Searchable Abstracts Document

ACM-SIAM Symposium on
Discrete Algorithms
(SODA25)

January 12–15, 2025

SIAM Symposium on
Algorithm Engineering and Experiments
(ALENEX25)

January 12–13, 2025

SIAM Symposium on
Simplicity in Algorithms
(SOSA25)

January 13–15, 2025

This document was current as of December 9, 2024. Abstracts appear as submitted.

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

IP1

Fully Dynamic Matching, Matching Sparsifiers, and (Ordered) Ruzsa-Szemerédi Graphs

The fully dynamic matching problem involves efficiently maintaining a near-optimal matching in a graph undergoing edge insertions and deletions. Despite significant effort, the goal of designing highly efficient fully dynamic matching algorithms has remained elusive. A promising natural avenue for obtaining faster algorithms is to maintain a matching sparsifier, that is, a sparse subgraph that approximately preserves large matchings in every induced subgraph of the original graph. The success of this approach hinges on a positive resolution of two challenges: showing the existence of sparse matching sparsifiers and designing efficient algorithms for constructing them. As it turns out, the first challenge is intimately connected to the question of determining the density of Ruzsa-Szemerédi (RS) graphs, namely, graphs whose edges can be partitioned into induced matchings of linear size. However, even an optimistic resolution of the RS graph density question would still leave open the second challenge. An exciting recent insight has highlighted a relaxed notion of RS graphs, called ordered RS graphs, that surprisingly allows folding both challenges above into a single combinatorial question on the density of such graphs. This talk will explore the fascinating interplay between these ideas, and will show how this interplay has led to a conditional dream result for fully dynamic matching.

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

IP2

When and Why do Efficient Algorithms Exist (for Constraint Satisfaction and Beyond)?

Computational problems exhibit a diverse range of behaviors in terms of how quickly and effectively they can be solved. What underlying mathematical structure (or lack thereof) in a computational problem leads to an efficient algorithm for solving it (or dictates its intractability)? Given the vast landscape of problems and algorithmic approaches, it would be simplistic to hope for a universal theory explaining the underpinnings of easiness/hardness. Yet, in the realm of constraint satisfaction problems (CSP), the algebraic dichotomy theorem gives a definitive answer: a polynomial time algorithm exists when there are nontrivial local operations called polymorphisms under which the solution space is closed; otherwise the problem is NP-complete. Inspired and emboldened by this, one might speculate a polymorphic principle in more general contexts, with appropriate notions of symmetries governing the existence of efficient algorithms. Beginning with some background on CSP and the polymorphic approach to understand their complexity, the talk will discuss some extensions beyond CSP where the polymorphic principle seems promising (yet far from understood). In particular, we will discuss "promise CSP' where one is allowed to satisfy a relaxed version of the constraints, a framework that includes problems such as approximate graph coloring and discrepancy minimization. We will provide glimpses of the emerging theory characterizing the (in)tractability of interesting classes of promise CSP, touch upon connections to optimization (linear and affine relaxations), and highlight some of the many challenges that remain.

Venkatesan Guruswami
UC Berkeley
venkatg@berkeley.edu

IP3

Learning in Environments with Carryover Effects

We will consider learning evolving environments. In the last decades we have developed good understanding about impact of learning behavior in repeated games (including online auctions) with the same game repeated each iteration. We will consider learning environments that evolve due to the actions of the player. A classical such environment is budgeted auctions, where the remaining budget affects the learners ability to bid on later items. In this talk, we will consider settings where the participants care not only about how frequently they win, but also about how their winnings are distributed over time, making the value of a win depend on the history. In joint work with Giannis Fikioris, Bobby Kleinberg, Yoav Kolumbus, Raunak Kumar, and Yishay Mansour, we study the optimal policies for a simple model such auctions and offer learning algorithms for the bidders that achieve low regret against the optimal bidding policy in a Bayesian online setting.

