SIAM 2026 Searchable Abstracts Document

ACM-SIAM Symposium on Discrete Algorithms (SODA26)

January 11–14, 2026 • Hyatt Regency Vancouver, Vancouver, Canada

SIAM Symposium on Algorithm Engineering and Experiments (ALENEX26)

January 11, 2026 • Hyatt Regency Vancouver, Vancouver, Canada

SIAM Symposium on Simplicity in Algorithms (SOSA26)

January 12–14, 2026 • Hyatt Regency Vancouver, Vancouver, Canada

This document was current as of December 23, 2025. Abstracts appear as submitted.

SIAM
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

Can We Speed Safely?

Often, algorithmic tasks can be greatly sped up for inputs that are promised to have certain structural properties, such as inputs that are assumed to be random, or to come from restricted classes of graphs. However, in practice, we rarely know if these promises hold, and verifying them can cost more than using a worst case algorithm. This talk surveys emerging lines of work that build trustworthy fast algorithms, by pairing speedups with weaker notions of testability. Our new algorithms, which either give an answer that is certified to be correct, or flag the input as one that does not satisfy the promised conditions, have complexities that are significantly faster than that of algorithms that do not rely on promises. This talk will discuss works that are joint with Arsen Vasilyan, Talya Eden, Cassandra Marcussen and Madhu Sudan.

Ronitt Rubinfeld
Massachusetts Institute of Technology
U.S.
ronitt@csail.mit.edu

IP2

The Four Color Theorem: Generalizations and Faster Algorithms

The celebrated Four Color Theorem (4CT) was conjectured by Francis Guthrie in 1852, and remained open for more than 100 years, until Appel and Haken found a proof in 1976. The proof is computer-assisted, and a simplified version was found by Robertson et al. in 1997. It has been known since 1880 (by Tait) that the 4CT is equivalent to stating that every 2-connected cubic planar graph is 3-edge-colorable. In our ongoing project, we aim to utilize today's machine power to extend the computer-assisted proof (and computer-free extension) of 4CT, giving significantly stronger results and providing a much faster algorithm. Topics include: 1) Three-edge cubic graphs on a surface: Here, we completely characterize all non-three-edge colorable cubic graphs (i.e., snarks) that can be embedded in a projective plane or in a torus (while the 4CT has no extension to other surfaces). 2) A much faster algorithm to 4-color planar graphs, namely: we can 4-color a planar graph of order n in $O(n \log n)$ time. This generalizes the $O(n^2)$ algorithm by Robertson et al. in 1997. 3) 4-coloring almost all vertices of graphs on a surface (including Albretson's conjecture from 1980).

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

IP3

Some Incremental Results, and the Future of TCS

I will begin by examining several recent theoretical results whose proofs were developed (almost entirely) by large language models. I will then discuss what this moment might mean for the theory community, and how we can thoughtfully navigate the opportunities and risks inherent in the age of LLM-accelerated research.

Gregory Valiant
Stanford
valiant@stanford.edu

CP0

Best Paper Award 2: Is Nasty Noise Actually Harder Than Malicious Noise?

We consider the relative abilities and limitations of computationally efficient algorithms for learning in the presence of noise, under two well-studied and challenging adversarial noise models for learning Boolean functions: - malicious noise, in which an adversary can arbitrarily corrupt a random subset of examples given to the learner; and - nasty noise, in which an adversary can arbitrarily corrupt an adversarially chosen subset of examples given to the learner. We consider both the distribution-independent and fixed-distribution settings. Our main results highlight a dramatic difference between these two settings: 1. For distribution-independent learning, we prove a strong equivalence between the two noise models: If a class $\mathcal{C}$ of functions is efficiently learnable in the presence of η-rate malicious noise, then it is also efficiently learnable in the presence of η-rate nasty noise. 2. In sharp contrast, for the fixed-distribution setting we show an arbitrarily large separation: Under a standard cryptographic assumption, for any arbitrarily large value r there exists a concept class for which there is a ratio of r between the rate $\eta_{\mathrm{malicious}}$ of malicious noise that polynomial-time learning algorithms can tolerate, versus the rate η_{nasty} of nasty noise that such learning algorithms can tolerate. To offset the negative result given in (2) for the fixed-distribution setting, we define a broad and natural class of algorithms, namely those that ignore contradictory examples (ICE). We show that for these algorithms, malicious noise and nasty noise are equivalent up to a factor of two in the noise rate: Any efficient ICE learner that succeeds with η-rate malicious noise can be converted to an efficient learner that succeeds with $\eta/2$-rate nasty noise. We further show that the above factor of two is necessary, again under a standard cryptographic assumption. As a key ingredient in our proofs, we show that it is possible to efficiently amplify the success probability of nasty noise learners in a black-box fashion. Perhaps surprisingly, this was not previously known; it turns out to be an unexpectedly non-obvious result which we believe may be of independent interest.

Guy Blanc
Stanford University
gblanc@stanford.edu

Yizhi Huang, Tal Malkin, Rocco A. Servedio
Columbia University
yizhi@cs.columbia.edu, tal@cs.columbia.edu, rocco@cs.columbia.edu

CP0

Best Paper Award 1: Existence of Fair and Efficient Allocation of Indivisible Chores

We study the problem of allocating indivisible chores among agents with additive cost functions in a fair and efficient manner. A major open question in this area is whether there always exists an allocation that is envy-free up to one chore (EF1) and Pareto optimal (PO). Our main contribution is to provide a positive answer to this question by proving the existence of such an allocation for indivisible chores under additive cost functions. This is achieved by a novel combination of a fixed point argument and a discrete algorithm, providing a significant methodological advance in this area. Our additional key contributions are as follows. We show that there always exists an allocation that is EF1 and fractional Pareto optimal (fPO), where fPO is a stronger efficiency concept than PO. We also show that an EF1 and PO allocation can be computed in polynomial time when the number of agents is constant. Finally, we extend all of these results to the more general setting of weighted EF1 (wEF1), which accounts for the entitlements of agents.

Ryoga Mahara
University of Tokyo
mahara@mist.i.u-tokyo.ac.jp

CP0

Best Student Paper Award: Online Resource Allocation with Concave, Diminishing-Returns Objectives

Online resource allocation problems are central challenges in economics and computer science, modeling situations in which n items arriving one at a time must each be immediately allocated among m agents. In such problems, our objective is to maximize a monotone reward function $f(\mathbf{x})$ over the allocation vector $\mathbf{x} = (x_{ij})_{i,j}$, which describes the amount of each item given to each agent. When f is concave and has "diminishing returns" (monotone decreasing gradient), several lines of work over the past two decades have had great success designing constant-competitive algorithms. Notably, while a greedy algorithm is $\frac{1}{2}$-competitive, prior works have shown a competitive ratio of $1 - \frac{1}{e}$ in many settings where items are divisible (i.e. allowing fractional allocations). However, these works use a variety of problem-specific techniques, leaving open the general question: Does a $(1 - \frac{1}{e})$-competitive fractional algorithm always exist for online resource allocation problems with concave, diminishing-returns objectives? In this work, we answer this question affirmatively, thereby unifying and generalizing prior results for special cases. Our algorithm is one which makes continuous greedy allocations with respect to an auxiliary objective $U(x)$ defined in terms of f. We then analyze this algorithm using a new, general perspective on classic primal-dual methods.

Kalen R. Patton
Georgia Institute of Technology
kpatton33@gatech.edu

CP1

Tight Differentially Private PCA via Matrix Coherence

We revisit the task of computing the span of the top r singular vectors $u_1, \ldots, u_r$ of a matrix under differential privacy. We show that a simple and efficient algorithmbased on singular value decomposition and standard perturbation mechanismsreturns a private rank-r approximation whose error depends only on the rank-r coherence of $u_1, \ldots, u_r$ and the spectral gap $\sigma_r - \sigma_{r+1}$. This resolves a question posed by Hardt and Roth. Our estimator outperforms the state of the artsignificantly so in some regimes. In particular, we show that in the dense setting, it achieves the same guarantees for single-spike PCA in the Wishart model as those attained by optimal non-private algorithms, whereas prior private algorithms failed to do so. In addition, we prove that (rank-r) coherence does not increase under Gaussian perturbations. This implies that any estimator based on the Gaussian mechanismincluding ourspreserves the coherence of the input. We conjecture that similar behavior holds for other structured models, including planted problems in graphs. We also explore applications of coherence to graph problems. In particular, we present a differentially private algorithm for Max-Cut and other constraint satisfaction problems under low coherence assumptions.

Tommaso D'Orsi
Bocconi
tommaso.dorsi@unibocconi.it

Gleb Novikov
Lucerne School of Computer Science
gleb.novikov@hslu.ch

CP1

Stochastic Embedding of Digraphs into Dags

Given a weighted digraph $G = (V, E, w)$, a stochastic embedding into DAGs is a distribution $\mathcal{D}$ over pairs of DAGs (D_1, D_2) such that for every u, v: (1) the reachability is preserved: $u \rightsquigarrow_G v$ (i.e., v is reachable from u in G) implies that $u \rightsquigarrow_{D_1} v$ or $u \rightsquigarrow_{D_2} v$ (but not both), and (2) distances are dominated: $d_G(u,v) \le \min\{d_{D_1}(u,v), d_{D_2}(u,v)\}$. The stochastic embedding $\mathcal{D}$ has expected distortion t if for every $u, v \in V$,

$$\mathbb{E}_{(D_1,D_2)\sim\mathcal{D}}\left[d_{D_1}(u,v)\cdot\mathbb{1}[u \rightsquigarrow_{D_1} v] + d_{D_2}(u,v)\cdot\mathbb{1}[u \rightsquigarrow_{D_2} v]\right] \le t \cdot d_G(u,v)\,.$$

Finally, the sparsity of $\mathcal{D}$ is the maximum number of edges in any of the DAGs in its support. Given an n vertex digraph with m edges, we construct a stochastic embedding into DAGs with expected distortion $\tilde{O}(\log n)$ and $\tilde{O}(m)$ sparsity, improving a previous result by Assadi, Hoppenworth, and Wein [STOC 25], which achieved expected distortion $\tilde{O}(\log^3 n)$. Further, we can sample DAGs from this distribution in $\tilde{O}(m)$ time.

Arnold Filtser
Bar-Ilan University
filtsea@cs.biu.ac.il

CP1

Lower Bounds for the Universal TSP on the Plane

We show a lower bound for the universal traveling salesman heuristic on the plane: for any linear order on the unit square $[0,1]^2$, there are finite subsets $S \subset [0,1]^2$ of arbitrarily large size such that the path visiting each element of S according to the linear order has length $\ge C\sqrt{\log|S|/\log\log|S|}$ times the length of the shortest path visiting each element in S. ($C > 0$ is a constant that depends only on the linear order.) This improves the previous lower bound $\ge C\sqrt[6]{\log|S|/\log\log|S|}$ of Hajiaghayi, Kleinberg and Leighton (SODA 2006). The proof establishes a dichotomy about any long walk on a cycle: the walk either zig-zags between two far away points, or else for a large amount of time it stays inside a set of small diameter.

Cosmas Kravaris
Princeton University
ck6221@princeton.edu

CP1

Online Connectivity Augmentation

The Connectivity Augmentation Problem (CAP) is a fundamental problem in fault-tolerant network design. In this work, we consider CAP in the online setting: given a k-edge-connected graph $G = (V, E)$ and a set L of additional edges (links), online requests arrive, each specifying vertices $\{u, v\}$ that need to be $(k+1)$-edge-connected. For each request, an algorithm must irrevocably pick links from L so that u and v are $(k+1)$-edge-connected, minimizing the total number (or weight) of picked links. The performance of the algorithm is measured by the competitive ratio against the optimal offline solution. Prior work gives a randomized $\tilde{O}(k \log^3 n)$ upper bound for online CAP; and for $k = 1$ (the online Tree Augmentation Problem, TAP), a deterministic $O(\log n)$ bound and an $\Omega(\log n)$ lower bound. We obtain the following results: (1) A deterministic $O(\log n)$-competitive algorithm for online CAP, improving over the previous randomized bound and matching the lower bound up to constants. (2) An $\Omega(\log n)$ lower bound in the random-order model, showing that random arrival order does not improve competitiveness, even for TAP. (3) A deterministic $O(\log^2 n)$-competitive algorithm for the weighted setting, improving upon the $\tilde{O}(k \log^3 n)$ bound. (4) A simple combinatorial $O(\log n)$-competitive algorithm for online TAP.

Aditya Subramanian
Indian Institute of Science
adityasubram@iisc.ac.in

Mohit Garg
Eindhoven University of Technology
University of Bremen
garg@uni-bremen.de

CP1

Improved Approximation for Ranking on General Graphs

In this paper, we study Ranking, a well-known randomized greedy matching algorithm, for general graphs. The algorithm was introduced by Karp et al.[STOC 1990] for the online bipartite matching problem, where it achieves a tight approximation ratio of $1 - 1/e$. The Ranking algorithm for general graphs is as follows: a permutation π over the vertices is chosen uniformly at random. The vertices are then processed sequentially according to π, with each vertex being matched to the first available neighbor according to the same permutation. While the algorithm is quite well-understood for bipartite graphs, its approximation ratio for general graphs remains less well-characterized,lying between 0.526 by Chan et al.[TALG 2018] and 0.727 by Karande et al.[STOC 2011]. In this work, we improve the approximation ratio of Ranking for general graphs to 0.5469. This also surpasses the best-known approximation ratio of 0.531 by Tang et al.[JACM 2023] for the oblivious matching problem. Our approach builds on the standard primal-dual analysis used for online matching. The novelty of our work lies in proving new structural properties by introducing the notion of the backup for vertices matched by the algorithm. For a fixed permutation, a vertex's backup is its potential match if its current match is removed. This concept helps characterize the rank distribution of the match of each vertex, enabling us to eliminate certain bad events that constrained previous work.

Tao Yu, Mahsa Derakhshan
Northeastern University
yu.tao2@northeastern.edu,
m.derakhshan@northeastern.edu

Mohammad Roghani
Stanford University
roghani@stanford.edu

Mohammad Saneian
Northeastern University
saneian.m@northeastern.edu

CP1

Low-Sensitivity Matching Via Sampling from Gibbs Distributions

In this work, we study the maximum matching problem from the perspective of sensitivity. The sensitivity of an algorithm A on a graph G is defined as the maximum Wasserstein distance between the output distributions of A on G and on $G - e$, where $G - e$ is the graph obtained by deleting an edge e from G. The maximum is taken over all edges e, and the underlying metric for the Wasserstein distance is the Hamming distance. We first show that for any $\varepsilon > 0$, there exists a polynomial-time $(1-\varepsilon)$-approximation algorithm with sensitivity $\Delta^{O(1/\varepsilon)}$. The algorithm is based on sampling from the Gibbs distribution over matchings and runs in time $O_{\varepsilon,\Delta}(m \log m)$.This result significantly improves the previously known sensitivity bounds. Next, we present significantly faster algorithms for planar and bipartite graphs as a function of ε and Δ, which run in time poly(n/ε). This improvement is achieved by designing a more efficient algorithm for sampling matchings from the Gibbs distribution in these graph classes, improving upon the previous best running time. Finally, for general graphs with potentially unbounded maximum degree, we show that there exists a polynomial-time $(1-\varepsilon)$-approximation algorithm with sensitivity $\sqrt{n} \cdot (\varepsilon^{-1} \log n)^{O(1/\varepsilon)}$, improving upon the previous best bound of $O(n^{1/(1+\varepsilon^2)})$.

Yuichi Yoshida, Zhang Zihan
National Institute of Informatics
yyoshida@nii.ac.jp, zihan@nii.ac.jp

CP2

Determinization of Min-Plus Weighted Automata Is Decidable

We show that the determinization problem for min-plus (tropical) weighted automata is decidable, thus resolving this long-standing open problem. In doing so, we develop a new toolbox for analyzing and reasoning about the run-structure of nondeterministic automata.

Shaull Almagor, Guy Arbel, Sarai Sheinvald
Technion Israel Institute of Technology
Haifa, Israel
shaull@technion.ac.il, guy.arbel@cs.technion.ac.il,
surke@cs.technion.ac.il

CP2

From Unweighted to Weighted Dynamic Matching in Non-Bipartite Graphs: A Low-Loss Reduction

We study the approximate maximum weight matching (MWM) problem in a fully dynamic graph subject to edge insertions and deletions. We design meta-algorithms that reduce the problem to the unweighted approximate maximum cardinality matching (MCM) problem. Despite recent progress on bipartite graphs – Bernstein-Dudeja-Langley (STOC 2021) and Bernstein-Chen-Dudeja-Langley-Sidford-Tu (SODA 2025) – the only previous meta-algorithm that applied to non-bipartite graphs suffered a $\frac{1}{2}$ approximation loss (Stubbs-Williams, ITCS 2017). We significantly close the weighted-and-unweighted gap by showing the first low-loss reduction that transforms any fully dynamic $(1-\varepsilon)$-approximate MCM algorithm on bipartite graphs into a fully dynamic $(1-\varepsilon)$approximate MWM algorithm on general (not necessarily bipartite) graphs, with only a $\text{poly}(\log n/\varepsilon)$ overhead in the update time. Central to our approach is a new primaldual framework that reduces the computation of an approximate MWM in general graphs to a sequence of approximate induced matching queries on an auxiliary bipartite extension. In addition, we give the first conditional lower bound on approximate partially dynamic matching with worst-case update time.

Aaron Bernstein
New York University
aaron.bernstein@nyu.edu

Jiale Chen
Stanford University
jialec@stanford.edu

CP2

On the Computation of Schrijver's Kernels

The geometry of a graph G embedded on a closed oriented surface S can be probed by counting the intersections of G with closed curves on S. Of special interest is the map $c \mapsto \mu_G(c)$ counting the minimum number of intersections between G and any curve freely homotopic to a given curve c. Schrijver [On the uniqueness of kernels, 1992] calls G a kernel if for any proper graph minor H of G we have $\mu_H < \mu_G$. Hence, G admits a minor H which is a kernel and such that $\mu_G = \mu_H$. We show how to compute such a minor kernel of G in $O(n^3 \log n)$ time where n is the number of edges of G, and $g \geq 2$ is the genus of S. Our algorithm leverages a tight bound on the size of minimal bigons in a system of closed curves. It also relies on several subroutines of independent interest including the computation of the area enclosed by a curve and a test of simplicity for the lift of a curve in the universal covering of S. As a consequence of our minor kernel algorithm and a recent result of Dubois [Making multicurves cross minimally on surfaces, 2024], after a preprocessing that takes $O(n^3 \log n)$ time and $O(n)$ space, we are able to compute $\mu_G(c)$ in $O(g(n+\ell)\log(n+\ell))$ time given any closed walk c with ℓ edges. This improves the state-of-the-art algorithm by Colin de Verdire and Erickson [Tightening non-simple paths and cycles on surfaces, 2010].

Vincent Delecroix, Oscar Fontaine
Université de Bordeaux
vincent.delecroix@u-bordeaux.fr,
oscar.fontaine@u-bordeaux.fr

Francis Lazarus
G-SCOP/Institut Fourier, CNRS, Université Grenoble Alpes
francis.lazarus@grenoble-inp.fr

CP2

Centered Colorings in Minor-Closed Graph Classes

A vertex coloring of a graph G is p-centered if for every connected subgraph H of G, either the coloring uses more than p colors on H, or there is a color that appears exactly once on H. We prove that for every fixed positive integer t, every K_t-minor-free graph admits a p-centered coloring using $\mathcal{O}(p^{t-1})$ colors.

Jedrzej Hodor
Jagiellonian University
jedrzej.hodor@gmail.com

Hoang La
Université Paris-Saclay
CNRS
hoang.la.research@gmail.com

Piotr Micek
Jagiellonian University
piotr.micek@uj.edu.pl

Clément Rambaud
Université Côte d'Azur
CNRS
clement.rambaud@inria.fr

CP2

From Incremental Transitive Cover to Strongly Polynomial Maximum Flow

We provide faster strongly polynomial time algorithms solving maximum flow in structured n-node m-arc networks. Our results imply an $n^{\omega+o(1)}$-time strongly polynomial time algorithms for computing a maximum bipartite b-matching where ω is the matrix multiplication constant. Additionally, they imply an $m^{1+o(1)}W$-time algorithm for solving the problem on graphs with a given tree decomposition of width W. We obtain these results by strengthening and efficiently implementing an approach in Orlin's (STOC 2013) state-of-the-art $O(mn)$ time maximum flow algorithm. We develop a general framework that reduces solving maximum flow with arbitrary capacities to (1) solving a sequence of maximum flow problems with polynomial bounded capacities and (2) dynamically maintaining a size-bounded supersets of the transitive closure under arc additions; we call this problem *incremental transitive cover*. Our applications follow by leveraging recent weakly polynomial, almost linear time algorithms for maximum flow due to Chen, Kyng, Liu, Peng, Gutenberg, Sachdeva (FOCS 2022) and Brand, Chen, Kyng, Liu, Peng, Gutenberg, Sachdeva, Sidford (FOCS 2023), and by developing incremental transitive cover data structures.

Daniel Dadush
Centrum Wiskunde & Informatica
dadush@cwi.nl

James Orlin
Massachusetts Institute of Technology
jorlin@mit.edu

Aaron Sidford
Stanford University
sidford@stanford.edu

László A. Végh
University of Bonn
lvegh@uni-bonn.de

CP2

Minimum s-t Cuts with Fewer Cut Queries

We study the problem of computing a minimum s–t cut in an unweighted, undirected graph via *cut queries*. In this model, the input graph is accessed through an oracle that, given a subset of vertices $S \subseteq V$, returns the size of the cut $(S, V \setminus S)$. This line of work was initiated by Rubinstein, Schramm, and Weinberg (ITCS 2018), who gave a randomized algorithm that computes a minimum s–t cut using $\widetilde{O}(n^{5/3})$ queries. A recent result by Anand, Saranurak, and Wang (SODA 2025) also matched this upper bound via a deterministic algorithm based on blocking flows. In this work, we present a new randomized algorithm that improves the cut-query complexity to $\widetilde{O}(n^{8/5})$. At the heart of our approach is a query-efficient subroutine that incrementally reveals the graph edge-by-edge while increasing the maximum s–t flow in the learned subgraph at a rate faster than classical augmenting-path methods. Notably, our algorithm is simple, purely combinatorial, and can be naturally interpreted as a recursive greedy procedure. As a further consequence, we obtain a *deterministic* and *combinatorial* two-party communication protocol for computing a minimum s–t cut using $\widetilde{O}(n^{11/7})$ bits of communication. This improves upon the previous best bound of $\widetilde{O}(n^{5/3})$, which was obtained via reductions from the aforementioned cut-query algorithms.

Yonggang Jiang
MPI-INF
yjiang@mpi-inf.mpg.de

Danupon Nanongkai
Max Planck Institute for Informatics
Saarland Informatics Campus
danupon@gmail.com

Pachara Sawettamalya
Princeton University
ps3122@princeton.edu

CP3

A Classification of Long-Refinement Graphs for Colour Refinement

The Colour Refinement algorithm is a classical procedure to detect symmetries in graphs, whose most prominent application is in graph-isomorphism tests. The algorithm and its generalisation, the Weisfeiler–Leman algorithm, iteratively evaluate local information to compute a colouring for the vertices. Different final colours of two vertices certify that no isomorphism can map one onto the other. The number of iterations that Colour Refinement takes to terminate is its central complexity parameter. For a long time, it was an open question whether graphs that take the maximum theoretically possible number of iterations actually exist. Kiefer and McKay found infinite families of such long-refinement (LR) graphs with degrees 2 and 3, thereby showing that the trivial upper bound on the iteration number is tight. In this work, we provide a complete characterisation of the LR graphs with small (or, equivalently, large) degrees. We show that, with one exception, the aforementioned families are the only LR graphs with maximum degree at most 3, and we fully classify the LR graphs with maximum degree 4 via a reverse-engineering approach. Incidentally, this yields that all LR graphs with degrees 2 and 3 can be represented as compact strings over a very small alphabet. Since the class of LR graphs is closed under taking edge complements, our analysis for small degrees also yields a classification of LR graphs with large degrees. Kiefer and McKay initiated a search for LR graphs that are only distinguished in the last iteration of Colour Refinement before termination. We conclude it in this submission by showing that such graphs cannot exist.

Sandra Kiefer, T. Devini De Mel
University of Oxford
sandra.kiefer@cs.ox.ac.uk, devini.demel@gmail.com

CP3

Near-Optimal Centerpoints in Polynomial Time in the Ambient Dimension

An α-centerpoint of a set of n input points in R^d, is a point such that any halfspace containing it, also contains an α-fraction of the input points. A recent breakthrough of Cherapanamjeri [FOCS 2024] gave the first polynomial time randomized algorithm to compute approximate $\Omega(1/d)$-centerpoints. We give a practical and efficient polynomial time randomized algorithm for computing $\Omega(\frac{1}{d\log^2 d})$-centerpoints of arbitrary pointsets. This improves on the longstanding $d^{O(d)}$ running time of Clarkson, Eppstein, Miller, Sturtivant and Teng [*IJCGA*, 1996] for obtaining such centerpoints. For low-dimensional pointsets, we also give an algorithm with complexity subexponential in the number of dimensions, and independent of the bit-size of the inputs. On the way, we also get a simple algorithm to compute $O(\sqrt{\frac{n}{d}})$-Tverberg partitions using a lifting trick of Sarkaria [*Isr. J. Math.* 1991], together with sampling ideas from Soberón [*Combinatorica*, 2019]. Again using linear programming, this can also be used to compute a point in the intersection of the parts contributing to the s-Tverberg partition. Lastly, we also give applications of our improved bounds to problems such as the Resilient Vector Consensus problem, lower-bounding functions with gradient oracles and deciding the lightness of a convex body.

Kunal Dutta
Institute of Informatics, University of Warsaw
kdutta@mimuw.edu.pl

Karol Pisula
Institute of Informatics, University of Warsaw
Poland
k.pisula@student.uw.edu.pl

CP3

Finite Pinwheel Scheduling: the K-Visits Problem

Pinwheel Scheduling is a fundamental scheduling problem, in which each task i is associated with a positive integer d_i, and the objective is to schedule one task per time slot, ensuring each task perpetually appears at least once in every d_i time slots. Although conjectured to be PSPACE-complete, it remains open whether Pinwheel Scheduling is NP-hard (unless a compact input encoding is used) or even contained in NP. We introduce k-Visits, a finite version of Pinwheel Scheduling, where given n tasks, the goal is to schedule each task exactly k times. While we observe that the 1-Visit problem is trivial, we prove that 2-Visits is strongly NP-complete through a reduction from Numerical 3-Dimensional Matching (N3DM). We further extend our strong NP-hardness result to a generalization of k-Visits in which the deadline of each task may vary throughout the schedule, as well as to a similar generalization of Pinwheel Scheduling, thus making progress towards settling the complexity of the latter. Additionally, we prove that 2-Visits can be solved in linear time if all deadlines are distinct. We achieve this through a Turing reduction from 2-Visits to a variation of N3DM, which we call Position

Matching. Based on this reduction, we also present an FPT algorithm for 2-Visits parameterized by a value related to how close the input deadlines are to each other, as well as a linear-time algorithm for instances with up to two distinct deadlines.

Sotiris Kanellopoulos, Christos Pergaminelis
National Technical University of Athens
Archimedes, Athena Research Center
sotkanel99@gmail.com, chrispergaminelis@gmail.com

Maria Kokkou
Aix Marseille Univ, CNRS, LIS
maria.kokkou@lis-lab.fr

Euripides Markou
University of Ioannina
Archimedes, Athena Research Center
emarkou@cse.uoi.gr

Aris Pagourtzis
National Technical University of Athens
Archimedes, Athena Research Center
pagour@cs.ntua.gr

CP3

An Unconditional Lower Bound for the Active-Set Method in Convex Quadratic Maximization

We prove that the active-set method needs an exponential number of iterations in the worst-case to maximize a convex quadratic function subject to linear constraints, regardless of the pivot rule used. This substantially improves over the best previously known lower bound [IPCO 2025], which needs objective functions of polynomial degrees $\omega(\log d)$ in dimension d, to a bound using a convex polynomial of degree 2. In particular, our result firmly resolves the open question [IPCO 2025] of whether a constant degree suffices, and it represents significant progress towards linear objectives, where the active-set method coincides with the simplex method and a lower bound for all pivot rules would constitute a major breakthrough. Our result is based on a novel extended formulation, recursively constructed using deformed products. Its key feature is that it projects onto a polygonal approximation of a parabola while preserving all of its exponentially many vertices. We define a quadratic objective that forces the active-set method to follow the parabolic boundary of this projection, without allowing any shortcuts along chords corresponding to edges of its full-dimensional preimage.

Eleon Bach
TU München
eleon.bach@proton.me

Yann Disser
TU Darmstadt
disser@mathematik.tu-darmstadt.de

Sophie Huiberts
CNRS
sophie@huiberts.me

Nils Mosis
TU Darmstadt
mosis@mathematik.tu-darmstadt.de

CP3

Interaction Between Skew-Representability, Tensor Products, Extension Properties, and Rank Inequalities

Skew-representable matroids form a fundamental class in matroid theory, bridging combinatorics and linear algebra. They play an important role in areas such as coding theory, optimization, and combinatorial geometry, where linear structure is crucial for both theoretical insights and algorithmic applications. Since deciding skew-representability is computationally intractable, much effort has been focused on identifying necessary or sufficient conditions for a matroid to be skew-representable. In this paper, we introduce a novel approach to studying skew-representability and structural properties of matroids and polymatroid functions via tensor products. We provide a characterization of skew-representable matroids, as well as of those representable over skew fields of a given prime characteristic, in terms of tensor products. As an algorithmic consequence, we show that deciding skew-representability, or representability over a skew field of fixed prime characteristic, is co-recursively enumerable: that is, certificates of non-skew-representability – in general or over a fixed prime characteristic – can be verified. Finally, as an application of the tensor product framework, we derive the first known linear rank inequality for folded skew-representable matroids that does not follow from the common information property.

Kristóf Bérczi
Eötvös Loránd University
Hungary
kristof.berczi@ttk.elte.hu

Boglárka Gehér
HUN-REN Alfréd Rényi Institute of Mathematics
HUN-REN Alfréd Rényi Institute of Mathematics
bogigeher@gmail.com

András Imolay
ELTE Eötvös Loránd University
andras.imolay@ttk.elte.hu

László Lovász
HUN-REN Alfréd Rényi Institute of Mathematics
laszlo.lovasz@ttk.elte.hu

Carles Padró
Universitat Politécnica de Catalunya
carles.padro@upc.edu

Tamás Schwarcz
London School of Economics and Political Science
t.b.schwarcz@lse.ac.uk

CP3

Cell-Probe Lower Bounds Via Semi-Random Csp Refutation: Simplified and the Odd-Locality Case

A recent work (Korten, Pitassi, and Impagliazzo, FOCS 2025) established an insightful connection between static data structure lower bounds, range avoidance of NC^0 circuits, and the refutation of pseudorandom CSP instances, leading to improvements to some longstanding lower bounds in the cell-probe/bit-probe models. Here, we

improve these lower bounds in certain cases via a more streamlined reduction to XOR refutation, coupled with handling the odd-arity case. Our result can be viewed as a complete derandomization of the state-of-the-art semi-random k-XOR refutation analysis (Guruswami, Kothari and Manohar STOC 2022, Hsieh, Kothari and Mohanty, SODA 2023), which complements the derandomization of the even-arity case obtained by Korten et al. As our main technical statement, we show that for any multi-output constant-depth circuit that substantially stretches its input, its output is very likely far from strings sampled from distributions with sufficient independence, and further this can be efficiently certified. Via suitable shifts in perspectives, this gives applications to cell-probe lower bounds and range avoidance algorithms for NC^0 circuits.

Weiqiang Yuan
EPFL
weiqiang.yuan@epfl.ch

Venkatesan Guruswami, Xin Lyu
UC Berkeley
venkatg@berkeley.edu, xinlyu@berkeley.edu

CP4

Short Circuit Walks in Fixed Dimension

Circuit augmentation schemes are a family of combinatorial algorithms for linear programming that generalize the simplex method. To solve the linear program, they construct a monotone circuit walk: They start at an initial vertex of the feasible region and traverse a discrete sequence of points on the boundary, while moving along certain allowed directions (circuits) and improving the objective function at each step until reaching an optimum. Since the existence of short circuit walks has been conjectured, several works have investigated how well one can efficiently approximate shortest monotone circuit walks towards an optimum. A first result addressing this question was given by De Loera, Kafer, and Sanit [SIAM J. Opt., 2022], who showed that finding a 2-approximation for this problem is NP-hard. Cardinal and the third author [Math. Prog. 2023] gave a stronger lower bound showing an approximation factor of $O(\frac{\log m}{\log \log m})$ is intractable for a linear program with m inequalities. Both reductions used high dimensional degenerate polytopes from combinatorial optimization. In this paper, we strengthen these hardness results by showing that for every fixed $\varepsilon > 0$ approximating the problem on polygons with m edges to within a factor of $O(m^{1-\varepsilon})$ is NP-hard. This result is essentially best-possible, as it cannot be improved beyond $o(m)$. In particular, this implies hardness for simple polytopes and in fixed dimension.

Alexander E. Black
Bowdoin College
a.black@bowdoin.edu

CP4

Feature Selection and Junta Testing Are Statistically Equivalent

For a function $f\colon \{0,1\}^n \to \{0,1\}$, the junta testing problem asks whether f depends on only k variables. If f depends on only k variables, the feature selection problem asks to find those variables. We prove that these two tasks are statistically equivalent. Specifically, we show that the "brute-force" algorithm, which checks for any set of k variables consistent with the sample, is simultaneously sample-optimal for both problems, and the optimal sample size is $\Theta\left(\frac{1}{\varepsilon}\left(\sqrt{2^k \log \binom{n}{k}} + \log \binom{n}{k}\right)\right)$.

Nathan Harms
University of British Columbia
nharms@cs.ubc.ca

Lorenzo Beretta
IBM
UC Santa Cruz
lorenzo2beretta@gmail.com

Caleb Koch
Stanford University
ckoch@stanford.edu

CP4

You (Almost) Can't Beat Brute Force for 3-Matroid Intersection

The ℓ-matroid intersection (ℓ-MI) problem asks if ℓ given matroids share a common basis. Already for $\ell = 3$, notable canonical NP-complete special cases are 3-Dimensional Matching and Hamiltonian Path on directed graphs. However, while these problems admit exponential-time algorithms that improve the simple brute force significantly, the fastest known algorithm for 3-MI on general matroids is exactly brute force with runtime $2^n/\mathrm{poly}(n)$, where n is the number of elements. Our main result shows that, in fact, brute force cannot be significantly improved, by ruling out an algorithm for ℓ-MI with runtime $o\left(2^{n-5\cdot n^{\frac{1}{\ell-1}}\cdot\log(n)}\right)$, for any fixed $\ell \geq 3$. For 3-MI, this gives a lower bound of $o\left(2^{n-5\cdot\sqrt{n}\cdot\log(n)}\right)$. We complement this result with (i) an algorithm which solves ℓ-MI faster than brute force in time $2^{n-\Omega\left(\log^2(n)\right)}$ for any $\ell \geq 3$, and (ii) a parameterized running time lower bound of $2^{(\ell-2)\cdot k\cdot\log k}\cdot\mathrm{poly}(n)$ for ℓ-MI, for any $\ell \geq 3$. We obtain these results by generalizing the monotone local search technique of [Fomin-Gaspers-Lokshtanov-Saurabh, J. ACM'19].

Ilan Doron-Arad
Technion - Israel Institute of Technology
Ilan Hayim
idoron-arad@cs.technion.ac.il

Ariel Kulik
Ben Gurion University
kulik@bgu.ac.il

Hadas Shachnai
Computer Science Department, Technion
hadas@cs.technion.ac.il

CP4

Halfspaces Are Hard to Test with Relative Error

Several recent works have studied a model of property testing of Boolean functions under a relative-error criterion. In this model, the distance from a target function $f : \{0,1\}^n \to \{0,1\}$ that is being tested to a function g is defined relative to the number of inputs x for which $f(x) = 1$; moreover, testing algorithms in this model have access both to a black-box oracle for f and to independent uniform satisfying assignments of f. The motivation for this model is that it provides a natural framework for testing sparse Boolean functions f that have few satisfying assignments, analogous to well-studied models for property testing of sparse graphs. The main result of this paper is a lower bound for testing halfspaces (i.e. linear threshold functions) in the relative error model: we show that $\tilde{\Omega}(\log n)$ oracle calls are required for any relative-error halfspace testing algorithm over the Boolean hypercube $\{0, 1\}^n$. This stands in sharp contrast both with the constant-query testability (independent of n) of halfspaces in the standard model, and with the positive results for relative-error testing given in recent works. Our lower bound for halfspaces gives the first example of a natural class of functions for which relative-error testing is provably more difficult than standard-model testing.

Xi Chen
Columbia University
xichen@cs.columbia.edu

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

Yizhi Huang
Columbia University
yizhi@cs.columbia.edu

Shivam Nadimpalli
MIT
shivamn@mit.edu

Rocco A. Servedio, Tianqi Yang
Columbia University
rocco@cs.columbia.edu, tianqi@cs.columbia.edu

CP4

Computational Complexity in Property Testing

We initiate a systematic study of the computational complexity of property testing, focusing on the relationship between query and time complexity. While traditional work in property testing has emphasized query complexityoften via information-theoretic techniquesrelatively little is known about the computational hardness of property testers. Our first contribution is a pair of *time-query hierarchy theorems* for property testing. For all suitable nondecreasing functions $q(n)$ and $t(n)$ with $t(n) \geq q(n)$, we construct properties with query complexity $\widetilde{\Theta}(q(n))$ and time complexity $\widetilde{\Omega}(t(n))$. We then turn to halfspaces in $\mathbb{R}^d$, a fundamental class in property testing and learning theory. We study the problem of approximating the distance from the input function to the nearest halfspace. For the distribution-free distance approximation problem, known algorithms achieve query complexity $O(d/\epsilon^2)$, but run in time $\widetilde{\Theta}(1/\epsilon^d)$. We provide a fine-grained justification for this gap: assuming the (integer) k-SUM conjecture, any algorithm must have running time $\Omega(1/\epsilon^{d/2})$. We also prove that any randomized Statistical Query (SQ) algorithm under the standard Gaussian distibution requires $(1/\epsilon)^{\Omega(d)}$ queries if the queries are answered with additive error up to $\epsilon^{\Omega(d)}$, revealing a fundamental barrier even in the distribution-specific setting.

Renato Ferreira Pinto JR.
Columbia University
r4ferrei@uwaterloo.ca

Diptaksho Palit, Sofya Raskhodnikova
Boston University
dpalit@bu.edu, sofya@bu.edu

CP4

Space-Efficient Population Protocols for Exact Majority on General Graphs

We study exact majority consensus in the population protocol model. In this model, the system is described by a graph $G = (V, E)$ with n nodes, and in each time step, a scheduler samples uniformly at random a pair of adjacent nodes to interact. In the exact majority consensus task, each node is given a binary input, and the goal is to design a protocol that almost surely reaches a stable configuration, where all nodes output the majority input value. We give improved upper and lower bounds for exact majority in general graphs. Specifically, our upper bound is parameterized by the relaxation time τ_{rel} of the random walk on G in the population protocol model.

Joel Rybicki, Jakob Solnerzik, Olivier Stietel
Humboldt University of Berlin
joel.rybicki@hu-berlin.de, jakob.solnerzik@hu-berlin.de, olivier.stietel@hu-berlin.de

Robin Vacus
CNRS, IRIF, Paris
robin.vacus@gmail.com

CP5

Combinatorial Selection with Costly Information

We consider a class of optimization problems over stochastic variables where the algorithm can learn information about the value of any variable through a series of costly steps; we model this information acquisition process as a Markov Decision Process (MDP). The algorithm's goal is to minimize the cost of its solution plus the cost of information acquisition, or alternately, maximize the value of its solution minus the cost of information acquisition. Such bandit superprocesses have been studied previously but solutions are known only for fairly restrictive special cases. We develop a framework for approximate optimization of bandit superprocesses that applies to arbitrary acyclic MDPs with a matroid feasibility constraint. Our framework establishes a bound on the optimal cost through a novel cost amortization; it then couples this bound with a notion of local approximation that allows approximate solutions for each component MDP in the superprocess to be composed without loss into a global approximation. We use this framework to obtain approximately optimal solutions for several variants of bandit superprocesses for both maximization and minimization. We obtain new approximations for combinatorial versions of the previously studied Pandora's Box with Optional Inspection and Pandora's Box with Partial Inspection; the less-studied Additive Pandora's Box problem; as well as a new problem that we call the Weighing Scale problem.

Dimitrios Christou
University of Texas at Austin
christou@cs.utexas.edu

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

Amit Harlev, Ziv Scully
Cornell University

ah843@cornell.edu, zivscully@cornell.edu

CP5

Deterministic and Exact Fully-Dynamic Minimum Cut of Superpolylogarithmic Size in Subpolynomial Time

We present an exact fully-dynamic minimum cut algorithm that runs in $n^{o(1)}$ deterministic update time when the minimum cut size is at most $2^{\Theta(\log^{3/4-c} n)}$ for any $c > 0$, improving on the previous algorithm of Jin, Sun, and Thorup (SODA 2024) whose minimum cut size limit is $(\log n)^{o(1)}$. Combined with graph sparsification, we obtain the first $(1+\epsilon)$-approximate fully-dynamic minimum cut algorithm on *weighted* graphs, for any $\epsilon \geq 2^{-\Theta(\log^{3/4-c} n)}$, in $n^{o(1)}$ randomized update time. Our main technical contribution is a deterministic local minimum cut algorithm, which replaces the randomized LocalKCut procedure from El-Hayek, Henzinger, and Li (SODA 2025).

Antoine El-Hayek
Institute of Science and Technology, Austria
antoine.el-hayek@ist.ac.at

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

Jason M. Li
Carnegie Mellon University
jmli@cs.cmu.edu

CP5

A Broader View on Clustering under Cluster-Aware Norm Objectives

We revisit the (f, g)-clustering problem that we introduced in a recent work [SODA'25], and which subsumes fundamental clustering problems such as k-Center, k-Median, Min-Sum of Radii, and Min-Load k-Clustering. This problem assigns each of the k clusters a cost determined by the monotone, symmetric norm f applied to the vector distances in the cluster, and aims at minimizing the norm g applied to the vector of cluster costs. Previously, we focused on certain special cases for which we designed constant-factor approximation algorithms. Our bounds for more general settings left, however, large gaps to the known bounds for the basic problems they capture. In this work, we provide a clearer picture of the approximability of these more general settings. First, we design an $O(\log^2 n)$-approximation algorithm for (f, L_1)-clustering for any f. This improves upon our previous $\widetilde{O}(\sqrt{n})$-approximation. Second, we provide an $O(k)$-approximation for the general (f, g)-clustering problem, which improves upon our previous $\widetilde{O}(\sqrt{kn})$-approximation algorithm and matches the best-known upper bound for Min-Load k-Clustering. We then design an approximation algorithm for (f, g)-clustering that interpolates, up to polylog factors, between the best known bounds for k-Center, k-Median, Min-Sum of Radii, Min-Load k-Clustering, (Top, L_1)-clustering, and (L_∞, g)-clustering based on a newly defined parameter of f and g.

Martin G. Herold
MPI Informatics, SIC
mherold@mpi-inf.mpg.de

Evangelos Kipouridis
Max Planck Institute for Informatics, SIC
kipoujr@hotmail.com

Joachim Spoerhase
University of Liverpool, United Kingdom
joachim.spoerhase@liverpool.ac.uk

CP5

Near-Optimal Min-Sum Multi-Robot Motion Planning in a Planar Polygonal Environment

Let $W \subset \mathbb{R}^2$ be a planar polygonal environment with n vertices, and let $[k] = \{1, \ldots, k\}$ denote k unit-square robots translating in W. Given source and target placements $s_1, t_1, \ldots, s_k, t_k \in W$ for each robot, we wish to compute a collision-free motion plan π, i.e., a coordinated motion for each robot i along a continuous path from s_i to t_i so that robot i does not leave W or collide with any other j. Moreover, we additionally require that π minimizes the sum of the path lengths; this variant is known as min-sum motion planning. For $r > 0$, let $\mathsf{opt}(\mathbf{s}, \mathbf{t}, r)$ denote the cost of a min-sum motion plan for k square robots of radius r each from $\mathbf{s} = (s_1, \ldots, s_k)$ to $\mathbf{t} = (t_1, \ldots, t_k)$. Given a parameter $\epsilon > 0$, we present an algorithm for computing a coordinated motion plan for k unit radius square robots of cost at most $(1+\epsilon)\mathsf{opt}(\mathbf{s}, \mathbf{t}, 1+\epsilon) + \epsilon$, that runs in time $f(k, \epsilon)n^{O(k)}$, where $f(k, \epsilon) = (k/\epsilon)^{O(k^2)}$. Our result is the first polynomial-time bicriteria $(1+\epsilon)$-approximation algorithm for any optimal multi-robot motion planning problem amidst obstacles for a constant value of $k > 2$. The algorithm also works if robots are modeled as congruent disks.

Benjamin Holmgren, Alex Steiger, Pankaj Agarwal
Duke University
ben.holmgren@duke.edu, alexander.steiger@duke.edu, pankaj@cs.duke.edu

CP5

Coloring 3-Colorable Graphs with Low Threshold Rank

We present a new algorithm for finding large independent sets in 3-colorable graphs with small 1-sided threshold rank. Specifically, given an n-vertex 3-colorable graph whose uniform random walk matrix has at most r eigenvalues larger than ε, our algorithm finds a proper 3-coloring on at least $(\frac{1}{2} - O(\varepsilon))n$ vertices in time $n^{O(r/\varepsilon^2)}$. This extends and improves upon the result of Bafna, Hsieh, and Kothari on 1-sided expanders. Furthermore, an independent work by Buhai, Hua, Steurer, and Vri-Kakas shows that it is UG-hard to properly 3-color more than $(\frac{1}{2} + \varepsilon)n$ vertices, thus establishing the tightness of our result.

Jun-Ting Hsieh
Carnegie Mellon University
juntingh@mit.edu

CP5

An Optimal Algorithm for Average Distance in Typical Regular Graphs

We design a deterministic algorithm that, given n points in a *typical* constant degree regular graph, queries $O(n)$ distances to output a constant factor approximation to the average distance among those points, thus answering a question posed in [?]. Our algorithm uses the method of [?] to

construct a sequence of constant degree graphs that are expanders with respect to certain nonpositively curved metric spaces, together with a new rigidity theorem for metric transforms of nonpositively curved metric spaces. The fact that our algorithm works for typical (uniformly random) constant degree regular graphs rather than for all constant degree graphs is unavoidable, thanks to the following impossibility result that we obtain: For every fixed $k \in \mathbb{N}$, the approximation factor of any algorithm for average distance that works for all constant degree graphs and queries $o(n^{1+1/k})$ distances must necessarily be at least $2(k+1)$. This matches the upper bound attained by the algorithm that was designed for general finite metric spaces in [?].

Alexandros Eskenazis
CNRS
Sorbonne Université
alexandros.eskenazis@imj-prg.fr

Manor Mendel
The Open University of Israel
mendelma@gmail.com

Assaf Naor
Princeton University
naor@math.princeton.edu

CP6

Distributed Interactive Proofs for Planarity with Log-Star Communication

We provide new communication-efficient distributed interactive proofs for planarity. The notion of a *distributed interactive proof (DIP)* was introduced by Kol, Oshman, and Saxena (PODC 2018). In a DIP, the *prover* is a single centralized entity whose goal is to prove a certain claim regarding an input graph G. To do so, the prover communicates with a distributed *verifier* that operates concurrently on all n nodes of G. A DIP is measured by the amount of prover-verifier communication it requires. Namely, the goal is to design a DIP with a small number of interaction rounds and a small *proof size*, i.e., a small amount of communication per round. In prior work, Naor, Parter, and Yogev (SODA 2020) presented a 3-round DIP protocol for planarity with a proof size of $O(\log n)$. Later on, Feuilloley et al. (PODC 2020) showed that the same proof size can be accomplished with a non-interactive protocol. In a very recent work by Gil and Parter (GP25), a 5-round protocol with a proof size of $O(\log\log n)$ is presented for *embedded planarity*, which is defined such that an embedding of the graph is given (e.g., each node holds a clockwise ordering of its incident edges) and the goal is to decide if it is a valid planar embedding. In addition, (GP25) presented a protocol with a proof size of $O(\log\log n + \log\Delta)$ for (non-embedded) planarity, where Δ is the maximum degree of the graph.

Yuval Gil
Reykjavik University
yuvalg@ru.is

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

CP6

Algorithmic Improvements to List Decoding of Folded Reed-Solomon Codes

Folded Reed-Solomon (FRS) codes are a well-studied family of codes, known for achieving list decoding capacity. In this work, we give improved deterministic and randomized algorithms for list decoding FRS codes of rate R up to radius $1-R-\varepsilon$. We present a deterministic decoder that runs in near-linear time $\widetilde{O}_\varepsilon(n)$, improving upon the best-known runtime $n^{\Omega(1/\varepsilon)}$ for decoding FRS codes. Prior to our work, no capacity achieving code was known whose deterministic decoding could be done in time $\widetilde{O}_\varepsilon(n)$. We also present a randomized decoder that runs in fully polynomial time $\mathrm{poly}(1/\varepsilon)\cdot\widetilde{O}(n)$, improving the best-known runtime $\exp(1/\varepsilon)\cdot\widetilde{O}(n)$ for decoding FRS codes. Again, prior to our work, no capacity achieving code was known whose decoding time depended polynomially on $1/\varepsilon$. Our results are based on improved pruning procedures for finding the list of codewords inside a constant-dimensional affine subspace.

Mursalin Habib
Rutgers University
mursalin.habib@rutgers.edu

Vikrant Ashvinkumar
Rutgers
va264@cs.rutgers.edu

Shashank Srivastava
DIMACS, Rutgers University
Institute for Advanced Study
shashanks@ias.edu

CP6

Excluding a Line Minor Via Design Matrices and Column Number Bounds for the Circuit Imbalance Measure

For a real matrix $A \in \mathbb{R}^{d\times n}$ with non-collinear columns, we show that $n \le O(d^4\kappa_A)$ where κ_A is the circuit imbalance measure of A. The circuit imbalance measure κ is a real analogue of Δ-modularity for integer matrices, satisfying $\kappa_A \le \Delta_A$ for integer A. The circuit imbalance measure has numerous applications in the context of linear programming (see Ekbatani, Natura and Vegh (2022) for a survey). Our result generalizes the $O(d^4\Delta_A)$ bound of Averkov and Schymura (2023) for integer matrices and provides the first polynomial bound holding for all parameter ranges on real matrices. To derive our result, similar to the strategy of Geelen, Nelson and Walsh (2021) for Δ-modular matrices, we show that real representable matroids induced by κ-bounded matrices are minor closed and exclude a rank 2 uniform matroid on $O(\kappa)$ elements as a minor (also known as a line of length $O(\kappa)$). As our main technical contribution, we show that any simple rank d complex representable matroid which excludes a line of length l has at most $O(d^4 l)$ elements. This complements the tight bound of $(l-3)\binom{d}{2}+d$ for $l \ge 4$, of Geelen, Nelson and Walsh which holds when the rank d is sufficiently large compared to l (at least doubly exponential in l).

