], Stahl developed the foundation of "random topological graph theory". Most of his results have been unsurpassed until today. In our work, we analyze the expected number of faces of random embeddings (equivalently, the average genus) of a graph G. It was very recently shown that for any graph G, the expected number of faces is at most linear. We show that the actual expected number of faces $F(G)$ is almost always much smaller. In particular, we prove: 1) $\frac{1}{2}\ln n - 2 < \mathbb{E}[F(K_n)] \leq 3.65 \ln n + o(1)$. 2) For random graphs $G(n,p)$ ($p = p(n)$), we have $\mathbb{E}[F(G(n,p))] \leq \ln^2 n + \frac{1}{p}$. 3) For random models $B(n,\Delta)$ containing only graphs, whose maximum degree is at most Δ, we get stronger bounds $\mathbb{E}[F(B(n,\Delta))] = \Theta(\ln n)$.

Jesse Campion Loth, Kevin C. Halasz
Simon Fraser University
jesse_campion_loth@sfu.ca, khalasz@sfu.ca

Tomáš Masaríik
University of Warsaw
masarik@mimuw.edu.pl

Bojan Mohar
Simon Fraser University
mohar@sfu.ca

Robert Šámal
IÚUK MFF UK – CSI of Charles University in Prague
samal@iuuk.mff.cuni.cz

CP9

The Grid-Minor Theorem Revisited

We prove that for every planar graph X of treedepth h, there exists a positive integer c such that for every X-minor-free graph G, there exists a graph H of treewidth at most $f(h)$ such that G is isomorphic to a subgraph of HK_c. This is a qualitative strengthening of the Grid-Minor Theorem of Robertson and Seymour (JCTB, 1986), and treedepth is the optimal parameter in such a result. As an example application, we use this result to improve the upper bound for weak coloring numbers of graphs excluding a given graph as a minor.

Pat Morin
Carleton University
morin@scs.carleton.ca

Vida Dujmovic
University of Ottawa
vida.dujmovic@uottawa.ca

Robert Hickingbotham
Monash University
robert.hickingbotham@monash.edu

Jedrzej Hodor
Jagiellonian University
jedrzej.hodor@gmail.com

Gwenael Joret
Universite libre de Bruxelles
gwenael.joret@ulb.be

Hoang La
Jagiellonian University
hoang.la.research@gmail.com

Piotr Micek
Jagiellonian University, Kraków, Poland
piotr.micek@uj.edu.pl

Clement Rambaud
Ecole Normale Superieure
clement.rambaud@inria.fr

David Wood
Monash University, Australia
david.wood@monash.edu

CP9

Triangulations Admit Dominating Sets of Size $2n/7$.

We show that every planar triangulation on $n > 10$ vertices has a dominating set of size $2n/7 = n/3.5$. This approaches the $n/4$ bound conjectured by Matheson and Tarjan, and improves significantly on the previous best bound of $17n/53 \approx n/3.117$ by pacapan. From our proof it follows that every 3-connected n-vertex near-triangulation (except for 3 sporadic examples) has a dominating set of size $n/3.5$. On the other hand, for 3-connected near-triangulations, we show a lower bound of $3(n-1)/11 \approx n/3.666$, demonstrating that the conjecture by Matheson and Tarjan cannot be strengthened to 3-connected near-triangulations. Our proof uses a penalty function that, aside from the number of vertices, penalises vertices of degree 2 and specific constellations of neighbours of degree 3 along the boundary of the outer face. To facilitate induction, we not only consider near-triangulations, but a wider class of graphs (skeletal triangulations), allowing us to delete vertices more freely. Our main technical contribution is a set of attachments, that are small graphs we inductively attach to our graph, in order both to remember whether existing vertices are already dominated, and that serve as a tool in a divide and conquer approach. We complement our proof with a constructive algorithm that returns a dominating set of size $\leq 2n/7$. Our algorithm has a quadratic running time.

Aleksander B. Christiansen, Eva Rotenberg
Technical University of Denmark
abgch@dtu.dk, erot@dtu.dk

Daniel P. Rutschmann
DTU Compute, Technical University of Denmark

daru@dtu.dk

CP9

On the Extremal Functions of Acyclic Forbidden 0–1 Matrices

The extremal theory of forbidden 0–1 matrices studies the asymptotic growth of the function $\mathrm{Ex}(P,n)$, which is the maximum weight of a matrix $A \in \{0,1\}^{n\times n}$ whose submatrices avoid a fixed pattern $P \in \{0,1\}^{k\times l}$. This theory has been wildly successful at resolving problems in combinatorics, discrete and computational geometry, structural graph theory, and the analysis of data structures, particularly corollaries of the dynamic optimality conjecture. All these applications use acyclic patterns, meaning that when P is regarded as the adjacency matrix of a bipartite graph, the graph is acyclic. The biggest open problem in this area is to bound $\mathrm{Ex}(P,n)$ for acyclic P. Prior results have only ruled out the strict $O(n\log n)$ bound conjectured by Furedi and Hajnal. At the two extremes, it is consistent with prior results that $\forall P.\mathrm{Ex}(P,n) \le n\log^{1+o(1)} n$, and also consistent that $\forall \epsilon > 0.\exists P.\mathrm{Ex}(P,n) \ge n^{2-\epsilon}$. In this paper we establish a stronger lower bound on the extremal functions of acyclic P. Specifically, for any $t \ge 1$ we give a new construction of relatively dense 0–1 matrices with $\Theta(n(\log n/\log\log n)^t)$ 1s that avoid a certain acyclic pattern X_t. Pach and Tardos have conjectured that this type of result is the best possible, i.e., no acyclic P exists for which $\mathrm{Ex}(P,n) \ge n(\log n)^{\omega(1)}$.

Seth Pettie
University of Michigan, Ann Arbor
seth@pettie.net

Gabor Tardos
Simon Fraser University, Canada and Renyi Institute, Hungary
tardos@renyi.hu

CP10

Approximating Subset Sum Ratio Faster Than Subset Sum

Subset Sum Ratio is the following optimization problem: Given a set of n positive numbers I, find disjoint subsets $X, Y \subseteq I$ minimizing the ratio $\max\{\Sigma(X)/\Sigma(Y), \Sigma(Y)/\Sigma(X)\}$, where $\Sigma(Z)$ denotes the sum of all elements of Z. Subset Sum Ratio is an optimization variant of the Equal Subset Sum problem. It was introduced by Woeginger and Yu in '92 and is known to admit an FPTAS [Bazgan, Santha, Tuza '98]. The best approximation schemes before this work had running time $O(n^4/\varepsilon)$ [Melissinos, Pagourtzis '18], $\widetilde{O}(n^{2.3}/\varepsilon^{2.6})$ and $\widetilde{O}(n^2/\varepsilon^3)$ [Alonistiotis et al. '22]. In this work, we present an improved approximation scheme for Subset Sum Ratio running in time $O(n/\varepsilon^{0.9386})$. Here we assume that the items are given in sorted order, otherwise we need an additional running time of $O(n\log n)$ for sorting. Our improved running time simultaneously improves the dependence on n to linear and the dependence on $1/\varepsilon$ to sublinear. For comparison, an approximation scheme achieving the same running time for Subset Sum would falsify the Strong Exponential Time Hypothesis [Abboud, Bringmann, Hermelin, Shabtay '19] as well as the Min-Plus-Convolution Hypothesis [Bringmann, Nakos '21]. We thus establish that Subset Sum Ratio admits faster approximation schemes than Subset Sum.

Karl Bringmann
Saarland University and Max Planck Institute for Informatics
Saarland Informatics Campus, Saarbrücken, Germany
kbringma@mpi-inf.mpg.de

CP10

A Parameterized Family of Meta-Submodular Functions

Submodular function maximization has found a wealth of new applications in recent years. The related supermodular maximization models also offer an abundance of applications, but they appeared to be highly intractable even under simple cardinality constraints. Hence, while there are well-developed tools for maximizing a submodular function subject to a matroid constraint, there is much less work on the corresponding supermodular maximization problems. We introduce a parameterized family of monotone functions which contains submodular functions and a broad spectrum of supermodular functions including the well-studied diversity functions. Functions in this parameterized family are called γ-meta-submodular. We develop approximation algorithms for maximization of these functions under a general matroid constraint. The approximation factors we establish only depend on the parameter γ. We show that the γ-meta-submodular families model previously studied applications arising from classes of functions such as meta-submodular functions ($\gamma = 0$), metric diversity functions and proportionally submodular functions (both with $\gamma = 1$), diversity functions based on negative-type distances or Jensen-Shannon divergence (both with $\gamma = 2$), and σ-semi metric diversity functions ($\gamma = \sigma$). Importantly, meta-submodular functions are not restricted to pairwise structures and are closed under composition with non-negative concave functions.

Mehrdad Ghadiri
MIT
mehrdadg@mit.edu

Richard Santiago
ETH Zurich
richard.santiago@ifor.math.ethz.ch

Bruce Shepherd
University of British Columbia
fbrucesh@cs.ubc.ca

CP10

Tight Approximability for MAX 2-SAT and Relatives, Under UGC

Austrin showed that the approximation ratio $\beta \simeq 0.94016567$ obtained by the MAX 2-SAT approximation algorithm of Lewin, Livnat and Zwick (LLZ) is optimal modulo the Unique Games Conjecture (UGC) and modulo a Simplicity Conjecture that states that the worst performance of the algorithm is obtained on so called *simple* configurations. We prove Austrin's conjecture, thereby showing the optimality of the LLZ approximation algorithm, relying only on the Unique Games Conjecture. Our proof uses a combination of analytic and computational tools. We also present new approximation algorithms for two restrictions of the MAX 2-SAT problem. For MAX HORN-$\{1,2\}$-SAT, i.e., MAX CSP($\{x\vee y, \bar{x}\vee y, x, \bar{x}\}$), in which clauses are not allowed to contain two negated literals, we

obtain an approximation ratio of 0.94615981. For MAX CSP($\{x \vee y, x, \bar{x}\}$), i.e., when 2-clauses are not allowed to contain negated literals, we obtain an approximation ratio of 0.95397990. By adapting Austrin's and our arguments for the MAX 2-SAT problem we show that these two approximation ratios are also tight, modulo only the UGC conjecture. This completes a full characterization of the approximability of the MAX 2-SAT problem and its restrictions.

Joshua Brakensiek
Stanford University
jbrakens@stanford.edu

Neng Huang
University of Chicago
nenghuang@uchicago.edu

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

CP10

Approximation Algorithms for the Weighted Nash Social Welfare via Convex and Non-Convex Programs

In an instance of the weighted Nash Social Welfare problem, we are given a set of m indivisible items, G, and n agents, A, where each agent $i \in A$ has a valuation $v_{ij} \geq 0$ for each item $j \in G$. In addition, every agent i has a non-negative weight w_i such that the weights collectively sum up to 1. The goal is to find an assignment $\sigma : G \to A$ that maximizes $\prod_{i \in A} \left(\sum_{j \in \sigma^{-1}(i)} v_{ij} \right)^{w_i}$. When all the weights equal to $\frac{1}{n}$, the problem reduces to the classical Nash Social Welfare problem, which has recently received much attention. In this work, we present a $5 \cdot \exp\left(2 \cdot D_{\mathrm{KL}}(w \,||\, \frac{\vec{1}}{n})\right) = 5 \cdot \exp\left(2 \log n + 2 \sum_{i=1}^{n} w_i \log w_i\right)$-approximation algorithm for the weighted Nash Social Welfare problem, where $D_{\mathrm{KL}}(w \,||\, \frac{\vec{1}}{n})$ denotes the KL-divergence between the distribution w and the uniform distribution on $[n]$. We generalize the convex programming relaxations for the symmetric variant of Nash Social Welfare introduced by Cole et al. 2017 and Anari et al. 2017 to two different mathematical programs, one convex, and the second non-convex. The approximation factor derives from the difference in the objective values of the convex and non-convex relaxation.

Aditi Laddha
Yale University
aditi.laddha@yale.edu

Adam Brown
Georgia Institute of Technology
ajmbrown@gatech.edu

Madhusudhan Reddy Pittu
Carnegie Mellon University
mpittu@andrew.cmu.edu

Mohit Singh
Georgia Institute of Technology
msingh94@gatech.edu

CP10

Gap Amplification for Reconfiguration Problems

In this paper, we demonstrate gap amplification for reconfiguration problems. In particular, we prove an explicit factor of PSPACE-hardness of approximation for three popular reconfiguration problems only assuming Reconfiguration Inapproximability Hypothesis (RIH) due to Ohsaka (STACS 2023). Our main result is that under RIH, Maxmin Binary CSP Reconfiguration is PSPACE-hard to approximate within a factor of 0.9942. Moreover, the same result holds even if the constraint graph is restricted to (d, λ)-expander for arbitrarily small $\frac{\lambda}{d}$. The crux of its proof is an alteration of the gap amplification technique due to Dinur (J. ACM, 2007), which amplifies the 1 vs. $1-\epsilon$ gap for arbitrarily small $\epsilon > 0$ up to the 1 vs. $1 - 0.0058$ gap. As an application of the main result, we demonstrate that Minmax Set Cover Reconfiguration and Minmax Dominating Set Reconfiguration are PSPACE-hard to approximate within a factor of 1.0029 under RIH. Our proof is based on a gap-preserving reduction from Label Cover to Set Cover due to Lund and Yannakakis (J. ACM, 1994). However, unlike Lund–Yannakakis' reduction, the expander mixing lemma is essential to use. We finally complement the main result by showing that it is NP-hard to approximate Maxmin Binary CSP Reconfiguration within a factor better than $\frac{3}{4}$.

Naoto Ohsaka
NEC Corporation
naoto.ohsaka@gmail.com

CP11

AG Codes Have No List-Decoding Friends: Approaching the Generalized Singleton Bound Requires Exponential Alphabets

A simple, recently observed generalization of the classical Singleton bound to list-decoding asserts that rate R codes are not list-decodable using list-size L beyond an error fraction $\frac{L}{L+1}(1-R)$ (the Singleton bound being the case of $L = 1$, i.e., unique decoding). We prove that in order to approach this bound for any fixed $L > 1$, one needs exponential alphabets. Specifically, for every $L > 1$ and $R \in (0, 1)$, if a rate R code can be list-of-L decoded up to error fraction $\frac{L}{L+1}(1 - R - \varepsilon)$, then its alphabet must have size at least $\exp(\Omega_{L,R}(1/\varepsilon))$. This is in sharp contrast to the situation for unique decoding where certain families of rate R algebraic-geometry (AG) codes over an alphabet of size $O(1/\varepsilon^2)$ are unique-decodable up to error fraction $(1 - R - \varepsilon)/2$. Our lower bound is tight up to constant factors in the exponent–with high probability random codes (or, as shown recently, even random linear codes) over $\exp(O_L(1/\varepsilon))$-sized alphabets, can be list-of-L decoded up to error fraction $\frac{L}{L+1}(1 - R - \varepsilon)$.

Omar Alrabiah, Venkatesan Guruswami, Ray Li
UC Berkeley
oalrabiah@berkeley.edu, venkatg@berkeley.edu, rayyli@berkeley.edu

CP11

Optimal Thresholds for Latin Squares, Steiner Triple Systems, and Edge Colorings

Given a graph G, a random (k, n)-list assignment L for

edges of G is an assignment of an independent, uniformly random set $L(e) \in \binom{[n]}{k}$ of colors to each edge e. We show that for a random $(O(\log n), n)$-list assignment L for edges of the complete bipartite graph $K_{n,n}$, there is a proper L-list coloring of $K_{n,n}$ with high probability. We also prove analogous results for the thresholds of Steiner triple systems and Latin squares in random (binomial) hypergraphs. All of our results are optimal up to absolute constants, and resolve several related conjectures of Johansson, Luria-Simkin, Casselgren-Haggkvist, Simkin, and Kang-Kelly-Kuhn-Methuku-Osthus. A key contribution of our work is to show that in natural settings, the Lovasz Local Lemma – a central tool in probabilistic combinatorics to establish the existence of objects with desired properties – can also be used to design optimally "spread' distributions on such objects. This is made possible by carefully exploiting the local uniformity property of the so-called Lovasz Local Lemma distribution, an important observation that has recently been utilized in finding efficient algorithms for sampling approximately uniformly random solutions to constraint satisfaction problems. In conjunction with the recently proved Kahn-Kalai conjecture, this opens the door to obtaining optimal threshold results for the appearance of many interesting objects.

Vishesh Jain
University of Illinois Chicago
visheshj@uic.edu

Huy T. Pham
Stanford University
huypham@stanford.edu

CP11

Conflict Checkable and Decodable Codes and Their Applications

Let C be an error-correcting code over a large alphabet q of block length n. Assume that, a possibly corrupted, codeword c is distributed to n servers such that the ith entry is being held by the ith server, and that every pair of servers publicly announce whether the corresponding entries are "consistent' with some legal codeword or "conflicted'. What type of information about c can be inferred from this consistency graph? Can we check whether errors occurred and if so, can we find the error locations and effectively decode? We initiate the study of conflict-checkable and conflict-decodable codes and prove the following main results: - For every distance $d \leq n$, there exists a code that supports conflict-based error-detection whose dimension k almost achieves the singleton bound, i.e., $k \geq n - d + 0.99$. This yields an n-partite graph over $[q]^n$ that contains q^k cliques of size n whose pairwise intersection is at most $n - d \leq k - 0.99$ vertices. - For every distance $d \leq n$ there exists a linear code that supports conflict-based error-decoding up to half of the distance and satisfies $k = (n - d + 2)/2$. We extend our result to the byzantine setting at the expense of decreasing the dimension to $k = (n - d + 3)/3$, and use the resulting code to construct the first polynomial-time statistical three-round general MPC protocol that remains secure in the presence of an active adversary that corrupts up to $t < n/3.001$ of the parties.

Benny Applebaum, Eliran Kachlon
Tel Aviv University
benny.applebaum@gmail.com, elirn.chalon@gmail.com

CP11

The Hierarchy of Hereditary Sorting Operators

We consider the following general model of a sorting procedure: we fix a hereditary permutation class C, which corresponds to the operations that the procedure is allowed to perform in a single step. The input of sorting is a permutation π of the set $[n] = \{1, 2, \ldots, n\}$, i.e., a sequence where each element of $[n]$ appears once. In every step, the sorting procedure picks a permutation σ of length n from C, and rearranges the current permutation of numbers by composing it with σ. The goal is to transform the input π into the sorted sequence $1, 2, \ldots, n$ in as few steps as possible. This model of sorting captures not only classical sorting algorithms, like insertion sort or bubble sort, but also sorting by series of devices, like stacks or parallel queues, as well as sorting by block operations commonly considered, e.g., in the context of genome rearrangement. Our goal is to describe the possible asymptotic behavior of the worst-case number of steps needed when sorting with a hereditary permutation class. As the main result, we show that any hereditary permutation class C falls into one of five distinct categories. Disregarding trivial extreme cases, a worst-case sorting time of a hereditary class is either $\Theta(n^2)$, or a function between $O(n)$ and $\Omega(\sqrt{n})$, or $O(\log^2 n)$, and for each of these cases we provide a structural characterization of the corresponding hereditary classes.

Vít Jelínek
Charles University in Prague
jelinek@iuuk.mff.cuni.cz

Michal Opler
Czech Technical University in Prague
michal.opler@fit.cvut.cz

Jakub Pekárek
Charles University
pekarej@kam.mff.cuni.cz

CP11

Cliquewidth and Dimension

We prove that every poset with bounded cliquewidth and with sufficiently large dimension contains the standard example of dimension k as a subposet. This applies in particular to posets whose cover graphs have bounded treewidth, as the cliquewidth of a poset is bounded in terms of the treewidth of the cover graph. For the latter posets, we prove a stronger statement: every such poset with sufficiently large dimension contains the Kelly example of dimension k as a subposet. Using this result, we obtain a full characterization of the minor-closed graph classes $\mathcal{C}$ such that posets with cover graphs in $\mathcal{C}$ have bounded dimension: they are exactly the classes excluding the cover graph of some Kelly example. Finally, we consider a variant of poset dimension called Boolean dimension, and we prove that posets with bounded cliquewidth have bounded Boolean dimension. The proofs rely on Colcombet's deterministic version of Simon's factorization theorem, which is a fundamental tool in formal language and automata theory, and which we believe deserves a wider recognition in structural and algorithmic graph theory.

Gwenaël Joret
Université libre de Bruxelles, Belgium
gwenael.joret@ulb.be

Piotr Micek
Jagiellonian University, Kraków, Poland
piotr.micek@uj.edu.pl

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

Bartosz Walczak
Jagiellonian University, Kraków, Poland
bartosz.walczak@uj.edu.pl

CP12

Beyond the Quadratic Time Barrier for Network Unreliability

Karger (STOC 1995) gave the first FPTAS for the network (un)reliability problem, setting in motion research over the next three decades that obtained increasingly faster running times, eventually leading to a $\tilde{O}(n^2)$-time algorithm (Karger, STOC 2020). This represented a natural culmination of this line of work because the algorithmic techniques used can enumerate $\Theta(n^2)$ (near)-minimum cuts. In this paper, we go beyond this quadratic barrier and obtain a faster FPTAS for the network unreliability problem. Our algorithm runs in $m^{1+o(1)}+\tilde{O}(n^{1.5})$ time. Our main contribution is a new estimator for network unreliability in very reliable graphs. These graphs are usually the bottleneck for network unreliability since the disconnection event is elusive. Our estimator is obtained by defining an appropriate importance sampling subroutine on a dual spanning tree packing of the graph. To complement this estimator for very reliable graphs, we use recursive contraction for moderately reliable graphs. We show that an interleaving of sparsification and contraction can be used to obtain a better parametrization of the recursive contraction algorithm that yields a faster running time matching the one obtained for the very reliable case.

Ruoxu Cen
Duke University
ruoxu.cen@duke.edu

William He
Computer Science Department, Carnegie Mellon Universit
wrhe@andrew.cmu.edu

Jason Li
Simons Institute, UC Berkeley
jmli@alumni.cmu.edu

Debmalya Panigrahi
Duke University
debmalya@cs.duke.edu

CP12

Cactus Representations in Polylogarithmic Max-Flow via Maximal Isolating Mincuts

A cactus representation is an edge sparsifier of $O(n)$ size that captures all global minimum cuts of the graph [DKL76]. It is a key ingredient for the connectivity augmentation problems and for maintaining mincuts under edge insertions (e.g. [Naor97], [CLP22], [Henzinger95]). This sparsifier was generalized to Steiner cactus for a vertex set T, which can be seen as a vertex sparsifier of $O(|T|)$ size that captures all partitions of T corresponding to a T-Steiner minimum cut, and also hypercactus, an analogous concept in hypergraphs. A near-linear time construction of cactus was shown by [KP09], but their technique based on tree packing inherently does not generalize. The state-of-the-art algorithms for Steiner cactus and hypercactus are still slower than linear time by a factor of $O(|T|)$ [DV94] and $O(n)$ [CX17], respectively. We give the first almost-linear time algorithms for constructing Steiner cactus and hypercactus using polylogarithmic calls to max flow. The constructions imply almost-linear-time connectivity augmentation algorithms in the Steiner and hypergraph settings, as well as speed up the incremental algorithm for maintaining mincuts in hypergraphs by a factor of n. The key technique behind our result is a novel variant of the influential isolating mincut technique [LP20, AKL+21] which we called maximal isolating mincuts. This technique makes the isolating mincuts to be more balanced which we believe is useful in future applications.

Zhongtian He
Princeton University
hezhongtian1@gmail.com

Shang-En Huang
Boston College
huangaul@bc.edu

Thatchaphol Saranaurak
University of Michigan
thsa@umich.edu

CP12

Cactus Representation of Minimum Cuts

Given an undirected weighted graph with n vertices and m edges, we give the first deterministic $m^{1+o(1)}$-time algorithm for constructing the cactus representation of *all* global minimum cuts. This improves the current $n^{2+o(1)}$-time state-of-the-art deterministic algorithm, which can be obtained by combining ideas implicitly from three papers [Karger 2000, Li 2021, and Gabow 2016] The known explicitly stated deterministic algorithm has a runtime of $\tilde{O}(mn)$ [Fleischer 1999, Nagamochi and Nakao 2000] Using our technique, we can even speed up the fastest randomized algorithm of [Karger and Panigrahi 2009] whose running time is at least $\Omega(m \log^4 n)$ to $O(m \log^3 n)$.

Zhongtian He
Princeton University
hezhongtian1@gmail.com

Shang-En Huang
Boston College
huangaul@bc.edu

Thatchaphol Saranurak
University of Michigan
thsa@umich.edu

CP12

Max s,t-Flow Oracles and Negative Cycle Detection in Planar Digraphs

We study the maximum s,t-flow oracle problem on planar directed graphs where the goal is to design a data structure answering max s,t-flow value queries for arbitrary source-target pairs (s,t). For the case of polynomially bounded integer edge capacities, we describe an exact max s,t-flow oracle with truly subquadratic space and preprocessing,

and sublinear query time. Moreover, if $(1-\epsilon)$-approximate answers are acceptable, we obtain a static oracle with near-linear preprocessing and $\tilde{O}(n^{3/4})$ query time and a dynamic oracle supporting edge capacity updates and queries in $\tilde{O}(n^{6/7})$ worst-case time. To the best of our knowledge, for directed planar graphs, no (approximate) max s,t-flow oracles have been described even in the unweighted case, and only trivial tradeoffs involving either no preprocessing or precomputing all the n^2 possible answers have been known. One key technical tool we develop on the way is a sublinear (in the number of edges) algorithm for finding a negative cycle in so-called dense distance graphs. By plugging it in earlier frameworks, we obtain improved bounds for other fundamental problems on planar digraphs. In particular, we show (1) a deterministic $O(n\log(nC))$ time algorithm for negatively-weighted SSSP in planar digraphs with integer edge weights at least $-C$, and (2) an improved $O(n\log n)$ bound on finding a perfect matching in a bipartite planar graph.

Adam Karczmarz
University of Warsaw
a.karczmarz@mimuw.edu.pl

CP12

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

Given a large graph G with a subset $|T|=k$ of its vertices called terminals, a quality-q flow sparsifier is a small graph G' that contains T and preserves all multicommodity flows that can be routed between terminals in T, to within factor q. A natural approach of constructing $O(1)$-quality flow sparsifiers, which was adopted in most previous constructions, is contraction. Andoni, Krauthgamer, and Gupta constructed a sketch of size $f(k,\epsilon)$ that stores all feasible multicommodity flows up to a factor of $(1+\epsilon)$, raised the question of constructing quality-$(1+\epsilon)$ flow sparsifiers whose size only depends on k,ϵ (but not the number of vertices in the input graph G), and proposed a contraction-based framework towards it using their sketch result. In this paper, we settle their question for contraction-based flow sparsifiers, by showing that quality-$(1+\epsilon)$ contraction-based flow sparsifiers with size $f(\epsilon)$ exist for all 5-terminal graphs, but not for all 6-terminal graphs. Our hardness result on 6-terminal graphs improves upon a recent hardness result by Krauthgamer and Mosenzon on exact (quality-1) flow sparsifiers, for contraction-based constructions. Our construction and proof utilize the notion of *tight spans* in metric geometry, which we believe is a powerful tool for future work.

Zihan Tan
Rutgers University
zihantan1993@gmail.com

Yu Chen
University of Pennsylvania
yu.chen@epfl.ch

CP13

A $(3+\epsilon)$-Approximation Algorithm For The Minimum Sum Of Radii Problem With Outliers And Extensions For Generalized Lower Bounds

Clustering is a fundamental problem setting with applications in many different areas. For a given set of points in a metric space and an integer k, we seek to partition the given points into k clusters. For each computed cluster, one typically defines one point as the center of the cluster. A natural objective is to minimize the sum of the cluster center's radii, where we assign the smallest radius r to each center such that each point in the cluster is at a distance of at most r from the center. The best-known polynomial time approximation ratio for this problem is 3.389. In the setting with outliers, i.e., we are given an integer m and allow up to m points that are not in any cluster, the best-known approximation factor is 12.365. In this paper, we improve both approximation ratios to $3+\epsilon$. Our algorithms are primal-dual algorithms that use fundamentally new ideas to compute solutions and to guarantee the claimed approximation ratios. For example, we replace the classical binary search to find the best value of a Lagrangian multiplier λ by a primal-dual routine in which λ is a variable that is raised. Also, we show that for each connected component due to almost tight dual constraints, we can find one single cluster that covers all its points and we bound its cost via a new primal-dual analysis. We remark that our approximation factor of $3+\epsilon$ is a natural limit for the known approaches in the literature.

Moritz Buchem, Katja Ettmayr, Hugo Rosado, Andreas Wiese
Technical University of Munich
m.buchem@tum.de, katja.ettmayr@tum.de, hugo.rosado@tum.de, andreas.wiese@tum.de

CP13

On Deterministically Approximating Total Variation Distance

Total variation distance (TV distance) is an important measure for the difference between two distributions. Recently, there has been progress in approximating the TV distance between product distributions: a deterministic algorithm for a restricted class of product distributions (Bhattacharyya, Gayen, Meel, Myrisiotis, Pavan and Vinodchandran 2023) and a randomized algorithm for general product distributions (Feng, Guo, Jerrum and Wang 2023). We give a deterministic fully polynomial-time approximation algorithm (FPTAS) for the TV distance between product distributions. Given two product distributions $\mathbb{P}$ and $\mathbb{Q}$ over $[q]^n$, our algorithm approximates their TV distance with relative error ε in time $O\big(\frac{qn^2}{\varepsilon}\log q\log\frac{n}{\varepsilon\Delta_{\mathrm{TV}}(\mathbb{P},\mathbb{Q})}\big)$. Our algorithm is built around two key concepts: 1) The likelihood ratio as a distribution, which captures sufficient information to compute the TV distance. 2) We introduce a metric between likelihood ratio distributions, called the minimum total variation distance. Our algorithm computes a sparsified likelihood ratio distribution that is close to the original one w.r.t. the new metric. The approximated TV distance can be computed from the sparsified likelihood ratio. Our technique also implies deterministic FPTAS for the TV distance between Markov chains.

Weiming Feng
ETH Zurich
fwm1994@gmail.com

Liqiang Liu, Tianren Liu
Peking University
lql@pku.edu.cn, trl@pku.edu.cn

CP13

On Approximability of Steiner Tree in ℓ_p-Metrics

In the Continuous Steiner Tree problem (CST), we are given as input a set of points (called terminals) in a metric

space and ask for the minimum-cost tree connecting them. Additional points (called Steiner points) from the metric space can be introduced as nodes in the solution. In the Discrete Steiner Tree problem (DST), we are given in addition to the terminals, a set of facilities, and any solution tree connecting the terminals can only contain the Steiner points from this set of facilities. Trevisan [SICOMP'00] showed that CST and DST are APX-hard when the input lies in the ℓ_1-metric (and Hamming metric). Chlebnd ChlebvCS'08] showed that DST is NP-hard to approximate to factor of $96/95 \approx 1.01$ in the graph metric (and consequently ℓ_∞-metric). Prior to this work, it was unclear if CST and DST are APX-hard in essentially every other popular metric! In this work, we prove that DST is APX-hard in every ℓ_p-metric. We also prove that CST is APX-hard in the ℓ_∞-metric. Finally, we relate CST and DST, showing a general reduction from CST to DST in ℓ_p-metrics. As an immediate consequence, this yields a 1.39-approximation polynomial time algorithm for CST in ℓ_p-metrics.

Henry Fleischmann
University of Michigan
henryfl@umich.edu

Karthik Srikanta, Surya Teja Gavva
Rutgers University
karthik0112358@gmail.com, suryatejagavva@gmail.com

CP13

Improved Approximations for Ultrametric Violation Distance

We study the Ultrametric Violation Distance problem introduced by Cohen-Addad, Fan, Lee, and Mesmay [FOCS, 2022]. Given pairwise distances $x \in \mathbb{R}_{>0}^{\binom{[n]}{2}}$ as input, the goal is to modify the minimum number of distances so as to make it a valid ultrametric. In other words, this is the problem of fitting an ultrametric to given data, where the quality of the fit is measured by the ℓ_0 norm of the error; variants of the problem for the ℓ_∞ and ℓ_1 norms are well-studied in the literature. Our main result is a 5-approximation algorithm for Ultrametric Violation Distance, improving the previous best large constant factor (≥ 1000) approximation algorithm. We give an $O(\min\{L, \log n\})$-approximation for weighted Ultrametric Violation Distance where the weights satisfy triangle inequality and L is the number of distinct values in the input. We also give a 16-approximation algorithm for the problem on k-partite graphs, where the input is specified on pairs of vertices that form a complete k-partite graph. All our results use a unified algorithmic framework with small modifications for the three cases.