Éva Tardos
Cornell University, U.S.
eva.tardos@cornell.edu

CP1

Optimal Neighborhood Exploration for Dynamic Independent Sets

A dynamic graph algorithm is a data structure that supports edge insertions, deletions, and problem specific queries. While extensive research exists on dynamic algorithms for graph problems solvable in polynomial time, most of these algorithms have not been implemented or empirically evaluated. This work addresses the NP-complete maximum weight as well as the maximum cardinality independent set problem in a dynamic setting, applicable to areas like dynamic map-labeling and vehicle routing. Exact solvers can find optimal solutions but have exponential worst-case runtimes. Conversely, heuristic algorithms use local search techniques to improve solutions by optimizing vertices. In this work, specifically we introduce a novel local search technique called *optimal neighborhood exploration*. This technique creates independent subproblems that are solved to optimality, leading to improved overall solutions. Through numerous experiments, we assess the effectiveness of our approach and compare it with other state-of-the-art dynamic solvers. Our algorithm features a parameter, the subproblem size, that balances running time and solution quality.

Jannick Borowitz, Ernestine Großmann
Heidelberg University
jannick.borowitz@uni-heidelberg.de,
e.grossmann@informatik.uni-heidelberg.de

Christian Schulz
Heidelberg University, Germany
christian.schulz@informatik.uni-heidelberg.de

CP1

Engineering Fully Dynamic Exact Δ-Orientation Algorithms

A (fully) dynamic graph algorithm is a data structure that supports edge insertions, edge deletions, and answers spe-

cific queries pertinent to the problem at hand. In this work, we address the fully dynamic edge orientation problem, also known as the fully dynamic Δ-orientation problem. The objective is to maintain an orientation of the edges in an undirected graph such that the out-degree of any vertex remains low. When edges are inserted or deleted, it may be necessary to reorient some edges to prevent vertices from having excessively high out-degrees. In this paper, we introduce the first algorithm that maintains an optimal edge orientation during both insertions and deletions.In experiments comparing with recent nearly exact algorithms, we achieve a 32% lower running time. The update time of our algorithm is up to 6 orders of magnitude faster than static exact algorithms.

Christian Schulz
Heidelberg University, Germany
christian.schulz@informatik.uni-heidelberg.de

Ernestine Großmann, Henrik Reinstädtler, Fabian Walliser
Heidelberg University
e.grossmann@informatik.uni-heidelberg.de,
henrik.reinstaedtler@informatik.uni-heidelberg.de,
fabian.walliser@stud.uni-heidelberg.de

CP1

A Simpler Approach for Monotone Parametric Minimum Cut: Finding the Breakpoints in Order

We present parametric breadth-first search (PBFS), a new algorithm for solving the parametric minimum cut problem in a network with source-sink-monotone capacities. The objective is to find the set of breakpoints, i.e., the points at which the minimum cut changes. It is well known that this problem can be solved in the same asymptotic runtime as the static minimum cut problem. However, existing algorithms that achieve this runtime bound involve fairly complicated steps that are inefficient in practice. PBFS uses a simpler approach that discovers the breakpoints in ascending order, which allows it to achieve the desired runtime bound while still performing well in practice. We evaluate our algorithm on benchmark instances from polygon aggregation and computer vision. Polygon aggregation was recently proposed as an application for parametric minimum cut, but the monotonicity property has not been exploited fully. PBFS outperforms the state of the art on most benchmark instances, usually by a factor of 2–3. It is particularly strong on instances with many breakpoints, which is the case for polygon aggregation. Compared to the existing min-cut-based approach for polygon aggregation, PBFS scales much better with the instance size. On large instances with millions of vertices, it is able to compute all breakpoints in a matter of seconds.

Jonas Sauer
University of Bonn
jsauer1@uni-bonn.de

Arne Beines
Formerly of University of Bonn
beines-arne@t-online.de

Michael Kaibel, Philip Mayer, Petra Mutzel
University of Bonn
s6mikaib@uni-bonn.de, pmayer@cs.uni-bonn.de, petra.mutzel@cs.uni-bonn.de