Daniel Dadush
Centrum Wiskunde & Informatica
dadush@cwi.nl

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

Rom Pinchasi
Math Dept. Technion
Haifa 32000 Israel
room@math.technion.ac.il

Thomas Rothvoss
University of Washington
rothvoss@uw.edu

Neta Singer
EPFL
neta.singer@epfl.ch

CP6

Generating Pivot Gray Codes for Spanning Trees of Complete Graphs in Constant Amortized Time

We present the first known pivot Gray code for spanning trees of complete graphs, listing all spanning trees such that consecutive trees differ by pivoting a single edge around a vertex. This pivot Gray code thus addresses an open problem posed by Knuth in The Art of Computer Programming, Volumn 4 (Exercise 101, Section 7.2.1.6, [Knuth, 2011]), rated at a difficulty level of 46 out of 50, and imposes stricter conditions than existing revolving-door or edge-exchange Gray codes for spanning trees of complete graphs. Our recursive algorithm generates each spanning tree in constant amortized time using $O(n^2)$ space. In addition, we provide a novel proof of Cayleys formula, n^{n-2}, for the number of spanning trees in a complete graph, derived from our recursive approach. We extend the algorithm to generate edge-exchange Gray codes for general graphs with n vertices, achieving $O(n^2)$ time per tree using $O(n^2)$ space. For specific graph classes, the algorithm can be optimized to generate edge-exchange Gray codes for spanning trees in constant amortized time per tree for complete bipartite graphs, $O(n)$-amortized time per tree for fan graphs, and $O(n)$-amortized time per tree for wheel graphs, all using $O(n^2)$ space.

Dennis Wong, Bowie Liu, Chan-Tong Lam, Sio-Kei Im
Macao Polytechnic University
cwong@uoguelph.ca, bowen.liu@mpu.edu.mo, ctlam@mpu.edu.mo, marcusim@mpu.edu.mo

CP6

Unbounded Error Correcting Codes

Traditional error-correcting codes (ECCs) assume a fixed message length, but many scenarios involve ongoing or indefinite transmissions where the message length is not known in advance. For example, when streaming a video, the user should be able to fix a fraction of errors that occurred before any point in time. We introduce unbounded error-correcting codes (unbounded codes), a natural generalization of ECCs that supports arbitrarily long messages without a predetermined length. An unbounded code with rate R and distance ε ensures that for every sufficiently large k, the message prefix of length Rk can be recovered from the code prefix of length k even if an adversary corrupts up to an ε fraction of the symbols in this code prefix.

Or Zamir
Tel Aviv University
orzamir90@gmail.com

Klim Efremenko
Ben Gurion University
klim@bgu.ac.il

CP6

Language Generation in the Limit: Noise, Loss, and Feedback

Kleinberg and Mullainathan (2024) recently proposed a formal framework called language generation in the limit and showed that given a sequence of example strings from an unknown target language drawn from any countable collection, an algorithm can correctly generate unseen strings from the target language within finite time. Li, Raman, and Tewari (2024) further defined stricter categories of non-uniform and uniform generation, and asked if a finite union of non-uniformly generatable collections is generatable in the limit. We resolve this question in the negative: we give a uniformly generatable collection and a non-uniformly generatable collection whose union is not generatable in the limit. We then use facets of this construction to further our understanding of several variants of language generation. The first two, generation with noise and without samples, were introduced by Raman and Raman (2025) and Li, Raman, and Tewari (2024) respectively. We show the equivalence of these models for uniform and non-uniform generation, and provide a characterization of non-uniform noisy generation. We also show a separation between noisy and non-noisy generation in the limit with just a single noisy string. Finally, we study the framework of generation with feedback, introduced by Charikar and Pabbaraju (2025), where the algorithm is strengthened by allowing it to ask membership queries. We show finite queries add no power, but infinite queries yield a strictly more powerful model.

Yannan Bai, Debmalya Panigrahi, Ian Zhang
Duke University
yannan.bai@duke.edu, debmalya@cs.duke.edu, ian.zhang@duke.edu

CP7

Faster Estimation of the Average Degree of a Graph Using Random Edges and Structural Queries

We revisit the problem of designing sublinear algorithms for estimating the average degree of an n-vertex graph. The standard access model for graphs allows for the following queries: sampling a uniform random vertex, the degree of a vertex, sampling a uniform random neighbor of a vertex, and “pair queries' which determine if a pair of vertices form an edge. In this model, original results on this problem prove that the complexity of getting $(1+\varepsilon)$-multiplicative approximations to the average degree, ignoring ε-dependencies, is $\Theta(\sqrt{n})$. When random edges can be sampled, it is known that the average degree can estimated in $\widetilde{O}(n^{1/3})$ queries, even without pair queries. We give a nearly optimal algorithm in the standard access model with random edge samples. Our algorithm makes $\widetilde{O}(n^{1/4})$ queries exploiting the power of pair queries. We also analyze the “full neighborhood access” model wherein the entire adjacency list of a vertex can be obtained with a single query; this model is relevant in many practical applications. In a weaker version of this model, we give an algorithm that makes $\widetilde{O}(n^{1/5})$ queries. Both these results underscore the power of *structural queries*, such as pair queries and full neighborhood access queries, for estimating the average degree. We give nearly matching lower bounds, ignoring ε-dependencies, for all our results. [Abstract shortened]

Deeparnab Chakrabarty
Dartmouth College
deeparnab@dartmouth.edu

Lorenzo Beretta
IBM
UC Santa Cruz
lorenzo2beretta@gmail.com

C Seshadhri
UC Santa Cruz
sesh@ucsc.edu

CP7

On Sampling Two Spin Models Using the Local Connective Constant

This work establishes novel optimum mixing bounds for the Glauber dynamics on the Hard-core and Ising models. These bounds are expressed in terms of the *local connective constant* of the underlying graph G. This is a notion of effective degree for G, and as such, it allows us to obtain bounds which are inherently less restrictive than those obtained using other graph invariants, e.g., the maximum degree. Our results have some interesting consequences for bounded degree graphs: (A) They include the max-degree bounds as a special case. (B) They improve on the running time of the FPTAS considered in [Sinclair, Srivastava, Stefankonic, Yin: PTRF 2017] for general graphs. (C) They allow us to obtain mixing bounds in terms of the spectral radius of the adjacency matrix and improve on [Hayes: FOCS 2006]. We obtain our results using tools from the theory of high-dimensional expanders and, in particular, the Spectral Independence method [Anari, Liu, Oveis-Gharan: FOCS 2020]. We explore a new direction by utilising the notion of the k-non-backtracking matrix $H_{G,k}$ in our analysis with the Spectral Independence. The results with $H_{G,k}$ are interesting in their own right. The novelties in the analysis amount to relating the maximum singular value (of a sufficiently large power) of $H_{G,k}$ and the spectral radius of the Influence matrix.

Charilaos Efthymiou
Warwick University, United Kingdom
efthymiou@gmail.com

CP7

#cfg and #dnnf Admit Fpras

We provide the first fully polynomial-time randomized approximation scheme for the following two counting problems: 1. Given a Context-Free Grammar G over alphabet Σ, count the number of words of length exactly n generated by G. 2. Given a circuit φ in Decomposable Negation Normal Form (DNNF) over the set of Boolean variables X, compute the number of assignments to X such that φ evaluates to 1. Finding polynomial time algorithms for the aforementioned problems has been a longstanding open problem. Prior work could either only obtain a quasi-polynomial runtime or a polynomial-time randomized approximation scheme for non-deterministic finite automata and non-deterministic tree automata.

Kuldeep S. Meel
Georgia institute of Technology
meel@cs.toronto.edu

Alexis De Colnet
TU Wien
decolnet@ac.tuwien.ac.at

CP7

Does Block Size Matter in Randomized Block Krylov Low-Rank Approximation?

We study the problem of computing a rank-k approximation of a matrix using randomized block Krylov iteration. Prior work has shown that, for block size $b = 1$ or $b = k$, a $(1+\varepsilon)$-factor approximation to the best rank-k approximation can be obtained after $\tilde{O}(k/\sqrt{\varepsilon})$ matrix-vector products with the target matrix. On the other hand, when b is between 1 and k, the best known bound on the number of matrix-vector products scales with $b(k-b)$, which could be as large as $O(k^2)$. Nevertheless, in practice, the performance of block Krylov methods is often optimized by choosing a block size $1 \ll b \ll k$. We address this theory-practice gap by proving that randomized block Krylov iteration produces a $(1+\varepsilon)$-factor approximate rank-k approximation using $\tilde{O}(k/\sqrt{\varepsilon})$ matrix-vector products for *any* block size $1 \le b \le k$. Our analysis relies on new bounds for the minimum singular value of a random block Krylov matrix, which may be of independent interest. Similar bounds are central to recent breakthroughs on faster algorithms for sparse linear systems [Peng & Vempala, SODA 2021; Nie, STOC 2022].

Raphael A. Meyer, Tyler Chen
New York University
ram900@berkeley.edu, tyler.chen@nyu.edu

Ethan Epperly
UC Berkeley
eepperly@berkeley.edu

Christopher Musco
New York University, U.S.
cmusco@nyu.edu

Akash Rao
Washington State University
akashgopinath.rao@wsu.edu

CP7

Optimal Randomized Clustering for Subsets of L_p When $p > 2$

We resolve multiple fundamental open questions about the bi-Lipschitz geometry of subsets of L_p for $2 \le p < \infty$ via a novel multiscale and localization framework. Specifically, we prove that the separation modulus of any n-point subset of L_p is $\Theta_p(\sqrt{\log n})$. If that subset has doubling constant λ, then we obtain the improved bound $O_p(\sqrt{\log \lambda})$, which is new even for the Euclidean space $p = 2$. We also break the longstanding $O(\log n)$ barrier for embedding every n-point subset of L_p into Euclidean space for all $p > 2$, as well as the longstanding $O((\log n)/\log\log n)$ barrier for their Lipschitz extension modulus.

Assaf Naor, Kevin Ren
Princeton University
naor@math.princeton.edu, kr5621@princeton.edu

CP7

Approaching Optimality for Solving Dense Linear

Systems with Low-Rank Structure

We provide new high-accuracy randomized algorithms for solving linear systems and regression problems that are well-conditioned except for k large singular values. For solving such $d \times d$ positive definite system our algorithms succeed whp. and run in time $\tilde{O}(d^2 + k^\omega)$. For solving such regression problems in a matrix $A \in \mathbb{R}^{n \times d}$ our methods succeed whp. and run in time $\tilde{O}(\mathrm{nnz}(A) + d^2 + k^\omega)$ where ω is the matrix multiplication exponent and $\mathrm{nnz}(A)$ is the number of non-zeros in A. Our methods nearly-match a natural complexity limit under dense inputs for these problems and improve upon a trade-off in prior approaches that obtain running times of either $\tilde{O}(d^{2.065} + k^\omega)$ or $\tilde{O}(d^2 + dk^{\omega-1})$ for $d \times d$ systems. Moreover, we show how to obtain these running times even under the weaker assumption that all but k of the singular values have a suitably bounded generalized mean. Consequently, we give the first nearly-linear time algorithm for computing a multiplicative approximation to the nuclear norm of an arbitrary dense matrix. Our algorithms are built on three general recursive preconditioning frameworks, where matrix sketching and low-rank update formulas are carefully tailored to the problems' structure.

Michal Derezinski
University of Michigan
derezin@umich.edu

Aaron Sidford
Stanford University
sidford@stanford.edu

CP8

Deterministic Dynamic Edge Colouring

Given a dynamic graph G with n vertices and m edges subject to insertions and deletions of edges, we show how to maintain a $(1+\varepsilon)\Delta$-edge-colouring of G without the use of randomisation. More specifically, we show a deterministic dynamic algorithm with an amortised update time of $2^{\tilde{O}_{\log \varepsilon^{-1}}(\sqrt{\log n})}$ using $(1+\varepsilon)\Delta$ colours. If $\varepsilon^{-1} \in 2^{O(\log^{0.49} n)}$, then our update time is sub-polynomial in n. While there exists randomised algorithms maintaining colourings with the same number of colours [Bhattacharya, Costa, Panski, Solomon SODA'24, Christiansen STOC'23, Duan, He, Zhang SODA'19] in polylogarithmic and even constant update time, this is the first efficient deterministic algorithm to go below the greedy threshold of $2\Delta - 1$ colours for all input graphs. On the way to our main result, we show how to dynamically maintain a shallow hierarchy of degree-splitters with both recourse and update time in $n^{o(1)}$. We believe that this algorithm might be of independent interest.

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

CP8

Dynamic Hierarchical J-Tree Decomposition and Its Applications

We develop a new algorithmic framework for designing approximation algorithms for cut-based optimization problems on capacitated undirected graphs that undergo edge insertions and deletions. Specifically, our framework dynamically maintains a variant of the hierarchical j-tree decomposition of [Madry FOCS10], achieving a polylogarithmic approximation factor to the graph's cut structure and supporting edge updates in $O(n^\epsilon)$ amortized update time, for any arbitrarily small constant $\epsilon \in (0,1)$. Consequently, we obtain new trade-offs between approximation and update/query time for fundamental cut-based optimization problems in the fully dynamic setting, including all-pairs minimum cuts, sparsest cut, multi-way cut, and multi-cut. For the last three problems, these trade-offs give the first fully-dynamic algorithms achieving polylogarithmic approximation in sub-linear time per operation. The main technical ingredient behind our dynamic hierarchy is a dynamic cut-sparsifier algorithm that can handle vertex splits with low recourse. This is achieved by white-boxing the dynamic cut sparsifier construction of [Abraham et al. FOCS16], based on forest packing, together with new structural insights about the maintenance of these forests under vertex splits. Given the versatility of cut sparsification in both the static and dynamic graph algorithms literature, we believe this construction may be of independent interest.

Zöcklein Gernot
Institute of Science and Technology Austria (ISTA)
gernotzoecklein@gmail.com

Gramoz Goranci
University of Vienna
gernotzoecklein@gmail.com

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

Peter Kiss, Ali Momeni
University of Vienna
peter.kiss@univie.ac.at, ali.momeni@univie.ac.at

CP8

Learning Packing and Covering from Samples

We consider mixed packing and covering problems in the online setting: at each timestep t, the algorithm faces a collection of K choices. Each choice consumes resources, and gives benefits; both resources and benefits are d-dimensional vectors. This general problem captures the question of load-balancing on unrelated machines, where the choices are allocations of jobs to machines. We study these problems in the *Online with a Sample* model, where the algorithm is given a random p-fraction of the instance up-front, and needs to perform well on the entire instance presented in arbitrary order. We give online algorithms in this model for the multiple-choice packing/covering problem (where we are allowed to violate the capacities slightly), and also for the load-balancing problem (where we exceed the optimal load B slightly). Previous works used samples to learn a good "dual' vector, and use it to make online decisions. Instead, we use the samples to run a full-fledged offline algorithm that (implicitly) computes decisions for all online items; when an online item actually arrives, we reconstruct the execution of the algorithm for this item. We show how to implement this for parallel algorithms having a certain local-global structure and are robust under sampling the local data.

Anupam Gupta
New York University, USA
anupam.g@nyu.edu

Marco Molinaro
Microsoft Research (Redmond) and Computer Science Department
mmolinaro@microsoft.com

CP8

Tree Embedding in High Dimensions: Dynamic and Massively Parallel

Tree embedding has been a fundamental method in algorithm design with wide applications. We focus on the efficiency of building tree embedding in various computational settings under high-dimensional Euclidean $\mathbb{R}^d$. We devise a new tree embedding construction framework that operates on an arbitrary metric decomposition with bounded diameter, offering a tradeoff between distortion and the locality of its algorithmic steps. This framework works for general metric spaces and may be of independent interest beyond the Euclidean setting. Using this framework, we obtain a dynamic algorithm that maintains an $O_\epsilon(\log n)$-distortion tree embedding with update time $\tilde{O}(n^\epsilon + d)$ subject to point insertions/deletions, and a massively parallel algorithm that achieves $O_\epsilon(\log n)$-distortion in $O(1)$ rounds and total space $\tilde{O}(n^{1+\epsilon})$ (for constant $\epsilon \in (0, 1)$). These new tree embedding results allow for a wide range of applications. Notably, under a similar performance guarantee as in our tree embedding algorithms, i.e., $\tilde{O}(n^\epsilon + d)$ update time and $O(1)$ rounds, we obtain $O_\epsilon(\log n)$-approximate dynamic and MPC algorithms for k-median and earth-mover distance in $\mathbb{R}^d$.

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

Shaofeng Jiang
Peking University
shaofeng.jiang@pku.edu.cn

Peter Kiss
University of Vienna
peter.kiss@univie.ac.at

Qihao Kong, Yi Qian
Peking University
qh.kong@stu.pku.edu.cn, qianyi@stu.pku.edu.cn

Eva Szilagyi
University of Vienna
eva.szilagyi@univie.ac.at

CP8

History-Independent Load Balancing

We give a (strongly) history-independent two-choice balls-and-bins algorithm on n bins that supports both insertions and deletions on a set of up to m balls, while guaranteeing a maximum load of $m/n + O(1)$ with high probability, and achieving an expected recourse of $O(\log \log(m/n))$ per operation. To the best of our knowledge, this is the first history-independent solution to achieve nontrivial guarantees of any sort for $m/n \geq \omega(1)$ and is the first fully dynamic solution (history independent or not) to achieve $O(1)$ overload with $o(m/n)$ expected recourse.

Rose Silver
Carnegie Mellon University
rosesilv@andrew.cmu.edu

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

William Kuszmaul
CMU
kuszmaul@cmu.edu

Elaine Shi
Carnegie Mellon University
rshi@andrew.cmu.edu

CP8

An Optimal Online Algorithm for Robust Flow Time Scheduling

The problem of minimizing the total flow time on a single machine is one of the few problems for which we can give an optimal online algorithm: just schedule the job with the shortest remaining processing time (SRPT). However, this requires knowledge of the true running time p_j of each job j. Azar, Leonardi, and Touitou recently asked: what if we are given estimates $\widehat{p}_j$ for each job, such that the multiplicative error between p_j and $\widehat{p}_j$ (called the distortion) is at most μ? It is easy to construct examples where no algorithm can be $o(\mu)$ competitive; can we get $O(\mu)$ competitiveness? We show how to achieve this asymptotically optimal result, improving on the previous best result of $O(\mu \log \mu)$. Moreover, we give a very simple algorithm to get this tight bound of $O(\mu)$; the previous ZigZag algorithm was relatively more involved. Our proof is via a dual-fitting argument based on the idea of a reduced instance: we consider an LP relaxation based on knapsack-cover inequalities, and show a solution with a large dual value. Our ideas can also be extended to give alternative dual-fitting arguments for previously analyzed algorithms (like SRPT, ZigZag, and others for minimizing the total flow time, and the Bansal-Dhamdhere algorithm for weighted flow-time minimization).

Zhaozi Wang
New York University
zw4460@nyu.edu

Anupam Gupta
New York University, USA
anupam.g@nyu.edu

Amit Kumar
IIT Delhi
amit.kumar@cse.iitd.ac.in

Debmalya Panigrahi
Duke University
debmalya@cs.duke.edu

CP9

Bounding the Asymptotic Quantum Value of All Multipartite Compiled Non-Local Games

Non-local games are a powerful tool to distinguish between correlations possible in classical and quantum worlds. Kalai et al. (STOC'23) proposed a compiler that converts multipartite non-local games into interactive protocols with a single prover, relying on cryptographic tools to remove the assumption of physical separation of the players. While quantum completeness and classical soundness of the construction have been established for all multipar-

tite games, quantum soundness is known only in the special case of bipartite games. In this paper, we prove that the Kalai et al.'s compiler indeed achieves quantum soundness for all multipartite compiled non-local games, by showing that any correlations that can be generated in the asymptotic case correspond to quantum commuting strategies. Our proof uses techniques from the theory of operator algebras, and relies on a characterisation of sequential operationally no-signalling strategies as quantum commuting operator strategies in the multipartite case, thereby generalising several previous results. On the way, we construct universal C*-algebras of sequential PVMs and prove a new chain rule for Radon-Nikodym derivatives of completely positive maps on C*-algebras which may be of independent interest.

Matilde Baroni
LIP6, Sorbonne University
matilde.baroni@lip6.fr

Dominik Leichtle
University of Edinburgh
dominik.leichtle@ed.ac.uk

Siniša Jankovic
University of Belgrade
jansimke@gmail.com

Ivan Šupic
Université Grenoble Alpes
ivan.supic@univ-grenoble-alpes.fr

CP9

Distributed Quantum Advantage in Locally Checkable Labeling Problems

In this paper, we present the first known example of a locally checkable labeling problem (LCL) that admits asymptotic distributed quantum advantage in the LOCAL model of distributed computing: our problem can be solved in $O(\log n)$ communication rounds in the quantum-LOCAL model, but it requires $\Omega(\log n \cdot \log^{0.99} \log n)$ communication rounds in the classical randomized-LOCAL model. We also show that distributed quantum advantage cannot be arbitrarily large: if an LCL problem can be solved in $T(n)$ rounds in the quantum-LOCAL model, it can also be solved in $\tilde{O}(\sqrt{nT(n)})$ rounds in the classical randomized-LOCAL model. In particular, a problem that is strictly global classically is also almost-global in quantum-LOCAL. Our second result also holds for $T(n)$-dependent probability distributions. As a corollary, if there exists a finitely dependent distribution over valid labelings of some LCL problem Π, then the same problem Π can also be solved in $\tilde{O}(\sqrt{n})$ rounds in the classical randomized-LOCAL and deterministic-LOCAL models. That is, finitely dependent distributions cannot exist for global LCL problems.

Alkida Balliu, Filippo Casagrande, Francesco D'Amore
Gran Sasso Science Institute
alkida.balliu@gssi.it, filippo.casagrande@gssi.it, francesco.damore95@gmail.com

Massimo Equi, Barbara Keller, Henrik Lievonen
Aalto University
massimo.equi@aalto.fi, barbara.keller@aalto.fi, henrik.lievonen@aalto.fi

Dennis Olivetti
Gran Sasso Science Institute
dennis.olivetti@gssi.it

Gustav Schmid
University of Freiburg
schmidg@informatik.uni-freiburg.de

Jukka Suomela
Aalto University
jukka.suomela@aalto.fi

CP9

Quantum Hamiltonian Certification

We formalize and study the Hamiltonian certification problem, a fundamental task in quantum physics, crucial for verifying the accuracy of quantum simulations and quantum-enhanced technologies. Given access to e^{-iHt} for an unknown Hamiltonian H, the goal of the problem is to determine whether H is ε_1-close to or ε_2-far from a target Hamiltonian H_0. This work introduces a direct and efficient framework for Hamiltonian certification, which distinguishes whether an unknown Hamiltonian matches a target specification within given precision bounds. Our approach achieves *optimal* total evolution time $\Theta((\varepsilon_2-\varepsilon_1)^{-1})$ for certification under the normalized Frobenius norm, without prior structural assumptions. This approach also extends to certify Hamiltonians with respect to all Pauli norms and normalized Schatten p-norms for $1 \leq p \leq 2$ in the one-sided error setting ($\varepsilon_1 = 0$), where the optimality is consistently maintained. Notably, the result in Pauli 1-norm suggests a quadratic advantage of our approach over all possible Hamiltonian learning approaches. We also establish matching lower bounds to show the optimality of our approach across all the above settings. We complement our result by showing that the certification problem with respect to normalized Schatten ∞-norm is coQMA-hard, and therefore unlikely to have efficient solutions.

Minbo Gao
Institute of Software, Chinese Academy of Sciences
University of Chinese Academy of Sciences
gaomb@ios.ac.cn

Zhengfeng Ji
Tsinghua University
jizhengfeng@tsinghua.edu.cn

Qisheng Wang
University of Edinburgh
qishengwang1994@gmail.com

Wenjun Yu, Qi Zhao
The University of Hong Kong
wenjunyus@gmail.com, zhaoqcs@hku.hk

CP9

Quantum Advantage Via Solving Multivariate Polynomials

In this work, we propose a new way to (non-interactively) demonstrate quantum advantage by solving the average-case NP search problem of finding a solution to a system of (underdetermined) constant degree multivariate equations over the finite field $\mathbb{F}_2$ drawn from a specified distribution. In particular, for any $d \geq 2$, we design a distribution of degree up to d polynomials $\{p_i(x_1, \dots, x_n)\}_{i\in[m]}$ for $m < n$ over $\mathbb{F}_2$ for which we show that there is a expected polynomial-time quantum algorithm that provably simultaneously solves $\{p_i(x_1, \ldots, x_n) = y_i\}_{i\in[m]}$ for a random vector $(y_1, \ldots, y_m)$. On the other hand, while solutions exist with high probability, we conjecture that for constant $d > 2$, it is classically hard to find one based on a thorough review of existing classical cryptanalysis. Our work thus posits that degree three functions are enough to instantiate the random oracle to obtain non-relativized quantum advantage. Our key technical contribution is a new analysis on the Fourier spectra of distributions induced by a general family of distributions over $\mathbb{F}_2$ multivariate polynomials-those that satisfy 2-wise independence and shift-invariance. This family of distributions includes the distribution of uniform random degree at most d polynomials for any constant $d \geq 2$. Our analysis opens up potentially new directions for quantum cryptanalysis of other multivariate systems.

Pierre Briaud
Simula UiB
pierre@simula.no

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

Riddhi Ghosal
UCLA
ghosalriddhi@ucla.edu

Aayush Jain
CMU
aayushja@andrew.cmu.edu

Paul Lou, Amit Sahai
UCLA
pslou@cs.ucla.edu, sahai@cs.ucla.edu

CP9

Polynomial-Time Classical Simulation of Noisy Quantum Circuits with Naturally Fault-Tolerant Gates

We construct a polynomial-time classical algorithm that samples from the output distribution of noisy geometrically local Clifford circuits with any product-state input and single-qubit measurements in any basis. Our results apply to circuits with nearest-neighbor gates on an $O(1)$-D architecture with depolarizing noise after each gate. Importantly, we assume that the circuit does not contain qubit resets or mid-circuit measurements. This class of circuits includes Clifford-magic circuits and Conjugated-Clifford circuits, which are important candidates for demonstrating quantum advantage using non-universal gates. Additionally, our results can be extended to the case of IQP circuits augmented with CNOT gates, which is another class of non-universal circuits that are relevant to current experiments. Importantly, these results do not require randomness assumptions over the circuit families considered (such as anticoncentration properties) and instead hold for *every* circuit in each class as long as the depth is above a constant threshold. This allows us to rule out the possibility of fault-tolerance in these circuit models. As a key technical step, we prove that interspersed noise causes a decay of long-range entanglement at depths beyond a critical threshold. To prove our results, we merge techniques from percolation theory and Pauli path analysis.

Jon Nelson, Joel Rajakumar
University of Maryland
nelson1@umd.edu, joeldraj4@gmail.com

Dominik Hangleiter
ETH Zeurich
mail@dhangleiter.eu

Michael Gullans
University of Maryland, NIST
mgullans@umd.edu

CP10

Unsplittable Flow Cut Gap in Undirected Graphs

We study multicommodity flows in undirected graphs. Each instance consists of a capacitated supply graph G and a set of source-sink demand pairs forming a demand graph H. An instance is feasible if there exists a flow in G that routes all demands while respecting the edge capacities. A flow is said to be unsplittable if each demand is routed along a single path. We study conditions under which the existence of a feasible (splittable) flow implies the existence of an unsplittable flow that does not significantly violate edge capacities. We focus on instances where feasibility has a natural characterization, the cut-condition: for every cut, the total demand crossing the cut is at most the total capacity of the edges crossing it. It is well known that the cut-condition is (approximately) sufficient for the existence of feasible flows when $G + H$ is planar [Seymour, 81], and when G is series-parallel [Gupta-Newman-Rabinovich-Sinclair, 04]. We prove analogous results in the unsplittable setting and show that if the cut-condition is satisfied, then: 1. If $G + H$ is planar, there exists an unsplittable flow that violates no edge capacity (in additive sense) by more than twice the maximum demand value. This result is almost tight. We will focus on these instances during the talk. 2. If G is series-parallel, there exists an unsplittable flow such that the flow on any edge is at most four times its capacity plus eight times the maximum demand value.

David Aleman Espinosa
University of Waterloo
dalemanespinosa@uwaterloo.ca

Nikhil Kumar
Tata Institute of Fundamental Research
kumar.nikhil@tifr.res.in

Joseph Poremba, Bruce Shepherd
University of British Columbia
jporemba@cs.ubc.ca, fbrucesh@cs.ubc.ca

CP10

Augmenting Packing Dynamic Programs to Handle (Many) Additional Budget Constraints

In a packing problem, we are given a set I of n items, each with a profit, and the goal is to find a subset of items with maximum profit, satisfying some given constraints that depend on the problem at hand. Many approximation algorithms for NP-hard packing problems reduce to an auxiliary packing problem solved via dynamic programming (DP), as in Knapsack, Geometric Knapsack, Independent Set of Rectangles, and Maximum Throughput Scheduling. Often, additional global constraints arise, e.g., for fairness, quotas, or resource limits. Suppose we have M extra cardinality or budget constraints, where each constraint h is of the form $\sum_{i\in S} a_{i,h} \leq b_h$ for every feasible solution $S \subseteq I$. We study how to handle such constraints efficiently. Our main result is a meta-algorithm that, given an exact DP for a packing problem, yields a randomized PTAS, i.e., a

$(1+\varepsilon)$-approximation algorithm for any constant $\varepsilon > 0$ when up to $M = O(\log n/\log\log n)$ cardinality constraints are added. We extend the result to budget constraints if $a_{i,h}$ are small relative to b_h or if we are allowed to use resource augmentation. Also, we show that under Gap-ETH, no polynomial-time $(1+\varepsilon)$-approximation is possible for $M = \omega(\log n)$ cardinality constraints or for $M = \omega(1)$ budget constraints without resource augmentation, making our results almost tight.

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

Fabrizio Grandoni, Antoine Tinguely
IDSIA, USI-SUPSI, Switzerland
fabrizio.grandoni@supsi.ch, antoine.tinguely@supsi.ch

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

CP10

A Better-Than-5/4-Approximation for Two-Edge Connectivity

The 2-Edge-Connected Spanning Subgraph Problem (2ECSS) is a fundamental problem in survivable network design. Given an undirected 2-edge-connected graph, the goal is to find a 2-edge-connected spanning subgraph with the minimum number of edges; a graph is 2-edge-connected if it is connected after the removal of any single edge. 2ECSS is APX-hard and has been extensively studied in the context of approximation algorithms. Very recently, Bosch-Calvo, Garg, Grandoni, Hommelsheim, Jabal Ameli, and Lindermayr showed the currently best-known approximation ratio of $\frac{5}{4}$ [STOC 2025]. This factor is tight for many of their techniques and arguments, and it was not clear whether $\frac{5}{4}$ can be improved. We break this natural barrier and present a $(\frac{5}{4} - \eta)$-approximation algorithm, for some constant $\eta \geq 10^{-6}$. On a high level, we follow the approach of previous works: take a triangle-free 2-edge cover and transform it into a 2-edge-connected spanning subgraph by adding only a few additional edges. For $\frac{5}{4}$-approximations, one can heavily exploit that a 4-cycle in the 2-edge cover can buy one additional edge. This enables simple and nice techniques, but immediately fails for our improved approximation ratio. To overcome this, we design two complementary algorithms that perform well for different scenarios: one for few 4-cycles and one for many 4-cycles.

Felix Hommelsheim
Universidad de Chile
felix.hommelsheim@gmail.com

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

Zhenwei Liu
Zhejiang University
lzw98@zju.edu.cn

CP10

A Better-Than-2 Approximation for the Directed Tree Augmentation Problem

We introduce and study a directed analogue of the weighted Tree Augmentation Problem (WTAP). In the weighted Directed Tree Augmentation Problem (WDTAP), we are given an oriented tree $T = (V, A)$ and a set of directed links $L \subseteq V \times V$ with positive costs. The goal is to select a minimum cost set of links which enters each fundamental dicut of T (cuts with one leaving and no entering tree arc). WDTAP captures the problem of covering a cross-free set family with directed links. It can also be used to solve weighted multi 2-TAP, in which we must cover the edges of an undirected tree at least twice. WDTAP can be approximated to within a factor of 2 using standard techniques. We provide an improved $(1.75+\varepsilon)$-approximation algorithm for WDTAP in the case where the links have bounded costs, a setting that has received significant attention for WTAP. To obtain this result, we discover a class of instances, called "willows', for which the natural set covering LP is an integral formulation. We further introduce the notion of "visibly k-wide' instances which can be solved exactly using dynamic programming. Finally, we show how to leverage these tractable cases to obtain an improved approximation ratio via an elaborate structural analysis of the tree.

Meike Neuwohner
University of Bonn
work@meike-neuwohner.de

Olha Silina
Carnegie Mellon University
osilina@andrew.cmu.edu

Michael Zlatin
Pomona College
michael.zlatin@pomona.edu

CP10

Sparse Navigable Graphs for Nearest Neighbor Search: Algorithms and Hardness

We study approximation algorithms for constructing sparse α-navigable graphs [IX23, DGM24], an important principle behind recent advances in graph-based nearest neighbor search. Given an n-point dataset P with an associated metric d and a parameter $\alpha \geq 1$, the goal is to efficiently build the sparsest graph $G = (P, E)$ that is α-navigable: for every distinct $s, t \in P$, there is an edge $(s, u) \in E$ with $d(u, t) < d(s, t)/\alpha$. We begin with a negative result: the slow-preprocessing version of DiskANN (analyzed in [IX23]) can yield solutions whose sparsity is $\widetilde{\Omega}(n)$ times larger than optimal. We then show a tight approximation-preserving equivalence between the Sparsest Navigable Graph problem and the classic Set Cover problem, obtaining an $O(n^3)$-time $(\ln n+1)$-approximation algorithm, as well as establishing NP-hardness of achieving an $o(\ln n)$-approximation. Building on this equivalence, we develop faster $O(\ln n)$-approximation algorithms. The first runs in $\widetilde{O}(n \cdot OPT)$ time and is therefore much faster when the optimal solution is sparse. The second, based on fast matrix multiplication, is a bicriteria algorithm that computes an $O(\ln n)$-approximation to the sparsest 2α-navigable graph, running in $\widetilde{O}(n^\omega)$ time. Finally, we show a query complexity lower bound, showing that any $o(n)$-approximation requires examining $\Omega(n^2)$ distances.

Ashwin Padaki
Columbia University

apadaki@seas.upenn.edu

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

Erik Waingarten
Simons Institute
ewaingar@seas.upenn.edu

CP10

Efficiently Constructing Sparse Navigable Graphs

Graph-based nearest neighbor search methods have seen a surge of popularity in recent years. Central to these methods is the task of constructing a *sparse navigable search graph* for a given dataset. Unfortunately, doing so is computationally expensive, so heuristics are universally used in practice. In this work, we initiate the study of fast algorithms with provable guarantees for search graph construction. For a dataset with n points, constructing an optimally sparse navigable graph can be framed as n minimum set cover instances. This yields a naive $O(n^3)$ time greedy algorithm that returns an $O(\log n)$-approximate sparsest navigable graph. We improve significantly on this, taking advantage of correlation between set cover instances and introducing problem-specific pre-processing techniques to obtain an $\tilde{O}(n^2)$ time algorithm for constructing an $O(\log n)$-approximate sparsest navigable graph under any distance function. The runtime of our method is optimal up to logarithmic factors under the Strong Exponential Time Hypothesis. Moreover, we prove that obtaining better than an $O(\log n)$-approximation is NP-hard. Finally, we show that our approach can also beat cubic time for the closely related and practically important problems of constructing α-shortcut reachable and τ-monotonic graphs, for which we obtain $\tilde{O}(n^{2.5})$ time or better algorithms.

Alex Conway
Cornell University
ajc473@cornell.edu

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

Martin Farach-Colton
New York University, U.S.
mlf9579@nyu.edu

Rob Johnson
VMware Research
rob.johnson@broadcom.com

Ben Landrum
Cornell Tech
btl46@cornell.edu

Christopher Musco, Yarin Shechter, Torsten Suel
New York University, U.S.
cmusco@nyu.edu, ys6378@nyu.edu, torsten.suel@nyu.edu

Richard Wen
University of Maryland, College Park
rwen1@umd.edu

CP11

Approximately Dominating Sets in Elections

Condorcet's paradox is a fundamental result in social choice theory which states that there exist elections in which, no matter which candidate wins, a majority of voters prefer a different candidate. In fact, even if we can select any k winners, there still may exist another candidate that would beat each of the winners in a majority vote. That is, elections may require arbitrarily large dominating sets. We show that approximately dominating sets of constant size always exist. In particular, for every $\varepsilon > 0$, every election (irrespective of the number of voters or candidates) can select $O(\frac{1}{\varepsilon^2})$ winners such that no other candidate beats each of the winners by a margin of more than ε fraction of voters. Our proof uses a simple probabilistic construction using samples from a maximal lottery, a well-studied distribution over candidates derived from the Nash equilibrium of a two-player game. In stark contrast to general approximate equilibria, which may require support logarithmic in the number of pure strategies, we show that maximal lotteries can be approximated with *constant* support size. These approximate maximal lotteries may be of independent interest.

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

Prasanna Ramakrishnan
Stanford University
pras1712@stanford.edu

Kangning Wang
Rutgers University
kn.w@rutgers.edu

CP11

Compatibility of Fairness and Nash Welfare under Subadditive Valuations

We establish a compatibility between fairness and efficiency, captured via Nash Social Welfare (NSW), under the broad class of subadditive valuations.. We prove that there always exists a *partial* allocation that is envy-free up to the removal of any good (EFx) and has NSW at least half of the optimal; here, optimality is considered across all allocations, fair or otherwise. We also prove, for subadditive valuations, the universal existence of complete allocations that are envy-free up to one good (EF1) and also achieve a factor 1/2 approximation to the optimal NSW. In addition, we develop a polynomial-time algorithm which, given an arbitrary allocation $\widetilde{A}$ as input, returns an EF1 allocation with NSW at least $\frac{1}{e^{2/e}} \approx \frac{1}{2.08}$ times that of $\widetilde{A}$. Therefore, our results imply that the EF1 criterion can be attained simultaneously with a constant-factor approximation to optimal NSW in polynomial time (with demand queries), for subadditive valuations. The previously best-known approximation factor for optimal NSW, under EF1 and among n agents, was $O(n)$ – we improve this bound to $O(1)$. It is known that EF1 and exact Pareto efficiency (PO) are incompatible with subadditive valuations. Complementary to this negative result, the current work shows that we regain compatibility by considering a factor 1/2 approximation: EF1 can be achieved in conjunction with

$\frac{1}{2}$-PO under subadditive valuations.

Siddharth Barman
Indian Institute of Science
barman@iisc.ac.in

Mashbat Suzuki
UNSW Sydney
mashbat.suzuki@unsw.edu.au

CP11

Likelihood of the Existence of Average Justified Representation

We study the approval-based multi-winner election problem where n voters jointly decide a committee of k winners from m candidates. We focus on the axiom *average justified representation* (AJR) proposed by Fernández, Elkind, Lackner, García, Arias-Fisteus, Basanta-Val, and Skowron (2017). AJR postulates that every group of voters with a common preference should be sufficiently represented in that their average satisfaction should be no less than their Hare quota. Formally, for every group of $\lceil \ell \cdot \frac{n}{k} \rceil$ voters with ℓ common approved candidates, the average number of approved winners for this group should be at least ℓ. It is well-known that a winning committee satisfying AJR is not guaranteed to exist for all multi-winner election instances. In this paper, we study the likelihood of the existence of AJR under the Erdős–Rényi model. We consider the Erdős–Rényi model parameterized by $p \in [0, 1]$ that samples multi-winner election instances from the distribution where each voter approves each candidate with probability p (and the events that voters approve candidates are independent), and we provide a clean and complete characterization of the existence of AJR committees in the case where m is a constant and n tends to infinity.

Qishen Han
Rutgers University
hnickc2017@gmail.com

Biaoshuai Tao
Shanghai Jiao Tong University
bstao@sjtu.edu.cn

Lirong Xia, Chengkai Zhang
Rutgers University
xialirong@gmail.com, chengkai.zhang@rutgers.edu

Houyu Zhou
UNSW Sydney
houyu.zhou@unsw.edu.au

CP11

Optimal Type-Dependent Liquid Welfare Guarantees for Autobidding Agents with Budgets

Online advertising systems have recently transitioned to autobidding, allowing advertisers to delegate bidding decisions to automated agents. Each advertiser directs their agent to optimize an objective function subject to return-on-investment (ROI) and budget constraints. Given their practical relevance, this shift has spurred a surge of research on the liquid welfare price of anarchy (POA) of fundamental auction formats under autobidding, most notably simultaneous first-price auctions (FPA). One of the main challenges is to understand the efficiency of FPA in the presence of heterogeneous agent types. We introduce a type-dependent smoothness framework that enables a unified analysis of the POA in such complex autobidding environments. In our approach, we derive type-dependent smoothness parameters which we carefully balance to obtain POA bounds. This balancing gives rise to a POA-revealing mathematical program, which we use to determine tight bounds on the POA of coarse correlated equilibria (CCE). Additionally, we develop a reduction technique that transforms budget-constrained agents into budget-unconstrained ones and obtain tight bounds on the POA of CCE in the general hybrid agent model with both ROI and budget constraints. Our bounds uncover an intriguing threshold phenomenon showing that the POA depends intricately on the smallest and largest agent types. We also extend our study to FPAs with reserve prices, which can be interpreted as predictions of agents' values.

Riccardo Colini-Baldeschi
Meta
rickuz@meta.com

Sophie Klumper, Twan Kroll
CWI and University of Amsterdam
s.j.klumper@cwi.nl, t.kroll@cwi.nl

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

Guido Schäfer
CWI and University of Amsterdam
g.schaefer@cwi.nl

Artem Tsikiridis
Technical University of Munich
artem.tsikiridis@gmail.com

CP11

Contextual Search in Principal-Agent Games: The Curse of Degeneracy

In this work, we introduce and study contextual search in general principal-agent games, where a principal repeatedly interacts with agents by offering contracts based on contextual information and historical feedback, without knowing the agents' true costs or rewards. Our model generalizes classical contextual pricing by accommodating richer agent action spaces. Over T rounds with d-dimensional contexts, we establish an asymptotically tight exponential $T^{1-\Theta(1/d)}$ bound in terms of the pessimistic Stackelberg regret, benchmarked against the best utility for the principal that is consistent with the observed feedback. We also establish a lower bound of $\Omega(T^{\frac{1}{2}-\frac{1}{2d}})$ on the classic Stackelberg regret for principal-agent games, demonstrating a surprising double-exponential hardness separation from the contextual pricing problem (a.k.a, principal-agent game with two actions), which is known to admit a near-optimal $O(d \log \log T)$ regret bound. In particular, this double-exponential hardness separation occurs even in the special case with three actions and two-dimensional context. We identify that this significant increase in learning difficulty arises from a structural phenomenon that we call contextual action degeneracy, where adversarially chosen contexts can make some actions strictly dominated (and unincentivizable), blocking the principal's ability to explore or learn about them, and fundamentally limiting learning progress.

Yiding Feng

Hong Kong University of Science and Technology
ydfeng@ust.hk

Mengfan Ma
Central China Normal University
mengfanma1@gmail.com

Bo Peng
Shanghai University of Finance and Economics
ahqspb@163.sufe.edu.cn

Zongqi Wan
Great Bay University
zongqiwan98@gmail.com

CP12

The Complexity of Dynamic Lz77 Is $\tilde{\Theta}(n^{2/3})$

The Lempel-Ziv 77 (LZ77) factorization is a fundamental compression scheme widely used in text processing and data compression. In this work, we investigate the time complexity of maintaining the LZ77 factorization of a dynamic string. By establishing matching upper and lower bounds, we fully characterize the complexity of this problem. We present an algorithm that efficiently maintains the LZ77 factorization of a string S undergoing edit operations, including character substitutions, insertions, and deletions. Our data structure can be constructed in $\tilde{O}(n)$ time for an initial string of length n and supports updates in $\tilde{O}(n^{2/3})$ time, where n is the current length of S. Additionally, we prove that no algorithm can achieve an update time of $O(n^{2/3-e})$ unless the Strong Exponential Time Hypothesis fails. This lower bound holds even in the restricted setting where only substitutions are allowed and only the length of the LZ77 factorization is maintained.

Itai Boneh
University of Wroclaw
itai.boneh@cs.uni.wroc.pl

Shay Golan
Ariel University
shayg@ariel.ac.il

Matan Kraus
Bar Ilan University
matan3@gmail.com

CP12

Optimal Random Access and Conditional Lower Bounds for 2D Compressed Strings

Compressed indexing is a powerful technique that enables efficient querying over data stored in compressed form. While extensive progress has been made for one-dimensional strings, many real-world datasets (such as images, maps, and adjacency matrices) are inherently two-dimensional and highly compressible. Unfortunately, naively applying 1D techniques to 2D data leads to suboptimal results, as fundamental structural repetition is lost during linearization. We present three main contributions that advance the theory of compressed indexing for 2D strings: (1) We design the first data structure that supports optimal-time random access to a 2D string compressed by a 2D grammar. Specifically, for a 2D string $T \in \Sigma^{r\times c}$ compressed by a 2D grammar G and any constant $\epsilon > 0$, we achieve $O(\log n/\log\log n)$ query time and $O(|G|\log^{2+\epsilon} n)$ space, where $n = \max(r, c)$. (2) We prove conditional lower bounds for pattern matching over 2D-grammar compressed strings. Assuming the Orthogonal Vectors Conjecture, no algorithm can solve this problem in time $O(|G|^{2-\epsilon}\cdot|P|^{O(1)})$ for any $\epsilon > 0$. (3) We show that several fundamental 2D queries, such as the 2D longest common extension, rectangle sum, and equality, cannot be supported efficiently under hardness assumptions for rank and symbol occurrence queries on 1D grammar-compressed strings.

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

CP12

Tight Lower Bounds for Central String Queries in Compressed Space

In this work, we study the limits of compressed data structures, i.e., structures that support various queries on an input text $T \in \Sigma^n$ using space proportional to the size of T in compressed form. Nearly all fundamental queries can currently be efficiently supported in $O(\delta(T)\log^{O(1)} n)$ space, where $\delta(T)$ is the substring complexity, a strong compressibility measure that lower-bounds the optimal space to represent the text [Kociumaka, Navarro, Prezza, IEEE Trans. Inf. Theory 2023]. However, optimal query time has been characterized only for random access. We address this gap by developing tight lower bounds for nearly all other fundamental queries: (1) We prove that suffix array (SA), inverse suffix array (SA^{-1}), longest common prefix (LCP) array, and longest common extension (LCE) queries all require $\Omega(\log n/\log\log n)$ time within $O(\delta(T)\log^{O(1)} n)$ space, matching known upper bounds. (2) We further show that other common queries, currently supported in $O(\log\log n)$ time and $O(\delta(T)\log^{O(1)} n)$ space, including the Burrows-Wheeler Transform (BWT), permuted longest common prefix (PLCP) array, Last-to-First (LF), inverse LF, lexicographic predecessor (Φ), and inverse Φ queries, all require $\Omega(\log\log n)$ time, yielding another set of tight bounds. Our lower bounds hold even for texts over a binary alphabet.

Dominik Kempa
Stony Brook University
kempa@cs.stonybrook.edu

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

CP12

Explaining the Inherent Tradeoffs for Suffix Array Functionality: Equivalences Between String Problems and Prefix Range Queries

We study the fundamental question of how efficiently suffix array entries can be accessed when the array cannot be stored explicitly. The suffix array $SA_T[1..n]$ of a text T of length n encodes the lexicographic order of its suffixes. Previous work established one-way reductions showing how suffix array queries can be answered using, for example, rank queries on the Burrows-Wheeler Transform. More recently, a new class of prefix queries was introduced, together with reductions that, among others, transform a simple tradeoff for prefix-select queries into a suffix array tradeoff matching state-of-the-art space and query-time bounds, while achieving sublinear construction time. For binary texts, the resulting data structure achieves space

$O(n)$ bits, preprocessing time $O(n/\sqrt{\log n})$, preprocessing space of $O(n)$ bits, and query time $O(\log^{\epsilon} n)$ for any constant $\epsilon > 0$. However, whether these bounds could be improved using different techniques has remained open. We present the first bidirectional reduction showing that suffix array queries are, up to an additive $O(\log\log n)$ term in query time, equivalent to prefix-select queries in all parameters. Moreover, we prove analogous equivalences for inverse suffix array queries, pattern ranking, lexicographic range, and SA-interval queries, identifying six core problem pairs that connect string and prefix query models.

Dominik Kempa
Stony Brook University
kempa@cs.stonybrook.edu

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

CP12

Space-Efficient k-Mismatch Text Indexes

Cole, Gottlieb, and Lewenstein (STOC 2004) proposed k-errata trees, a family of text indexes supporting approximate pattern matching queries of several types. In particular, they yield a solution to k-mismatch queries, where we are to compute all substrings of the text that match the pattern with up to k mismatches. The resulting k-mismatch index for a length-n string text uses $O(n \log^k n)$ space and answers a query for a length-m string pattern in $O(\log^k n \log\log n + m + \mathsf{occ})$ time, where occ is the number of approximate occurrences. From the perspective of time, k-errata trees appear very well optimized: even though they have been adapted to various settings, the original time-space trade-off for k-mismatch indexing has not been improved in general. We present the first such improvement, a k-mismatch index with $O(n \log^{k-1} n)$ space and the same query time as k-errata trees. Previously, due to a result of Chan, Lam, Sung, Tam, and Wong (Algorithmica 2010), such an $O(n \log^{k-1} n)$-size index has been known only for texts of alphabets of $O(1)$ size. In this setting, however, we obtain an index of size only $O(n \log^{k-2+\varepsilon+\frac{2}{k+2-(k \bmod 2)}} n) \subseteq O(n \log^{k-1.5+\varepsilon} n)$ for $2 \le k \le O(1)$ and any constant $\varepsilon > 0$. Along the way, we develop improved indexes for the practically relevant special case of short patterns.

Jakub Radoszewski
University of Warsaw
jrad@mimuw.edu.pl

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

CP13

The Parameterised Complexity of Counting Small Sub-Hypergraphs

Subgraph counting is a fundamental problem whose computational complexity is well understood. Quite surprisingly, the hypergraph version of subgraph counting has been almost ignored. In this work, we address this gap by investigating the most basic sub-hypergraph counting problem: given two hypergraphs H and G, compute the number of sub-hypergraphs of G isomorphic to H. Formally, for a family C of hypergraphs, let #Sub(C) be the restriction of the problem to H in C; the induced variant #IndSub(C) is defined analogously. Our main contribution is a complete classification of the fixed-parameter tractability of these problems. Assuming the Exponential Time Hypothesis, we prove that #Sub(C) is fixed-parameter tractable if and only if C has bounded fractional co-independent edge-cover number, a novel graph parameter introduced in this work, and that #IndSub(C) is fixed-parameter tractable if and only if C has bounded fractional edge-cover number. Both results subsume pre-existing results for graphs as special cases. We also show that the fixed-parameter tractable cases of #Sub(C) and #IndSub(C) are unlikely to be in polynomial time, unless respectively #P=P and Graph Isomorphism is contained in P. This shows a separation with the special case of graphs, where the fixed-parameter tractable cases are known to actually be in polynomial time.

Marc Roth
Merton College
Oxford University, United Kingdom
m.roth@qmul.ac.uk

Marco Bressan
University of Milan
marco.bressan@unimi.it

Julian Brinkmann
Goethe University Frankfurt
j.brinkmann@em.uni-frankfurt.de

Holger Dell
Basic Algorithms Research and IT University of Copenhagen
Goethe University Frankfurt
hold@itu.dk

Philip Wellnitz
National Institute of Informatics
Tokyo
wellnitz@nii.ac.jp

CP13

A Logic-Based Algorithmic Meta-Theorem for Treedepth: Single Exponential Fpt Time and Polynomial Space

For a graph G, the parameter treedepth measures the minimum depth among all forests F, called elimination forests, such that G is a subgraph of the ancestor-descendant closure of F. We introduce the neighborhood operator logic with acyclicity, connectivity and clique constraints ($\mathrm{NEO}_2[\mathrm{FRec}] + \mathrm{ACK}$), that captures all NP-hard problems—like Independent Set or Hamiltonian Cycle—that are known to be tractable in time $2^{O(td)} n^{O(1)}$ and space $n^{O(1)}$ on n-vertex graphs provided with elimination forests of depth td. We provide a model checking algorithm for our logic with such complexity that unifies and extends these results. For the fragment of the above logic without acyclicity and connectivity constraints, we get a strengthening of this result, where the space complexity is reduced to $O(td \log(n))$. With a similar mechanism as the distance neighborhood logic [Bergougnoux, Dreier and Jaffke, SODA 2023], our logic is an extension of the fully-existential MSO_2 with predicates for (1) querying generalizations of the neighborhoods of vertex sets, (2) verifying the connectivity and acyclicity of vertex and edge sets, and (3) verifying that a vertex set induces a clique. Our logic captures CNF-SAT via the incidence graphs associated to

CNF formulas and several modulo counting problems like Odd Dominating Set for which the existence of such algorithms was previously unknown.