Moses Charikar
Stanford University, California
moses@cs.stanford.edu

Ruiquan Gao
Stanford University
ruiquan@cs.stanford.edu

CP13

Composition of Nested Embeddings with An Application to Outlier Removal

We study the design of embeddings into Euclidean space with outliers. Given a metric space (X, d) and an integer k, the goal is to embed all but k points in X (called the "outliers") into ℓ_2 with the smallest possible distortion c. Finding the optimal distortion c for a given outlier set size k, or alternately the smallest k for a given target distortion c are both NP-hard problems. In fact, it is UGC-hard to approximate k to within a factor smaller than 2 even when the metric sans outliers is isometrically embeddable into ℓ_2. We consider bi-criteria approximations. Our main result is a polynomial time algorithm that approximates the outlier set size to within an $O(\log^2 k)$ factor and the distortion to within a constant factor. The main technical component in our result is an approach for constructing Lipschitz extensions of embeddings into Banach spaces (such as ℓ_p spaces). We consider a stronger version of Lipschitz extension that we call a *nested composition of embeddings*: given a low distortion embedding of a subset S of the metric space X, our goal is to extend this embedding to all of X such that the distortion over S is preserved, whereas the distortion over the remaining pairs of points in X is bounded by a function of the size of $X \setminus S$.

Shuchi Chawla
UT-Austin
shuchi@cs.utexas.edu

Kristin Sheridan
University of Texas at Austin
kristin@cs.utexas.edu

CP14

Count on CFI Graphs for #P-Hardness

Many interesting counting problems are finite linear combinations $p(\cdot) = \sum_H \alpha_H \hom(H, \cdot)$ of homomorphism counts from fixed graphs H; this includes subgraph and induced subgraph counts for fixed patterns. It is known that such finite linear combinations are as hard to evaluate as their hardest terms, and this fact was used in several works to establish hardness results for counting problems. More precisely, given a linear combination p as above, and a graph S with $\alpha_S \neq 0$, it is possible to reduce counting homomorphisms from S into an n-vertex input graph G to the evaluation of p. Given an oracle for evaluating p, this known reduction runs in time $2^{|E(S)|}\text{poly}(n, s)$, where s is the maximum number of vertices among graphs in the linear combination. This running time is acceptable when S and s are small compared to n, as considered in previous works. We show that the above reduction can actually be performed in $\text{poly}(n, s)$ time when S has constant maximum degree, even if S is part of the input. Thus, for constant-degree graphs S, the running time dependency on the size of S can be made polynomial instead of exponential. Using this new reduction, we obtain #P-hardness results for problems that could previously only be studied under parameterized complexity assumptions that are a priori stronger than "classical" assumptions. This includes (induced) subgraph counting problems.

Radu Curticapean
University of Regensburg
IT University of Copenhagen
racu@itu.dk

CP14

Randomized Communication and Implicit Representations for Matrices and Graphs of Small Sign-Rank

We prove a characterization of the structural conditions on matrices of sign-rank 3 and unit disk graphs (UDGs) which

permit constant-cost public-coin randomized communication protocols. Therefore, under these conditions, these graphs also admit implicit representations. The sign-rank of a matrix $M \in \{-1,+1\}^{N\times N}$ is the smallest rank of a matrix R such that $M_{i,j} = \mathrm{sign}(R_{i,j})$ for all $i,j \in [N]$; equivalently, it is the smallest dimension d in which M can be represented as a point-halfspace incidence matrix with halfspaces through the origin, and it is essentially equivalent to the unbounded-error communication complexity. Matrices of sign-rank 3 can achieve the maximum possible bounded-error randomized communication complexity $\Theta(\log N)$, and meanwhile the existence of implicit representations for graphs of bounded sign-rank (including UDGs, which have sign-rank 4) has been open since at least 2003. We prove that matrices of sign-rank 3, and UDGs, have constant randomized communication complexity if and only if they do not encode arbitrarily large instances of the Greater-Than communication problem, or, equivalently, if they do not contain large half-graphs as semi-induced subgraphs. This also establishes the existence of implicit representations for these graphs under the same conditions.

Nathaniel Harms
EPFL
nathaniel.harms@epfl.ch

Viktor Zamaraev
Department of Computer Science, University of Liverpoo
viktor.zamaraev@liverpool.ac.uk

CP14

On the Hardness of Posslp

The problem PosSLP involves determining whether an integer computed by a given straight-line program is positive. This problem has attracted considerable attention within the field of computational complexity as it provides a complete characterization of the complexity associated with numerical computation. However, non-trivial lower bounds for PosSLP remain unknown. In this paper, we demonstrate that PosSLP $\in$ BPP would imply that NP $\subseteq$ BPP, under the assumption of a conjecture concerning the complexity of the radical of a polynomial proposed by Dutta, Saxena, and Sinhababu (STOC'2018). Our proof builds upon the established NP-hardness of determining if a univariate polynomial computed by an SLP has a real root, as demonstrated by Perrucci and Sabia (JDA'2005).

Gorav Jindal
Aalto University
gorav.jindal@gmail.com

Peter Bürgisser
Technische Universität Berlin
Germany
pbuerg@math.tu-berlin.de

CP14

Computations with Polynomial Evaluation Oracle: Ruling Out Superlinear SETH-based Lower Bounds

One of the most popular and successfully used assumptions, Strong Exponential Time Hypothesis (SETH), implies that SAT cannot be solved in $2^{(1-\epsilon)n}$ time. In recent years, it has been proved that known algorithms for many problems are optimal under SETH. Despite wide applicability of SETH, for many problems, there are no known SETH-based lower bounds. In this paper, we prove that for a range of problems from P, proving superlinear lower bounds under SETH is challenging as it implies new circuit lower bounds. To this end, we show that these problems can be solved in nearly linear time with oracle calls to evaluating a polynomial of constant degree. Then, we introduce a strengthening of SETH stating that solving SAT in time $2^{(1-\varepsilon)n}$ is difficult even if one has constant degree polynomial evaluation oracle calls. This hypothesis is stronger than SETH, but refuting it is still challenging: we show that this implies circuit lower bounds. Finally, by considering computations that make oracle calls to evaluating constant degree polynomials depending on a small number of variables, we show connections between nondeterministic time lower bounds and arithmetic circuit lower bounds. Namely, we prove that if any of MAX-k-SAT, Binary Permanent, or a variant of Set Cover problems cannot be solved in co-nondeterministic time $2^{(1-\varepsilon)n}$, for any $\varepsilon > 0$,then one gets arbitrary large polynomial arithmetic circuit lower bounds.

Tatiana Belova
Steklov Institute of Mathematics at St. Petersburg
yukikomodo@gmail.com

Alexander Kulikov
Steklov Institute of Mathematics at St. Petersburg
Saint Petersburg University
alexander.s.kulikov@gmail.com

Ivan Mihajlin
University of CaliforniaSan Diego
ivmihajlin@gmail.com

Olga Ratseeva
St. Petersburg State University
ratsionalno@mail.ru

Grigory Reznikov
Steklov Institute of Mathematics at St. Petersburg
gritukan@gmail.com

Denil Sharipov
Saint Petersburg University
d.i.sharipov@yandex.ru

CP14

The Sharp Power Law of Local Search on Expanders

Local search is a powerful heuristic in optimization, which has been studied in the white box and black box models. In the black box model, we get a graph $G = (V, E)$ and oracle access to a function $f : V \to \mathbb{R}$. The problem is to find a vertex v that is a local minimum, with as few queries to the oracle as possible. We show the query complexity of local search on constant-degree expanders is $\Omega\left(\frac{\sqrt{n}}{\log n}\right)$, where n is the number of vertices of the graph. This matches within a logarithmic factor the upper bound of $O(\sqrt{n})$ from Aldous (1983), implying that steepest descent with a warm start is an optimal algorithm for expanders. The previous best lower bound was $\Omega\left(\frac{\sqrt[8]{n}}{\log n}\right)$, shown by Santha and Szegedy (2004). We obtain this result by considering broader graph features such as vertex congestion and separation number. We show that for each graph, the randomized query complexity of local search is $\Omega\left(\frac{n^{1.5}}{g}\right)$, where g is the vertex congestion of the graph;

and $\Omega\left(\sqrt[4]{\frac{s}{\Delta}}\right)$, where s is the separation number and Δ is the maximum degree. We also show a variant of the relational adversary method from Aaronson (2006). Our variant is asymptotically at least as strong for randomized algorithms and strictly stronger on some problems.

Nicholas Recker, Simina Branzei
Purdue University
nrecker@purdue.edu, simina.branzei@gmail.com

Davin Choo
National University of Singapore
davin@u.nus.edu

CP15

Edge-Coloring Algorithms for Bounded Degree Multigraphs

In this paper, we consider algorithms for edge-coloring multigraphs G of bounded maximum degree, i.e., $\Delta(G) = O(1)$. Shannon's theorem states that any multigraph of maximum degree Δ can be properly edge-colored with $\lfloor 3\Delta/2 \rfloor$ colors. Our main results include algorithms for computing such colorings. We design deterministic and randomized sequential algorithms with running time $O(n \log n)$ and $O(n)$, respectively. This is the first improvement since the $O(n^2)$ algorithm in Shannon's original paper, and our randomized algorithm is optimal up to constant factors. We also develop distributed algorithms in the LOCAL model of computation. Namely, we design deterministic and randomized LOCAL algorithms with running time $\tilde{O}(\log^5 n)$ and $O(\log^2 n)$, respectively. The deterministic sequential algorithm is a simplified extension of earlier work of Gabow et al. in edge-coloring simple graphs. The other algorithms apply the entropy compression method in a similar way to recent work by the author and Bernshteyn, where the authors design algorithms for Vizings theorem for simple graphs. We also extend those results to Vizings theorem for multigraphs.

Abhishek Dhawan
Georgia Institute of Technology
abhishek.dhawan@math.gatech.edu

CP15

A Faster Combinatorial Algorithm for Maximum Bipartite Matching

The maximum bipartite matching problem is among the most fundamental and well-studied problems in combinatorial optimization. A beautiful and celebrated combinatorial algorithm of Hopcroft and Karp (1973) shows that maximum bipartite matching can be solved in $O(m\sqrt{n})$ time on a graph with n vertices and m edges. For the case of very dense graphs, a fast matrix multiplication based approach gives a running time of $O(n^{2.371})$. These results represented the fastest known algorithms for the problem until 2013, when Madry introduced a new approach based on continuous techniques achieving much faster runtime in sparse graphs. This line of research has culminated in a spectacular recent breakthrough due to Chen *et al.* that gives an $m^{1+o(1)}$ time algorithm for maximum bipartite matching (and more generally, for min cost flows). This raises a natural question: are continuous techniques essential to obtaining fast algorithms for the bipartite matching problem? Our work makes progress on this question by presenting a new, purely combinatorial algorithm for bipartite matching, that runs in $\tilde{O}(m^{1/3}n^{5/3})$ time, and hence outperforms both Hopcroft-Karp and the fast matrix multiplication based algorithms on moderately dense graphs. Using a standard reduction, we also obtain an $\tilde{O}(m^{1/3}n^{5/3})$ time deterministic algorithm for maximum vertex-capacitated s-t flow in directed graphs when all vertex capacities are identical.

Julia Chuzhoy
Toyota Technological Institute at Chicago
cjulia@ttic.edu

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

CP15

Computing the 5-Edge-Connected Components in Linear Time

We provide a deterministic algorithm for computing the 5-edge-connected components of an undirected multigraph in linear time. There were probably good indications that this computation can be performed in linear time, but no such algorithm was actually known prior to this work. Thus, our paper answers a theoretical question, and sheds light on the possibility that a solution may exist for general k. A key component in our algorithm is an oracle for answering connectivity queries for pairs of vertices in the presence of at most four edge-failures. Specifically, the oracle can be constructed in linear time, and it answers connectivity queries in the presence of at most four edge-failures in worst-case constant time. We note that this is a result of independent interest. Our paper can be considered as a follow-up of recent work on computing the 4-edge-connected components in linear time. However, in dealing with the computation of the 5-edge-connected components, we are faced with unique challenges that do not appear when dealing with lower connectivity. The problem is that the 4-edge cuts in 3-edge-connected graphs are entangled in various complicated ways, that make it difficult to organize them in a compact way. Here we provide a novel analysis of those cuts, that reveals the existence of various interesting structures. These can be exploited so that we can disentangle and collect only those cuts that are essential in computing the 5-edge-connected components.

Evangelos Kosinas
University of Ioannina
ekosinas@cs.uoi.gr

CP15

Exact Community Recovery in the Geometric SBM

We study the problem of exact community recovery in the Geometric Stochastic Block Model (GSBM), where each vertex has an unknown community label as well as a known position, generated according to a Poisson point process in $\mathbb{R}^d$. Edges are formed independently conditioned on the community labels and positions, where vertices may only be connected by an edge if they are within a prescribed distance of each other. The GSBM thus favors the formation of dense local subgraphs, which commonly occur in real-world networks, a property that makes the GSBM qualitatively very different from the standard Stochastic Block Model (SBM). We propose a linear-time algorithm for exact community recovery, which succeeds down to the information-theoretic threshold, confirming a conjecture of Abbe, Baccelli, and Sankararaman. The algorithm involves two phases. The first phase exploits the density of local subgraphs to propagate estimated community la-

bels among sufficiently occupied subregions, and produces an almost-exact vertex labeling. The second phase then refines the initial labels using a Poisson testing procedure. Thus, the GSBM enjoys local to global amplification just as the SBM, with the advantage of admitting an information-theoretically optimal, linear-time algorithm.

Julia Gaudio, Xiaochun Niu, Ermin Wei
Northwestern University
julia.gaudio@northwestern.edu, xiaochunniu2024@u.northwestern.edu, ermin.wei@northwestern.edu

CP15

Higher-Order Cheeger Inequality for Partitioning with Buffers

We prove a new generalization of the higher-order Cheeger inequality for partitioning with buffers. Consider a graph $G = (V, E)$. The buffered expansion of a set $S \subseteq V$ with a buffer $B \subseteq V \setminus S$ is the edge expansion of S after removing all the edges from set S to its buffer B. An ε-buffered k-partitioning is a partitioning of a graph into disjoint components P_i and buffers B_i, in which the size of buffer B_i for P_i is small relative to the size of P_i: $|B_i| \leq \varepsilon|P_i|$. The buffered expansion of a buffered partition is the maximum of buffered expansions of the k sets P_i with buffers B_i. Let $h_G^{k,\varepsilon}$ be the buffered expansion of the optimal ε-buffered k-partitioning, then for every $\delta > 0$,

$$h_G^{k,\varepsilon} \leq O_\delta(1) \cdot \left(\frac{\log k}{\varepsilon}\right) \cdot \lambda_{\lfloor(1+\delta)k\rfloor},$$

where $\lambda_{\lfloor(1+\delta)k\rfloor}$ is the $\lfloor(1+\delta)k\rfloor$-th smallest eigenvalue of the normalized Laplacian of G. Our inequality is constructive and avoids the "square-root loss' that is present in the standard Cheeger inequalities (even for $k = 2$). We also provide a complementary lower bound, and a novel generalization to the setting with arbitrary vertex weights and edge costs.

Liren Shan
TTIC
lirenshan@ttic.edu

Konstantin Makarychev
Northwestern University
konstantin@northwestern.edu

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

Aravindan Vijayaraghavan
Northwestern University
aravindv@northwestern.edu

CP16

New Approximation Bounds for Small-Set Vertex Expansion

The vertex expansion of a graph is a fundamental graph parameter. Given a graph $G = (V, E)$ and a parameter $\delta \in (0, 1/2]$, its δ-Small-Set Vertex Expansion (SSVE) is defined as

$$\min_{S:|S|=\delta|V|} \frac{|\partial^V(S)|}{\min\{|S|, |S^c|\}}$$

where $\partial^V(S)$ is the vertex boundary of a set S. The SSVE problem is a natural graph partitioning problem, and is also of interest due to its connections to the StrongUniqueGames problem. We give a randomized algorithm running in time $n^{\mathsf{poly}(1/\delta)}$, which outputs a set S of size $\Theta(\delta n)$, having vertex expansion at most

$$\max\left(O(\sqrt{\phi^* \log d \log(1/\delta)}), \tilde{O}(d \log(1/\delta)) \cdot \phi^*\right),$$

where d is the largest vertex degree of the graph, and ϕ^* is the optimal δ-SSVE. The previous best known guarantees for this were the bi-criteria bounds of $\tilde{O}(1/\delta)\sqrt{\phi^* \log d}$ and $\tilde{O}(1/\delta)\phi^*\sqrt{\log n}$ due to Louis-Makarychev [TOC'16]. Our algorithm uses the basic SDP relaxation of the problem augmented with poly$(1/\delta)$ rounds of SoS. Our rounding algorithm is a combination of rounding algorithms of Raghavendra-Tan [SODA'12] and Austrin-Benabbas-Georgiou [SODA'13]. A key component of our analysis is novel Gaussian rounding lemma for hyperedges which might be of independent interest.

Suprovat Ghoshal
Northwestern University and TTIC at Chicago, U.S.
suprovat.ghoshal@gmail.com

Anand Louis
Indian Institute of Science
Bangalore
anandl@iisc.ac.in

CP16

Dependent Rounding with Strong Negative-Correlation, and Scheduling on Unrelated Machines to Minimize Completion Time

We describe a new dependent-rounding algorithmic framework for bipartite graphs. Given a fractional assignment y of values to edges of graph $G = (U \cup V, E)$, the algorithms return an integral solution Y such that each right-node $v \in V$ has at most one neighboring edge f with $Y_f = 1$, and where the variables Y_e also satisfy broad nonpositive-correlation properties. In particular, for any edges e_1, e_2 sharing a left-node $u \in U$, the variables Y_{e_1}, Y_{e_2} have strong negative-correlation properties, i.e. the expectation of $Y_{e_1}Y_{e_2}$ is significantly below $y_{e_1}y_{e_2}$. This algorithm is based on generating negatively-correlated Exponential random variables and using them in a contention-resolution scheme inspired by an algorithm Im & Shadloo (2020). Our algorithm gives stronger and much more flexible negative correlation properties. Dependent rounding schemes with negative correlation properties have been used for approximation algorithms for job-scheduling on unrelated machines to minimize weighted completion times (Bansal, Srinivasan, Svensson (2021), Im & Shadloo (2020), Im & Li (2023)). Using our new dependent-rounding algorithm, among other improvements, we obtain a 1.4-approximation for this problem. This significantly improves over the prior 1.45-approximation ratio of Im & Li (2023).

David G. Harris
University of Maryland
davidgharris29@gmail.com

CP16

New SDP Roundings and Certifiable Approximation for Cubic Optimization

We give new rounding schemes for SDP relaxations for the problems of maximizing cubic polynomials over the unit sphere and the hypercube. In both cases, the resulting algorithms yield a $O(\sqrt{n/k})$ multiplicative approximation in

$2^{O(k)}$poly(n) time. For the unit sphere, this improves on the rounding algorithms of Bhattiprolu et al. that need quasi-polynomial time to obtain a similar approximation guarantee. Over the hypercube, our results match the guarantee of a search algorithm of Khot and Naor that obtains a similar approximation ratio via techniques from convex geometry. Unlike their method, our algorithm obtains an upper bound on the integrality gap of SDP relaxations for the problem and as a result, also yields a certificate on the optimum value of the input instance. Our results naturally generalize to homogeneous polynomials of higher degree and imply improved algorithms for approximating satisfiable instances of Max-3SAT. Our main motivation is the stark lack of rounding techniques for SDP relaxations of higher degree polynomial optimization in sharp contrast to a rich theory of SDP roundings for the quadratic case. Our rounding algorithms introduce two new ideas: 1) a new polynomial reweighting based method to round sum-of-squares relaxations of higher degree polynomial maximization problems, and 2) a general technique to compress such relaxations down to substantially smaller SDPs by relying on an explicit construction of certain hitting sets.

Jun-Ting Hsieh
Carnegie Mellon University
juntingh@cs.cmu.edu

Pravesh Kothari
CMU
praveshk@cs.cmu.edu

Lucas Pesenti, Luca Trevisan
Bocconi University
lucas.pesenti@phd.unibocconi.it,
l.trevisan@unibocconi.it

CP16

Faster Exact and Approximation Algorithms for Packing and Covering Matroids via Push-Relabel

Input your abstract, including TeX commands, here. The abstract should be no longer than 1500 characters, including spaces. Only input the abstract text. Don't include title or author information here.

Kent Quanrud
Purdue University
krq@purdue.edu

CP16

On the Hardness of Finding Balanced Independent Sets in Random Bipartite Graphs

We consider the algorithmic problem of finding large *balanced* independent sets in sparse random bipartite graphs, and more generally the problem of finding independent sets with specified proportions of vertices on each side of the bipartition. In a bipartite graph it is trivial to find an independent set of density at least half (take one of the partition classes). In contrast, in a random bipartite graph of average degree d, the largest balanced independent sets (containing equal number of vertices from each class) are typically of density $(2+o_d(1))\frac{\log d}{d}$. Can we find such large balanced independent sets in these graphs efficiently? By utilizing the overlap gap property and the low-degree algorithmic framework, we prove that local and low-degree algorithms (even those that know the bipartition) cannot find balanced independent sets of density greater than $(1+\epsilon)\frac{\log d}{d}$ for any $\epsilon > 0$ fixed and d large but constant. This factor 2 *statistical–computational gap* between what exists and what local algorithms can achieve is analogous to the gap for finding large independent sets in (non-bipartite) random graphs. A particularly striking aspect of the gap in bipartite graphs is that the algorithm achieving the lower bound is extremely simple and can be implemented as a 1-local algorithm and a degree-1 polynomial (a linear function).

Will Perkins, Yuzhou Wang
Georgia Tech
william.perkins@gmail.com, ywang3694@gatech.edu

CP17

Quantum Worst-Case to Average-Case Reductions for All Linear Problems

We study the problem of designing efficient worst-case to average-case reductions. Prior to this work, such reductions were only known for a small number of specific problems or restricted computational models. In contrast, we show that for quantum computation, all linear problems admit worst-case to average-case reductions. Specifically, we provide an explicit and efficient transformation of quantum algorithms that are only correct on a small (even sub-constant) fraction of their inputs into ones that are correct on all inputs. En route, we obtain a tight $\Omega(n^2)$ lower bound on the average-case quantum query complexity of the Matrix-Vector Multiplication problem. Our techniques strengthen and generalise the recently introduced additive combinatorics framework for classical worst-case to average-case reductions (STOC 2022). We rely on quantum singular value transformations to construct quantum algorithms for linear verification in superposition and learning Bogolyubov subspaces from noisy quantum oracles, which could be of independent interest. We use these tools to prove a quantum local correction lemma, which lies at the heart of our reductions, based on a noise-robust probabilistic generalisation of Bogolyubov's lemma from additive combinatorics.

Vahid Asadi
University of Waterloo
vrasadi@uwaterloo.ca

Alexander Golovnev
Georgetown University
alexgolovnev@gmail.com

Tom Gur
University of Warwick
tom.gur@warwick.ac.uk

Igor Shinkar
Simon Fraser University
ishinkar@sfu.ca

Sathyawageeswar Subramanian
University of Cambridge
sathya.subramanian@warwick.ac.uk

CP17

Recovering the Original Simplicity: Succinct and Deterministic Quantum Algorithm for the Welded Tree Problem

This work revisits quantum algorithms for the well-known

welded tree problem, proposing a very succinct quantum algorithm based on the simplest coined quantum walks. It simply iterates the naturally defined coined quantum walk operator for a predetermined time and finally measure, where the predetermined time can be efficiently computed on classical computers. Then, the algorithm returns the correct answer deterministically, and achieves exponential speedups over any classical algorithm. The significance of the results may be seen as follows. (i) Our algorithm is rather simple compared with the one in (Jeffery and Zur, STOC'2023), which not only breaks the stereotype that coined quantum walks can only achieve quadratic speedups over classical algorithms, but also demonstrates the power of the simplest quantum walk model. (ii) Our algorithm theoretically achieves zero-error, which is not possible with existing methods. Thus, it becomes one of the few examples that exhibit exponential separation between deterministic (exact) quantum and randomized query complexities, which may also change people's perception that since quantum mechanics is inherently probabilistic, it impossible to have a deterministic quantum algorithm with exponential speedups for the weled tree problem.

Guanzhong Li, Lvzhou Li, Jingquan Luo
Sun Yat-sen University
ligzh9@mail2.sysu.edu.cn, lilvzh@mail.sysu.edu.cn, luojq25@mail2.sysu.edu.cn

CP17

Viderman's Algorithm for Quantum LDPC Codes

Quantum low-density parity-check (LDPC) codes are an important class of quantum error correcting codes. A critical ingredient is the associated decoding algorithm that allows for the corrections of errors. In this talk I am going to describe a decoding algorithm for a specific family of quantum LDPC codes called hypergraph product codes (Tillich, Zémor, IEEE T-IT, 2013). Our decoding algorithm, a variant of the classical Viderman's algorithm (Viderman, TOCT 2013), protects against errors by producing a subset of qubits that contains the error. This is a linear time algorithm (i.e. in time O(N) where N is the total number of qubits). This then allows an erasure decoding algorithm to later correct the errors. Our decoding algorithm is different from all previous decoding algorithms, that are based on flipping small sets of qubits.

Inbal R. Livni Navon, Anirudh Krishna, Mary Last Nwoottersame
Stanford University
inballn@stanford.edu, anirudhk@stanford.edu, marykw@stanford.edu

CP17

Efficient Quantum State Synthesis with One Query

We present a polynomial-time quantum algorithm making a single query (in superposition) to a classical oracle, such that for every state $|\psi\rangle$ there exists a choice of oracle that makes the algorithm construct an exponentially close approximation of $|\psi\rangle$. Previous algorithms for this problem either used a linear number of queries and polynomial time, or a constant number of queries and polynomially many ancillae but no nontrivial bound on the runtime. As corollaries we do the following: - We simplify the proof that statePSPACE $\subseteq$ stateQIP (a quantum state analogue of PSPACE $\subseteq$ IP) and show that a constant number of rounds of interaction suffices. - We show that QAC^0_f lower bounds for constructing explicit states would imply breakthrough circuit lower bounds for computing explicit boolean functions. - We prove that every n-qubit state can be constructed to within 0.01 error by an $O(2^n/n)$-size circuit over an appropriate finite gate set. More generally we give a size-error tradeoff which, by a counting argument, is optimal for any finite gate set.

Gregory A. Rosenthal
University of Cambridge
University of Warwick
gregrosent@gmail.com

CP17

An Improved Classical Singular Value Transformation for Quantum Machine Learning

We study quantum speedups in quantum machine learning (QML) by analyzing the quantum singular value transformation (QSVT) framework. QSVT, introduced by [GSLW, STOC'19], unifies all major types of quantum speedup; in particular, a wide variety of QML proposals are applications of QSVT on low-rank classical data. We challenge these proposals by providing a classical algorithm that matches the performance of QSVT in this regime up to a small polynomial overhead. We show that, given a matrix $A \in \mathbb{C}^{m\times n}$, a vector $b \in \mathbb{C}^n$, a bounded degree-d polynomial p, and linear-time pre-processing, we can output a description of a vector v such that $\|v - p(A)b\| \le \varepsilon\|b\|$ in $\widetilde{\mathcal{O}}(d^{11}\|A\|_{\mathrm{F}}^4/(\varepsilon^2\|A\|^4))$ time. This improves upon the best known classical algorithm [CGLLTW, STOC'20], which requires $\widetilde{\mathcal{O}}(d^{22}\|A\|_{\mathrm{F}}^6/(\varepsilon^6\|A\|^6))$ time, and narrows the gap with QSVT, which, after linear-time pre-processing to load input into a quantum-accessible memory, can estimate the magnitude of an entry $p(A)b$ to $\varepsilon\|b\|$ error in $\widetilde{\mathcal{O}}(d\|A\|_{\mathrm{F}}/(\varepsilon\|A\|))$ time. Our key insight is to combine the Clenshaw recurrence, an iterative method for computing matrix polynomials, with sketching techniques to simulate QSVT. We introduce techniques for such polynomials which may be of independent interest.

Ainesh Bakshi
MIT
ainesh@mit.edu

Ewin Tang
University of California at Berkeley
ewin@berkeley.edu

CP18

Fault-Tolerant Spanners Against Bounded-Degree Edge Failures: Linearly More Faults, Almost for Free

We study a new and stronger notion of fault-tolerant graph structures whose size bounds depend on the degree of the failing edge set, rather than the total number of faults. For a subset of faulty edges $F \subseteq G$, the faulty-degree $\deg(F)$ is the largest number of faults in F incident to any given vertex. We design new fault-tolerant structures with size comparable to previous constructions, but which tolerate every fault set of small faulty-degree $\deg(F)$, rather than only fault sets of small size $|F|$. Our main results are: New FT-Certificates: For every n-vertex graph G and degree threshold f, one can compute a connectivity certificate $H \subseteq G$ with $|E(H)| = \widetilde{O}(fn)$ edges that has the following guarantee: for any edge set F with faulty-degree $\deg(F) \le f$ and every vertex pair u, v, it holds that u and v are connected in $H \setminus F$ iff they are connected in $G \setminus F$.

This bound on $|E(H)|$ is nearly tight. New FT-Spanners: We show that every n-vertex graph G admits a $(2k-1)$-spanner H with $|E(H)| = O_k(f^{1-1/k}n^{1+1/k})$ edges, which tolerates any fault set F of faulty-degree at most f. This bound on $|E(H)|$ optimal up to its hidden dependence on k.

Greg Bodwin
University of Michigan
bodwin@umich.edu

Bernhard Haeupler
ETH Zurich
haeuplb@ethz.ch

Merav Parter
Weizmann Institute
merav.parter@weizmann.ac.il

CP18

Distances and Shortest Paths on Graphs of Bounded Highway Dimension: Simple, Fast, Dynamic

Dijkstra's algorithm is the standard method for computing shortest paths on arbitrary graphs. However, it is slow for large graphs, taking at least linear time. It has been long known that for real world road networks, creating a hierarchy of well-chosen shortcuts allows fast distance and path computation, with exact distance queries seemingly being answered in logarithmic time. However, these methods were but heuristics until the work of Abraham et al. [JACM 2016], where they defined a graph parameter called highway dimension which is constant for real-world road networks, and showed that in graphs of constant highway dimension, a shortcut hierarchy exists that guarantees shortest distance computation takes $O(\log(U + |V|))$ time and $O(V \log(U + |V|))$ space, where U is the ratio of the smallest to largest edge, and $|V|$ is the number of vertices. The problem is that they were unable to efficiently compute the hierarchy of shortcuts. Here we present a simple and efficient algorithm to compute the needed hierarchy of shortcuts in time and space $O(V \log(U + |V|))$, as well as supporting updates in time $O(\log(U + |V|))$.

Sébastien Collette
Synapsis Group
sco@synapsis-group.com

John Iacono
Univerité libre de Bruxelles and Synapsis Group
johniacono@gmail.com

CP18

Exact Shortest Paths with Rational Weights on the Word Ram

Exact computation of shortest paths in weighted graphs has been traditionally studied in one of two settings. First, one can assume that the edge weights are real numbers and all the performed operations on reals (typically comparisons and additions) take constant time. Classical Dijkstra's and Bellman-Ford algorithms have been described in this setting. More efficient exact shortest paths algorithms have been obtained for integer-weighted graphs. Integrality assumption not only enables faster algorithms but also allows implementing the aforementioned algorithms in a much more realistic word RAM model where only arithmetic operations on $O(\log n)$-bit integers are performed in constant time. On the word RAM one can as efficiently exactly encode even rational-weighted instances with $O(\log n)$-bit numerators and denominators, however no known algorithms can be directly applied to get satisfactory results in such setting. In this paper, we design randomized exact single-source shortest paths algorithms for rational-weighted graphs on the word RAM. Most importantly, in the non-negative case, we obtain a near-linear time algorithm matching Dijkstra's algorithm running time up to polylogarithmic factors. In presence of negative weights, we give an $\widetilde{O}(n^{2.5})$-time algorithm breaking through the best known strongly polynomial bound attained by Bellman-Ford for sufficiently dense graphs.

Adam Karczmarz, Wojciech Nadara, Marek Sokolowski
University of Warsaw
a.karczmarz@mimuw.edu.pl, w.nadara@mimuw.edu.pl, marek.sokolowski@mimuw.edu.pl

CP18

Simpler and Higher Lower Bounds for Shortcut Sets

We study the well-known shortcut set problem: how much can one decrease the diameter of a directed graph on n vertices and m edges by adding $O(n)$ or $O(m)$ of shortcuts from the transitive closure of the graph. Our results are based on a vast simplification of the recent construction of Bodwin and Hoppenworth [FOCS 2023] which was used to show an $\tilde{\Omega}(n^{1/4})$ lower bound for the $O(n)$-sized shortcut set problem. We highlight that our simplification completely removes the use of the convex sets by Bárány and Larman [Math. Ann. 1998] used in all previous lower bound constructions. Our simplification also removes the need for randomness and further removes some log factors. This allows us to generalize the construction to higher dimensions, which in turn can be used to show an $\tilde{\Omega}(n^{1/5})$ lower bound for $O(m)$-sized shortcut sets, improving on the previous best $\tilde{\Omega}(n^{1/8})$ lower bound.