CP1

Parallel Cluster-BFS and Applications to Shortest Paths

Breadth-first Search (BFS) is one of the most important graph processing subroutines, especially for computing the unweighted distance. Many applications may require running BFS from multiple sources. Sequentially, when running BFS on a cluster of nearby vertices, a known optimization is using bit-parallelism. Given a subset of vertices with size k and the distance between any pair of them is no more than d, BFS can be applied to all of them in total work $O(dm(k/w + 1))$, where w is the length of a word in bits and m is the number of edges. We will refer to this approach as cluster-BFS (C-BFS). Such an approach has been studied and shown effective both in theory and in practice in the sequential setting. However, it remains unknown how this can be combined with thread-level parallelism. In this paper, we focus on designing efficient parallel C-BFS based on BFS to answer unweighted distance queries. Our solution combines the strengths of bit-level parallelism and thread-level parallelism, and achieves significant speedup over the plain sequential solution. We also apply our algorithm to real-world applications. In particular, we identified another application (landmark-labeling for the approximate distance oracle) that can take advantage of parallel C-BFS. Under the same memory budget, our new solution improves accuracy and/or time on all the 18 tested graphs.

Letong Wang
University of California, Riverside
lwang323@ucr.edu

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

Yan Gu
University of California, Riverside
ygu@cs.ucr.edu

Yihan Sun
UC Riverside
yihans@cs.ucr.edu

CP2

Massively Parallel Minimum Spanning Tree in General Metric Spaces

We study the minimum spanning tree (MST) problem in the massively parallel computation (MPC) model, in the *strictly sublinear* regime where the space per machine is $O(n^\delta)$. Here n is the number of vertices and $\delta \in (0, 1)$ an arbitrary constant. The MST problem admits a simple and folklore $O(\log n)$-round algorithm in the MPC model. For arbitrary weights, this matches a conditional lower-bound of $\Omega(\log n)$ which follows from the well-known 1vs2-Cycle conjecture. As such, much of the literature focuses on breaking the logarithmic barrier in more structured variants of the problem, such as when the vertices correspond to points in Euclidean spaces [ANOY'14, JMNZ'24]. We study general metric spaces. Namely, all pairwise weights are provided and are guaranteed to satisfy the triangle inequality, but are otherwise unconstrained. We show that

for any $\epsilon > 0$, a $(1+\epsilon)$-approximate MST can be found in $O(\log \frac{1}{\epsilon} + \log\log n)$ rounds, which is the first $o(\log n)$-round algorithm for finding any constant approximation in this setting. On the lowerbound side, we prove that under the 1vs2-Cycle conjecture, $\Omega(\log \frac{1}{\epsilon})$ rounds are needed for finding a $(1+\epsilon)$-approximate MST in general metrics. It is worth noting that while many existing lowerbounds in the MPC model under the 1vs2-Cycle conjecture only hold against "component stable" algorithms, our lowerbound applies to *all* algorithms.

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

Rajesh Jayaram, Jakub Lacki
Google Research
rkjayaram@google.com, jlacki@google.com

Vahab Mirrokni
Google Inc.
mirrokni@google.com

Peilin Zhong
Google Research
peilinz@google.com

CP2

Embedding Planar Graphs into Graphs of Treewidth $O(\log^3 n)$

Cohen-Addad, Le, Pilipczuk, and Pilipczuk [CLPP23] recently constructed a stochastic embedding with expected $1+\varepsilon$ distortion of n-vertex planar graphs (with polynomial aspect ratio) into graphs of treewidth $O(\varepsilon^{-1}\log^{13} n)$. Their embedding is the first to achieve polylogarithmic treewidth. However, there remains a large gap between the treewidth of their embedding and the treewidth lower bound of $\Omega(\log n)$ shown by Carroll and Goel [CG04]. In this work, we substantially narrow the gap by constructing a stochastic embedding with treewidth $O(\varepsilon^{-1}\log^3 n)$. We obtain our embedding by improving various steps in the CLPP construction. First, we streamline their embedding construction by showing that one can construct a low-treewidth embedding for any graph from (i) a stochastic hierarchy of clusters and (ii) a stochastic balanced cut. We shave off some logarithmic factors in this step by using a single hierarchy of clusters. Next, we construct a stochastic hierarchy of clusters with optimal separating probability and hop bound based on shortcut partition [CCLMST23, CCLMST24]. Finally, we construct a stochastic balanced cut with an improved trade-off between the cut size and the number of cuts. This is done by a new analysis of the contraction sequence introduced by [CLPP23]; our analysis gives an optimal treewidth bound for graphs admitting a contraction sequence.