Benjamin Bergougnoux
LIS, Aix-Marseille Université, France
benjamin.bergougnoux@lis-lab.fr

Vera Chekan
Humboldt-Universität zu Berlin
vera.chekan@hu-berlin.de

Giannos Stamoulis
Université Paris Cité, CNRS, IRIF
stamoulis@irif.fr

CP13

Finding Sparse Induced Subgraphs on Graphs of Bounded Induced Matching Treewidth

The induced matching width of a tree decomposition of a graph G is the cardinality of a largest induced matching M of G, such that there exists a bag that intersects every edge in M. The induced matching treewidth of G, denoted by tree-$\mu(G)$, is the minimum induced matching width of a tree decomposition of G. The parameter tree-μ was introduced by Yolov [SODA '18], who showed that, for example, Maximum-Weight Independent Set can be solved in polynomial-time on graphs of bounded tree-μ. Lima, Milanic, Muric, Okrasa, Rzazewski, and torgel [ESA '24] conjectured that this algorithm can be generalized to a meta-problem called Maximum-Weight Induced Subgraph of Bounded Treewidth, where we are given a vertex-weighted graph G, an integer w, and a CMSO_2-sentence Φ, and are asked to find a maximum-weight set $X \subseteq V(G)$ so that $G[X]$ has treewidth at most w and satisfies Φ. They proved the conjecture for some special cases, such as for the problem Maximum-Weight Induced Forest. In this paper, we prove the general case of the conjecture. In particular, we show that Maximum-Weight Induced Subgraph of Bounded Treewidth is polynomial-time solvable when tree-$\mu(G)$, w, and $|\Phi|$ are bounded. The running time of our algorithm for n-vertex graphs G with tree-$\mu(G) \le k$ is $f(k, w, |\Phi|) \cdot n^{O(kw^2)}$ for a computable function f.

Hans Bodlaender
Utrecht University
h.l.bodlaender@uu.nl

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

Tuukka Korhonen
University of Copenhagen
tuko@di.ku.dk

CP13

Augmenting to 4-vertex Connectivity is Fixed-parameter Tractable

We present fixed-parameter algorithms (FPT algorithms) for the λ-vertex connectivity augmentation (λ-VCA) problem for all values of $\lambda \le 4$; that is, we give an algorithm that given a graph G, a set L of non-edges, and an integer k, determines in time $k^{O(k)} \cdot n^c$ (for some constant c independent of k) whether G can be made λ-vertex connected by adding at most k elements from L. A key technical contribution of our work is a new structure theorem characterizing graphs that can be made λ-vertex connected (for any $\lambda \le 4$) with at most k additional edges. Our central algorithmic result is the first FPT algorithm for 4-VCA parameterized by the standard parameter, the solution size k. Notably, our result also yields the first FPT algorithms parameterized by the solution size, even for λ'-VCA ($\lambda' \le 3$) without any connectivity assumptions on the input graph G. Previously, the problem was known to be FPT parameterized by the solution size only in restricted cases: for $\lambda = 2$ when G is connected (Guo and Uhlmann, Networks 2010; Marx and Végh, TALG 2015) and for $\lambda = 3$ when G is 2-connected (Nutov, ESA 2024).

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

Johannes Carmesin
TU Bergakademie Freiberg
johannes.carmesin@gmail.com

Jan Kurkofka, Jan Kurkofka
TU Freiberg
jan.kurkofka@math.tu-freiberg.de, jan.kurkofka@math.tu-freiberg.de

CP13

K-Sum Hardness Implies Treewidth-Seth

We show that if k-SUM is hard, in the sense that the standard algorithm is essentially optimal, then a variant of the SETH called the Primal Treewidth SETH is true. Formally: if there is an $\varepsilon > 0$ and an algorithm which solves SAT in time $(2-\varepsilon)^{tw}|\phi|^{O(1)}$, where tw is the width of a given tree decomposition of the primal graph of the input, then there exists a randomized algorithm which solves k-SUM in time $n^{(1-\delta)\frac{k}{2}}$ for some δ and all sufficiently large k. We also establish an analogous result for the k-XOR problem, where integer addition is replaced by component-wise addition modulo 2. As an application of our reduction we are able to revisit tight lower bounds on the complexity of several fundamental problems parameterized by treewidth (Independent Set, Max Cut, k-Coloring). Our results imply that these bounds, which were initially shown under the SETH, also hold if one assumes the k-SUM or k-XOR Hypotheses, arguably increasing our confidence in their validity.

Michael Lampis
Université Paris-Dauphine
michail.lampis@dauphine.fr

CP13

Circuits and Backdoors: Five Shades of the Seth

The Strong Exponential Time Hypothesis (SETH) is a standard assumption in (fine-grained) parameterized complexity and many tight lower bounds are based on it. We consider a number of reasonable weakenings of the SETH, with sources from (i) circuit complexity (ii) backdoors for SAT-solving (iii) graph width parameters and (iv) weighted satisfiability problems. Our goal is to arrive at formulations which are simultaneously more plausible as hypotheses, but also capture interesting and robust notions of complexity. Using several tools from classical complexity theory we are able to consolidate these numerous hypotheses into a hierarchy of five main equivalence classes of increasing solidity. This framework serves as a step towards structurally classifying a variety of SETH-based lower bounds into intermediate equivalence classes.

Michael Lampis
Université Paris-Dauphine
michail.lampis@dauphine.fr

CP14

Prophet Inequality from Samples: Is the More the Merrier?

We study a variant of the single-choice prophet inequality problem where the decision-maker does not know the underlying distributions and has only access to a set of samples from the distributions. Rubinstein et al. [2020] showed that the optimal competitive ratio of $\frac{1}{2}$ can surprisingly be obtained by observing a set of n samples, one from each of the distributions. In this paper, we prove that this competitive ratio of $\frac{1}{2}$ becomes unattainable when the decision-maker is provided with a set of more samples (for sufficiently many samples). We then examine the natural class of ordinal static threshold algorithms, where the algorithm selects the i-th highest ranked sample, sets this sample as a static threshold, and then chooses the first value that exceeds this threshold. We show that the best possible algorithm within this class achieves a competitive ratio of $0.433 - o(1)$ (where the $o(1)$ is an expression that decreases as the number of samples increases), for which we provide a matching upper bound of 0.433. Along the way, we utilize the tools developed in the paper and provide an alternative proof of the main result of Rubinstein et al. [2020].

Tomer Ezra
Harvard University
tomer@cmsa.fas.harvard.edu

CP14

Hallucinating Flows for Optimal Mechanisms

Myersons seminal characterization of the revenue-optimal auction for a single item remains a cornerstone of mechanism design. However, generalizing this framework to multi-item settings has proven exceptionally challenging. Even under restrictive assumptions, closed-form characterizations of optimal mechanisms are rare. In this work, we build upon the bi-valued setting introduced by Yao (EC 2017), where each item's value has support 2 and lies in $\{a, b\}$. Yao's result provides the only known closed-form optimal mechanism for multiple agents. We extend this line of work along three natural axes, establishing the first closed-form optimal mechanisms in each of the following settings: (i) n i.i.d. agents and m i.i.d. items (ii) n non-i.i.d. agents and two i.i.d. items and (iii) n i.i.d. agents and two non-i.i.d. items. Our results lie at the limit of what is considered possible, since even with a single agent and m bi-valued non-i.i.d. items, finding the optimal mechanism is $\#P$-Hard. We finally generalize the discrete analog of a result from Daskalakis et al. (Econometrica 2017), showing that for a single agent with m items drawn from arbitrary (non-identical) discrete distributions, grand bundling is optimal when all item values are sufficiently large. We further show that for any continuous product distribution, grand bundling achieves $\mathrm{OPT}-\epsilon$ revenue for large enough values.

Marios Mertzanidis, Athina Terzoglou
Purdue University
mmertzan@purdue.edu, aterzogl@purdue.edu

CP14

The Communication Complexity of Combinatorial Auctions with Additional Succinct Bidders

We study the communication complexity of welfare maximization in combinatorial auctions with bidders from either a standard valuation class (which require exponential communication to explicitly state, e.g., subadditive or XOS), *or* arbitrary succinct valuations (which can be fully described in poly communication, such as single-minded). Although succinct valuations can be efficiently communicated, we show that additional succinct bidders have a nontrivial impact on communication complexity of combinatorial auctions. Specifically, let n be the number of subadditive/XOS bidders. We show that for $\mathsf{SA} \cup \mathsf{Succ}$ (the union of subadditive and succinct valuations): - There is a poly communication 3-approximation algorithm. - As $n \to \infty$, there is a matching 3-hardness of approximation, which is larger than the optimal approximation ratio of 2 for SA. - For all $n \geq 3$, there is a separation between the optimal approximation ratios for $\mathsf{SA} \cup \mathsf{Succ}$ and SA. Similarly, we show that for $\mathsf{XOS} \cup \mathsf{Succ}$: - There is a poly communication 2-approximation algorithm. - As $n \to \infty$, there is a matching 2-hardness of approximation, which is larger than the optimal approximation ratio of $e/(e-1)$ for XOS. - For all $n \geq 2$, there is a separation between the optimal approximation ratios for $\mathsf{XOS} \cup \mathsf{Succ}$ and XOS.

Frederick V. Qiu, S. Matthew Weinberg, Qianfan Zhang
Princeton University
frdrckq@gmail.com, smweinberg@princeton.edu, qianfan@princeton.edu

CP14

Contract Design Beyond Hidden-Actions

In classical contract design, a principal delegates the execution of a costly task to an agent. To complete the task, the agent chooses an action from a set of actions, each associated with a cost and a success probability to accomplish the task. To incentivize the agent to exert effort, the principal can commit to a contract, which is the amount of payment based on the task's success but not on the hidden-action chosen by the agent. In this work, we study contract design under binary outcomes relaxing the hidden-action assumption. We introduce new models where the principal is allowed to inspect subsets of actions at some cost that depends on the inspected subset. If the principal discovers that the agent did not select the agreed-upon action through the inspection, the principal can withhold payment. This relaxation introduces a broader strategy space for the principal, who now faces a tradeoff between positive incentives (increasing payment) and negative incentives (increasing inspection). We devise algorithms for finding the best deterministic and randomized incentive-compatible inspection schemes for various assumptions on the inspection cost function. In particular, we show tractability in the case of submodular inspection costs. We complement our results by showing that it is impossible to efficiently find the optimal randomized inspection scheme for the more general case of XOS inspection costs, and that there is no PTAS for subadditive

inspection costs.

Tomer Ezra
Harvard University
tomer@cmsa.fas.harvard.edu

Stefano Leonardi
Sapienza Università di Roma, Italy
leonardi@diag.uniroma1.it

Matteo Russo
Sapienza, University of Rome, Italy
mrusso@diag.uniroma1.it

CP14

The Power of Matching for Online Fractional Hedonic Games

We study coalition formation in the framework of fractional hedonic games (FHGs). The objective is to maximize social welfare in an online model where agents arrive one by one and must be assigned to coalitions immediately and irrevocably. A recurrent theme in online coalition formation is that online matching algorithms, where coalitions are restricted to size at most 2, yield good competitive ratios. For example, computing maximal matchings achieves the optimal competitive ratio for general online FHGs. However, this ratio is bounded only if agents' valuations are themselves bounded. We identify optimal algorithms with constant competitive ratios in two related settings, independent of the range of agent valuations. First, under random agent arrival, we present an asymptotically optimal $(\frac{1}{3} - \frac{1}{n})$-competitive algorithm, where n is the number of agents. This result builds on our identification of an optimal matching algorithm in a general model of online matching with edge weights and an unknown number of agents. In this setting, we also achieve an asymptotically optimal competitive ratio of $\frac{1}{3} - \frac{1}{n}$. Second, when agents arrive in an arbitrary order but algorithms are allowed to irrevocably and entirely dissolve coalitions, we show that another matching-based algorithm achieves an optimal competitive ratio of $\frac{1}{6+4\sqrt{2}}$.

Martin Bullinger
University of Bristol
School of Engineering Mathematics and Technology
martin.bullinger@bristol.ac.uk

René Romen, Alexander Schlenga
Technical University of Munich
School of Computation, Information, and Technology
rene.romen@tum.de, alexander.schlenga@tum.de

CP14

Improved Maximin Share Guarantee for Additive Valuations

The maximin share (MMS) is the most prominent share-based fairness notion in the fair allocation of indivisible goods. Recent years have seen significant efforts to improve the approximation guarantees for MMS for different valuation classes, particularly for additive valuations. For the additive setting, it has been shown that for some instances, no allocation can guarantee a factor better than $1 - \frac{1}{n^4}$ of maximin share value to all agents. However, the best currently known algorithm achieves an approximation guarantee of $\frac{3}{4} + \frac{3}{3836}$ for MMS. In this work, we narrow this gap and improve the best-known approximation guarantee for MMS to $\frac{10}{13}$.

Ehsan Heidari, Alireza Kaviani, Masoud Seddighin, Amirmohammad Shahrezaei
Sharif University of Technology
eh3an1383@gmail.com, akaviani05@gmail.com, seddighin.masood@gmail.com, s.a.m.shahrezaei@gmail.com

CP15

A Csp Approach to Graph Sandwich Problems

The Sandwich Problem (SP) for a graph class $\mathcal{C}$ is the following computational problem. The input is a pair of graphs (V, E_1) and (V, E_2) where $E_1 \subseteq E_2$, and the task is to decide whether there is an edge set E where $E_1 \subseteq E \subseteq E_2$ such that the graph (V, E) belongs to $\mathcal{C}$. In this paper we show that many SPs correspond to the constraint satisfaction problem (CSP) of an infinite 2-edge-coloured graph H. We then notice that several known complexity results for SPs also follow from general complexity classifications of infinite-domain CSPs, suggesting a fruitful application of the theory of CSPs to complexity classifications of SPs. We strengthen this evidence by using basic tools from constraint satisfaction theory to propose new complexity results of the SP for several graph classes including line graphs of multigraphs, line graphs of bipartite multigraphs, K_k-free perfect graphs, and classes described by forbidding finitely many induced subgraphs, such as $\{I_4, P_4\}$-free graphs, settling an open problem of Alvarado, Dantas, and Rautenbach (2019). We also construct a graph sandwich problem which is in coNP, but neither in P nor coNP-complete (unless P = coNP).

Santiago Guzmán Pro
TU Dresden
santiago.guzman_pro@tu-dresden.de

Manuel Bodirsky
Institut für Algebra, TU Dresden, Germany.
manuel.bodirsky@tu-dresden.de

CP15

Faster Negative Length Shortest Paths by Bootstrapping Hop Reducers

The textbook algorithm for real-weighted single-source shortest paths takes $O(mn)$ time on a graph with m edges and n vertices. The breakthrough algorithm by [Fin24] takes $\tilde{O}(mn^{8/9})$ randomized time. The running time was subsequently improved to $\tilde{O}(mn^{4/5})$ [HJQ25]. We build on [Fin24, HJQ25] to obtain an $\tilde{O}(mn^{3/4} + m^{4/5}n)$ randomized running time. (Equivalently, $\tilde{O}(mn^{3/4})$ for $m \geq n^{5/4}$, and $\tilde{O}(m^{4/5}n)$ for $m \leq n^{5/4}$.) The main new technique replaces the hop-reducing auxiliary graph from [Fin24] with a bootstrapping process where constant-hop reducers for small subgraphs of the input graph are iteratively amplified and expanded until the desired polynomial-hop reduction is achieved over the entire graph.

Yufan Huang
Purdue University
huan1754@purdue.edu

Peter Jin
University of Illinois Urbana-Champaign
peterj5@illinois.edu

Kent Quanrud

Purdue University
krq@purdue.edu

CP15

Disjoint Paths in Expanders in Deterministic Almost-Linear Time Via Hypergraph Perfect Matching

We design efficient deterministic algorithms for finding short edge-disjoint paths in expanders. Specifically, given an n-vertex m-edge expander G of conductance ϕ and minimum degree δ, and a set of pairs $\{(s_i,t_i)\}_i$ such that each vertex appears in at most k pairs, our algorithm deterministically computes a set of edge-disjoint paths from s_i to t_i, one for every i, 1. each of length at most $18\log(n)/\phi$ and in $mn^{1+o(1)}\min\{k,\phi^{-1}\}$ total time, assuming $\phi^3\delta \geq (35\log n)^3 k$, or 2. each of length at most $n^{o(1)}/\phi$ and in total $m^{1+o(1)}$ time, assuming $\phi^3\delta \geq n^{o(1)}k$. Before our work, deterministic polynomial-time algorithms were known only for expanders with constant conductance and were significantly slower. To obtain our result, we give an almost-linear time algorithm for ¡i¿hypergraph perfect matching¡/i¿ under generalizations of Hall-type conditions [Haxell 1995], a powerful framework with applications in various settings, which until now has only admitted large polynomial-time algorithms [Annamalai 2018].

Matija Bucic
Princeton University
Institute for Advanced Study
mb5225@princeton.edu

Zhongtian He
Princeton University
hezhongtian1@gmail.com

Shang-En Huang, Thatchaphol Saranurak
University of Michigan
sehuang@csie.ntu.edu.tw, thsa@umich.edu

CP15

Balanced Spanning Tree Distributions Have Separation Fairness

Sampling-based methods such as ReCom are widely used to audit redistricting plans for fairness, with the balanced spanning tree distribution playing a central role since it favors compact, contiguous, and population-balanced districts. However, whether such samples are truly representative or exhibit hidden biases remains an open question. In this work, we introduce the notion of separation fairness, which asks whether adjacent geographic units are separated with at most a constant probability (bounded away from one) in sampled redistricting plans. Focusing on grid graphs and two-district partitions, we prove that a smooth variant of the balanced spanning tree distribution satisfies separation fairness. Our results also provide theoretical support for popular MCMC methods like ReCom, suggesting that they maintain fairness at a granular level in the sampling process. Along the way, we prove a novel local-interchangeability lemma for 2-partitions on grids, showing that any separation of adjacent vertices can be undone via a constant-sized modification. This lemma, along with our other tools for analyzing the structure of partitions and loop-erased random walks, may be of independent interest.

Harry Chen
MIT
harryc60@mit.edu

Kamesh Munagala, Govind S. Sankar
Duke University
munagala@duke.edu, govind.subash.sankar@duke.edu

CP15

The Directed Disjoint Paths Problem with Congestion

The classic result by Fortune, Hopcroft, and Wyllie[TCS '80] states that the directed disjoint paths problem is NP-complete even for two pairs of terminals. Extending this well-known result, we show that the directed disjoint paths problem is NP-complete for any constant congestion $c \geq 1$ and $k \geq 3c-1$ pairs of terminals. This refutes a conjecture by Giannopoulou et al. [SODA '22], which says that the directed disjoint paths problem with congestion two is polynomial-time solvable for any constant number k of terminal pairs. We then consider the cases that are not covered by this hardness result. The first nontrivial case is $c = 2$ and $k = 3$. Our second main result is to show that this case is polynomial-time solvable.

Stephan Kreutzer, Matthias Bentert, Dario Cavallaro, Amelie Heindl
Technical University Berlin
stephan.kreutzer@tu-berlin.de, bentert@tu-berlin.de, d.cavallaro@tu-berlin.de, a.heindl@tu-berlin.de

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

Johannes Schröder
Technical University Berlin
j.schroeder.1@tu-berlin.de

CP15

Perfect Matchings in Random Sparsifications of Dense Hypergraphs

The decision problem of perfect matchings in uniform hypergraphs is famously an NP-complete problem. It has been shown by Keevash–Knox–Mycroft [STOC, 2013] that for every $\varepsilon > 0$, such decision problem restricted to k-uniform hypergraphs H satisfying that every $(k-1)$-set of vertices is in at least $(1/k+\varepsilon)|H|$ edges is tractable, and the quantity $1/k$ is best possible. In this paper we study the existence of perfect matchings in the random p-sparsification of such k-uniform hypergraphs, that is, for $p = p(n) \in [0,1]$, every edge is kept with probability p independent of others. As a consequence, we give a polynomial-time algorithm that with high probability solves the decision problem; we also derive effective bounds on the number of perfect matchings in such hypergraphs. At last, similar results are obtained for the F-factor problem in graphs. The key ingredients of the proofs are a strengthened partition lemma for the lattice-based absorption method, and the random redistribution method developed recently by Kelly, Müyesser and Pokrovskiy, based on the spread method.

Jie Han
Department of Mathematics
University of Rhode Island
hanjie@bit.edu.cn

Jingwen Zhao
School of Mathematics and Statistics
Beijing Institute of Technology
jingwen.zhao@bit.edu.cn

CP16

Explicit Min-Wise Hash Families with Optimal Size

We study explicit constructions of min-wise hash and their extension to k-min-wise hash. A min-wise hash family guarantees that for any fixed subset $X \subseteq [N]$, every element in X has an equal chance to have the smallest value among all elements in X; a k-min-wise hash family guarantees this for every subset of size k in X. Min-wise hash is widely used in many areas of computer science such as sketching, web page detection, and ℓ_0 sampling. The classical works by Indyk and Pătraşcu and Thorup have shown $\Theta(\log(1/\delta))$-wise independent families give min-wise hash of multiplicative error δ, resulting in constructions of $\Theta(\log(1/\delta) \log N)$ random bits. Based on a reduction by Saks, Srinivasan, Zhou and Zuckerman, Gopolan and Yehudayoff improved the number of bits to $O(\log N \log\log N)$ for polynomially small errors δ. However, no construction with $O(\log N)$ bits and sub-constant error was known before. We continue and extend the study of constructing min-wise hash families from pseudorandomness. Our result gives the first explicit min-wise hash families that use an optimal number of random bits and achieve a sub-constant error, an explicit family of k-min-wise hash with $O(k \log N)$ bits and $2^{-O(\frac{\log N}{\log\log N})}$ error. Our main techniques involve new ideas to adapt the Nisan-Zuckerman pseudorandom generator to fool min-wise hashing with a multiplicative error.

Xue Chen, Shengtang Huang
University of Science and Technology of China
xuechen1989@ustc.edu.cn, peanuttang@mail.ustc.edu.cn

Xin Li
Johns Hopkins University
lixints@cs.jhu.edu

CP16

Efficient Online Random Sampling Via Randomness Recycling

This article studies the fundamental problem of using i.i.d. coin tosses from an entropy source to efficiently generate random variables $X_i \sim P_i$ $(i \geq 1)$, where $(P_1, P_2, \dots)$ is a random sequence of rational discrete probability distributions subject to an *arbitrary* stochastic process. Our method achieves an amortized expected entropy cost within $\varepsilon > 0$ bits of the information-theoretically optimal Shannon lower bound using $O(\log(1/\varepsilon))$ space. This result holds both pointwise in terms of the Shannon information content conditioned on X_i and P_i, and in expectation to obtain a rate of $\mathbb{E}[H(P_1)+\cdots+H(P_n)]/n+\varepsilon$ bits per sample as $n \to \infty$ (where H is the Shannon entropy). The combination of space, time, and entropy properties of our method improves upon the Knuth and Yao (1976) entropy-optimal algorithm and Han and Hoshi (1997) interval algorithm for online sampling, which require unbounded space. It also uses exponentially less space than the more specialized methods of Kozen and Soloviev (2022) and Shao and Wang (2025) that generate i.i.d. samples from a fixed distribution. Our online sampling algorithm rests on a powerful algorithmic technique called *randomness recycling*, which reuses a fraction of the random information consumed by a probabilistic algorithm to reduce its amortized entropy cost. We apply randomness recycling to accelerate several prominent sampling algorithms.

Thomas L. Draper, Feras Saad
Carnegie Mellon University
tdraper@cmu.edu, fsaad@cmu.edu

CP16

Comparison Theorems for the Mixing Times of Systematic and Random Scan Dynamics

A popular method for sampling from high-dimensional distributions is the Gibbs sampler, which iteratively resamples sites from the conditional distribution of the desired measure given the values of the other coordinates. It is natural to ask to what extent does the order of site updates matter in the mixing time? Two popular choices are (i) standard, or random scan, Glauber dynamics where the updated variable is chosen uniformly at random, and (ii) the systematic scan dynamics where variables are updated in a fixed, cyclic order. We first show that for systems of dimension n, one round of the systematic scan dynamics has spectral gap at most a factor of order n worse than the corresponding spectral gap of a single step of Glauber dynamics, tightening existing bounds in the literature by He, et al. [NeurIPS '16] and Chlebicka, Latuszyński, and Miasodejow [Ann. Appl. Probab. '25]. The corresponding bound on mixing times is sharp even for simple spin systems by an explicit example of Roberts and Rosenthal [Int. J. Statist. Prob. '15]. We complement this with a converse statement: if all, or even just one scan order rapidly mixes, the Glauber dynamics has a polynomially related mixing time, resolving a question of Chlebicka, Latuszyński, and Miasodejow.

Jason Gaitonde
Duke University
jason.gaitonde@duke.edu

Elchanan Mossel
MIT
elmos@mit.edu

CP16

A Classical Quadratic Speedup for Planted Kxor

A recent work of Schmidhuber et al. (QIP, SODA, & Phys. Rev. X 2025) exhibited a quantum algorithm for the noisy planted $k$xor problem running quartically faster than all known classical algorithms. In this work, we design a new classical algorithm that is quadratically faster than the best previous one, in the case of large constant k. Thus for such k, the quantum speedup of Schmidhuber et al. becomes only quadratic (though it retains a space advantage). Our algorithm, which also works in the semirandom case, combines tools from sublinear-time algorithms (essentially, the birthday paradox) and polynomial anticoncentration.

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

Meghal Gupta
UC Berkeley
meghal@berkeley.edu

Noah Singer
Carnegie Mellon University

ngsinger@cs.cmu.edu

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

CP16
On Distributed Colouring of Hyperbolic Random Graphs

We analyse simple distributed colouring algorithms on hyperbolic random graphs (HRGs), a generative model capturing properties of real-world networks such as power-law degrees and large clustering. We focus on the number of rounds and the colour space needed to colour HRGs in the distributed setting. In the simplest algorithm, each vertex selects a random colour and keeps it if no neighbour picks the same; otherwise it retries next round. We show this algorithm terminates in two rounds a.a.s. when using $\varepsilon \cdot \Delta$ colours for any constant $\varepsilon > 0$, while it fails w.h.p. if given only $\chi \cdot n^{\delta}$ colours for small constant $\delta > 0$. We then analyse a standard variation that breaks conflicts by smaller IDs. It also succeeds with $\varepsilon \cdot \Delta$ colours in constant rounds, but fails w.h.p. if only $\Delta / \log^{\Omega(1)} \Delta \gg \chi$ colours are available. Lastly, we propose a version prioritising high-degree vertices, exploiting HRG structure. It achieves a valid colouring in two rounds a.a.s. using $\Delta^{1-\delta}$ colours for some constant $\delta > 0$. This is asymptotically tight up to polylogarithmic factors for some power-law exponents. All three algorithms are very simple and run in CONGEST. These results show constant-time algorithms can outperform the worst-case $\Omega(\log^* n)$ lower bound for $\Delta + 1$ colouring from [Linial 87] and [Naor 91]. Our proofs rely on new structural insights into HRGs.

Yannic Maus
TU Graz
yannic.maus@ist.tugraz.at

Janosch Ruff
Hasso Plattner Institute
janosch.ruff@hpi.de

CP16
Improving Algorithmic Efficiency Using Cryptography: Trapdoored Matrices and Applications

Cryptographic primitives have been used for various non-cryptographic objectives, such as eliminating or reducing randomness and interaction. We show how to use cryptography to improve the time complexity of solving computational problems. Specifically, we show that under standard cryptographic assumptions, we can design algorithms that are asymptotically faster than existing ones while maintaining correctness. As a demonstration, we construct a distribution of trapdoored matrices with the following properties: (a) computationally bounded adversaries cannot distinguish a random matrix from one drawn from this distribution (under computational hardness assumptions), and (b) given a trapdoor, we can multiply such an $n \times n$ matrix with any vector in near-linear (in n) time. We provide constructions both over finite fields and over the reals. This enables a broad speedup technique: any algorithm relying on a random matrixsuch as those that use various notions of dimensionality reductioncan replace it with a matrix from our distribution, achieving computational speedups while preserving correctness. Using these, we present the first uniform reduction from worst-case to approximate and average-case matrix multiplication with optimal parameters (improving on HS2025, albeit under computational assumptions), the first WC to average-case reductions for matrix inversion and other linear operations, as well as a speedup of inference time in classification models.

Or Zamir
Tel Aviv University
orzamir90@gmail.com

Vinod Vaikuntanathan
MIT CSAIL
vinodv@mit.edu

CP17
Peeling Rotten Potatoes for a Faster Approximation of Convex Cover

The minimum convex cover problem seeks to cover a polygon P with the fewest convex polygons that lie within P. This problem is $\exists\mathbb{R}$-complete, and the best previously known algorithm, due to Eidenbenz and Widmayer (2001),achieves an $O(\log n)$-approximation in $O(n^{29} \log n)$ time, where n is the complexity of P. In this work we present a novel approach that preserves the $O(\log n)$ approximation guarantee while significantly reducing the running time. By discretizing the problem and formulating it as a set cover problem, we focus on efficiently finding a convex polygon that covers the largest number of uncovered regions, in each iteration of the greedy algorithm. This core subproblem, which we call the rotten potato peeling problem, is a variant of the classic potato peeling problem. We solve it by finding maximum weighted paths in Directed Acyclic Graphs (DAGs) that correspond to visibility polygons, with the DAG construction carefully constrained to manage complexity. Our approach yields a substantial improvement in the overall running time and introduces techniques that may be of independent interest for other geometric covering problems.

Omrit Filtser
Stony Brook University
omrit.filtser@gmail.com

Tzalik Maimon, Ofir Yomtovyan
The Open University of Israel
post.tmx@gmail.com, ofyotb1@gmail.com

CP17
Online Joint Replenishment Problem with Arbitrary Holding and Backlog Costs

In their seminal paper [Moseley, Niaparast, and Ravi, "Putting Off the Catching Up: Online Joint Replenishment Problem with Holding and Backlog Costs"] they introduced the Joint Replenishment Problem (JRP) with holding and backlog costs that models the trade-off between ordering costs, holding costs, and backlog costs in supply chain planning systems. Their model generalized the classical make-to-order version as well make-to-stock version. For the case where holding costs function of all items are the same and all backlog costs are the same, they provide a constant competitive algorithm, leaving designing a constant competitive algorithm for arbitrary functions open. Moreover, they noticed that their algorithm does not work for arbitrary (request dependent) holding costs and backlog costs functions. We resolve their open problem and design a constant competitive algorithm that works for arbitrary request dependent functions. Specifically, we establish a

4-competitive algorithm for the single-item case and a 16-competitive for the general (multi-item) version. The algorithm of Moseley, Niaparast, and Ravi is based on fixed priority on the requests to items, and request to an item are always served by order of deadlines. In contrast, we design an algorithm with dynamic priority over the requests such that instead of servicing a prefix by deadline of requests, we may need to service a general subset of the requests.

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

Shahar Lewkowicz
Tel Aviv University
shaharlewko22@gmail.com

CP17

Breaching the 2-Approximation Barrier for Euclidean Capacitated Vehicle Routing

In the (Unit Demand) Euclidean Capacitated Vehicle Routing problem (CVRP), we are given a collection of n points in the Euclidean plane (the clients), one extra point (the depot), and one integer $Q \geq 1$ (the vehicle capacity). A feasible solution is a collection of tours, where each tour contains the depot and at most Q clients, such that each client belongs to at least one such tour. Our goal is to minimize the total length of the tours. This models, e.g., the problem of delivering identical items stored at the depot to clients using a single vehicle that can carry at most Q items at a time. The current best (polynomial-time) approximation algorithm [Haimovich and Kan, '85] for CVRP works as follows. First one computes a (TSP) tour, and then one splits it into tours containing at most Q clients each, where this step increases the cost of the solution by the so-called radial lower bound $rlb \leq opt$. Altogether this gives a $2+\epsilon$ approximation. We break the long-standing 2-approximation barrier for CVRP. Let $rlb = (1-\gamma)opt$ for some $\gamma \in [0,1]$. The above algorithm provides a $2+\epsilon-\gamma$ approximation, hence bad when the radial lower bound is large. We complement this with a new and drastically different algorithm, based on a careful use of matroid intersection, which achieves a $1+\epsilon+O(\sqrt[5]{\gamma})$ approximation. Hence the approximation factor of our new algorithm approaches 1 when rlb approaches opt.

Zachary Friggstad
University of Alberta
zacharyf@ualberta.ca

Fabrizio Grandoni
IDSIA
fabrizio.grandoni@gmail.com

Ramin Mousavi
IDSIA, USI-SUPSI
ram.mousavi@gmail.com

CP17

Approximating Matroid Basis Testing for Partition Matroids Using Budget-In-Expectation

We consider the following Stochastic Boolean Function Evaluation problem, which is closely related to several problems from the literature. A matroid $mathcalM$ (in compact representation) on ground set E is given, and each element $i \in E$ is active independently with known probability $p_i \in (0,1)$. The elements can be queried, upon which it is revealed whether the respective element is active or not. The goal is to find an adaptive querying strategy for determining whether there is a basis of $\mathcal{M}$ in which all elements are active, with the objective of minimizing the expected number of queries. When $\mathcal{M}$ is a uniform matroid, this is the problem of evaluating a k-of-n function, first studied in the 1970s. This problem is well-understood, and has an optimal adaptive strategy that can be computed in polynomial time. Taking $\mathcal{M}$ to instead be a partition matroid, we show that previous approaches fail to give a constant-factor approximation. Our main result is a randomized polynomial-time constant-factor approximation algorithm for this partition matroid problem, for which we combine a new technique with several well-established techniques. Our algorithm adaptively interleaves solutions to several instances of a novel type of stochastic querying problem, with a constraint on the *expected* cost. We believe that this type of problem is of independent interest, will spark follow-up work, and has the potential for additional applications.

Lisa Hellerstein
NYU
lisa.hellerstein@nyu.edu

Benedikt M. Plank
independent
b.plank@mail.de

Kevin Schewior
University of Cologne, Köln, Germany
University of Southern Denmark, Odense, Denmark
k.schewior@uni-koeln.de

CP17

Improved Online Algorithms for Inventory Management Problems with Holding and Delay Costs: Riding the Wave Makes Things Simpler, Stronger, and More General

The Joint Replenishment Problem (JRP) is a classical inventory management problem, that aims to model the trade-off between coordinating orders for multiple commodities (and their cost) with holding costs incurred by meeting demand in advance. Recently, Moseley, Niaparast and Ravi introduced a natural online generalization of the JRP in which inventory corresponding to demands may be replenished late, for a delay cost, or early, in which case there is a holding cost associated with storing it until the desired service time. They established that when the holding and delay costs are monotone and uniform across demands, there is a 30-competitive algorithm. We develop a 5-competitive algorithm that handles arbitrary monotone demand-specific holding and delay cost functions. Our primal-dual algorithm is in the sprit of the work Buchbinder, Kimbrel, Levi, Makarychev, and Sviridenko, which maintains a wavefront dual solution to decide when to place an order and which items to order. The main twist is in deciding which requests to serve early. An important special case of the JRP is the single-item lot-sizing problem. Here, Moseley et al. gave a 3-competitive algorithm when the holding and delay costs are uniform across demands. We provide a new algorithm for which the competitive ratio is $\phi + 1 \approx 2.681$, where ϕ is the golden ratio.

Varun Suriyanarayana, David B. Shmoys
Cornell University
vs478@cornell.edu, david.shmoys@cornell.edu

William Umboh
University of Melbourne
william.umboh@unimelb.edu.au

CP17

Smooth Trade-off for Tensor Pca Via Sharp Bounds for Kikuchi Matrices

In this work, we revisit algorithms for Tensor PCA: given an order-r tensor of the form $T = G + \lambda \cdot v^{\otimes r}$ where G is a random symmetric Gaussian tensor with unit variance entries and v is an unknown boolean vector in $\{\pm 1\}^n$, what's the minimum λ at which one can distinguish T from a random Gaussian tensor and more generally, recover v? As a result of a long line of work, we know that for any $\ell \in$, there is a $n^{O(\ell)}$ time algorithm that succeeds when the signal strength $\lambda \gtrsim \sqrt{\log n} \cdot n^{-r/4} \cdot \ell^{1/2-r/4}$. The question of whether the logarithmic factor is necessary turns out to be crucial to understanding whether larger polynomial time allows recovering the signal at a lower signal strength. Such a smooth trade-off is necessary for tensor PCA being a candidate problem for quantum speedups [**?**]. It was first conjectured by [**?**] and then, more recently, with an eye on smooth trade-offs, reiterated in a blogpost of Bandeira. In this work, we resolve these conjectures and show that spectral algorithms based on the Kikuchi hierarchy [**?**] succeed whenever $\lambda \geq \Theta_r(1) \cdot n^{-r/4} \cdot \ell^{1/2-r/4}$ where $\Theta_r(1)$ only hides an absolute constant independent of n and ℓ. A sharp bound such as this was previously known only for $\ell \leq 3r/4$ via non-asymptotic techniques in random matrix theory inspired by free probability.

Jeff Xu
TTIC
jeffxusichao@cmu.edu

CP17

Approximate Light Spanners in Planar Graphs

In their seminal work, Althfer et al. (DCG 1993) showed that the greedy $(1 + \epsilon)$-spanner of any weighted planar graph G has weight at most $(1+2/\epsilon) \cdot \omega(\mathsf{MST}(G))$, and this bound is existentially tight. Moreover, even as a bicriteria approximation, the greedy spanner can be far from optimal: for some planar graphs, its weight is $\Omega\big(\frac{1}{\epsilon x^2}\big) \cdot \omega(G_{\text{opt},\epsilon})$ when allowing stretch $1 + x\epsilon$ (with $x = O(\epsilon^{-1/2})$). Despite decades of research, no algorithm has beaten this existential bound for light spanners in planar graphs. We present the first polynomial-time algorithm that constructs, for any weighted planar graph G, a $\big(1 + \epsilon \cdot 2^{O(\log^* 1/\epsilon)}\big)$-spanner of weight $O(1) \cdot \omega(G_{\text{opt},\epsilon})$. Our key innovation is *iterative planar pruning*: we iteratively replace heavy edge sets with light paths, using planarity to establish a laminar structure on removable edges, which enables dynamic programming directly on the input graphan approach distinct from prior planar network design techniques and of potential independent interest.

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

Shay Solomon
Tel Aviv University
solo.shay@gmail.com

Cuong Than
University of Massachusetts Amherst
cthan@umass.edu

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

Tianyi Zhang
ETH Zurich
tianyiz25@nju.edu.cn

CP18

Detecting Correlation Efficiently in Very Supercritical Stochastic Block Models: Breaking the Otter's Threshold Barrier

Consider a pair of sparse correlated stochastic block models $\mathcal{S}(n, \frac{\lambda}{n}, \epsilon; s)$ subsampled from a common parent stochastic block model with two symmetric communities, average degree $\lambda = O(1)$, divergence parameter $\epsilon \in (0, 1)$ and subsampling probability s. For all $\epsilon \in (0, 1)$, we construct a statistic based on the combination of two low-degree polynomials and show that there exists a sufficiently small constant $\delta = \delta(\epsilon) > 0$ and a sufficiently large constant $\Delta = \Delta(\epsilon, \delta)$ such that when $\lambda > \Delta$ and $s > \sqrt{\alpha} - \delta$ where $\alpha \approx 0.338$ is Otter's constant, this statistic can distinguish this model and a pair of independent stochastic block models $\mathcal{S}(n, \frac{\lambda s}{n}, \epsilon)$ with probability $1 - o(1)$. We also provide an efficient algorithm that approximates this statistic in polynomial time. Our result is the first detection or matching type algorithm that breaks the Otter's threshold in sparse correlated random graphs. The crux of our statistic's construction lies in a carefully curated family of multigraphs called decorated trees, which enables effective aggregation of the community signal and graph correlation from the counts of the same decorated tree while suppressing the undesirable correlations among counts of different decorated trees. We believe such construction may be of independent interest.

Zhangsong Li, Guanyi Chen, Jian Ding, Shuyang Gong
School of Mathematical Sciences, Peking University
ramblerlzs@pku.edu.cn, cgy7698@stu.pku.edu.cn, dingjian@math.pku.edu.cn, gongshuyang@stu.pku.edu.cn

CP18

Temporal Exploration of Random Spanning Tree Models

The Temporal Graph Exploration problem (TEXP) takes as input a temporal graph, i.e., a sequence of graphs $(G_i)_{i\in\mathbb{N}}$ on the same vertex set, and asks for a walk of shortest length visiting all vertices, where the i-th step uses an edge from G_i or stays put. If each such G_i is connected, then an exploration of length n^2 exists, and this is known to be the best possible up to a constant. More fine-grained lower and upper bounds have been obtained for restricted temporal graph classes, however, a large gap persists between known bounds. We study the problem in a randomised setting. We introduce the Random Spanning Tree (RST) model, which consists of a set of n-vertex trees together with an arbitrary probability distribution μ over this set. A random temporal graph generated by the RST model is a sequence of independent samples drawn from μ. We initiate a systematic study of the Temporal Graph Exploration problem in such random temporal graphs and establish tight general bounds on exploration time. Our first main result proves that any RST model can, w.h.p., be explored in $O(n^{3/2})$ time, and we show that this bound is tight up to a constant factor. This demonstrates a fundamental difference between the adversarial and random settings. Our second main result shows that if all trees of an RST are subgraphs of a fixed graph with m edges then, w.h.p., it can be explored in $O(m)$ time.

Samuel Baguley
Hasso-Plattner-Institute
samuel.baguley@hpi.de

Andreas Göbel
Hasso Plattner Institute
University of Potsdam
andreas.goebel@hpi.de

Nicolas Klodt, George Skretas
Hasso-Plattner-Institute
nicolas.klodt@hpi.de, georgios.skretas@hpi.de

John Sylvester
University of Liverpool
john.sylvester@liverpool.ac.uk

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

CP18

Expander Pruning with Polylogarithmic Worst-Case Recourse and Update Time

Expander graphs are known to be robust to edge deletions in the following sense: for any online sequence of edge deletions $e_1, e_2, \dots, e_k$ to an m-edge graph G that is initially a ϕ-expander, the algorithm can grow a set $P \subseteq V$ such that at any time t, $G[V \setminus P]$ is an expander of the same quality as the initial graph G up to a constant factor and the set P has volume at most $O(t/\phi)$. However, currently, there is no algorithm to grow P with low worst-case recourse that achieves any non-trivial guarantee. In this work, we present an algorithm that achieves near-optimal guarantees: we give an algorithm that grows P only by $\tilde{O}(1/\phi^2)$ vertices per time step and ensures that $G[V \setminus P]$ remains $\tilde{\Omega}(\phi)$-expander at any time. Even more excitingly, our algorithm is extremely efficient: it can process each update in near-optimal worst-case update time $\tilde{O}(1/\phi^2)$. This affirmatively answers the main open question posed in SW'19 whether such an algorithm exists.

Simon Meierhans, Maximilian Probst Gutenberg
ETH Zurich
mesimon@inf.ethz.ch, maximilian.probst@inf.ethz.ch

Thatchaphol Saranurak
University of Michigan
thsa@umich.edu

CP18

An Optimal Density Bound for Discretized Point Patrolling

The pinwheel problem is a real-time scheduling problem that asks, given n tasks with periods $a_i \in \mathbb{N}$, whether it is possible to schedule the tasks infinitely, one per time unit, so that every task i is scheduled in every interval of a_i time units. We study a corresponding version of this packing problem in the covering setting, stylized as the discretized point patrolling problem in the literature. Specifically, given n tasks with periods a_i, the problem asks whether it is possible to assign each day to a task so that every task i is scheduled at most once every a_i days. The density of an instance in either case is defined as the sum of the inverses of task periods. Recently, the long-standing 5/6 density bound conjecture in the packing setting was resolved affirmatively. We resolve a corresponding 10-year-old conjecture in the covering setting by proving that every discretized point patrolling instance with density at least $\sum_{i=0}^{\infty} 1/(2^i+1) \approx 1.264$ is schedulable. This significantly improves upon the current best-known density bound of 1.546 and is, in fact, optimal. We also study the bamboo garden trimming problem, an optimization variant of the pinwheel problem. We achieve an efficient 9/7-approximation algorithm for this problem, improving on the current best known approximation factor of 4/3.

Ahan Mishra
Cornell University
ahanbmishra@gmail.com

CP18

Faster Negative-Weight Shortest Paths and Directed Low-Diameter Decompositions

We present a faster algorithm for low-diameter decompositions on directed graphs, matching the $O(\log n \log\log n)$ loss factor from Bringmann, Fischer, Haeupler, and Latypov (ICALP 2025) and improving the running time to $O((m + n\log\log n)\log n \log\log n)$ in expectation. We then apply our faster low-diameter decomposition to obtain an algorithm for negative-weight single source shortest paths on integer-weighted graphs in $O((m + n\log\log n)\log(nW)\log n\log\log n)$ time, a nearly log-factor improvement over the algorithm of Bringmann, Cassis, and Fischer (FOCS 2023).

Jason M. Li, Connor Mowry
Carnegie Mellon University
jmli@cs.cmu.edu, cmowry@andrew.cmu.edu

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

CP18

New Oracles and Labeling Schemes for Vertex Cut Queries

We study succinct representations of vertex cuts by centralized oracles and labeling schemes. For an n-vertex graph $G = (V, E)$ and parameter $f \geq 1$, the goal is supporting vertex cut queries: Given $F \subseteq V$ with $|F| \leq f$, determine if F is a vertex cut in G. In the centralized setting, it is required to preprocess G into an f-vertex cut oracle that can answer such queries quickly, while occupying only small space. In the labeling setting, one should assign a short label to each vertex in G, so that a cut query F can be answered by merely inspecting the labels of the vertices in F. We provide the first significant progress on these problems: f-Vertex Cut Labels: Every n-vertex graph admits an f-vertex cut labeling scheme, where the labels have length of $\tilde{O}(n^{1-1/f})$ bits (when f is polylogarithmic in n). This nearly matches the recent lower bound given by Long, Pettie and Saranurak (SODA 2025). f-Vertex Cut Oracles: For $f = O(\log n)$, every n-vertex graph G admits f-vertex cut oracle with $\tilde{O}(n)$ space and $\tilde{O}(2^f)$ query time. We also show that our f-vertex cut oracles for every $1 \leq f \leq n$ are optimal up to $n^{o(1)}$ factors (conditioned on fine-grained complexity conjectures). If G is f-connected, i.e., when one is interested in minimum vertex cuts queries, the query time improves to $\tilde{O}(f^2)$, for any $1 \leq f \leq n$.

Asaf Petruschka
Weizmann Institute
asaf.petruschka@weizmann.ac.il

Yonggang Jiang
MPI-INF
yjiang@mpi-inf.mpg.de

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

CP19
Constructive L2-Discrepancy Minimization with Additive Deviations

The *signed series* problem in the ℓ_2 norm asks, given a set of vectors $v_1, \ldots, v_n \in R^d$ having at most unit ℓ_2 norm, does there always exist a series $(\varepsilon_i)_{i\in[n]}$ of ± 1 signs such that for all $i \in [n]$, $\max_{i\in[n]} \| \sum_{j=1}^{i} \varepsilon_i v_i \|_2 = O(\sqrt{d})$. We give a polynomial-time randomized algorithm to find signs $x(i) \in \{-1, 1\}$, $i \in [n]$ such that

$$\max_{i\in[n]} \| \sum_{j=1}^{i} x(i) v_i \|_2 = O(\sqrt{d + \log n}).$$

By Harvey and Samadi (2014), this also yields a constructive bound of $O(\sqrt{d + \log n})$ for the Steinitz problem in the ℓ_2-norm, matching Banaszczyk's (2012) existential bounds for both problems, and improving upon the bounds of Bansal and Garg (2017). Our algorithm is based on the framework on Bansal and Garg, together with a new analysis involving (i) additional linear orthogonality constraints during the construction of the covariance matrix of the random walk steps, alongwith (ii) a "Freedman-like" version of the Hanson-Wright concentration inequality, for filtration-dependent sums of subgaussian chaoses. As further applications of our method, we also obtain some new discrepancy bounds, including an ℓ_2-discrepancy bound for combinatorial vector balancing over set systems of bounded VC dimension, and a *local mean discrepancy* bound for the Komlos discrepancy problem.

Kunal Dutta
Institute of Informatics, University of Warsaw
kdutta@mimuw.edu.pl

CP19
A Tutte-type Canonical Decomposition of 3- and 4-connected Graphs

We provide a unique decomposition of every 4-connected graph into parts that are either quasi-5-connected, cycles of triangle-torsos and 3-connected torsos on ≤ 5 vertices, generalised double-wheels, or thickened $K_{4,m}$'s. The decomposition can be described in terms of a tree-decomposition but with edges allowed in the adhesion-sets. Our construction is explicit, canonical, and exhibits a defining property of the Tutte-decomposition. As a corollary, we obtain a new Tutte-type canonical decomposition of 3-connected graphs into parts that are either quasi-4-connected, generalised wheels or thickened $K_{3,m}$'s. This decomposition is similar yet different from the tri-separation decomposition of Carmesin and Kurkofka. As an application of the decomposition for 4-connectivity, we obtain a new theorem characterising all vertex-transitive finite connected graphs as essentially quasi-5-connected or on a short explicit list of graphs.

Jan Kurkofka, Jan Kurkofka, Tim Planken
TU Freiberg
jan.kurkofka@math.tu-freiberg.de, jan.kurkofka@math.tu-freiberg.de, tim.planken@math.tu-freiberg.de

CP19
Helly-Type Theorems for Splitting Point Sets

Let $0 < \alpha \leq 1/2$. We say that a finite point set P in $\mathbb{R}^d$ is *α-split* by a hyperplane h if each of the closed half-spaces determined by h, contains at least $\alpha|P|$ of the points of P. We further say P is α-split by a k-dimensional flat τ if P is α-split by *any* hyperplane through τ. In the standard notation (which coincides with Tukey depth for $k = 0$), the k-flat τ has depth α with respect to P. We establish interesting Helly-type theorems for splitting families of finite point sets in $\mathbb{R}^d$. Unlike the classical sufficient Helly-type criteria for transversals to families of compact convex sets, which exist only for point and hyperplanes, our results extend to splitting families of point sets by collections of k-flats of arbitrary dimensionality $0 \leq k \leq d-1$.

Lidor Portal
Ben Gurion University of the Negev
lidorpor@post.bgu.ac.il

Natan Rubin
Ben Gurion University
rubinnat.ac@gmail.com

CP19
On the Edge Expansion of Random Polytopes

A 0/1-polytope in $\mathbb{R}^n$ is the convex hull of a subset of $\{0,1\}^n$. The graph of a polytope P is the graph whose vertices are the zero-dimensional faces of P and whose edges are the one-dimensional faces of P. A conjecture of Mihail and Vazirani states that the edge expansion of the graph of every 0/1-polytope is at least one. In this talk, we study a random version of the problem, where the polytope is generated by selecting vertices of $\{0,1\}^n$ independently at random with probability $p \in (0,1)$. Improving earlier results, we show that, for any $p \in (0,1)$, with high probability the edge expansion of the random 0/1-polytope is bounded from below by an absolute constant.