Zixuan Xu, Virginia Vassilevska Williams
MIT
zixuanxu@mit.edu, virgi@mit.edu

Yinzhan Xu
Massachusetts Institute of Technology
xyzhan@mit.edu

CP18

Nearly Optimal Approximate Dual-Failure Replacement Paths

Given a directed graph $G = (V, E, \omega)$ on n vertices with positive edge weights as well as two designated terminals $s, t \in V$, our goal is to compute the shortest path from s to t avoiding any pair of presumably failed edges $f_1, f_2 \in E$, which is a natural generalization of the classical replacement path problem which considers single edge failures only. This dual failure replacement paths problem was recently studied by [Vassilevska Williams, Woldeghebriela and Xu, 2022] who designed a $\tilde{O}(n^3)$ time algorithm for general weighted digraphs which is conditionally optimal. In the same paper, they also showed that the cubic time barrier can be bypassed by a $\tilde{O}(M^{2/3}n^{2.9146})$ time algorithm for input graphs with small integer edge weights from $\{-M, -M+1, \cdots, M-1, M\}$. In this paper, we study the natural question whether sub-cubic time algo-

rithms exist for general weighted digraphs when approximation is allowed. As our main result, we show that $(1+\epsilon)$-approximations of all dual-failure replacement paths can be computed in $\tilde{O}_\epsilon(n^2)$ time which is nearly optimal.

Shiri Chechik
Tel-Aviv University, Israel
shiri.chechik@gmail.com

Tianyi Zhang
IIIS, Tsinghua
tianyi-z16@tsinghua.org.cn

CP19

Oracle Efficient Online Multicalibration and Omniprediction

A recent line of work has shown a connection between multicalibration, a multi-group fairness notion, and omniprediction, a learning paradigm that provides simultaneous loss minimization guarantees for a family of loss functions. We initiate the study of omniprediction in the online adversarial setting. Although there exist algorithms for obtaining multicalibration in the online setting, they work only for small finite classes of benchmark functions F, because they require enumerating every function $f \in F$ at every round. In contrast, omniprediction is most interesting for learning theoretic hypothesis classes F, which are generally continuously large. We develop a new online multicalibration algorithm for infinite benchmark classes F and is oracle efficient (i.e. there exist a no-regret learning algorithm for F). The result is the first efficient online omnipredictor — an oracle efficient prediction algorithm that can be used to simultaneously obtain no regret guarantees to all Lipschitz convex loss functions. And we show upper and lower bounds on the extent to which our rates can be improved. Our oracle efficient algorithm promises a stronger swap-omniprediction guarantee. We show that $O(\sqrt{T})$ bounds for swap-omniprediction is impossible in the online setting and give a non-oracle efficient algorithm that gets the optimal $O(\sqrt{T})$ omniprediction bounds without going through multicalibration, giving an information theoretic separation.

Christopher Jung
University of Pennsylvania
csj93@stanford.edu

CP19

Santa Claus Meets Makespan and Matroids: Algorithms and Reductions

In this paper we study the relation of two fundamental problems in scheduling and fair allocation: makespan minimization on unrelated parallel machines and max-min fair allocation, also known as the Santa Claus problem. For both of these problems the best approximation factor is a notorious open question; more precisely, whether there is a better-than-2 approximation for the former problem and whether there is a constant approximation for the latter. While the two problems are intuitively related and history has shown that techniques can often be transferred between them, no formal reductions are known. We first show that an affirmative answer to the open question for makespan minimization implies the same for the Santa Claus problem by reducing the latter problem to the former, and, that for instances with only two input values both questions are equivalent. We then move to a special case called "restricted assignment", which is well studied in both problems. Although our reductions do not maintain the characteristics of this special case, we show that they hold in a slight generalization, where the jobs or resources are assigned to multiple machines or players subject to a matroid constraint and in addition we have only two values. To complete the picture, we give an algorithm for our new matroid variant of the Santa Claus problem using a nontrivial extension of the local search method from restricted assignment.

Etienne Bamas
ETH AI Center, Switzerland
etienne.bamas@inf.ethz.ch

Alexander Lindermayr
Faculty of Mathematics and Computer Science,
University of Bremen
lindera@uni-bremen.de

Nicole Megow
Universität Bremen, Germany
nicole.megow@uni-bremen.de

Lars Rohwedder
Maastricht University, The Netherlands
l.rohwedder@maastrichtuniversity.nl

Jens Schlöter
Faculty of Mathematics and Computer Science
University of Bremen
jschloet@uni-bremen.de

CP19

Fair Price Discrimination

A seller is pricing identical copies of a good to a stream of unit-demand buyers. Each buyer has a value on the good as his private information. The seller only knows the empirical value distribution of the buyer population and chooses the revenue-optimal price. We consider a widely studied third-degree price discrimination model where an information intermediary with perfect knowledge of the arriving buyer's value sends a signal to the seller, hence changing the seller's posterior and inducing the seller to set a personalized posted price. Prior work of Bergemann, Brooks, and Morris (2015) has shown the existence of a signaling scheme that preserves seller revenue, while always selling the item, hence maximizing consumer surplus. Departing from this, we ask whether the consumer surplus generated is fairly distributed among buyers with different values. To this end, we aim to maximize welfare functions that reward more balanced surplus allocations. Our main result is the surprising existence of a novel signaling scheme that simultaneously 8-approximates all welfare functions that are non-negative, monotonically increasing, symmetric, and concave, compared with any other signaling scheme. Such a guarantee cannot be given by consumer-surplus-maximizing schemes – which are the ones typically studied in the literature. In addition, our scheme is socially efficient, and has the fairness property that buyers with higher values enjoy higher expected surplus.

Siddhartha Banerjee
Cornell University
sbanerjee@cornell.edu

Kamesh Munagala
Duke University, USA

kamesh@cs.duke.edu

Yiheng Shen
Duke University
yiheng.shen@duke.edu

Kangning Wang
Stanford University
knwang@stanford.edu

CP19

Improved Approximation Algorithms for the Joint Replenishment Problem with Outliers, and with Fairness Constraints

We study a natural generalization of the Joint Replenishment Problem (JRP), which is a classical inventory management problem: we model outliers, and may reject a subset of demand points. In this paper, we are motivated by issues of fairness - we wish to spread out the rejection pain in a balanced way among customers, communities, or with respect to any specified market segmentation. In our most general setting, we consider a set of C features, where each demand point has an associated rejection cost for each feature, and we are given a bound on the allowed rejection cost incurred in total for each feature. This generalizes a notion of fairness introduced earlier for the Colorful k-Center problem in which (analogously) each demand point has a given color, and the number of rejections of each color class is bounded. We give the first constant approximation algorithms for the fairness-constrained JRP with a constant number of features: we give a 2.86-approximation algorithm. Even for the special case where we simply bound the total (weighted) number of outliers, this performance guarantee improves upon previously known results. Our approach is an LP-based algorithm that splits the instance into two subinstances. One is solved by a novel iterative rounding approach and the other by pipage-based rounding. The standard LP relaxation has an unbounded integrality gap, therefore, we strengthen the relaxation by correctly guessing a few key attributes of the optimal solution.

Varun Suriyanarayana
Cornell University
vs478@cornell.edu

Varun Sivashankar, Siddharth Gollapudi
Microsoft Research
varunsiva@ucla.edu, sgollapu@berkeley.edu

David B. Shmoys
Cornell University
david.shmoys@cornell.edu

CP19

School Redistricting: Wiping Unfairness Off the Map

We introduce and study the problem of designing an equitable school redistricting map, which we formalize as that of assigning n students to school attendance zones in a way that is fair to various demographic groups. Drawing on methodology from fair division, we consider the demographic groups as players and seats in schools as homogeneous goods. Due to geographic constraints, not every school can be assigned to every student. This raises new obstacles, rendering some classic fairness criteria infeasible. Nevertheless, we show that it is always possible to find an almost proportional allocation among g demographic groups if we are allowed to add $O(g \log g)$ extra seats. For any fixed g, we show that such an allocation can be found in polynomial time, obtaining a runtime of $O(n^2 \log n)$ in the special (but practical) case where $g \leq 3$.

Jamie Tucker-Foltz, Ariel Procaccia
Harvard University
jtuckerfoltz@gmail.com, arielpro@gmail.com

Isaac Robinson
University of Oxford
isaac_robinson@college.harvard.edu

CP20

Time-Space Lower Bounds for Bounded-Error Computation in the Random-Query Model

The random-query model was introduced by Raz and Zhan at ITCS 2020 as a new model of space-bounded computation. In this model, a branching program of length T and width 2^S attempts to compute a function $f : \{0,1\}^n \to \{0,1\}$. The input $(x_1, \ldots, x_n)$ is given in pairs of the form $(i_j, x_{i_j}) \in \{1, \ldots, n\} \times \{0,1\}$ for $j = 1, 2, \ldots, T$, where $i_1, \ldots, i_T$ are uniform and independent indices. Raz and Zhan proved that any branching program in the random-query model that computes f with sensitivity k satisfies $T \cdot (S + \log n) \geq \Omega(n \cdot k)$. This gives a quadratic time-space lower bound for many natural functions which have sensitivity $\Omega(n)$, such as XOR and Majority. The bound was proved in the zero-error regime. Furthermore, Raz and Zhan conjectured that (up to logarithmic factors) a quadratic time-space lower bound still holds for the XOR function in the more conventional bounded-error regime. In this paper, we prove this conjecture. More generally, let $f : \{0,1\}^n \to \{0,1\}$ have average sensitivity (or total influence) $\mathrm{I}[f]$. We prove that any branching program in the random-query model that computes f in the bounded-error regime satisfies $T \cdot S \geq \tilde{\Omega}(n) \cdot \mathrm{I}[f]$ (where $\tilde{\Omega}$ hides logarithmic factors in n). Our proof is based on a reduction from a communication complexity problem.

Itai Dinur
Ben-Gurion University
dinuri@bgu.ac.il

CP20

Robust 1-Bit Compressed Sensing with Iterative Hard Thresholding

In 1-bit compressed sensing, the aim is to estimate a k-sparse unit vector $x \in S^{n-1}$ within an ϵ error (in ℓ_2) from minimal number of linear measurements that are quantized to just their signs, i.e., from measurements of the form $y = \mathrm{Sign}(\langle a, x \rangle)$. In this paper, we study a noisy version where a fraction of the measurements can be flipped, potentially by an adversary. In particular, we analyze the popular Binary Iterative Hard Thresholding (BIHT) algorithm in this noisy setting. It is known from recent results that, with $\tilde{O}(\frac{k}{\epsilon})$ noiseless measurements, BIHT provides an estimate within ϵ error. This result is optimal and universal, meaning one set of measurements work for all sparse vectors. In this paper, we show that BIHT also provides better results than all known methods for the noisy setting. We show that when up to τ-fraction of the sign measurements are incorrect (adversarial error), with the same number of measurements as before, BIHT agnostically provides an estimate of x within an $\tilde{O}(\epsilon + \tau)$ error, maintaining the

universality of measurements. This establishes stability of iterative hard thresholding in the presence of measurement error. To obtain the result, we use the restricted approximate invertibility of Gaussian matrices, as well as a tight analysis of the high-dimensional geometry of the adversarially corrupted measurements.

Arya Mazumdar, Namiko Matsumoto
University of California San Diego
arya@ucsd.edu, nmatsumo@ucsd.edu

CP20

An Unconditional Lower Bound for Two-Pass Streaming Algorithms for Maximum Matching Approximation

We present the first unconditional space lower bound for two-pass streaming algorithms for Maximum Bipartite Matching approximation. We show that every randomized two-pass streaming algorithm that computes a $(\frac{8}{9}+\epsilon)$-approximation to Maximum Bipartite Matching, for any constant $\epsilon > 0$, requires space $n^{1+\Omega(\frac{1}{(\log\log n)^2})}$, where n is the number of vertices of the input graph. This improves on the previous conditional lower bound by Assadi [SODA'22] that, in the best case, only rules out approximation ratios better than 0.98. Our lower bound makes use of the information cost trade-off of the Index problem in the two-party communication setting established by Jain et al. [JACM'09]. To the best of our knowledge, our work is the first that exploits this trade-off result in the context of lower bounds for multi-pass graph streaming algorithms.

Christian Konrad, Kheeran K. Naidu
University of Bristol
christian.konrad@bristol.ac.uk,
kheeran.naidu@bristol.ac.uk

CP20

A $(3+\varepsilon)$-Approximate Correlation Clustering Algorithm in Dynamic Streams

Grouping together similar elements in datasets is a common task in data mining and machine learning. In this paper, we study streaming and parallel algorithms for correlation clustering, where each pair of elements is labeled either similar or dissimilar. The task is to partition the elements and the objective is to minimize disagreements, that is, the number of dissimilar elements grouped together and similar elements that get separated. Our main contribution is a semi-streaming algorithm that achieves a $(3+\varepsilon)$-approximation to the minimum number of disagreements using a single pass over the stream. In addition, the algorithm also works for dynamic streams. Our approach builds on the analysis of the PIVOT algorithm by Ailon, Charikar, and Newman [JACM'08] that obtains a 3-approximation in the centralized setting. Our design allows us to sparsify the input graph by ignoring a large portion of the nodes and edges without a large extra cost as compared to the analysis of PIVOT. Our work improves on the approximation ratio of the recent single-pass 5-approximation algorithm and on the number of passes of the recent $O(1/\varepsilon)$-pass $(3+\varepsilon)$-approximation algorithm [Behnezhad, Charikar, Ma, Tan FOCS'22, SODA'23]. Our algorithm is also more robust and can be applied in dynamic streams. Furthermore, it is the first single pass $(3+\varepsilon)$-approximation algorithm that uses polynomial post-processing time.

Mélanie Cambus
Aalto University
melanie.cambus@aalto.fi

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

Etna Lindy, Shreyas Pai, Jara Uitto
Aalto University
etna.lindy@aalto.fi, shreyas.pai@aalto.fi,
jara.uitto@aalto.fi

CP20

Robust Sparsification for Matroid Intersection with Applications

Matroid intersection is a classical optimization problem where, given two matroids over the same ground set, the goal is to find the largest common independent set. In this paper, we show that there exists a certain "sparsifer': a subset of elements, of size $O(|S^{opt}| \cdot 1/\varepsilon)$, where S^{opt} denotes the optimal solution, that is guaranteed to contain a $3/2+\varepsilon$ approximation, while guaranteeing certain robustness properties. We call such a small subset a Density Constrained Subset (DCS), which is inspired by the Edge-Degree Constrained Subgraph (EDCS) [Bernstein and Stein, 2015], originally designed for the maximum cardinality matching problem in a graph. Our proof is constructive and hinges on a greedy decomposition of matroids, which we call the density-based decomposition. We show that this sparsifier has certain robustness properties that can be used in one-way communication and random-order streaming models.

Chien-Chung Huang
DI ENS, CNRS, PSL
chien-chung.huang@ens.fr

Francois Sellier
DI ENS, PSL & Mines Paris, PSL
francois.sellier0359@gmail.com

CP21

Fully Dynamic Shortest Path Reporting Against An Adaptive Adversary

Algebraic data structures are the main subroutine for maintaining distances in fully dynamic graphs in subquadratic time. However, these dynamic algebraic algorithms generally cannot maintain the shortest paths against adaptive adversaries. We present the first fully dynamic algorithm that maintains the shortest paths against an adaptive adversary in subquadratic update time. This is obtained via a combinatorial reduction that allows reconstructing the shortest paths with only a few distance estimates. Using this reduction, we obtain the following: On weighted directed graphs with real edge weights in $[1, W]$, we maintain $(1+\epsilon)$-approximate shortest paths in $\tilde{O}(n^{1.816}\epsilon^{-2}\log W)$ update and $\tilde{O}(n^{1.741}\epsilon^{-2}\log W)$ query time. This improves upon the approximate distance data structures from [v.d.Brand, Nanongkai; FOCS'19], which only returned a distance estimate, by matching their complexity and returning an approximate shortest path. On unweighted directed graphs, we maintain exact shortest paths in $\tilde{O}(n^{1.823})$ update and $\tilde{O}(n^{1.747})$ query time. This

improves upon [Bergamaschi, Henzinger, P.Gutenberg, V.Williams, Wein; SODA'21] who could report the path only against oblivious adversaries. We improve both their update and query time while also handling adaptive adversaries.

Anastasiia Alokhina, Jan Van Den Brand
Georgia Tech
aalokhina3@gatech.edu, vdbrand@gatech.edu

CP21

Fully Dynamic Matching: (2-v2)-Approximation in Polylog Update Time

We study maximum matchings in fully dynamic graphs, which are graphs that undergo both edge insertions and deletions. Our focus is on algorithms that estimate the size of maximum matching after each update while spending a small time. An important question studied extensively is the best approximation achievable via algorithms that only spend $\text{poly}(\log n)$ time per update, where n is the number of vertices. The current best bound is a $(1/2+\epsilon_0)$-approximation for a small constant $\epsilon_0 > 0$, due to recent works of Behnezhad [SODA'23] ($\epsilon_0 \sim 0.001$) and Bhattacharya, Kiss, Saranurak, Wajc [SODA'23] ($\epsilon_0 \sim 0.006$) who broke the long-standing 1/2-approximation barrier. These works also showed that for any fixed $\epsilon > 0$, the approximation can be further improved to $(2-\sqrt{2}-\epsilon) \sim .585$ for bipartite graphs, leaving a huge gap between general and bipartite graphs. In this work, we close this gap. We show that for any fixed $\epsilon > 0$, a $(2-\sqrt{2}-\epsilon)$ approximation can be maintained in $\text{poly}(\log n)$ time per update *even in general graphs.* Our techniques also lead to the same approximation for general graphs in two passes of the semi-streaming setting, removing a similar gap in that setting.

Mohammad Roghani
Stanford University
roghani@stanford.edu

Amir Azarmehr, Soheil Behnezhad
Northeastern University
azarmehr.a@northeastern.edu,
s.behnezhad@northeastern.edu

CP21

Incremental Approximate Maximum Flow on Undirected Graphs in Subpolynomial Update Time

We provide an algorithm which, with high probability, maintains a $(1+\epsilon)$-approximate maximum flow on an undirected graph undergoing m-edge additions in $m^{o(1)}\epsilon^{-3}$ time per update. To obtain this result, we provide a more general algorithm that solves what we call the *incremental, thresholded, p-norm flow problem* that asks to determine the first edge-insertion in an undirected graph that causes the minimum ℓ_p-norm flow to decrease below a given threshold in value. Consequently, we also obtain improved algorithms for dynamically maintaining the effective resistance between a pair of vertices in an undirected graph undergoing edge additions. Our algorithm builds upon previous dynamic algorithms for approximately solving the minimum-ratio cycle that underlies previous advances on the maximum flow problem [Chen-Kyng-Liu-Peng-Probst Gutenberg-Sachdeva, FOCS '22]. Our work departs from this result as well as previous dynamic maximum flow algorithms [v.d.Brand-Liu-Sidford, STOC '23] in how it reasons adaptive queries and updates to the minimum ratio cycle problem as well as how the dynamic algorithm is used in optimization frameworks to solve p-norm flow.

Jan Van Den Brand
Georgia Tech
vdbrand@gatech.edu

Li Chen
Georgia Institute of Technology
lichen@gatech.edu

Rasmus Kyng
ETH Zurich
kyng@inf.ethz.ch

Yang Liu
Stanford University
yangpatil@gmail.com

Richard Peng
University of Waterloo
y5peng@uwaterloo.ca

Maximilian Probst Gutenberg
ETH Zürich
maxprobst@ethz.ch

Sushant Sachdeva
University of Toronto
sachdeva@cs.toronto.edu

Aaron Sidford
Stanford University
sidford@stanford.edu

CP21

Fully Dynamic Min-Cut of Superconstant Size in Subpolynomial Time

We present a deterministic fully dynamic algorithm with subpolynomial worst-case time per graph update such that after processing each update of the graph, the algorithm outputs a minimum cut of the graph if the graph has a cut of size at most c for some $c = (\log n)^{o(1)}$. Previously, the best update time was $\widetilde{O}(\sqrt{n})$ for any $c > 2$ and $c = O(\log n)$ [Thorup 2007].

Xiaorui Sun, Wenyu Jin
University of Illinois at Chicago
sunsirius@gmail.com, wjin9@uic.edu

Mikkel Thorup
University of Copenhagen
mikkel2thorup@gmail.com

CP21

Adaptive Out-Orientations with Applications

We maintain edge-orientations of a fully-dynamic graph, such that the out-degree of each vertex is bounded. On one hand, we show how to orient the edges such that the out-degree of each vertex is proportional to the arboricity α of the graph, in, either, an amortised update time of $O(\log^2 n \log \alpha)$, or a worst-case update time of $O(\log^3 n \log \alpha)$. We obtain another trade-off, namely either $O(\log n \log \alpha)$, amortised, or $O(\log^2 n \log \alpha)$, worst-case time, for the problem of maintaining an edge-orientation

with at most $O(\alpha + \log n)$ out-edges per vertex. Our algorithms adapt to the current arboricity of the graph, and yield improvements over previous work: We obtain an $O(\varepsilon^{-6} \log^3 n \log \rho)$ worst-case update time algorithm for maintaining a $(1+\varepsilon)$ approximation of the maximum subgraph density, ρ. We obtain an $O(\varepsilon^{-6} \log^3 n \log \alpha)$ worst-case update time algorithm for maintaining a $(1+\varepsilon)\cdot OPT+2$ approximation of the optimal out-orientation of a graph with adaptive arboricity α. This yields the first worst-case polylogarithmic dynamic algorithm for decomposing into $O(\alpha)$ forests. We obtain arboricity-adaptive fully-dynamic deterministic algorithms for a variety of problems including maximal matching, $\Delta+1$ coloring, and matrix vector multiplication. All update times are worst-case $O(\alpha + \log^2 n \log \alpha)$.

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

Aleksander Christiansen
DTU Compute, Technical University of Denmark
abgch@dtu.dk

Jacob Holm
University of Copenhagen
University of Copenhagen
jaho@di.ku.dk

Ivor Van Der Hoog
Technical University of Denmark
idjva@dtu.dk

Kent Quanrud
Purdue University
krq@purdue.edu

Eva Rotenberg
Technical University of Denmark
erot@dtu.dk

Chris Schwiegelshohn
Aarhus University
schwiegelshohn@cs.au.dk

CP22

Learning Hard-Constrained Models with One Sample

We consider the problem of estimating the parameters of a Markov Random Field with *hard*-constraints using a single sample. As our main running examples, we use the k-SAT, the proper coloring models and the general H-coloring models, for which we obtain both positive and negative results. In contrast to the soft-constrained case, we show in particular that single-sample estimation is not always possible, and that the existence of an estimator is related to the existence of non-satisfiable instances. Our algorithms are based on the pseudo-likelihood estimator. We show variance bounds for this estimator using coupling techniques inspired, in the case of k-SAT, by Moitra's sampling algorithm (JACM, 2019); our positive results for colorings build on this new coupling approach. For q-colorings on graphs with maximum degree d, we give a linear-time estimator exists when $q > d+1$, whereas the problem is non-identifiable when $q \leq d+1$. For general H-colorings, we show that standard conditions that guarantee sampling, such as Dobrushin's condition, are insufficient for one-sample learning; on the positive side, we provide a general condition that is sufficient to guarantee linear-time learning and obtain applications for proper colorings and permissive models. For the k-SAT model on formulas with maximum degree d, we provide a linear-time estimator when $k6.45 \log d$, whereas the problem becomes non-identifiable when $k \log d$.

Anthimos Vardis Kandiros
MIT
kandiros@mit.edu

Alkis Kalavasis
Yale University
alvertos.kalavasis@yale.edu

Andreas Galanis
University of Oxford
andreas.galanis@cs.ox.ac.uk

CP22

Online Robust Mean Estimation

We study the problem of high-dimensional robust mean estimation in an online setting. Specifically, we consider a scenario where n sensors are measuring some common, ongoing phenomenon. At each time step $t = 1, 2, \ldots, T$, the i^{th} sensor reports its readings $x_t^{(i)}$ for that time step. The algorithm must then commit to its estimate μ_t for the true mean value of the process at time t. We assume that most of the sensors observe independent samples from some common distribution X, but an ϵ-fraction of them may instead behave maliciously. The algorithm wishes to compute a good approximation μ to the true mean $\mu^* := \mathbf{E}[X]$. We prove two main results about online robust mean estimation in this model. First, if the uncorrupted samples satisfy the standard condition of (ϵ, δ)-stability, we give an efficient online algorithm that outputs estimates μ_t, $t \in [T]$, such that with high probability it holds that $\|\mu - \mu^*\|_2 = O(\delta \log(T))$, where $\mu = (\mu_t)_{t\in[T]}$. We note that this error bound is nearly competitive with the best offline algorithms, which would achieve ℓ_2-error of $O(\delta)$. Our second main result shows that with additional assumptions on the input (most notably that X is a product distribution) there are inefficient algorithms whose error does not depend on T at all.

Daniel Kane
University of California San Diego
dakane@ucsd.edu

Ilias Diakonikolas
UW-Madison
ilias@cs.wisc.edu

Hanshen Xiao
Massachusetts Institute of Technology
hsxiao@mit.edu

Sihan Liu
UC San Diego
sil046@ucsd.edu

CP22

Optimal Rates for Ranking a Permuted Isotonic Matrix in Polynomial Time

We consider a ranking problem where we have noisy observations from a matrix with isotonic columns whose rows

have been permuted by some permutation π^*. This encompasses many models, including crowd-labeling and ranking in tournaments by pair-wise comparisons. In this work, we provide an optimal and polynomial-time procedure for recovering π^*, settling an open problem in [Flammarion et al, Optimal rates of statistical seriation]. As a byproduct, our procedure is used to improve the state-of-the art for ranking problems in the stochastically transitive model (SST). Our approach is based on iterative pairwise comparisons by suitable data-driven weighted means of the columns. These weights are built using a combination of spectral methods with new dimension-reduction techniques. In order to deal with the important case of missing data, we establish a new concentration inequality for sparse and centered rectangular Wishart-type matrices.

Emmanuel Pilliat
Univ. Montpellier
emmanuel.pilliat@umontpellier.fr

Alexandra Carpentier
Institüt für Mathematik
Universität Potsdam
carpentier@uni-potsdam.de

Nicolas Verzelen
MISTEA, INRAE
Univ. Montpellier
nicolas.verzelen@inrae.fr

CP22

How Many Neurons Does It Take to Approximate the Maximum?

We study the size of a neural network needed to approximate the maximum function over d inputs, in the most basic setting of approximating with respect to the L_2 norm, for continuous distributions, for a network that uses ReLU activations. We provide new lower and upper bounds on the width required for approximation across various depths. Our results establish new depth separations between depth 2 and 3, and depth 3 and 5 networks, as well as providing a depth $\mathcal{O}(\log(\log(d)))$ and width $\mathcal{O}(d)$ construction which approximates the maximum function. Our depth separation results are facilitated by a new lower bound for depth 2 networks approximating the maximum function over the uniform distribution, assuming an exponential upper bound on the size of the weights. Furthermore, we are able to use this depth 2 lower bound to provide tight bounds on the number of neurons needed to approximate the maximum by a depth 3 network. Our lower bounds are of potentially broad interest as they apply to the widely studied and used max function, in contrast to many previous results that base their bounds on specially constructed or pathological functions and distributions.

Itay M. Safran
Purdue University
isafran@purdue.edu

Daniel Reichman
Cornell University
daniel.reichman@gmail.com

Paul Valiant
Purdue University
pvaliant@purdue.edu

CP22

The Cost of Parallelizing Boosting

We study the cost of parallelizing weak-to-strong boosting algorithms for learning, following the recent work of Karbasi and Larsen. First, we prove a tight lower bound, showing that even a "slight" parallelization of boosting requires an exponential blow-up in the complexity of training. Specifically, let γ be the weak learner's advantage over random guessing. The famous AdaBoost algorithm produces an accurate hypothesis by interacting with the weak learner for $\tilde{O}(1/\gamma^2)$ rounds, where each round runs in polynomial time. We show any boosting algorithm either has $\Omega(1/\gamma^2)$ rounds of interaction. Otherwise, the boosting algorithm must send $\exp(\Omega(d))$ queries in each round, where d denotes the VC dimension of the hypothesis class. This improves over the $\Omega(1/\gamma)$ lower bound by Karbasi and Larsen. Complementing our lower bound, we show that there exists a boosting algorithm using $\tilde{O}(1/(t\gamma^2))$ rounds, and sending $\exp(O(dt^2))$ queries in each round. This provides the first trade-off between the parallelism and the total work required for boosting. Our lower bound follows from a novel interpretation of parallel boosting as a variant of "coin game". The upper bound is inspired by the "bagging" technique in machine learning and the Advanced Composition theorem from differential privacy.

Hongxun Wu
EECS
UC Berkeley
wuhx@berkeley.edu

Xin Lyu
UC Berkeley
xinlyu@berkeley.edu

Junzhao Yang
Institute for Interdisciplinary Information Sciences
Tsinghua University
yang-jz20@mails.tsinghua.edu.cn

CP23

Grammar Boosting: A New Technique for Proving Lower Bounds for Computation over Compressed Data

Grammar compression is a general framework in which a string T of length N is represented as a context-free grammar of size n whose language contains only T. We focus on studying the limitations of algorithms and data structures operating on grammar-compressed strings. Previous work focused on algorithms that achieve the approximation ratio $\rho = \mathcal{O}(\text{polylog } N)$. Unfortunately, ρ is often either unknown or satisfies $\rho = \omega(\text{polylog } N)$. In their seminal paper, Charikar et al. [IEEE Trans. Inf. Theory 2005] studied seven popular grammar compression algorithms: RePair, Greedy, LongestMatch, Sequential, Bisection, LZ78, and α-Balanced. Only one of them (α-Balanced) is known to achieve $\rho = \mathcal{O}(\text{polylog } N)$. We develop the first technique for proving lower bounds that does not depend on ρ. We prove that $\Omega(\log N/\log\log N)$ time is required for random access on RePair, Greedy, LongestMatch, Sequential, and Bisection, while $\Omega(\log\log N)$ time is required for LZ78. These lower bounds match the existing upper bounds. We then generalize this technique to conditional lower bounds for compressed computation; e.g., we prove for all $c > 0$

and $\epsilon > 0$, that unless the k-Clique Conjecture fails, there is no algorithm for CFG parsing on Bisection ($\rho = \tilde{\Theta}(N^{1/2})$) that runs in $\mathcal{O}(n^c \cdot N^{\omega-\epsilon})$ time. Previously, this was known only for $c < 2\epsilon$.

Rajat De, Dominik Kempa
Stony Brook University
rde@cs.stonybrook.edu, kempa@cs.stonybrook.edu

CP23

Deterministic Sparse Pattern Matching via the Baur-Strassen Theorem

How fast can you test whether a constellation of stars appears in the night sky? This question can be modeled as the computational problem of testing whether a set of points P can be moved into (or close to) another set Q under some prescribed group of transformations. Consider, as a simple representative, the following problem: Given two sets of at most n integers $P, Q \subseteq [N]$, determine whether there is some shift s such that $P + s = \{p + s : p \in P\} \subseteq Q$. This problem, to which we refer as the Constellation problem, can be solved in near-linear time $O(n \log n)$ by a Monte Carlo randomized algorithm [Cardoze, Schulman; FOCS'98] and time $O(n \log^2 N)$ by a Las Vegas randomized algorithm [Cole, Hariharan; STOC'02] and in time $n \cdot 2^{O(\sqrt{\log n \log\log N})}$ by a deterministic algorithm [Chan, Lewenstein;STOC'15]. An interesting open question is whether Constellation is in deterministic near-linear time. We answer this question positively by giving an $n(\log N)^{O(1)}$-time deterministic algorithm for the Constellation problem. Our algorithm extends to various more complex Point Pattern Matching problems in higher dimensions, under translations and rigid motions, and possibly with mismatches, and also to a near-linear-time derandomization of the Sparse Wildcard Matching problem on strings. Interestingly, our algorithm relies an innovative application of the Baur-Strassen theorem from algebraic complexity theory.

Nick Fischer
Weizmann Institute of Science
nick.fischer@weizmann.ac.il

CP23

Near-Optimal Quantum Algorithms for Bounded Edit Distance and Lempel-Ziv Factorization

Classically, the edit distance of two length-n strings can be computed in $O(n^2)$ time, whereas an $O(n^{2-\epsilon})$-time procedure would falsify the Orthogonal Vectors Hypothesis. If the edit distance value does not exceed k, the running time can be improved to $O(n + k^2)$, which is near-optimal (conditioned on OVH) as a function of n and k. Our first main contribution is a quantum $\tilde{O}(\sqrt{nk} + k^2)$-time algorithm that uses $\tilde{O}(\sqrt{nk})$ queries, where the $\tilde{O}(\cdot)$ notation hides polylogarithmic factors. This query complexity is unconditionally optimal, and any significant improvement in the time complexity would resolve the long-standing open question of whether edit distance admits an $O(n^{2-\epsilon})$-time quantum algorithm. The Lempel-Ziv factorization problem can be classically solved in $O(n)$ time, which is unconditionally optimal in the quantum setting. We can, however, hope for a quantum speedup if we parameterize the complexity in terms of z. Already a generic oracle identification algorithm yields the optimal query complexity of $\tilde{O}(\sqrt{nz})$ at the price of exponential running time. Our second main contribution is a quantum algorithm that also achieves the optimal time complexity of $\tilde{O}(\sqrt{nz})$. We can then obtain the string's run-length encoded Burrows-Wheeler Transform (BWT), construct the r-index, and solve many fundamental string processing problems in time $\tilde{O}(\sqrt{nz})$.