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

Vincent Cohen-Addad
Google Research
cohenaddad@google.com

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

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

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

CP2

Connectivity Labeling Schemes for Edge and Vertex Faults Via Expander Hierarchies

We consider the problem of assigning short labels to the vertices and edges of a graph G so that given any query $\langle s,t,F\rangle$ with $|F| \le f$, we can determine whether s and t are still connected in $G-F$, given only the labels of $F \cup \{s,t\}$. This problem has been considered when $F \subset E$ (edge faults) or $F \subset V$ (vertex faults), where correctness is guaranteed with high probability or deterministically. Our main results are as follows. 1. We give a new deterministic labeling scheme for edge faults that uses $\tilde{O}(\sqrt{f})$-bit labels, which can be constructed in polynomial time. This improves on Dory and Parter's [PODC 2021] existential bound of $O(f\log n)$ (requiring exponential time to compute) and the efficient $\tilde{O}(f^2)$-bit scheme of Izumi, Emek, Wadayama, and Masuzawa [PODC 2023]. Our construction uses an improved edge-expander hierarchy and a distributed coding technique based on Reed-Solomon codes. 2. We improve Parter, Petruschka, and Pettie's [STOC 2024] deterministic $O(f^7\log^{13} n)$-bit labeling scheme for vertex faults to $O(f^4\log^{7.5} n)$ bits. 3. We improve the size of Dory and Parter's [PODC 2021] randomized edge fault labeling scheme from $O(\min\{f+\log n, \log^3 n\})$ bits to $O(\min\{f+\log n, \log^2 n\log f\})$ bits. We also improve the size of Parter, Petruschka, and Pettie's [STOC 2024] randomized vertex fault labeling scheme from $O(f^3\log^5 n)$ bits to $O(f^2\log^6 n)$ bits.

Yaowei Long
Tsinghua University
yaoweil@umich.edu

Seth Pettie, Thatchaphol Saranurak
University of Michigan
pettie@umich.edu, thsa@umich.edu

CP2

A Dichotomy Hierarchy for Linear Time Subgraph Counting in Bounded Degeneracy Graphs

Subgraph and homomorphism counting are fundamental algorithmic problems. Given a pattern graph H and an input graph G, we wish to count the number of H-homomorphisms/subgraphs in G. Given the massive sizes of real-world graphs and the practical importance of counting problems, we focus on when linear time algorithms are possible. Recent works (Bera et al, SODA 2021, JACM 2022) show a dichotomy theorem characterizing the patterns for which homomorphism counting is possible in linear time, for bounded degeneracy inputs G. At the other end, Neetril and Ossona de Mendez used their deep theory of sparsity to define bounded expansion graphs. They prove that, for all H, homomorphism counting can be done in linear time for bounded expansion inputs. What lies between? For a specific H, can we characterize input classes where H-homomorphism counting is possible in

linear time? We discover a hierarchy of dichotomy theorems that answer the above questions. We show the existence of an infinite sequence of graph classes $G_0 \supseteq G_1 \supseteq \ldots \supseteq G_\infty$. Fix any constant sized pattern graph H. Let $LICL(H)$ denote the length of the longest induced cycle in H. We prove the following. If $LICL(H) < 3(r+2)$, then H-homomorphisms can be counted in linear time for inputs in $\mathcal{G}_r$. If $LICL(H) \geq 3(r+2)$, then (assuming fine-grained complexity conjectures) H-homomorphism counting on inputs from $\mathcal{G}_r$ takes $\Omega(m^{1+\gamma})$ time.

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

CP2

Deterministic Edge Connectivity and Max Flow Using Subquadratic Cut Queries

We give the first deterministic algorithm that makes subquadratic queries to find the global min-cut of a simple graph in the cut query model. Given an n-vertex graph G, our algorithm makes $\widetilde{O}(n^{\frac{5}{3}})$ queries to compute the global min-cut in G. As a key ingredient, we also show an algorithm for finding s-t max-flows of size $\widetilde{O}(n)$ in $\widetilde{O}(n^{\frac{5}{3}})$ queries. We also show efficient cut-query implementations of versions of expander decomposition and isolating cuts, which may be of independent interest.