Marcelo Sales
Emory University
mtsales@uci.edu

Asaf Ferber
Department of Mathematics, University of California, Irvine.
asaff@uci.edu

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

Wojciech Samotij
Tel Aviv University

Isreal
samotij@tauex.tau.ac.il

CP19

Time-Biased Random Walks and Robustness of Expanders

Random walks on expanders play a crucial role in Markov Chain Monte Carlo algorithms, derandomization, graph theory, and distributed computing. A desirable property is that they are rapidly mixing, which is equivalent to having a spectral gap γ bounded away from 0. First, we establish a dichotomy for the robustness of mixing times of random walks on edge-weighted d-regular n-vertex expanders subject to a Lipschitz condition, which bounds the ratio of adjacent weights by $\beta \geq 1$. -If $\beta > 1$ is sufficiently small, then for all edge-weightings γ is bounded away from 0 and the mixing time is logarithmic in n. -If $\beta \geq 2d$, then there is an edge-weighting such that γ is polynomial in $1/n$ and the mixing time is polynomial in n. Second, we apply our robustness result to a time-dependent version of the so-called ε-biased random walk. -We show that, for any constant $\varepsilon > 0$, a bias strategy can be chosen adaptively so that the ε-biased random walk covers any bounded-degree regular expander in $\Theta(n)$ expected time. This is optimal and improves over the previous-best bound of $O(n \log\log n)$. -We prove the first non-trivial lower bound on the cover time of the ε-biased random walk, showing that, on bounded-degree regular expanders, it is $\omega(n)$ whenever $\varepsilon = o(1)$.

Thomas Sauerwald
University of Cambridge
tms41@cam.ac.uk

John Sylvester
University of Liverpool
john.sylvester@liverpool.ac.uk

Sam Olesker Taylor
University of Warwick
oleskertaylor.sam@gmail.com

CP19

On a Clique Game and the Erdos-Hajnal Problem on High-Chromatic High-Girth Subgraphs

For a fixed positive integer k, two players, Builder and Chooser, alternate turns playing the following game on a dynamically changing graph that is initially empty. In each round, Builder introduces a new vertex with edges to all previous vertices and then partitions the entire edge set into two subsets, after which Chooser deletes one of the two. Builder attempts to build a clique of size k, while Chooser attempts to prevent that. We prove tower-type upper and lower bounds on how many rounds Builder needs to guarantee a k-clique. Using this game as a key tool, we provide the first non-trivial lower bound on the well-known conjecture of Erdos and Hajnal that every graph with sufficiently large chromatic number contains a subgraph with large girth and large chromatic number. Namely, we prove that the mth Burling graph, which has chromatic number m, has no subgraph with girth 5 and chromatic number greater than k, where m is a tower of 2s of height linear in k. We also analyze several variants of the clique game, and we use Builder's strategies in these variants to construct high-chromatic subgraphs of Burling graphs with girth 5 and high-chromatic ordered graphs that avoid the pattern $\{1{-}3, 1{-}5, 2{-}4\}$ as an ordered subgraph.

Seth Pettie
University of Michigan
pettie@umich.edu

Gábor Tardos
Rényi Institute, Budapest
tardos@renyi.hu

Bartosz Walczak
Jagiellonian University
bartosz.m.walczak@gmail.com

CP20

Testing Forbidden Order-Pattern Properties on Hypergrids

We study testing π-freeness of functions $f : [n]^d \to \mathbb{R}$, where f is π-free if no k points $x_1 \prec \cdots \prec x_k$ realize the relative order prescribed by π. For $d = 2$ and permutations of size $k = 3$, we give an adaptive, one-sided tester with query complexity $O(n^{4/5+o(1)})$ that works for all π. We also prove general lower bounds for $k = 3$: any nonadaptive tester requires $\Omega(n)$ queries, and any adaptive tester needs $\Omega(\sqrt{n})$, giving the first super-logarithmic lower bounds for π-freeness. For the monotone patterns $\pi = (1, 2, 3)$ and $(3, 2, 1)$ we design a nonadaptive tester with polylogarithmic queries, exhibiting an exponential gap between monotone and nonmonotone patterns (unlike the 1-D case). As a key tool, we develop erasure-resilient (δ-ER) ϵ-testers for monotonicity over $[n]^d$ with query complexity $O\left(\frac{\log^{O(d)} n}{\epsilon(1-\delta)}\right)$, improving prior ER testers that required $\delta = O(\epsilon/d)$; our nonadaptive tester is nearly optimal via a matching lower bound. Finally, we show that current techniques cannot yield sublinear-query testers for patterns of length 4 even on two-dimensional hypergrids.

Harish Chandramouleeswaran
Chennai Mathematical Institute
harishc@cmi.ac.in

Ilan I. Newman
Haifa University
ilan@cs.haifa.ac.il

Tomer Pelleg
Univeristy of Haifa
tomer@pelleg.org

Nithin Varma
University of Cologne
varma@cs.uni-koeln.de

CP20

Streaming and Massively Parallel Algorithms for Euclidean Max-Cut

Given a set of vectors $X = \{x_1, \ldots, x_n\} \subset \mathbb{R}^d$, the Euclidean max-cut problem asks to partition the vectors into two parts so as to maximize the sum of Euclidean distances which cross the partition. We design new algorithms for Euclidean max-cut in models for massive datasets: • We give a fully-scalable constant-round MPC algorithm using $O(nd) + n \cdot \mathrm{poly}(\log(n)/\epsilon)$ total space which gives a $(1 + \epsilon)$-approximate Euclidean max-cut. • We give a dynamic streaming algorithm using $\mathrm{poly}(d \log \Delta/\epsilon)$ space

when $X \subseteq [\Delta]^d$, which provides oracle access to a $(1+\epsilon)$-approximate Euclidean max-cut. Recently, Chen, Jiang, and Krauthgamer [STOC '23] gave a dynamic streaming algorithm with space $\mathrm{poly}(d \log \Delta/\epsilon)$ to approximate the value of the Euclidean max-cut, but could not provide oracle access to an approximately optimal cut. This was left open in that work, and we resolve it here. Both algorithms follow from the same framework, which analyzes a "parallel" and "subsampled" (Euclidean) version of a greedy algorithm of Mathieu and Schudy [SODA '08] for dense max-cut.

Nicolas Menand
University of Pennsylvania
nmenand@seas.upenn.edu

Erik Waingarten
Simons Institute
ewaingar@seas.upenn.edu

CP20

Sublinear Time Low-Rank Approximation of Hankel Matrices

Hankel matrices are an important class of highly-structured matrices, arising across computational mathematics, engineering, and theoretical computer science. It is well-known that positive semidefinite (PSD) Hankel matrices are always approximately low-rank. In particular, a highly celebrated result of Beckermann and Townsend [BT17] shows that, for any PSD Hankel matrix $H \in \mathbb{R}^{n\times n}$ and any $\epsilon > 0$, letting H_k be the best rank-k approximation of H (obtained via truncated singular value decomposition), $\|H - H_k\|_F \leq \epsilon\|H\|_F$ for $k = O(\log n \log(1/\epsilon))$. As such, PSD Hankel matrices are natural targets for low-rank approximation algorithms. We give the first such algorithm that runs in *sublinear time*. In particular, we show how to compute, in $polylog(n, 1/\epsilon)$ time, a factored representation of a rank-$O(\log n \log(1/\epsilon))$ Hankel matrix $\widehat{H}$ matching the error guarantee of Beckermann and Townsend up to constant factors. We further show that our algorithm is *robust* – given input $H + E$ where $E \in \mathbb{R}^{n\times n}$ is an arbitrary non-Hankel noise matrix, we obtain error $\|H - \widehat{H}\|_F \leq O(\|E\|_F) + \epsilon\|H\|_F$. Towards this algorithmic result, our first contribution is a *structure-preserving* existence result: we show that there exists a rank-k *Hankel* approximation to H matching the error bound of Beckermann and Townsend.

Michael Kapralov
EPFL
michael.kapralov@epfl.ch

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

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

CP20

One Attack to Rule Them All: Tight Quadratic Bounds for Adaptive Queries on Cardinality Sketches

Cardinality sketches are compact data structures for representing sets or vectors. These sketches are space-efficient, typically requiring only logarithmic storage in the input size, and enable approximation of cardinality. A crucial property in applications is *composability*, meaning that the sketch of a union of sets can be computed from individual sketches. Existing designs provide strong statistical guarantees, ensuring that a randomly sampled sketching map remains robust for an exponential number of queries in terms of the sketch size k. However, these guarantees degrade to quadratic in k when queries are *adaptive*. Prior works on statistical queries (Steinke and Ullman, 2015) and specific MinHash cardinality sketches (Ahmadian and Cohen, 2024) established that this is tight in that they can be compromised using a quadratic number of adaptive queries. In this work, we develop a universal attack framework that applies to broad classes of cardinality sketches. We show any union-composable sketching map can be compromised with $\tilde{O}(k^4)$ adaptive queries and this improves to a tight bound of $\tilde{O}(k^2)$ for "monotone" maps. Similarly, any linear sketching map over the reals $\mathbb{R}$ and finite fields $\mathbb{F}_p$ can be compromised using $\tilde{O}(k^2)$ adaptive queries, which is optimal and strengthens some of the recent results (Gribelyuk et al., 2024), who established a weaker polynomial bound.

Jelani Nelson
University of California, Berkeley
minilek@berkeley.edu

Edith Cohen, Tamas Sarlos
Google Research
edith@alumni.stanford.edu, stamas@google.com

Mihir Singhal
UC Berkeley
mihir.a.singhal@gmail.com

Uri Stemmer
Tel Aviv University and Google Research
u@uri.co.il

CP20

Spectral Clustering in Birthday Paradox Time

Given a vertex in a (k, ϕ, ϵ)-clusterable graph, i.e. a graph whose vertex set can be partitioned into ϕ-expanders of size $\approx n/k$ with outer conductance bounded by ϵ, can one quickly tell which cluster it belongs to? This question goes back to the expansion testing problem of Goldreich and Ron'11. For $k = 2$, $\approx n^{1/2+O(\epsilon/\phi^2)}$ logarithmic-length walks from a vertex suffice to determine its cluster by the birthday paradox. The study of the general case $k > 2$ was initiated by Czumaj, Peng and Sohler [STOC'15], and the work of Chiplunkar et al. [FOCS'18], Gluch et al. [SODA'21] showed that $\approx poly(k) \cdot n^{1/2+O(\epsilon/\phi^2)}$ samples suffice. This matches the $k = 2$ result up to polynomial factors in k, but creates a conceptual inconsistency: if the birthday paradox guides the query complexity, then the query complexity should decrease, as opposed to increase, with k! We design a novel vertex representation in a (k, ϕ, ϵ)-clusterable graph using mixtures of logarithmic length walks, which uses the optimal $\approx (n/k)^{1/2+O(\epsilon/\phi^2)}$ number of walks per vertex, and supports a fast nearest neighbor search: given $\approx k$ representative vertices, we can find the cluster of a query vertex in nearly linear time. This gives a clustering oracle with query time $\approx (n/k)^{1/2+O(\epsilon/\phi^2)}$ and space $k \cdot (n/k)^{1/2+O(\epsilon/\phi^2)}$, match-

ing the birthday paradox bound.

Michael Kapralov, Ekaterina Kochetkova, Weronika Wrzos-Kaminska
EPFL
michael.kapralov@epfl.ch, ekaterina.kochetkova@epfl.ch, weronika.wrzos-kaminska@epfl.ch

CP20

Online Learning with Limited Information in the Sliding Window Model

We consider the experts problem in the sliding window model. The sliding window model is a well-studied model that captures applications such as traffic monitoring, epidemic tracking, and automated trading, where recent information is more valuable than older data. Formally, we have n experts, T days, the ability to query the predictions of q experts on each day, a limited amount of memory, and should achieve the (near-)optimal regret $\sqrt{nW}\text{polylog}(nT)$ regret over any window of the last W days. While it is impossible to achieve such regret with 1 query, we show that with 2 queries we can achieve such regret and with only polylog(nT) bits of memory. Not only are our algorithms optimal for sliding windows, but we also show for every interval $\mathcal{I}$ of days that we achieve $\sqrt{n|\mathcal{I}|}\text{polylog}(nT)$ regret with 2 queries and only polylog(nT) bits of memory, providing an exponential improvement on the memory of previous interval regret algorithms. Building upon these techniques, we address the bandit problem in data streams, where $q = 1$, achieving $nT^{2/3}\text{polylog}(T)$ regret with polylog(nT) memory, which is the first sublinear regret in the streaming model in the bandit setting with polylogarithmic memory; this can be further improved to the optimal $\mathcal{O}(\sqrt{nT})$ regret if the best expert's losses are in a random order.

Vladimir Braverman
Johns Hopkins University
vova@cs.jhu.edu

Sumegha Garg
Rutgers University
sumegha.garg@rutgers.com

Chen Wang
Rensselaer Polytechnic Institute
wangc33@rpi.edu

David Woodruff, Samson Zhou
Carnegie Mellon University
dwoodruf@andrew.cmu.edu, samsonzhou@gmail.com

CP21

Rapid Mixing for Gibbs States Within a Logical Sector: a Dynamical View of Self-Correcting Quantum Memories

Self-correcting quantum memories store logical quantum information for exponential time in thermal equilibrium at low temperatures. By definition, these systems are slow mixing. This raises the question of how the memory state, which we refer to as the Gibbs state within a logical sector, is created in the first place. In this paper, we show that for a broad class of self-correcting quantum memories on lattices with parity check redundancies, a quasi-local Gibbs sampler rapidly converges to the corresponding low-temperature Gibbs state within a logical sector when initialized from a ground state. This illustrates a dynamical view of self-correcting quantum memories, where the "syndrome sector" rapidly converges to thermal equilibrium, while the "logical sector" remains metastable. As a key application, when initialized from a random ground state, this gives a rapid Gibbs state preparation algorithm for the 4D toric code in polylog(n) depth. The main technical ingredients behind our approach are new, low-temperature decay-of-correlation properties for these metastable states.

Thiago Bergamaschi
MIT
thiagob@berkeley.edu

Reza Gheissari
Northwestern University
gheissari@northwestern.edu

Yunchao Liu
Harvard
yunchaoliu@berkeley.edu

CP21

On the Quantum Chromatic Gap

The largest known gap between quantum and classical chromatic number of graphs, obtained via quantum protocols for colouring Hadamard graphs based on the Deutsch–Jozsa algorithm and the quantum Fourier transform, is exponential. We put forth a quantum pseudo-telepathy version of Khot's d-to-1 Games Conjecture and prove that, conditional on its validity, the gap is unbounded: There exist graphs whose quantum chromatic number is 3 and whose classical chromatic number is arbitrarily large. Furthermore, we show that the existence of a certain form of pseudo-telepathic XOR games would imply the conjecture and, thus, the unboundedness of the quantum chromatic gap. As two technical steps of our proof that might be of independent interest, we establish a quantum adjunction theorem for Pultr functors between categories of relational structures, and we prove that the Dinur–Khot–Kindler–Minzer–Safra reduction, recently used for proving the 2-to-2 Games Theorem, is quantum complete.

Lorenzo Ciardo
TU Graz
lorenzo.ciardo@tugraz.at

CP21

A Post-Quantum Lower Bound for the Distributed Lovsz Local Lemma

In this work, we study the Lovsz local lemma (LLL) problem in the area of distributed quantum computing, which has been the focus of attention of recent advances in quantum computing [STOC'24, STOC'25, STOC'25]. We prove a lower bound of $2^{\Omega(\log^* n)}$ for the complexity of the distributed LLL in the quantum-LOCAL model. More specifically, we obtain our lower bound already for a very well-studied special case of the LLL, called sinkless orientation, in a stronger model than quantum-LOCAL, called the randomized online-LOCAL model. As a consequence, we obtain the same lower bounds for sinkless orientation and the distributed LLL also in a variety of other models studied across different research communities. Our work provides the first superconstant lower bound for sinkless orientation and the distributed LLL in all of these models, addressing recently stated open questions. Moreover, to obtain our results, we develop an entirely new lower bound tech-

nique that we believe has the potential to become the first generic technique for proving post-quantum lower bounds for many of the most important problems studied in the context of locality.

Tim Göttlicher
Aalto University
timgo@mailbox.org

Sebastian Brandt
CISPA Helmholtz Center for Information Security
brandt@cispa.de

CP21

Quantum State Preparation with Optimal T-Count

How many T gates are needed to approximate an arbitrary n-qubit quantum state to within error ε? Improving prior work of Low, Kliuchnikov, and Schaeffer, we show that the optimal asymptotic scaling is $\Theta\left(\sqrt{2^n \log(1/\varepsilon)} + \log(1/\varepsilon)\right)$ if we allow ancilla qubits. We also show that this is the optimal T-count for implementing an arbitrary diagonal n-qubit unitary to within error ε. We describe applications in which a tensor product of many single-qubit unitaries can be synthesized in parallel for the price of one.

Robin Kothari
Google Research
robin@robinkothari.com

David Gosset
IQC Waterloo
dngosset@gmail.com

Kewen Wu
University of California at Berkeley
shlw_kevin@hotmail.com

CP21

Beating Full State Tomography For Unentangled Spectrum Estimation

How many copies of a mixed state $\rho \in \mathbb{C}^{d\times d}$ are needed to learn its spectrum? To date, the best-known algorithms for spectrum estimation require as many copies as full state tomography, suggesting that spectrum estimation might be as hard as tomography. We show that this is not the case in the setting of unentangled measurements, by giving a spectrum estimation algorithm that uses $n = O(d^3 \cdot (\log\log(d)/\log(d))^4)$ copies, asymptotically fewer than the $n = \Omega(d^3)$ copies for full state tomography. Our algorithm is inspired by the classical technique of local moment matching and shows how to apply it in the quantum setting. As an important subroutine in our spectrum estimation algorithm, we give an estimator of the k-th moment $\mathrm{tr}(\rho^k)$ which performs unentangled measurements and uses $O(d^{3-2/k})$ copies of ρ in order to achieve a constant multiplicative error. This directly translates to an additive-error estimator of quantum Renyi entropy of order k with the same number of copies. Finally, we present numerical evidence that the sample complexity of spectrum estimation can only improve over full state tomography by a sub-polynomial factor. Specifically, for spectrum learning with fully entangled measurements, we run simulations which suggest a lower bound of $\Omega(d^{2-\gamma})$ copies for any constant $\gamma > 0$. From this, we conclude the current best lower bound of $\Omega(d)$ is likely not tight.

Angelos Pelecanos
UC Berkeley
apelecan@berkeley.edu

Xinyu Tan
MIT
norahtan@mit.edu

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

John Wright
The University of California, Berkeley
jswright@berkeley.edu

CP22

A $(2+\varepsilon)$-Approximation Algorithm for the General Scheduling Problem in Quasipolynomial Time

We study the general scheduling problem (GSP) which generalizes and unifies several well-studied preemptive single-machine scheduling problems, such as weighted flow time, weighted sum of completion time, and minimizing the total weight of tardy jobs. We are given a set of jobs with their processing times and release times and seek to compute a (possibly preemptive) schedule for them on one machine. Each job incurs a cost that depends on its completion time in the computed schedule, as given by a separate job-dependent cost function for each job, and our objective is to minimize the total resulting cost of all jobs. The best known result for GSP is a polynomial time $O(\log\log P)$-approximation algorithm [Bansal and Pruhs, FOCS 2010, SICOMP 2014]. We give a quasi-polynomial time $(2+\varepsilon)$-approximation algorithm for GSP, assuming that the jobs' processing times are quasi-polynomially bounded integers. For the special case of the weighted tardiness objective, we even obtain an improved approximation ratio of $1+\varepsilon$. For this case, no better result had been known than the mentioned $O(\log\log P)$-approximation for the general case of GSP.

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

Lars Rohwedder
University of Southern Denmark
rohwedder@sdu.dk

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

CP22

Shortcuts and Transitive-Closure Spanners Approximation

For a directed unweighted graph $G = (V, E)$ and an integer d, a set of edges $E' \subseteq V \times V$ is called a d-TC spanner of G if the graph $H := (V, E')$ has (i) the same transitive-closure as G and (ii) diameter at most d. The set $E' \subseteq V \times V$ is a d-shortcut of G if $E \cup E'$ is a d-TC spanner of G. Our focus is on the following (α_D, α_S)-approximation algorithm: given a directed graph G and integers d and s such that G admits a d-shortcut (respectively d-TC spanner) of size

s, find a $(d\alpha_D)$-shortcut (resp. $(d\alpha_D)$-TC spanner) with $s\alpha_S$ edges, for as small α_S and α_D as possible. These problems are important special cases of graph sparsification and arise naturally in the context of reachability problems across computational models. As our main result, we show that, under the Projection Game Conjecture (PGC), there exists a small constant $\epsilon > 0$, such that no polynomial-time (n^ϵ, n^ϵ)-approximation algorithm exists for finding d-shortcuts as well as d-TC spanners of size s. Previously, super-constant lower bounds were known only for d-TC spanners with constant d and $\alpha_D = 1$ [Bhattacharyya, Grigorescu, Jung, Raskhodnikova, Woodruff 2009]. Similar lower bounds for super-constant d were previously known only for a more general case of directed spanners [Elkin, Peleg 2000].

Parinya Chalermsook
University of Sheffield
chalermsook@gmail.com

Yonggang Jiang
MPI-INF
yjiang@mpi-inf.mpg.de

Danupon Nanongkai
Max Planck Institute for Informatics
Saarland Informatics Campus
danupon@gmail.com

Sagnik Mukhopadhyay
University of Birmingham
s.mukhopadhyay@bham.ac.uk

CP22

Tree Covers of Size 2 for the Euclidean Plane

For a given metric space (P, ϕ), a tree cover of stretch t is a collection of trees on P such that edges (x, y) of trees receive length $\phi(x, y)$, and such that for any pair of points $u, v \in P$ there is a tree T in the collection such that the induced graph distance in T between u and v is at most $t\phi(u, v)$. In this paper, we show that, for any set of points P on the Euclidean plane, there is a tree cover consisting of two trees and with stretch $O(1)$. Although the problem in higher dimensions remains elusive, we manage to prove that for a slightly stronger variant of a tree cover problem we must have at least $(d+1)/2$ trees in any constant stretch tree cover in $\mathbb{R}^d$.

Andrey Kupavskii, Artur Bikeev
Moscow Institute of Physics and Technology
kupavskii@yandex.ru, bikeev99@mail.ru

Maksim Turevskii
Saint Petersburg State University
turmax20052005@gmail.com

CP22

(α, β)-Spanners and Hybrid Spanners with Nearly Tight Bounds

An (α, β)-spanner of an undirected, unweighted graph $G = (V, E)$ is a subgraph H such that for all $u, v \in V$,

$$d_H(u, v) \leq \alpha \cdot d_G(u, v) + \beta.$$

Classical results show that for any $k \in \mathbb{N}$ there exists a $(2k-1, 0)$-spanner with $O(n^{1+1/k})$ edges, conditionally optimal under Erdos' girth conjecture. However, this lower bound concerns only adjacent pairs, leaving room for better stretch on larger distances. We construct spanners nearly optimal for all pairs (v, u) of original distances $d_G(v, u) = d \leq k$. Our spanner $H \subseteq G$ has $O(n^{1+1/k} + d \log d\, n)$ edges and satisfies $d_H(u, v) \leq 2k + O(d \log d)$ for all those pairs. Thus, the multiplicative stretch for distance-d pairs is $2k/d + O(\log d)$. For $d = k/\log k$, this yields an $(O(\log k), O(k))$-spanner with $O(n^{1+1/k} + kn)$ edges, improving the $(O(k^\epsilon), O(k))$-spanner of Ben-Levy and Parter (SODA20), exponentially improving the multiplicative stertch. Finally, we improve Parters k-hybrid spanners, which guarantee multiplicative stretch $2k-1$ for neighbors and k otherwise with $O(k^2 n^{1+1/k})$ edges: our construction achieves the same guarantees with only $O(n^{1+1/k} + kn)$ edges, eliminating the quadratic dependence on k and remaining nearly tight under Erdos' conjecture.

Gur Lifshitz
Tel Aviv University
gurlifshitz@gmail.com

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

CP22

Covering the Euclidean Plane by a Pair of Trees

A t-stretch tree cover of a metric space $M = (X, \delta)$, for a parameter $t \geq 1$, is a collection of trees such that every pair of points has a t-stretch path in one of the trees. Tree covers provide an important sketching tool that has found various applications over the years. The celebrated Dumbbell Theorem by Arya et al. [STOC'95] states that any set of points in the Euclidean plane admits a $(1+\epsilon)$-stretch tree cover with $O_\epsilon(1)$ trees. This result extends to any (constant) dimension and was also generalized for arbitrary doubling metrics by Bartal et al. [ICALP'19]. Although the number of trees provided by the Dumbbell Theorem is constant, this constant is not small, even for a stretch significantly larger than $1+\epsilon$. At the other extreme, any single tree on the vertices of a regular n-polygon must incur a stretch of $\Omega(n)$. Using known results of ultrametric embeddings, one can easily get a stretch of $\tilde{O}(\sqrt{n})$ using two trees. The question of whether a low stretch can be achieved using two trees has remained illusive, even in the Euclidean plane. In this work, we resolve this fundamental question in the affirmative by presenting a constant-stretch cover with a pair of trees, for any set of points in the Euclidean plane. Our main technical contribution is a surprisingly simple Steiner construction, for which we provide a tight stretch analysis of $\sqrt{26}$.

...

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

Tianyi Zhang
ETH Zurich

tianyiz25@nju.edu.cn

CP22

Weighted Pseudorandom Generators for Read-Once Branching Programs Via Weighted Pseudorandom Reductions

We study weighted pseudorandom generators (WPRGs) and derandomizations for read-once branching programs (ROBPs). Denote n and w as the length and the width of a ROBP. We have the following results. For standard ROBPs, there exists an explicit ε-WPRG with seed length $O\left(\frac{\log n \log(nw)}{\max\{1, \log\log w - \log\log n\}} + \log w \left(\log\log\log w - \log\log \max\left\{2, \frac{\log nw}{\log n/\varepsilon}\right\}\right) + \log(1/\varepsilon)\right)$. Further, by using this result in a black-box way, we attain a WPRG for regular ROBPs with seed length $O\left(\log n\left(\sqrt{\log(1/\varepsilon)} + \log w + \log\log n\right) + \log(1/\varepsilon)\right)$. For permutation ROBPs with unbounded widths and single accept nodes, we give an explicit ε-WPRG with seed length $O\left(\log n\left(\log\log n + \sqrt{\log(1/\varepsilon)}\right) + \log(1/\varepsilon)\right)$. For regular ROBPs with $n \leq 2^{O(\sqrt{\log w})}$, $\varepsilon = 1/\text{poly} w$, we give a derandomization within space $O(\log w)$, i.e. in **L** exactly. All our results are based on iterative weighted pseudorandom reductions, which can iteratively reduce fooling long ROBPs to fooling short ones.

Kuan Cheng, Ruiyang Wu
Peking University
ckkcdh@pku.edu.cn, wuruiyang@stu.pku.edu.cn

CP23

A Truly Subcubic Combinatorial Algorithm for Induced 4-Cycle Detection

In this paper, we present the first truly subcubic, combinatorial algorithm for detecting an induced 4-cycle in a graph. The running time is $O(n^{2.84})$ on n-node graphs, thus separating the task of detecting induced 4-cycles from detecting triangles, which requires $n^{3-o(1)}$ time combinatorially under the popular Boolean Matrix Multiplication hypothesis. Significant work has gone into characterizing the time complexity of induced H-detection, relative to the complexity of detecting cliques of various sizes. Prior work identified the question of whether induced 4-cycle detection is triangle-hard as the only remaining case towards completing the lowest level of the classification [Dalirrooyfard, Vassilevska W., FOCS 2022]. Our result can be seen as a negative resolution of this question. Our algorithm deviates from previous techniques in the large body of subgraph detection algorithms and instead falls into the trendy topic of graph decompositions. While our algorithm is slower than the (non-combinatorial) state-of-the-art $\tilde{O}(n^{\omega})$-time algorithm based on polynomial identity testing [Vassilevska W., Wang, Williams, Yu, SODA 2014], combinatorial advancements often come with other benefits. In particular, we give the first nontrivial *deterministic* algorithm for detecting induced 4-cycles.

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

Shyan Akmal, Nick Fischer
Max Planck Institute for Informatics
shyan.akmal@gmail.com, nfischer@mpi-inf.mpg.de

CP23

Derandomizing Pseudopolynomial Algorithms for Subset Sum

We reexamine the classical *subset sum* problem: given a set X of n positive integers and a number t, decide whether there exists a subset of X that sums to t; or more generally, compute the set out of all numbers $y \in \{0, \ldots, t\}$ for which there exists a subset of X that sums to y. Standard dynamic programming solves the problem in $O(tn)$ time. In SODA'17, two papers appeared giving the current best deterministic and randomized algorithms, ignoring polylogarithmic factors: Koiliaris and Xu's deterministic algorithm runs in $\widetilde{O}(t\sqrt{n})$ time, while Bringmann's randomized algorithm runs in $\widetilde{O}(t)$ time. We present the first deterministic algorithm running in $\widetilde{O}(t)$ time. Our technique has a number of other applications: for example, we can also derandomize the more recent output-sensitive algorithms by Bringmann and Nakos [STOC'20] and Bringmann, Fischer, and Nakos [SODA'25] running in $\widetilde{O}(|\text{out}|^{4/3})$ and $\widetilde{O}(|\text{out}|\sqrt{n})$ time, and we can derandomize a previous fine-grained reduction from 0-1 knapsack to min-plus convolution by Cygan et al. [ICALP'17].

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

CP23

Improved Additive Approximation Algorithms for APSP

The All-Pairs Shortest Paths (APSP) is a foundational problem in theoretical computer science. Approximating APSP in undirected unweighted graphs has been studied for many years, beginning with the work of Dor, Halperin and Zwick [SICOMP'01]. Many recent works have attempted to improve these original algorithms using the algebraic tools of fast matrix multiplication. We improve on these results for the following problems. - For +2-approximate APSP, the state-of-the-art algorithm runs in $O(n^{2.259})$ time [Drr, IPL 2023; Deng, Kirkpatrick, Rong, Vassilevska Williams, and Zhong, ICALP 2022]. We give an improved algorithm in $O(n^{2.2255})$ time. - For +4 and +6-approximate APSP, we achieve time complexities $O(n^{2.1462})$ and $O(n^{2.1026})$ respectively, improving the previous $O(n^{2.155})$ and $O(n^{2.103})$ achieved by [Saha and Ye, SODA 2024]. In contrast to previous works, we do not use the big hammer of bounded-difference (min, +)-product algorithms. Instead, our algorithms are based on a simple technique that decomposes the input graph into a small number of clusters of constant diameter and a remainder of low degree vertices, which could be of independent interest in the study of shortest paths problems. We then use only standard fast matrix multiplication to obtain our improvements.

Ce Jin
EECS
MIT
cejin@berkeley.edu

Yael Kirkpatrick
Massachusetts Institute of Technology
yaelkirk@mit.edu

Michal Stawarz
ETH Zurich
m.casi321@gmail.com

Virginia Vassilevska Williams
MIT
virgi@mit.edu

CP23

Near-Linear Time Subhypergraph Counting in Bounded Degeneracy Hypergraphs

Counting small patterns in a large dataset is a fundamental algorithmic task. The most common version of this task is subgraph/homomorphism counting, wherein we count the number of occurrences of a small pattern graph H in an input graph G. The study of this problem is a field in and of itself. Recently, both in theory and practice, there has been an interest in hypergraph algorithms, where $G = (V, E)$ is a hypergraph. One can view G as a set system where hyperedges are subsets of the universe V. Counting patterns H in hypergraphs is less studied, although thereare many applications in network science and database algorithms. Inspired by advances in the graph literature, we study when linear time algorithms are possible. We focus on input hypergraphs G that have bounded degeneracy, a well-studied concept for graph algorithms. We give a spectrum of definitions for hypergraph degeneracy that cover all existing notions. For each such definition, we give a precise characterization of the patterns H that can be counted in (near) linear time. Specifically, we discover a set of "obstruction patterns'. If H does not contain an obstruction, then the number of H-subhypergraphs can be counted exactly in $O(n \log n)$ time (where n is the number of vertices in G). If H contains an obstruction, then, there is a constant $\gamma > 0$, such that there is no $o(n^{1+\gamma})$ time algorithm for counting H-subhypergraphs.

Daniel Paul-Pena, C. Seshadhri
University of California, Santa Cruz
dpaulpen@ucsc.edu, sesh@ucsc.edu

CP23

Long Arithmetic Progressions in Sparse Subset Sums: A Computational Perspective

Recently Chen, Mao and Zhang [STOC, 2025] studied arithmetic progressions from a computational perspective: instead of merely knowing the existence of arithmetic progressions, they aim to construct it explicitly and find out how its terms can be represented using integers from the corresponding set. They show that both can be done in near-linear time for long arithmetic progressions in kA, the k-fold sum of an integer set A, and $\mathcal{S}(A)$, the set of all subset sums of A, where A is a set of nonnegative integers and $|A|$ is relatively large comparing to $\max(A)$ (the largest element in A). They left as an open problem whether the same thing can be achieved for long arithmetic progressions in the sumset of different sets, i.e., $A_1+A_2+\cdots+A_k$. This paper complements the work of Chen et al. by giving the following two results: 1. A constructive proof for long arithmetic progressions in $\mathcal{S}(A)$ when $|A|$ is relatively small comparing to $\max(A)$, which gives an algorithmic version of a theorem by Erdős and Sárközy [Discrete Mathematics, 1992]. 2. A constructive proof for long arithmetic progressions in $A_1 + A_2 + \cdots + A_k$, which gives an algorithmic version of a theorem by Szemerédi and Vu [Annals of Mathematics, 2006], and answers affirmatively the open problem raised by Chen, Mao and Zhang [STOC, 2025]. Our result builds upon the recent breakthrough in Sunflower lemma.

Lin Chen
Texas Tech University
chenlin198662@gmail.com

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

CP24

Computing the Heaviest Disk and Related Problems

We present an algorithm that, given m points and n disks in $\mathbb{R}^2$, computes the disk that contains the maximum number of points. The algorithm runs in $O^*(m^{2/3}n^{2/3} + m^{32/59}n^{145/177}+m+n)$ expected time, where the $O^*(\cdot)$ notation hides factors of the form n^ϵ, for an arbitrarily small $\epsilon > 0$, and coefficients that depend on ϵ. The algorithm is faster than existing algorithms for $m < n^{5/4}$, and it has similar performance bounds for $m \geq n^{5/4}$. As a matter of fact, except for disks that are fully contained in other disks, the algorithm counts the number of input points in each disk. Our approach also gives an improved bound of $O^*(m^{2/3}n^{2/3}+m^{1/2}n^{5/6}+m+n)$ on the maximum number of incidences between m points and n *non-nested* circles in $\mathbb{R}^2$ (i.e., no circle lies in the interior of another circle).

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

Pankaj Agarwal
Duke University
pankaj@cs.duke.edu

Micha Sharir
Tel Aviv University
michas@tauex.tau.ac.il

CP24

Dynamic 3D Convex Hulls Revisited and Applications

Chan [J. ACM, 2010] developed a data structure to maintain the convex hull of a dynamic set of points in 3D under insertions and deletions to answer certain queries on the convex hulls. The algorithm has been slightly without altering the main algorithmic framework. The current best result supports each insertion in $O(\log^2 n)$ amortized time, each deletion in $O(\log^4 n)$ amortized time, and each extreme query in $O(\log^2 n)$ worst-case time (along with other query types). These results have numerous applications, notably in dynamic Euclidean nearest neighbor searching in 2D. This framework was later extended to dynamic nearest neighbor searching under general distance functions. In this paper, we revisit Chans algorithmic framework and propose a modified version that reduces the deletion time to $O(\log^3 n \log\log n)$, while retaining the $O(\log^2 n)$ insertion time at the cost of an increased query time of $O(\log^3 n/ \log\log n)$. Our result is particularly appealing in scenarios where the overall running time is dominated by the operation with the highest time complexity; in such cases, our approach offers an improvement of roughly a logarithmic factor over prior work based on Chans original framework. This improvement translates into faster algorithms for several fundamental problems, such as maintaining the dynamic 2D bichromatic closest pair, computing shortest paths in weighted disk graphs, among others.

Haitao Wang
Kahlert School of Computing, University of Utah
haitao.wang@utah.edu

CP24

A Well-Separated Pair Decomposition for Low Density Graphs

Low density graphs are considered to be a realistic graph class for modelling road networks. It has advantages over other popular graph classes for road networks, such as planar graphs, bounded highway dimension graphs, and spanners. We believe that low density graphs have the potential to be a useful graph class for road networks, but until now, its usefulness is limited by a lack of available tools. In this paper, we develop two fundamental tools for low density graphs, that is, a well-separated pair decomposition and an approximate distance oracle. We believe that by expanding the algorithmic toolbox for low density graphs, we can help provide a useful and realistic graph class for road networks, which in turn, may help explain the many efficient and practical heuristics available for road networks.

Sampson Wong, Joachim Gudmundsson
University of Sydney
sampson.wong123@gmail.com,
joachim.gudmundsson@sydney.edu.au

CP24

A Simple and Fast Reduction from Gomory-Hu Trees to Polylog Maxflows

Given an undirected graph $G = (V, E, w)$, a Gomory-Hu tree T (Gomory and Hu, 1961) is a tree on V that preserves all-pairs mincuts of G exactly. We present a simple and efficient randomized reduction from Gomory-Hu trees to polylog maxflow computations. On **unweighted** graphs, our reduction reduces to maxflow computations on graphs of total instance size $\tilde{O}(m)$ and the algorithm requires only $\tilde{O}(m)$ additional time. Our reduction is the first that is **tight up to polylog factors**. The reduction also seamlessly extends to weighted graphs; however, instance sizes and runtime increase to $\tilde{O}(n^2)$. Finally, we show how to extend our reduction to reduce Gomory-Hu trees for unweighted hypergraphs to maxflow in hypergraphs. Again, our reduction is the first that is tight up to polylog factors.

Rasmus Kyng, Maximilian Probst Gutenberg
ETH Zurich
kyng@inf.ethz.ch, maximilian.probst@inf.ethz.ch

Weixuan Yuan
ETH Zurich
ETH Zurich
weyuan@inf.ethz.ch

Wuwei Yuan
ETH Zurich
wuyuan@ethz.ch

CP24

Succinct Dynamic Rank/Select: Bypassing the Tree-Structure Bottleneck

We show how to construct a dynamic ordered dictionary, supporting insert/delete/rank/select on a set of n elements from a universe of size U, that achieves the optimal amortized expected time complexity of $O(1 + \log n/\log\log U)$, while achieving a nearly optimal space consumption of $\log\binom{U}{n} + n/2^{(\log n)^{\Omega(1)}} + \text{polylog } U$ bits in the regime where $U = \text{poly}(n)$. This resolves an open question by Pibiri and Venturini as to whether a redundancy (a.k.a. space overhead) of $o(n)$ bits is possible, and is the first dynamic solution to bypass the so-called tree-structure bottleneck, in which the bits needed to encode some dynamic tree structure are themselves enough to force a redundancy of $\tilde{\Omega}(n)$ bits. Our main technical building block is a dynamic balanced binary search tree, which we call the *compressed tabulation-weighted treap*, that itself achieves a surprising time/space tradeoff. The tree supports polylog n-time operations and requires a static lookup table of size $\text{poly}(n) + \text{polylog } U$—but, in exchange for these, the tree is able to achieve a remarkable space guarantee. Its total space redundancy is $O(\log U)$ bits. In fact, if the tree is given n and U for free, then the redundancy further drops to $O(1)$ bits.

William Kuszmaul
CMU
kuszmaul@cmu.edu

Jingxun Liang, Renfei Zhou
Carnegie Mellon University
jingxunl@andrew.cmu.edu, renfeiz@andrew.cmu.edu

CP25

A Deterministic Polylogarithmic Competitive Algorithm for Matching with Delays

In the online Min-cost Perfect Matching with Delays (MPMD) problem, m requests in a metric space are submitted at different times by an adversary. The goal is to match all requests while (i) minimizing the sum of the distances between matched pairs as well as (ii) how long each request remained unmatched after it appeared. While there exist almost optimal algorithms when the metric space is finite and known a priori, this is not the case when the metric space is infinite or unknown. In this latter case, the best known algorithm, due to Azar and Jacob-Fanani, has competitiveness $\mathcal{O}(m^{0.59})$ which is exponentially worse than the best known lower bound of $\Omega(\log m/\log\log m)$ by Ashlagi et al. We present a $\mathcal{O}(\log^5 m)$-competitive algorithm for the MPMD problem. This algorithm is deterministic and does not need to know the metric space or m in advance. This is an exponential improvement over previous results and only a polylogarithmic factor away from the lower bound.

Marc Dufay, Roger Wattenhofer
ETH Zurich
mdufay@ethz.ch, wattenhofer@ethz.ch

CP25

Selfish, Local and Online Scheduling via Vector Fitting

We present a unified dual fitting technique on a semidefinite program yielding simple proofs of tight bounds for the robust price of anarchy of several congestion and scheduling games under the sum of weighted completion times objective. The same approach also allows to bound the

approximation ratio of local search algorithms and the competitive ratio of online algorithms for the scheduling problem $R||\sum w_jC_j$. As our main application, we show that the known coordination ratio bounds of respectively $4, (3+\sqrt{5})/2 \approx 2.618$, and $32/15 \approx 2.133$ for the scheduling game $R||\sum w_jC_j$ under the coordination mechanisms Smith's Rule, Proportional Sharing and Rand (STOC 2011) can be extended to congestion games and obtained through this approach. As a further application of this technique in a game theoretic setting, we show that it recovers the tight bound of $(3+\sqrt{5})/2$ for the price of anarchy of weighted affine congestion games and the Kawaguchi-Kyan bound of $(1+\sqrt{2})/2$ for the pure price of anarchy of $P||\sum w_jC_j$. Moreover, this approach allows to analyze a simple local search algorithm for $R||\sum w_jC_j$, the best known combinatorial approximation algorithm for this problem achieving an approximation ratio of $(5+\sqrt{5})/4+\varepsilon \approx 1.809+\varepsilon$, and an online algorithm which is 4-competitive.

Danish Kashaev
Centrum Wiskunde & Informatica, Amsterdam
danich.1997@gmail.com

CP25

Hardness of Approximation for Shortest Path with Vector Costs

We obtain hardness of approximation results for the ℓ_p-Shortest Path problem, a variant of the classic Shortest Path problem with vector costs. For every integer $p \in [2,\infty)$, we show a hardness of $\Omega(p(\log n/\log^2\log n)^{1-1/p})$ for both polynomial- and quasi-polynomial-time approximation algorithms. This nearly matches the approximation factor of $O(p(\log n/\log\log n)^{1-1/p})$ achieved by a quasi-polynomial-time algorithm of Makarychev, Ovsiankin, and Tani (ICALP 2025). No hardness of approximation results were previously known for any $p < \infty$. We also present results for the case where p is a function of n. For $p = \infty$, we establish a hardness of $\tilde{\Omega}(\log^2 n)$, improving upon the previous $\tilde{\Omega}(\log n)$ hardness result. Our result nearly matches the $O(\log^2 n)$ approximation guarantee of the quasi-polynomial-time algorithm by Li, Xu, and Zhang (ICALP 2025). Finally, we present asymptotic bounds on higher-order Bell numbers, which might be of independent interest.

Ron Mosenzon
Toyota Technological Institute of Science
ron.mosenzon@gmail.com

Charlie Carlson
State University of New York at Buffalo
cc387@buffalo.edu

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

CP25

New Algorithms and Hardness Results for Robust Satisfiability of (Promise) CSPs

In this paper, we continue the study of robust satisfiability of promise CSPs (PCSPs), initiated in (Brakensiek, Guruswami, Sandeep, STOC 2023), and obtain the following results: For the PCSP 1-in-3-SAT vs NAE-SAT with negations, we prove that it is hard, under the Unique Games conjecture (UGC), to satisfy $1-\Omega(1/\log(1/\epsilon))$ constraints in a $(1-\epsilon)$-satisfiable instance. For any Boolean PCSP that admits Majority polymorphisms, we give an algorithm satisfying $1-O(\sqrt{\epsilon})$ fraction of the weaker constraints when promised the existence of an assignment satisfying $1-\epsilon$ fraction of the stronger constraints. This significantly generalizes the CharikarMakarychevMakarychev algorithm for 2-SAT, and matches the optimal trade-off possible under the UGC. The algorithm also extends, with the loss of an extra $\log(1/\epsilon)$ factor, to PCSPs on larger domains with a certain structural condition. We prove that assuming the UGC, robust satisfiability is preserved under the addition of equality constraints. As a consequence, we can extend the rich algebraic techniques for decision/search PCSPs to robust PCSPs. The methods involve the development of a correlated and robust version of the general SDP rounding algorithm for CSPs due to (Brown-Cohen, Raghavendra, ICALP 2016), which might be of independent interest.

Joshua Brakensiek
UC Berkeley
josh.brakensiek@berkeley.edu

Lorenzo Ciardo
TU Graz
lorenzo.ciardo@tugraz.at

Venkatesan Guruswami
UC Berkeley
venkatg@berkeley.edu

Aaron Potechin
University of Chicago
potechin@uchicago.edu

Stanislav Zivny
University of Oxford, United Kingdom
standa.zivny@cs.ox.ac.uk

CP25

Combinatorial Philosopher Inequalities

In online combinatorial allocation, agents arrive sequentially and items are allocated in an online manner. The algorithm designer only knows the distribution of each agent's valuation, while the actual realization of the valuation is revealed only upon her arrival. Against the offline benchmark, Feldman, Gravin, and Lucier (SODA 2015) designed an optimal 0.5-competitive algorithm for XOS agents. An emerging line of work focuses on designing approximation algorithms against the (computationally unbounded) optimal online algorithm. The primary goal is to design algorithms with approximation ratios strictly greater than 0.5, surpassing the impossibility result against the offline optimum. Positive results are established for unit-demand agents (Papadimitriou, Pollner, Saberi, Wajc, MOR 2024), and for k-demand agents (Braun, Kesselheim, Pollner, Saberi, EC 2024). In this paper, we extend the existing positive results for agents with submodular valuations by establishing a 0.5+O(1) approximation against a newly constructed online configuration LP relaxation for the combinatorial allocation setting. Meanwhile, we provide negative results for agents with XOS valuations by providing a 0.5 integrality gap for the online configuration LP, showing an obstacle of existing approaches.

Yifan Wang
Georgia Tech
ywang3782@gatech.edu

Enze Sun
The University of Hong Kong
sunenze@connect.hku.hk

Zhihao Gavin Tang
Shanghai University of Finance and Economics
tang.zhihao@mail.shufe.edu.cn

CP25

PageRank Centrality in Directed Graphs with Bounded In-Degree

We study the complexity of locally estimating a node's PageRank centrality in a directed graph G. For any node t, its PageRank centrality $\pi(t)$ is defined as the probability that a random walk in G, starting from a uniformly chosen node, terminates at t, where each step terminates with a constant probability α. To obtain a multiplicative $(1 \pm O(1))$-approximation of $\pi(t)$ with probability $\Omega(1)$, the previously best upper bound is $O\left(n^{1/2} \min\{\Delta_{in}^{1/2}, \Delta_{out}^{1/2}, m^{1/4}\}\right)$, where n and m denote the number of nodes and edges in G, and Δ_{in} and Δ_{out} upper bound the in-degrees and out-degrees of G, respectively. Using a refinement of the proof in the same paper, we establish a lower bound of $\Omega\left(n^{1/2} \min\left\{\Delta_{in}^{1/2}/n^{\gamma}, \Delta_{out}^{1/2}/n^{\gamma}, m^{1/4}\right\}\right)$, where $\gamma = \frac{1}{2}\left(2\max\left\{\log_{1/(1-\alpha)} \Delta_{in}, 1\right\} - 1\right)^{-1}$. There remains a gap of $\Omega\left(n^{\gamma}\right)$ when considering the maximum in-degree Δ_{in}, and this gap is large when Δ_{in} is small. In this paper, we present a new algorithm that achieves the above lower bound (up to logarithmic factors). Our key technique is a novel randomized backwards propagation process that only propagates selectively based on Monte Carlo estimated PageRank scores.

Mikkel Thorup
University of Copenhagen
mikkel2thorup@gmail.com

Hanzhi Wang
BARC, University of Copenhagen
hanzhi.hzwang@gmail.com

Zhewei Wei, Mingji Yang
Renmin University of China
zhewei@ruc.edu.cn, kyleyoung@ruc.edu.cn

CP26

Faster Algorithms for Packing Forests in Graphs and Related Problems

We consider several problems related to packing forests in graphs. The first one is to find k edge-disjoint forests in a directed graph G of maximal size such that the indegree of each vertex in these forests is at most k. We describe a min-max characterization for this problem and show that it can be solved in almost linear time for fixed k, extending the algorithm of [Gabow, 1995]. Specifically, the complexity is $O(k\delta m \log n)$, where n, m are the number of vertices and edges in G respectively, and $\delta = \max\{1, k - k_G\}$, where k_G is the edge connectivity of the graph. Using our solution to this problem, we improve complexities for two existing applications: (1) k-forest problem: find k forests in an undirected graph G maximizing the number of edges in their union. We show how to solve this problem in $O(k^3 \min\{kn, m\} \log^2 n + k \cdot \text{MAXFLOW}(m, m) \log n)$ time, breaking the $O_k(n^{3/2})$ complexity barrier of previously known approaches. (2) Directed edge-connectivity augmentation problem: find a smallest set of directed edges whose addition to the given directed graph makes it strongly k-connected. We improve the deterministic complexity for this problem from $O(k\delta(m + \delta n) \log n)$ [Gabow, STOC 1994] to $O(k\delta m \log n)$. A similar approach with the same complexity also works for the undirected version of the problem.

Pavel Arkhipov
Institute of Science and Technology Austria
pavel.arkhipov@ist.ac.at

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

CP26

Burling Graphs in Graphs with Large Chromatic Number

A graph class is χ-bounded if the only way to force large chromatic number in graphs from the class is by forming a large clique. The class of string graphs is not χ-boundedPawlik et al. (2014) showed another way of forcing large chromatic number in this class, by triangle-free graphs B_k with $\chi(B_k) = k$ constructed by Burling (1965). This disproved the celebrated conjecture of Scott (1997) that classes of graphs excluding induced subdivisions of a fixed graph are χ-bounded. We prove that in broad classes of graphs excluding induced subdivisions of a fixed graph, including the increasingly more general classes of string graphs, region intersection graphs, and hereditary classes of graphs with finite asymptotic dimension, large chromatic number can be forced only by large cliques or large graphs B_k. One corollary is that the hereditary closure of $\{B_k : k \geq 1\}$ forms a minimal hereditary graph class with unbounded chromatic numberthe second known graph class with this property after the class of complete graphs. Another corollary is that the decision variant of approximate coloring in the aforesaid graph classes can be solved in polynomial time by exhaustively searching for a sufficiently large clique or copy of B_k. We also discuss how our results can be turned into polynomial-time algorithms for the search variant of approximate coloring in string graphs (with intersection model in the input) and other aforesaid graph classes.

Tara Abrishami
Stanford University
tara.abrishami@stanford.edu

Marcin Brianski
Jagiellonian University
marcin.brianski@doctoral.uj.edu.pl

James Davies
Leipzig University
jgdavies@uwaterloo.ca

Xiying Du
Georgia Institute of Technology
xdu90@gatech.edu

Jana Masaríková
University of Warsaw
jnovotna@mimuw.edu.pl

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

Bartosz Walczak
Jagiellonian University
bartosz.m.walczak@gmail.com

CP26

On Independent Spanning Trees in Random Graphs

A central challenge in network design is ensuring resilience: how can we guarantee multiple, independent, communication pathways between nodes, even when some connections fail in a network? In 1989, Zehavi and Itai formulated a graph-theoretic conjecture that captures the essence of this problem. They proposed that any k-vertex-connected graph contains k independent spanning trees rooted at any given root r, which means that for every vertex v in the graph, the unique r–v paths within these k spanning trees are entirely disjoint, apart from their endpoints r and v. Despite decades of effort, this conjecture has only been proven for $k \leq 4$ and for specific graph families using their underlying topological structure, leaving the general case as an open problem in graph theory with substantial consequences in the field of distributed algorithms. We make significant progress on the Zehavi-Itai conjecture by proving it holds for almost all graphs of relevant densities. More precisely, we show that there exists some constant $C > 1$ such that for all $C \log n/n \leq p < 0.99$, the binomial random graph $G(n,p)$ contains a family of $\delta(G)$ independent spanning trees rooted at any given vertex r with high probability. Note that the lower bound on p up to the constant C matches the standard threshold for connectivity for $G(n,p)$, thus we establish an essentially best possible result for random graphs.