Daniel Gibney
University of Texas at Dallas
daniel.j.gibney@gmail.com

Ce Jin
EECS
MIT
cejin@mit.edu

Tomasz Kociumaka
Max Planck Institute for Informatics
tomasz.kociumaka@mpi-inf.mpg.de

Sharma Thankachan
North Carolina State University
sharmavt.nitc@gmail.com

CP23

Sparse Regular Expression Matching

A regular expression specifies a set of strings formed by single characters combined with concatenation, union, and Kleene star operators. Given a regular expression R and a string Q, the regular expression matching problem is to decide if Q matches any of the strings specified by R. Regular expressions are a fundamental concept in formal languages and regular expression matching is a basic primitive for searching and processing data. A standard textbook solution [Thompson, CACM 1968] constructs and simulates a NFA, leading to an $O(nm)$ time algorithm, where $n = |Q|$ and $m = |R|$. Only polylogarithmic improvements of this bound are known. Backurs and Indyk [FOCS 2016] proved that, assuming SETH, regular expression matching cannot be solved in $O((nm)^{1-\epsilon})$, for any constant $\epsilon > 0$. We take a new approach and introduce a *density* parameter, Δ, that captures the amount of nondeterminism in the NFA simulation on Q. The density is at most $nm + 1$ but can be significantly smaller. Our main result is a new algorithm that solves regular expression matching in

$$O\left(\Delta \log\log \frac{nm}{\Delta} + n + m\right)$$

time. This essentially replaces nm with Δ in the complexity of regular expression matching. We complement our upper bound by a matching conditional lower bound that proves that we cannot solve regular expression matching in time $O(\Delta^{1-\epsilon})$ for any constant $\epsilon > 0$ assuming SETH.

Inge Li Gørtz
Danish Technical University, Denmark
inge@dtu.dk

Philip Bille
Technical University of Denmark
phbi@dtu.dk

CP23

Faster Sublinear-Time Edit Distance

We study the fundamental problem of approximating the edit distance of two strings. After the development of a constant-factor approximation in almost-linear time, recent years have witnessed a notable shift towards sublinear-time

algorithms. Here, the task is formalized as the (k, K)-gap edit distance problem: Distinguish whether the edit distance of two strings is at most k or more than K. Surprisingly, it is still possible to obtain meaningful results in this regime. Nevertheless, in almost all previous work, truly sublinear running time of $O(n^{1-\epsilon})$ (for a constant $\epsilon > 0$) comes at the price of polynomial gap $K \geq k \cdot n^{\Omega(\epsilon)}$. Only recently, [Bringmann, Cassis, Fischer, and Nakos; STOC 2022] broke through this barrier and solved the $(k, k^{1+o(1)})$-gap edit distance problem in time $O(n/k + k^{4+o(1)})$, which is truly sublinear if $n^{\Omega(1)} \leq k \leq n^{\frac{1}{4}-\Omega(1)}$. The n/k term is inevitable, but it remains an important task to optimize the poly(k) term and, in general, solve the $(k, k^{1+o(1)})$-gap edit distance problem in sublinear time for larger k. In this work, we design an $O(n/k + k^{2+o(1)})$-time algorithm for $(k, k^{1+o(1)})$-gap edit distance, yielding a significant quadratic speed-up over the previous algorithm. Notably, our algorithm is unconditionally almost-optimal (up to subpolynomial factors) for $k \leq n^{\frac{1}{3}}$ and improves upon the state of the art for $k \leq n^{\frac{1}{2}-o(1)}$.

Karl Bringmann
Saarland University and Max Planck Institute for Informatics
Saarland Informatics Campus, Germany
bringmann@cs.uni-saarland.de

Alejandro Cassis
Saarland University and Max Planck Institute for Informatics
acassis@cs.uni-saarland.de

Nick Fischer
Weizmann Institute of Science
nick.fischer@weizmann.ac.il

Tomasz Kociumaka
Max Planck Institute for Informatics
tomasz.kociumaka@mpi-inf.mpg.de

CP24

Dynamic Algorithms for Matroid Submodular Maximization

Submodular maximization under matroid and cardinality constraints are classical problems with a wide range of applications in machine learning, auction theory, and combinatorial optimization. In this paper, we consider these problems in the dynamic setting where (1) we have oracle access to a monotone submodular function $f : 2^V \to \mathbb{R}^+$ and (2) we are given a sequence $\mathcal{S}$ of insertions and deletions of elements of an underlying ground set V. We develop the first fully dynamic $(4 + \epsilon)$-approximation algorithm for the submodular maximization problem under the matroid constraint using an expected worst-case $O(k \log(k) \log^3(k/\epsilon))$ query complexity where $0 < \epsilon \leq 1$. This resolves an open problem of Chen and Peng (STOC'22) and Lattanzi et al. (NeurIPS'20). As a byproduct, for the submodular maximization under the cardinality constraint k, we propose a parameterized (by the cardinality constraint k) dynamic algorithm that maintains a $(2 + \epsilon)$-approximate solution of the sequence $\mathcal{S}$ at any time t using an expected worst-case query complexity $O(k\epsilon^{-1} \log^2(k))$. This is the first dynamic algorithm for the problem that has a query complexity independent of the size of ground set V.

Kiarash Banihashem
University of Maryland
kiarash@umd.edu

Leyla Biabani
Eindhoven University of Technology
l.biabani@tue.nl

Samira Goudarzi, MohammadTaghi Hajiaghayi, Peyman Jabbarzade
University of Maryland
samirag@umd.edu, hajiagha@umd.edu, peymanj@umd.edu

Morteza Monemizadeh
Computer Science Institute
Charles University, Prague, Czech Republic
m.monemizadeh@gmail.com

CP24

Fully Dynamic Consistent K-Center Clustering

We study the consistent k-center clustering problem. In this problem, the goal is to maintain a constant factor approximate k-center solution during a sequence of n point insertions and deletions while minimizing the recourse, i.e., the number of changes made to the set of centers after each point insertion or deletion. Previous works by Lattanzi and Vassilvitskii [ICML 12] and Fichtenberger, Lattanzi, NorouziFard, and Svensson [SODA 21] showed that in the incremental setting, where deletions are not allowed, one can obtain k polylog(n)/n amortized recourse for both k-center and k-median, and demonstrated a matching lower bound. However, no algorithm for the fully dynamic setting achieves less than the trivial O(k) changes per update, which can be obtained by simply reclustering the full dataset after every update. In this work, we give the first algorithm for consistent k-center clustering for the fully dynamic setting, i.e., when both point insertions and deletions are allowed, and improves upon a trivial O(k) recourse bound. Specifically, our algorithm maintains a constant factor approximate solution while ensuring worstcase constant recourse per update, which is optimal in the fully dynamic setting. Moreover, our algorithm is deterministic and is therefore correct even if an adaptive adversary chooses the insertions and deletions.

Christoph Grunau
ETH Zurich
cgrunau@student.ethz.ch

Jakub Lacki
Google Research
jlacki@google.com

Bernhard Haeupler, Václav Rozhon
ETH Zurich
bernhard.haeupler@inf.ethz.ch, rozhonv@ethz.ch

Rajesh Jayaram
Google Research
rkjayaram@google.com

CP24

On Dynamic Graph Algorithms with Predictions

Dynamic algorithms operate on inputs undergoing updates, e.g., insertions or deletions of edges or vertices. After processing each update, the algorithm has to an-

swer queries regarding the current state of the input data. We study dynamic algorithms in the model of algorithms with predictions (also known as learning-augmented algorithms). We assume the algorithm is given imperfect predictions regarding future updates, and we ask how such predictions can be used to improve the running time. In other words, we study the complexity of dynamic problems parameterized by the prediction accuracy. This can be seen as a model interpolating between classic online dynamic algorithms – which know nothing about future updates – and offline dynamic algorithms with the whole update sequence known upfront, which is similar to having perfect predictions. Our results give smooth tradeoffs between these two extreme settings.

Adam Polak
Bocconi University, Milan, Italy
pl.adam.polak@gmail.com

Jan Van Den Brand
Georgia Tech
vdbrand@gatech.edu

Sebastian Forster
University of Salzburg
forster@cs.sbg.ac.at

Yasamin Nazari
VU Amsterdam
y.nazari@vu.nl

CP24

Dynamic Algorithms for k-Center on Graphs

In this paper we give the first efficient algorithms for the k-center problem on dynamic graphs undergoing edge updates. In this problem, the goal is to partition the input into k sets by choosing k centers such that the maximum distance from any data point to the closest center is minimized. It is known that it is NP-hard to get a better than 2 approximation for this problem. While in many applications the input may naturally be modeled as a graph, all prior works on k-center problem in dynamic settings are on metrics. In this paper, we give a deterministic decremental $(2+\epsilon)$-approximation algorithm and a randomized incremental $(4+\epsilon)$-approximation algorithm, both with amortized update time $kn^{o(1)}$ for weighted graphs. Moreover, we show a reduction that leads to a fully dynamic $(2+\epsilon)$-approximation algorithm for the k-center problem, with worst-case update time that is within a factor k of the state-of-the-art upper bound for maintaining $(1+\epsilon)$-approximate single-source distances in graphs. Matching this bound is a natural goalpost because the approximate distances of each vertex to its center can be used to maintain a $(2+\epsilon)$-approximation of the graph diameter and the fastest known algorithms for such a diameter approximation also rely on maintaining approximate single-source distances.

Emilio Cruciani, Sebastian Forster
University of Salzburg
emilio.cruciani@plus.ac.at, forster@cs.sbg.ac.at

Gramoz Goranci
University of Vienna
gramoz.goranci@univie.ac.at

Yasamin Nazari
VU Amsterdam
y.nazari@vu.nl

Antonis Skarlatos
University of Salzburg
antonis.skarlatos@plus.ac.at

CP24

Nibbling at Long Cycles: Dynamic (and Static) Edge Coloring in Optimal Time

We consider the problem of maintaining a $(1+\epsilon)\Delta$-edge coloring in a dynamic graph G with n nodes and maximum degree at most Δ. The state-of-the-art update time is $O_\epsilon(\texttt{polylog}(n))$, by Duan, He and Zhang [SODA'19] and by Christiansen [STOC'23], and more precisely $O(\log^7 n/\epsilon^2)$, where $\Delta = \Omega(\log^2 n/\epsilon^2)$. The following natural question arises: What is the best possible update time of an algorithm for this task? More specifically, **can we bring it all the way down to some constant** (for constant ϵ)? This question coincides with the *static* time barrier for the problem: Even for $(2\Delta-1)$-coloring, there is only a naive $O(m\log\Delta)$-time algorithm. We answer this fundamental question in the affirmative, by presenting a dynamic $(1+\epsilon)\Delta$-edge coloring algorithm with $O(\log^4(1/\epsilon)/\epsilon^9)$ update time, provided $\Delta = \Omega_\epsilon(\texttt{polylog}(n))$. As a corollary, we also get the first linear time (for constant ϵ) *static* algorithm for $(1+\epsilon)\Delta$-edge coloring; in particular, we achieve a running time of $O(m\log(1/\epsilon)/\epsilon^2)$. We obtain our results by carefully combining a variant of the Nibble algorithm from Bhattacharya, Grandoni and Wajc [SODA'21] with the subsampling technique of Kulkarni, Liu, Sah, Sawhney and Tarnawski [STOC'22].

Sayan Bhattacharya
University of Warwick
Coventry
S.Bhattacharya@warwick.ac.uk

Martín Costa
University of Warwick
martin.costa@warwick.ac.uk

Nadav Panski, Shay Solomon
Tel Aviv University
nadavpanski@mail.tau.ac.il, solo.shay@gmail.com

CP25

Fast Algorithms for Separable Linear Programs

A fundamental topic of study in numerical linear algebra is the design of efficient solvers for linear systems that exhibit structural properties. A prominent success story is the method of generalized nested dissection [LRT'79] for separable matrices. On the other hand, most recent developments in linear programs (LP) have not leveraged the ideas underlying linear system solvers, nor exploited the separable structure of the constraint matrix. In this work, we consider LPs of the form $\min_{Ax=b,\ell\le x\le u} c^\top x$, where the graphical support of the constraint matrix $A\in\mathbb{R}^{n\times m}$ is n^α-separable. We present an $\tilde{O}((m+m^{1/2+2\alpha})\log(1/\epsilon))$-time algorithm to solve such LPs to ϵ relative accuracy. Our solver has two important implications: for the k-multicommodity flow problem on planar graphs, we obtain an $\tilde{O}(k^{5/2}m^{3/2}\log(1/\epsilon))$-time algorithm; and when the support of A is n^α-separable with $\alpha\le 1/4$, our runtime of $\tilde{O}(m)$ is nearly optimal, in contrast with the natural approach of combining interior point methods and nested

dissection, whose time complexity is lower bounded by $\Omega(\sqrt{m}(m+m^{\alpha\omega}))=\Omega(m^{3/2})$, where ω is the matrix multiplication constant. Lastly, our solver can be applied to low-treewidth LPs to recover the results of [DLY21,GS22] while using significantly simpler data structure machinery.

Sally Dong
University of Washington
sallyqd@uw.edu

Gramoz Goranci
University of Vienna
gramoz.goranci@univie.ac.at

Lawrence Li, Sushant Sachdeva
University of Toronto
lawrenceli@cs.toronto.edu, sachdeva@cs.toronto.edu

Guanghao Ye
Massachusetts Institute of Technology
ghye@mit.edu

CP25

Integer Programming with GCD Constraints

We study the non-linear extension of integer programming with greatest common divisor constraints of the form $\gcd(f,g)\sim d$, where f and g are linear polynomials, d is a positive integer, and $\sim$ is a relation among $\leq,=,\neq$ and $\geq$. We show that the feasibility problem for these systems is in NP, and that an optimal solution minimizing a linear objective function, if it exists, has polynomial bit length. To show these results, we identify an expressive fragment of the existential theory of the integers with addition and divisibility that admits solutions of polynomial bit length. It was shown by Lipshitz [Trans. Am. Math. Soc., 235, pp. 271-283, 1978] that this theory adheres to a local-to-global principle in the following sense: a formula Φ is equi-satisfiable with a formula Ψ in this theory such that Ψ has a solution if and only if Ψ has a solution modulo every prime p. We show that in our fragment, only a polynomial number of primes of polynomial bit length need to be considered, and that the solutions modulo prime numbers can be combined to yield a solution to Φ of polynomial bit length. As a technical by-product, we establish a Chinese-remainder-type theorem for systems of congruences and non-congruences showing that solution sizes do not depend on the magnitude of the moduli of non-congruences.

Rémy Défossez
IMDEA Software Institute and École Normale supérieure
remy.defossez@ens.psl.eu

Christoph Haase
Oxford University
christoph.haase@cs.ox.ac.uk

Alessio Mansutti
IMDEA Software Institute
alessio.mansutti@imdea.org

Guillermo A. Pérez
University of Antwerp Flanders Make, Belgium
guillermo.perez@uantwerpen.be

CP25

Sub-Exponential Lower Bounds for Branch-and-Bound with General Disjunctions via Interpolation

We investigate linear programming based branch-and-bound using general disjunctions, also known as stabbing planes, for solving integer programs. We derive the first sub-exponential lower bound (in the encoding length L of the integer program) for the size of a general branch-and-bound tree for a particular class of (compact) integer programs, namely $2^{\Omega(L^{1/12-\epsilon})}$ for every $\epsilon>0$. This is achieved by showing that general branch-and-bound admits quasi-feasible monotone real interpolation, which allows us to utilize sub-exponential lower-bounds for monotone real circuits separating the so-called clique-coloring pair. The same ideas also prove that refuting $\Theta(\log(n))$-CNFs requires size $2^{n^{\Omega(1)}}$ branch-and-bound trees with high probability by considering the closely related notion of infeasibility certificates recently introduced by Hrube and Pudl One important ingredient of the proof of our interpolation result is that for every general branch-and-bound tree proving integer-freeness of a product $P\times Q$ of two polytopes P and Q, there exists a closely related branch-and-bound tree for showing integer-freeness of P or one showing integer-freeness of Q. Moreover, we prove that monotone real circuits can perform binary search efficiently.

Maximilian Gläser
TU Darmstadt
glaeser@mathematik.tu-darmstadt.de

Marc Pfetsch
Research Group Optimization, Department of Mathematics
TU Darmstadt
pfetsch@mathematik.tu-darmstadt.de

CP25

A Whole New Ball Game: A Primal Accelerated Method for Matrix Games and Minimizing the Maximum of Smooth Functions

We design algorithms for minimizing $\max_{i\in[n]} f_i(x)$ over a d-dimensional Euclidean or simplex domain. When each f_i is 1-Lipschitz and 1-smooth, our method computes an ϵ-approximate solution using $\widetilde{O}(n\epsilon^{-1/3}+\epsilon^{-2})$ gradient and function evaluations, and $\widetilde{O}(n\epsilon^{-4/3})$ additional runtime. For large n, our evaluation complexity bound is optimal up to polylogarithmic factors. In the special case where each f_i is linear—which corresponds to finding a near-optimal primal strategy in a matrix game –our method finds an ϵ-approximate solution in runtime $\widetilde{O}(n(d/\epsilon)^{2/3}+\epsilon^{-2})$. For $n>d$ and $\epsilon=1/\sqrt{n}$ this improves over all existing first-order methods. When additionally $d=\omega(n^{8/11})$ our runtime also improves over all known interior point methods. Our algorithm combines three novel primitives: (1) A dynamic data structure which enables efficient stochastic gradient estimation in small ℓ_2 or ℓ_1 balls. (2) A mirror descent algorithm tailored to our data structure implementing an oracle which minimizes the objective over said ball. (3) A simple ball oracle acceleration framework suitable for non-Euclidean geometry.

Arun Jambulapati
University of Washington
jmblpati@gmail.com

Yair Carmon
Tel Aviv University
Computer Science

ycarmon@tauex.tau.ac.il

Aaron Sidford, Yujia Jin
Stanford University
sidford@stanford.edu, yujiajin@stanford.edu

CP25

Arborescences, Colorful Forests, and Popularity

Our input is a directed, rooted graph $G = (V \cup \{r\}, E)$ where each vertex in V has a partial order preference over its incoming edges. The preferences of a vertex extend naturally to preferences over arborescences rooted at r. We seek a *popular* arborescence in G, i.e., one for which there is no "more popular" arborescence. The *popular arborescence* problem is to decide if a given input instance admits a popular arborescence or not. We present the first polynomial-time algorithm for this problem. Our algorithm is combinatorial, and can be regarded as a primal-dual algorithm. It searches for an arborescence along with its dual certificate, a chain of subsets of E, witnessing its popularity. In fact, our algorithm solves the more general *popular common base* problem in the intersection of two matroids, where one matroid is the partition matroid and the other is an arbitrary matroid. We also study the related *popular colorful forest* (or more generally, the *popular common independent set*) problem where edges are partitioned into color classes, and the task is to find a colorful forest that is popular within the set of all colorful forests. For the case with weak rankings, we formulate the popular colorful forest polytope, and thus show that a minimum-cost popular colorful forest can be computed efficiently. By contrast, we prove that it is NP-hard to compute a minimum-cost popular arborescence, even when rankings are strict.

Telikepalli Kavitha
Tata Institute of Fundamental Research
Mumbai, India
kavitha@tifr.res.in

Kazuhisa Makino
Kyoto University
makino@kurims.kyoto-u.ac.jp

Ildikó Schlotter
Centre for Economic and Regional Studies
schlotter.ildiko@krtk.hun-ren.hu

Yu Yokoi
Tokyo Institute of Technology
yokoi@c.titech.ac.jp

CP25

Convex Minimization with Integer Minima in $\widetilde{O}(n^4)$ Time

Given a convex function f on $\mathbb{R}^n$ with an integer minimizer, we show how to find an exact minimizer of f using $O(n^2 \log n)$ calls to a separation oracle and $O(n^4 \log n)$ time. The previous best polynomial time algorithm for this problem given in [Jiang, SODA 2021, JACM 2022] achieves $O(n^2 \log\log n/\log n)$ oracle complexity. However, the overall runtime of Jiang's algorithm is at least $\widetilde{\Omega}(n^8)$, due to expensive sub-routines such as the Lenstra-Lenstra-Lovasz (LLL) algorithm [Lenstra, Lenstra, Lovasz, Math. Ann. 1982] and random walk based cutting plane method [Bertsimas, Vempala, JACM 2004]. Our significant speedup is obtained by a nontrivial combination of a faster version of the LLL algorithm due to [Neumaier, Stehle, ISSAC 2016] that gives similar guarantees, the volumetric center cutting plane method (CPM) by [Vaidya, FOCS 1989] and its fast implementation given in [Jiang, Lee, Song, Wong, STOC 2020]. For the special case of submodular function minimization (SFM), our result implies a strongly polynomial time algorithm for this problem using $O(n^3 \log n)$ calls to an evaluation oracle and $O(n^4 \log n)$ additional arithmetic operations. Both the oracle complexity and the number of arithmetic operations of our more general algorithm are better than the previous best-known runtime algorithms for this specific problem given in [Lee, Sidford, Wong, FOCS 2015] and [Dadush, Vegh, Zambelli, SODA 2018, MOR 2021].

Haotian Jiang
University of Washington
jhtdavid96@gmail.com

Yin Tat Lee
Microsoft Research
University of Washington
yintat@uw.edu

Zhao Song
Adobe Research
zsong@adobe.com

Lichen Zhang
Massachusetts Institute of Technology
lichenz@mit.edu

CP26

Fast Fourier Transform via Automorphism Groups of Rational Function Fields

The Fast Fourier Transform (FFT) over a finite field $\mathbb{F}_q$ computes evaluations of a given polynomial of degree less than n at a specifically chosen set of n distinct evaluation points in $\mathbb{F}_q$. If q or $q-1$ is a smooth number, then the divide-and-conquer approach leads to the fastest known FFT algorithms. Depending on the type of group that the set of evaluation points forms, these algorithms are classified as multiplicative (Math of Comp. 1965) and additive (FOCS 2014) FFT algorithms. In this work, we provide a unified framework for FFT algorithms that include both multiplicative and additive FFT algorithms as special cases, and beyond: our framework also works when $q+1$ is smooth, while all known results require q or $q-1$ to be smooth. For the new case where $q+1$ is smooth (this new case was not considered before in literature as far as we know), we show that if n is a divisor of $q+1$ that is B-smooth for a real $B > 0$, then our FFT needs $O(Bn \log n)$ arithmetic operations in $\mathbb{F}_q$. Our unified framework is a natural consequence of introducing the algebraic function fields into the study of FFT.

Songsong Li, Chaoping Xing
Shanghai Jiao Tong University
songsli@sjtu.edu.cn, xingcp@sjtu.edu.cn

Omar Alrabiah
UC Berkeley
oalrabiah@berkeley.edu

CP26

The Identity Problem in Nilpotent Groups of

Bounded Class

Let G be a unitriangular matrix group of nilpotency class at most ten. We show that the Identity Problem (does a semigroup contain the identity matrix?) and the Group Problem (is a semigroup a group?) are decidable in polynomial time for finitely generated subsemigroups of G. Our decidability results also hold when G is an arbitrary finitely generated nilpotent group of class at most ten. This extends earlier work of Babai et al. on commutative matrix groups (SODA'96) and work of Bell et al. on $\mathsf{SL}(2,\mathbb{Z})$ (SODA'17). Furthermore, we formulate a sufficient condition for the generalization of our results to nilpotent groups of class $d > 10$. For every such d, we exhibit an effective procedure that verifies this condition in case it is true.

Ruiwen Dong
University of Oxford
ruiwen.dong@kellogg.ox.ac.uk

CP26

Faster Rectangular Matrix Multiplication by Combination Loss Analysis

Duan, Wu and Zhou (FOCS 2023) recently obtained the improved upper bound on the exponent of square matrix multiplication $\omega < 2.3719$ by introducing a new approach to quantify and compensate the "combination loss" in prior analyses of powers of the Coppersmith-Winograd tensor. In this paper we show how to use this new approach to improve the exponent of rectangular matrix multiplication as well. Our main technical contribution is showing how to combine this analysis of the combination loss and the analysis of the fourth power of the Coppersmith-Winograd tensor in the context of rectangular matrix multiplication developed by Le Gall and Urrutia (SODA 2018).

François Le Gall
Nagoya University
legall@math.nagoya-u.ac.jp

CP26

New Nearly Optimal Black Box Polynomial Root-Finders

Univariate polynomial root-finding has been studied for four millennia and very intensively in the last decades. Our novel nearly optimal Las Vegas randomized root-finders approximate all zeros of a polynomial almost as fast as one accesses its coefficients with the precision required for the solution within a prescribed error bound. Moreover, our root-finders can be applied to a black box polynomial, defined by an oracle (that is, black box subroutine) for its evaluation rather than by its coefficients. Such root-finders are particularly fast for polynomials that can be evaluated fast, e.g., the sum of a few shifted monomials, but the only other known black box root-finder is the pioneering one by Louis and Vempala at FOCS 2016, and it only approximates the absolutely largest root of a real-rooted polynomial. Our divide and conquer deterministic algorithm of ACM STOC 1995 is the only other known nearly optimal polynomial root-finder, but it extensively uses the coefficients, is quite involved, and has never been implemented, while according to extensive numerical experiments with standard test polynomials, our new root-finders competes with user's choice package MPSolve and supersedes it more and more significantly where the polynomial degree grows large. We also approximate matrix eigenvalues in a record Las Vegas bit operation time. Our auxiliary algorithms and techniques for computations with black box polynomials can be of independent interest.

Victor Pan
CUNY
Lehman
victor.pan@lehman.cuny.edu

CP26

Deterministic Algorithms for Low-Degree Factors of Constant-Depth Circuits

For every constant d, we design a deterministic algorithm that takes as input a multivariate polynomial given f as a constant depth algebraic circuit over the field of rational numbers, and outputs all irreducible factors of f of degree at most d together with their respective multiplicities. The algorithm runs in subexponential time for every constant d. Moreover, if f is a sparse polynomial, then the algorithm runs in quasipolynomial time. Our results are based on a more fine grained connection between polynomial identity testing (PIT) and polynomial factorization in the context of constant-degree factors and rely on a clean connection between divisibility testing of polynomials and PIT due to Forbes and on subexponential time deterministic PIT algorithms for constant-depth algebraic circuits from the recent work of Limaye, Srinivasan and Tavenas.

Varun Ramanathan, Mrinal Kumar
Tata Institute of Fundamental Research, Mumbai
varun.ramanathan@tifr.res.in, mrinal@tifr.res.in

Ramprasad Saptharishi
Tata Institute of Fundamental Research, Mumbai, India
ramprasad@tifr.res.in

CP26

New Bounds for Matrix Multiplication: from Alpha to Omega

The main contribution of this paper is a new improved variant of the laser method for designing matrix multiplication algorithms. Building upon the recent techniques of [Duan, Wu, Zhou, FOCS 2023], the new method introduces several new ingredients that not only yield an improved bound on the matrix multiplication exponent ω, but also improve the known bounds on rectangular matrix multiplication by [Le Gall and Urrutia, SODA 2018]. In particular, the new bound on ω is $\omega \leq 2.371552$ (improved from $\omega \leq 2.371866$). For the dual matrix multiplication exponent α defined as the largest α for which $\omega(1,\alpha,1) = 2$, we obtain the improvement $\alpha \geq 0.321334$ (improved from $\alpha \geq 0.31389$). Similar improvements are obtained for various other exponents for multiplying rectangular matrices.

Virginia Vassilevska Williams
MIT
virgi@mit.edu

Yinzhan Xu
Massachusetts Institute of Technology
xyzhan@mit.edu

Zixuan Xu
MIT
zixuanxu@mit.edu

Renfei Zhou
Tsinghua University

zhourf20@mails.tsinghua.edu.cn

CP27

A Distributed Palette Sparsification Theorem

The celebrated palette sparsification result of [Assadi, Chen, and Khanna SODA'19] shows that to compute a $\Delta + 1$ coloring of the graph, where Δ denotes the maximum degree, it suffices if each node limits its color choice to $O(\log n)$ independently sampled colors in $\{1, 2, \ldots, \Delta + 1\}$. They showed that it is possible to color the resulting sparsified graph and obtain a $\Delta+1$ coloring for the original graph. However, to compute the actual coloring, that information must be gathered at a single location for centralized processing. We seek instead a local algorithm to compute such a coloring in the sparsified graph in poly$(\log n)$ distributed rounds with small messages. Our main result is an algorithm that computes a $\Delta + 1$-coloring after palette sparsification with $O(\log^2 n)$ random colors per node and runs in $O(\log^2 \Delta + \log^3 \log n)$ rounds on the sparsified graph, using $O(\log n)$-bit messages. We show that this is close to the best possible: any distributed $\Delta + 1$-coloring algorithm that runs in the LOCAL model on the sparsified graph, given by palette sparsification, for any poly$(\log n)$ colors per node, requires $\Omega(\log \Delta / \log \log n)$ rounds. This distributed palette sparsification result leads to the first poly$(\log n)$-round algorithms for $\Delta+1$-coloring in two previously studied distributed models: the Node Capacitated Clique, and the cluster graph model.

Maxime Flin
Reykjavik University
maximef@ru.is

Mohsen Ghaffari
MIT
ghaffari@mit.edu

Magnús M. Halldórsson
Reykjavik University
mmh@ru.is

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

Alexandre Nolin
CISPA Helmholtz Center for Information Security
alexandre.nolin@cispa.de

CP27

A Nearly Linear-Time Distributed Algorithm for Exact Maximum Matching

In this paper, we propose a randomized $\tilde{O}(\mu(G))$-round algorithm for the maximum cardinality matching problem in the CONGEST model, where $\mu(G)$ means the maximum size of a matching of the input graph G. The proposed algorithm substantially improves the current best worst-case running time. The key technical ingredient is a new randomized algorithm of finding an augmenting path of length ℓ with high probability within $\tilde{O}(\ell)$ rounds, which positively settles an open problem left in the prior work by Ahmadi and Kuhn [DISC'20].The idea of our augmenting path algorithm is based on a recent result by Kitamura and Izumi [IEICE Trans.'22], which efficiently identifies a sparse substructure of the input graph containing an augmenting path, following a new concept called alternating base trees. Their algorithm, however, resorts to a centralized approach of collecting the entire information of the substructure into a single vertex for constructing an augmenting path. The technical highlight of this paper is to provide a fully-decentralized counterpart of such a centralized method. To develop the algorithm, we prove several new structural properties of alternating base trees, which are of independent interest.

Taisuke Izumi, Naoki Kitamura, Yutaro Yamaguchi
Osaka University
izumi.taisuke.ist@osaka-u.ac.jp, n-kitamura@ist.osaka-u.ac.jp, yutaro.yamaguchi@ist.osaka-u.ac.jp

CP27

Parallel Approximate Maximum Flows in Near-Linear Work and Polylogarithmic Depth

We give a parallel algorithm for the $(1 - \epsilon)$-approximate maximum flow problem in capacitated, undirected graphs with n vertices and m edges, achieving $\mathrm{O}(\epsilon^{-3}\mathrm{polylog}\ n)$ depth and $O(m\epsilon^{-3}\mathrm{polylog}\ n)$ work in the PRAM model. Although near-linear time sequential algorithms for this problem have been known for almost a decade, no parallel algorithms that simultaneously achieved polylogarithmic depth and near-linear work were known. At the heart of our result is a polylogarithmic depth, near-linear work recursive algorithm for computing congestion approximators, which only requires solving maximum flows on subgraphs obtained by contracting vertices, as opposed to vertex-induced subgraphs used in Re, Shah, and Tig [SODA14]. Along the way, we also develop a parallel flow-decomposition algorithm that is crucial to achieving polylogarithmic depth and may be of independent interest. We extend our results to related graph problems such as sparsest and balanced sparsest cuts, fair and isolating cuts, approximate Gomory-Hu trees, and hierarchical clustering. All algorithms achieve polylogarithmic depth and near-linear work. Finally, our PRAM results also imply the first polylogarithmic round, near-linear total space MPC algorithms for approximate undirected maximum flows, as well as all its aforementioned applications in the fully scalable regime where the local machine memory is $O(n^{\delta})$ for any constant $\delta > 0$.