Aditya Anand, Thatchaphol Saranurak
University of Michigan
adanand@umich.edu, thsa@umich.edu

Yunfan Wang
Tsinghua University
yunfan-w20@mails.tsinghua.edu.cn

CP3

Breaking the Two Approximation Barrier for Various Consensus Clustering Problems

Consensus clustering combines the outcomes of multiple clustering algorithms on a single dataset. Given a set of clusterings on a fixed dataset, finding a clustering that minimizes total pairwise disagreements with the input clusters is known as the median partition problem. Wakabayashi [Resenhas 1998] showed that computing an exact median partition is NP-hard, while Ailon et al. [STOC 2005] developed a 4/3-approximation algorithm. This work generalizes the median partition problem, where the input includes multiple metric spaces on a fixed set. The goal is to identify a metric space minimizing the total distance to the inputs. We show that for collections of metric spaces meeting a specific proximity condition, a $(2-\zeta)$ approximation, for constant $\zeta > 0$, is achievable, providing a strict < 2 approximation for the one-median problem on ultrametrics (or tree metrics). We also address the more complex 1-center problem in the context of consensus clustering, where given m clusterings on n elements, the aim is to find a clustering that minimizes the maximum pairwise disagreements with input clusterings. For this problem, nothing better than a folklore 2-approximation was known; we present the first algorithm that improves upon this 2-approximation barrier.

Debarati Das
Pennsylvania State University
debaratix710@gmail.com

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

CP3

Clustering to Minimize Cluster-Aware Norm Objectives

We initiate the study of the following general clustering problem. We seek to partition a given set P of data points into k clusters by finding a set X of k centers and assigning each data point to one of the centers. The cost of a cluster, represented by a center $x \in X$, is a monotone, symmetric norm f (inner norm) of the vector of distances of points assigned to x. The goal is to minimize a norm g (outer norm) of the vector of cluster costs. This problem, which we call (f,g)-Clustering, generalizes many fundamental clustering problems such as k-Center, k-Median , Min-Sum of Radii, and Min-Load k-Clustering . A recent line of research (Chakrabarty, Swamy [STOC'19]) studies norm objectives that are oblivious to the cluster structure such as k-Median and k-Center. In contrast, our problem models cluster-aware objectives including Min-Sum of Radii and Min-Load k-Clustering. Our main results are as follows. First, we design a constant-factor approximation algorithm for (Top_ℓ, L_1)-Clustering where the inner norm (Top_ℓ) sums over the ℓ largest distances. Second, we design a constant-factor approximation for (L_∞, Ord)-Clustering where the outer norm is a convex combination of Top_ℓ norms (ordered weighted norm). To obtain our results, we unify various algorithmic techniques originally designed for the cluster-oblivious k-Median objective and for the cluster-aware Min-Sum Radii objective separately.

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

Evangelos Kipouridis
Saarland University and MPI Informatics, SIC
kipoujr@hotmail.com

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

CP3

Beyond 2-Approximation for K-Center in Graphs

We consider the classical k-Center problem in undirected graphs. The problem is known to have a polynomial-time 2-approximation. The conventional wisdom is that the problem is closed, as $(2-\epsilon)$-approximation is NP-hard when k is part of the input, and for constant $k \geq 2$ it requires $n^{k-o(1)}$ time under the Strong Exponential Time Hypothesis (SETH). Our first set of results show that one can beat the multiplicative factor of 2 in undirected unweighted graphs if one is willing to allow additional small additive error, obtaining $(2-\epsilon, O(1))$ approximations. We provide several algorithms that achieve such approximations for all integers k with running time $O(n^{k-\delta})$ for $\delta > 0$. For instance, for every $k \geq 2$, we obtain an $O(mn + n^{k/2+1})$ time $\left(2-\frac{1}{2k-1}, 1-\frac{1}{2k-1}\right)$-approximation to k-Center. Our second set of results are strong fine-grained lower bounds for k-Center. We give a time/approximation trade-off: under SETH, for any integer t, $n^{k/t^2-1-o(1)}$ time is needed for

any $(2-1/(2t-1), O(1))$-approximation algorithm for k-Center. This explains why our $(2-\epsilon, O(1))$ approximation algorithms have k appearing in the exponent of the running time. Our reductions also imply that, assuming ETH, the approximation ratio 2 of the known near-linear time algorithms cannot be improved by any algorithm whose running time is a polynomial independent of k, even if one allows additive error.