Keith Frankston
Rutgers University
k.frankston@fastmail.com

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

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

Alexey Pokrovskiy
University College London
a.pokrovskiy@ucl.ac.uk

Liana Yepremyan
Emory University
lyeprem@emory.edu

CP26

All-Pairs Minimum Cut Using $\tilde{O}(n^{7/4})$ Cut Queries

We present the first non-trivial algorithm for the all-pairs minimum cut problem in the cut-query model. Given cut-query access to an unweighted graph $G = (V,E)$ with n vertices, our randomized algorithm constructs a Gomory-Hu tree of G, and thus solves the all-pairs minimum cut problem, using $\tilde{O}(n^{7/4})$ cut queries.

Yotam Kenneth-Mordoch, Robert Krauthgamer
Weizmann Institute of Science
yotam.kenneth@weizmann.ac.il,
robert.krauthgamer@weizmann.ac.il

CP26

History-Independent Maximal Matchings can Be Surprisingly Efficient, and Lead to Better Worst-Case Guarantees

One of the most basic problems in dynamic graph algorithms is to maintain a maximal matching as edges are inserted and deleted over time. In a line of work started by Baswana, Gupta and Sen, and then completed by Solomon, it was shown how to solve this problem in amortized expected time $O(1)$ per update. (Interestingly, achieving worst-case expected $O(1)$ remains open.) Recent work by Behnezhad et al. suggests that very simple alternative algorithm may also perform well: simply maintain a randomized greedy matching. They show that this solution achieves $O(\log^4 n)$ expected time per update. In this paper, we show that the randomized greedy solution is even much more efficient than Behnezhad et al.'s result might seem to suggest, achieving an expected time bound of $O((\log \log n)^2)$ per update. This new bound is within a quadratic factor of tight. Our result comes with several consequences. We obtain a nearly exponential time improvement for maintaining a history-independent dynamic maximal matching; a nearly exponential time improvement for maintaining a dynamic maximal matching with **worst-case** expected time guarantees; and a polynomial improvement for maintaining a dynamic maximal matching with **high-probability** time guarantees. Our result on history independence is of special interest given the major push in the data-structures community in recent years to study history independence as a privacy guarantee for data structures.

Rathish Das
University of Houston
rathish@uh.edu

William Kuszmaul
CMU
kuszmaul@cmu.edu

CP26

Sensitivity Lower Bounds for Approximation Algorithms

Sensitivity measures how much the output of an algorithm changes, in terms of Hamming distance, when part of the input is modified. While approximation algorithms with low sensitivity have been developed for many problems, no sensitivity lower bounds were previously known for approximation algorithms. In this work, we establish the first polynomial lower bound on the sensitivity of (randomized) approximation algorithms for constraint satisfaction problems (CSPs) by adapting the probabilistically checkable proof (PCP) framework to preserve sensitivity lower bounds. From this, we derive polynomial sensitivity lower bounds for approximation algorithms for a variety of problems, including maximum clique, minimum vertex cover, and maximum cut. Leveraging the connection between sensitivity and locality in the non-signaling model, which subsumes the LOCAL, quantum-LOCAL, and bounded dependence models, we establish locality lower bounds for

several graph problems in the non-signaling model.

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

Noah Fleming
Lund University & Columbia University
noah.fleming@cs.lth.se

CP27

Optimal Mass Estimation in the Conditional Sampling Model

The conditional sampling model, introduced by Cannone, Ron, Servedio (SODA 14, SIAM J. Comput. 15) and independently by Chakraborty, Fischer, Goldhirsh, Matsliah (ITCS 13, SIAM J. Comput. 16), is a common framework for a number of studies concerning strengthened models of distribution testing. A core task in these investigations is that of estimating the mass of individual elements. The above mentioned works, and the improvement of Kumar, Meel, Pote (AISTATS 25), provided polylogarithmic algorithms for this task. In this work we shatter the polylogarithmic barrier, and provide an estimator for the mass of individual elements that uses only $O(\log\log N + poly(1/\epsilon))$ conditional samples. We complement this result with an $\Omega(\log\log N)$ lower bound. We then show that our mass estimator provides an improvement (and in some cases a unifying framework) for a number of related tasks, such as testing by learning of any label-invariant property, and distance estimation between two (unknown) distributions. In light of some known lower bounds for common restricted models, our results imply that the full power of the conditional model is indeed required for the doubly-logarithmic upper bound. Finally, we exponentially improve the previous lower bound on testing by learning of label-invariant properties from double-logarithmic to $\Omega(\log N)$ conditional samples, whereas our testing by learning algorithm provides an upper bound of $O(poly(1/\epsilon)\log N \log\log N)$.

Tomer Adar
Technion
tomer-adar@campus.technion.ac.il

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

Amit Levi
University of Haifa
alevi@cs.haifa.ac.il

CP27

Balls and Bins and the Infinite Process with Random Deletions

We consider an infinite balls-into-bins process with deletions where in each discrete step t a coin is tossed as to whether, with probability $\beta(t) \in (0,1)$, a new ball is allocated using the Greedy[2] strategy or, with remaining probability $1-\beta(t)$, a ball is deleted from a non-empty random bin. Let n be the number of bins. We are interested in bounding the discrepancy current maximum load relative to current average) and the overload (current maximum load relative to highest average observed so far). We prove that at an arbitrarily chosen time t the total number of balls above the average is $O(n)$ and that the discrepancy is $O(\log(n))$. For the discrepancy, we provide a matching lower bound. Furthermore we prove that at an arbitrarily chosen time t the overload is $\log\log(n)+O(1)$. For "good" insertion probability sequences we show that even the discrepancy is bounded by $\log\log(n)+O(1)$. One of our main analytical tools is a layered induction, as per [Azar, Broder, Karlin, Upfal; Balanced Allocations]. Since our model allows for rather more general scenarios than what was previously considered, the formal analysis requires some extra ingredients as well, in particular a detailed potential analysis. Furthermore, we simplify the setup by applying probabilistic couplings to obtain certain "recovery" properties, which eliminate much of the need for intricate and careful conditioning elsewhere in the analysis.

Peter Kling
University of Hamburg
peter.kling@h-da.de

Petra Berenbrink
Universität Hamburg
petra.berenbrink@uni-hamburg.de

Tom Friedetzky
Durham University
tom.friedetzky@durham.ac.uk

Lars Nagel
Loughborough University
l.nagel@lboro.ac.uk

CP27

Weighted k-Server Admits An Exponentially Competitive Algorithm

The weighted k-server is a variant of the k-server problem, where the cost of moving a server is the server's weight times the distance through which it moves. The problem is famous for its intriguing properties and for evading standard techniques for designing and analyzing online algorithms. Even on uniform metric spaces with sufficiently many points, the deterministic competitive ratio of weighted k-server is known to increase doubly exponentially with respect to k, while the behavior of its randomized competitive ratio is not fully understood. Specifically, no upper bound better than doubly exponential is known, while the best known lower bound is singly exponential in k. In this paper, we close the exponential gap between these bounds by giving an $\exp(O(k^2))$-competitive randomized online algorithm for the weighted k-server problem on uniform metrics, thus breaking the doubly exponential barrier for deterministic algorithms for the first time. This is achieved by a recursively defined notion of a *phase* which, on the one hand, forces a lower bound on the cost of any offline solution, while, on the other hand, also admits a randomized online algorithm with bounded expected cost. The algorithm is also recursive; it involves running several algorithms virtually and in parallel and following the decisions of one of them in a random order.

Ashish Chiplunkar
Department of Computer Science and Engineering
Indian Institute of Technology Delhi
ashishc@iitd.ac.in

Adithya Bijoy
National University of Singapore
bijoyadithya02@gmail.com

Ankit Mondal
Indian Institute of Technology Delhi
ankmon25@gmail.com

CP27

Phase Transition of the Sinkhorn-Knopp Algorithm

Matrix scaling via the SinkhornKnopp algorithm is widely used and typically attains high-quality approximations in few iterations, yet theory still gives only polynomial-in-ε upper bounds on the iteration count, and sharp lower bounds remain scarce. For an $n \times n$ matrix, its normalized version is obtained by dividing each entry by the largest entry in the matrix. We say that a normalized matrix has a density γ if there exists a constant $\rho > 0$ such that one row or column has exactly $\lceil \gamma n \rceil$ entries with values at least ρ, and every other row and column has at least $\lceil \gamma n \rceil$ such entries. For the upper bound, we show that the SinkhornKnopp algorithm produces a nearly doubly stochastic matrix in $O(\log n - \log \varepsilon)$ iterations for all nonnegative square matrices whose normalized version has a density $\gamma > 1/2$. For the lower bound, we establish a tight bound of $\Omega(\sqrt{n}/\varepsilon)$ iterations for positive matrices under the ℓ_2-norm error measure. Moreover, for every $\gamma < 1/2$, there exists a matrix with density γ for which the SinkhornKnopp algorithm requires $\Omega(\sqrt{n}/\varepsilon)$ iterations. In summary, our results reveal a sharp phase transition in the SinkhornKnopp algorithm at the density threshold $\gamma = 1/2$.

Kun He
DEKE
Renmin University of China
hekun.threebody@foxmail.com

CP27

Dichotomy for Orderings?

We prove that ordering problems for graphs defined by finitely many forbidden ordered subgraphs capture the full power of the class NP, that is, any language in the class NP is polynomially equivalent to an ordering problem. In particular, we refute a conjecture of Hell, Mohar and Rafiey that dichotomy holds for this class. An interesting feature appeared: while the full power is reached by disconnected structures, and one can even guarantee the connectivity of all patterns, this is no longer the case for biconnected patterns. We confirm the Hell-Mohar-Rafiey conjecture if all the patterns are biconnected. Duffus, Ginn and Rdl conjectured that ordering problems defined by a single obstruction which is a biconnected ordered graph are NP-complete if the graph is not complete. We settle their conjecture. We prove a general phenomenon: for finite sets of biconnected patterns (which may be colored structures or ordered structures) dichotomy holds, while for general patterns we have full power. A principal tool for obtaining these results is the Sparse Incomparability Lemma in many of its variants, which are classical results in the theory of homomorphisms of graphs and structures. We prove it for ordered stuctures as a Temporal Sparse Incomparability Lemma. We confirm the Duffus-Ginn-Rdl conjecture by bringing together most of the techniques developed in the paper, and we also use the results of Bodirsky and Kra on the complexity of temporal CSPs.

Gabor Kun
Hungarian Academy of Sciences
kungabor@renyi.hu

Jaroslav Nesetril
department of Applied Mathematics
Charles University
nesetril@kam.mff.cuni.cz

CP27

(Almost) Perfect Discrete Iterative Load Balancing

We consider discrete, iterative load balancing via matchings on arbitrary graphs. Initially each node holds a certain number of tokens (load), and the objective is to redistribute the tokens such that eventually each node has approximately the same number of tokens. We present results for a general class of simple local balancing schemes where the tokens are balanced via matchings. In each round the process averages the tokens of any two matched nodes. If the sum of their tokens is odd, the node to receive the one excess token is selected at random. Our class covers three popular models: in the matching model a new matching is generated randomly in each round, in the balancing circuit model a fixed sequence of matchings is applied periodically, and in the asynchronous model the load is balanced over a randomly chosen edge. We measure the quality of a load vector by its discrepancy, defined as the difference between the maximum and minimum load across all nodes. As our main result we show that with high probability our discrete balancing scheme reaches a discrepancy of 3 in a number of rounds which asymptotically matches the spectral bound for continuous load balancing with fractional load. The result demonstrates that in the general model we consider, discrete load balancing is no harder than continuous load balancing.

Hamed Hosseinpour, Petra Berenbrink
Universität Hamburg
hamed.hosseinpour@uni-hamburg.de,
petra.berenbrink@uni-hamburg.de

Robert Elsaesser
University of Salzburg
elsa@cs.sbg.ac.at

Tom Friedetzky
Durham University
tom.friedetzky@durham.ac.uk

Dominik Kaaser
Technical University of Hamburg
dominik.kaaser@tuhh.de

Peter Kling
University of Hamburg
peter.kling@h-da.de

Thomas Sauerwald
University of Cambridge
tms41@cam.ac.uk

CP28

Faster Distributed Delta-Coloring Via a Reduction to MIS

Recent improvements on the deterministic complexities of fundamental graph problems in the LOCAL model of distributed computing have yielded state-of-the-art upper bounds of $\tilde{O}(\log^{5/3} n)$ rounds for maximal independent set (MIS) and $(\Delta + 1)$-coloring [Ghaffari, Grunau, FOCS'24] and $\tilde{O}(\log^{19/9} n)$ rounds for the more restrictive

Δ-coloring problem [Ghaffari, Kuhn, FOCS'21; Ghaffari, Grunau, FOCS'24; Bourreau, Brandt, Nolin, STOC'25]. In our work, we show that Δ-coloring can be solved deterministically in $\tilde{O}(\log^{5/3} n)$ rounds as well, matching the currently best bound for $(\Delta+1)$-coloring. We achieve our result by developing a reduction from Δ-coloring to MIS that guarantees that the (asymptotic) complexity of Δ-coloring is at most the complexity of MIS, unless MIS can be solved in sublogarithmic time, in which case, due to the $\Omega(\log n)$-round Δ-coloring lower bound from [BFHKLRSU, STOC'16], our reduction implies a tight complexity of $\Theta(\log n)$ for Δ-coloring. In particular, any improvement on the complexity of the MIS problem will yield the same improvement for the complexity of Δ-coloring (up to the true complexity of Δ-coloring). Our reduction also yields improvements for Δ-coloring in the randomized LOCAL model and when complexities are parameterized by both n and Δ.

Yann Bourreau, Sebastian Brandt, Alexandre Nolin
CISPA Helmholtz Center for Information Security
yann.bourreau@cispa.de, brandt@cispa.de, alexandre.m.p.nolin@gmail.com

CP28

Optimal Subspace Embeddings: Resolving Nelson-Nguyen Conjecture Up to Sub-Polylogarithmic Factors

We give a proof of the conjecture of Nelson and Nguyen [FOCS 2013] on the optimal dimension and sparsity of oblivious subspace embeddings, up to sub-polylogarithmic factors: For any $n \geq d$ and $\varepsilon \geq d^{-O(1)}$, there is a random $\tilde{O}(d/\varepsilon^2) \times n$ matrix Π with $\tilde{O}(\log(d)/\varepsilon)$ non-zeros per column such that for any $A \in \mathbb{R}^{n\times d}$, with high probability, $(1-\varepsilon)\|Ax\| \leq \|\Pi Ax\| \leq (1+\varepsilon)\|Ax\|$ for all $x \in \mathbb{R}^d$, where $\tilde{O}(\cdot)$ hides only sub-polylogarithmic factors in d. Our result in particular implies a new fastest sub-current matrix multiplication time reduction of size $\tilde{O}(d/\varepsilon^2)$ for a broad class of $n \times d$ linear regression tasks. A key novelty in our analysis is a matrix concentration technique we call *iterative decoupling*, which we use to fine-tune the higher-order trace moment bounds attainable via existing random matrix universality tools [Brailovskaya and van Handel, GAFA 2024].

Shabarish Chenakkod
University of Michigan
Ann Arbor
shabari@umich.edu

Michal Derezinski
University of Michigan
derezin@umich.edu

Xiaoyu Dong
National University of Singapore
Department of Mathematics
xdong@nus.edu.sg

CP28

Spectral Clustering with Side Information

In the graph clustering problem with a planted solution, the input is a graph on n vertices partitioned into k clusters, and the task is to recover these clusters from graph structure. A standard assumption is that clusters induce $\Omega(1)$-expanders and form ϵ-sparse cuts. The planted solution is then unique up to $\approx \epsilon$ misclassification rate, and efficient recovery algorithms achieving this rate are known. In practice, however, vertices of the graph are typically equipped with labels, or features, providing additional clues about their cluster membership - for example, a cluster label corrupted independently with probability δ. Using either the graph or labels alone leads to misclassification rate $\min\{\epsilon, \delta\}$; can the two be combined to achieve $\approx \epsilon\delta$? We give an affirmative answer, showing that this misclassification rate can be achieved in sublinear time in n. Our key insight is a new observation on "spectrally ambiguous' vertices in a well-clusterable graph. While our classifier achieves a nearly optimal misclassification rate, the recovered clusters do not necessarily induce expanders. In our second result, we give a polynomial-time edge-reweighting algorithm that transforms a $(k, \epsilon, \Omega(1))$-clusterable graph into a $(k, \widetilde{O}(\epsilon\delta), \Omega(1))$-clusterable one (for constant k), nearly optimally improving cut sparsity while preserving cluster expansion.

Ekaterina Kochetkova
EPFL
ekaterina.kochetkova@epfl.ch

Hendrik Fichtenberger
Google Research
fichtenberger@google.com

Michael Kapralov
EPFL
michael.kapralov@epfl.ch

Silvio Lattanzi
Google Zurich
silviol@google.com

Davide Mazzali, Weronika Wrzos-Kaminska
EPFL
davide.mazzali@epfl.ch, weronika.wrzos-kaminska@epfl.ch

CP28

Near-Optimal Four-Cycle Counting in Graph Streams

We study four-cycle counting in arbitrary order graph streams. We present a 3-pass algorithm for $(1 + \epsilon)$-approximating the number of four-cycles using $\widetilde{O}(m/\sqrt{T})$ space, where m is the number of edges and T the number of four-cycles in the graph. This improves upon a 3-pass algorithm by Vorotnikova using space $\widetilde{O}(m/T^{1/3})$ and matches a multi-pass lower bound of $\Omega(m/\sqrt{T})$ by McGregor and Vorotnikova.

Sebastian Lüderssen, Stefan Neumann
TU Wien
sebastian.luederssen@tuwien.ac.at,
stefan.neumann@tuwien.ac.at

Pan Peng
University of Science and Technology of China
ppeng@ustc.edu.cn

CP28

L_p Sampling in Distributed Data Streams with Ap-

plications to Adversarial Robustness

In the distributed monitoring model, a data stream over a universe of size n is distributed over k servers, who must continuously provide certain statistics of the overall dataset, while minimizing communication with a central coordinator. In such settings, collecting a random sample from the global stream is a powerful primitive. Of particular interest is the task of producing a perfect L_p sample, which given a frequency vector $f \in \mathbb{R}^n$, outputs an index i with probability $\frac{f_i^p}{\|f\|_p^p} + \frac{1}{\mathrm{poly}(n)}$. In this paper, we resolve the problem of perfect L_p sampling for all $p \geq 1$ in the distributed monitoring model. Specifically, our algorithm runs in $k^{p-1} \cdot \mathrm{polylog}(n)$ bits of communication, which is optimal up to polylogarithmic factors. Utilizing our perfect L_p sampler, we achieve adversarially-robust distributed monitoring protocols for the F_p moment estimation problem, where the goal is to provide a $(1+\varepsilon)$-approximation to $f_1^p + \ldots + f_n^p$. Our algorithm uses $\frac{k^{p-1}}{\varepsilon^2} \cdot \mathrm{polylog}(n)$ bits of communication for all $p \geq 2$ and achieves optimal bounds up to polylogarithmic factors, matching lower bounds by Woodruff and Zhang (STOC 2012) in the non-robust setting. Finally, we achieve near-optimal adversarially robust distributed protocols for central problems such as counting, frequency estimation, heavy-hitters, and distinct element estimation.

Honghao Lin
Carnegie Mellon University
honghaol@andrew.cmu.edu

Zhao Song
University of California, Berkeley
magic.linuxkde@gmail.com

David Woodruff
Carnegie Mellon University
dwoodruf@cs.cmu.edu

Shenghao Xie
Texas A&M University
xsh1302@gmail.com

Samson Zhou
Carnegie Mellon University
samsonzhou@gmail.com

CP28

Distribution Testing in the Presence of Arbitrarily Dominant Noise with Verification Queries

We study distribution testing without direct access to a source of relevant data, but rather to one where only a tiny fraction is relevant. To enable this, we introduce the following verification query model. The goal is to perform a statistical task on distribution $\boldsymbol{p}$ given sample access to a mixture $\boldsymbol{r} = \lambda\boldsymbol{p} + (1-\lambda)\boldsymbol{q}$ and the ability to query whether a sample was generated by $\boldsymbol{p}$ or by $\boldsymbol{q}$. In general, if m_0 samples from $\boldsymbol{p}$ suffice for a task, then $O(m_0/\lambda)$ samples and queries always suffice in our model. Are there tasks for which the number of queries can be significantly reduced? We study canonical problems in distribution testing. For all $m \leq n$, we obtain (i) a uniformity and identity tester using $O(m + \sqrt{n}\varepsilon^{-2}\lambda^{-1})$ samples and $O(nm^{-1}\varepsilon^{-4}\lambda^{-2})$ queries, and (ii) a closeness tester using $O(m + n^{2/3}\varepsilon^{-4/3}\lambda^{-1} + \varepsilon^{-4}\lambda^{-3})$ samples and $O(n^2m^{-2}\varepsilon^{-4}\lambda^{-3})$ queries. Moreover, we show that these query complexities are tight for all testers using $m \ll n$ samples. Next, we show that for all three problems using $\tilde{O}(n\varepsilon^{-2}\lambda^{-1})$ samples we can achieve query complexity $\tilde{O}(\varepsilon^{-2}\lambda^{-1})$ which is nearly optimal even for the basic task of bias estimation with unbounded samples.

Hadley Black, Christopher Ye
University of California, San Diego
hadleyblack92@gmail.com, czye@ucsd.edu

CP29

Tight Parameterized (In)tractability of Layered Crossing Minimization: Subexponential Algorithms and Kernelization

The starting point of our work is a decade-old open question concerning the subexponential parameterized complexity of 2-Layer Crossing Minimization. In this problem, the input is an n-vertex graph G whose vertices are partitioned into two independent sets V_1 and V_2, and a non-negative integer k. The question is whether G admits a 2-layered drawing with at most k crossings, where each V_i lies on a distinct line parallel to the x-axis, and all edges are straight lines. We resolve this open question by giving the first subexponential fixed-parameter algorithm for this problem, running in time $2^{O(\sqrt{k}\log k)} + n \cdot k^{O(1)}$. We then ask whether the subexponential phenomenon extends beyond two layers. In the general h-Layer Crossing Minimization problem, the goal is to decide whether an h-layered drawing with at most k crossings exists. We present a subexponential FPT algorithm with running time $2^{O(k^{2/3}\log k)} + n \cdot k^{O(1)}$ for $h = 3$ layers. In contrast, we show that for all $h \geq 5$, no algorithm with running time $2^{o(k/\log k)} \cdot n^{O(1)}$ exists unless the Exponential-Time Hypothesis fails. Finally, we address polynomial kernelization. While a polynomial kernel was already known for $h = 2$, we design a new polynomial kernel for $h = 3$. These kernels are essential ingredients in our subexponential algorithms. Finally, we rule out polynomial kernels for all $h \geq 4$ unless the polynomial hierarchy collapses.

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

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

Tanmay Inamdar
Indian Institute of Technology Jodhpur
taninamdar@gmail.com

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

Meirav Zehavi
Ben-Gurion University
zehavimeirav@gmail.com

CP29

Catching Rats in H-Minor-Free Graphs

We show that every H-minor-free graph excluding a $(k \times k)$-grid as a minor has treewidth/branchwidth bounded from above by a function $f(t, k)$ that is linear in k and poly-

nomial in $t := |V(H)|$. Such a result was proven originally by [Demaine & Hajiaghayi, Combinatorica, 2008], where f was indeed linear in k. However the dependency in t in this result was non-explicit. Later, [Kawarabayashi & Kobayashi, JCTB, 2020] showed that this bound can be estimated to be $f(t,k) \in 2^{\mathcal{O}(t \log t)} \cdot k$. Wood recently asked whether f can be pushed further to be polynomial, while maintaining the linearity on k. We answer this in a strong sense, by showing that the treewidth/branchwidth of G is in $\mathcal{O}(gk + t^{2304})$, where g is the Euler genus of H. This yields $f(t,k) = \mathcal{O}(t^2 k + t^{2304})$. We build on techniques for branchwidth and new bounds and insights for the Graph Minor Structure Theorem (GMST) due to [Gorsky, Seweryn & Wiederrecht, arXiv, 2025]. We employ our methods to provide approximation algorithms for the treewidth/branchwidth of H-minor-free graphs. For every $\varepsilon > 0$ and every t-vertex graph H with Euler genus g, we give a $(g+\varepsilon)$-approximation algorithm for the branchwidth of H-minor-free graphs running in $2^{\mathbf{poly}(t)/\varepsilon} \cdot \mathbf{poly}(n)$-time. Our algorithms explicitly return either an appropriate branch-decomposition or a grid-minor certifying a negative answer.

Maximilian Gorsky
Institute for Basic Science
m.gorsky@pm.me

Giannos Stamoulis
Université Paris Cité, CNRS, IRIF
stamoulis@irif.fr

Dimitrios Thilikos
CNRS, LIRMM, France
sedthilk@thilikos.info

Sebastian Wiederrecht
KAIST
wiederrecht@kaist.ac.kr

CP29

A Quasi-Polynomial Bound for the Minimal Excluded Minors for a Surface

As part of their graph minor project, Robertson and Seymour showed in 1990 that the class of graphs that can be embedded in a given surface can be characterized by a finite set of minimal excluded minors. However, their proof, because existential, does not provide any information on these excluded minors. Seymour proved in 1993 the first and, until now, only known upper bound on the order of the minimal excluded minors for a given surface. This bound is double exponential in the Euler genus g of the surface and, therefore, very far from the $\Omega(g)$ lower bound on the maximal order of minimal excluded minors for a surface and most likely far from the best possible bound. More than thirty years later, this paper finally makes progress in lowering this bound to a quasi-polynomial in the Euler genus of the surface. The main catalyzer to reach a quasi-polynomial bound is a breakthrough on the characteristic size of a forbidden structure for a minimal excluded minor G for a surface of Euler genus g: although it is not hard to show that G does not contain $O(g)$ disjoint cycles that are contractible and nested in some embedding of G as demonstrated by Seymour, this bound can be lowered to $O(\log g)$ which is essential to obtain the quasi-polynomial bound in this paper. As subsidiary results, we also improve the current bound on the treewidth of a minimal excluded minor G for a surface by improving the first and, until now, only known bound provided by Seymour.

Sarah Houdaigoui
Graduate university for advanced studies (SOKENDAI)
National Institute of Informatics (Tokyo)
shoudaigoui@nii.ac.jp

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

CP29

$\mathcal{H}$-Planarity and Parametric Extensions: When Modulators Act Globally

We introduce a series of graph decompositions based on the modulator/target scheme of modification problems that enable several algorithmic applications that parametrically extend the algorithmic potential of planarity. In the core of our approach is a polynomial time algorithm for computing planar $\mathcal{H}$-modulators. Given a graph class $\mathcal{H}$, a planar $\mathcal{H}$-modulator of a graph G is a vertex set X such that the torso of X is planar and all connected components of $G - X$ belong to $\mathcal{H}$. We introduce $\mathcal{H}$-Planarity as the problem of deciding whether a graph G has a planar $\mathcal{H}$-modulator. We prove that, if $\mathcal{H}$ is hereditary, CMSO-definable, and decidable in polynomial time, then $\mathcal{H}$-Planarity is solvable in polynomial time. Additionally, we introduce two parametric extensions of $\mathcal{H}$-Planarity by defining the notions of $\mathcal{H}$-planar treedepth and $\mathcal{H}$-planar treewidth, which generalize the concepts of elimination distance and tree decompositions to the class $\mathcal{H}$. By leveraging our polynomial-time algorithm for $\mathcal{H}$-Planarity, we prove the following theorem. Suppose that $\mathcal{H}$ is a hereditary, CMSO-definable, and union-closed graph class, and that Vertex Deletion to $\mathcal{H}$ is solvable in FPT-time parameterized by the solution size. Then deciding whether a graph G has $\mathcal{H}$-planar treedepth or $\mathcal{H}$-planar treewidth at most k is also FPT when parameterized by k.

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

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

Laure Morelle
University of Bergen
laure_morelle@orange.fr

Dimitrios Thilikos
CNRS, LIRMM, France
sedthilk@thilikos.info

CP29

A Parameterized Linear Formulation of the Integer Hull

Let $A \in \mathbb{Z}^{m \times n}$ be an integer matrix with components bounded by Δ in absolute value. Cook et al. (1986) have shown that there exists a universal matrix $B \in \mathbb{Z}^{m' \times n}$ with the following property: For each $b \in \mathbb{Z}^m$, there exists $t \in \mathbb{Z}^{m'}$ such that the integer hull of the polyhedron $P = \{x \in \mathbb{R}^n : Ax \leq b\}$ is described by $P_I = \{x \in \mathbb{R}^n : Bx \leq t\}$. Our main result is that t is an affine function of b as long as b is from a fixed equivalence class of the lattice $D \cdot \mathbb{Z}^m$. Here $D \in \mathbb{N}$ is a number that depends on n and Δ only. Furthermore, D as well as the matrix B can be computed in time depending on Δ and n only. An application of this result is the solution of an open problem posed by Cslovjecsek et al. (SODA 2024) concerning the complexity of 2-stage-stochastic integer programming problems. The main tool of our proof is the classical theory of Chvatal-Gomory cutting planes and the elementary closure of rational polyhedra.

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

Thomas Rothvoss
University of Washington
rothvoss@uw.edu

CP29

Planar Disjoint Shortest Paths Is Fixed-Parameter Tractable

In the Disjoint Shortest Paths problem one is given a graph G and a set $\mathcal{T} = \{(s_1, t_1), \ldots, (s_k, t_k)\}$ of k vertex pairs. The question is whether there exist vertex-disjoint paths $P_1, \ldots, P_k$ in G so that each P_i is a shortest path between s_i and t_i. While the problem is known to be W[1]-hard in general, we show that it is fixed-parameter tractable on planar graphs with positive edge weights. Specifically, we propose an algorithm for Planar Disjoint Shortest Paths with running time $2^{O(k \log k)} \cdot n^{O(1)}$. Notably, our parameter dependency is better than state-of-the-art $2^{O(k^2)}$ for the Planar Disjoint Paths problem, where the sought paths are not required to be shortest paths.

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

Giannos Stamoulis
Université Paris Cité, CNRS, IRIF
stamoulis@irif.fr

Michal Wlodarczyk
Institute of Informatics
University of Warsaw, Poland
michal.wloda@gmail.com

CP30

Nearly Tight Sample Complexity for Matroid Online Contention Resolution

Due to their numerous applications, Prophet Inequalities have experienced a surge of interest. They describe competitive ratios for basic stopping time problems where random variables get revealed sequentially. A key drawback in the classical setting is the assumption of full distributional knowledge of the involved random variables. A natural way to address this is via sample-based approaches. Recently, Fu, Lu, Gavin Tang, Wu, Wu, and Zhang (EC 2024) showed that sample-based Online Contention Resolution Schemes (OCRS) are a powerful tool to obtain sample-based Prophet Inequalities. They presented the first sample-based OCRS for matroid constraints. This allowed them to get the first sample-based Matroid Prophet Inequality, using $O(\log^4 n)$ many samples (per random variable), while obtaining a constant competitiveness of $1/4 - \varepsilon$. We present a nearly optimal sample-based OCRS for matroid constraints, which uses only $O(\log \rho \cdot \log^2 \log \rho)$ many samples, almost matching a known lower bound of $\Omega(\log \rho)$, where ρ is the rank of the matroid. Through the above-mentioned connection to Prophet Inequalities, this yields a sample-based Matroid Prophet Inequality using only $O(\log n + \log \rho \cdot \log^2 \log \rho)$ samples, matching the competitiveness of $1/4 - \varepsilon$, which is the best known competitiveness for the considered almighty adversary setting even when the distributions are fully known.

Moran Feldman
Open University of Israel
moranfe@cs.haifa.ac.il

Ola Svensson
EPFL
ola.svensson@epfl.ch

Rico Zenklusen
ETH Zurich
ricoz@ethz.ch

CP30

Persuasive Calibration

We introduce and study the persuasive calibration problem, where a principal aims to provide trustworthy predictions about underlying events to a downstream agent to make desired decisions. We adopt the standard calibration framework that regulates predictions to be unbiased conditional on their own value and thus can reliably be interpreted at face value by the agent. Allowing a small calibration error budget, we aim to ask: what the optimal predictor is and how to compute it under this calibration error budget, especially when there exists incentive misalignment between the principal and the agent? We focus on the standard ℓ_t-norm Expected Calibration Error (ECE) metric. We develop a general framework by viewing predictors as post-processed versions of perfectly calibrated predictors. Using this, we first characterize the structure of the optimal predictor. Specifically, when the principal's utility is outcome-independent, we show: (1) the optimal predictor is over-(under-) confident for high (low) true expected outcomes, while remaining perfectly calibrated in the middle; (2) the miscalibrated predictions exhibit a collinearity structure with the principal's utility function. On the algorithmic side, we provide an FPTAS for computing approximately optimal predictor for general principal utility and general ℓ_t-norm ECE. Moreover, for the ℓ_1- and ℓ_∞-norm ECE, we provide polynomial-time algorithms that compute the exact optimal predictor.

Yiding Feng
Hong Kong University of Science and Technology
ydfeng@ust.hk

Wei Tang
Chinese University of Hong Kong
weitang@cuhk.edu.hk

CP30

Online Proportional Apportionment

Traditionally, apportioning the seats of a legislative body has been treated as a one-shot problem without dynamic considerations. Yet, in many contexts, dynamic aspects play an important role. We initiate the study of online apportionment, introducing an algorithmic framework that

handles proportional seat allocation without knowledge of future votes. Time proceeds in discrete steps, and n parties receive a share of votes at each step. An online method must irrevocably assign a fixed number of seats per step, ensuring that each partys cumulative allocation remains close to its cumulative share. We analyze deterministic and randomized methods. For deterministic methods, we construct a family of adversarial instances yielding a lower bound linear in n on the worst-case deviation between a partys allocated and ideal cumulative seats. We show that this bound is tight and matched by a natural greedy method. Consequently, a method ensuring that cumulative allocations always equal cumulative shares rounded up or down (the global quota property) exists if and only if $n \leq 3$. For randomized allocations, we show that when $n \leq 3$, one can randomize over methods satisfying global quota so that each party receives its proportional share in expectation at every step. The proof is constructive: such methods can be obtained from a flow on a recursively defined network.

Javier Cembrano
Max Planck Institute for Informatics
jcembran@mpi-inf.mpg.de

Jose Correa
Universidad de Chile, Chile
correa@uchile.cl

Svenja M. Griesbach
Centro de Modelamiento Matemático
Universidad de Chile
sgriesbach@cmm.uchile.cl

Victor Verdugo
Pontificia Universidad Católica de Chile
victor.verdugo@uc.cl

CP30

A Few Good Choices

Condorcet winning sets address the Condorcet paradox by selecting a small set of candidatesrather than a single winnersuch that a majority prefers no unselected alternative over all members of the set. This notion extends to α-undominated sets, which require the same property to hold for any α-fraction of voters. Such sets are guaranteed to exist with constant size for any fixed α. However, the requirement that an outsider be preferred to every member of the set can be overly restrictive and difficult to justify in many applications. Motivated by this, we introduce a more flexible notion: (t, α)-undominated sets. Here, each voter compares an outsider to their t-th favorite member of the set, and the set is undominated if no outsider is preferred by more than an α-fraction of voters. This framework subsumes prior definitions and introduces a tunable notion of collective acceptability for $t > 1$. We establish three main results: 1) For all values of t and α, there exists a (t, α)-undominated set of size $(O(t/\alpha)$. 2) As t grows large, the minimum size of such sets approaches $\frac{t}{\alpha}$, establishing the asymptotic optimality of our bound. 3) In the special case $t = 1$, we improve the upper bound on the size of α-undominated sets given by Charikar et al. (STOC 2025). In particular, we show that a Condorcet winning set of five candidates always exists, improving their bound of six.

Haoyu Song
Purdue University
song522@prudue.edu

Thanh Nguyen
Purdue University, Daniel School of Business
nguye161@purdue.edu

Young-San Lin
Melbourne School of Business
y.lin@mbs.edu

CP30

Networked Information Aggregation Via Machine Learning

We study a distributed learning problem in which agents are embedded in a directed acyclic graph (DAG). There is a fixed, arbitrary distribution over real-valued features/labels, and each agent is able to directly observe only a subset of features — potentially different for every agent. The agents learn sequentially in an order consistent with a topological sort of the DAG, committing to a model mapping observations to predictions. Each agent observes the predictions of their parents and trains their model using both the features of the instance they directly observe and these predictions as additional features. We ask when this process achieves *information aggregation*, i.e. some agent learns a model whose error is competitive with the best model that could have been learned (in some hypothesis class) with direct access to all features, despite the fact that no single agent has such access. We give upper and lower bounds for both linear and general hypothesis classes. Our results identify the *depth* of the DAG as the key parameter: information aggregation occurs over sufficiently long paths, assuming that all of the relevant features are well represented, and there are distributions over which information aggregation cannot occur even in the linear case, and even in arbitrarily large DAGs that do not have sufficient depth (e.g. hub-and-spokes topology in which the spokes collectively see all features). We complement our theory with comprehensive experiments.

Michael Kearns
University of Pennsylvania
mkearns@cis.upenn.edu

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

Emily Ryu
Cornell University
eryu@cs.cornell.edu

CP30

Collaborative Prediction: Tractable Information Aggregation Via Agreement

We give efficient "collaboration protocols" through which two parties, who observe different features about the same instances, can interact to arrive at predictions that are more accurate than either could have obtained on their own. The parties only need to iteratively share and update their own label predictions. Our protocols are efficient reductions to the problem of learning on each party's feature space alone, and so can be used even in settings in which each party's feature space is illegible to the other-which arises in models of human/AI interaction and in multimodal learning. In an online adversarial setting we show how to give regret bounds on the predictions that the parties arrive at with respect to a class of benchmark policies defined on the joint feature space of the two parties, despite the fact that neither party has access to this joint feature space. Our theorems give a computationally and statistically tractable generalization of past work on information aggregation amongst Bayesians who share a common prior, as part of a literature studying “agreement” in the style of Aumann’s agreement theorem. Our results require no knowledge of (or even the existence of) a prior distribution and are computationally efficient. Nevertheless we show how to lift our theorems back to this Bayesian setting, and in doing so, give new information aggregation theorems for Bayesian agreement.

Natalie Collina
University of Pennsylvania
ncollina@seas.upenn.edu

Ira Globus-Harris
Cornell University
University of Pennsylvania
igh@seas.upenn.edu

Surbhi Goel
University of Pennsylvania
surbhig@seas.upenn.edu

Varun Gupta
Vector Institute
University of Pennsylvania
vgup@seas.upenn.edu

Aaron Roth, Mirah Shi
University of Pennsylvania
aaroth@gmail.com, mirahshi@seas.upenn.edu

CP31

Traversing Regions of Supersolvable Hyperplane Arrangements and Their Lattice Quotients

For an arrangement $\mathcal{H}$ of hyperplanes in $\mathbb{R}^n$ through the origin, a region is a connected subset of $\mathbb{R}^n \setminus \mathcal{H}$. The graph of regions $G(\mathcal{H})$ has a vertex for every region and an edge between any two regions separated by a single hyperplane. We aim to compute a Hamiltonian path or cycle in $G(\mathcal{H})$. Our first result is that if $\mathcal{H}$ is supersolvable, then $G(\mathcal{H})$ has a Hamiltonian cycle. More generally, we consider quotients of lattice congruences of the poset of regions $P(\mathcal{H}, R_0)$, obtained by orienting $G(\mathcal{H})$ away from a particular base region R_0. Our second result is that if $\mathcal{H}$ is supersolvable, then for any lattice congruence $\equiv$ on $P(\mathcal{H}, R_0) =: L$, the cover graph of the quotient lattice $L/\equiv$ has a Hamiltonian path. When applying our results to well-known supersolvable arrangements, we recover a number of known Gray code algorithms. For the type B Coxeter arrangement and subarrangements, we obtain new Gray codes for (pattern-avoiding) signed permutations, symmetric triangulations, acyclic orientations of certain signed graphs, and in general for combinatorial families of Coxeter type B, which generalizes the theory of zigzag languages to signed permutations. Our approach also yields new Hamiltonicity results for large classes of polytopes, in particular signed graphic zonotopes and type B quotientopes.

Sofia Brenner
University of Kassel
sbrenner@mathematik.uni-kassel.de

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

Thomas McConville
Department of Mathematics
Kennesaw State University
tmcconvi@kennesaw.edu

Arturo Merino
Universidad de OHiggins, Rancagua
arturo.merino@uoh.cl

Torsten Mütze
University of Kassel
tmuetze@mathematik.uni-kassel.de

CP31

Nearly Optimal Bounds for Stochastic Online Sorting

In the *online sorting* problem, we have an array A of n cells, and receive a stream of n items $x_1, \ldots, x_n \in [0, 1]$. When an item arrives, we need to immediately and irrevocably place it into an empty cell. The goal is to minimize the sum of absolute differences between adjacent items, which is called the *cost* of the algorithm. It has been shown by Aamand, Abrahamsen, Beretta, and Kleist (SODA 2023) that when the stream $x_1, \ldots, x_n$ is generated adversarially, the optimal cost bound for any deterministic algorithm is $\Theta(\sqrt{n})$. In this paper, we study the *stochastic* version of online sorting, where the input items $x_1, \ldots, x_n$ are sampled uniformly at random. Despite the intuition that the stochastic version should yield much better cost bounds, the previous best algorithm for stochastic online sorting by Abrahamsen, Bercea, Beretta, Klausen and Kozma (ESA 2024) only achieves $\tilde{O}(n^{1/4})$ cost, which seems far from optimal. We show that stochastic online sorting indeed allows for much more efficient algorithms, by presenting an algorithm that achieves expected cost $\log n \cdot 2^{O(\log^* n)}$. We also prove a cost lower bound of $\Omega(\log n)$, thus show that our algorithm is nearly optimal.

Yang Hu
Tsinghua University
yanghu2004@outlook.com

CP31

On Lines Crossing Pairwise Intersecting Convex Sets in Three Dimensions

The 1913 Helly’s theorem states that any family $\mathcal{K}$ of $n \geq d+1$ convex sets in $\mathbb{R}^d$ can be pierced by a single point if and only if any $d+1$ of $\mathcal{K}$’s elements can. In 2002 Alon, Kalai, Matoušek and Meshulam ruled out the possibility of similar criteria for the existence of *lines* crossing multiple convex sets in dimension $d \geq 3$ – for any $k \geq 3$, they described arbitrary large families $\mathcal{K}$ of convex sets in $\mathbb{R}^3$ so that any k elements of $\mathcal{K}$ can be crossed by a line yet no $k+4$ of them can. Let $\mathcal{K}$ be a family of n pairwise intersecting convex sets in $\mathbb{R}^3$. We show that there exists a line crossing $\Theta(n)$ elements of $\mathcal{K}$. This resolves the most extensively studied variant of a problem by Martinez, Roldán-Pensado and Rubin (*Discrete Comput. Geom. 2020*) which was highlighted by Bárány and Kalai (*Bull. Amer. Math. Soc. 2021*). Our result adds to the very few sufficient (and non-trivial) conditions that have been known for the existence of line transversals to large families of convex sets. Our argument is based on a Ramsey-type result of independent interest for families of pairwise intersecting convex sets in

$\mathbb{R}^2$, and the structure of line arrangements in $\mathbb{R}^3$.

Natan Rubin
Ben Gurion University
rubinnat.ac@gmail.com

CP31

The Erdos-Psa Property for Circle Graphs as Vertex-minors

We prove that for any circle graph H with at least one edge and for any positive integer k, there exists an integer $t = t(k, H)$ so that every graph G either has a vertex-minor isomorphic to the disjoint union of k copies of H, or has a t-perturbation with no vertex-minor isomorphic to H. Using the same techniques, we also prove that for any planar multigraph H, every binary matroid either has a minor isomorphic to the cycle matroid of kH, or is a low-rank perturbation of a binary matroid with no minor isomorphic to the cycle matroid of H.

Sebastian Wiederrecht, Rutger Campbell
KAIST
wiederrecht@kaist.ac.kr, rutger@ibs.re.kr

J. Pascal Gollin
FAMNIT, University of Primorska,
pascal.gollin@famnit.upr.si

Meike Hatzel
TU Darmstadt
research@meikehatzel.com

O-Joung Kwon
Hanyang University
ojoungkwon@hanyang.ac.kr

Rose McCarty
Georgia Institute of Technology
rmccarty3@gatech.edu

Sang-Il Oum
Institute for Basic Science / KAIST
sangil@ibs.re.kr

CP31

Universal Connection Schedules for Reconfigurable Networking

Reconfigurable networks are a novel communication paradigm in which the pattern of connectivity between hosts varies rapidly over time. Prior theoretical work explored the inherent tradeoffs between throughput (or, hop-count) and latency. Existing Pareto-optimal designs use a connection schedule which is fine-tuned to the desired hop-count h, permitting lower latency as h increases. However, in reality datacenter workloads contain a mix of low-latency and high-latency requests. Using a design fine-tuned for one request type leads to inefficiencies when serving other types. A more flexible alternative is a universal schedule, a single connection schedule capable of attaining many Pareto-optimal tradeoff points simultaneously by varying the choice of routing paths. We present the first universal schedules for oblivious routing, which are near-optimal for all possible hop-counts h. The key idea is to specialize a connection schedule based on cyclic permutations and to develop a novel Fourier-analytic method for analyzing randomized routing on these connection schedules. We first show a uniformly random connection schedule suffices with multiplicative error in throughput, and latency optimal up to a $\log N$ factor. We then analyze a more carefully designed random connection schedule with additive error in throughput and improved latency optimal up to constant factors. Finally, we show that our first randomized construction can be made deterministic.

Tegan Wilson, Shaleen Baral, Robert Kleinberg, Sylvan Martin, Henry Rogers, Ruogu Zhang
Cornell University
te.wilson@northeastern.edu, sb969@cornell.edu, rdk@cs.cornell.edu, slm338@cornell.edu, hwr34@cornell.edu, rz234@cornell.edu

CP31

Rapid Mixing of Glauber Dynamics for Monotone Systems Via Entropic Independence

We study the mixing time of Glauber dynamics on monotone systems. For monotone systems satisfying the entropic independence condition, we prove a new mixing time comparison result for Glauber dynamics. For concrete applications, we obtain $\tilde{O}(n)$ mixing time for the random cluster model induced by the ferromagnetic Ising model with consistently biased external fields, and $\tilde{O}(n^2)$ mixing time for the bipartite hardcore model under the one-sided uniqueness condition, where n is the number of variables in corresponding models, improving the best known results in [Chen and Zhang, SODA'23] and [Chen, Liu, and Yin, FOCS'23], respectively.

Minji Yang, Weiming Feng
The University of Hong Kong
ymjessen02@connect.hku.hk, wfeng@hku.hk

CP32

On the Universality of Round Elimination Fixed Points

Recent work on distributed graph algorithms has drawn attention to the following open question: are round elimination fixed points a universal technique for proving lower bounds? That is, given a locally checkable problem Π that requires at least $\Omega(\log n)$ rounds in the deterministic LOCAL model, can we always find a relaxation Π' of Π that is a nontrivial fixed point for the round elimination technique? If yes, then a key part of distributed computational complexity would be also decidable. The key obstacle so far has been a certain family of homomorphism problems, which require $\Omega(\log n)$ rounds, but the only known proof is based on Marks' technique. We develop a new technique for constructing round elimination lower bounds systematically. Using so-called tripotent inputs we show that the aforementioned homomorphism problems indeed admit a lower bound proof that is based on round elimination fixed points. Hence we eliminate the only known obstacle for the universality of round elimination. Yet we also present a new obstacle: we show that there are some problems with inputs that require $\Omega(\log n)$ rounds, yet there is no proof that is based on relaxations to nontrivial round elimination fixed points. Hence round elimination cannot be a universal technique for problems with inputs (but it might be universal for problems without inputs).

Sebastian Brandt
CISPA Helmholtz Center for Information Security
brandt@cispa.de

Alkida Balliu
Gran Sasso Science Institute
alkida.balliu@gssi.it

Ole Gabsdil
CISPA Helmholtz Center for Information Security
ole.gabsdil@gmail.com

Dennis Olivetti
Gran Sasso Science Institute
dennis.olivetti@gssi.it

Jukka Suomela
Aalto University
jukka.suomela@aalto.fi

CP32

Computational Barriers for Permutation-Based Problems, and Cumulants of Weakly Dependent Random Variables

In many high-dimensional problems, polynomial-time algorithms fall short of achieving the statistical limits attainable without computational constraints. A powerful approach to probe the limits of polynomial-time algorithms is to study the performance of low-degree polynomials. The seminal work of [Schramm,Wein 22] connects low-degree lower bounds to multivariate cumulants. Prior works [Luo et.al 23, Even et.al 25] leverage independence among latent variables to bound cumulants. However, such approaches break down for problems with latent structure lacking independence, such as those involving random permutations. To address this important restriction, we develop a technique to upper-bound cumulants under weak dependencies – such as those arising from sampling without replacement or random permutations. To show-case the effectiveness of our approach, we uncover evidence of statistical–computational gaps in multiple feature matching and in seriation problems.

Bertrand Even, Christophe Giraud
Université Paris-Saclay
bertrand.even@universite-paris-saclay.fr,
christophe.giraud@universite-paris-saclay.fr

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

CP32

Downward Self-reducibility in the Total Function Polynomial Hierarchy

A problem P is considered downward self-reducible, if there exists an efficient algorithm for P that is allowed to make queries to only strictly smaller instances of P. Downward self-reducibility has been well studied in the case of decision problems, and it is well known that any downward self-reducible problem must lie in PSPACE. Harsha, Mitropolsky and Rosen [ITCS, 2023] initiated the study of downward self reductions in the case of search problems. They showed the following interesting collapse: if a problem is in TFNP and also downward self-reducible, then it must be in PLS. Moreover, if the problem admits a unique solution then it must be in UEOPL. We demonstrate that this represents just the tip of a much more general phenomenon, which holds for even harder search problems that lie higher up in the total function polynomial hierarchy ($TF\Sigma_i^P$). In fact, even if we allow our downward self-reduction to be much more powerful, such a collapse will still occur. We show that any problem in $TF\Sigma_i^P$ which admits a randomized downward self-reduction with access to a Σ_{i-1}^P oracle must be in $PLS^{\Sigma_{i-1}^P}$. If the problem has essentially unique solutions then it lies in $UEOPL^{\Sigma_{i-1}^P}$. As one (out of many) application of our framework, we get new upper bounds for the problems RangeAvoidance and LinearOrderingPrinciple and show that they are both in $UEOPL^{NP}$.

Karthik Gajulaplli
Georgetown University
kg816@georgetown.edu

Surendra Ghentiyala
Cornell University
sg974@cornell.edu

Zeyong Li
National University of Singapore
li.zeyong@u.nus.edu

Sidhant Saraogi
Georgetown University
ss4456@georgetown.edu

CP32

Better Bounds for Semi-Streaming Single-Source Shortest Paths

In the semi-streaming model, an algorithm must process any n-vertex graph by making one or few passes over a stream of its edges, use $\widetilde{O}(n) := O(n \cdot \text{polylog}(n))$ space, and at the end of the last pass, output a solution to the problem at hand. Approximating shortest paths on *undirected* graphs is a longstanding open question in this model. In this work, we make progress on this question from both upper and lower bound fronts: 1. We present a simple randomized algorithm that for any $\epsilon > 0$, with high probability computes $(1+\epsilon)$-approximate shortest paths from a given source vertex in $O(\epsilon^{-1} \cdot n \log n)$ space and $O(\epsilon^{-1} \cdot (\frac{\log n}{\log \log n})^2)$ passes. The algorithm can also be derandomized and made to work on dynamic streams at a cost of some extra $\text{poly}(\log(n), 1/\epsilon)$ factors only in the space. Previously, the best known algorithms for this problem required $\epsilon^{-1} \cdot \log^c(n)$ passes, for an unspecified large constant c. 2. We prove that any semi-streaming algorithm that with large constant probability outputs any constant approximation to shortest paths from a given source vertex requires $\Omega\left(\frac{\log n}{\log \log n}\right)$ passes. Our results collectively reduce the gap in the pass complexity of approximating single-source shortest paths in the semi-streaming model from polylog(n) vs $\omega(1)$ to only a quadratic gap.