Prathamesh Patil, Huan Li, Sanjeev Khanna, Nathan White
University of Pennsylvania
pprath@seas.upenn.edu, huanli@seas.upenn.edu, sanjeev@cis.upenn.edu, nathanlw@seas.upenn.edu

Arpit Agarwal
Columbia University
aa4931@columbia.edu

Chen Wang
Rutgers University
chen.wang.cs@rutgers.edu

Peilin Zhong
Google Research
peilinz@google.com

CP27

Breaking 3-Factor Approximation for Correlation Clustering in Polylogarithmic Rounds

In this paper, we study parallel algorithms for the correlation clustering problem, where every pair of two different

entities is labeled with similar or dissimilar. The goal is to partition the entities into clusters to minimize the number of disagreements with the labels. Currently, all efficient parallel algorithms have an approximation ratio of at least 3. In comparison with the $1.994 + \epsilon$ ratio achieved by polynomial-time sequential algorithms [CLN22], a significant gap exists. We propose the first poly-logarithmic round parallel algorithm that achieves a better approximation ratio than 3. Specifically, our algorithm computes a $(2.4+\epsilon)$-approximate solution and uses $\tilde{O}(m^{1.5})$ work. Additionally, it can be translated into a $\tilde{O}(m^{1.5})$-time sequential algorithm and a poly-logarithmic rounds sublinear-memory MPC algorithm with $\tilde{O}(m^{1.5})$ total memory. Our approach is inspired by Awerbuch, Khandekar, and Rao's [AKR12] length-constrained multi-commodity flow algorithm, where we develop an efficient parallel algorithm to solve a truncated correlation clustering linear program of Charikar, Guruswami, and Wirth [CGW05]. Then we show the solution of the truncated linear program can be rounded with a factor of at most 2.4 loss by using the framework of [CMSY15]. Such a rounding framework can then be implemented using parallel pivot-based approaches.

Hsin-Hao Su
Boston College
Computer Science Department
suhx@bc.edu

Nairen Cao
Boston College
Department of Computer Science
caonc@bc.edu

Shang-En Huang
Boston College
huangaul@bc.edu

CP27

Massively Parallel Algorithms for High-Dimensional Euclidean Minimum Spanning Tree

We study the classic Euclidean Minimum Spanning Tree (MST) problem in the Massively Parallel Computation (MPC) model. Given a set $X \subset \mathbb{R}^d$ of n points, the goal is to produce a spanning tree for X with weight within a small factor of optimal. Euclidean MST is one of the most fundamental hierarchical geometric clustering algorithms, and with the proliferation of enormous high-dimensional data sets, such as massive transformer-based embeddings, there is now a critical demand for efficient distributed algorithms to cluster such data sets. In this paper, we give a constant factor approximation in $\tilde{O}(\log \log n)$ rounds of the MPC model. In contrast to tree-embedding-based approaches, which necessarily must pay $\Omega(\log n)$-distortion, our algorithm is based on a new combination of graph-based distributed MST algorithms and geometric space partitions. Additionally, although the approximate MST we return can have a large depth, we show that it can be modified to obtain a $\tilde{O}(\log \log n)$-round constant factor approximation to the Euclidean Traveling Salesman Problem (TSP) in the MPC model. Previously, only a $O(\log n)$ round was known for the problem.

Peilin Zhong, Rajesh Jayaram
Google Research
peilinz@google.com, rkjayaram@google.com

Vahab Mirrokni
Google Inc.
mirrokni@google.com

Shyam Narayanan
Massachusetts Institute of Technology
shyamsn@mit.edu

CP27

The Minority Dynamics and the Power of Synchronicity

We study the minority-opinion dynamics over a complete network of n nodes with binary opinions. Upon activation, a node receives a sample of opinions from a limited number of neighbors chosen uniformly at random. Each activated node then adopts the opinion that is least common within the received sample. Unlike all other known consensus dynamics, we prove that this elementary protocol behaves in dramatically different ways, depending on whether activations occur sequentially or in parallel. Specifically, we show that its expected consensus time is exponential in n under asynchronous models, such as asynchronous GOSSIP. On the other hand, despite its chaotic nature, we show that it converges within $O(\log^2 n)$ rounds w.h.p. under synchronous models, such as synchronous GOSSIP. Finally, our results shed light on the bit-dissemination problem, previously introduced to model the spread of information in biological scenarios. Our analysis implies that the minority-opinion dynamics is the first stateless solution to this problem, in the parallel passive-communication setting, achieving convergence within a polylogarithmic number of rounds. This, together with a known lower bound for sequential stateless dynamics, implies a parallel-vs-sequential gap for this problem that is nearly quadratic in the number n of nodes. This is in contrast to all known results for problems in this area, which exhibit a linear gap between the parallel and the sequential setting.

Isabella Ziccardi
Bocconi University
isabella.ziccardi@unibocconi.it

Luca Becchetti
Sapienza University of Rome
becchetti@diag.uniroma1.it

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

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

Luca Trevisan
Bocconi University
l.trevisan@unibocconi.it

Robin Vacus
CNRS, IRIF, Paris
rvacus@irif.fr

CP28

Breaking the $k/\log k$ Barrier in Collective Tree Exploration via Tree-Mining

In collective tree exploration, a team of k mobile agents is assigned to go through all edges of an unknown tree as

fast as possible. An edge of the tree is revealed to the team when one agent becomes adjacent to that edge. The exploration algorithm's guarantee is typically compared to the number of rounds required to traverse all edges if the agents had known the tree in advance. This quantity is at least $\max\{2n/k, 2D\}$ where n is the number of nodes and D is the tree depth. Since the introduction of the problem by [11], two types of guarantees have emerged: the first takes the form $r(k)(n/k+D)$, where $r(k)$ is called the competitive ratio, and the other takes the form $2n/k + f(k, D)$, where $f(k, D)$ is called the competitive overhead. In this paper, we present the first algorithm with linear-in-D competitive overhead, thereby reconciling both approaches. Specifically, our bound is in $2n/k + O(k^{\log_2(k)-1}D)$ and leads to a competitive ratio in $O(k/\exp(\sqrt{\ln 2 \ln k}))$. This is the first improvement over $\mathcal{O}(k/\ln k)$ since the introduction of the problem, twenty years ago. Our algorithm is developed for an asynchronous generalization of collective tree exploration (ACTE). It belongs to a broad class of locally-greedy exploration algorithms that we define. We show that the analysis of locally-greedy algorithms can be seen through the lens of a 2-player game that we call the tree-mining game and which could be of independent interest.

Romain Cosson
Inria
romain.cosson@inria.fr

CP28

Maintaining Matroid Intersections Online

Maintaining a maximum bipartite matching online while minimizing augmentations is a well studied problem, motivated by content delivery, job scheduling, and hashing. A breakthrough result of Bernstein, Holm, and Rotenberg (*SODA 2018*) resolved this problem up to a logarithmic factors. However, to model other problems in scheduling and resource allocation, we may need a richer class of combinatorial constraints (e.g., matroid constraints). We consider the problem of maintaining a maximum independent set of an arbitrary matroid $\mathcal{M}$ and a partition matroid $\mathcal{P}$. Specifically, at each timestep t one part P_t of the partition matroid is revealed: we must now select at most one newly-revealed element, but may exchange some previously selected elements, to maintain a maximum independent set on the elements seen thus far. The goal is to minimize the number of augmentations. If $\mathcal{M}$ is also a partition matroid, we recover the problem of maintaining a maximum bipartite matching online with recourse as a special case. Our main result is an $O(n \log^2 n)$-competitive algorithm, where n is the rank of the largest common base; this matches the current best quantitative bound for the bipartite matching special case. Our result builds substantively on the result of Bernstein, Holm, and Rotenberg: a key contribution of our work is to make use of market equilibria and prices in submodular utility allocation markets.

Niv Buchbinder
Tel Aviv University
nivb@tauex.tau.ac.il

Anupam Gupta, Daniel Hathcock
Carnegie Mellon University
anupamg@cs.cmu.edu, dhathcoc@andrew.cmu.edu

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

Sherry Sarkar
Carnegie Mellon University
sherrys@andrew.cmu.edu

CP28

Bin Packing under Random-Order: Breaking the Barrier of 3/2

Best-Fit is one of the most prominent and practically used algorithms for the bin packing problem, where a set of items with associated sizes needs to be packed in the minimum number of unit-capacity bins. Kenyon [SODA '96] studied online bin packing under random-order arrival, where the adversary chooses the list of items, but the items arrive one by one according to an arrival order drawn uniformly at random from the set of all permutations of the items. Kenyon's seminal result established an upper bound of 1.5 and a lower bound of 1.08 on the random-order ratio of Best-Fit, and it was conjectured that the true ratio is ≈ 1.15. The conjecture, if true, will also imply that Best-Fit (on randomly permuted input) has the best performance guarantee among all the widely-used simple algorithms for (offline) bin packing. This conjecture has remained one of the major open problems in the area, as highlighted in the recent survey on random-order models by Gupta and Singla [Beyond the Worst-Case Analysis of Algorithms '20].

In this paper, after almost three decades, we make the first progress towards the conjecture, by showing that Best-Fit achieves a random-order ratio of at most $1.5 - \varepsilon$, for a small constant $\varepsilon > 0$. Furthermore, we establish an improved lower bound of 1.144 on the random-order ratio of Best-Fit, nearly reaching the conjectured ratio.

Anish S. Hebbar
Duke University
anishshripad.hebbar@duke.edu

Arindam Khan, K. V. N. Sreenivas
Indian Institute of Science, Bengaluru, India
arindamkhan@iisc.ac.in, venkatanaga@iisc.ac.in

CP28

Controlling Tail Risk in Online Ski-Rental

The classical ski-rental problem admits a textbook 2-competitive deterministic algorithm, and a simple randomized algorithm that is $\frac{e}{e-1}$-competitive in expectation. The randomized algorithm, while optimal in expectation, has a large variance in its performance: it has more than a 37% chance of the competitive ratio exceeding 2, and a $\Theta(1/n)$ chance of the competitive ratio exceeding n. We ask what happens to the optimal solution if we insist that the tail risk, i.e., the chance of the competitive ratio exceeding a specific value, is bounded by some constant δ. We find that this additional modification significantly changes the structure of the optimal solution. The probability of purchasing skis on a given day becomes non-monotone, discontinuous, and arbitrarily large (for sufficiently small tail risk δ and large purchase cost n).

Thomas Lavastida
The University of Texas at Dallas
thomas.lavastida@utdallas.edu

Michael Dinitz
Johns Hopkins University
mdinitz@cs.jhu.edu

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

Benjamin Moseley
Carnegie Mellon University
moseleyb@andrew.cmu.edu

Sergei Vassilvitskii
Google Research
sergeiv@google.com

CP28

Poly-Logarithmic Competitiveness for the k-Taxi Problem

The online k-taxi problem generalizes the k-server problem, requiring servers to move between source-sink pairs in an n-point metric space, and the cost is the overhead incurred. In the deterministic setting, the problem has a lower bound on the competitiveness of $\Omega(2^k)$, a qualitative separation from k-server. Randomized algorithms are known with competitiveness $O(2^k \log n)$ (by Coester and Koutsoupias, even against adaptive adversaries), $O(2^{O(\sqrt{\log \Delta \log k})} \log_\Delta n)$ (by Buchbinder, Coester and Naor, where Δ is the aspect ratio of the metric space), and $O((n \log k)^2 \log n)$ (by Bubeck, Buchbinder, Coester, and Sellke). The best lower bound known is $\Omega(\log^2 k)$ which is inherited from the k-server problem, obtained in a recent breakthrough by Bubeck, Coester, and Rabani, showing a large gap in our understanding of problems that go slightly beyond the metrical task system framework. An open question left by these works was whether there is a randomized algorithm for the k-taxi problem with a competitive ratio that is poly-logarithmic in all the parameters. We answer this question in the affirmative in this paper. We give a covering relaxation for k-taxi on HSTs, which is obtained from the (non-covering) min-cost flow formulation of the problem. The constraints of our LP have compositionality properties that we use to develop a hierarchical primal-dual algorithm defined on the subtrees of the HST.

Debmalya Panigrahi
Duke University
debmalya@cs.duke.edu

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

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

CP29

Tight Lower Bound on Equivalence Testing in Conditional Sampling Model

We study the equivalence testing problem where the goal is to determine if the given two unknown distributions on $[n]$ are equal or ϵ-far in the total variation distance in the conditional sampling model (CFGM, SICOMP16; CRS, SICOMP15) wherein a tester can get a sample from the distribution conditioned on any subset. Equivalence testing is a central problem in distribution testing, and there has been a plethora of work on this topic in various sampling models. Despite significant efforts over the years, there remains a gap in the current best-known upper bound of $\tilde{O}(\log \log n)$ [FJOPS, COLT 2015] and lower bound of $\Omega(\sqrt{\log \log n})$[ACK, RANDOM 2015, Theory of Computing 2018]. Closing this gap has been repeatedly posed as an open problem (listed as problems 66 and 87 at sublinear.info). In this paper, we completely resolve the query complexity of this problem by showing a lower bound of $\tilde{\Omega}(\log \log n)$. For that purpose, we develop a novel and generic proof technique that enables us to break the $\sqrt{\log \log n}$ barrier, not only for the equivalence testing problem but also for other distribution testing problems, such as uniblock property.

Gunjan Kumar
Tata Institute of Fundamental Research, Mumbai
gunjan.kg1024@gmail.com

Diptarka Chakraborty
National University of Singapore
diptarka@comp.nus.edu.sg

Sourav Chakraborty
Indian Statistical Institute, Kolkata
sourav@isical.ac.in

CP29

Mildly Exponential Lower Bounds on Tolerant Testers for Monotonicity, Unateness, and Juntas

We give the first super-polynomial (in fact, mildly exponential) lower bounds for tolerant testing (equivalently, distance estimation) of monotonicity, unateness, and juntas with a *constant* separation between the "yes" and "no" cases. Specifically, we give * A $2^{\Omega(n^{1/4}/\sqrt{\epsilon})}$-query lower bound for non-adaptive, two-sided tolerant monotonicity testers and unateness testers when the "gap' parameter $\epsilon_2 - \epsilon_1$ is equal to ϵ, for any $\epsilon \geq 1/\sqrt{n}$; * A $2^{\Omega(k^{1/2})}$-query lower bound for non-adaptive, two-sided tolerant junta testers when the gap parameter is an absolute constant. In the constant-gap regime no non-trivial prior lower bound was known for monotonicity, the best prior lower bound known for unateness was $\tilde{\Omega}(n^{3/2})$ queries, and the best prior lower bound known for juntas was poly(k) queries.

Xi Chen
Columbia University
xichen@cs.columbia.edu

Anindya De
University of Pennsylvania
anindyad@seas.upenn.edu

Yuhao Li, Shivam Nadimpalli, Rocco A. Servedio
Columbia University
yuhaoli@cs.columbia.edu, sn2855@columbia.edu, rocco@cs.columbia.edu

CP29

Uniformity Testing over Hypergrids with Subcube Conditioning

We give an algorithm for testing uniformity of distributions supported on hypergrids $[m_1] \times \cdots \times [m_n]$, which makes $\widetilde{O}(\text{poly}(m)\sqrt{n}/\epsilon^2)$ many queries to a subcube conditional sampling oracle with $m = \max_i m_i$. When m is a constant, our algorithm is nearly optimal and strengthens

the algorithm of Canonne et al. (2021) which has the same query complexity but works for hypercubes $\{\pm 1\}^n$ only. A key technical contribution behind the analysis of our algorithm is a proof of a robust version of Pisier's inequality for functions over hypergrids using Fourier analysis.

Xi Chen
Columbia University
xichen@cs.columbia.edu

Cassandra Marcussen
Harvard University
cmarcussen@g.harvard.edu

CP29

A Tight Bound for Testing Partition Properties

A *partition property* of order k asks if a graph can be partitioned into k vertex sets of prescribed sizes so that the densities between any pair of sets falls within a prescribed range. This family of properties has been extensively studied in various areas of research ranging from theoretical computer science to statistical physics. Our main result is that every partition property of order k is testable with query complexity $\text{poly}(k/\epsilon)$. We thus obtain an exponential improvement (in k) over the $(1/\epsilon)^{O(k)}$ bound obtained by Goldreich, Goldwasser and Ron in their seminal FOCS 1996 paper. We further prove that our bound is tight in the sense that it cannot be made sub-polynomial in either k or ϵ. Besides the intrinsic interest in obtaining a tight bound for the above well studied family of properties, our improved bound has several combinatorial and algorithmic implications, stemming from the fact that it remains polynomial even when testing partition properties of order $k = \text{poly}(1/\epsilon)$.

Asaf Shapira
Tel Aviv University
asafico@tau.ac.il

Henrique Stagni
University of São Paulo
stagni@gmail.com

CP29

Adversarial Low Degree Testing

In the t-online-erasure model in property testing, an adversary is allowed to erase t values of a queried function for each query the tester makes. This model was recently formulated by Kalemaj, Raskhodnikova and Varma, who showed that the properties of linearity of functions as well as quadraticity can be tested in $O_t(1)$ many queries: $O(\log(t))$ for linearity and $2^{2^{O(t)}}$ for quadraticity. They asked whether the more general property of low-degreeness can be tested in the online erasure model, whether better testers exist for quadraticity, and if similar results hold when "erasures' are replaced with "corruptions'. We show that, in the t-online-erasure model, for a prime power q, given query access to a function $f : \mathbb{F}_q^n \to \mathbb{F}_q$, one can distinguish in $\log^{O(d+q)}(t)/\delta)$ queries between the case that f is degree at most d, and the case that f is δ-far from any degree d function (with respect to the fractional hamming distance). This answers the aforementioned questions and brings the query complexity to nearly match the query complexity of low-degree testing in the classical property testing model.

Kai Zhe Zheng, Dor Minzer
MIT
kzzheng@mit.edu, minzer.dor@gmail.com

CP30

Solving Fréchet Distance Problems by Algebraic Geometric Methods

We study several polygonal curve problems under the Fréchet distance via algebraic geometric methods. Let $\mathbb{X}_m^d$ and $\mathbb{X}_k^d$ be the spaces of all polygonal curves of m and k vertices in $\mathbb{R}^d$, respectively. We assume that $k \leq m$. Let $\mathcal{R}_{k,m}^d$ be the set of ranges in $\mathbb{X}_m^d$ for all possible metric balls of polygonal curves in $\mathbb{X}_k^d$ under the Fréchet distance. We prove a nearly optimal bound of $O(dk \log(km))$ on the VC dimension of the range space $(\mathbb{X}_m^d, \mathcal{R}_{k,m}^d)$, improving on the previous $O(d^2k^2 \log(dkm))$ upper bound and approaching the current $\Omega(dk \log k)$ lower bound. Our upper bound also holds for the weak Fréchet distance. We also obtain exact solutions that are hitherto unknown for curve simplification, range searching, nearest neighbor search, and distance oracle.

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

Haoqiang Huang
The Hong Kong University of Science and Technology
The Hong Kong University of Science and Technology
haoqiang.huang@connect.ust.hk

CP30

Improved Bounds for Point Selections and Halving Hyperplanes in Higher Dimensions

Let (P, E) be a $(d+1)$-uniform geometric hypergraph, where P is an n-point set in general position in R^d and $E \subseteq \binom{P}{d+1}$ is a collection of $\epsilon\binom{n}{d+1}$ d-dimensional simplices with vertices in P, for $0 < \epsilon \leq 1$. We show that there is a point $x \in R^d$ that pierces $\Omega\left(\epsilon^{(d^4+d^2+d)(d+1)+\delta}\binom{n}{d+1}\right)$ simplices in E, for any fixed $\delta > 0$. This is a dramatic improvement in all dimensions $d \geq 3$, over the previous lower bounds of the form $\epsilon^{(cd)^{d+1}} n^{d+1}$. As a by-product, we show that any n-point set in general position in R^d admits only $O\left(n^{d-\frac{1}{d(d-1)^4+d(d-1)^2+d(d-1)}+\delta}\right)$ halving hyperplanes, for any $\delta > 0$, which constitutes an improvement over the previously best known bound $O\left(n^{d-\frac{1}{(2d)^d}}\right)$ in all dimensions $d \geq 5$. An essential ingredient of our proof is the following semi-algebraic Turán-type result: Let $(V_1, \ldots, V_k, E)$ be a hypergraph of bounded semi-algebraic description complexity, whose vertices are in general position in R^d. Suppose that $|E| \geq \varepsilon|V_1| \cdot \ldots \cdot |V_k|$ holds for some $\varepsilon > 0$, then there exist subsets $W_i \subseteq V_i$ so that $|W_i| = \Omega\left(\varepsilon^{d+1+\delta}|V_i|\right)$ for $1 \leq i \leq d-1$, $|W_k| = \Omega\left(\varepsilon|V_i|\right)$, and $W_1 \times \ldots \times W_k \subseteq E$.

Natan Rubin
Ben Gurion University

rubinnat.ac@gmail.com

CP30

Fast and Accurate Approximations of the Optimal Transport in Semi-Discrete and Discrete Settings

Given a d-dimensional continuous (resp. discrete) probability distribution μ and a discrete distribution ν, the semi-discrete (resp. discrete) Optimal Transport (OT) problem asks for computing a minimum-cost plan to transport mass from μ to ν; we assume n to be the size of the support of the discrete distributions, and we assume we have access to an oracle outputting the mass of μ inside a constant-complexity region in $O(1)$ time. In this paper, we present three approximation algorithms for the OT problem. (i) Semi-discrete additive approximation: For any $\epsilon > 0$, we present an algorithm that computes a semi-discrete transport plan with ϵ-additive error in $n^{O(d)} \log \frac{C_{\max}}{\epsilon}$ time; here, $C_{\max}$ is the diameter of the supports of μ and ν. (ii) Semi-discrete relative approximation: For any $\epsilon > 0$, we present an algorithm that computes a $(1+\epsilon)$-approximate semi-discrete transport plan in $n\epsilon^{-O(d)} \log(n) \log^{O(d)}(\log n)$ time; here, we assume the ground distance is any L_p norm. (iii) Discrete relative approximation: For any $\epsilon > 0$, we present a Monte-Carlo $(1+\epsilon)$-approximation algorithm that computes a transport plan under any L_p norm in $n\epsilon^{-O(d)} \log(n) \log^{O(d)}(\log n)$ time; here, we assume that the spread of the supports of μ and ν is polynomially bounded.

Pankaj Agarwal
Duke University
pankaj@cs.duke.edu

Sharath Raghvendra
North Carolina State University
r.sharath@gmail.com

Pouyan Shirzadian
Virginia Tech University
pshirzadian@vt.edu

Keegan Yao
Duke University
keegan.yao@duke.edu

CP30

An Optimal Algorithm for Higher-Order Voronoi Diagrams in the Plane: The Usefulness of Nondeterminism

We present the first optimal randomized algorithm for constructing the order-k Voronoi diagram of n points in two dimensions. The expected running time is $O(n \log n + nk)$, which improves the previous, two-decades-old result of Ramos (SoCG'99) by a $2^{O(\log^* k)}$ factor. To obtain our result, we (i) use a recent decision-tree technique of Chan and Zheng (SODA'22) in combination with Ramos's cutting construction, to reduce the problem to *verifying* an order-k Voronoi diagram, and (ii) solve the verification problem by a new divide-and-conquer algorithm using planar-graph separators. We also describe a deterministic algorithm for constructing the k-level of n lines in two dimensions in $O(n \log n + nk^{1/3})$ time, and constructing the k-level of n planes in three dimensions in $O(n \log n + nk^{3/2})$ time. These time bounds (ignoring the $n \log n$ term) match the current best upper bounds on the combinatorial complexity of the k-level. Previously, the same time bound in two dimensions was obtained by Chan (1999) but with randomization.

Timothy M. Chan
University of Illinois at Urbana-Champaign
tmc@illinois.edu

Pingan Cheng
Department of Computer Science
Aarhus University
pingancheng@cs.au.dk

Da Wei Zheng
Department of Computer Science
University of Illinois at Urbana Champaign
dwzheng2@illinois.edu

CP30

Fully Scalable Massively Parallel Algorithms for Embedded Planar Graphs

We consider the massively parallel computation (MPC) model. In this model, assuming the widely believed 1-vs-2-cycles conjecture, many basic graph problems in $O(1)$ rounds with a strongly sublinear memory size per machine is impossible. We improve on the recent work of Holm and Tětek [SODA 2023] that bypass this barrier when a planar embedding of the graph is given. We extend their work to the fully scalable regime, where the memory size per machine can be $\mathcal{S} = n^{\delta}$ for any constant $0 < \delta < 1$. We give the first constant round fully scalable algorithms for embedded planar graphs for the problems of (i) connectivity and (ii) minimum spanning tree (MST). Moreover, we show that the ε-emulator can be incorporated into our recursive framework to obtain constant-round $(1+\varepsilon)$-approximation algorithms for the problems of computing (iii) single source shortest path (SSSP), (iv) global min-cut, and (v) st-max flow. All previous results on cuts and flows required linear memory in the MPC model. We also give as corollaries constant round fully scalable algorithms for other problems involve embedded planar graphs such as: (vi) 2D Euclidean MST using $O(n)$ total memory and (vii) $(1+\varepsilon)$-approximate weighted edit distance using $\widetilde{O}(n^{2-\delta})$ memory.

Yi-Jun Chang
National University of Singapore
cyijun@umich.edu

Da Wei Zheng
Department of Computer Science
University of Illinois at Urbana Champaign
dwzheng2@illinois.edu

CP31

Power of Posted-Price Mechanisms for Prophet Inequalities

We study the power of posted pricing mechanisms for Bayesian online optimization problems subject to combinatorial feasibility constraints. When the objective is to maximize social welfare, the problem is widely studied in the literature on prophet inequalities. While most (though not all) existing algorithms for the problem are implemented using a pricing mechanism, whether or not this can be done in general is unknown, and was formally left as an open question by Dtting, Feldman, Kesselheim, and

Lucier (FOCS 2017, SICOMP 2020). We show that any prophet inequality has an implementation using a posted price mechanism, thereby resolving the open question of Dtting et al. Given an algorithm for Bayesian online optimization, we show that it can be transformed, in a black-box manner, to a posted price algorithm that has the same or higher expected social welfare and preserves the distribution over the assigned outcomes. We further show how to implement our reduction efficiently under standard assumptions using access to a sampling oracle. As an immediate consequence, we obtain improved pricing-based prophet inequalities for maximum weight matching, resolving an open problem of Ezra, Feldman, Gravin and Tang (EC 2020, MOR 2022). We additionally obtain a posted pricing based implementation of the prophet inequality algorithm of Correa and Cristi (STOC 2023) for combinatorial auctions with subadditive valuations, resolving their open question as well.

Kiarash Banihashem, MohammadTaghi Hajiaghayi
University of Maryland
kiarash@umd.edu, hajiagha@umd.edu

Dariusz Kowalski, Piotr Krysta
Augusta University
dkowalski@augusta.edu, pkrysta@augusta.edu

Jan Olkowski
University of Maryland
olkowski@umd.edu

CP31

Online Duet Between Metric Embeddings and Minimum-Weight Perfect Matchings

Low-distortional metric embeddings are a crucial component in the modern algorithmic toolkit. In an online metric embedding, points arrive sequentially and the goal is to embed them into a simple space irrevocably, while minimizing the distortion. We show that there is a deterministic online embedding of a general metric into Euclidean space with distortion $O(\log n) \cdot \min\{\sqrt{\log \Phi}, \sqrt{n}\}$ (or, $O(d) \cdot \min\{\sqrt{\log \Phi}, \sqrt{n}\}$ if the metric has doubling dimension d), solving a conjecture by Newman and Rabinovich (2020). We show there is a stochastic embedding of a metric space into trees with expected distortion $O(d \cdot \log \Phi)$. Next, we study the online minimum-weight perfect matching problem, where the goal is to return a perfect matching that approximates the minimum-weight perfect matching at all times, while minimizing the recourse. Our third result is a randomized algorithm with competitive ratio $O(d \cdot \log \Phi)$ and recourse $O(\log \Phi)$ against an oblivious adversary. Next, we show there is a deterministic algorithm against an adaptive adversary, using $O(\log^2 n)$ recourse, that maintains a matching of weight at most $O(\log n)$ times the weight of the MST. We complement our upper bounds with a strategy for an oblivious adversary that, with recourse r, establishes a lower bound of $\Omega(\frac{\log n}{r \log r})$ for both competitive ratio and lightness.

Sujoy Bhore
Indian Institute of Technology Bombay
sujoy.bhore@gmail.com

Arnold Filtser
Columbia University
arnold273@gmail.com

Csaba D. Toth
California State University Northridge
csaba.toth@csun.edu

CP31

Edge-Disjoint Paths in Expanders: Online with Removals

We consider the problem of finding edge-disjoint paths between given pairs of vertices in a sufficiently strong d-regular expander graph G with n vertices. In particular, we describe a deterministic, polynomial time algorithm which maintains an initially empty collection of edge-disjoint paths $\mathcal{P}$ in G and fulfills any series of two types of requests: ¡ol¿ ¡li¿ Given two vertices a and b such that each appears as an endpoint in $O(d)$ paths in $\mathcal{P}$ and, additionally, $|\mathcal{P}| = O(nd/\log n)$, the algorithm finds a path of length at most $\log n$ connecting a and b which is edge-disjoint from all other paths in $\mathcal{P}$, and adds it to $\mathcal{P}$. ¡li¿ Remove a given path $P \in \mathcal{P}$ from $\mathcal{P}$. ¡/ol¿ Importantly, each request is processed before seeing the next one. The upper bound on the length of found paths and the constraints are the best possible up to a constant factor. This establishes the first online algorithm for finding edge-disjoint paths in expanders which also allows removals, significantly strengthening a long list of previous results on the topic.

Rajko Nenadov
University of Auckland
rajko.nenadov@auckland.ac.nz

Nemanja Draganic
University of Oxford
nemanja.draganic@maths.ox.ac.uk

CP31

Set Covering with Our Eyes Wide Shut

In the stochastic set cover problem [Grandoni et al., FOCS '08], we are given a collection S of m sets over a universe U of size N, and a distribution D over elements of U. The algorithm draws n elements one-by-one from D and must buy a set to cover each element on arrival; the goal is to minimize the total cost of sets bought during this process. Grandoni et al. gave an $O(\log mN)$-competitive universal algorithm for this stochastic set cover problem. We improve unilaterally upon this result by giving a simple, polynomial time $O(\log mn)$-competitive universal algorithm for the more general prophet version, in which U is formed by drawing from n different distributions $D_1, \ldots, D_n$. Furthermore, we show that we do not need full foreknowledge: in fact, a single sample from each distribution suffices. We show similar results for the 2-stage prophet and online-with-a-sample settings. We obtain our results via a generic reduction from the single-sample prophet setting to the random-order setting (for which Gupta et al. [FOCS '21] provide an algorithm); this reduction holds for a broad class of minimization problems that includes all covering problems. We take advantage of this framework by giving random-order algorithms for non-metric facility location and set multicover; using our framework, these automatically translate to universal prophet algorithms.

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

Gregory Kehne
University of Texas at Austin
gkehne@utexas.edu

Roie Levin
Carnegie Mellon University
roiel@tauex.tau.ac.il

CP31

Combinatorial Stationary Prophet Inequalities

Numerous recent papers have studied the tension between thickening and clearing a market in (uncertain, online) long-time horizon Markovian settings. In particular, (Aouad and Sarita'20, Collina et al. WINE'20, Kessel et al. EC'22) studied what the latter referred to as the Stationary Prophet Inequality Problem, due to its similarity to the classic finite-time horizon prophet inequality problem. These works all consider unit-demand buyers. Mirroring the long line of work on the classic prophet inequality problem subject to combinatorial constraints, we initiate the study of the stationary prophet inequality problem subject to combinatorially-constrained buyers. Our results can be summarized succinctly as unearthing an algorithmic connection between contention resolution schemes (CRS) and stationary prophet inequalities. While the classic prophet inequality problem has a tight connection to online CRS (Feldman et al. SODA'16, Lee and Singla ESA'18), we show that for the stationary prophet inequality problem, offline CRS play a similarly central role. We show that, up to small constant factors, the best (ex-ante) competitive ratio achievable for the combinatorial prophet inequality equals the best possible balancedness achievable by offline CRS for the same combinatorial constraints.

Neel B. Patel
University of Southern California
neelbpat@usc.edu

David Wajc
Google Research
david.wajc@gmail.com

CP32

Faster Algorithms for Bounded Knapsack and Bounded Subset Sum via Fine-Grained Proximity Results

We investigate pseudopolynomial-time algorithms for Bounded Knapsack and Bounded Subset Sum. Recent years have seen a growing interest in settling their fine-grained complexity with respect to various parameters. For Bounded Knapsack, the number of items n and the maximum item weight $w_{\max}$ are two of the most natural parameters that have been studied extensively in the literature. The previous best running time in terms of n and $w_{\max}$ is $O(n + w_{\max}^3)$ [Polak, Rohwedder, Wegrzycki '21]. There is a conditional lower bound of $(n + w_{\max})^{2-o(1)}$ based on $(\min, +)$-convolution hypothesis [Cygan, Mucha, Wegrzycki, Wlodarczyk '17]. We narrow the gap significantly by proposing an $\widetilde{O}(n + w_{\max}^{12/5})$-time algorithm. For Bounded Subset Sum, we give two algorithms running in $\widetilde{O}(nw_{\max})$ and $\widetilde{O}(n + w_{\max}^{3/2})$ time, respectively. These results match the currently best running time for 0-1 Subset Sum. Prior to our work, the best running times (in terms of n and $w_{\max}$) for Bounded Subset Sum are $\widetilde{O}(n + w_{\max}^{5/3})$ [Polak, Rohwedder, Wegrzycki '21] and $\widetilde{O}(n + \mu_{\max}^{1/2} w_{\max}^{3/2})$ [implied by Bringmann '19 and Bringmann, Wellnitz '21], where $\mu_{\max}$ refers to the maximum multiplicity of item weights.