Sepehr Assadi
University of Waterloo
sepehr@assadi.info

Gary Hoppenworth
University of Michigan
garytho@umich.edu

Janani Sundaresan
University of Waterloo

jsundaresan@uwaterloo.ca

CP32

Plane Vs. Plane Low Degree Test

In this work, we give an optimal analysis of the plane versus plane test of Raz and Safra (STOC'97). More specifically, consider a table $\mathcal{T}$ that assigns every plane P from $\mathbb{F}_q^m$ a bivariate degree d polynomial. The goal is to check if these polynomials are restrictions of a global degree d polynomial $f : \mathbb{F}_q^m \to \mathbb{F}_q$. Raz and Safra introduced the following natural test: sample two random planes P, P' intersecting in a line ℓ and check if $\mathcal{T}(\mathsf{P})|_\ell = \mathcal{T}(\mathsf{P}')|_\ell$, i.e., the two table entries agree on the points on ℓ. We show that if the test passes with probability at least $\epsilon = \Omega(d/q)$, then there is a global degree d polynomial f such that for at least $\Omega(\epsilon)$ fraction of the planes P, $\mathcal{T}(\mathsf{P}) = f|_\mathsf{P}$. This improves on the previous best analysis of the test by Moshkovitz and Raz (STOC'06), where they proved the soundness of the test is at least $(\mathsf{poly}(d)/q)^{1/8}$. With $\Omega(1/q)$ as a natural lower bound on the soundness of this test, our result gets the optimal dependence on the field size, while also working for large degree parameters $d = \Omega(q)$. Our proof combines algebraic aspects from prior work on the lines vs lines test, with combinatorial aspects of recent works on the cubes vs cubes test.

Silas Richelson, Amey Bhangale
UC Riverside
silas@cs.ucr.edu, ameyb@ucr.edu

CP32

Sublinear-Time Lower Bounds for Approximating Matching Size Using Non-Adaptive Queries

We study the problem of estimating the size of the maximum matching in sublinear-time. This problem has been extensively studied, with several known upper and lower bounds. A notable result by Behnezhad (FOCS 2021) established a 2-approximation in $\widetilde{O}(n)$ time. However, all known bounds are in the adaptive query model, where each query can depend on previous answers. In contrast, non-adaptive querieswhere the distribution over all queries must be fixed in advanceare widely studied in property testing, often revealing fundamental gaps between adaptive and non-adaptive complexities. This raises the natural question: is adaptivity also necessary for approximating the maximum matching size in sublinear time? This motivates the goal of achieving a constant or a polylogarithmic approximation using $\widetilde{O}(n)$ non-adaptive adjacency list queries, similar to what was done by Behnezhad using adaptive queries. We show that this is not possible by proving that any randomized non-adaptive algorithm achieving an $n^{1/3-\gamma}$-approximation, for any constant $\gamma > 0$, with probability at least 2/3, must make $\Omega(n^{1+\varepsilon})$ adjacency list queries, for some constant $\varepsilon > 0$. This result highlights the necessity of adaptivity in achieving strong approximations. However, non-trivial upper bounds are still achievable: we present a simple randomized algorithm that achieves an $n^{1/2}$-approximation in $O(n \log^2 n)$ queries.

Vihan Shah
Rutgers University
vihanshah98@gmail.com

CP33

Algebraic Closure of Matrix Sets Recognized by 1-Vass

It is known how to compute the Zariski closure of a finitely generated monoid of matrices and, more generally, of a set of matrices specified by a regular language. This result was recently used to give a procedure to compute all polynomial invariants of a given affine program. Decidability of the more general problem of computing all polynomial invariants of affine programs with recursive procedure calls remains open. Mathematically speaking, the core challenge is to compute the Zariski closure of a set of matrices defined by a context-free language. In this talk, we approach the problem from two sides: Towards decidability, we give a procedure to compute the Zariski closure of sets of matrices given by one-counter languages (that is, languages accepted by one-dimensional vector addition systems with states and zero tests), a proper subclass of context-free languages. On the other side, we show that the problem becomes undecidable for indexed languages, a natural extension of context-free languages corresponding to nested pushdown automata. One of our main technical tools is a novel adaptation of Simon's factorization forests to infinite monoids of matrices.

Rida Ait El Manssour
Department of Computer Science, University of Oxford
rida.aitelmanssour@cs.ox.ac.uk

Mahsa Naraghi
IRIF
mnaraghi@irif.fr

Mahsa Shirmohammadi
Université Paris Cité, CNRS, IRIF
mahsa@irif.fr

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

CP33

On the Complexity of the Skolem Problem at Low Orders

The Skolem Problem asks to determine whether a given linear recurrence sequence (LRS) $\langle u_n \rangle_{n=0}^{\infty}$ over the integers has a zero term. Decidability of the problem is open in general, with the most notable positive result being a decision procedure for LRS of order at most 4. In this paper we consider a bounded version of the Skolem Problem, in which the input consists of an LRS $\langle u_n \rangle_{n=0}^{\infty}$ and a bound $N \in \mathbb{N}$ (with all integers written in binary), and the task is to determine whether there exists $n \in \{0, \dots, N\}$ such that $u_n = 0$. We give a randomised algorithm for this problem that, for all $d \in \mathbb{N}$, runs in polynomial time on the class of LRS of order at most d. As a corollary we show that the (unrestricted) Skolem Problem for LRS of order at most 4 lies in coRP, improving the best previous upper bound of $\mathsf{NP}^{\mathsf{RP}}$. The running time of our algorithm is exponential in the order of the LRS—a dependence that appears necessary in view of the NP-hardness of the Bounded Skolem Problem. However, even for LRS of a fixed order, the problem involves detecting zeros within an exponentially large range. For this, our algorithm relies on results from p-adic analysis to isolate polynomially many candidate zeros and then test in randomised polynomial time whether each candidate is an actual zero by reduction to arithmetic-circuit identity testing.

Piotr Bacik
University of Oxford
Max Planck Institute for Software Systems
piotr.bacik@stcatz.ox.ac.uk

Joel Ouaknine
Max Planck Institute for Software Systems
joel@mpi-sws.org

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

CP33

Entrywise Approximation for Matrix Inversion and Linear Systems

We study the bit complexity of inverting diagonally dominant matrices, which correspond to random-walk quantities such as hitting times and escape probabilities. Although these quantities can be exponentially small, their nonnegativity allows meaningful entrywise approximationa stronger notion than norm-based error. Under this notion, existing Laplacian solvers and fast matrix-multiplication methods require mn^2 and $n^{\omega+1}$ bit operations, respectively, where m is the number of nonzeros, n the dimension, and ω the matrix-multiplication exponent. We present algorithms for computing entrywise $\exp(\epsilon)$-approximate inverses of row-diagonally-dominant L-matrices (RDDL) under two types of input numbers. When the entries are given in floating-point representation, we obtain a cubic-time algorithm and prove its optimality under the all-pairs shortest-paths conjecture. When the entries are given in fixed-point representation, we design high-probability algorithms with entrywise guarantees: (1) for SDDM matrices, solving linear systems in $\tilde{O}(m\sqrt{n})$ and inverting in $\tilde{O}(mn)$ bit operations; (2) for RDDL matrices, solving in $\tilde{O}(mn^{1+o(1)})$ and inverting in $\tilde{O}(n^{\omega+0.5})$ or $\tilde{O}(mn^{1.5+o(1)})$ bit operations.

Mehrdad Ghadiri
MIT
mehrdadg@mit.edu

Hoai-An Nguyen
Carnegie Mellon University
hnnguyen@andrew.cmu.edu

Junzhao Yang
Institute for Interdisciplinary Information Sciences
Tsinghua University
junzhaoy@andrew.cmu.edu

CP33

Optimization Modulo Integer Linear-Exponential Programs

This paper presents the first study of the complexity of the optimization problem for integer linear-exponential programs which extend integer linear programs with the functions $x \mapsto 2^x$ and $(x, y) \mapsto (x \bmod 2^y)$. The problem of deciding if such a program has a solution was recently shown to be NP-complete in [Chistikov et al., ICALP'24]. The optimization problem instead asks for a solution that maximizes (or minimizes) a linear-exponential objective function, subject to the constraints of an integer linear-exponential program. We establish the following results: - When optimal solutions exist, one of them can be succinctly represented as an integer linear-exponential straight-line program (ILESLP): an arithmetic circuit whose gates output an integer value (by construction) and implement the operations of addition, exponentiation, and multiplication by rational numbers. - There is an algorithm that runs in polynomial time, given access to an integer factoring oracle, which determines whether an ILESLP encodes a solution to an integer linear-exponential program. This algorithm can also be used to compare the values taken by the objective function on two given solutions. Building on these results, we place the optimization problem for integer linear-exponential programs within an extension of the optimization class NPO that lies within $\mathrm{FNP}^{\mathrm{NP}}$. In essence, this extension forgoes determining the optimal solution via binary search.

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

S Hitarth
Hong Kong University of Science and Technology
hitarth.singh@connect.ust.hk

Gurupreran Shabadi
University of Pennsylvania
shabadi@seas.upenn.edu

CP33

The Division Barrier: Optimal Bounds and Structural Limits in Toom-Cook Interpolation

Toom-Cook-k ($2 \leq k \in \mathbb{N}$) is a family of fast algorithms for multiplying long integers using $O\left(n^{\log_k (2k-1)}\right)$ arithmetic operations, offering asymptotical improvement over the naïve quadratic-time schoolbook approach. Despite this advantage, Toom-Cook algorithms often involve nontrivial divisions, divisions by elements that are not powers of 2, which can be both computationally expensive and numerically unstable, especially in cryptography or quantum computing applications. Reducing or eliminating these divisions is therefore of significant theoretical and practical interest. Although previous work has sought to minimize the number of such divisions, Bodrato and Zanoni conjectured that any Toom-Cook-k algorithm must perform at least $2k-5$ nontrivial divisions. In this work, we prove a stronger and more general result; any linear interpolation scheme for a Toom-Cook-(k_1, k_2) algorithm with integer evaluation points must include at least (k_1+k_2-2-p) divisions by every prime $p \leq 2k-3$. Our findings confirms and extends the BodratoZanoni division conjecture. Furthermore, we establish the optimality of several known algorithms Toom-Cook-3, 3.5, 4, 4.5, 5, and 8. We also extend our bounds to polynomial multiplication over finite fields and provide arithmetic lower bounds in division-aware cost models.

Roy Nissim
The Hebrew University of Jerusalem
rnissim@redhat.com

Oded Schwartz
Hebrew University of Jerusalem
odedsc@cs.huji.ac.il

Yuval Spiizer
Tel Aviv University
yuvalspiizer@mail.tau.co.il

CP34

Additive Approximation Schemes for Low-Dimensional Embeddings

We consider the task of fitting low-dimensional embeddings to high-dimensional data. In particular, we study the k-Euclidean Metric Violation problem (k-EMV), where the input is $D \in \mathbb{R}_{\geq 0}^{\binom{n}{2}}$ and the goal is to find the closest vector $X \in \mathbb{M}_k$, where $\mathbb{M}_k \subset \mathbb{R}_{\geq 0}^{\binom{n}{2}}$ is the set of all k-dimensional Euclidean metrics on n points, and closeness is formulated as the following optimization problem:

$$\mathsf{OPT}_{\mathrm{EMV}} = \min_{X \in \mathbb{M}_k} \|D - X\|_2^2 .$$

Cayton and Dasgupta [CD'06] showed that this problem is NP-Hard, even when $k = 1$. Dhamdhere [Dha'04] obtained a $O(\log(n))$-approximation for 1-EMV and leaves finding a PTAS for it as an open question. Although k-EMV has been studied in the statistics community for over 70 years, under the name "multi-dimensional scaling", there are no known efficient approximation algorithms for k ¿ 1, to the best of our knowledge. We provide the first polynomial-time additive approximation scheme for k-EMV. In particular, we obtain an embedding with objective value $\mathsf{OPT}_{\mathrm{EMV}} + \varepsilon\|D\|_2^2$ in $(n \cdot B)^{\mathsf{poly}(k,\varepsilon^{-1})}$ time, where each entry in D can be represented by B bits. We also show that our techniques allow us to obtain additive approximation schemes for two related problems: a weighted variant of k-EMV and ℓ_p low-rank approximation for p¿2.

Prashanti Anderson
MIT
paanders@mit.edu

Ainesh Bakshi
NYU
aineshbakshi@nyu.edu

Samuel Hopkins
MIT
samhop@mit.edu

CP34

Approximating Asymmetric A Priori Tsp Beyond the Adaptivity Gap

In Asymmetric A Priori TSP (with independent activation probabilities) we are given an instance of the Asymmetric Traveling Salesman Problem together with an activation probability for each vertex. The task is to compute a tour that minimizes the expected length after short-cutting to the randomly sampled set of active vertices. We prove a polynomial lower bound on the adaptivity gap for Asymmetric A Priori TSP. Moreover, we show that a poly-logarithmic approximation ratio, and hence an approximation ratio below the adaptivity gap, can be achieved by a randomized algorithm with quasi-polynomial running time. To achieve this, we provide a series of polynomial-time reductions. First we reduce to a novel generalization of the Asymmetric Traveling Salesman Problem, called Hop-ATSP. Next, we use directed low-diameter decompositions to obtain structured instances, for which we then provide a reduction to a covering problem. Eventually, we obtain a polynomial-time reduction of Asymmetric A Priori TSP to a problem of finding a path in an acyclic digraph minimizing a particular objective function, for which we give an $O(\log n)$-approximation algorithm in quasi-polynomial time.

Manuel Christalla
ETH Zürich
mchristalla@ethz.ch

Luise Puhlmann
University of Bonn
puhlmann@dm.uni-bonn.de

Vera Traub
Department of Computer Science
ETH Zurich
vtraub@ethz.ch

CP34

Spanning Tree Embeddings Are Not Much Harder Than Hierarchical Partitions

Embedding graph distances into trees is an important technique in algorithms. Often, applications of these embeddings require that the trees into which distances are embedded are spanning trees of the input graph. However, this requirement makes embedding more difficult. This is largely due to the fact that embedding graph distances into hierarchical partitions is well-understood and, in turn, embedding hierarchical partitions into non-spanning trees is trivial. On the other hand, no such embedding of hierarchical partitions into spanning trees is known. In this work, we show that, generally speaking, hierarchical partitions of graphs embed into spanning trees. In particular, we show that every hierarchical partition embeds into a spanning tree with $O(1)$-distortion as long as the diameter of each level of the hierarchy increases by $\Omega(\log n)$. As an immediate consequence of our embedding, we improve the best-known approximation for universal Steiner trees (USTs) from $O(\log^7 n)$ to $O(\log^5 n)$. Likewise, our embedding gives a new and extremely simple probabilistic tree embedding into spanning trees with distortion $O(\log^2 n)$ whose analysis is essentially the same as the classic (non-spanning) tree embedding of Bartal (FOCS'96, ESA'04). Our embeddings also give the first construction of Ramsey spanning trees from the Ramsey partitions of Mendel and Naor (FOCS'06) and the first non-trivial embeddings of hop-constrained distances into spanning trees.

Willem Fletcher, D Ellis Hershkowitz
Brown University
willemrfletcher@gmail.com, delhersh@gmail.com

CP34

Max Bisection Might Be Harder to Approximate Than Max Cut

The MAX BISECTION problem seeks a maximum-size cut that evenly divides the vertices of a given undirected graph. An open problem raised by Austrin, Benabbas, and Georgiou is whether MAX BISECTION can be approximated as well as MAX CUT, i.e., to within $\alpha_{\mathrm{GW}} \approx 0.8785672\ldots$, which is the approximation ratio achieved by the celebrated Goemans-Williamson algorithm for MAX CUT, which is best possible assuming the Unique Games Conjecture. They conjectured that the answer is yes. In this paper, we prove limitations on the state of the art algorithmic methods for MAX BISECTION due to Raghavendra and Tan [SODA'12] and Austrin, Benabbas, and Georgiou. More precisely, for any $\varepsilon > 0$, we construct an explicit instance of MAX BISECTION for which the ratio between the value of the optimal integral solution and the value of some ε-uncorrelated solution of the Basic SDP relaxation is less than $0.87853 < \alpha_{\mathrm{GW}}$. Our instances are also integrality gaps for the Basic SDP relaxation of MAX BISECTION. As a corollary, we also obtain a family of MAX BISECTION-based dictatorship tests with the same ratio.

Neng Huang
University of Chicago
nengh@umich.edu

Joshua Brakensiek
UC Berkeley
josh.brakensiek@berkeley.edu

Aaron Potechin
University of Chicago
potechin@uchicago.edu

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

CP34

An Efficient Massively Parallel Constant-Factor Approximation Algorithm for the k-Means Problem

In this paper, we present an efficient massively parallel approximation algorithm for the k-means problem. Specifically, we provide an MPC algorithm that computes a constant-factor approximation to an arbitrary k-means instance in $O(\log\log n \cdot \log\log\log n)$ rounds. The algorithm uses $O(n^\sigma)$ bits of memory per machine, where $\sigma > 0$ is a constant that can be made arbitrarily small. The global memory usage is $O(n^{1+\varepsilon})$ bits for an arbitrarily small constant $\varepsilon > 0$, and is thus only slightly superlinear. Recently, Czumaj, Gao, Jiang, Krauthgamer, and Vesel showed that a constant-factor bicriteria approximation can be computed in $O(1)$ rounds in the MPC model. However, our algorithm is the first constant-factor approximation for the general k-means problem that runs in $o(\log n)$ rounds in the MPC model. Our approach builds upon the foundational framework of Jain and Vazirani. The core component of our algorithm is a constant-factor approximation for the related facility location problem. While such an approximation was already achieved in constant time in the work of Czumaj et al. mentioned above, our version additionally satisfies the so-called Lagrangian Multiplier Preserving (LMP) property. This property enables the transformation of a facility location approximation into a comparably good k-means approximation.

Vincent Cohen-Addad
Google Research
cohenaddad@google.com

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

Zahra Parsaeian
University of Freiburg
zahra.parsaeian@cs.uni-freiburg.de

CP35

Algorithms and Lower Bounds for the Maximum Overlap of Two Polygons Under Translation

A fundamental problem in shape matching and geometric similarity is computing the maximum area overlap between two polygons under translation. For general simple polygons, the best-known algorithm runs in $O((nm)^2 \log(nm))$ time [Mount, Silverman, Wu 96], where n and m are the complexities of the input polygons. In a recent breakthrough, Chan and Hair gave a linear-time algorithm for the special case when both polygons are convex. A key challenge in computational geometry is to design improved algorithms for other natural classes of polygons. We address this by presenting an $O((nm)^{3/2} \log(nm))$-time algorithm for the case when both polygons are orthogonal. This is the first algorithm for polygon overlap on orthogonal polygons that is faster than the almost 30 years old algorithm for simple polygons. We provide k-SUM lower bounds for problems on simple polygons with only orthogonal and diagonal edges. First, we establish that there is no algorithm for polygon overlap with running time $O(\max(n^2, nm^2)^{1-\varepsilon})$, where $m \leq n$, unless the k-SUM hypothesis fails. This matches the running time of our algorithm when $n = m$. We use part of the above construction to also show a lower bound for the polygon containment problem, a popular special case of the overlap problem. Concretely, there is no algorithm for polygon containment with running time $O(n^{2-\varepsilon})$ under the 3-SUM hypothesis, even when the polygon to be contained has $m = O(1)$ vertices.

Mikkel Abrahamsen
University of Copenhagen
miab@di.ku.dk

Sujoy Bhore
IIT Bombay
sujoy@cse.iitb.ac.in

Maike Buchin
Faculty of Mathematics
Ruhr-University Bochum
maike.buchin@rub.de

Jacobus Conradi
University of Bonn, Germany
University of Copenhagen, Denmark
jacobus.conradi@gmx.de

Ce Jin
EECS
MIT
cejin@berkeley.edu

André Nusser
Université Côte dAzur, CNRS, Inria, France
andre.nusser@inria.fr

Carolin Rehs
Technical University of Dortmund, Germany
Technical University of Eindhoven, The Netherlands

carolin.rehs@tu-dortmund.de

CP35

Vizings Theorem in Deterministic Almost-Linear Time

Vizing's theorem states that any n-vertex m-edge graph of maximum degree Δ can be edge colored using at most $\Delta+1$ different colors. Vizing's original proof is easily translated into a deterministic $O(mn)$ time algorithm. This deterministic time bound was subsequently improved to $\tilde{O}(m\sqrt{n})$ time, independently by [Arjomandi, 1982] and by [Gabow et al., 1985]. A series of recent papers improved the time bound of $\tilde{O}(m\sqrt{n})$ using randomization, culminating in the randomized near-linear time $(\Delta+1)$-coloring algorithm by [Assadi, Behnezhad, Bhattacharya, Costa, Solomon, and Zhang, 2025]. At the heart of all of these recent improvements, there is some form of a sublinear time algorithm. Unfortunately, sublinear time algorithms as a whole almost always require randomization. This raises a natural question: can the deterministic time complexity of the problem be reduced below the $\tilde{O}(m\sqrt{n})$ barrier? In this paper, we answer this question in the affirmative. We present a deterministic almost-linear time $(\Delta+1)$-coloring algorithm, namely, an algorithm running in $m \cdot 2^{O(\sqrt{\log \Delta})} \cdot \log n = m^{1+o(1)}$ time. Our main technical contribution is to entirely forego sublinear time algorithms. We do so by presenting a new deterministic color-type sparsification approach that runs in almost-linear (instead of sublinear) time, but can be used to color a much larger set of edges.

Sepehr Assadi
University of Waterloo
sepehr@assadi.info

Soheil Behnezhad
Northeastern University
s.behnezhad@northeastern.edu

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

Martin Costa
University of Warwick
martincosta2000@outlook.com

Shay Solomon
Tel Aviv University
solo.shay@gmail.com

Tianyi Zhang
ETH Zurich
tianyiz25@nju.edu.cn

CP35

Online Orthogonal Vectors Revisited

We prove new upper and lower bounds for the Online Orthogonal Vectors Problem ($OnlineOV_{n,d}$). In this problem, a preprocessing algorithm receives n vectors $x_1, \ldots, x_n \in \{0,1\}^d$ and constructs a data structure of size S. A query algorithm subsequently receives a query vector $q \in \{0,1\}^d$ and in time T decides whether q is orthogonal to any of the input vectors x_i. We design a new deterministic data structure for $OnlineOV_{n,d}$. In low dimensions ($d = c\log n$), our data structure matches the performance of the best known randomized algorithm due to Chan [SoCG 2017]. Furthermore, in moderate dimensions ($d = n^{\varepsilon}$), we give the first improvement since Charikar, Indyk and Panigrahy [ICALP 2002]. This data structure also extends to a number of problems, including Partial Match, Orthogonal Range Search, and DNF Evaluation. Under the Non-Uniform Strong Exponential Time Hypothesis, we also prove arbitrarily large polynomial space lower bounds for any OnlineOV data structure with sublinear query time even with computationally unbounded preprocessing. These lower bounds extend to several other problems, including Polynomial Evaluation, Partial Match, Orthogonal Range Search, and Approximate Nearest Neighbors. We also prove similar lower bounds for 3-SUM with preprocessing under the Non-Uniform Hamiltonian Path Conjecture.

Karthik Gajulapalli, Alexander Golovnev, Samuel King, Sidhant Saraogi
Georgetown University
kg816@georgetown.edu, alexgolovnev@gmail.com, sik29@georgetown.edu, ss4456@georgetown.edu

CP35

Separations Between Oblivious and Adaptive Adversaries for Natural Dynamic Graph Problems

We establish the first update-time separation between dynamic algorithms against oblivious adversaries and those against adaptive adversaries in natural dynamic graph problems. Under the combinatorial BMM hypothesis, we show that every combinatorial algorithm against an adaptive adversary for the incremental maximal independent set problem requires $n^{1-o(1)}$ amortized update time. Further, assuming either the 3SUM or APSP hypotheses, every algorithm for the decremental maximal clique problem needs $\Delta/n^{o(1)}$ amortized update time when the initial maximum degree is $\Delta \leq \sqrt{n}$. These lower bounds are matched by existing algorithms against adaptive adversaries. In contrast, both problems admit algorithms against oblivious adversaries that achieve polylog(n) amortized update time. Therefore, our separations are exponential. Previously known separations for dynamic algorithms were either engineered for contrived problems and relied on strong cryptographic assumptions, or worked for problems whose inputs are not explicitly given but are accessed through oracle calls. As a byproduct, we also provide a separation between incremental and decremental algorithms for the triangle detection problem: we show a decremental algorithm with $\tilde{O}(n^{\omega})$ total update time, while every incremental algorithm requires $n^{3-o(1)}$ total update time, assuming the OMv hypothesis. To our knowledge this is the first separation of this kind.

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

Aaron Bernstein
Rutgers University
bernstei@gmail.com

Nick Fischer
Max Planck Institute for Informatics
nfischer@mpi-inf.mpg.de

Peter Kiss

University of Vienna
peter.kiss@univie.ac.at

Thatchaphol Saranurak
University of Michigan
thsa@umich.edu

CP35

Faster Algorithms for Global Minimum Vertex-Cut in Directed Graphs

We study the directed global minimum vertex-cut problem: given a directed vertex-weighted graph G, compute a vertex-cut (L, S, R) in G of minimum value, which is defined to be the total weight of all vertices in S. The fastest currently known algorithm for directed global minimum vertex-cut (Henzinger, Rao and Gabow, FOCS 1996 and J. Algorithms 2000) has running time $\tilde{O}(mn)$, where m and n denote the number of edges and vertices, respectively. A long line of work over the past decades led to faster algorithms for other main versions of the problem, including the undirected edge-based setting (Karger, STOC 1996 and J. ACM 2000), directed edge-based setting (Cen et al., FOCS 2021), and undirected vertex-based setting (Chuzhoy and Trabelsi, STOC 2025). However, for the vertex-based version in directed graphs, the 29 year-old $\tilde{O}(mn)$-time algorithm of Henzinger, Rao and Gabow remains the state of the art to this day, in all edge-density regimes. In this paper we break the $\Theta(mn)$ running time barrier for the first time, by providing a randomized algorithm for directed global minimum vertex-cut, with running time $O(mn^{0.976} \cdot \text{polylog} W)$ where W is the ratio of largest to smallest vertex weight. Additionally, we provide a randomized $O\left(\min\left\{m^{1+o(1)} \cdot k, n^{2+o(1)}\right\}\right)$-time algorithm for the unweighted version of directed global minimum vertex-cut, where k is the value of the optimal solution.

Ron Mosenzon
Toyota Technological Institute of Science
ron.mosenzon@gmail.com

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

Ohad Trabelsi
University of Haifa, Israel
otrabelsi@cs.haifa.ac.il

CP36

Matroids Are Equitable

We show that if the ground set of a matroid can be partitioned into $k \geq 2$ bases, then for any given subset S of the ground set, there is a partition into k bases such that the sizes of the intersections of the bases with S may differ by at most one. This settles the matroid equitability conjecture by Fekete and Szabo (Electron. J. Comb. 2011) in the affirmative. We also investigate equitable splittings of two disjoint sets S_1 and S_2, and show that there is a partition into k bases such that the sizes of the intersections with S_1 may differ by at most one and the sizes of the intersections with S_2 may differ by at most two; this is the best one can hope for arbitrary matroids. We also derive applications of this result into matroid constrained fair division problems. We show that there exists a matroid-constrained fair division that is envy-free up to one item if the valuations are identical and tri-valued additive. We also show that for bi-valued additive valuations, there exists a matroid-constrained allocation that provides everyone their maximin share.

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

Roshan Raj
TIFR Mumbai
amritanshus128@gmail.com

Laszlo Vegh
University of Bonn
lvegh@uni-bonn.de

CP36

Sample-Efficient Replicable Median in Polynomial Time

Replicable algorithm design has emerged as a central notion in algorithmic stability and the reliability of scientific findings. Introduced by Impagliazzo, Lei, Pitassi, and Sorrell [STOC22] and by Ghazi, Kumar, and Manurangsi [NeurIPS21], it has since been studied extensively in connection with privacy, generalization, and high-dimensional learning. An algorithm is *replicable* if, with high probability, it produces the same output on independent datasets when run with the same randomnessstrengthening both differential privacy and adaptive generalization. A central open problem in this area is *replicable median estimation*. Prior works are either computationally inefficient or require sample complexity exponential in $\log^* |\mathcal{X}|$, even though the information-theoretic lower bound is polynomial. We show that replicable median estimation reduces to the *replicable interior point* problem, for which we design a polynomial-time algorithm with sample complexity $\text{poly}(\log^* |\mathcal{X}|)$. This yields a polynomial-time replicable algorithm for median estimation, PAC learning of thresholds, and distribution learning under Kolmogorov distance. A new technical tool we develop, which may be of independent interest, is *semi-replicable recursion*, which allows recursive algorithms to maintain replicability despite data-dependent subcalls by controlling distributional drift across recursion.

Kiarash Banihashem
University of Maryland
kiarash@umd.edu

MohammadHossein Bateni, Hossein Esfandiari
Google Research
bateni@google.com, esfandiari@google.com

Samira Goudarzi, MohammadTaghi Hajiaghayi
University of Maryland
samirag@umd.edu, hajiaghayi@gmail.com

CP36

On the Structure of Replicable Hypothesis Testers

A hypothesis testing algorithm is *replicable* (Impagliazzo, Lei, Pitassi, and Sorell [STOC'22]) if, when run on two different samples from the same distribution, it produces the same output with high probability. We build general tools to prove lower and upper bounds on the sample complexity of replicable testers, unifying and quantitatively improving upon existing results. We identify a set of natural canonical properties, and prove that any replicable testing algorithm can be modified to satisfy these properties without worsening accuracy or sample complexity. We prove new lower bounds for uniformity, identity, and closeness testing by reducing to the case where the replicable algorithm satisfies these canonical properties, thus resolving an open question by Liu and Ye [NeurIPS'24]. We systematize and improve upon a common strategy for replicable algorithm design based on test statistics with known expectation and bounded variance. Our framework allows extensively studied non-replicable testers to be made replicable with minimal overhead. In particular, we can apply our framework to obtain constant-factor optimal bounds for coin testing and closeness testing and get replicability for free for uniformity testing in a large parameter regime. We also give new algorithms and lower bounds for replicable Gaussian mean testing, improving over the previous best sample complexity of Bun, Gaboardi, Hopkins, Impagliazzo, Lei, Pitassi, Sivakumar, and Sorrell [STOC'23].

Anders Aamand
Massachusetts Institute of Technology, U.S.
andersaamanda@gmail.com

Maryam Aliakbarpour
Rice University
maryama@rice.edu

Justin Y. Chen
MIT
justc@mit.edu

Shyam Narayanan
Citadel Securities
shyam.s.narayanan@gmail.com

Sandeep Silwal
University of Wisconsin-Madison
silwal@cs.wisc.edu

CP36

Differentially Private Quasi-Concave Optimization: Bypassing the Lower Bound and Application to Geometric Problems

We study the sample complexity of differentially private optimization of quasi-concave functions. For a fixed input domain $\mathcal{X}$, Cohen et al. (STOC 2023) proved that any generic private optimizer for low sensitive quasi-concave functions must have sample complexity $\Omega(2^{\log^* |X|})$. We show that the lower bound can be bypassed for a series of "natural' problems. We define a new class of *approximated* quasi-concave functions, and present a generic differentially private optimizer for approximated quasi-concave functions with sample complexity $\tilde{O}(\log^* |X|)$. As applications, we use our optimizer to privately select a center point of points in d dimensions and *probably approximately correct* (PAC) learn d-dimensional halfspaces. In previous works, Bun et al. (FOCS 2015) proved a lower bound of $\Omega(\log^* |X|)$ for both problems. Beimel et al. (COLT 2019) and Kaplan et al. (NeurIPS 2020) gave an upper bound of $\tilde{O}(d^{2.5} \cdot 2^{\log^* |X|})$ for the two problems, respectively. We improve the dependency of the upper bounds on the cardinality of the domain by presenting a new upper bound of $\tilde{O}(d^{5.5} \cdot \log^* |X|)$ for both problems. To the best of our understanding, this is the first work to reduce the sample complexity dependency on $|X|$ for these two problems from exponential in $\log^* |X|$ to $\log^* |X|$.

Kobbi Nissim
Georgetown University
kobbi.nissim@georgetown.edu

Eliad Tsfadia
Bar-Ilan University
eliadtsfadia@gmail.com

Chao Yan
Georgetown University
Georgetown University
cy399@georgetown.edu

CP36

Differentially Private Algorithms for Graph Cuts: A Shifting Mechanism Approach and More

We address the challenge of differential privacy in the context of graph cuts, specifically focusing on the multiway cut and the minimum k-cut. We introduce edge-differentially private algorithms that achieve nearly optimal performance for these problems. Motivated by multiway cut, we propose the shifting mechanism, a general framework for private combinatorial optimization problems. This framework allows us to develop an efficient private algorithm with a multiplicative approximation ratio that matches the state-of-the-art non-private algorithm, improving over previous private algorithms that have provably worse multiplicative loss. We then provide a tight information-theoretic lower bound on the additive error, demonstrating that for constant k, our algorithm is optimal in terms of the privacy cost. The shifting mechanism also allows us to design private algorithm for the multicut and max-cut problems, with runtimes determined by the best non-private algorithms for these tasks. For the minimum k-cut problem we combine the exponential mechanism with bounds on the number of approximate k-cuts to get the first private algorithm with optimal additive error of $O(k \log n)$ (for a fixed privacy parameter). We also establish an information-theoretic lower bound that matches this additive error. Furthermore, we provide an efficient private algorithm even for non-constant k.

Zongrui Zou
Nanjing University
zou.zongrui@smail.nju.edu.cn

Rishi Chandra, Michael Dinitz
Johns Hopkins University
rishichandra159@gmail.com, mdinitz@cs.jhu.edu

Chenglin Fan
Sorbonne Université
fanchenglin@gmail.com

CP37

A Near-Complete Resolution of the Exponential-Time Complexity of k-opt for the Traveling Salesman Problem

The k-opt algorithm is one of the simplest and most widely used heuristics for solving the traveling salesman problem. Starting from an arbitrary tour, the k-opt algorithm im-

proves the current tour in each iteration by exchanging up to k edges. The algorithm continues until no further improvement of this kind is possible. For a long time, it remained an open question how many iterations the k-opt algorithm might require for small values of k, assuming the use of an optimal pivot rule. In this paper, we resolve this question for the cases $k = 3$ and $k = 4$ by proving that in both these cases an exponential number of iterations may be needed even if an optimal pivot rule is used. Combined with a recent result [Heimann, Hoang, Hougardy, ICALP 2024], this provides a complete answer for all $k \geq 3$ regarding the number of iterations the k-opt algorithm may require under an optimal pivot rule. In addition we establish an analogous exponential lower bound for the 2.5-opt algorithm, a variant that generalizes 2-opt and is a restricted version of 3-opt. All our results hold for both the general and the metric traveling salesman problem.

Sophia Heimann
University of Bonn
sheimann@uni-bonn.de

Hung Hoang
TU Wien
phoang@ac.tuwien.ac.at

Stefan Hougardy
University of Bonn
hougardy@or.uni-bonn.de

CP37

Local Search for Clustering in Almost-Linear Time

We propose the first local search algorithm for Euclidean clustering that attains an $O(1)$-approximation in almost-linear time. Specifically, for Euclidean k-Means, our algorithm achieves an $O(c)$-approximation in $\tilde{O}(n^{1+1/c})$ time, for any constant $c \geq 1$, maintaining the same running time as the previous (non-local-search-based) approach [la Tour and Saulpic, arXiv'2407.11217] while improving the approximation factor from $O(c^6)$ to $O(c)$. The algorithm generalizes to any metric space with sparse spanners, delivering efficient constant approximation in ℓ_p metrics, doubling metrics, Jaccard metrics, etc. This generality derives from our main technical contribution: a local search algorithm on general graphs that obtains an $O(1)$-approximation in almost-linear time. We establish this through a new 1-swap local search framework featuring a novel swap selection rule. At a high level, this rule "scores" every possible swap, based on both its modification to the clustering and its improvement to the clustering objective, and then selects those high-scoring swaps. To implement this, we design a new data structure for maintaining approximate nearest neighbors with amortized guarantees tailored to our framework.

Shaofeng Jiang
Peking University
shaofeng.jiang@pku.edu.cn

YAONAN Jin
Huaweis Taylor Lab
jinyaonan@huawei.com

Jianing Lou
Peking University
loujn@pku.edu.cn

Pinyan Lu
Shanghai University of Finance and Economics
lu.pinyan@mail.shufe.edu.cn

CP37

Approximate Counting of Permutation Patterns

We consider the problem of counting the copies of a length-k pattern σ in a sequence $f\colon [n] \to \mathbb{R}$, where a copy is a subset of indices $i_1 < \ldots < i_k \in [n]$ such that $f(i_j) < f(i_\ell)$ if and only if $\sigma(j) < \sigma(\ell)$. This problem has a range of applications in ranking, nonparametric statistics, combinatorics, and fine-grained complexity. Recent advances have significantly improved our understanding of counting and detecting patterns. Guillemot and Marx [2014] obtained an $O(n)$ time algorithm for the detection variant for any fixed k. Counting, in contrast, is harder: it has a conditional lower bound of $n^{\Omega(k/\log k)}$ [Berendsohn, Kozma, and Marx, 2019] and is expected to be polynomially harder than detection as early as $k = 4$, given its equivalence to counting 4-cycles in graphs [Dudek and Gawrychowski, 2020]. We design a deterministic near-linear time $(1+\varepsilon)$-approximation algorithm for counting σ-copies in f for all $k \leq 5$. Combined with the conditional lower bound for $k = 4$, this establishes the first known separation between approximate and exact pattern counting. Interestingly, while neither the sequence f nor the pattern σ are monotone, our algorithm makes extensive use of coresets for monotone functions [Har-Peled, 2006]. Along the way, we develop a near-optimal data structure for $(1+\varepsilon)$-approximate increasing pair range queries in the plane.

Omri Ben-Eliezer
Technion
omribene@cs.technion.ac.il

Slobodan Mitrovic
UC Davis
smitrovic@ucdavis.edu

Pranjal Srivastava
Massachusetts Institute of Technology
pranjal0@mit.edu

CP37

Optimal Rounding for Two-Stage Bipartite Matching

We study two-stage bipartite matching, in which the edges of a bipartite graph on nodes $(B_1 \cup B_2, I)$ are revealed in two batches. In stage one, a matching is selected among edges $E \subseteq B_1 \times I$. In stage two, edges $E^\theta \subseteq B_2 \times I$ are sampled from a known distribution, and a second matching is selected between B_2 and unmatched vertices in I. The objective is to maximize the weight of the combined matching. We design polynomial-time approximations to the optimal online algorithm, achieving guarantees of 7/8 for vertex-weighted and $2\sqrt{2}-2 \approx 0.828$ for edge-weighted graphs under arbitrary distributions. Both ratios match known upper bounds on the integrality gap of the natural fractional relaxation, improving upon the 0.767 approximation of Feng, Niazadeh, and Saberi [Two-Stage Stochastic Matching and Pricing with Applications to Ride Hailing, 2021] for unweighted graphs whose second batch consists of independently arriving nodes. Our results are achieved via an algorithm that rounds a fractional matching revealed in two stages, matching offline nodes (or edges) with probability proportional to their fractional weights up to a constant-factor loss. We leverage negative association

(NA) of offline node availabilitiesa property induced by dependent roundingand derive new lower bounds on the expected size of the max weight matching in random bipartite graphs where one side is realized via NA binary random variables.

Anders Wikum, Tristan Pollner
Stanford University
wikum@stanford.edu, tpollner@stanford.edi

Amin Saberi
Management Science and Engineering
Stanford University
saberi@stanford.edu

CP37

Sublinear Metric Steiner Forest Via Maximal Independent Set

In this work we consider the Metric Steiner Forest problem in the sublinear time model. Given a set V of n points in a metric space where distances are provided by means of query access to an $n \times n$ distance matrix, along with a set of k terminal pairs $(s_1, t_1), \ldots, (s_k, t_k) \in V \times V$, the goal is to find a minimum-weight subset of edges that connects each terminal pair. Although sublinear time algorithms have been studied for estimating the weight of a minimum spanning tree in both general and metric settings, as well as for the metric Steiner Tree problem, no sublinear time algorithm was known for the metric Steiner Forest problem. Here, we give an $O(\log k)$-approximation algorithm for the problem that runs in time $\widetilde{O}(n^{3/2})$. Along the way, we provide the first sublinear-time algorithm for estimating the size of a Maximal Independent Set (MIS). Our algorithm runs in time $\widetilde{O}(n^{3/2}/\varepsilon^2)$ under the adjacency matrix oracle model and obtains a purely multiplicative $(1+\varepsilon)$-approximation. Previously, sublinear-time algorithms for MIS were only known for bounded-degree graphs.

Mohammad Roghani
Stanford University
roghani@stanford.edu

Sepideh Mahabadi, Jakub Tarnawski
Microsoft Research
smahabadi@microsoft.com,
jakub.tarnawski@microsoft.com

Ali Vakilian
Virginia Tech
vakilian@vt.edu

CP37

Local Gibbs Sampling Beyond Local Uniformity

Local samplers are algorithms that generate random samples based on local queries to high-dimensional distributions, ensuring the samples follow the correct induced distributions while maintaining time complexity that scales locally with the query size. These samplers have broad applications, including deterministic approximate counting [He, Wang, Yin, SODA '23; Feng et.al., FOCS '23], sampling from infinite or high-dimensional Gibbs distributions [Anand, Jerrum, SICOMP '22; He, Wang, Yin, FOCS '22], and providing local access to large random objects [Biswas, Rubinfield, Yodpinyanee, ITCS '20]. Specifically, we design linear-time local samplers for: - spin systems with soft constraints, including the first local sampler for near-critical Ising models; - truly repulsive spin systems, represented by the first local sampler for uniform proper q-colorings, with $q = O(\Delta)$ colors on graphs with maximum degree Δ. These local samplers are efficient beyond the "local uniformity' threshold, which imposes unconditional marginal lower bounds — a key assumption required by all prior local samplers. Our results show that, in general, local sampling is not significantly harder than global sampling for spin systems. As an application, our results also imply local algorithms for probabilistic inference in the same near-critical regimes.

Chunyang Wang
Nanjing University
Nanjing University
wcysai@smail.nju.edu.cn

Hongyang Liu, Yitong Yin
Nanjing University
liuhongyang@smail.nju.edu.cn, yinyt@nju.edu.cn

CP38

Sparsifying Sums of Positive Semidefinite Matrices

In this paper, we revisit spectral sparsification for sums of arbitrary positive semidefinite (PSD) matrices. This generalizes spectral sparsification of graphs which corresponds to the matrices being the set of Laplacians of edges. It also captures sparsifying Cayley graphs by choosing a subset of generators. The former has been extensively studied with optimal sparsifiers known. The latter has received attention recently and was solved forCayley graph over $\mathbb{F}_2^n$. Prior work shows any sum of PSD matrices can be sparsified down to $O(n)$ elements. This bound however yields no non-trivial bound for building Cayley sparsifiers for Cayley graphs. In this work, we develop a new, instance-specific theory of PSD matrix sparsification based on a new parameter $N^*(\mathcal{A})$, where $\mathcal{A}$ is our instance of PSD matrices, called connectivity threshold that generalizes the threshold of the number of edges required to make a graph connected. Our main result gives a sparsifier that uses at most $O(\varepsilon^{-2} N^*(\mathcal{A})(\log n)(\log r))$ matrices and is constructible in randomized polynomial time. We also show that we need $N^*(\mathcal{A})$ elements to sparsify for any $\varepsilon < 0.99$. As the main application of our framework, we prove that any Cayley graph can be sparsified to $O(\varepsilon^{-2} \log^4 N)$ generators. Previously, a non-trivial bound on Cayley sparsifiers was known only in the case when the group is $\mathbb{F}_2^n$.

Arpon Basu, Pravesh Kothari
Princeton University
arpon.basu@princeton.edu, kothari@cs.princeton.edu

Yang Liu
CMU
yangl7@andrew.cmu.edu

Raghu Meka
University of California, Los Angeles
raghum@cs.ucla.edu

CP38

Recognizing Leaf Powers and Pairwise Compatibility Graphs is NP-Complete

A graph $G = (V, E)$ is a k-leaf power if there is a tree T whose leaves are the vertices of G with the property that a pair of leaves u and v induce an edge in G if and only if they are distance at most k apart in T. For $k \leq 4$, it

is known that there exists a finite set $\mathcal{F}_k$ of graphs such that the class $\mathcal{L}(k)$ of k-leaf power graphs is characterized as the set of strongly chordal graphs that do not contain any graph in $\mathcal{F}_k$ as an induced subgraph. We prove no such characterization holds for $k \geq 5$. That is, for any $k \geq 5$, there is no finite set $\mathcal{F}_k$ of graphs such that $\mathcal{L}(k)$ is equivalent to the set of strongly chordal graphs that do not contain as an induced subgraph any graph in $\mathcal{F}_k$.

Max Dupré La Tour, Ndiamé Ndiaye
McGill University
maxduprelatour@gmail.com,
ndiame.ndiaye@mail.mcgill.ca

Manuel Lafond
University of Sherbrooke
manuel.lafond@USherbrooke.ca

CP38

Three-Edge-Coloring (Tait Coloring) Cubic Graphs and Nowhere-Zero 4-Flow for Graphs on the Torus

We prove that every cyclically 4-edge-connected cubic graph that can be embedded in the torus, with the exceptional graph class called "Petersen-like", is 3-edge-colorable. This means every (non-trivial) toroidal snark can be obtained from several copies of the Petersen graph using the dot product operation. The Petersen-like cubic graphs fall into two infinite families, one that was discovered by Vodopivec in 2008 and one obtained by Belcastro and Kaminski. This proves a strengthening of the well-known, long-standing conjecture of Grnbaum [Grnbaum, Conjecture 6, in Recent Progress in Combinatorics] from 1968. This implies that a 2-connected cubic (multi)graph that can be embedded in the torus is not 3-edge-colorable if and only if it can be obtained from a dot product of copies of the Petersen graph by replacing its vertices with 2-edge-connected planar cubic (multi)graphs. This result is a nontrivial generalization of the Four Color Theorem, and its proof requires a combination of extensive computer verification and computer-free extension of existing proofs on colorability. An important consequence is a strong version of the Tutte 4-Flow Conjecture for toroidal graphs. We show that a 2-edge connected graph embedded in the torus admits a nowhere-zero 4-flow unless it is Petersen-like (in which case it does not admit nowhere-zero 4-flows).

Yuta Inoue
The University of Tokyo
yutainoue@is.s.u-tokyo.ac.jp

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

Miyashita Atsuyuki
The University of Tokyo
miyashita-atsuyuki869@is.s.u-tokyo.ac.jp

Bojan Mohar
Simon Fraser University
mohar@sfu.ca

Tomohiro Sonobe
National Institute of Informatics
tomohiro_sonobe@nii.ac.jp

CP38

Dynamic Connectivity with Expected Polylogarithmic Worst-Case Update Time

Whether a graph $G = (V, E)$ is connected is arguably its most fundamental property. Naturally, connectivity was the first characteristic studied for dynamic graphs, i.e. graphs that undergo edge insertions and deletions. While connectivity algorithms with polylogarithmic amortized update time have been known since the 90s, achieving worst-case guarantees has proven more elusive. Two recent breakthroughs have made important progress on this question: (1) Kapron, King and Mountjoy [SODA'13; Best Paper] gave a Monte-Carlo algorithm with polylogarithmic worst-case update time, and (2) Nanongkai, Saranurak and Wulff-Nilsen [STOC'17, FOCS'17] obtained a Las-Vegas data structure, however, with subpolynomial worst-case update time. Their algorithm was subsequently derandomized [FOCS'20]. In this article, we present a new dynamic connectivity algorithm based on the popular core graph framework that maintains a hierarchy interleaving vertex and edge sparsification. Previous dynamic implementations of the core graph framework required subpolynomial update time. In contrast, we show how to implement it for dynamic connectivity with polylogarithmic expected worst-case update time.

Simon Meierhans, Maximilian Probst Gutenberg
ETH Zurich
mesimon@inf.ethz.ch, maximilian.probst@inf.ethz.ch

CP38

Sparsifying Cayley Graphs on Every Group

A classic result in graph theory, due to Batson, Spielman, and Srivastava (STOC 2009) shows that every graph admits a $(1 \pm \varepsilon)$ cut (or spectral) sparsifier which preserves only $O(n/\varepsilon^2)$ reweighted edges. However, when applying this result to *Cayley graphs*, the resulting sparsifier is no longer necessarily a Cayley graph — it can be an arbitrary subset of edges. Thus, a recent line of inquiry, and one which has only seen minor progress, asks: for any group G, do all Cayley graphs over the group G admit sparsifiers which preserve only $\text{polylog}(|G|)/\varepsilon^2$ many reweighted generators? As our primary contribution, we answer this question in the affirmative, presenting a proof of the existence of such Cayley graph spectral sparsifiers, along with an efficient algorithm for finding them. Our algorithm even extends to *directed* Cayley graphs, if we instead ask only for cut sparsification instead of spectral sparsification. We additionally study the sparsification of linear equations over non-abelian groups. In contrast to the abelian case, we show that for non-abelian valued equations, super-polynomially many linear equations must be preserved in order to approximately preserve the number of satisfied equations for any input. Together with our Cayley graph sparsification result, this provides a formal separation between Cayley graph sparsification and sparsifying linear equations.

Jun-Ting Hsieh, Daniel Z. Lee
MIT
juntingh@mit.edu, lee_d@mit.edu

SIDHANTH Mohanty
Northwestern University
sidm@mit.edu

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

RACHEL Yun Zhang
MIT
rachelyz@mit.edu

CP38

Strongly Polynomial Parallel Work-Depth Tradeoffs for Directed SSSP

In this paper, we show new strongly polynomial work-depth tradeoffs for computing single-source shortest paths (SSSP) in non-negatively weighted directed graphs in parallel. Most importantly, we prove that directed SSSP can be solved within $\widetilde{O}(m + n^{2-\epsilon})$ work and $\widetilde{O}(n^{1-\epsilon})$ depth for some positive $\epsilon > 0$. In particular, for dense graphs with non-negative real weights, we provide the first nearly work-efficient strongly polynomial algorithm with sublinear depth. Our result immediately yields improved strongly polynomial parallel algorithms for min-cost flow and the assignment problem. It also leads to the first non-trivial strongly polynomial dynamic algorithm for minimum mean cycle. Moreover, we develop efficient parallel algorithms in the Word RAM model for several variants of SSSP in graphs with exponentially large edge weights.

Marek Sokolowski
Max Planck Institute for Informatics
Saarland Informatics Campus, Germany
msokolow@mpi-inf.mpg.de

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

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

CP38

Coloring Graphs with Few Colors in the Streaming Model

We study graph coloring problems in the streaming model, wherein the goal is to process an n-vertex graph whose edges arrive in a stream, using a limited space that is much smaller than the trivial $O(n^2)$ bound. While prior work has largely focused on coloring graphs with a large number of colors we explore the opposite end of the spectrum: deciding whether the input graph can be colored using only a few, say, a constant number of colors. We are interested in each of the adversarial, random order, or dynamic streams, and focus solely on the space complexity rather than running time. Some of our main results include: - Adversarial: for distinguishing between q- vs $2^{\Omega(q)}$-colorable graphs, lower bounds of $n^{2-o(1)}$ space for q up to $(\log n)^{1/2-o(1)}$, and $n^{1+\Omega(1/\log\log n)}$ space for q further up to $(\log n)^{1-o(1)}$. - Random order: for distinguishing between q- vs q^t-colorable graphs for $q, t \geq 2$, an upper bound of $\widetilde{O}(n^{1+1/t})$ space. - Dynamic: for distinguishing between q- vs $q \cdot t$-colorable graphs for any $q \geq 3$ and $t \geq 1$, nearly optimal upper and lower bounds of $\widetilde{\Theta}(n^2/t^2)$ space. We develop several new technical tools along the way: cluster packing graphs, a generalization of Ruzsa-Szemerdi graphs; a player elimination framework based on cluster packing graphs; and new edge and vertex sampling lemmas tailored to graph coloring.

Sepehr Assadi, Janani Sundaresan, Helia Yazdanyar
University of Waterloo
sepehr@assadi.info, jsundaresan@uwaterloo.ca, hyazdanyar@uwaterloo.ca

CP39

Listing Faces of Polytopes

This paper investigates the problem of listing faces of combinatorial polytopes, such as hypercubes, permutahedra, associahedra, and their generalizations. First, we consider the face lattice, which is the inclusion order of all faces of a polytope, and we seek a Hamiltonian cycle in its cover graph. We construct such Hamiltonian cycles for hypercubes, permutahedra, B-permutahedra, associahedra, cyclic polytopes, 3-dimensional polytopes, graph associahedra of chordal graphs, and quotientopes. Secondly, we consider facet-Hamiltonian cycles, which are cycles on the skeleton of a polytope that enter and leave every facet exactly once. This notion was recently introduced by Akitaya et al., where the authors conjectured that B-permutahedra admit a facet-Hamiltonian cycle for all dimensions. We construct such facet-Hamiltonian cycles in this paper, thus establishing their conjecture as a theorem. A key tool are so-called rhombic strips, certain planar spanning subgraphs of the cover graph of the face lattice. Specifically, we construct a rhombic strip in the face lattice of the hypercube of any dimension, and characterize the existence of rhombic strips in the face lattice of 3-dimensional polytopes. Our constructions yield time- and space-efficient algorithms for computing the aforementioned cycles and thus for listing the corresponding combinatorial objects, including ordered set partitions and dissections of a convex polygon.

Sofia Brenner
University of Kassel
sbrenner@mathematik.uni-kassel.de

Nastaran Behrooznia
University of Warwick
nastaran.behrooznia@warwick.ac.uk

Arturo Merino
Universidad de OHiggins, Rancagua
arturo.merino@uoh.cl

Torsten Mütze, Christian Rieck, Francesco Verciani
University of Kassel
tmuetze@mathematik.uni-kassel.de,
christian.rieck@mathematik.uni-kassel.de,
francesco.verciani@mathematik.uni-kassel.de

CP39

A Near-Optimal Quadratic Goldreich-Levin Algorithm

We present a quadratic Goldreich-Levin algorithm that is nearly optimal in the following ways. Given a bounded function f on the Boolean hypercube $\mathbb{F}_2^n$ and any $\varepsilon > 0$, the algorithm returns a quadratic polynomial $q : \mathbb{F}_2^n \to \mathbb{F}_2$ so that the correlation of f with the function $(-1)^q$ is within an additive ε of the maximum possible correlation with a quadratic phase function. The algorithm runs in $O_\varepsilon(n^3)$ time and makes $O_\varepsilon(n^2 \log n)$ queries to f, which matches the information-theoretic lower bound of $\Omega(n^2)$ queries up to a logarithmic factor. As a result, we obtain a number of corollaries: - A near-optimal self-corrector of quadratic

Reed-Muller codes. - An algorithmic polynomial inverse theorem for the Gowers U^3-norm. - An efficient structure-vs-randomness decomposition algorithm. Our algorithm is obtained using ideas from recent work on quantum learning theory. Its construction deviates from previous approaches based on algorithmic proofs of the inverse theorem for the Gowers U^3-norm (and in particular does not rely on the recent resolution of the polynomial Freiman-Ruzsa conjecture).

Jop Briët
CWI
jop@cwi.nl

Davi Castro-Silva
Centrum Wiskunde & Informatica (CWI)
davisilva15@gmail.com

CP39

Braiding Vineyards

In this work, we link two areas in computational topology, namely topological data analysis (TDA) and knot theory. For a manifold $\mathcal{M} \subset \mathbb{R}^d$ and a closed curve or loop $\gamma \subset \mathbb{R}^d$, we consider the vineyard of the Euclidean distance to $\gamma(t)$, restricted to the manifold $\mathcal{M}$, i.e. $d_{\mathbb{E}}(\cdot, \gamma(t))|_{\mathcal{M}}$. Because γ is a loop, we can identify the ends of the vineyard, yielding what we call a closed vineyard. We show that given a knot (and even a link) and an $l \in \mathbb{Z}_{\geq 0}$, we can construct a manifold and loop such that the closed vineyard (in l-persistence) contains the given knot (or link). This shows that vineyards are topologically as rich as one could possibly hope. Importantly, it has at least two immediate consequences we explicitly point out:

- Monodromy of any periodicity can occur in a l-vineyard for any l. This answers a variant of a question by Arya and collaborators. To exhibit this as a consequence of our first main result we also reformulate monodromy in a more geometric way, which may be of interest in itself.
- Topologically distinguishing closed vineyards is likely to be difficult (from a complexity theoretical as well as practical perspective) because of the difficulty of knot and link recognition, which have strong connections to many NP-hard problems.

Erin Chambers
Department of Computer Science and Engineering
University of Notre Dame
echambe2@nd.edu

Christopher Fillmore
Institute of Science and Technology Austria
cdfillmore@gmail.com

Elizabeth Stephenson
Orteliu
elizasteprene@gmail.com

Mathijs Wintraecken
Centre Inria d'Université Côte d'Azur , Sophia Antipolis
mathijs.wintraecken@inria.fr

CP39

Rumour Spreading Depends on the Latent Geometry and Degree Distribution in Social Network Models

We study push-pull rumour spreading in ultra-small-world models for social networks where the degrees follow a power-law distribution. In a non-geometric setting, Fountoulakis, Panagiotou and Sauerwald have shown that rumours always spread ultra-fast (SODA 2012). On the other hand, Janssen and Mehrabian have found that rumours spread slowly in a spatial preferential attachment model (SIDMA 2017). We study the question systematically for the model of Geometric Inhomogeneous Random Graphs (GIRGs), which has been found to be a good theoretical and empirical fit for social networks. We show that slow, fast and ultra-fast (i.e., polynomial, polylogarithmic and doubly logarithmic number of rounds) rumour spreading may occur, depending on the exponent of the power law and the strength of the geometry in the network, and we fully characterise the phase boundaries between these regimes. The regimes do not coincide with the graph distance regimes, i.e., polylogarithmic or even polynomial rumour spreading may occur even if graph distances are doubly logarithmic. We expect these results to hold with little effort for related models, e.g. Scale-Free Percolation.

Marc Kaufmann, Kostas Lakis, Johannes Lengler, Raghu Raman Ravi, Ulysse Schaller, Konstantin Sturm
ETH Zurich
marc.kaufmann@inf.ethz.ch, konstantinos.lakis@inf.ethz.ch, johannes.lengler@inf.ethz.ch, raghu.ravi@inf.ethz.ch, ulysse.schaller@inf.ethz.ch, konstantin.sturm@inf.ethz.ch

CP39

Evasive Sets, Twisted Varieties, and Container-Clique Trees

In the affine space $\mathbb{F}_q^n$ over the finite field of order q, a point set S is said to be (d,k,r)-evasive if the intersection between S and any variety, of dimension k and degree at most d, has cardinality less than r. As q tends to infinity, the size of a (d,k,r)-evasive set in $\mathbb{F}_q^n$ is at most $O\left(q^{n-k}\right)$ by a simple averaging argument. We exhibit the existence of such evasive sets of sizes at least $\Omega\left(q^{n-k}\right)$ for much smaller values of r than previously known constructions, and establish an enumerative upper bound $2^{O(q^{n-k})}$ for the total number of such evasive sets. The existence result is based on our study of twisted varieties. In the projective space $\mathbb{P}^n$ over an algebraically closed field, a variety V is said to be d-twisted if the intersection between V and any variety, of dimension $n-\dim(V)$ and degree at most d, has dimension zero. We prove an upper bound on the smallest possible degree of twisted varieties which is best possible in a mild sense. The enumeration result includes a new technique for the container method which we believe is of independent interest. To illustrate the potential of this technique, we give a simpler proof of a result by Chen–Liu–Nie–Zeng that characterizes the maximum size of a collinear-triple-free subset in a random sampling of $\mathbb{F}_q^2$ up to polylogarithmic factors.

Jeck Lim
California Institute of Technology
limjeck@yahoo.com.sg

Jiaxi Nie
Georgia Institute of Technology
jnie47@gatech.edu

Ji Zeng

Alfréd Rényi Institute of Mathematics
zeng.ji@renyi.hu

CP39

Extended Vc-Dimension, and Radon and Tverberg Type Theorems for Unions of Convex Sets

We define and study an extension of the notion of the VC-dimension of a hypergraph and apply it to establish a Tverberg type theorem for unions of convex sets. We also prove a new Radon type theorem for unions of convex sets and settle a well-known open problem posed by Kalai in the 1970s.

Shakhar Smorodinsky
Ben Gurion University
shakhar@bgu.ac.il

Noga Alon
Princeton University
nalon@math.princeton.edu

CP40

On the Usefulness of Promises

A Boolean predicate A is defined to be promise-useful if PCSP(A, B) is tractable for some non-trivial Boolean predicate B and otherwise it is promise-useless. We initiate investigations of this notion and derive sufficient conditions for both promise-usefulness and promise-uselessness (assuming P $\neq$ NP). While we do not obtain a complete characterization, our conditions are sufficient to classify all predicates of arity at most 4 and almost all predicates of arity 5. We also derive asymptotic results to show that for large arities a vast majority of all predicates are promise-useless. Our results are primarily obtained by a thorough study of the "Promise-SAT" problem, in which we are given a k-SAT instance with the promise that there is a satisfying assignment for which the literal values of each clause satisfy some additional constraint. The algorithmic results are based on the basic LP + affine IP algorithm of Brakensiek et al. (SICOMP, 2020) while we use a number of novel criteria to establish NP-hardness.

Per Austrin, Johan Håstad, Björn Martinsson
KTH Royal Institute of Technology
austrin@kth.se, johanh@kth.se, bmart@kth.se

CP40

Problems from Optimization and Computational Algebra Equivalent to Hilberts Nullstellensatz

Efficient algorithms for many problems in optimization and computational algebra often arise from casting them as systems of polynomial equations. Blum, Shub, and Smale formalized this as Hilbert's Nullstellensatz Problem HN_R: given multivariate polynomials over a ring R, decide whether they have a common solution in R. We can also view HN_R as a complexity class by taking the downward closure of the problem HN_R under polynomial-time many-one reductions. In this work, we show that many important problems from optimization and algebra are complete or hard for this class. We first consider the Affine Polynomial Projection Problem: given polynomials f, g, does an affine projection of the variables transform f into g? We show that this problem is at least as hard as HN_F for any field F. Then we consider the Sparse Shift Problem: given a polynomial, can its number of monomials be reduced by an affine shift of the variables? Prior HN_R-hardness for this problem was known for non-field integral domains R, which we extend to fields. For the special case of the real field, HN captures the existential theory of the reals and its complement captures the universal theory of the reals. We prove that the problems of deciding real stability, convexity, and hyperbolicity of a given polynomial are all complete for the universal theory of the reals, thereby pinning down their exact complexity.

Sagnik Dutta
Max Planck Institut für Informatik
duttasagnik436@gmail.com

Markus Bläser
Saarland University
mblaeser@cs.uni-saarland.de

Gorav Jindal
University of Warsaw
gjindal@mimuw.edu.pl

CP40

Lower Bounds for Csp Hierarchies Through Ideal Reduction

We present a generic way to obtain level lower bounds for (promise) CSP hierarchies from degree lower bounds for algebraic proof systems. More specifically, we show that pseudo-reduction operators in the sense of Alekhnovich and Razborov [Proc. Steklov Inst. Math. 2003] can be used to fool the cohomological k-consistency algorithm. As applications, we prove optimal level lower bounds for c vs. ℓ-coloring for all $\ell \geq c \geq 3$, and give a simplified proof of the lower bounds for lax and null-constraining CSPs of Chan and Ng [STOC 2025].

Jonas Conneryd
Lund University
jonas.conneryd@cs.lth.se

Yassine Ghannane
University of Copenhagen
Lund University
yagh@di.ku.dk

Shuo Pang
University of Bristol
s.pang@bristol.ac.uk

CP40

Numerical Linear Algebra in Linear Space

We present a randomized linear-space solver for general linear systems $\mathbf{Ax} = \mathbf{b}$ with $\mathbf{A} \in \mathbb{Z}^{n \times n}$ and $\mathbf{b} \in \mathbb{Z}^n$, without any assumption on the condition number of $\mathbf{A}$. For matrices whose entries are bounded by poly(n), the solver returns a $(1 + \epsilon)$-multiplicative entry-wise approximation to vector $\mathbf{x} \in \mathbb{Q}^n$ using $\widetilde{O}(n^2 \cdot \mathrm{nnz}(\mathbf{A}))$ bit operations and $O(n \log n)$ bits of working space (i.e., linear in the size of a vector), where nnz denotes the number of nonzero entries. Our solver works for right-hand vector $\mathbf{b}$ with entries up to $n^{O(n)}$. To our knowledge, this is the first linear-space linear system solver over the rationals that runs in $\widetilde{O}(n^2 \cdot \mathrm{nnz}(\mathbf{A}))$ time. We also present several applications of our solver to numerical linear algebra problems, for which we provide algorithms with efficient polynomial running time and near-linear space. In particular, we present results for linear

regression, linear programming, eigenvalues and eigenvectors, and Singular Value Decomposition.

Yiping Liu
Tsinghua University
liu-yp23@mails.tsinghua.edu.cn

Hoai-An Nguyen
Carnegie Mellon University
hnnguyen@andrew.cmu.edu

Junzhao Yang
Institute for Interdisciplinary Information Sciences
Tsinghua University
junzhaoy@andrew.cmu.edu

CP40

On Deterministically Finding An Element of High Order Modulo a Composite

We give a deterministic algorithm that, given a composite number N and a target order $D \geq N^{1/6}$, runs in time $D^{1/2+o(1)}$ and finds either an element $a \in \mathbb{Z}_N^*$ of multiplicative order at least D, or a nontrivial factor of N. Our algorithm improves upon an algorithm of Hittmeir [Hit18], who designed a similar algorithm under the stronger assumption $D \geq N^{2/5}$. Hittmeir's algorithm played a crucial role in the recent breakthrough deterministic integer factorization algorithms of Hittmeir and Harvey [Hit21,Har21,HH22b]. When N is assumed to have an r-power divisor with $r \geq 2$, our algorithm provides the same guarantees assuming $D \geq N^{1/6r}$.

Ziv Oznovich, Ben Lee Volk
Reichman University
oznovich.ziv@gmail.com, benlee.volk@runi.ac.il

CP41

Contract Design for Sequential Actions

We introduce a novel model of contracts with combinatorial actions that captures sequential and adaptive agent behavior. As in the standard setting, a principal delegates a costly project to an agent and incentivizes them via a contract specifying payments for each possible outcome. The novelty of our model lies in allowing the agent to select actions sequentially: after each action, they observe the outcome and decide whether to stop or continue. This framework captures common scenarios in which agents can make multiple attempts to achieve a desired outcome. We study the optimal contract problem in this new setting, namely the contract that maximizes the principal's utility. We first observe that the agent's problem — (adaptively) finding the sequence of actions that maximizes his utility for a given contract — is equivalent to the well-known Pandora's Box problem. We then provide algorithms and hardness results for the optimal contract problem, under both independent and correlated actions, and for both linear and general contracts. For independent actions, we give a polynomial-time algorithm for computing the optimal linear contract and show that finding the optimal general contract is NP-hard. When the number of outcomes is constant, we provide a polynomial-time algorithm even for general contracts. In the case of correlated actions, we show that, for both linear and general contracts, approximating the optimal contract within any constant factor is NP-hard.

Tomer Ezra
Harvard University
tomer@cmsa.fas.harvard.edu

Michal Feldman, Maya Schlesinger
Tel Aviv University
mfeldmantau@gmail.com, maya.schlesinger@gmail.com

CP41

When Contracts Get Complex: Information-Theoretic Barriers

In the combinatorial-action contract model a principal delegates the execution of a complex project to an agent, who can choose any subset from a given set of actions. Each set incurs a cost to the agent, given by a function c, and induces an expected reward to the principal, given by f. To incentivize the agent, the principal designs a contract that specifies the payment upon success, with the optimal contract being the one that maximizes the principal's utility. It is known that with access to value queries no constant-approximation is possible for submodular f and additive c. A fundamental open problem is: does the problem become tractable with demand queries? We answer this question to the negative, by establishing that finding an optimal contract for this setting requires exponentially many demand queries. We extend and strengthen this result to different combinations of submodular/supermodular f and c; while allowing the principal to access f and c using arbitrary communication protocols. Our results are driven by novel equal-revenue constructions when one of the functions is additive, immediately implying value query hardness. We then identify a new property — sparse demand — which allows us to strengthen these results to demand query hardness. Finally, by augmenting a perturbed version of these constructions with one additional action, thereby making both functions combinatorial, we establish exponential communication complexity.

Paul Dütting
Google Research
duetting@google.com

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

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

Aviad Rubinstein
Stanford
aviad@cs.stanford.edu

CP41

Robust Equilibria in Shared Resource Allocation Via Strengthening Borders Theorem

We consider the repeated allocation of a shared resource, where a single item is allocated to one of multiple agents each round. We assume agents have i.i.d. values and additive utilities. Past work offers two types of guarantees: (i) approximate Bayes-Nash equilibria (BNE) via linkage-based mechanisms requiring extensive knowledge of value

distributions, or (ii) simple, distribution-agnostic mechanisms with robust utility guarantees for each agent. These robust guarantees hold against arbitrary, even collusive, behavior but are worse than the Nash outcome. Recent work hinted at a barrier to achieving both simultaneously. Our work establishes this is not the case. We propose the first mechanism where a natural strategy for each agent is simultaneously a BNE and provides strong robust guarantees for individual utilities. Our mechanism stems from a surprising connection between online resource allocation and implementation theory, using a novel strengthening of Borders theorem. We show that establishing robust equilibria in this setting reduces to proving a particular subset of the Border polytope is non-empty. We establish this via a novel joint Schur-convexity argument. This strengthening of Border's criterion is of independent technical interest and may prove useful in other settings.

David X. Lin, Siddhartha Banerjee, Giannis Fikioris
Cornell University
dxl2@cornell.edu, sbanerjee@cornell.edu,
gfikioris@cs.cornell.edu

Éva Tardos
Cornell University
Department of Computer Science
eva.tardos@cornell.edu

CP41

Better Regret Rates in Bilateral Trade Via Sublinear Budget Violation

Bilateral trade is a central problem in algorithmic economics, and recent work has explored how to design trading mechanisms that use no-regret learning algorithms. However, no-regret learning is impossible when budget balance has to be enforced at each time step. Bernasconi et al. 2024 show how this impossibility result can be circumvented by relaxing the budget balance constraint to hold only globally over all time steps. In particular, they design an algorithm achieving regret of the order of $\tilde{O}(T^{3/4})$ and provide a lower bound of $\Omega(T^{5/7})$. In this work, we interpolate between these two extremes by studying how the optimal regret rate varies as a function of the allowed violation of the global budget balance constraint. Specifically, we design an algorithm that, by violating the constraint by at most T^{β} for any given $\beta \in [3/4, 6/7]$, attains regret $\tilde{O}(T^{1-\beta/3})$. We complement this result with a matching lower bound, thus fully characterizing the trade-off between regret and budget violation. Our results show that both the $\tilde{O}(T^{3/4})$ upper bound in the global budget balance case and the $\Omega(T^{5/7})$ lower bound under unconstrained budget balance violation obtained by Bernasconi et al.2024 are tight.

Anna Lunghi, Matteo Castiglioni, Alberto Marchesi
Politecnico di Milano
anna.lunghi@polimi.it, matteo.castiglioni@polimi.it, alberto.marchesi@polimi.it

CP41

Fair Division Beyond Monotone Valuations with Applications to Equitable Graph Partitioning

We will focus on fair division among agents whose cardinal preferences are not necessarily monotone. For indivisible items, we will establish the existence of nearly envy-free (EFE3) allocations for nonnegative (non-monotone) valuations. We also consider equitability under nonnegative (non-monotone) valuations and prove that near equitable (EQE3) allocations always exist. In addition, we develop an approximation algorithm that finds allocations that are equitable within additive margins for nonnegative valuations. Our results have combinatorial implications. The developed results imply the following novel results: (i) Given any graph $G = (V, E)$ and integer k (at most $|V|$), there always exists a partition $V_1, V_2, \ldots, V_k$ of the vertex set such that the edge densities within the parts, V_i, are additively within four of each other, and (ii) Given any graph $G = (V, E)$ and integer k (at most $|V|$), there always exists a partition $V_1, V_2, \ldots, V_k \neq \emptyset$ of the vertex set such that the cut function values of the parts, V_i, are additively within $5\Delta + 1$ of each other; here, Δ is the maximum degree of G. Further, such partitions can be computed efficiently. In addition to being interesting in and of itself, this result highlights the reach of the developed guarantees beyond fair division and algorithmic game theory.

Paritosh Verma
Purdue University
paritoshverma97@gmail.com

CP42

Learning in An Echo Chamber: Online Learning with Replay Adversary

As machine learning systems increasingly train on self-annotated data, they risk reinforcing errors and becoming echo chambers of their own beliefs. To model this phenomenon we introduce Online Learning in the Replay Setting. In round t, the learner outputs a hypothesis $\hat{h}_t$; the adversary then reveals either the true label $f^*(x_t)$ or a replayed label $\hat{h}_i(x_t)$ from an earlier round i. A mistake is counted only when the true label is shown, yet classical algorithms such as the SOA or the halving algorithm are easily misled by the replayed errors. We introduce the Extended Threshold dimension, ExThD(H), the exact measure of learnability in this model: a closure-based learner makes at most ExThD(H) mistakes against any adaptive adversary, and no algorithm can perform better. For stochastic adversaries, we prove a similar bound for every intersection-closed class. The replay setting is harder than the mistake bound setting: some classes have constant Littlestone dimension but arbitrarily large ExThD(H). Proper learning exhibits an even sharper separation: a class is properly learnable under replay if and only if it is (almost) intersection-closed. Otherwise, every proper learner suffers $\Omega(T)$ errors, whereas our improper algorithm still achieves the ExThD(H) bound. These results give the first tight analysis of learning against replay adversaries, based on new results for closure-type algorithms.

Daniil Dmitriev
ETH Zurich
dmitrievdaniil97@gmail.com

Harald Franck, Carolin Heinzler
University of Copenhagen
hafr@di.ku.dk, chei@di.ku.dk

Amartya Sanyal
Department of Computer Science, University of Copenhagen

amartya18x@gmail.com

CP42

Online 3-Taxi on General Metrics

The online k-taxi problem, introduced in 1990 by Fiat, Rabani and Ravid, is a generalization of the k-server problem where k taxis must serve a sequence of requests in a metric space. Each request is a pair of two points, representing the pick-up and drop-off location of a passenger. In the interesting "hard" version of the problem, the cost is the total distance that the taxis travel without a passenger. The problem is known to be substantially harder than the k-server problem, and prior to this work even for $k = 3$ taxis it has been unknown whether a finite competitive ratio is achievable on general metric spaces. We present an $O(1)$-competitive algorithm for the 3-taxi problem.

Christian Coester, Tze-Yang Poon
University of Oxford
christian.coester@cs.ox.ac.uk, poontzeyang@gmail.com

CP42

Online Conformal Prediction with Efficiency Guarantees

We study the problem of conformal prediction in a novel online framework that directly optimizes efficiency. In our problem, we are given a target miscoverage rate $\alpha > 0$, and a time horizon T. On each day $t \leq T$ an algorithm must output an interval $I_t \subseteq [0, 1]$, then a point $y_t \in [0, 1]$ is revealed. The goal of the algorithm is to achieve coverage, that is, $y_t \in I_t$ on (close to) a $(1 - \alpha)$-fraction of days, while maintaining efficiency, that is, minimizing the average volume (length) of the intervals played. We study this problem over arbitrary and exchangeable (random order) input sequences. For exchangeable sequences, we show that it is possible to achieve optimal coverage and volume, compared to the best fixed interval in hindsight. For arbitrary sequences, we show that any algorithm that achieves a μ-approximation in volume compared to the best fixed interval in hindsight, must make a multiplicative factor more mistakes than αT, where the factor depends on μ and the aspect ratio of the problem. We give a matching algorithm that recovers all Pareto-optimal settings of μ and number of mistakes. We show that no single algorithm can simultaneously be Pareto-optimal for arbitrary sequences and optimal for exchangeable sequences, ruling out a "best of both worlds" guarantee. On the algorithmic side, we give an algorithm that achieves the near-optimal tradeoff between the two cases.

Vaidehi Srinivas
Northwestern University
vaidehi@u.northwestern.edu

CP42

Nearly Tight Bounds for the Online Sorting Problem

In the online sorting problem, a sequence of n numbers in $[0, 1]$ (including $0, 1$) must be inserted into an array of size $m \geq n$ to minimize the sum of absolute differences between pairs of numbers occupying consecutive non-empty cells. Aamand *et al.* (SODA 2023) gave a deterministic $2^{\sqrt{\log n}\sqrt{\log\log n+\log(1/\varepsilon)}}$-competitive algorithm when $m = (1 + \varepsilon)n$ for $\varepsilon \geq \Omega(\log n/n)$, and showed a lower bound: with $m = \gamma n$, any deterministic algorithm has ratio at least $\frac{1}{\gamma} \cdot \Omega(\log n/ \log\log n)$. This left an exponential gap between upper and lower bounds. We bridge this gap and almost completely resolve the problem. We give a deterministic $O(\log^2 n/\varepsilon)$-competitive algorithm for $m = (1+\varepsilon)n$ and, for $m = \gamma n$ where $\gamma \in [O(1), O(\log^2 n)]$, a deterministic $O(\log^2 n/\gamma)$-competitive algorithm. In particular, this yields an $O(1)$-competitive algorithm using $O(n \log^2 n)$ space, within an $O(\log n \cdot \log\log n)$ factor of the lower bound $\Omega(n \log n/ \log\log n)$. Combined, these results give an almost tight tradeoff: the product of competitiveness and γ is $O(\log^2 n)$, while the lower bound is $\Omega(\log n/ \log\log n)$. The results also extend to unknown numeric ranges, with an additional $O(\log n)$ factor.

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

Debmalya Panigrahi
Duke University
debmalya@cs.duke.edu

Or Vardi
Tel Aviv University
orvardi@mail.tau.ac.il

CP43

Adaptive Set Intersection on the Word Ram

The set intersection problem is fundamental across many domains, yet its analysis has mostly relied on the comparison model. We study the indexed-set intersection problem on the word RAM model, to exploit modern processors beyond comparison-based bounds. Implementing optimal comparison-model algorithms, such as Barbay and Kenyons, atop word-RAM successor data structures yields immediate gains, we ask whether the model allows for further improvement. Let $S = \{S_1, \ldots, S_N\}$ be integer sets, $S_i \subseteq [0, 2^w)$, $|S_i| = n_i$, w being the machine-word size, and queries $Q = \{i_1, \ldots, i_k\} \subseteq [1, N]$ asking for $INT(Q) = \cap_{i\in Q} S_i$. For $1 \leq c \leq \lceil w/ \lg w\rceil$, let $S_i^{(c)}$ be S_i truncated to its $w - c \lg w$ most significant bits, with $|S_i^{(c)}| = n_i^{(c)}$. We combine range reduction, trie certificates, and a novel data structure for constant-time intersections on small-universe sets into an approach we call recursive range reduction. We represent each S_i using $\Theta(c \cdot n_i^{(1)})$ words, and compute $INT(Q)$ in time $O(\delta^{(c)} \sum_{i\in Q^{(c)}} \lg\lg \frac{n_i^{(c)}}{\delta^{(c)}} + k \sum_{j=1}^{c-1} |INT(Q^{(j)})| + |INT(Q)|)$, which in the worst case is $O(\delta \sum_{i\in Q} \lg\lg \frac{n_i}{\delta})$, where δ is Barbay and Kenyon's alternation of Q and $Q^{(j)}$ denotes $\cap_{i\in Q} S_j^{(i)}$. Experiments suggest our approach achieves significantly better practical performance.

Diego Arroyuelo
Pontificia Universidad Católica de Chile
Millennium Institute for Foundational Research on Data
diego.arroyuelo@uc.cl

José M M. Cazorla
Universidad Técnica Federico Santa María
Millennium Institute for Foundational Research on Data
jcazorla@usm.cl

CP43

WarpSpeed: A High-Performance Library for Con-

current Gpu Hash Tables

GPU hash tables are increasingly used to accelerate data processing, but their limited functionality restricts adoption in large-scale data processing applications. Current limitations include incomplete concurrency support and missing compound operations such as upserts. This paper presents WarpSpeed, a library of highperformance concurrent GPU hash tables with a unified benchmarking framework for performance analysis. WarpSpeed implements eight state-of-the-art Nvidia GPU hash table designs and provides a rich API designed for modern GPU applications. Our evaluation uses diverse benchmarks to assess both correctness and scalability, and we demonstrate real-world impact by integrating these hash tables into three downstream applications. We propose several optimization techniques to reduce concurrency overhead, including fingerprint-based metadata to minimize cache line probes and specialized Nvidia GPU instructions for lock-free queries. Our findings provide new insights into concurrent GPU hash table design and offer practical guidance for developing efficient, scalable data structures on modern GPUs.

Hunter Mccoy, Prashant Pandey
Northeastern University
hunterm@berkeley.edu, p.pandey@northeastern.edu

CP43

Practical Insertion-Only Convex Hull

We study insertion-only data structures for convex hulls of a planar point set P of size n, supporting various containment and intersection queries. We investigate a variety of methods tailored to the insertion-only setting. We explore a broad selection of trade-offs involving robustness, memory access patterns, and space usage, providing an extensive evaluation of both existing and novel techniques. We observe that all logarithmic-time methods rely on pointer-based tree structures, which suffer in practice due to poor memory locality. Motivated by this, we develop a vector-based solution inspired by Overmars' logarithmic method. Our structure has worse asymptotic bounds, supporting queries in $O(\log^2 n)$ time, but stores data in $O(\log n)$ contiguous vectors. Through empirical evaluation on real-world and synthetic data sets, we uncover surprising trends. Let h denote the size of the convex hull. We show that a nave $O(h)$ insertion-only algorithm based on Graham scan consistently outperforms both theoretical and practical state-of-the-art methods under realistic workloads, even on data sets with rather large convex hulls. While tree-based methods with $O(\log h)$ update times offer solid theoretical guarantees, they are never optimal in practice. In contrast, our vector-based logarithmic method, despite its theoretically inferior bounds, is highly competitive across all tested scenarios. It is optimal whenever the convex hull becomes large.

Ivor Van Der Hoog
IT University of Copenhagen
ivva@itu.dk

Henrik Reinstädtler
Heidelberg University
henrik.reinstaedtler@informatik.uni-heidelberg.de

Eva Rotenberg
IT University of Copenhagen
erot@itu.dk

CP43

Phast - Perfect Hashing Made Fast

Perfect hash functions give unique "names" to arbitrary keys requiring only a few bits per key. This is an essential building block in applications like static hash tables, databases, or bioinformatics. The lecture presents the PHast approach that combines the fastest available queries, very fast construction, and good space consumption (below 2 bits per key). PHast improves bucket-placement which first hashes each key k to a bucket, and then looks for the bucket seed s such that a placement function maps pairs (s, k) in a collision-free way. PHast can use small-range hash functions with linear mapping, fixed-width encoding of seeds, and parallel construction. This is achieved using small overlapping slices of allowed values and bumping to handle unsuccessful seed assignment. A variant we called PHast$^+$ uses additive placement, which enables bit-parallel seed searching, speeding up the construction by an order of magnitude.

Piotr Beling
University of Lódz
Faculty of Mathematics and Computer Science
piotr.beling@wmii.uni.lodz.pl

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

CP43

Smaller and More Flexible Cuckoo Filters

Cuckoo filters are space-efficient approximate set membership data structures with a controllable false positive rate (FPR) and no false negatives. Existing Cuckoo filters use fingerprints of $(k+3)$ bits per key and an additional space overhead factor of at least 1.05 to achieve an FPR of 2^{-k}. The number of hash table buckets is restricted to a power of two, which may lead to much larger space overheads. We present two improvements of Cuckoo filters. First, we remove the restriction that the number of buckets must be a power of two by using a different placement strategy. Second, we reduce the space overhead factor of Cuckoo filters by using overlapping windows instead of disjoint buckets. Thereby, we maintain a comparable load threshold of the hash table, while reducing the number of alternative slots where any fingerprint may be found and decrease the fingerprint size. A detailed evaluation demonstrates that the alternative memory layout based on overlapping windows decreases the size of Cuckoo filters not only in theory, but also in practice. A comparison with other state-of-the art filter types shows that the reduced space overhead makes windowed Cuckoo filters the smallest filters supporting online insertions, with similarly fast queries, but longer insertion times.

Sven Rahmann, Johanna Schmitz, Jens Zentgraf
Saarland University
rahmann@cs.uni-saarland.de, jschmitz@cs.uni-

saarland.de, zentgraf@cs.uni-saarland.de

CP43

BlockFIFO & MultiFIFO: Scalable Relaxed Queues

FIFO queues are a fundamental data structure used in a wide range of applications. Concurrent FIFO queues allow multiple execution threads to access the queue simultaneously. Maintaining strict FIFO semantics in concurrent queues leads to low throughput due to high contention at the head and tail of the queue. By relaxing the FIFO semantics to allow some reordering of elements, it becomes possible to achieve much higher scalability. This work presents two orthogonal designs for relaxed concurrent FIFO queues, one derived from the MultiQueue and the other based on ring buffers. We evaluate both designs extensively on various micro-benchmarks and a breadth-first search application on large graphs. Both designs outperform state-of-the-art relaxed and strict FIFO queues, achieving higher throughput and better scalability.

Marvin Williams, Stefan Koch
Karlsruhe Institute of Technology
williams@kit.edu, stefan.koch@student.kit.edu

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

CP44

Distributed Reductions for the Maximum Weight Independent Set Problem

Finding maximum-weight independent sets in graphs is an important NP-hard optimization problem. Given a vertex-weighted graph G, the task is to find a subset of pairwise non-adjacent vertices of G with maximum weight. Most recently published practical exact algorithms and heuristics for this problem use a variety of data-reduction rules to compute (near-)optimal solutions. Applying these rules results in an equivalent instance of reduced size. An optimal solution to the reduced instance can be easily used to construct an optimal solution for the original input. In this work, we present the first distributed-memory parallel reduction algorithms for this problem, targeting graphs beyond the scale of previous sequential approaches. Furthermore, we propose the first distributed reduce-and-greedy and reduce-and-peel algorithms for finding a maximum weight independent set heuristically. Our experiments on up to 1024 processors demonstrate good scalability of our distributed reduce algorithms while maintaining good reduction impact. Our asynchronous reduce-and-peel approach achieves an average speedup of 33x over a sequential state-of-the-art reduce-and-peel approach on 36 real-world graphs with a solution quality close to the sequential algorithm. Our reduce-and-greedy algorithms even achieve average speedups of up to 50x at the cost of a lower solution quality. Our approach allows us to consider graphs with more than one billion vertices and 17 billion edges.

Jannick Borowitz
Karlsruhe Institute of Technology
jannick.borowitz@kit.edu

Ernestine Großmann
Heidelberg University
e.grossmann@informatik.uni-heidelberg.de

Matthias Schimek
Karlsruhe Institute of Technology
schimek@kit.edu

CP44

A Customized Sat-Based Solver for Graph Coloring

We introduce ZykovColor, a novel SAT-based algorithm to solve the graph coloring problem working on top of an encoding that mimics the Zykov tree. Our method is based on an approach of Hbrard and Katsirelos (2020) that employs a propagator to enforce transitivity constraints, incorporate lower bounds for search tree pruning, and enable inferred propagations. We leverage the recently introduced IPASIR-UP interface for CaDiCaL to implement these techniques with a SAT solver. Furthermore, we propose new features that take advantage of the underlying SAT solver. These include modifying the integrated decision strategy with vertex domination hints and using incremental bottom-up search that allows to reuse learned clauses from previous calls. Additionally, we integrate a more effective clique computation and an algorithm for computing the fractional chromatic number to improve the lower bounds used for pruning during the search. We validate the effectiveness of each new feature through an experimental analysis. ZykovColor outperforms other state-of-the-art graph coloring implementations on the DIMACS benchmark set. Further experiments on random Erdos-Rnyi graphs show that our new approach matches or outperforms state-of-the-art SAT-based methods for both very sparse and highly dense graphs. We give an additional configuration of ZykovColor that dominates other SAT-based methods on the Erdos-Rnyi graphs.

Timo Brand
Technical University of Munich
Timo.Brand@tum.de

Daniel Faber, Stephan Held, Petra Mutzel
University of Bonn
s6dafabe@uni-bonn.de, held@dm.uni-bonn.de, pmutzel@uni-bonn.de

CP44

On Computing Top-k Simple Shortest Paths from a Single Source

We investigate the problem of computing the top-k simple shortest paths in weighted digraphs. While the single-pair variant has been extensively studied over the past decades, with Yen's algorithm and its heuristic improvements emerging as the most effective solving strategies, relatively little attention has been devoted to the more general single-source version, which asks to determine top-k simple shortest paths from a source vertex to all other vertices. Motivated by the numerous practical applications of ranked shortest paths, in this paper, we provide new insights and algorithmic contributions to this problem. In particular, we first present a characterization of the structural properties of its solutions. Then, we introduce the first polynomial-time algorithm specifically designed to handle it. On the one hand, we prove our new algorithm is on par, in terms of time complexity, with the only polynomial-time approach known to solve the problem, that is, applying the fastest single-pair algorithm independently to each vertex pair formed by the source and the remaining vertices. On the other hand, through an extensive experimental evaluation, we demonstrate that our algorithm outperforms the latter in terms of running time,

achieving speed-ups of up to several orders of magnitude. These results establish our new algorithm as the solution to be preferred for computing k simple shortest paths from a single source in practical settings.

Mattia D'Emidio, Gabriele Di Stefano
University of L'Aquila
mattia.demidio@univaq.it, gabriele.distefano@univaq.it

CP44

Efficient Algorithms for Temporal Balanced Graph Partitioning of Datacenter Workloads

The popularity of distributed machine learning applications and hardware training imposes increasingly stringent performance requirements on the interconnecting communication network. A clever scheduling of the computational workload has the potential to greatly improve datacenter resource utilization, by keeping frequently communicating nodes topologically close. A fundamental underlying optimization problem is known as (static) balanced graph partitioning: How to partition a graph (describing a workload) into equally-sized subgraphs ("clusters") to minimize the number of inter-cluster edges? We study a temporal version where edges between nodes denote requests, and where clusters can be adjusted over time. Over the last years, several fundamental results have been obtained in the online setting. Motivated by the fact that machine learning workloads are often repetitive and fairly predictable, we focus on the dynamic offline problem variant. In particular, we analyze the border of tractability, obtaining hardness and polynomial-time algorithms for special cases. We further develop effective heuristics based on a novel approximation algorithm that is efficient on static graphs of low treewidth. We complement our theoretical insights with an empirical study of our algorithms on real-world datacenter workloads, and shed light on how much knowledge about future demands is required to improve the performance of online algorithms.

Aleksander Figiel, André Nichterlein, Stefan Schmid
TU Berlin
a.figiel@tu-berlin.de, andre.nichterlein@tu-berlin.de, schmiste@gmail.com

CP44

Revisiting a Successful Reduction Rule for Dominating Set

Given a graph $G = (V, E)$ with n vertices and m edges, the Dominating Set problem asks for a set $\mathcal{D} \subseteq V$ of minimal cardinality such that every vertex either is in D or adjacent to a member of D. Although there is little hope for a kernelization algorithm on general graphs due to the W[2]-hardness of Dominating Set, data reduction rules are extensively used in practice. In this context, Rule 1 due to Alber, Fellows, and Niedermeier [JACM 2004] has been shown to be very powerful, yet its best-known running time is $\mathcal{O}(n^3)$ $(= \mathcal{O}(nm))$ for general graphs. In this work, we propose, to the best of our knowledge, the first $\mathcal{O}(n + m)$-time algorithm for Rule 1 on general graphs. We additionally propose simple, but practically significant, extensions to our algorithmic framework to further prune the input instances. We complement our theoretical claims with experiments that confirm the practicality of our approach. On average, we see significant speedups of over one order of magnitude while removing $59.8\times$ more nodes and $410.9\times$ more edges than the original formulation across a large dataset comprised of real-world and synthetic networks.

Lukas Geis
Goethe University, Frankfurt
Institute for Algorithm Engineering
lgeis@ae.cs.uni-frankfurt.de

Alexander Leonhardt
Goethe University, Frankfurt
aleonhardt@ae.cs.uni-frankfurt.de

Johannes Meintrup
University of Applied Sciences Mittelhessen
johannes.meintrup@mni.thm.de

Ulrich Meyer
Goethe University Frankfurt
umeyer@ae.cs.uni-frankfurt.de

Manuel Penschuck
Goethe University Frankfurt, Germany
mpenschuck@ae.cs.uni-frankfurt.de

Lukas Retschmeier
BARC, University of Copenhagen
lure@di.ku.dk

CP44

Engineering Dominating Patterns: A Fine-Grained Case Study

The *Dominating H-Pattern* problem generalizes the classical k-Dominating Set problem: for a fixed *pattern* H and a given graph G, the goal is to find an induced subgraph S of G such that (1) S is isomorphic to H, and (2) S forms a dominating set in G. Fine-grained complexity results show that on worst-case inputs, any significant improvement over the naive brute-force algorithm is unlikely, as this would refute the Strong Exponential Time Hypothesis. Nevertheless, a recent work by Dransfeld et al. (ESA 2025) reveals some significant improvement potential particularly in *sparse* graphs. We ask: Can algorithms with conditionally almost-optimal worst-case performance solve the Dominating H-Pattern, for selected patterns H, efficiently on practical inputs? We develop and experimentally evaluate several approaches on a large benchmark of diverse datasets, including baseline approaches using the Glasgow Subgraph Solver (GSS), the SAT solver Kissat, and the ILP solver Gurobi. Notably, while a straightforward implementation of the algorithms – with conditionally close-to-optimal worst-case guarantee – performs comparably to existing solvers, we propose a tailored Branch-&-Bound approach – supplemented with careful pruning techniques – that achieves improvements of up to two orders of magnitude on our test instances.

Jonathan Dransfeld, Marvin Künnemann, Mirza Redzic
Karlsruhe Institute of Technology
dransheld@outlook.de, marvin.kuennemann@kit.edu, mirza233@gmail.com

Marcus Wunderlich
University of Augsburg
marcus.wunderlich@uni-a.de

CP45

New Heuristic and Multivalued Decision Diagram-Based Exact Algorithms for Repetition-Free

Longest Common Subsequence Problems

The longest common subsequence (LCS) problem is one of the fundamental problems in string algorithms. A constrained version, the repetition-free longest common subsequence (RFLCS) problem, additionally requires that each character may appear at most once in the solution sequence. In sharp contrast to LCS, RFLCS is NP-hard and even APX-hard. Previous work has shown that multivalued decision diagrams (MDDs) provide an effective tool for solving the RFLCS problem to optimality. We introduce new randomized heuristic improvements based on a new dominating match graph encoding of the RFLCS instance. In addition, we advance the use of MDDs in various ways to obtain an improved MDD-based exact algorithm. Our extensive experimental evaluation shows that the new heuristic already solves the majority of instances in an established benchmark set to optimality, and our improved exact algorithm outperforms the state-of-the-art exact solver both with respect to running time and number of solved instances.

Kathrin Hanauer, Georg Braun
University of Vienna
kathrin.hanauer@univie.ac.at, g.braun@gmx.net

Maximilian Vötsch
Faculty of Computer Science
University of Vienna
maximilian.voetsch@univie.ac.at

CP45

Exact Minimum Cuts in Hypergraphs at Scale

The hypergraph minimum cut problem aims to partition the vertices of a hypergraph into two non-empty parts while minimizing the total weight of hyperedges crossing the cut. This problem lies at the core of many tasks in network reliability, VLSI placement, and community detection. We introduce HeiCut, the first algorithm that makes exact minimum cut computation feasible for both weighted and unweighted instances at scales of hundreds of millions of vertices. HeiCut presents seven exact reduction rules that provably preserve the minimum cut, and an optional heuristic contraction based on label propagation that shrinks complex and persistent structures. When no further reductions are possible, the remaining instance is solved exactly with a known algorithm. Our extensive evaluation on more than 500 real-world hypergraphs reveals that the exact reductions alone already expose the minimum cut (i.e., the residual collapses to a single vertex or has no hyperedges) in over 85% of instances. Across all instances, HeiCut solves over twice as many instances as the state-of-the-art within set computational limits, and is up to five orders of magnitude faster. Thus, HeiCut significantly advances hypergraph minimum cut computation in real-world, large-scale scenarios.

Adil Chhabra
Heidelberg University
adil.chhabra@informatik.uni-heidelberg.de

Christian Schulz
Heidelberg University, Germany
christian.schulz@informatik.uni-heidelberg.de

Bora Ucar
LIP, ENS-Lyon, France.
Bora.Ucar@ens-lyon.fr

Loris Wilwert
Heidelberg University
loris.wilwert@stud.uni-heidelberg.de

CP45

Efficient Heuristics and Exact Methods for Pairwise Interaction Sampling

We consider pairwise interaction sampling, an optimization problem that is fundamental to testing in modern configurable software systems, e.g., in automotive industries. We are given a configuration space, the dimensions of which are the Boolean features of a software system; valid configurations are the satisfying assignments of a given propositional formula. The objective is to find a minimum family of configurations that jointly tests each pair of features. Due to its relevance in Software Engineering, this problem has been studied extensively for over 20 years. Aside from new theoretical insights (we prove BH-hardness), we vastly improve the state of the art on the practical side. We devise and engineer an algorithm that can find solutions and identify the set of valid interactions in reasonable time, even for huge instances that previous approaches could not solve in such time. We present an enhanced approach for computing lower bounds. We present an exact algorithm to find optimal solutions based on an interaction selection heuristic driving an incremental SAT solver. For larger instances, we present an LNS-based solver solving most instances in a diverse benchmark library to provable optimality within an hour on commodity hardware. Remarkably, we can solve the largest instances in published benchmark sets (with about 500000000 feasible interactions) to provable optimality; previous approaches could not even find feasible solutions.

Phillip Keldenich
TU Braunschweig
keldenich@ibr.cs.tu-bs.de

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

Dominik M. Krupke, Michael Perk
TU Braunschweig
krupke@ibr.cs.tu-bs.de, perk@ibr.cs.tu-bs.de

CP45

Engineering Fast and Space-Efficient Recompression from SLP-Compressed Text

Compressed indexing enables powerful queries over massive and repetitive textual datasets using space proportional to the compressed input. While theoretical advances have led to highly efficient index structures, their practical construction remains a bottleneck, especially for complex components like recompression RLSLP - a grammar-based representation crucial for building powerful text indexes that support widely used suffix and LCP array queries. In this work, we present the first implementation of recompression RLSLP construction that runs in compressed time, operating on an LZ77-like approximation of the input. Compared to state-of-the-art uncompressed-time methods, our approach achieves up to 46× speedup and 17× lower RAM usage on large, repetitive inputs. These gains unlock scalability to larger datasets and affirm compressed computation as a practical path forward for fast index construction.

Ankith Reddy Adudodla, Dominik Kempa
Stony Brook University
aadudodla@cs.stonybrook.edu,
kempa@cs.stonybrook.edu

CP45

Distributed Computation of Persistent Cohomology

Persistent (co)homology is a central construction in topological data analysis, where it is used to quantify prominence of features in data to produce stable descriptors suitable for downstream analysis. Persistence is challenging to compute in parallel because it relies on global connectivity of the data. We propose a new algorithm to compute persistent cohomology in the distributed setting. It combines domain and range partitioning. The former is used to reduce and sparsify the coboundary matrix locally. After this initial local reduction, we redistribute the matrix across processors for the global reduction. We experimentally compare our cohomology algorithm with DIPHA, the only publicly available code for distributed computation of persistent (co)homology; our algorithm demonstrates a significant improvement in strong scaling.

Arnur Nigmetov, Dmitriy Morozov
Lawrence Berkeley National Laboratory
anigmetov@lbl.gov, dmorozov@lbl.gov

CP45

Collapse and Persistence of Directed Filtered Graphs

Pritam et al. initiated a line of work that introduces the concept of *dominated edges and vertices* to reduce the size of (bi)filtered *undirected* graphs, without compromising their topological features. In this article, we extend these ideas to the setting of *directed filtered graphs*. We provide a characterization of vertex and edge domination in terms of the local neighborhood structure of the directed graph. Building on this, we adapt the filtration reduction algorithm of [**?**] to the directed case, and demonstrate through experiments that this often yields considerable improvements in both runtime and memory consumption in the computational pipeline for homology and persistent homology. As with its undirected counterpart, our algorithm operates directly on the filtered directed graph.

Siddharth Pritam
Chennai Mathematical Institute, India
spritam@cmi.ac.in

Rohit Roy
Chennai Mathematical Institute
India
rohitroy@cmi.ac.in

CP46

A Simple Schnyder Drawing Algorithm for Cylindric and Toroidal Triangulations, with Grid Size O(n)xO(n)

We consider the problem of computing a straight-line crossing-free and periodic grid drawing of cylindric and toroidal triangulations. More precisely, given a cylindric simple triangulation G with n vertices, we design an algorithm based on the Schnyder face-counting principle that computes an x-periodic drawing of G on an integer grid of size $w \times h$, with $w \leq 2n, h \leq 6n$. As a byproduct, this yields an algorithm that computes an xy-periodic drawing of a simple toroidal triangulation with n vertices on a grid of size $w \times h$, with $w \leq 4n, h \leq 10n$. A vertex-counting variant improves the grid bounds to $w \leq n, h \leq 3n$ in the cylindric case and $w \leq 2n, h \leq 5n$ in the toroidal case. Our algorithm is simple to describe and implement, and runs in linear time.

Luca Castelli Aleardi
LIX, Ecole Polytechhnique, Institut Polytechnique de Paris
amturing@lix.polytechnique.fr

Giselle Feng
Department of Biomedical Informatics, Columbia University
LIX, Ecole Polytechhnique, Institut Polytechnique de Paris
fg2618@cumc.columbia.edu

Eric Fusy
LIGM, Université Gustave Eiffel, France
eric.fusy@univ-eiffel.fr

CP46

Computing the Frchet Distance When Just One Curve is c-Packed: A Simple Almost-Tight Algorithm

We study approximating the continuous Frchet distance of two curves with complexity n and m, under the assumption that only one of the two curves is c-packed. Driemel, Har-Peled and Wenk DCG'12 studied Frchet distance approximations under the assumption that both curves are c-packed. In $\mathbb{R}^d$, they prove a $(1+\varepsilon)$-approximation in $\tilde{O}(c\,\frac{n+m}{\varepsilon})$ time. Bringmann and Knnemann IJCGA'17 improved this to $\tilde{O}(c\,\frac{n+m}{\sqrt{\varepsilon}})$ time, which they showed is near-tight under SETH. Both algorithms have a hidden exponential dependency on the dimension d. Recently, Gudmundsson, Mai, and Wong ISAAC'24 studied our setting where only one of the curves is c-packed. They provide an involved $\tilde{O}((c+\varepsilon^{-1})(cn\varepsilon^{-2}+c^2m\varepsilon^{-7}+\varepsilon^{-2d-1}))$-time algorithm when the c-packed curve has n vertices and the arbitrary curve has m. In this paper, we show a simple technique to compute a $(1+\varepsilon)$-approximation in $\mathbb{R}^d$ in time $O(c\,\frac{n+m}{\varepsilon}\log\frac{n+m}{\varepsilon})$ when one of the curves is c-packed. Our approach is not only simpler than previous work, but also significantly improves the dependencies on c, ε, and d (which is only linear). Our algorithm is robust in the sense that it does not require knowledge of c, nor information about which of the two input curves is c-packed.