Lin Chen
Texas Tech University
chenlin198662@gmail.com

Jiayi Lian, Yuchen Mao, Guochuan Zhang
Zhejiang University
jiayilian@zju.edu.cn, maoyc@zju.edu.cn, zgc@zju.edu.cn

CP32

The Time Complexity of Fully Sparse Matrix Multiplication

What is the time complexity of matrix multiplication of sparse integer matrices with m_{in} nonzeros in the input and m_{out} nonzeros in the output? This paper provides improved upper bounds for this question for almost any choice of m_{in} vs. m_{out}, and provides evidence that these new bounds might be optimal up to further progress on fast matrix multiplication. Our main contribution is a new algorithm that reduces sparse matrix multiplication to dense (but smaller) rectangular matrix multiplication. Our running time thus depends on the optimal exponent $\omega(a, b, c)$ of multiplying dense $n^a \times n^b$ by $n^b \times n^c$ matrices. We discover that when $m_{out} = \Theta(m_{in}^r)$ the time complexity of sparse matrix multiplication is $O(m_{in}^{\sigma+\epsilon})$, for all $\epsilon > 0$, where σ is the solution to the equation $\omega(\sigma - 1, 2 - \sigma, 1 + r - \sigma) = \sigma$. No matter what $\omega(\cdot, \cdot, \cdot)$ turns out to be, and for all $r \in (0, 2)$, the new bound beats the state of the art, and we provide evidence that it is optimal. In particular, in terms of the input plus output size $m = m_{in} + m_{out}$ our algorithm runs in time $O(m^{1.3459})$. This improves over the previous $m^{\frac{2\omega}{\omega+1}+\epsilon} = O(m^{1.4071})$ bound [Amossen, Pagh; 2009], which was a natural barrier since it coincides with the longstanding bound of all-edge triangle in sparse graphs [Alon, Yuster, Zwick; 1994].

Amir Abboud
Weizmann Institute of Science
Israel
amir.abboud@weizmann.ac.il

Karl Bringmann
Saarland University and Max Planck Institute for Informatics
Saarland Informatics Campus, Saarbrücken, Germany
kbringma@mpi-inf.mpg.de

Nick Fischer
Weizmann Institute of Science
nick.fischer@weizmann.ac.il

Marvin Künnemann
Karlsruhe Institute of Technology
kuennemann@cs.uni-kl.de

CP32

The Effect of Sparsity on k-Dominating Set and Related First-Order Graph Properties

We revisit the classic k-Dominating Set, one of the first problems for which a tight $n^{k-o(1)}$ conditional lower bound (for $k \geq 3$), based on SETH, was shown (Patrascu et al., SODA 2007). However, the underlying reduction creates dense graphs, raising the question: how much does the sparsity of the graph affect its fine-grained complexity? We first settle the fine-grained complexity of k-Dominating Set in terms of both the number of nodes n and number of edges m. Specifically, we show an $mn^{k-2-o(1)}$ lower bound based on SETH, for any dependence of m on n. This

is complemented by an $mn^{k-2+o(1)}$-time algorithm for all $k \geq 3$. For the $k = 2$ case, we give a randomized algorithm that employs a Bloom-filter inspired hashing to improve the state of the art of $n^{\omega+o(1)}$ to $m^{\omega/2+o(1)} = O(m^{1.187})$. If $\omega = 2$, this yields a conditionally tight bound for all $k \geq 2$. To study if k-Dominating Set is special in its sensitivity to sparsity, we consider a class of very related problems. The k-Dominating Set problem belongs to a type of first-order definable graph properties that we call monochromatic basic problems. These problems are the natural monochromatic variants of the basic problems that were proven complete for the class FOP of first-order definable properties (Gao et al., TALG 2019). We show that among these problems, k-Dominating Set is the only one whose fine-grained complexity decreases in sparse graphs.

Mirza Redzic
Karlsruhe Institute of Technology
mirza233@gmail.com

Nick Fischer
Weizmann Institute of Science
nick.fischer@weizmann.ac.il

Marvin Künnemann
Karlsruhe Institute of Technology
kuennemann@cs.uni-kl.de

CP32

Improved Roundtrip Spanners, Emulators, and Directed Girth Approximation

We study approximation of the girth in directed graphs. For graphs on n vertices and m edges, it was previously known that one can get a 2-approximation of girth in $\tilde{O}(n^2, m\sqrt{n})$ time by Chechik and Lifshitz [SODA 2021], which is optimal for dense graphs. However, it is unknown whether one can faster algorithms for other approximation factors. Our main result is an algorithm that gives a 4-approximation and runs in $O(mn^{1/3})$ time, improving on the previous best known $(4+\varepsilon)$-approximation in $O(mn^{\sqrt{2}-1})$ time. In addition, we obtain better algorithms for related problems including the construction of roundtrip spanners and emulators, which are directed analogs of spanners and emulators where the roundtrip metric is used as a notion of distance. This talk will focus on the high level idea behind our algorithm for 4-approximation of directed girth.

Zixuan Xu, Alina Harbuzova
MIT
zixuanxu@mit.edu, hadought@mit.edu

Ce Jin
EECS
MIT
cejin@mit.edu

Virginia Vassilevska Williams
MIT
virgi@mit.edu

CP32

Fast 2-Approximate All Pairs Shortest Paths, Faster Approximate All Pairs Shortest Paths

The all pairs shortest path problem (APSP) is a foundational problem in computer science. For weighted dense graphs on n vertices, no truly sub-cubic algorithms exist to compute APSP exactly even for undirected graphs. This is popularly known as the APSP conjecture and has played a prominent role in developing the field of fine-grained complexity. Even for unweighted undirected graphs, computing a $(2-\varepsilon)$-approximation is at least as hard as boolean matrix multiplication, requiring $\Omega(n^\omega)$ time. In these works, we present a multitude of new approximation algorithms for the APSP problem. In particular, we provide new algorithms for 2-approximate APSP in unweighted graphs: a deterministic one that runs in $O\left(n^{2.072}\right)$ time and a randomized one that runs in $O\left(n^{2.032}\right)$ time improving upon the best known bound of $\tilde{O}\left(n^{2.25}\right)$ by Roditty (STOC, 2023). We also provide improved algorithms for weighted graphs and improved additive approximations, including obtaining an improvement for all additive factors over the bounds of Dor, Halperin, and Zwick (FOCS, 1996). Our techniques also lead to improved distance oracles for weighted graphs. In particular, for the sparse regime we construct a distance oracle in $\tilde{O}\left(mn^{2/3}\right)$ time that supports 2-approximate queries in constant time. To the best of our knowledge, this is the first 2-approximate distance oracle that has subquadratic preprocessing time in sparse graphs.

Michal Dory
University of Haifa
mdory@ds.haifa.ac.il

Sebastian Forster
University of Salzburg
forster@cs.sbg.ac.at

Yael Kirkpatrick
Massachusetts Institute of Technology
yaelkirk@mit.edu

Yasamin Nazari
VU Amsterdam
y.nazari@vu.nl

Virginia Vassilevska Williams
MIT
virgi@mit.edu

Tijn de Vos
University of Salzburg
tdevos@csd.sbg.ac.at

Barna Saha, Christopher Ye
University of California, San Diego
bsaha@ucsd.edu, czye@ucsd.edu

CP33

Delaunay Bifiltrations of Functions on Point Clouds

The Delaunay filtration $\mathcal{D}_\bullet(X)$ of a point cloud $X \subset \mathbb{R}^d$ is a central tool of computational topology. Its use is justified by the topological equivalence of $\mathcal{D}_\bullet(X)$ and the offset (i.e., union-of-balls) filtration of X. Given a function $\gamma : X \to \mathbb{R}$, we introduce a Delaunay bifiltration $\mathcal{DC}_\bullet(\gamma)$ that satisfies an analogous topological equivalence, ensuring that $\mathcal{DC}_\bullet(\gamma)$ topologically encodes the offset filtrations of all sublevel sets of γ, as well as the topological relations between them. $\mathcal{DC}_\bullet(\gamma)$ is of size $O(|X|^{\lceil \frac{d+1}{2} \rceil})$, which for d odd matches the worst-case size of $\mathcal{D}_\bullet(X)$. Adapting the Bowyer-Watson algorithm for computing Delaunay trian-

gulations, we give a simple, practical algorithm to compute $\mathcal{DC}_\bullet(\gamma)$ in time $O(|X|^{\lceil \frac{d}{2} \rceil + 1})$. Our implementation, based on CGAL, computes $\mathcal{DC}_\bullet(\gamma)$ with modest overhead compared to computing $\mathcal{D}_\bullet(X)$, and handles tens of thousands of points in $\mathbb{R}^3$ within seconds.

ngel Javier Alonso, Michael Kerber
University of Technology, Graz, Austria
alonsohernandez@tugraz.at, kerber@tugraz.at

Tung Lam
SUNY Albany, USA
tlam@albany.edu

Michael Lesnick
University at Albany, U.S.
mlesnick@albany.edu

CP33

Flip Graph Connectivity for Arrangements of Pseudolines and Pseudocircles

We study flip graphs of arrangements of pseudolines and of arrangements of pseudocircles, which are combinatorial generalizations of lines and circles, respectively. In both cases we consider triangle flips as local transformation and prove conjectures regarding their connectivity. In the case of n pseudolines we show that the connectivity of the flip graph equals its minimum degree, which is exactly $n-2$. For the proof we introduce the class of *shellable* line arrangements, which serve as reference objects for the construction of disjoint paths. In fact, shellable arrangements are elements of a flip graph of line arrangements which are vertices of a polytope (Felsner and Ziegler; DM 241 (2001), 301–312). This polytope forms a cluster of good connectivity in the flip graph of pseudolines. In the case of pseudocircles we show that triangle flips induce a connected flip graph on *intersecting* arrangements and also on *cylindrical* intersecting arrangements. The result for cylindrical arrangements is used in the proof for intersecting arrangements. We also show that in both settings the diameter of the flip graph is in $\Theta(n^3)$. Our constructions make essential use of variants of the sweeping lemma for pseudocircle arrangements (Snoeyink and Hershberger; Proc. SoCG 1989: 354–363). We finally study cylindrical arrangements in their own right and provide new combinatorial characterizations of this class.

Yan Alves Radtke
Technische Universität Berlin
alvesradtke@campus.tu-berlin.de

Stefan Felsner
Technical University of Berlin
Institute for Mathematics
felsner@math.tu-berlin.de

Johannes Obenaus
Freie Universität Berlin
johannes.obenaus@fu-berlin.de

Sandro Roch, Manfred Scheucher
Technische Universität Berlin
roch@math.tu-berlin.de, scheucher@math.tu-berlin.de

Birgit Vogtenhuber
Graz University of Technology
bvogt@ist.tugraz.at

CP33

Untangling Graphs on Surfaces

Consider a graph drawn on a surface, possibly with crossings. We provide an algorithm to decide whether such a drawing can be untangled; in other words, whether the drawing is homotopic to an embedding. While the problem boils down to planarity testing when the surface is the sphere or the disk, the other cases have never been studied before, except when the input graph is a cycle, in an abundant literature in topology and more recently by Desprd Lazarus [SoCG 2017, J. ACM 2019]. Our algorithm runs in O(m + poly(g+b) n log n) time, where g ¿= 0 and b ¿= 0 are the genus and the number of boundary components of the input orientable surface S, and n is the size of the input graph drawing, lying on some fixed graph of size m cellularly embedded on S. We use various tools. Most notably, we introduce reducing triangulations, a novel discrete analog of hyperbolic surfaces in the spirit of systems of quads by Lazarus and Rivaud [FOCS 2012] and Erickson and Whittlesey [SODA 2013], which have the additional benefit that reduced paths are unique and stable upon reversal; they are likely of independent interest. Tailored data structures are needed to achieve certain homotopy tests efficiently on these triangulations. As a key subroutine, we rely on an algorithm to test the weak simplicity of a graph drawn on a surface by Akitaya, Fulek, and TSODA 2018, TALG 2019].

Loïc Dubois
LIGM, CNRS, Univ Gustave Eiffel, F-77454
Marne-la-Vallée
loic.dubois@univ-eiffel.fr

Éric Colin de Verdière
LIGM, CNRS, Univ Gustave Eiffel, F-77454
Marne-la-Vallée
France
eric.colin-de-verdiere@univ-eiffel.fr

Vincent Despré
Université de Lorraine, CNRS, Inria, LORIA, F-54000
Nancy
France
vincent.despre@inria.fr

CP33

Fast Approximation Algorithms for Piercing Boxes by Points

Let $\mathcal{B} = \{\mathsf{b}_1, \ldots, \mathsf{b}_n\}$ be a set of n axis-aligned boxes in $\mathbb{R}^d$ where $d \geq 2$ is a constant. The *piercing problem* is to compute a smallest set of points $\mathcal{N} \subset \mathbb{R}^d$ that hits every box in $\mathcal{B}$. Let $\kappa := \kappa(\mathcal{B})$, the *piercing number* be the minimum size of a piercing set of $\mathcal{B}$. We first present a randomized $O(\log\log\kappa)$-approximation algorithm with expected running time $O(n^{d/2}\,\mathrm{polylog}(n))$. Next, we show that the expected running time can be improved to near-linear using a sampling-based technique, if $\kappa = O(n^{1/(d-1)})$. Specifically, in the plane, the improved running time is $O(n \log \kappa)$, assuming $\kappa < n/\log^{\Omega(1)} n$. Finally, we study the dynamic version of the piercing problem where boxes can be inserted or deleted. For boxes in $\mathbb{R}^2$, we obtain a randomized $O(\log\log\kappa)$-approximation algorithm with $O(n^{1/2}\,\mathrm{polylog}(n))$ amortized expected update time. Our algorithms are based on the multiplicative

weight-update (MWU) method and require the construction of a weak ε-net for a point set with respect to boxes. A key idea of our work is to exploit the duality between the piercing set and independent set (for boxes) to speed up our MWU. We also present a simpler algorithm for constructing a weak ε-net than in previous work.

Pankaj Agarwal
Duke University
pankaj@duke.edu

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

Rahul Raychaudhury
Duke University
rahul.raychaudhury@duke.edu

Stavros Sintos
University of Illinois at Chicago
stavros@uic.edu

CP33

Near-Optimal Min-Sum Motion Planning for Two Square Robots in a Polygonal Environment

Let $\mathcal{W} \subset \mathbb{R}^2$ be a planar polygonal environment (i.e., a polygon potentially with holes) with a total of n vertices, and let A, B be two robots, each modeled as an axis-aligned unit square, that can translate inside $\mathcal{W}$. Given source and target placements $s_A, t_A, s_B, t_B \in \mathcal{W}$ of A and B, respectively, the goal is to compute a collision-free motion plan π^*, i.e., a motion plan that continuously moves A from s_A to t_A and B from s_B to t_B so that A and B remain inside $\mathcal{W}$ and do not collide with each other during the motion. Furthermore, if such a plan exists, then we wish to return a plan that minimizes the sum of the lengths of the paths traversed by the robots, $|\pi^*|$. Given $\mathcal{W}, s_A, t_A, s_B, t_B$ and a parameter $\varepsilon > 0$, we present an $n^2\varepsilon^{-O(1)} \log n$-time $(1+\varepsilon)$-approximation algorithm for this problem. We are not aware of any polynomial time algorithm for this problem, nor do we know whether the problem is NP-Hard. Our result is the first polynomial-time $(1+\varepsilon)$-approximation algorithm for an optimal motion planning problem involving two robots moving in a polygonal environment.

Alex Steiger, Pankaj Agarwal
Duke University
asteiger@cs.duke.edu, pankaj@cs.duke.edu

Dan Halperin, Micha Sharir
Tel Aviv University
danha@tauex.tau.ac.il, michas@tauex.tau.ac.il

CP34

Combinatorial Approach for Factorization of Variance and Entropy in Spin Systems

We present a simple combinatorial framework for establishing approximate tensorization of variance and entropy in the setting of spin systems (a.k.a. undirected graphical models) based on balanced separators of the underlying graph. Such approximate tensorization results immediately imply as corollaries many important structural properties of the associated Gibbs distribution, in particular rapid mixing of the Glauber dynamics for sampling. We prove approximate tensorization by recursively establishing block factorization of variance and entropy with a small balanced separator of the graph. Our approach goes beyond the classical canonical path method for variance and the recent spectral independence approach, and allows us to obtain new rapid mixing results. As applications of our approach, we show that: 1. On graphs of treewidth t, the mixing time of the Glauber dynamics is $n^{O(t)}$, which recovers the recent results of Eppstein and Frishberg with improved exponents and simpler proofs; 2. On bounded-degree planar graphs, strong spatial mixing implies $\tilde{O}(n)$ mixing time of the Glauber dynamics, which gives a faster algorithm than the previous deterministic counting algorithm by Yin and Zhang.

Zongchen Chen
University at Buffalo
zchen83@buffalo.edu

CP34

Fast Sampling of b-Matchings and b-Edge Covers

For an integer $b \geq 1$, a b-matching (resp. b-edge cover) of a graph $G = (V, E)$ is a subset $S \subseteq E$ of edges such that every vertex is incident with at most (resp. at least) b edges from S. We prove that for any $b \geq 1$ the simple Glauber dynamics for sampling (weighted) b-matchings and b-edge covers mixes in $O(n \log n)$ time on all n-vertex bounded-degree graphs. This significantly improves upon previous results which have worse running time and only work for b-matchings with $b \leq 7$ and for b-edge covers with $b \leq 2$. More generally, we prove spectral independence for a broad class of binary symmetric Holant problems with log-concave signatures, including b-matchings, b-edge covers, and antiferromagnetic 2-spin edge models. We hence deduce optimal mixing time of the Glauber dynamics from spectral independence. The core of our proof is a recursive coupling inspired by (Chen and Zhang, SODA '23) which upper bounds the Wasserstein W_1 distance between distributions under different pinnings. Using a similar method, we also obtain the optimal $O(n \log n)$ mixing time of the Glauber dynamics for the hardcore model on n-vertex bounded-degree claw-free graphs, for any fugacity λ. This improves over previous works which have at least cubic dependence on n.

Zongchen Chen
University at Buffalo
zchen83@buffalo.edu

Yuzhou Gu
Institute for Advanced Study
yuzhougu@ias.edu

CP34

Optimality of Glauber Dynamics for General-Purpose Ising Model Sampling and Free Energy Approximation

Recently, Eldan, Koehler, and Zeitouni (2020) showed that Glauber dynamics mixes rapidly for general Ising models so long as the difference between the largest and smallest eigenvalues of the coupling matrix is at most $1 - \epsilon$ for any fixed $\epsilon > 0$. We give evidence that Glauber dynamics is in fact optimal for this "general-purpose sampling" task. Namely, we give an average-case reduction from hypothesis testing in a Wishart negatively-spiked matrix model to approximately sampling from the Gibbs measure of a general Ising model for which the difference between the largest

and smallest eigenvalues of the coupling matrix is at most $1+\epsilon$ for any fixed $\epsilon > 0$. Combined with results of Bandeira, Kunisky, and Wein (2019) that analyze low-degree polynomial algorithms to give evidence for the hardness of the former spiked matrix problem, our results in turn give evidence for the hardness of general-purpose sampling improving on Glauber dynamics. We also give a similar reduction to approximating the free energy of general Ising models, and again infer evidence that simulated annealing algorithms based on Glauber dynamics are optimal.

Dmitriy Kunisky
Yale University
dmitriy.kunisky@yale.edu

CP34

Smoothed Complexity of Swap in Local Graph Partitioning

We give the first quasipolynomial upper bound $\phi n^{\mathrm{polylog}(n)}$ for the smoothed complexity of the SWAP algorithm for local Graph Partitioning (also known as Bisection Width) under the full perturbation model, where n is the number of nodes in the graph and ϕ is a parameter that measures the magnitude of perturbations applied on its edge weights. More generally, we show that the same quasipolynomial upper bound holds for the smoothed complexity of the 2-FLIP algorithm for any binary Maximum Constraint Satisfaction Problem, including local Max-Cut, for which similar bounds were only known for 1-FLIP. Our results are based on an analysis of a new notion of useful cycles in the multigraph formed by long sequences of double flips, showing that it is unlikely for every double flip in a long sequence to incur a positive but small improvement in the cut weight.

Emmanouil Vlatakis-Gkaragkounis
UC Berkeley
Simons Institute
emvlatakis@berkeley.edu

CP34

Universality of Spectral Independence with Applications to Fast Mixing in Spin Glasses

We study Glauber dynamics for sampling from discrete distributions μ on the hypercube $\{\pm 1\}^n$. Recently, techniques based on spectral independence have successfully yielded optimal $O(n)$ relaxation times for a host of different distributions μ. We show that spectral independence is universal: a relaxation time of $O(n)$ implies spectral independence. We then study a notion of tractability for μ, defined in terms of smoothness of the multilinear extension of its Hamiltonian – $\log \mu$ – over $[-1,+1]^n$. We show that Glauber dynamics has relaxation time $O(n)$ for such μ, and using the universality of spectral independence, we conclude that these distributions are also fractionally log-concave and consequently satisfy modified log-Sobolev inequalities. We sharpen our estimates and obtain approximate tensorization of entropy and the optimal $\widetilde{O}(n)$ mixing time for random Hamiltonians, i.e. the classically studied mixed p-spin model at sufficiently high temperature. These results have significant downstream consequences for concentration of measure, statistical testing, and learning.

Nima Anari
Stanford University
Department of Computer Science
anari@cs.stanford.edu

Vishesh Jain
University of Illinois Chicago
visheshj@uic.edu

Frederic Koehler, Huy Tuan Pham, Thuy-Duong Vuong
Stanford University
fkoehler@stanford.edu, huypham@stanford.edu, tdvuong@stanford.edu

CP35

Linear-Sized Sparsifiers via Near-Linear Time Discrepancy Theory

Discrepancy theory has provided powerful tools for producing higher-quality objects which "beat the union bound' in fundamental settings throughout combinatorics and computer science. However, this quality has often come at the price of more computationally-expensive algorithms. We introduce a new framework for bridging this gap, by allowing for the efficient implementation of discrepancy-theoretic primitives. Our framework repeatedly solves regularized optimization problems to low accuracy to approximate the partial coloring method of [**?**], and simplifies and generalizes recent work of [**?**] on fast algorithms for Spencer's theorem. As a first application, we use our framework to obtain an $\tilde{O}(m \cdot \epsilon^{-3.5})$-time algorithm for constructing an ϵ-approximate spectral sparsifier of an m-edge graph, matching the sparsity of [**?**] up to constant factors and improving upon the $\tilde{O}(m \cdot \epsilon^{-6.5})$ runtime of [**?**]. We further give a state-of-the-art algorithm for constructing graph ultrasparsifiers and an almost-linear time algorithm for constructing linear-sized degree-preserving sparsifiers via discrepancy theory; in the latter case, such sparsifiers were not known to exist previously.

Arun Jambulapati, Victor Reis
University of Washington
jmblpati@gmail.com, voreis@cs.washington.edu

Kevin Tian
Microsoft Rsearch
kjtian@cs.utexas.edu

CP35

Code Sparsification and Its Applications

We introduce a notion of code sparsification that generalizes the notion of cut sparsification in graphs. For a (linear) code $\mathcal{C} \subseteq \mathbb{F}_q^n$ of dimension k a $(1 \pm \epsilon)$-sparsification of size s is given by a weighted set $S \subseteq [n]$ with $|S| \leq s$ such that for every codeword $c \in \mathcal{C}$ the projection $c|_S$ of c to the set S has (weighted) hamming weight which is a $(1 \pm \epsilon)$ approximation of the hamming weight of c. We show that for every code there exists a $(1 \pm \epsilon)$-sparsification of size $s = \widetilde{O}(k \log(q)/\epsilon^2)$. This immediately implies known results on graph and hypergraph cut sparsification up to polylogarithmic factors (with a simple unified proof). One application of our result is near-linear size sparsifiers for constraint satisfaction problems (CSPs) over $\mathbb{F}_p$-valued variables whose unsatisfying assignments can be expressed as the zeros of a linear equation modulo a prime p. Building on this, we obtain a complete characterization of ternary Boolean CSPs that admit near-linear size sparsification. Finally, by connections between the eigenvalues of the Laplacians of Cayley graphs over $\mathbb{F}_2^k$ to the weights of codewords, we also give the first proof of the existence of spectral Cayley graph sparsifiers over $\mathbb{F}_2^k$ by Cayley graphs, i.e., where we sparsify the set of generators

to nearly-optimal size.

Aaron L. Putterman
Harvard University
aputterman@g.harvard.edu

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

Madhu Sudan
Harvard University
madhusudan@g.harvard.edu

CP35

Quotient Aparsification for Submodular Functions

Input your abstract, including TeX commands, here. The abstract should be no longer than 1500 characters, including spaces. Only input the abstract text. Don't include title or author information here.

Kent Quanrud
Purdue University
krq@purdue.edu

CP35

Sublinear Time Low-Rank Approximation of Toeplitz Matrices

We present a sublinear time algorithm for computing a near optimal low-rank approximation to any positive semidefinite (PSD) Toeplitz matrix $T \in \mathbb{R}^{d\times d}$, given noisy access to its entries. In particular, given entrywise query access to $T + E$ for an arbitrary noise matrix $E \in \mathbb{R}^{d\times d}$, integer rank $k \leq d$, and error parameter $\delta > 0$, our algorithm runs in time $\text{poly}(k, \log(d/\delta))$ and outputs (in factored form) a Toeplitz matrix $\widetilde{T} \in \mathbb{R}^{d\times d}$ with rank $\text{poly}(k, \log(d/\delta))$ satisfying, for some fixed constant C,

$$\|T - \widetilde{T}\|_F \leq C \cdot \max\{\|E\|_F, \|T - T_k\|_F\} + \delta \cdot \|T\|_F.$$

Here $\|\cdot\|_F$ is the Frobenius norm and T_k is the best (not necessarily Toeplitz) rank-k approximation to T in the Frobenius norm, given by projecting T onto its top k eigenvectors. Our result can be applied in several settings. When $E = 0$, we obtain the first sublinear time near-relative-error low-rank approximation algorithm for PSD Toeplitz matrices, resolving the main open problem of Kapralov et al. SODA '23. Our algorithm can also be applied to approximate the unknown Toeplitz covariance matrix of a multivariate Gaussian distribution, given sample access to this distribution, resolving an open question of Eldar et al. SODA '20.

Cameron Musco
University of Massachusetts Amherst
cmusco@cs.umass.edu

Kshiteej Sheth
School of Computer and Communication Sciences,
EPFL, Switzerland
kshiteej.sheth@epfl.ch

CP35

A Quasi-Monte Carlo Data Structure for Smooth Kernel Evaluations

In the kernel density estimation (KDE) problem one is given a kernel $K(x, y)$ and a dataset P of points in a high dimensional Euclidean space, and must prepare a small space data structure that can quickly answer density queries: given a point q, output a $(1+\epsilon)$-approximation to $\mu := \frac{1}{|P|}\sum_{p\in P} K(p, q)$. In this work we give a data structure with $\approx \text{polylog}(1/\mu)/\epsilon$ query time for smooth kernel KDE. Our main insight is a new way to combine discrepancy theory with randomized space partitioning inspired by, but significantly more efficient than, that of the fast multipole methods. We hope that our techniques will find further applications to linear algebra for kernel matrices.

Erik Waingarten
Simons Institute
ewaingar@seas.upenn.edu

Moses Charikar
Stanford University, California
moses@cs.stanford.edu

Michael Kapralov
EPFL
michael.kapralov@epfl.ch

CP36

Shortcut Partitions in Minor-Free Graphs: Steiner Point Removal, Distance Oracles, Tree Covers, and More

The notion of shortcut partition, introduced recently by Chang, Conroy, Le, Milenkovic, Solomon, and Than [CCLMST23], is a new type of graph partition into low-diameter clusters. Roughly speaking, the shortcut partition guarantees that for every two vertices u and v in the graph, there exists a path between u and v that intersects only a few clusters. They proved that any planar graph admits a shortcut partition and gave several applications, including a construction of tree cover for arbitrary planar graphs with stretch $1+\varepsilon$ and $O(1)$ many trees for any fixed $\varepsilon \in (0, 1)$. However, the construction heavily exploits planarity in multiple steps, and is thus inherently limited to planar graphs. In this work, we breach the "planarity barrier" to construct a shortcut partition for K_r-minor-free graphs for any r. To this end, we take a completely different approach our key contribution is a novel deterministic variant of the cop decomposition in minor-free graphs [And86, AGG+14]. As a highlight of our work, we employ our shortcut partition to resolve a major open problem the Steiner point removal (SPR) problem: Given any set K of terminals in an arbitrary edge-weighted planar graph G, it is possible to construct a minor M of G whose vertex set is K, which preserves the shortest-path distances between all pairs of terminals in G up to a constant factor.

Hsien-Chih Chang
Department of Computer Science
Dartmouth College
hsien-chih.chang@dartmouth.edu

Jonathan Conroy
Dartmouth College
Jonathan.Conroy.GR@dartmouth.edu

Hung Le
University of Massachusetts Amherst
hungle@cs.umass.edu

Lazar Milenkovic, Shay Solomon
Tel Aviv University

milenkovic.lazar@gmail.com, solo.shay@gmail.com

Cuong Than
Manning CICS, UMass Amherst
cthan@cs.umass.edu

CP36

Induced-Minor-Free Graphs: Separator Theorem, Subexponential Algorithms, and Improved Hardness of Recognition

A graph G contains a graph H as an induced minor if H can be obtained from G by vertex deletions and edge contractions. The class of H-induced-minor-free graphs generalizes the class of H-minor-free graphs, but unlike H-minor-free graphs, it can contain dense graphs. We show that if an n-vertex m-edge graph G does not contain a graph H as an induced minor, then it has a balanced vertex separator of size $O_H(\sqrt{m})$, where the $O_H(\cdot)$-notation hides factors depending on H. We give an algorithm for finding either an induced minor model of H in G or such a separator in randomized polynomial-time. We apply this to obtain subexponential $2^{O_H(n^{2/3}\log n)}$ time algorithms on H-induced-minor-free graphs for a large class of problems including maximum independent set, minimum feedback vertex set, 3-coloring, and planarization. For graphs H where every edge is incident to a vertex of degree at most 2, our results imply a $2^{O_H(n^{2/3}\log n)}$ time algorithm for testing if G contains H as an induced minor. Our second main result is that there exists a fixed tree T, so that there is no $2^{o(n/\log^3 n)}$ time algorithm for testing for T as an induced minor unless the Exponential Time Hypothesis (ETH) fails. This solves an open problem asked by Fellows, Kratochvil, Middendorf, and Pfeiffer [Algorithmica, 1995], who asked if there exists a fixed planar graph H so that testing for H as an induced minor is NP-hard.

Tuukka Korhonen
University of Bergen
tuukka.korhonen@uib.no

Daniel Lokshtanov
UCSB
daniello@ucsb.edu

CP36

VC Set Systems in Minor-Free (Di)Graphs and Applications

A recent line of work on VC set systems in minor-free (undirected) graphs, starting from Li and Parter, who constructed a new VC set system for planar graphs, has given surprising algorithmic results. In this work, we initialize a more systematic study of VC set systems for minor-free graphs and their applications in both undirected graphs and directed graphs (a.k.a digraphs). More precisely: We propose a new variant of Li-Parter set system for undirected graphs. Our set system settles two weaknesses of Li-Parter set system: the terminals can be anywhere, and the graph can be K_h-minor-free for any fixed h. We extend our set system to K_h-minor-free digraphs and show that its VC dimension is $O(h^2)$. We use this result to design the first subquadratic time algorithm for computing (unweighted) diameter and all-vertices eccentricities in K_h-minor-free digraphs. We show that the system of directed balls in minor-free digraphs has VC dimension at most $h-1$. On the negative side, we show that VC set system constructed from shortest path trees of planar digraphs does not have a bounded VC dimension. The highlight of our work is the results for digraphs, as we are not aware of known algorithmic work on constructing and exploiting VC set systems for digraphs.

Hung Le
University of Massachusetts Amherst
hungle@cs.umass.edu

Christian Wulff-Nilsen
University of Copenhagen
cwulnil@gmail.com

CP36

Sparse Induced Subgraphs in P_6-Free Graphs

We prove that a number of computational problems that ask for the largest sparse induced subgraph satisfying some property definable in CMSO_2 logic, most notably Feedback Vertex Set, are polynomial-time solvable in the class of P_6-free graphs. This generalizes the work of Grzesik, Klimoovilipczuk, and Pilipczuk on the Maximum Weight Independent Set problem in P_6-free graphs [SODA 2019, TALG 2022], and of Abrishami, Chudnovsky, Pilipczuk, Rzazewski, and Seymour on problems in P_5-free graphs [SODA 2021]. The key step is a new generalization of the framework of potential maximal cliques. We show that instead of listing a large family of potential maximal cliques, it is sufficient to only list their carvers: vertex sets that contain the same vertices from the sought solution and have similar separation properties.

Maria Chudnovsky
Princeton University
mchudnov@math.princeton.edu

Rose McCarty
Princeton University, USA
rm1850@math.princeton.edu

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

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

Pawel Rzazewski
Warsaw University of Technology, Poland
Institute of Informatics, University of Warsaw, Poland
pawel.rzazewski@pw.edu.pl

CP36

Odd Cycle Transversal on P_5-Free Graphs in Quasi-Polynomial Time

An independent set in a graph G is a set of pairwise non-adjacent vertices. A graph G is bipartite if its vertex set can be partitioned into two independent sets. In the Odd Cycle Transversal problem, the input is a graph G along with a weight function w associating a real weight with each vertex, and the task is to find a smallest weight vertex subset S in G such that G - S is bipartite; the weight of S, $w(S) = \sum_{v\in S} w(v)$. We show that Odd Cycle Transversal admits an algorithm with running time $n^{O(\log^2 n)}$ on graphs excluding P_5 (a path on five vertices) as an induced subgraph. The

problem was previously known to be polynomial time solvable on P_4-free graphs and NP-hard on P_6-free graphs [Dabrowski, Feghali, Johnson, Paesani, Paulusma and Rzażewski, Algorithmica 2020]. Bonamy, Dabrowski, Feghali, Johnson and Paulusma [Algorithmica 2019] posed the existence of a polynomial time algorithm on P_5 as an open problem, this was later restated by Rzażewski [Dagstuhl Reports, 9(6) : 2019] and by Chudnovsky, King, Pilipczuk, Rzażewski, and Spirkl [SIDMA 2021], who gave an algorithm with running time $n^{O(\sqrt{n})}$. While our $n^{O(\log^2 n)}$ time algorithm falls short of completely resolving the complexity status of Odd Cycle Transversal on P_5-free graphs it shows that the problem is not NP-hard unless every problem in NP is solvable in quasi-polynomial time.

Roohani Sharma
Max Planck Institute of Informatics
roohani.sharma90@gmail.com

Akanksha Agrawal
IIT Madras
India
akanksha@cse.iitm.ac.in,

Paloma Lima
ITU Copenhagen
Denmark
palt@itu.dk,

Daniel Lokshtanov
University of California, Santa Barbara
USA
daniello@ucsb.edu,

Saket Saurabh
IMSc +UiB
saket@imsc.res.in

CP37

Breaking the Metric Voting Distortion Barrier

Suppose we have an election with n voters and m candidates who lie in a metric space. We would like to design a voting rule that chooses a candidate whose average distance to the voters is small. However, instead of having direct access to the distances in the metric space, each voter ranks the candidates in order of distance. Can we design a rule that regardless of the election instance and underlying metric space, chooses a candidate whose cost differs from the true optimum by only a small factor (known as the distortion)? A long line of work culminated in finding deterministic voting rules with metric distortion 3, which is the best possible for deterministic rules and many other classes of voting rules. However, without any restrictions, there is a significant gap in our understanding: Though the best lower bound is 2.112, the best upper bound is still 3. Finding a rule that guarantees distortion $3-\epsilon$ for some constant ϵ has been a major challenge in computational social choice. In this work, we give a rule that guarantees distortion less than 2.753. To do so we study a handful of voting rules that are new to the problem. One is Maximal Lotteries, a rule which dates back to the 60's. The others are novel rules that can be thought of as hybrids of Random Dictatorship and the Copeland rule. Though none of these rules can beat distortion 3 alone, a randomization between Maximal Lotteries and any of the novel rules can.

Moses Charikar
Stanford University, California
moses@cs.stanford.edu

Prasanna Ramakrishnan, Kangning Wang
Stanford University
pras1712@stanford.edu, knwang@stanford.edu

Hongxun Wu
UC Berkeley
wuhx@berkeley.edu

CP38

Deterministic Near-Linear Time Minimum Cut in Weighted Graphs

In 1996, Karger gave a startling randomized algorithm that finds a minimum-cut in a (weighted) graph in time $O(m \log^3 n)$ which he termed near-linear time meaning linear (in the size of the input) times a polylogarthmic factor. In this paper, we give the first deterministic algorithm which runs in near-linear time for weighted graphs. In terms of techniques, we provide a structural theorem that says there exists a sparse clustering that preserves minimum cuts in a weighted graph with $o(1)$ error. In addition, we construct it deterministically in near linear time.

Jason Li
Simons Institute, UC Berkeley
jmli@alumni.cmu.edu

Monika Henzinger
Institute of Science and Technolology Austria
monika.henzinger@ist.ac.at

Satish Rao
University of California, Berkeley
satishr@berkeley.edu

Di Wang
Google
wadi@google.com

CP39

Edge-Weighted Online Stochastic Matching: Beating $1-\frac{1}{e}$

We study the edge-weighted online stochastic matching problem. Since Feldman, Mehta, Mirrokni, and Muthukrishnan introduced the online stochastic matching problem and proposed the $(1-\frac{1}{e})$-competitive Suggested Matching algorithm, there has been no improvement in the edge-weighted setting. In this paper, we introduce the first algorithm beating the $1-\frac{1}{e}$ barrier in this setting, achieving a competitive ratio of 0.645. Under the LP proposed by Jaillet and Lu, we design an algorithmic preprocessing, dividing all edges into two classes. Then we use different matching strategies to improve the performance on edges in one class in the early stage and on edges in another class in the late stage, while keeping the matching events of different edges highly independent. By balancing them, we finally guarantee the matched probability of every single edge.

Shuyi Yan
University of Copenhagen

shya@di.ku.dk

CP40

New Explicit Constant-Degree Lossless Expanders

We present a new explicit construction of onesided bipartite lossless expanders of constant degree, with arbitrary constant ratio between the sizes of the two vertex sets. Our construction is simpler to state and analyze than the only prior construction of Capalbo, Reingold, Vadhan, and Wigderson (2002), and achieves improved parameters. We construct our lossless expanders by imposing the structure of a constant-sized lossless expander "gadget" within the neighborhoods of a large bipartite spectral expander; similar constructions were previously used to obtain the weaker notion of unique-neighbor expansion. Our analysis simply consists of elementary counting arguments and an application of the expander mixing lemma.

Louis Golowich
UC Berkeley
lgolowich@berkeley.edu

CP41

Fast and Simple Unrooted Dynamic Forests

A dynamic forest data structure maintains a forest (and associated data like edge weights) under edge insertions and deletions. Dynamic forests are widely used to solve online and offline graph problems. Well-known examples of dynamic forest data structures are link-cut trees [Sleator and Tarjan '83] and top trees [Alstrup, Holm, de Lichtenberg, and Thorup '05], both of which need $O(\log n)$ time per operation. While top trees are more flexible and arguably easier to use, link-cut trees are faster in practice [Tarjan and Werneck '10]. In this paper, we propose an alternative to link-cut trees. Our data structure is based on search trees on trees (STTs, also known as elimination trees) and an STT algorithm [Berendsohn and Kozma '22] based on the classical Splay trees [Sleator and Tarjan '85]. While link-cut trees maintain a hierarchy of binary search trees, we maintain a single STT. Most of the complexity of our data structure lies in the implementation of the STT rotation primitive, which can easily be reused, simplifying the development of new STT-based approaches. We implement several variants of our data structure in the Rust programming language, along with an implementation of link-cut trees for comparison. Experimental evaluation suggests that our algorithms are faster when the dynamic forest is unrooted, while link-cut trees are faster for rooted dynamic forests.

Benjamin A. Berendsohn
Freie Universität Berlin
beab@zedat.fu-berlin.de

CP41

Practical Parallel Algorithms for Near-Optimal Densest Subgraphs on Massive Graphs

The densest subgraph problem has received significant attention, both in theory and in practice, due to its applications in problems such as community detection, social network analysis, and spam detection. Due to the high cost of obtaining exact solutions, much attention has focused on designing approximate densest subgraph algorithms. However, existing approaches are not able to scale to massive graphs with billions of edges. In this paper, we introduce a new framework that combines approximate densest subgraph algorithms with a pruning optimization. We design new parallel variants of the state-of-the-art sequential Greedy++ algorithm, and plug it into our framework in conjunction with a parallel pruning technique based on k-core decomposition to obtain parallel $(1+\epsilon)$-approximate densest subgraph algorithms. On a single thread, our algorithms achieve 2.6-34x speedup over Greedy++, and obtain up to 22.37x self-relative parallel speedup on a 30-core machine with two-way hyper-threading. Compared with the state-of-the-art parallel algorithm by Harb et al. [NeurIPS'22], we achieve up to a 114x speedup on the same machine. Finally, against the recent sequential algorithm of Xu et al. [PACMMOD'23], we achieve up to a 25.9x speedup.The scalability of our algorithms enables us to obtain near-optimal density statistics on the hyperlink2012 (with 113 billion edges) and clueweb (with 37 billion edges) graphs for the first time in the literature.

Pattara Sukprasert
Databricks, San Francisco, California
pattara.sk127@gmail.com

Quanquan C. Liu
MIT
quanquan@mit.edu

Laxman Dhulipala
University of Maryland, College Park, MD
laxman@umd.edu

Julian Shun
MIT
jshun@mit.edu

CP41

Constrained Planarity in Practice - Engineering the Synchronized Planarity Algorithm

In the constrained planarity setting, we ask whether a graph admits a planar drawing that additionally satisfies a given set of constraints. These constraints are often derived from very natural problems; prominent examples are Level Planarity, where vertices have to lie on given horizontal lines indicating a hierarchy, and Clustered Planarity, where we additionally draw the boundaries of clusters which recursively group the vertices in a crossing-free manner. Despite receiving significant amount of attention and substantial theoretical progress on these problems, only very few of the found solutions have been put into practice and evaluated experimentally. In this paper, we describe our implementation of the recent quadratic-time algorithm by Blus et al. [TALG Vol 19, No 4] for solving the problem Synchronized Planarity, which can be seen as a common generalization of several constrained planarity problems, including the aforementioned ones. Our experimental evaluation on an existing benchmark set shows that even our baseline implementation outperforms all competitors by at least an order of magnitude. We systematically investigate the degrees of freedom in the implementation of the Synchronized Planarity algorithm for larger instances and propose several modifications that further improve the performance. Altogether, this allows us to solve instances with up to 100 vertices in milliseconds and instances with up to 100 000 vertices within a few minutes.

Simon D. Fink, Ignaz Rutter
University of Passau

finksim@fim.uni-passau.de, rutter@fim.uni-passau.de

CP41

A Direct k-Way Hypergraph Partitioning Algorithm for Optimizing the Steiner Tree Metric

Minimizing wire-lengths is one of the most important objectives in circuit design. The process involves initially placing the logical units of a circuit onto a physical layout, and subsequently routing the wires to connect the cells. In this work, we present a novel HGP formulation that maps a hypergraph H, representing a logical circuit, onto a routing layout represented by a weighted graph G. The objective is to minimize the total length of all wires induced by the hyperedges of H on G. To capture wire-lengths, we compute minimum Steiner trees - a metric commonly used in routing algorithms. For this formulation, we present the first direct k-way multilevel mapping algorithm that incorporates techniques used by the highest-quality partitioning algorithms. We contribute a greedy mapping algorithm to compute an initial solution and three refinement algorithms to improve the initial mapping: Two move-based local search heuristics (based on label propagation and the FM algorithm) and a refinement algorithm based on max-flow min-cut computations. Our experiments demonstrate that our new algorithm achieves an improvement in the Steiner tree metric by 7% (median) on VLSI instances when compared to the best performing partitioning algorithm that optimizes the mapping in a postprocessing step. We achieve this improvement with only a 2-3 times slowdown in partitioning time compared to optimizing the connectivity metric.

Tobias Heuer
Karlsruhe Institute of Technology
tobias.heuer@gmx.net

CP41

Parallel Unconstrained Local Search for Partitioning Irregular Graphs

We present new refinement heuristics for the balanced graph partitioning problem that break with an age-old rule. Traditionally, local search only permits moves that keep the block sizes balanced (below a size constraint). In this work, we demonstrate that admitting *large* temporary balance violations drastically improves solution quality. The effects are particularly strong on irregular instances such as social networks. Designing efficient implementations of this general idea involves both careful selection of candidates for unconstrained moves as well as algorithms for rebalancing the solution later on. We explore a wide array of design choices to achieve this, in addition to our third goal of high parallel scalability. We present compelling experimental results, demonstrating that our parallel unconstrained local search techniques outperform the prior state of the art by a substantial margin. Compared with four state-of-the-art solvers, our new technique finds 75% of the best solutions on irregular graphs. We achieve a 9.6% improvement in edge cut over the next best competitor, while being only 7.7% slower in the geometric mean.

Nikolai Maas
Karlsruhe Institute of Technology (KIT)
nikolai.maas@kit.edu

Lars Gottesbüren, Daniel Seemaier
Karlsruhe Institute of Technology
lars.gottesbueren@kit.edu, daniel.seemaier@kit.edu

CP42

2-Fault-Tolerant Strong Connectivity Oracles

We study the problem of efficiently answering strong connectivity queries under two vertex failures. Given a directed graph G with n vertices, we provide a data structure with $O(nh)$ space and $O(h)$ query time, where h is the height of a decomposition tree of G into strongly connected subgraphs. This immediately implies data structures with $O(n \log n)$ space and $O(\log n)$ query time for graphs of constant treewidth, and $O(n^{3/2})$ space and $O(\sqrt{n})$ query time for planar graphs. For general directed graphs, we give a refined version of our data structure that achieves $O(n\sqrt{m})$ space and $O(\sqrt{m})$ query time, where m is the number of edges of the graph. We also provide some simple BFS-based heuristics that seem to work remarkably well in practice. In the experimental part, we first evaluate various methods to construct a decomposition tree with small height h in practice. Then we provide efficient implementations of our data structures, and evaluate their empirical performance by conducting an extensive experimental study on graphs taken from real-world applications.

Evangelos Kosinas, Loukas Georgiadis, Daniel Tsokaktsis
University of Ioannina
ekosinas@cs.uoi.gr, loukas@cse.uoi.gr, dtsokaktsis@cse.uoi.gr

CP42

Near-Optimal Coverage Path Planning with Turn Costs

Coverage path planning is a fundamental challenge in robotics, with diverse applications in aerial surveillance, manufacturing, cleaning, inspection, agriculture, and more. The main objective is to devise a trajectory for an agent that efficiently covers a given area, while minimizing time or energy consumption. Existing practical approaches often lack a solid theoretical foundation, relying on purely heuristic methods, or overly abstracting the problem to a simple Traveling Salesman Problem in Grid Graphs. Moreover, the considered cost functions only rarely consider turn cost, prize-collecting variants for uneven cover demand, or arbitrary geometric regions. In this paper, we describe an array of systematic methods for handling arbitrary meshes derived from intricate, polygonal environments. This adaptation paves the way to compute efficient coverage paths with a robust theoretical foundation for real-world robotic applications. Through comprehensive evaluations, we demonstrate that the algorithm also exhibits low optimality gaps, while efficiently handling complex environments. Furthermore, we showcase its versatility in handling partial coverage and accommodating heterogeneous passage costs, offering the flexibility to trade off coverage quality and time efficiency.

Dominik M. Krupke
TU Braunschweig
krupke@ibr.cs.tu-bs.de

CP42

Fast Many-to-Many Routing for Dynamic Taxi Sharing with Meeting Points

We introduce an improved algorithm for the dynamic taxi

sharing problem, i.e. a dispatcher that schedules a fleet of shared taxis as it is used by services like UberXShare and Lyft Shared. We speed up the basic online algorithm that looks for all possible insertions of a new customer into a set of existing routes, we generalize the objective function, and we efficiently support a large number of possible pick-up and drop-off locations. This lays an algorithmic foundation for taxi sharing systems with higher vehicle occupancy – enabling greatly reduced cost and ecological impact at comparable service quality. We find that our algorithm computes assignments between vehicles and riders several times faster than a previous state-of-the-art approach. Further, we observe that allowing meeting points for vehicles and riders can reduce the operating cost of vehicle fleets by up to 15% while also reducing rider wait and trip times.

Moritz Laupichler
Karlsruhe Institute of Technology
moritz.laupichler@kit.edu

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

CP42

Fast and Delay-Robust Multimodal Journey Planning

We study journey planning in the presence of vehicle delays for multimodal networks with public transit and an unrestricted transfer mode (e.g., walking, cycling). Delay-responsive algorithms receive a stream of delay updates, which are incorporated in an update phase. The query phase must return optimal journeys in the currently known delay scenario. The fastest known approach for multimodal journey planning is ULTRA, which precomputes shortcuts representing transfers between vehicles. This allows query algorithms to find optimal journeys without any performance loss compared to pure public transit networks. However, the precomputation phase does not account for delays and is too slow to rerun during the update phase. We present Delay-ULTRA, a delay-responsive variant of ULTRA. Since accounting for all theoretically possible delays is not feasible, our approach accounts for all delays up to a configurable limit. For delays above the limit, we run a heuristic search for missing shortcuts during the update phase. Our experimental evaluation on real-world data shows that Delay-ULTRA fails to find less than 0.02% of optimal journeys on metropolitan and mid-sized country networks, and 0.16% on the much larger Germany network. These error rates are negligible compared to errors in the input data in realistic applications. Query speed is at most twice as slow as ULTRA without delay information, and up to 8 times faster than the fastest algorithm with a preprocessing phase.

Jonas Sauer, Dominik Bez
Karlsruhe Institute of Technology
jonas.sauer2@kit.edu, dominik.bez@web.de

CP43

Counting Polyominoes, Revisited

A *polyomino* is an edge-connected set of squares on the square lattice. In this paper, we improve Jensen's algorithm for counting polyominoes by considering bounding boxes on the square lattice rotated by 45° instead of on the regular unrotated lattice. This allows us to extend significantly the count of polyominoes from 56 to 70 terms.

Gill Barequet
Technion - Israel Institute of Technology
barequet@cs.technion.ac.il

Gil Ben-Shachar
Technion - Israel Inst. of Technology
gilbe@cs.technion.ac.il

CP43

Simple and Robust Dynamic Two-Dimensional Convex Hull

The convex hull of a data set P is the smallest convex set that contains P. We present a new data structure for convex hull, that allows for efficient dynamic updates. In a dynamic convex hull implementation, the following traits are desirable: (1) algorithms for efficiently answering queries as to whether a specified point is inside or outside the hull, (2) adhering to geometric robustness, and (3) algorithmic simplicity.Furthermore, a specific but well-motivated type of two-dimensional data is rank-based data. Here, the input is a set of real-valued numbers Y where for any number $y \in Y$ its rank is its index in Y's sorted order. Each value in Y can be mapped to a point (rank,value) to obtain a two-dimensional point set. We present an efficient, geometrically robust, dynamic convex hull algorithm, that facilitates queries to whether a point is internal. Our improved solution is based on an algorithmic simplification of the classical convex hull data structure by Overmars and van Leeuwen [STOC'80], combined with new algorithmic insights. Our theoretical guarantees on the update time match those of Overmars and van Leeuwen, namely $O(\log^2 |P|)$, while we allow a wider range of functionalities (including rank-based data). Our algorithmic simplification includes simplifying an 11-case check down to a 3-case check that can be written in 20 lines of easily readable C-code.

Emil T. Gæde
Technical University of Denmark (DTU)
etoga@dtu.dk

Inge Li Gørtz
Danish Technical University, Denmark
inge@dtu.dk

Ivor Van Der Hoog
Technical University of Denmark
idjva@dtu.dk

Christoffer Krogh
Technical University of Denmark (DTU)
chrkro@dtu.dk

Eva Rotenberg
Technical University of Denmark
erot@dtu.dk

CP43

Covering Rectilinear Polygons with Area-Weighted Rectangles

Representing a polygon using a set of simple shapes has numerous applications in different use-case scenarios. We consider the problem of covering the interior of a rectilinear polygon with holes by a set of area-weighted, axis-aligned

rectangles such that the total weight of the rectangles in the cover is minimized. Already the unit-weight case is known to be $\mathcal{NP}$-hard and the general problem has, to the best of our knowledge, not been studied experimentally before. We show a new basic property of optimal solutions of the weighted problem. This allows us to speed up existing algorithms for the unit-weight case, obtain an improved ILP formulation for both the weighted and unweighted problem, and develop several approximation algorithms and heuristics for the weighted case. All our algorithms are evaluated in a large experimental study on 186837 polygons combined with six cost functions, which provides evidence that our algorithms are both fast and yield close-to-optimal solutions in practice.

Kathrin Hanauer, Martin Seybold, Julian Unterweger
University of Vienna
kathrin.hanauer@univie.ac.at, martin.seybold@univie.ac.at, julian.unterweger@univie.ac.at

CP43

Interactive Exploration Of The Temporal Alpha-Shape

Shape is a powerful tool to understand point sets. A formal notion of shape is given by α-shapes, which generalize the convex hull and provide adjustable level of detail. Many real-world point sets have an inherent temporal property as natural processes often happen over time, like lightning strikes during thunderstorms or moving animal swarms. To explore such point sets, where each point is associated with one timestamp, interactive applications may utilize α-shapes and allow the user to specify different time windows and α-values. We show how to compute the temporal α-shape α_T, a minimal description of all α-shapes over all time windows, in output-sensitive linear time. We also give complexity bounds on $|\alpha_T|$. We use α_T to interactively visualize α-shapes of user-specified time windows without having to constantly compute requested α-shapes. Experimental results suggest that our approach outperforms an existing approach by a factor of at least $\sim$52 and that the description we compute has reasonable size in practice. The basis for our algorithm is an existing algorithm which computes all Delaunay triangles over all time windows using $\mathcal{O}(1)$ time per triangle. Our approach generalizes to higher dimensions with the same runtime for fixed d.

Felix Weitbrecht
University of Stuttgart
weitbrfx@fmi.uni-stuttgart.de

CP44

Maintaining Discrete Probability Distributions in Practice

A classical problem in random number generation is the sampling of elements from a given discrete distribution. Formally, given a set of indices $S = \{1, \ldots, n\}$ and sequence of weights $w_1, \ldots, w_n \in \mathbb{R}^+$, the task is to provide samples from S with distribution $p(i) = w_i/W$ where $W = \sum_j w_j$. A commonly accepted solution is Walker's Alias Table, which allows for each sample to be drawn in constant time. However, some applications correspond to a dynamic setting, where elements are inserted or removed, or weights change over time. Here, the Alias Table is not efficient, as it needs to be re-built whenever the underlying distribution changes. In this paper, we engineer a simple data structure for maintaining discrete probability distributions in the dynamic setting. Construction of the data structure is possible in time $O(n)$, sampling is possible in expected time $O(1)$, and an update of size Δ can be processed in time $O(\Delta n/W)$. As a special case, we maintain an urn containing W marbles of n colors where with each update $O(W/n)$ marbles can be added or removed in $O(1)$ time per update. To evaluate the efficiency of the data structure in practice we conduct an empirical study. The results suggest that the dynamic sampling performance is competitive with the static Alias Table. Compared to existing more complex dynamic solutions we obtain a sampling speed-up of up to half an order of magnitude.

Daniel Allendorf
Goethe University Frankfurt
dallendorf@ae.cs.uni-frankfurt.de

CP44

ShockHash: Towards Optimal-Space Minimal Perfect Hashing Beyond Brute-Force

A minimal perfect hash function (MPHF) maps a set S of n keys to $[n]$ without collisions. The brute-force algorithm that tries random hash functions until stumbling on an MPHF and stores that function's seed matches the space lower bound. In expectation, $e^n\text{poly}(n)$ seeds need to be tested. The most space-efficient previous algorithms for constructing MPHFs all use such a brute-force approach as a basic building block. In this paper, we introduce ShockHash, which uses two hash functions h_0 and h_1, hoping for the existence of a function $f : S \to \{0, 1\}$ such that $x \mapsto h_{f(x)}(x)$ is an MPHF on S. In graph terminology, ShockHash generates n-edge random graphs until stumbling on a pseudoforest - a graph where each component contains as many edges as nodes. Using cuckoo hashing, ShockHash then derives an MPHF in linear time. It uses a 1-bit retrieval data structure to store f using $n + o(n)$ bits. We show that ShockHash needs to try only $(e/2)^n\text{poly}(n)$ hash function seeds in expectation. This reduces the space for storing the seed by roughly n bits (maintaining the asymptotically optimal space consumption) and speeds up construction by almost a factor of 2^n compared to brute-force. When using ShockHash as a building block within the RecSplit framework we obtain the currently most space efficient MPHFs, i.e., competing approaches need about two orders of magnitude more work to achieve the same space.

Hans-Peter Lehmann
Karlsruhe Institute of Technology
hans-peter.lehmann@kit.edu

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

Stefan Walzer
Karlsruhe Institute of Technology
stefan.walzer@kit.edu

CP44

Computing e-Th Roots in Number Fields

We describe algorithms for computing e-th roots of elements in a number field K, where e is an odd prime-power. In particular, we generalize Couveignes' and Thom algorithms to compute square-roots in the context of the General Number Field Sieve. Our algorithms cover most cases

of e and K and their complexity is better than general root finding algorithms. Our Python implementation compares extremely well to the one of these generic algorithms in well-known computer algebra softwares, allowing us to obtain reasonable timings even for large degree number fields and huge exponents e, which correspond to previously intractable cases. One important application of our algorithms consists in computing S-unit groups used for the cryptanalysis of the Ideal-SVP problem over cyclotomic fields in post-quantum cryptography. Indeed, in order to assess the efficiency of these so-called S-unit attacks, it is necessary to compute S-unit groups in large dimensions. As showed by Bernard al. in Asiacrypt 2022, this can be achieved by first considering S-units coming from the maximal totally real subfield and Stickelberger generators; then, a saturation step is required, which comes down to computing e-th roots where e is a prime-power factor of the relative class number, our largest example being a 93-bits prime. This paper tackles most cases, allowing for the explicit computation of S-unit groups even in dimension 200.

Andrea Lesavourey
INRIA
andrea.lesavourey@inria.fr

Olivier Bernard
Zama
olivier.bernard@normalesup.org

Pierre-Alain Fouque
Univ Rennes, IRISA
pa.fouque@gmail.com

CP44

Experimental Evaluation of Fully Dynamic K-Means via Coresets

For a set of points in ^{d}, the Euclidean k-means problems consists of finding k centers such that the sum of distances squared from each data point to its closest center is minimized. Coresets are one the main tools developed recently to solve this problem in a big data context. They allow to compress the initial dataset while preserving its structure: running any algorithm on the coreset provides a guarantee almost equivalent to running it on the full data. In this work, we study coresets in a fully-dynamic setting: points are added and deleted with the goal to efficiently maintain a coreset with which a solution can be computed. Based on an algorithm from Henzinger and Kale [ESA'20], we present an efficient and practical implementation of a fully dynamic coreset algorithm, that improves the running time by up to a factor of 20 compared to our non-optimized implementation of the algorithm by Henzinger and Kale, without sacrificing more than 7% on the quality of the solution.

David Saulpic
ISTA
david.saulpic@lip6.fr

Monika Henzinger
Institute of Science and Technolology Austria
monika.henzinger@ist.ac.at

Leonhard Sidl
University of Vienna
leonhard.sidl@gmx.at

CP45

An Alternate Proof of Near-Optimal Light Spanners

In 2016, a breakthrough result of Chechik and Wulff-Nilsen [SODA '16] established that every n-node graph G has a $(1+\varepsilon)(2k-1)$-spanner of lightness $O_\varepsilon(n^{1/k})$, and recent followup work by Le and Solomon [STOC '23] generalized the proof strategy and improved the dependence on ε. We give a new proof of this result, with the improved ε-dependence. Our proof is a direct analysis of the often-studied greedy spanner, and can be viewed as an extension of the folklore Moore bounds used to analyze spanner sparsity.

Greg Bodwin
University of Michigan
bodwin@umich.edu

CP45

Simple Linear-Size Additive Emulators

Given an input graph $G = (V, E)$, an additive emulator $H = (V, E', w)$ is a sparse weighted graph that preserves all distances in G with small additive error. A recent line of inquiry has sought to determine the best additive error achievable in the sparsest setting, when H has a linear number of edges. In particular, the work of [Kogan and Parter, ICALP 2023], following [Pettie, ICALP 2007], constructed linear size emulators with $+O(n^{0.222})$ additive error. It is known that the worst-case additive error must be at least $+\Omega(n^{2/29})$ due to [Lu, Vassilevska Williams, Wein, and Xu, SODA 2022]. We present a simple linear-size emulator construction that achieves additive error $+O(n^{0.191})$. Our approach extends the path-buying framework developed by [Baswana, Kavitha, Mehlhorn, and Pettie, SODA 2005] and [Vassilevska Williams and Bodwin, SODA 2016] to the setting of sparse additive emulators.

Gary T. Hoppenworth
University of Michigan
garyhoppenworth@gmail.com

CP45

Simpler Reductions from Exact Triangle

In this paper, we provide simpler reductions from Exact Triangle to two important problems in fine-grained complexity: Exact Triangle with Few Zero-Weight 4-Cycles and All-Edges Sparse Triangle. Exact Triangle instances with few zero-weight 4-cycles was considered by Jin and Xu [STOC 2023], who used it as an intermediate problem to show 3SUM hardness of All-Edges Sparse Triangle with few 4-cycles (independently obtained by Abboud, Bringmann and Fischer [STOC 2023]), which is further used to show 3SUM hardness of a variety of problems. We provide a simpler reduction from Exact Triangle to Exact Triangle with few zero-weight 4-cycles. Our new reduction also strengthens the conditional lower bounds from being under the 3SUM hypothesis to the even more believable Exact Triangle hypothesis. As a result, all conditional lower bounds shown by Jin and Xu [STOC 2023] and by Abboud, Bringmann and Fischer [STOC 2023] using All-Edges Sparse Triangle with few 4-cycles as an intermediate problem now also hold under the Exact Triangle hypothesis. We also provide two alternative proofs of the conditional lower bound of the All-Edges Sparse Triangle

problem under the Exact Triangle hypothesis, which was originally proved by Vassilevska Williams and Xu [FOCS 2020]. Both of our new reductions are simpler, and one of them is also deterministic—all previous reductions from Exact Triangle or 3SUM to All-Edges Sparse Triangle were randomized.

Timothy M. Chan
University of Illinois at Urbana-Champaign
tmc@illinois.edu

Yinzhan Xu
Massachusetts Institute of Technology
xyzhan@mit.edu

CP45

Listing 6-Cycles

Listing copies of small subgraphs (such as triangles, 4-cycles, small cliques) in the input graph is an important and well-studied problem in algorithmic graph theory. In this paper, we give a simple algorithm that lists t (non-induced) 6-cycles in an n-node undirected graph in $\tilde{O}(n^2 + t)$ time. This nearly matches the fastest known algorithm for detecting a 6-cycle in $O(n^2)$ time by Yuster and Zwick (1997). Previously, a folklore $O(n^2 + t)$-time algorithm was known for the task of listing 4-cycles.

Ce Jin
EECS
MIT
cejin@mit.edu

Virginia Vassilevska Williams
MIT
virgi@mit.edu

Renfei Zhou
Tsinghua University
zhourf20@mails.tsinghua.edu.cn

CP45

Linear-Sized Spectral Sparsifiers and the Kadison-Singer Problem

The Kadison-Singer Conjecture, as proved by Marcus, Spielman, and Srivastava (MSS) [Ann. Math. 182, 327-350 (2015)], has been informally thought of as a strengthening of Batson, Spielman, and Srivastavas theorem that every undirected graph has a linear-sized spectral sparsifier [SICOMP 41, 1704-1721 (2012)]. We formalize this intuition by using a corollary of the MSS result to derive the existence of spectral sparsifiers with a number of edges linear in their number of vertices for all undirected, weighted graphs. The proof consists of two steps. First, following a suggestion of Srivastava [Asia Pac. Math. Newsl. 3, 15-20 (2013)], we show the result in the special case of graphs with bounded leverage scores by repeatedly applying the MSS corollary to partition the graph, while maintaining an appropriate bound on the leverage scores of each subgraph. Then, we extend to the general case by constructing a recursive algorithm that repeatedly (i) divides edges with high leverage scores into multiple parallel edges and (ii) uses the bounded leverage score case to sparsify the resulting graph.