Jacobus Conradi
Bonn University
conradi@cs.uni-bonn.de

Ivor van Der Hoog
IT-University of Copenhagen
vanderhoog@gmail.com

Thijs Van Der Horst, Tim Ophelders
TU Eindhoven
Utrecht University

t.w.j.vanderhorst@uu.nl, t.a.e.ophelders@uu.nl

CP46

A Levelset Algorithm for 3D-Tarksi

We present a simple new algorithm for finding a Tarski fixed point of a monotone function $F : [N]^3 \to [N]^3$. Our algorithm runs in $O(\log^2 N)$ time and makes $O(\log^2 N)$ queries to F, matching the $\Omega(\log^2 N)$ query lower bound due to Etessami et al. as well as the existing state-of-the-art algorithm due to Fearnley et al.

Sebastian Haslebacher, Jonas Lill
Department of Computer Science
ETH Zurich
sebastian.haslebacher@inf.ethz.ch, jonas.lill@inf.ethz.ch

CP46

The Road to the Closest Point Is Paved by Good Neighbors

Given a set P of n points in $\mathbb{R}^d$, and a parameter $\varepsilon \in (0, 1)$, we present a new construction of a directed graph G, of size $O(n/\varepsilon^d)$, such that $(1 + \varepsilon)$-ANN queries can be answered by performing a greedy walk on G, repeatedly moving to a neighbor that is (significantly) better than the current point. To the best of our knowledge, this is the first construction of a linear size with no dependency on the spread of the point set. The resulting query time, is $O(\varepsilon^{-d} \log \Psi)$, where Ψ is the spread of P. The new construction is surprisingly simple and should be practical.

Benjamin A. Raichel
University of Texas at Dallas
benjamin.raichel@utdallas.edu

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

Eliot Robson
University of Illinois, Urbana Champaign
erobson2@illinois.edu

CP46

Redistricting in Triangular and Square Grids

We study the configuration space of connected graph partitions over sections of the triangular and square grids. In application domains, the partition of a geographical region into smaller regions is modeled by a graph $M = (V, E)$, called a map. A k-district map is a map M together with a partition $V = V_1 \cup \ldots \cup V_k$ into nonempty subsets, called districts, such that each district V_i induces a connected subgraph $M[V_i]$. We are interested in the configuration space, where the nodes are the k-district maps on M, and two nodes are adjacent if they are connected by a so-called recombination move, where two districts V_i, V_j are replaced by two new districts V_i', V_j' such that $V_i \cup V_j = V_i' \cup V_j'$. Navigating the configuration space, even for small and highly structured graphs, is computationally infeasible. Practitioners use Markov chains to analyze the quality of district maps. The theoretical underpinning of this approach relies on whether the configuration space is connected. Recently, Cannon (2024) proved that the configuration space of triangular maps is connected for $k = 3$ under various additional assumptions. Our main contribution is a significantly simpler proof for the same result. A simpler proof allows us to prove stronger results: (1) we can relax several assumptions made by Cannon in the triangular grid; and (2) we can extend the results to the square grid. Our results settle several open problems posed by Cannon (2024).

Frederick Stock, Hugo A. Akitaya
University of Massachusetts Lowell
frederick_stock@student.uml.edu, hugo_akitaya@uml.edu

Kyle Dituro, Andrei Gonczi
Tufts University
kditur01@cs.tufts.edu, andrei.gonczi@tufts.edu

Matias Korman
Siemens Electronic Design Automation
matias.korman@siemens.com

Diane Souvaine
Tufts University
diane.souvaine@tufts.edu

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

CP46

A Quasi-Polynomial Time Algorithm for 3-Coloring Circle Graphs

A graph G is a circle graph if it is an intersection graph of chords of a unit circle. We give an algorithm that takes as input an n vertex circle graph G, runs in time at most $n^{O(\log n)}$ and finds a proper 3-coloring of G, if one exists. As a consequence we obtain an algorithm with the same running time to determine whether a given ordered graph $(G, \prec)$ has a 3-page book embedding. This gives a partial resolution to the well known open problem of Dujmovic and Wood [Discret. Math. Theor. Comput. Sci. 2004], Eppstein [2014], and Bachmann, Rutter and Stumpf [J. Graph Algorithms Appl. 2024] of whether 3-Coloring on circle graphs admits a polynomial time algorithm.

Ajaykrishnan E S
University of California, Santa Barbara
aj66@ucsb.edu

Robert Ganian
TU Wien
rganian@gmail.com

Daniel Lokshtanov
UCSB
daniello@ucsb.edu

Vaishali Surianarayanan
UC Santa Barbara
vaishalisurianarayanan@gmail.com

CP47

Efficient Non-Adaptive Quantum Algorithms for Tolerant Junta Testing

We consider the problem of deciding whether an n-qubit unitary (or n-bit Boolean function) is ε_1-close to some k-junta or ε_2-far from every k-junta, where k-junta unitaries act non-trivially on at most k qubits and as the identity on the rest. For constant numbers $\varepsilon_1, \varepsilon_2$ such that $0 < \varepsilon_1 <$

$\varepsilon_2 < 1$, we show the following.

1. A non-adaptive $O(k \log k)$-query tolerant $(\varepsilon_1, \varepsilon_2)$-tester for k-junta unitaries when $2\sqrt{2}\varepsilon_1 < \varepsilon_2$.
2. A non-adaptive tolerant $(\varepsilon_1, \varepsilon_2)$-tester for k-junta Boolean functions with $O(k \log k)$ quantum queries when $4\varepsilon_1 < \varepsilon_2$.
3. A $2^{\widetilde{O}(k)}$-query tolerant $(\varepsilon_1, \varepsilon_2)$-tester for k-junta unitaries for any $\varepsilon_1, \varepsilon_2$.

The first algorithm provides an exponential improvement over the best-known quantum algorithms [CLL24, ADG25]. The second algorithm shows an exponential quantum advantage over any non-adaptive classical algorithm [CDL+24]. The third tester gives the first tolerant junta unitary testing result for an arbitrary gap. Besides, we adapt the first two quantum algorithms to be implemented using only single-qubit operations, thereby enhancing experimental feasibility, with a slightly more stringent requirement for the parameter gap.

Zongbo Bao
CWI
zongbo.bao@cwi.nl

Kunal Marwaha
University of Chicago
kmarw@uchicago.edu

Christopher Musco
New York University, U.S.
cmusco@nyu.edu

Xinchen Yang
University of Maryland
xcyang@cs.umd.edu

CP47

Most Juntas Saturate the Hardcore Lemma

Consider a function that is mildly hard for size-s circuits. For sufficiently large s, Impagliazzo's hardcore lemma provides a subset of inputs on which the same function is extremely hard for circuits of size $s' \ll s$. Blanc, Hayderi, Koch, and Tan [FOCS, 2024] proved that the size degradation from s to s' present in this lemma is quantitatively tight. We give a much simpler proof of this fact by showing that the hardcore lemma is tight for a random junta with high probability.

Vinayak M. Kumar
University of Texas at Austin
vmkumar@utexas.edu

CP47

An Exact Algorithm for the Unanimous Vote Problem

Consider n independent, biased coins, each with a known probability of heads. Presented with an ordering of these coins, flip (i.e., toss) each coin once, in that order, until we have observed both a head and a tail, or flipped all coins. The Unanimous Vote problem asks us to find the ordering that minimizes the expected number of flips. Gkenosis et al. gave a polynomial-time ϕ-approximation algorithm for this problem, where $\phi \approx 1.618$ is the golden ratio. They left open whether the problem was NP-hard. We answer this question by giving an exact algorithm that runs in time $O(n \log n)$. The Unanimous Vote problem is an instance of the more general Stochastic Boolean Function Evaluation problem: it thus becomes one of the only such problems known to be solvable in polynomial time. Our proof uses simple interchange arguments to show that the optimal ordering must be close to the ordering produced by a natural greedy algorithm. Beyond our main result, we compare the optimal ordering with the best adaptive strategy, proving a tight adaptivity gap of $1.2 \pm o(1)$ for the Unanimous Vote problem.

Feyza Duman Keles, Lisa Hellerstein
NYU
fd2135@nyu.edu, lisa.hellerstein@nyu.edu

CP47

The k-Fold Matroid Secretary Problem

In the matroid secretary problem, elements $N := [n]$ of a matroid $\mathcal{M} \subseteq 2^N$ arrive in random order. When an element arrives, its weight is revealed and a choice must be made to accept or reject the element, subject to the constraint that the accepted set $S \in \mathcal{M}$. [Kleinberg05] gives a $(1 - O(1/\sqrt{k}))$-competitive algorithm when $\mathcal{M}$ is a k-uniform matroid. We generalize their result, giving a $(1 - O(\sqrt{\log(n)/k}))$-competitive algorithm when $\mathcal{M}$ is a k-fold matroid union.

Rishi Gujjar, Kevin Hua, Robert Kleinberg
Cornell University
rg767@cornell.edu, kevinhua@cs.cornell.edu, rdk@cs.cornell.edu

Frederick V. Qiu
Princeton University
frdrckq@gmail.com

CP47

Trading Prophets with Initial Capital

Correa et al. [EC' 2023] introduced the following trading prophets problem. A trader observes a sequence of stochastic prices for a stock, each drawn from a known distribution, and at each time must decide whether to buy or sell. Unfortunately, they observed that in this setting it is impossible to compete with a prophet who knows all future stock prices. In this paper, we explore the trading prophets problem when we are given initial capital with which to start trading. We show that initial capital is enough to bypass the impossibility result and obtain a competitive ratio of 3 with respect to a prophet who knows all future prices (and who also starts with capital), and we show that this competitive ratio is best possible. We further study a more realistic model in which the trader must pay multiplicative and/or additive transaction costs for trading which model dynamics such as bid-ask spreads and broker fees.

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

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

Roie Levin, Or Vardi
Tel Aviv University

roie.levin@rutgers.edu, orvardi@mail.tau.ac.il

CP47

Static Pricing for Single Sample Multi-Unit Prophet Inequalities

In this paper, we study k-unit single sample prophet inequalities. A seller has k identical, indivisible items to sell. A sequence of buyers arrive one-by-one, with each buyer's private value for the item, X_i, revealed to the seller when they arrive. While the seller is unaware of the distribution from which X_i is drawn, they have access to a single sample, Y_i drawn from the same distribution as X_i. What strategies can the seller adopt for selling items so as to maximize social welfare? Previous work has shown that when $k = 1$, if the seller sets a price equal to the maximum of the samples, they achieve a competitive ratio of $\frac{1}{2}$ of the social welfare, and recently Pashkovich and Sayutina established an analogous result for $k = 2$. In this paper, we prove that for $k \geq 3$, setting a static price equal to the k^{th} largest sample also obtains a competitive ratio of $\frac{1}{2}$, resolving a conjecture Pashkovich and Sayutina pose. We also consider the case where k is large. We show that setting the $(k - \sqrt{2k \log k})^{\text{th}}$ largest sample as the price obtains a competitive ratio of $1-\sqrt{\frac{2\log k}{k}}-o\left(\sqrt{\frac{\log k}{k}}\right)$, and that this is optimal for static pricing with access to a single sample. This should be compared against $1-\sqrt{\frac{\log k}{k}}-o\left(\sqrt{\frac{\log k}{k}}\right)$, the ratio obtainable using static pricing with knowledge of the distributions of the values.

Pranav Nuti
University of Chicago
pranavn@uchicago.edu

Peter Westbrook
Citadel Securities
petewest@stanford.edu

CP48

Semi-Robust Communication Complexity of Maximum Matching

We study the one-way two-party communication complexity of Maximum Matching in the semi-robust setting where the edges of a maximum matching are randomly partitioned between Alice and Bob, but all remaining edges of the input graph are adversarially partitioned between the two parties. We show that the simple protocol where Alice solely communicates a lexicographically-first maximum matching of their edges to Bob is surprisingly powerful: We prove that it yields a 3/4-approximation in expectation and that our analysis is tight. The semi-robust setting is at least as hard as the fully robust setting. In this setting, all edges of the input graph are randomly partitioned between Alice and Bob, and the state-of-the-art result is a fairly involved 5/6-approximation protocol that is based on the computation of edge-degree constrained subgraphs [Azarmehr, Behnezhad, ICALP'23]. Our protocol also immediately yields a 3/4-approximation in the fully robust setting. One may wonder whether an improved analysis of our protocol in the fully robust setting is possible: While we cannot rule this out, we give an instance where our protocol only achieves a $0.832 < 5/6 = 0.8\overline{3}$ approximation. Hence, while our simple protocol performs surprisingly well, it cannot be used to improve over the state-of-the-art in the fully robust setting.

Christian Konrad, Gabriel Cipriani Huete,
Adithya Diddapur
University of Bristol
christian.konrad@bristol.ac.uk,
gabriel.cipriani.2021@bristol.ac.uk,
adi.diddapur@bristol.ac.uk

Pavel Dvorak
Charles University, Prague, Czech Republic
koblich@iuuk.mff.cuni.cz

CP48

Efficient Derandomization of Differentially Private Counting Queries

Differential privacy for the 2020 census required an estimated 90 terabytes of randomness [GL20], an amount which may be prohibitively expensive or entirely infeasible to generate. Motivated by these practical concerns, [CSV25] initiated the study of the randomness complexity of differential privacy, and in particular, the randomness complexity of d counting queries. This is the task of outputting the number of entries in a dataset that satisfy predicates $\mathcal{P}_1, \ldots, \mathcal{P}_d$ respectively. They showed the rather surprising fact that though any reasonably accurate, ε-differentially private mechanism for one counting query requires $1 - O(\varepsilon)$ bits of randomness in expectation, there exists a fairly accurate mechanism for d counting queries which requires only $O(\log d)$ bits of randomness in expectation. The mechanism of [CSV25] is inefficient (not polynomial time) and relies on a combinatorial object known as rounding schemes. Here, we give a polynomial time mechanism which achieves nearly the same randomness complexity versus accuracy tradeoff as that of [CSV25]. Our construction is based on the following simple observation: after a randomized shift of the answer to each counting query, the answer to many counting queries remains the same regardless of whether we add noise to that coordinate or not. This allows us to forgo the step of adding noise to the result of many counting queries.

Surendra Ghentiyala
Cornell University
sg974@cornell.edu

CP48

Simple Average-Case Analysis of Recursive Randomized Greedy Mis

We revisit the complexity analysis of the recursive version of the randomized greedy algorithm for computing a maximal independent set (MIS), originally analyzed by Yoshida, Yamamoto, and Ito (2009). They showed that, on average per vertex, the expected number of recursive calls made by this algorithm is upper-bounded by the average degree of the input graph. While their analysis is clever and intricate, we provide a significantly simpler alternative that achieves the same guarantee. Our analysis is inspired by the recent work of Dalirrooyfard, Makarychev, and Mitrovic (2024), who developed a potential-function-based argument to analyze a new algorithm for correlation clustering. We adapt this approach to the MIS setting, yielding a more direct and arguably more transparent analysis of the recursive randomized greedy MIS algorithm.

Mina Dalirrooyfard
Morgan Stanley

minad@mit.edu

Konstantin Makarychev
Northwestern University
konstantin@northwestern.edu

Slobodan Mitrovic
UC Davis
smitrovic@ucdavis.edu

CP48

A (Very) Nearly Optimal Sketch for k-Edge Connectivity Certificates

In this note, we present a simple algorithm for computing a *k-connectivity certificate* in dynamic graph streams. Our algorithm uses $O(n \log^2 n \cdot \max\{k, \log n \log k\})$ bits of space which improves upon the $O(kn \log^3 n)$-space algorithm of Ahn, Guha, and McGregor (SODA'12). For the values of k that are truly sublinear, our space usage *very nearly* matches the known lower bound $\Omega(n \log^2 n \cdot \max\{k, \log n\})$ established by Nelson and Yu (SODA'19; implicit) and Robinson (DISC'24). In particular, our algorithm fully settles the space complexity at $\Theta(kn \log^2 n)$ for $k = \Omega(\log n \log \log n)$, and bridges the gap down to only a doubly-logarithmic factor of $O(\log \log n)$ for a smaller range of $k = o(\log n \log \log n)$.

Pachara Sawettamalya, Huacheng Yu
Princeton University
ps3122@princeton.edu, yuhch123@gmail.com

CP48

Continual Release of Densest Subgraphs: Privacy Amplification and Sublinear Space Via Subsampling

We study the sublinear space continual release model for edge-differentially private (DP) graph algorithms, with a focus on the densest subgraph problem (DSG) in the insertion-only setting. Our main result is the first continual release DSG algorithm that matches the additive error of the best static DP algorithms and the space complexity of the best non-private streaming algorithms, up to constants. The key idea is a refined use of subsampling that simultaneously achieves privacy amplification and sparsification, a connection not previously formalized in graph DP. Via a simple black-box reduction to the static setting, we obtain both pure and approximate-DP algorithms with $O(\log n)$ additive error and $O(n \log n)$ space, improving both accuracy and space complexity over the previous state of the art. Along the way, we introduce graph densification in the graph DP setting, adding edges to trigger earlier subsampling, which removes the extra logarithmic factors in error and space incurred by prior work [Epasto, Liu, Mukherjee, Zhou, 2025]. We believe this simple idea may be of independent interest.

Felix Zhou
Yale University
felix.zhou@yale.edu

CP49

A Novel Reduction from #sat to #2sat Based on Symmetry: Simply Drop the Large Clauses

#2SAT is complete for #P under Turing (many-call) reductions, which dates back to the seminal work by Valiant from 1979. Arguably, this reduction is a sophisticated chain of transformations. In contrast, we give a simple reduction which makes only two calls instead of polynomially many calls. Our reduction is based on parity constraints, which we can encode using almost solely binary clauses and very few clauses of size four. However, thanks to the subtraction and the symmetry between the two formulas, these clauses of size four can simply be omitted without changing the overall result of the computation thus leading to a surprisingly simple reduction from #SAT to two calls of #2SAT. Apart from the fact that the reduction is simple, it also improves the construction by Valiant in various ways. First, it is computable either in logspace or linear time and, hence, also produces formulas that are linearly bounded by the input formula (whereas the construction of Valiant produces formulas of polynomial size). Furthermore, the construction can easily be adapted to preserve structural parameters like the input's treewidth, yielding a simple proof of the fact that #SAT can be solved in time $2^k \cdot poly(|\varphi|)$ on formulas of incidence treewidth k (without Mobius and Zeta transformations). Finally, our technique can be used to improve and simplify in the sparsification lemma from #dSAT to #2SAT.

Erik Demaine
MIT
edemaine@mit.edu

Max Bannach
European Space Agency
max.bannach@esa.int

Timothy Gomez
MIT
tagomez7@mit.edu

Markus Hecher
CNRS, UMR8188, CRIL, France
hecher@mit.edu

CP49

Simple Length-Constrained Expander Decompositions

Length-constrained expander decompositions are a new graph decomposition that has led to several recent breakthroughs in fast graph algorithms. Roughly, an (h, s)-length ϕ-expander decomposition is a small collection of length increases to a graph so that nodes within distance h can route flow over paths of length hs while using each edge to an extent at most $1/\phi$. Prior work showed that every n-node and m-edge graph admits an (h, s)-length ϕ-expander decomposition of size $\log n \cdot sn^{O(1/s)} \cdot \phi m$. In this work, we give a simple proof of the existence of (h, s)-length ϕ-expander decompositions with an improved size of $sn^{O(1/s)} \cdot \phi m$. Our proof is a straightforward application of the fact that the union of sparse length-constrained cuts is itself a sparse length-constrained cut. In deriving our result, we improve the loss in sparsity when taking the union of sparse length-constrained cuts from $\log^3 n \cdot s^3 n^{O(1/s)}$ to $s \cdot n^{O(1/s)}$.

Bernhard Haeupler
INSAIT, Sofia University "Kliment Ohridski'
ETH Zurich
bernhard.haeupler@inf.ethz.ch

Greg Bodwin
University of Michigan

bodwin@umich.edu

Ellis Hershkowitz
Brown University
delhersh@gmail.com

Zihan Tan
Rutgers University
zihantan1993@gmail.com

CP49

Tight Lower Bound for Multicolor Discrepancy

We prove the following asymptotically tight lower bound for k-color discrepancy: For any $k \geq 2$, there exists a hypergraph with n hyperedges such that its k-color discrepancy is at least $\Omega(\sqrt{n})$. This improves on the previously known lower bound of $\Omega(\sqrt{n/\log k})$ due to Caragiannis et al. (2025). As an application, we show that our result implies improved lower bounds for group fair division.

Pasin Manurangsi
Google Research
pasin@google.com

Raghu Meka
University of California, Los Angeles
raghum@cs.ucla.edu

CP49

The Sparsification Lemma Via Measure and Conquer

The Sparsification Lemma of Impagliazzo, Paturi, and Zane [FOCS 1998; JCSS 2001] is a foundational result which enables the extensive use of the Exponential Time Hypothesis (ETH) to give asymptotically tight lower bounds for exponential-time algorithms and parameterized algorithms. The lemma states that, for every positive integer d and every real $\varepsilon > 0$, every boolean d-CNF formula can be expressed as a disjunction of $2^{\varepsilon n}$ d-CNF formulas, each of which is $C_{(\varepsilon,d)}$-*sparse*, which means that every variable appears in at most $C_{(\varepsilon,d)}$ clauses. While the original proof establishes the lemma with $C_{(\varepsilon,d)}$ upper bounded by a double exponential function in d, subsequent refinements by Calabro et al. [CCC 2006] and Zeijlemaker [M.Sc. Thesis, Eindhoven University of Technology, 2020] improve the sparsity $C_{(\varepsilon,d)}$ to $(d/\varepsilon)^{\mathcal{O}(d)}$. The existing proofs of the Sparsification Lemma rely on intricate analysis of recursion trees. In this work, we present a short and self-contained proof of the Sparsification Lemma that matches the best known bounds. Our analysis avoids all non-trivial recursion-tree arguments used in prior work and instead applies an inductive argument based on the *Measure and Conquer* framework. This results in a modular and (in our opinion) conceptually simpler proof.

Daniel Lokshtanov
UCSB
daniello@ucsb.edu

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

Jie Xue
New York University Shanghai
jiexue@nyu.edu

CP49

A Simpler Exponential-Time Approximation Algorithm for MAX-k-SAT

We present an extremely simple polynomial-space exponential-time $(1 - \varepsilon)$-approximation algorithm for MAX-k-SAT that is (slightly) faster than the previous known polynomial-space $(1-\varepsilon)$-approximation algorithms by Hirsch (Discrete Applied Mathematics, 2003) and Escoffier, Paschos and Tourniaire (Theoretical Computer Science, 2014). Our algorithm repeatedly samples an assignment uniformly at random until finding an assignment that satisfies a large enough fraction of clauses. Surprisingly, we can show the efficiency of this simpler approach by proving that in any instance of MAX-k-SAT (or more generally any instance of MAXCSP), an exponential number of assignments satisfy a fraction of clauses close to the optimal value.

Harry Buhrman
Quantinuum and QuSoft
hbuhrman@gmail.com

Sevag Gharibian
Paderborn University
sevag.gharibian@upb.de

Zeph Landau
University of California, Berkeley
zeph.landau@gmail.com

Francois Le Gall
Nagoya University, Japan
legall@math.nagoya-u.ac.jp

Norbert Schuch
University of Vienna
norbert.schuch@univie.ac.at

Suguru Tamaki
University of Hyogo
suguru.tamaki@gmail.com

CP49

A Linear-Time Algorithm for the MCS End-Vertex Problem on Chordal Graphs: A Bonus-Driven Search Strategy

The end-vertex problem for Maximum Cardinality Search (MCS) on chordal graphs asks whether a given vertex can be the last one visited by an MCS on a chordal graph. We present a simple linear-time algorithm that solves this problem without relying on complex data structures. To this end, we introduce a novel variant of MCS, termed MCS-Ultra. The core idea of this variant lies in a carefully designed bonus allocation scheme, which assigns an integer value to each vertex in the graph. This scheme is incorporated into the MCS procedure to effectively guide the vertex visiting order on chordal graphs while preserving linear-time complexity.

Guozhen Rong, Biao Yuan
Changsha University of Science and Technology
rongguozhen@csust.edu.cn, yuanbiao@stu.csust.edu.cn

Yongjie Yang

Saarland University
yyongjiecs@gmail.com

Zhen Zhang
School of Advanced Interdisciplinary Studies
zz@hutb.edu.cn

CP50

Debiasing Polynomial and Fourier Regression

We study the problem of approximating an unknown function $f : \mathbb{R} \to \mathbb{R}$ by a degree-d polynomial using as few function evaluations as possible, where error is measured with respect to a probability distribution μ. Existing randomized algorithms achieve near-optimal sample complexities to recover a $(1 + e)$-optimal polynomial but produce biased estimates of the best polynomial approximation, which is undesirable. We propose a simple debiasing method based on a connection between polynomial regression and random matrix theory. Our method involves evaluating $f(?_1), , f(?_{d+1})$ where $?_1, , ?_{d+1}$ are the eigenvalues of a suitably designed random complex matrix tailored to the distribution μ. Our estimator is unbiased, has near-optimal sample complexity, and experimentally outperforms iid leverage score sampling. Additionally, our techniques enable us to debias existing methods for approximating a periodic function with a truncated Fourier series with near-optimal sample complexity.

Chris Camano
Caltech
ccamano@caltech.edu

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

Kevin Shu
California Institute of Technology
kshu@caltech.edu

CP50

Faster Convolutions: Yates and Strassen Revisited

Given vectors $u, v \in \mathbb{Q}^D$ over a finite domain D and a function $f : D \times D \to D$, the convolution problem asks to compute $w \in \mathbb{Q}^D$ with

$$w(d) = \sum_{\substack{x,y \in D \\ f(x,y)=d}} u(x)v(y).$$

In parameterized and exponential-time algorithms, product-domain convolutions are key: for fixed B and $h : B \times B \to B$, we set $D = B^k$ and $f = h^k$ applying h coordinate-wise. We give a multilinear-algebraic view of product-domain convolutions that unifies and simplifies prior work such as van Rooij et al. [ESA 2009]. Moreover, using established results from the theory of fast matrix multiplication, we derive improved $O^*(|B|^{2\omega/3 \cdot k}) = O(|D|^{1.582})$-time algorithms, surpassing bounds of the form $c^k|B|^{2k}$ for $c < 1$ by Esmer et al. [Algorithmica 86(1), 2024]. Under Strassens asymptotic rank conjecture, our framework yields quasi-linear $|D|^{1+o(1)}$-time algorithms. Using the setup described in this paper, Strassens asymptotic rank conjecture from algebraic complexity theory would imply quasi-linear $|D|^{1+o(1)}$-time algorithms. This conjecture has recently gained attention in the algorithms community, e.g., BjrklundKaski and Pratt [STOC 2024] and Bjrklund et al. [SODA 2025]. This paper provides a self-contained exposition for algorithms audiences, including necessary algebraic background in explicit coordinate form. In particular, we assume no knowledge in abstract algebra.

Cornelius Brand
University of Regensburg
cornelius.brand@ur.de

Radu Curticapean
University of Regensburg and IT University of Copenhagen
radu.curticapean@ur.de

Baitian Li
Columbia University
bl3052@columbia.edu

Kevin Pratt
New York University
kp2154@nyu.edu

CP50

A Simple Geometric Proof of the Optimality of the Sequential Probability Ratio Test for Symmetric Bernoulli Hypotheses

This paper revisits the classical problem of determining the bias of a weighted coin, where the bias is known to be either $p = 1/2 + \varepsilon$ or $p = 1/2 - \varepsilon$, while minimizing the expected number of coin tosses and the error probability. The optimal strategy for this problem is given by Wald's Sequential Probability Ratio Test (SPRT), which compares the log-likelihood ratio against fixed thresholds to determine a stopping time. Classical proofs of this result typically rely on analytical, continuous, and non-constructive arguments. In this paper, we present a discrete, self-contained proof of the optimality of the SPRT for this problem. We model the problem as a biased random walk on the two-dimensional (heads, tails) integer lattice, and model strategies as marked stopping times on this lattice. Our proof takes a straightforward greedy approach, showing how any arbitrary strategy may be transformed into the optimal, parallel-line "difference policy" corresponding to the SPRT, via a sequence of local perturbations that improve a Bayes risk objective.

Chirag Pabbaraju
Stanford University
cpabbara@stanford.edu

Gregory Valiant
Microsoft
gregory.valiant@gmail.com

Rishi Verma
Stanford University
verma.rishiraj@gmail.com

CP50

A Note on Ordered Ruzsa-Szemerdi Graphs

A recent breakthrough of Behnezhad and Ghafari (FOCS 2024) and subsequent work of Assadi, Khanna, and Kiss (SODA 2025) gave algorithms for the fully dynamic $(1 - \varepsilon)$-approximate maximum matching problem whose runtimes are determined by a purely combinatorial quantity:

the maximum density of *Ordered Ruzsa-Szemerédi* (ORS) graphs. We say a graph G is an (r,t)-ORS graph if its edges can be partitioned into t matchings $M_1, M_2, \ldots, M_t$ each of size r, such that for every i, M_i is an induced matching in the subgraph $M_i \cup M_{i+1} \cup \cdots \cup M_t$. This is a relaxation of the extensively-studied notion of a Ruzsa-Szemerédi (RS) graph, the difference being that in an RS graph each M_i must be an induced matching in G. In this note, we show that these two notions are roughly equivalent. Specifically, let $ORS(n)$ be the largest t such that there exists an n-vertex ORS-$(\Omega(n), t)$ graph, and define $RS(n)$ analogously. We show that if $ORS(n) \geq \Omega(n^c)$, then for any fixed $\delta > 0$, $RS(n) \geq \Omega(n^{c(1-\delta)})$.

Kevin Pratt
New York University
kp2154@nyu.edu

CP50

Tree Violation Distance under Constraints

We consider numerical taxonomy problems, where we are given an $n \times n$ distance matrix D and the goal is to compute the closest tree metric to D. Farach, Kannan, and Warnow [STOC 1993] suggested that constrained versions of these problems (where the input contains lower bound constraints on the tree distances) are relevant for applications. Surprisingly, one cannot hope for much better; handling both upper and lower bound constraints simultaneously is intractable. When the objective is to minimize the number of distance disagreements between D and the output tree (Tree Violation Distance), we show we can achieve a constant factor approximation, even in the constrained setting. Regarding techniques, we identify a possible reason for the lack of progress being a reduction used in all results in the unconstrained setting, which breaks in the constrained setting. In particular, we exhibit a simple structure (bad triplet) consisting of only 3 nodes that makes this reduction fail. Interestingly, handling these bad triplets is sufficient: we show that for Tree Violation Distance, one can design a very simple alternative reduction based on repeatedly identifying bad triplets. Additionally, we show that the reduced problem Ultrametric Violation Distance can be solved within a constant-factor approximation, even in the constrained setting. At the core of our solution is a simple LP rounding algorithm that rounds each variable to the closest integer.

Debarati Das
Pennsylvania State University
debaratix710@gmail.com

Evangelos Kipouridis
Max Planck Institute for Informatics, SIC
kipoujr@hotmail.com

Joachim Spoerhase
University of Liverpool, United Kingdom
joachim.spoerhase@liverpool.ac.uk

CP50

Quantum Search on Computation Trees

We show a simple generalization of the quantum walk algorithm for search in backtracking trees by Montanaro (ToC 2018) to the case where vertices can have different times of computation. If a vertex v in the tree of depth D is computed in t_v steps from its parent, then we show that detection of a marked vertex requires $O(\sqrt{TD})$ queries to the steps of the computing procedures, where $T = \sum_v t_v^2$. This framework provides an easy and convenient way to re-obtain a number of other quantum frameworks like variable time search, quantum divide & conquer and bomb query algorithms. The underlying algorithm is simple, explicitly constructed, and has low poly-logarithmic factors in the complexity. As a corollary, this gives a quantum algorithm for variable time search with unknown times with optimal query complexity $O(\sqrt{T \log \min(n, t_{\max})})$, where $T = \sum_i t_i^2$ and $t_{\max} = \max_i t_i$ if t_i is the number of steps required to compute the i-th variable. This resolves the open question of the query complexity of variable time search, as the matching lower bound was recently shown by Ambainis, Kokainis and Vihrovs (TQC'23). As another result, we obtain an $\tilde{O}(n)$ time algorithm for the geometric task of determining if any three lines among n given intersect at the same point, improving the $O(n^{1+o(1)})$ algorithm of Ambainis and Larka (TQC'20).

Jevgenijs Vihrovs
University of Latvia
jevgenijs.vihrovs@lu.lv

CP51

Directed and Undirected Vertex Connectivity Problems Are Equivalent for Dense Graphs

Vertex connectivity and its variants are among the most fundamental problems in graph theory, with decades of extensive study and numerous algorithmic advances. The directed variants of vertex connectivity are usually solved by manually extending fast algorithms for undirected graphs, which has required considerable effort. In this paper, we present an extremely simple reduction from directed to undirected vertex connectivity for dense graphs. As immediate corollaries, we vastly simplify the proof for directed vertex connectivity in $n^{2+o(1)}$ time [?], and obtain a parallel vertex connectivity algorithm for directed graphs with $n^{\omega+o(1)}$ work and $n^{o(1)}$ depth, via the undirected vertex connectivity algorithm of [?]. Our reduction further extends to the weighted, all-pairs and Steiner versions of the problem. By combining our reduction with the recent subcubic-time algorithm for undirected weighted vertex cuts [?], we obtain a subcubic-time algorithm for weighted directed vertex connectivity, improving upon a three-decade-old bound [?] for dense graphs. For the all-pairs version, by combining the conditional lower bounds on the all-pairs vertex connectivity problem for directed graphs [?], we obtain an alternate proof of the conditional lower bound for the all-pairs vertex connectivity problem on undirected graphs, vastly simplifying the proof by [?].

Olivier Fischer
ETH Zürich
olivier.fischer@inf.ethz.ch

Yonggang Jiang
MPI-INF
yjiang@mpi-inf.mpg.de

Sagnik Mukhopadhyay
University of Birmingham
schwagznikst@gmail.com

Sorrachai Yingchareonthawornchai
Aalto University

sorrachai.cp@gmail.com

CP51

Space-Efficient Hierholzer: Eulerian Cycles in O(m) Time and O(n) Space

We describe a simple variant of Hierholzer's algorithm that finds an Eulerian cycle in a (multi)graph with n vertices and m edges using $O(n \lg m)$ bits of working memory. This substantially improves the working space compared to standard implementations of Hierholzer's algorithm, which use $O(m \lg n)$ bits of space. Our algorithm runs in linear time, like the classical versions, but avoids an $O(m)$-size stack of vertices or storing information for each edge. To our knowledge, this is the first linear-time algorithm to achieve this space bound, and the method is very easy to implement. (The correctness argument, by contrast, is surprisingly subtle.) The space savings are particularly relevant for dense graphs or multigraphs with large edge multiplicities.

Ziad Ismaili Alaoui
University of Liverpool
ziad.ismaili-alaoui@liverpool.ac.uk

Detlef Plump
University of York
detlef.plump@york.ac.uk

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

CP51

Reducing Shortcut and Hopset Constructions to Shallow Graphs

We introduce a blackbox framework that simplifies all known parallel algorithms with near-linear work for single-source reachability and shortest paths in directed graphs. Specifically, existing reachability algorithms rely on constructing shortcuts; our blackbox allows these algorithms that construct shortcuts with hopbound h h to assume the input graph G G is "shallow', meaning if vertex s s can reach vertex t t, it can do so in approximately h h hops. This assumption significantly simplifies shortcut construction [**?**, **?**], resulting in simpler parallel reachability algorithms. Furthermore, our blackbox extends naturally to simplify parallel algorithms for constructing hopsets and, consequently, for computing shortest paths [**?**, **?**, **?**].

Bernhard Haeupler
ETH Zurich
bernhard.haeupler@inf.ethz.ch

Yonggang Jiang
MPI-INF
yjiang@mpi-inf.mpg.de

Thatchaphol Saranurak
University of Michigan
thsa@umich.edu

CP51

Simple Algorithms for Fully Dynamic Edge Connectivity

In the fully dynamic edge connectivity problem, the input is a simple graph G undergoing edge insertions and deletions, and the goal is to maintain its edge connectivity, denoted λ_G. We present two simple randomized algorithms solving this problem. The first algorithm maintains the edge connectivity in worst-case update time $\tilde{O}(n)$ per edge update, matching the known bound but with simpler analysis. Our second algorithm achieves worst-case update time $\tilde{O}(n/\lambda_G)$ and worst-case query time $\tilde{O}(n^2/\lambda_G^2)$, which is the first algorithm with worst-case update and query time $o(n)$ for large edge connectivity, namely, $\lambda_G = \omega(\sqrt{n})$.

Yotam Kenneth-Mordoch, Robert Krauthgamer
Weizmann Institute of Science
yotam.kenneth@weizmann.ac.il,
robert.krauthgamer@weizmann.ac.il

CP51

Lossless Derandomization for Undirected Single-Source Shortest Paths and Approximate Distance Oracles

A common step in algorithms related to shortest paths in undirected graphs is that, we select a subset of vertices as centers, then grow a ball around each vertex until a center is reached. We want balls to be small. A randomized algorithm can sample r centers to achieve the optimal expected ball size of $\Theta(n/r)$. A folklore derandomization is to use the $O(\log n)$ approximation for set cover in the hitting set version where we want to hit all the balls with the centers. However, the extra $O(\log n)$ factor is sometimes too expensive. For example, the recent $O(m\sqrt{\log n \log\log n})$ undirected single-source shortest path algorithm [DMSY23] beats Dijkstra's algorithm in sparse graphs, but the folklore derandomization would make it dominated by Dijkstra's. In this paper, we exploit the fact that the sizes of these balls can be adaptively chosen by the algorithm. We propose a simple deterministic algorithm achieving the optimal ball size of $\Theta(n/r)$ on average. Furthermore, given any polynomially large cost function of the ball size, we can still achieve the optimal cost on average. It allows us to derandomize [DMSY23], resulting in a deterministic $O(m\sqrt{\log n \log\log n})$ algorithm for undirected single-source shortest path. In addition, we show that the same technique can also be used to derandomize the seminal Thorup-Zwick approximate distance oracle [TZ05], also without any loss in the time/space complexity.

Shuyi Yan
University of Copenhagen
shya@di.ku.dk

CP52

Competitive Online Transportation Simplified

The setting for the online transportation problem is a metric space M, populated by m parking garages of varying capacities. Over time cars arrive in M, and must be irrevocably assigned to a parking garage upon arrival in a way that respects the garage capacities. The objective is to minimize the aggregate distance traveled by the cars. In 1998, Kalyanasundaram and Pruhs conjectured that there is a $(2m-1)$-competitive deterministic algorithm for the online transportation problem, matching the optimal competitive ratio for the simpler online metric matching problem. Recently, Harada and Itoh presented the first $O(m)$-competitive deterministic algorithm for the online transportation problem. Our contribution is an alternative algorithm design and analysis that we believe is simpler.

Stephen Arndt, Benjamin Moseley
Carnegie Mellon University
sarndt@andrew.cmu.edu, moseleyb@andrew.cmu.edu

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

Marc Uetz
University of Twente
m.uetz@utwente.nl

CP52

The Harmonic Policy for Online Buffer Sharing Is $(2 + \ln N)$-Competitive: A Simple Proof

The problem of online buffer sharing is expressed as follows. A switch with n output ports receives a stream of incoming packets. When an incoming packet is accepted by the switch, it is stored in a shared buffer and awaits its transmission through its destination output port. Each output port transmits one packet per time unit. The problem is to find an algorithm to accept or reject each packet in order to maximize the total number of transmitted packets. In TCS 2004, Kesselman and Mansour considered the problem of online buffer sharing which models most deployed internet switches. They presented the Harmonic policy and proved that it is $(2 + \ln n)$-competitive, which is the best known competitive ratio for this problem. The Harmonic policy unfortunately saw less practical relevance as it performs n threshold checks per packets which is deemed costly in practice. While the Harmonic policy is elegant, the original proof is also rather complex and involves a lengthy matching routine. This note presents a simplified Harmonic policy, both in terms of implementation and proof. First, we show that the Harmonic policy can be implemented with a constant number of threshold checks per packet, matching the widely deployed *Dynamic Threshold* policy. Second, we present a simple proof that shows the Harmonic policy is $(2 + \ln n)$-competitive. In contrast to the original proof, the current proof is direct and relies on a 3-partitioning of the packets.

Vamsi Addanki
Purdue University
vamsi.addanki666@gmail.com

Julien Dallot
TU Berlin
judafa@protonmail.com

Leon Kellerhals
TU Clausthal
leon.kellerhals@tu-clausthal.de

Maciej Pacut, Stefan Schmid, Aleksander Figiel
TU Berlin
maciej@inet.tu-berlin.de, stefan.schmid@inet.tu-berlin.de, a.figiel@tu-berlin.de

CP52

Text Indexing and Pattern Matching with Ephemeral Edits

A sequence $e_0, e_1, \ldots$ of edit operations in a string T is called ephemeral if operation e_i constructing string T^i, for all $i = 2k$ with $k \in \mathbb{N}$, is reverted by operation e_{i+1} that reconstructs T. Such a sequence arises when processing a stream of independent edits or testing hypothetical edits. We introduce text indexing with ephemeral substring edits, a new version of text indexing. The goal is to design a data structure over a given text that supports subsequent pattern matching queries with ephemeral substring insertions, deletions, or substitutions in the text; we require insertions and substitutions to be of constant length. In particular, we preprocess a text $T = T[0\,..\,n)$ over an integer alphabet $\Sigma = [0, \sigma)$ with $\sigma = n^{\mathcal{O}(1)}$ in $\mathcal{O}(n)$ time. Then, we preprocess any arbitrary pattern $P = P[0\,..\,m)$ given online in $\mathcal{O}(m \log\log m)$ time and $\mathcal{O}(m)$ space and support any ephemeral sequence of edit operations in T. Before reverting the ith operation, we report all Occ occurrences of P in T^i in $\mathcal{O}(\log\log n + Occ)$ time. We also give an optimal solution for pattern matching with ephemeral substring edits. Before reverting the ith operation, we report all Occ occurrences of a fixed P in T^i in $\mathcal{O}(Occ)$ time.

Solon P. Pissis
CWI Amsterdam
solon.pissis@cwi.nl

CP52

Maximal Palindromes in Mpc: Simple and Optimal

In the classical longest palindromic substring (LPS) problem, we are given a string S of length n, and the task is to output a longest palindromic substring in S. Gilbert, Hajiaghayi, Saleh, and Seddighin [SPAA 2023] showed how to solve the LPS problem in the Massively Parallel Computation (MPC) model in $\mathcal{O}(1)$ rounds using $\widetilde{\mathcal{O}}(n)$ total memory, with $\widetilde{\mathcal{O}}(n^{1-\epsilon})$ memory per machine, for any $\epsilon \in (0, 0.5]$. We present a simple and optimal algorithm to solve the LPS problem in the MPC model in $\mathcal{O}(1)$ rounds. The total time and memory are $\mathcal{O}(n)$, with $\mathcal{O}(n^{1-\epsilon})$ memory per machine, for any $\epsilon \in (0, 0.5]$. A key attribute of our algorithm is its ability to compute all maximal palindromes in the same complexities. Furthermore, our new insights allow us to bypass the constraint $\epsilon \in (0, 0.5]$ in the Adaptive MPC model. Our algorithms and the one proposed by Gilbert et al. for the LPS problem are randomized and succeed with high probability.

Solon P. Pissis
CWI Amsterdam
solon.pissis@cwi.nl

CP52

A Simple Analysis of Ranking in General Graphs

We provide a simple combinatorial analysis of the Ranking algorithm, originally introduced in the seminal work by Karp, Vazirani, and Vazirani [KVV90], demonstrating that it achieves a $(1/2 + c)$-approximate matching for general graphs for $c \geq 0.005$.

Tao Yu, Mahsa Derakhshan
Northeastern University
yu.tao2@northeastern.edu,
m.derakhshan@northeastern.edu

Mohammad Roghani
Stanford University
roghani@stanford.edu

Mohammad Saneian

Northeastern University
saneian.m@northeastern.edu

CP52

A Learning Perspective on Random-Order Covering Problems

In the random-order online set cover problem, the instance with m sets and n elements is chosen in a worst-case fashion, but then the elements arrive in a uniformly random order. Can this random-order model allow us to circumvent the bound of $O(\log m \log n)$-competitiveness for the adversarial arrival order model? This long-standing question was recently resolved by (Gupta et al., FOCS 2021), who gave an algorithm that achieved an $O(\log mn)$-competitive ratio. While their LearnOrCover was inspired by ideas in online learning (and specifically the multiplicative weights update method), the analysis proceeded by showing progress from first principles. In this work, we show a concrete connection between random-order set cover and stochastic mirror-descent/online convex optimization. In particular, we show how additive/multiplicative regret bounds for the latter translate into competitiveness for the former. Indeed, we give a clean recipe for this translation, allowing us to extend our results to covering integer programs, set multicover, and non-metric facility location in the random order model, matching (and giving simpler proofs of) the previous applications of the LearnOrCover framework.

Anupam Gupta
New York University, USA
anupam.g@nyu.edu

Marco Molinaro
Microsoft Research (Redmond) and Computer Science Department
mmolinaro@microsoft.com

Matteo Russo
Sapienza, University of Rome, Italy
mrusso@diag.uniroma1.it

CP53

Approximating Graphic Multi-Path Tsp and Graphic Ordered Tsp

The path version of the Traveling Salesman Problem is one of the most well-studied variants of the ubiquitous TSP. Its generalization, the Multi-Path TSP, has recently been used in the best known algorithm for path TSP by Traub and Vygen [Cambridge University Press, 2024]. The best known approximation factor for this problem is 2.214 by Bhm, Friggstad, Mmke and Spoerhase [SODA 2025]. In this paper we show that for the case of graphic metrics, a significantly better approximation guarantee of 2 can be attained. Our algorithm is based on sampling paths from a decomposition of the flow corresponding to the optimal solution to the LP for the problem, and connecting the left-out vertices with doubled edges. The cost of the latter is twice the optimum in the worst case; we show how the cost of the sampled paths can be absorbed into it without increasing the approximation factor. Furthermore, we prove that any below-2 approximation algorithm for the special case of the problem where each source is the same as the corresponding sink yields a below-2 approximation algorithm for Graphic Multi-Path TSP. We also show that our ideas can be utilized to give a factor 1.791-approximation algorithm for Ordered TSP in graphic metrics, for which the aforementioned paper [SODA 2025] and Armbruster, Mnich and Ngele [APPROX 2024] give a 1.868-approximation algorithm in general metrics.

Morteza Alimi
University of Augsburg
morteza.alimi@uni-a.de

Niklas Dahlmeier
RWTH Aachen
dahlmeier@gdm.rwth-aachen.de

Tobias Mömke
University of Augsburg, Germany
moemke@informatik.uni-augsburg.de

Philipp Pabst, Laura Vargas Koch
RWTH Aachen
philipp.pabst@oms.rwth-aachen.de,
vargaskoch@gdm.rwth-aachen.de

CP53

No SDP Needed for Efficiently Recovering Planted Regular Bipartite Graphs

We study exact recovery of an ¡i¿arbitrary¡/i¿ (adversarial) d-regular bipartite subgraph S of size k planted in a random Erdos-Renyi $G(n,p)$ graph. Prior algorithms by Kumar, Louis, and Paul [ICALP 2022], and Louis, Paul, and Raghavendra [COLT 2025] either solve a costly semidefinite program (SDP), or perform subspace enumeration that has a runtime exponential in a parameter $\ell \approx k/d$. We show that nothing exotic is needed: and our simple ¡i¿Eigenspace Thresholding algorithm¡/i¿ which, (i) computes the bottom few ($\ell \approx k/d$) many eigenvectors of the adjacency matrix, (ii) thresholds their coordinates at absolute value $c/\sqrt{\ell k}$ (for small constant c) and accumulates the selected vertices, and (iii) uses a standard matching-based clean-up step to fully recover the planted graph. We show that under standard assumptions, the procedure exactly recovers the planted bipartite graph up to the conjectured computational threshold given by Kothari, Vempala, Wein, and Xu (COLT 2023) of $d = \Omega_p(\sqrt{n})$, including the critical regime $d \approx pk$. It requires just a handful of eigenvector computations, followed by a simple cleanup step, making it both conceptually simple and practically efficient.

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

Rameesh Paul
Indian Institute of Science, Bangalore
rameeshpaul@iisc.ac.in

CP53

Unsplittable Cost Flows from Unweighted Error-Bounded Variants

A famous conjecture of Goemans on single-source unsplittable flows states that one can turn any fractional flow into an unsplittable one of no higher cost, while increasing the load on any arc by at most the maximum demand. Despite extensive work on the topic, only limited progress has been made. Recently, Morell and Skutella suggested an alternative conjecture, stating that one can turn any fractional flow into an unsplittable one without changing the load on any arc by more than the maximum demand. We show that their conjecture implies Goemans conjecture (with a violation of twice the maximum demand). To this end, we generalize a technique of Linhares and Swamy, used to obtain a low-cost chain-constrained spanning tree from an algorithm without cost guarantees. Whereas Linhares and Swamys proof relies on Langrangian duality, we provide a very simple elementary proof of a generalized version, which we hope to be of independent interest. Moreover, we show how this technique can also be used in the context of the weighted ring loading problem, showing that cost-unaware approximation algorithms can be transformed into approximation algorithms with additional cost guarantees.

Laura Vargas Koch
RWTH Aachen
vargaskoch@gdm.rwth-aachen.de

Chaitanya Swamy
University of Waterloo
cswamy@uwaterloo.ca

Vera Traub
Department of Computer Science
ETH Zurich
vtraub@ethz.ch

Rico Zenklusen
ETH Zurich
ricoz@ethz.ch

CP53

A Simple and Fast $(3+\varepsilon)$-Approximation for Constrained Correlation Clustering

In Constrained Correlation Clustering, the goal is to cluster a complete signed graph in a way that minimizes the number of negative edges inside clusters plus the number of positive edges between clusters, while respecting hard constraints on how to cluster certain friendly or hostile node pairs. Fischer et al. [FKKT25] recently developed a $\tilde{O}(n^3)$-time 16-approximation algorithm for this problem. We settle an open question posed by these authors by designing an algorithm that is equally fast but brings the approximation factor down to $(3+\varepsilon)$ for arbitrary constant $\varepsilon > 0$. Although several new algorithmic steps are needed to obtain our improved approximation, our approach maintains many advantages in terms of simplicity. In particular, it relies mainly on rounding a (new) covering linear program, which can be approximated quickly and combinatorially. Furthermore, the rounding step amounts to applying the very familiar Pivot algorithm to an auxiliary graph. Finally, we develop much simpler algorithms for instances that involve only friendly or only hostile constraints.

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

CP53

Yet Another Proof That BPP $\subseteq$ PH

We present a new, simplified proof that the complexity class BPP is contained in the Polynomial Hierarchy (PH), using k-wise independent hashing as the main tool. We further extend this approach to recover several other previously known inclusions between complexity classes. Our techniques are inspired by the work of Bellare, Goldreich, and Petrank (Information and Computation, 2000).

Ilya Volkovich
Boston College
ilya.volkovich@bc.edu