Phevos Paschalidis, Ashley Y. Zhuang
Harvard University
ppaschalidis@college.harvard.edu,
azhuang@college.harvard.edu

CP46

Simpler Constant Factor Approximation Algorithms For Weighted Flow Time – Now For Any p-Norm

A prominent problem in scheduling theory is the weighted flow time problem on one machine. We are given a machine and a set of jobs, each of them characterized by a processing time, a release time, and a weight. The goal is to find a (possibly preemptive) schedule for the jobs in order to minimize the sum of the weighted flow times, where the flow time of a job is the time between its release time and its completion time. Previously known approximation algorithms are quite complicated and involve for example a reduction to (geometric) covering problems, dynamic programs to solve those, and LP-rounding methods to reduce the running time. In this paper, we present a much simpler $(6 + \epsilon)$-approximation algorithm for the problem that works on the input jobs directly without any reduction. It even generalizes directly to an $O(1)$-approximation algorithm for minimizing the p-norm of the jobs' flow times, for any $0 < p < \infty$. Prior to our work, for $p > 1$ only a pseudopolynomial time $O(1)$-approximation algorithm was known, and no algorithm for $p < 1$. For the same objective functions, we present a very simple QPTAS for the setting of constantly many unrelated machines (assuming quasi-polynomially bounded input data). It works in the cases with and without the possibility to migrate a job to a different machine. This is the first QPTAS for the problem if migrations are allowed, and it is arguably simpler than the known QPTAS for the case without migration.

Alexander Armbruster
Technical University of Munich
alexander.armbruster@tum.de

Lars Rohwedder
Maastricht University
l.rohwedder@maastrichtuniversity.nl

Andreas Wiese
Technical University of Munich
andreas.wiese@tum.de

CP46

Improved Algorithms for Integer Complexity

Input your abstract, including TeX commands, here. The abstract should be no longer than 1500 characters, including spaces. Only input the abstract text. Don't include title or author information here.

Qizheng He
Department of Computer Science
University of Illinois at Urbana-Champaign
qizheng6@illinois.edu

CP46

Simple and Faster Algorithms for Knapsack

Input your abstract, including TeX commands, here. The abstract should be no longer than 1500 characters, including spaces. Only input the abstract text. Don't include

title or author information here.

Qizheng He
Department of Computer Science
University of Illinois at Urbana-Champaign
qizheng6@illinois.edu

Zhean Xu
Tsinghua University
xuza21@mails.tsinghua.edu.cn

CP46

Simple Approximation Algorithms for Minimizing the Total Weighted Completion Time of Precedence-Constrained Jobs

We consider the precedence-constrained scheduling problem to minimize the total weighted completion time. For a single machine, several 2-approximation algorithms are known, which are based on linear programming and network flows. We show that the same ratio is achieved by a simple weighted round-robin rule. Moreover, for preemptive scheduling on identical parallel machines, we give a strongly polynomial 3-approximation, which computes processing rates by solving a sequence of parametric flow problems. This matches the best known constant performance guarantee, previously attained only by a weakly polynomial LP-based algorithm. Our algorithms are both applicable in non-clairvoyant scheduling, where processing times are initially unknown. In this setting, our performance guarantees improve upon the best competitive ratio of 8 known so far.

Sven Jäger
RPTU Kaiserslautern-Landau
sven.jaeger@math.rptu.de

Philipp Warode
School of Business and Economics, HU Berlin
philipp.warode@hu-berlin.de

CP46

The Greedy Algorithm for the Shortest Common Superstring Problem Is a $\frac{1}{2}$-Approximation in Terms of Compression: a Simple Proof

In the shortest common superstring problem (SCS), one is given a set of strings and needs to find the shortest string that contains all of them as substrings. This problem is APX-hard, so approximation algorithms are of great interest. One of the simplest such algorithms is the greedy algorithm: while there is more than one string, find two of them with the largest overlap, merge them, and repeat. SCS can be viewed as a problem of minimizing the length of the solution as well as a problem of maximizing the *compression*, that is, the difference between the total length of input strings and the solution. While the approximation ratio of the greedy algorithm for the length of the superstring is still unknown with a lower bound of 2 and an upper bound of 3.425, it is known that for compression, the greedy algorithm is a $\frac{1}{2}$-approximation. This result was obtained in 1988 by Tarhio and Ukkonen through a rather complex analysis of the overlap graph, where vertices are the input strings and the weight of an edge between two strings is the length of their overlap. In this note, we present a simpler proof of this result, which consists of two parts: the first part is the classical observation that the weight of the maximum Hamiltonian path is at most twice the weight of the maximum matching, and the second part is that the weight of the greedy solution is at least the weight of the maximum matching.

Pavel Kalugin
Moscow Institute of Physics and Technology
kalugin.pe@gmail.com

Maksim Nikolaev
Steklov Institute of Mathematics at St. Petersburg
makc-nicko@yandex.ru

CP46

The Public University Secretary Problem

We consider a variation of the classic secretary problem where the goal is to hire the k cheapest secretaries. We use a beyond-worst-case approach where we assume that the online algorithm knows a priori an upper bound on the optimal cost, and we show this assumption is necessary. The main result is that the optimal competitive ratio is $\Theta(\log k)$. The upper bound holds even when candidates are interviewed in adversarial order, and is attained by a simple deterministic algorithm. The lower bound of $\Omega(\log k)$ holds against randomized algorithms and even when candidates are interviewed in random order.

Heather Newman, Benjamin Moseley
Carnegie Mellon University
hanewman@andrew.cmu.edu, moseleyb@andrew.cmu.edu

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

CP47

Ussr Is in P/poly

The Sum of Square Roots (SSR) problem is the following computational problem: Given positive integers $a_1, \ldots, a_k$, and signs $\delta_1, \ldots, \delta_k \in \{-1, 1\}$, check if $\sum_{i=1}^{k} \delta_i \sqrt{a_i} > 0$. The problem is known to have a polynomial time algorithm on the *real RAM* model of computation, however no sub-exponential time algorithm is known in the bit or Turing model of computation. The precise computational complexity of SSR has been a notorious open problem [**?**] over the last four decades. The problem is known to admit an upper bound in the third level of the *Counting Hierarchy*, i.e., CHtwoo and no non-trivial lower bounds are known. Even when the input numbers are *small*, i.e., given in *unary*, no better complexity bound was known prior to our work. In this paper, we show that the unary variant (USSR) of the sum of square roots problem is considerably easier by giving a P/poly upper bound.

Nikhil Balaji
IIT Delhi
nikhilrbalaji@gmail.com

Samir Datta
Chennai Mathematical Institute and UMI ReLaX,
Chennai, India
sdatta@cmi.ac.in

CP47

If Edge Coloring Is Hard under Seth, Then Seth Is False

The Edge Coloring problem is notoriously hard: it is still

unknown whether it can be solved in time $2^{o(n^2)}$ (let alone $2^{O(n)}$), where n is the number of nodes of the input graph. Can one explain the lack of such upper bounds by deriving a lower bound $2^{\Omega(n^2)}$ from a lower bound for SAT, 3-SUM, or APSP? In this note, we provide a negative answer for this question: if there is a reduction showing that Edge Coloring cannot be solved faster than in α^{n^2} (where $\alpha > 1$ is an explicit constant) under a hypothesis that known algorithms for one of the problems mentioned above are optimal, then the corresponding hypothesis is false.

Alexander Kulikov
Steklov Institute of Mathematics at St. Petersburg
Saint Petersburg University
alexander.s.kulikov@gmail.com

Ivan Mihajlin
University of CaliforniaSan Diego
ivmihajlin@gmail.com

CP47

Simple and Tight Complexity Lower Bounds for Solving Rabin Games

We give a simple proof that assuming the Exponential Time Hypothesis (ETH), determining the winner of a Rabin game cannot be done in time $2^{o(k \log k)} \cdot n^{O(1)}$, where k is the number of pairs of vertex subsets involved in the winning condition and n is the vertex count of the game graph. While this result follows from the lower bounds provided by Calude et al [SIAM J. Comp. 2022], our reduction is considerably simpler and arguably provides more insight into the complexity of the problem. In fact, the analogous lower bounds discussed by Calude et al, for solving Muller games and multidimensional parity games, follow as simple corollaries of our approach. Our reduction also highlights the usefulness of a certain pivot problem — Permutation SAT — which may be of independent interest.

Antonio Casares
LaBRI, Université de Bordeaux, France
antonio.casares-santos@labri.fr

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

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

Uéverton Souza
Universidade Federal Fluminense
University of Warsaw
ueverton@ic.uff.br

K. S. Thejaswini
University of Warwick, United Kingdom
Institute of Science and Technology, Austria
thejaswini.raghavan.1@warwick.ac.uk

CP47

A CS Guide to the Quantum Singular Value Transformation

We present a simplified exposition of some pieces of [Gily Su, Low, and Wiebe, STOC'19], which introduces a quantum singular value transformation (QSVT) framework for applying polynomial functions to block-encoded matrices. The QSVT framework has garnered substantial recent interest from the quantum algorithms community, as it was demonstrated by [GSLW19] to encapsulate many existing algorithms naturally phrased as an application of a matrix function. First, we posit that the lifting of quantum singular processing (QSP) to QSVT is better viewed not through Jordan's lemma (as was suggested by [GSLW19]) but as an application of the cosine-sine decomposition, which can be thought of as a more explicit and stronger version of Jordan's lemma. Second, we demonstrate that the constructions of bounded polynomial approximations given in [GSLW19], which use a variety of ad hoc approaches drawing from Fourier analysis, Chebyshev series, and Taylor series, can be unified under the framework of truncation of Chebyshev series, and indeed, can in large part be matched via a bounded variant of a standard meta-theorem from [Trefethen, 2013]. We hope this work finds use to the community as a companion guide for understanding and applying the powerful framework of [GSLW19].

Ewin Tang
University of California at Berkeley
ewin@berkeley.edu

Kevin Tian
Microsoft Research
tiankevin@microsoft.com

CP47

Quantum Logspace Computations Are Verifiable

In this note, we observe that quantum logspace computations are verifiable by classical logspace algorithms, with unconditional security. More precisely, every language in BQL has an (information-theoretically secure) streaming proof with a quantum logspace prover and a classical logspace verifier. The prover provides a polynomial-length proof that is streamed to the verifier. The verifier has a read-once one-way access to that proof and is able to verify that the computation was performed correctly. That is, if the input is in the language and the prover is honest, the verifier accepts with high probability, and, if the input is not in the language, the verifier rejects with high probability even if the prover is adversarial. Moreover, the verifier uses only $O(\log n)$ random bits.

Uma Girish, Ran Raz
Princeton University
ugirish@cs.princeton.edu, ranr@cs.princeton.edu

Wei Zhan
University of Chicago
jollwish@gmail.com

CP48

An Enumerative Perspective on Connectivity

Connectivity is an important measure in graph theory. Given a graph G with vertices s and t, the connectivity $\lambda(s,t)$ from s to t is defined to be the maximum number of edge-disjoint paths from s to t in G. Much research has gone into designing fast algorithms for computing connectivities in graphs. Previous work showed how to compute connectivities for all pairs of vertices in directed graphs with m edges in $\tilde{O}(m^\omega)$ time [Chueng, Lau, and Leung, FOCS 2011], where ω is the exponent of matrix multiplication. For the related problem of computing "small

connectivities," it is known that for any positive integer k, we can compute $\min(k, \lambda(s,t))$ for all pairs of vertices (s,t) in a directed graph with n nodes in $\tilde{O}((kn)^\omega)$ time [Akmal and Jin, ICALP 2023]. In this paper, we present alternate expositions of these $\tilde{O}(m^\omega)$ and $\tilde{O}((kn)^\omega)$ time algorithms, with simpler proofs of correctness. Earlier proofs were somewhat indirect, introducing an elegant but ad hoc "flow vector framework" for showing correctness of these algorithms. In contrast, we observe that these algorithms can be interpreted as testing whether certain generating functions enumerating families of edge-disjoint paths are nonzero. This new perspective yields more transparent proofs, and ties the approach for these problems more closely to the literature surrounding algebraic graph algorithms.

Shyan Akmal
EECS
MIT
naysh@mit.edu

CP48

Simple Analysis of Priority Sampling

We prove a tight upper bound on the variance of the priority sampling method (aka sequential Poisson sampling). Our proof is significantly shorter and simpler than the original proof given by Mario Szegedy at STOC 2006, which resolved a conjecture by Duffield, Lund, and Thorup.

Majid Daliri
New York University
majiddl.2099@gmail.com

CP48

A General Technique for Searching in Implicit Sets via Function Inversion

Given a function f from $[N]$ to a d-dimensional integer grid, we consider data structures that allow efficient orthogonal range queries in the image of f, without explicitly storing it. We show that, if f is of the form $[N] \to [2^w]^d$ for $w = \text{polylog}(N)$ and is computable in constant time, then, for any $0 < \alpha < 1$, we can obtain a data structure using $\tilde{O}(N^{1-\alpha/3})$ words of space s.t., for a given d-dimensional axis-aligned box B, we can search for $x \in [N]$ s.t. $f(x) \in B$ in time $\tilde{O}(N^\alpha)$. This is obtained by combining integer range searching with the Fiat-Naor function inversion scheme, which was already used in data-structure problems previously. We further obtain data structures for: range counting and reporting, predecessor, selection, ranking queries, and combinations thereof, on $f([N])$; preimage size and preimage selection queries; and selection and ranking queries on geometric quantities computed from tuples of points in d-space. These results unify and generalize previously known results on 3SUM-indexing and string searching, and are applicable as a black box to a variety of problems. In particular, we give a data structure for a generalized version of gapped string indexing, computing the Theil-Sen estimator of points contained in a given axis-aligned box, the kth largest area triangle, or the induced hyperplane that is the kth furthest from the origin.

Justin Dallant
Université libre de Bruxelles
justin.dallant@ulb.be

Boris Aronov
New York University
boris.aronov@nyu.edu

Jean Cardinal
ULB, Brussels
jean.cardinal@ulb.be

John Iacono
Université libre de Bruxelles
john.iacono@ulb.be

CP48

Finding the Saddlepoint Faster Than Sorting

A saddlepoint of an $n \times n$ matrix A is an entry of A that is a maximum in its row and a minimum in its column. Knuth (1968) gave several different algorithms for finding a saddlepoint. The worst-case running time of these algorithms is $\Theta(n^2)$, and Llewellyn, Tovey, and Trick (1988) showed that this cannot be improved, as in the worst case all entries of A may need to be queried. A strict saddlepoint of A is an entry that is the strict maximum in its row and the strict minimum in its column. The strict saddlepoint (if it exists) is unique, and Bienstock, Chung, Fredman, Scher, Shor, and Suri (1991) showed that it can be found in time $O(n \lg n)$, where a dominant runtime contribution is sorting the diagonal of the matrix. This upper bound has not been improved since 1991. In this paper we show that the strict saddlepoint can be found in $O(n \lg^* n) \subset o(n \lg n)$ time, where $\lg^*$ denotes the very slowly growing iterated logarithm function, coming close to the lower bound of $\Omega(n)$. In fact, we can also compute, within the same runtime, the value of a non-strict saddlepoint, assuming one exists. Our algorithm is based on a simple recursive approach, a feasibility test inspired by searching in sorted matrices, and a relaxed notion of saddlepoint.

Justin Dallant
Université libre de Bruxelles
justin.dallant@ulb.be

Frederik Haagensen
University of Copenhagen
haag@itu.dk

Riko Jacob
IT University Copenhagen
rikj@itu.dk

László Kozma
Freie Universität Berlin
laszlo.kozma@fu-berlin.de

Sebastian Wild
University of Liverpool
wild@liverpool.ac.uk

CP48

Sorting Signed Permutations by Reversals in Nearly-Linear Time

Given a signed permutation on n elements, we need to sort it with the fewest reversals. This is a fundamental algorithmic problem motivated by applications in comparative genomics, as it allows to accurately model rearrangements in small genomes. The first polynomial-time algorithm was given in the foundational work of Hannenhalli and Pevzner [J. ACM99]. Their approach was later streamlined and simplified by Kaplan, Shamir, and Tarjan [SIAM J. Com-

put.99] and their framework has eventually led to an algorithm that works in $O(n^{3/2}\sqrt{\log n})$ time given by Tannier, Bergeron, and Sagot [Discr. Appl. Math.07]. However, the challenge of finding a nearly-linear time algorithm remained unresolved. In this paper, we show how to leverage results on dynamic graph connectivity to obtain a surprisingly simple $O(n \log^2 n / \log\log n)$ time algorithm for this problem.

Bartlomiej Dudek, Pawel Gawrychowski
University of Wroclaw
bardek.dudek@gmail.com, gawry@cs.uni.wroc.pl

Tatiana Starikovskaya
Ecole Normale Supérieure
Paris
tat.starikovskaya@gmail.com

CP49

Convex Approximation and the Hilbert Geometry

Input your abstract, including TeX commands, here. The abstract should be no longer than 1500 characters, including spaces. Only input the abstract text. Don't include title or author information here.

Ahmed Abdelkader
University of Maryland
akader@cs.umd.edu

CP49

Dimension-Accuracy Tradeoffs in Contrastive Embeddings for Triplets, Terminals Top-K Nearest Neighbors

Metric embeddings traditionally study how to map n items to a target metric space such that distance lengths are not heavily distorted; but what if we only care to preserve the relative order of the distances (and not their length)? In this paper, we are motivated by the following basic question: given triplet comparisons of the form "item i is closer to item j than to item k,' can we find low-dimensional Euclidean representations for the n items that respect those distance comparisons? Such order-preserving embeddings naturally arise in important applications —recommendations, ranking, crowdsourcing, nearest-neighbor search. Our main results are: 1. Nearly-Tight Bounds on Triplet Dimension: We introduce the natural concept of triplet dimension of a dataset, and surprisingly, we show that in order for an ordinal embedding to be triplet-preserving, its dimension needs to grow as $\frac{n}{2}$ in the worst case. This is optimal (up to constant) as $n-1$ dimensions always suffice. 2. Tradeoffs for Dimension vs (Ordinal) Relaxation: We then relax the requirement that every triplet should be exactly preserved and present almost tight lower bounds for the maximum ratio between distances whose relative order was inverted by the embedding (ordinal relaxation). 3. New Bounds on Terminal and Top-k-NNs Embeddings: Going beyond triplets, we then study two well-motivated scenarios where we care about preserving specific sets of distances (not necessarily triplets).

Vaggos Chatziafratis
UC Santa Cruz
vaggos@ucsc.edu

Piotr Indyk
Massachusetts Institute of Technology
indyk@mit.edu

CP49

Fully Dynamic k-Center in Low Dimensions via Approximate Furthest Neighbors

Given point-set P, the k-center problem is to find a minimum radius r^* and set C of k points such that the distance from each point in P to its closest center in C is at most r^*. While NP-hard to exactly solve, there exists a famous $O(kn)$ time 2-approximation algorithm due to Gonzalez. This works by repeatedly adding the furthest neighbor in P to the current center set C until C contains exactly k centers. The dynamic version of the problem is to maintain P over insertions and deletions of points, in a way that permits efficiently solving the k center problem for the current P. There are various specialized $(2+\epsilon)$-approximation algorithms for solving this. The main new technical result of this paper is an algorithm for returning an approximate furthest neighbor in a dynamically maintained P from a query set C. We show that, for points in bounded doubling dimension, the approximate furthest neighbor problem can be solved using the known navigating nets data structure in a new way. This immediately provides a new algorithm for solving the dynamic k-center problem, by replacing the search for a furthest neighbor in P to C in Gonzalezs algorithm with this new approximate furthest neighbor search. Unlike some of the older algorithms, this new approach does not require knowing k or ϵ in advance. This new approach can also be used to solve the dynamic Euclidean k-center problem.

Mordecai Golin
Hong Kong University of Science & Technology, Hong Kong
golin@cse.ust.hk

Jinxiang Gan
Hong Kong University of Science & Technology
Hong Kong
jganad@connect.ust.hk

CP49

Revisiting Random Points: Combinatorial Complexity and Algorithms

Consider a set P of n points uniformly and independently from $[0,1]^d$ for a constant dimension d – such a point set is extremely well behaved in many aspects. For example, for a fixed $r \in [0,1]$, we prove a new concentration result on the number of pairs of points of P at a distance at most r – we show that this number lies in an interval that contains only $O(n \log n)$ numbers. We also present simple linear time algorithms to construct the Delaunay triangulation, Euclidean MST, and the convex hull of the points of P. The MST algorithm is an interesting divide-and-conquer algorithm which might be of independent interest. We also provide a new proof that the expected complexity of the Delaunay triangulation of P is linear – the new proof is simpler and more direct, and might be of independent interest. Finally, we present a simple $O(n^{4/3})polylog(n))$ time algorithm for the distance selection problem for $d = 2$.

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

Elfarouk Harb

University of Illinois at Urbana-Champaign (UIUC)
Eyfmharb@gmail.com

CP49

Insertion-Only Dynamic Connectivity in General Disk Graphs

Let $S \subseteq \mathbb{R}^2$ be a set of n sites in the plane, so that every site $s \in S$ has an associated radius $r_s > 0$. Let $D(S)$ be the disk intersection graph defined by S, i.e., the graph with vertex set S and an edge between two distinct sites $s, t \in S$ if and only if the disks with centers s, t and radii r_s, r_t intersect. Our goal is to design data structures that maintain the connectivity structure of $D(S)$ as S changes dynamically over time. We consider the incremental case, where new sites can be inserted into S. While previous work focuses on data structures whose running time depends on the ratio between the smallest and the largest site in S, we present a data structure with $O(\alpha(n))$ amortized query time and $O(\log^6 n)$ expected amortized insertion time. We also show that the same approach can be used for arbitrary intersection graphs.

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

Katharina Klost, Kristin Knorr, Wolfgang J. Mulzer
Freie Universität Berlin
kathklost@inf.fu-berlin.de, knorrkri@inf.fu-berlin.de, mulzer@inf.fu-berlin.de

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

CP49

Detecting Points in Integer Cones of Polytopes Is Double-Exponentially Hard

Let d be a positive integer. For a finite set $X \subseteq \mathbb{R}^d$, we define its integer cone as the set $\mathsf{IntCone}(X) := \{\sum_{x \in X} \lambda_x \cdot x \mid \lambda_x \in \mathbb{Z}_{\geq 0}\} \subseteq \mathbb{R}^d$. Goemans and Rothvoss showed that, given two polytopes $P, Q \subseteq \mathbb{R}^d$ with P being bounded, one can decide whether $\mathsf{IntCone}(P \cap \mathbb{Z}^d)$ intersects Q in time $\mathsf{enc}(P)^{2^{O(d)}} \cdot \mathsf{enc}(Q)^{O(1)}$ [J. ACM 2020], where $\mathsf{enc}(\cdot)$ denotes the number of bits required to encode a polytope through a system of linear inequalities. This result is the cornerstone of their XP algorithm for Bin Packing parameterized by the number of different item sizes. We complement their result by providing a conditional lower bound. In particular, we prove that, unless the ETH fails, there is no algorithm which, given a bounded polytope $P \subseteq \mathbb{R}^d$ and a point $q \in \mathbb{Z}^d$, decides whether $q \in \mathsf{IntCone}(P \cap \mathbb{Z}^d)$ in time $\mathsf{enc}(P, q)^{2^{o(d)}}$. Note that this does not rule out the existence of a fixed-parameter tractable algorithm for the problem, but shows that dependence of the running time on the parameter d must be at least doubly-exponential.

Marek Sokolowski
University of Warsaw
marek.sokolowski@mimuw.edu.pl

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

Alexandra Lassota
Max Planck Institute for Informatics, Saarbrücken, Germany
alassota@mpi-inf.mpg.de

Konrad Majewski
University of Warsaw
k.majewski@mimuw.edu.pl

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

CP50

A Simple (1 - ?)-Approximation Semi-Streaming Algorithm for Maximum (Weighted) Matching

We present a simple semi-streaming algorithm for $(1 - \epsilon)$-approximation of bipartite matching in $O(\log(n)/\epsilon)$ passes. This matches the performance of state-of-the-art "ϵ-efficient' algorithms—the ones with much better dependence on ϵ albeit with some mild dependence on n—while being considerably simpler. The algorithm relies on a direct application of the multiplicative weight update method with a self-contained primal-dual analysis that can be of independent interest. To show case this, we use the same ideas, alongside standard tools from matching theory, to present an equally simple semi-streaming algorithm for $(1-\epsilon)$-approximation of weighted matchings in general (not necessarily bipartite) graphs, again in $O(\log(n)/\epsilon)$ passes.

Sepehr Assadi
University of Waterloo
sepehr@assadi.info

CP50

Ssd Wear Leveling with Optimal Guarantees

Flash-based solid-state drives (SSDs) are fundamental building blocks of modern computer systems. However, due to the unique characteristics of flash media, each part of the SSD is limited in the number of times it can be written to. Maximizing the total number of write operations supported by the SSD is a challenging problem which has been addressed by the storage systems community for more than two decades. This optimization requires minimizing the internal garbage collection overheads of the SSD, and distributing the write operations evenly across the entire device. In this paper we address this challenge, considering it in both online and offline settings. In the online setting, we propose a simple randomized algorithm and prove its optimality. Compared to previous solutions, our algorithm both exhibits better performance and is effective for a wider (and more practical) range of SSD parameters. In the offline setting, we prove that the problem is NP-complete, and present an algorithm with near-optimal performance on any input.

Tomer Lange
Technion - Israel Institute of Technology
tlange@cs.technion.ac.il

Joseph Naor
Technion
naor@cs.technion.ac.il

Gala Yadgar
Technion - Israel Institute of Technology

gala@cs.technion.ac.il

CP50

When Stochastic Rewards Reduce to Deterministic Rewards in Online Bipartite Matching

We study the problem of vertex-weighted online bipartite matching with stochastic rewards where matches may fail with some known probability and the decision maker has to adapt to the sequential realization of these outcomes. Recent works have studied several special cases of this problem and it was known that the (randomized) Perturbed Greedy algorithm due to Aggarwal et al. (SODA, 2011) achieves the best possible competitive ratio guarantee of $(1-e^{-1})$ in some cases. We give a simple proof of these results by reducing (special cases of) the stochastic rewards problem to the deterministic setting of online bipartite matching (Karp, Vazirani, Vazirani (STOC, 1990)). More broadly, our approach gives conditions under which it suffices to analyze the competitive ratio of an algorithm for the simpler setting of deterministic rewards in order to obtain a competitive ratio guarantee for stochastic rewards. The simplicity of our approach reveals that the Perturbed Greedy algorithm has a competitive ratio of $(1-e^{-1})$ even in certain settings with correlated rewards, where no results were previously known. Finally, we show that without any special assumptions, the Perturbed Greedy algorithm has a competitive ratio strictly less than $(1-e^{-1})$ for vertex-weighted online matching with stochastic rewards.

Rajan Udwani
UC Berkeley
IEOR
rudwani@berkeley.edu

CP50

Growing a Random Maximal Independent Set Produces a 2-Approximate Vertex Cover

This paper proves that a simple greedy algorithm for finding a maximal independent set in a graph is also a randomized 2-approximation algorithm for weighted vertex cover. The unweighted version of the algorithm has existed for decades as a maximal independent set algorithm, but was not previously known to approximate vertex cover. This result leads to several new insights and simplified algorithms for different graph problems as corollaries. This includes a simple $O(\log n)$-round parallel algorithm for vertex cover, simplified approximation algorithms for certain edge deletion problems, connections to correlation clustering, and insights for list heuristic algorithms for vertex cover.

Nate Veldt
Texas A & M University
nveldt@tamu.edu

CP50

Simple and Asymptotically Optimal Online Bipartite Edge Coloring

We provide a simple online $\Delta(1+o(1))$-edge-coloring algorithm for bipartite graphs of maximum degree $\Delta = \omega(\log n)$ under adversarial vertex arrivals on one side of the graph. Our algorithm slightly improves the result of (Cohen, Peng and Wajc, FOCS19), which was the first, and currently only, to obtain an asymptotically optimal $\Delta(1+o(1))$ guarantee for an adversarial arrival model. More importantly, our algorithm provides a new, simpler approach for tackling online edge coloring.

Radu Vintan
EPFL
radu.vintan@epfl.ch

Joakim Blikstad
KTH Royal Institute of Technology
blikstad@kth.se

Ola Svensson
EPFL
ola.svensson@epfl.ch

David Wajc
Google Research
david.wajc@gmail.com

CP51

Conditional Lower Bounds for Sparse Parameterized 2-Csp: A Streamlined Proof

Assuming the Exponential Time Hypothesis (ETH), a result of Marx (ToC'10) implies that there is no $f(k) \cdot n^{o(k/\log k)}$ time algorithm that can solve 2-CSPs with k constraints (over a domain of arbitrary large size n) for any computable function f. This lower bound is widely used to show that certain parameterized problems cannot be solved in time $f(k)\cdot n^{o(k/\log k)}$ time (assuming the ETH). The purpose of this note is to give a streamlined proof of this result.

Karthik C. S.
Rutgers University
karthik.0112358@gmail.com

Dániel Marx
CISPA Helmholtz Center for Information Security
marx@cispa.de

Marcin Pilipczuk
University of Warsaw
m.pilipczuk@mimuw.edu.pl

Uéverton Souza
Universidade Federal Fluminense
University of Warsaw
ueverton@ic.uff.br

CP51

Minimum-Cost Paths for Electric Cars

An electric car equipped with a battery of a finite capacity travels on a road network with an infrastructure of charging stations. Each charging station has a possibly different cost per unit of energy. Traversing a given road segment requires a specified amount of energy that may be positive, zero or negative. The car can only traverse a road segment if it has enough charge to do so, and it cannot charge its battery beyond its capacity. More precisely, the charge in the battery can never drop below zero and it can never exceed its capacity. To travel from one point to another the car needs to choose a *travel plan* consisting of a path in the network and a recharging schedule that specifies how much energy to charge at each charging station on the path, making sure of having enough energy to reach the next charging station or the destination. The cost of the plan is the total charging cost along the chosen path. We reduce the prob-

lem of computing plans between every two junctions of the network to two problems: Finding optimal energetic paths when no charging is allowed and finding standard shortest paths. When there are no negative cycles in the network, we obtain an $O(n^3)$-time algorithm for computing all-pairs travel plans, where n is the number of junctions in the network. We obtain slightly faster algorithms under some further assumptions. We also consider the case in which a bound is placed on the number of rechargings allowed.

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

Dani Dorfman
Blavatnik school of computer science
Tel Aviv University
dani.i.dorfman@gmail.com

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

Robert Tarjan
Princeton University
ret@princeton.edu

Mikkel Thorup
University of Copenhagen
mikkel2thorup@gmail.com

CP51

Modern Hashing Made Simple

Modern work on hashing has led to hash tables with extraordinary guarantees. However, these data structures are too complex to be taught in (even an advanced) data structures course. In this paper, we show that this need not be the case: using standard machinery that we already teach, one can construct a simple hash table that offers guarantees much stronger than what are classically taught: Operations are $O(1)$-time with high probability; the hash table stores n k-bit items in $nk+O(n \log\log n)$ bits of space; and the hash table is dynamically resized, so the space bound holds with respect to the current size n at each time step.

William Kuszmaul
Massachusetts Institute of Technology
kuszmaul@mit.edu

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

Martin Farach-Colton
Rutgers University
martin@farach-colton.com

John Kuszmaul
Yale University
johnkuszmaul@gmail.com

CP51

Contention Resolution for the ℓ-Fold Union of a Matroid via the Correlation Gap

The correlation gap of a real-valued set function $f : 2^N \to \mathbb{R}_+$ [ADSY10] measures the worst-case ratio between two continuous extensions of f over all points in the unit cube; informally the gap measures the worst-case benefit of correlations between the variables. The correlation gap of any monotone submodular set function is known to be at most $(1 - 1/e)$, and this bound is tight even for the rank function of a uniform matroid of rank 1. Via a connection established in [CVZ14], this yields an optimal contention resolution scheme for rounding in a matroid polytope. In this paper, we study the correlation gap of the rank function of the ℓ-fold union of a matroid $\mathcal{M}$, denoted by $\mathcal{M}^\ell$, defined as the (matroid) union of ℓ-copies of $\mathcal{M}$. We prove that the correlation gap of $\mathcal{M}^\ell$, for *any* matroid $\mathcal{M}$, is at most $1 - \frac{\ell^\ell e^{-\ell}}{\ell!}$; this bound behaves as $1 - \frac{1}{\sqrt{2\pi\ell}}$ as ℓ grows. This generalizes the results in [Yan11, BFGG20, KS23]; they established this gap for the uniform matroid of rank ℓ which can be viewed as the ℓ-fold union of a uniform matroid of rank 1; moreover this bound is tight even for this special case. The correlation gap yields a corresponding contention resolution scheme for $\mathcal{M}^\ell$ which was the initial motivation for this work.

Junkai Song
The University of Hong Kong
dsgsjk@connect.hku.hk

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

Weizhong Zhang
Carnegie Mellon University
Tepper School of Business
weizhong@andrew.cmu.edu

CP51

Simpler Distribution Testing with Little Memory

We consider the question of distribution testing (specifically, uniformity and closeness testing) in the streaming setting, i.e., under stringent memory constraints. We improve on the results of Diakonikolas, Gouleakis, Kane, and Rao (2019) by providing considerably simpler algorithms, which remove some restrictions on the range of parameters and match their lower bounds.

Clément Canonne, Joy Qiping Yang
The University of Sydney
clement.canonne@sydney.edu.au, qyan6238@sydney.edu