Yael Kirkpatrick
Massachusetts Institute of Technology
yaelkirk@mit.edu

Ce Jin
EECS
MIT
cejin@mit.edu

Virginia Vassilevska Williams
MIT
virgi@mit.edu

Nicole Wein
University of Michigan
nswein@umich.edu

CP3

Clustering Mixtures of Bounded Covariance Distributions Under Optimal Separation

We study the clustering problem where given samples from the mixture $D = \sum_{i=1}^{k} w_i P_i$, where each $w_i \geq \alpha$ for some known α, and each P_i has unknown covariance $\Sigma_i \preceq \sigma_i^2 \cdot I_d$ for unknown σ_i, the goal is to cluster the samples assuming a pairwise mean separation in the order of $(\sigma_i+\sigma_j)/\sqrt{\alpha}$ for every pair P_i and P_j. 1. For the special case of nearly uniform mixtures, we give the first poly-time algorithm for this task. Prior work either required separation scaling with the maximum cluster standard deviation (i.e. $\max_i \sigma_i$) [DKKLT22] or required both additional structural assumptions and mean separation scaling as a large degree polynomial in $1/\alpha$ [BKK22]. 2. For general-weight mixtures, although accurate clustering is information-theoretically impossible, we introduce the notion of a *clustering refinement* — a list of not-too-small subsets satisfying a similar separation, that can be merged into a clustering approximating the ground truth — and show such a refinement can be efficiently computed. Furthermore, under a variant of the "no large sub-cluster" condition of [BKK22], our algorithm outputs an accurate clustering, not just a refinement. We finally show that the property holds whp for mixtures of well-conditioned high-dimensional log-concave distributions. Moreover, our algorithm is robust to a fraction of adversarial outliers comparable to α.

Thanasis Pittas
University of WisconsinMadison
pittas@wisc.edu

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

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

Jasper Lee
University of California, Davis
jasperlee@ucdavis.edu

CP3

Dynamic Consistent k-Center Clustering with Optimal Recourse

Given points from an arbitrary metric space and a sequence of point updates, what is the minimum recourse per update (i.e., the minimum number of changes needed to the set of centers after an update), in order to maintain a constant-factor approximation to a k-clustering problem? This question has received attention in recent years under the name consistent clustering. In this paper we study the k-center objective in the fully dynamic setting, where the update is either a point insertion or a point deletion. Before our work, Lacki, Haeupler, Grunau, Rozhon, and Jayaram [SODA '24] gave a deterministic fully dynamic constant-factor approximation algorithm for the k-center objective with worst-case recourse of 2 per point update. In this work, we prove that the k-center problem admits optimal recourse bounds by developing a deterministic fully dynamic constant-factor approximation algorithm with worst-case recourse of 1 per point update. Our algorithm performs simple choices based on light data structures, and thus is arguably more direct and faster than the previous one which uses a sophisticated structure. Additionally, we develop a new deterministic decremental algorithm and a new deterministic incremental algorithm, both of which maintain a 6-approximate k-center solution with worst-case recourse of 1 per point update. Our incremental algorithm improves over the 8-approximation algorithm by Charikar, Chekuri, Feder, and Motwani [STOC '97].

Sebastian Forster, Antonis Skarlatos
University of Salzburg
forster@cs.sbg.ac.at, antonis.skarlatos@plus.ac.at

CP4

Nearly Tight Bounds on Testing of Metric Properties

Given a non-negative $n \times n$ matrix viewed as a set of distances between n points, we consider the property testing problem of deciding if it is a metric. We also consider the same problem for two special classes of metrics tree metrics and ultrametrics. For general metrics, our paper is the first to consider these questions. We prove an upper bound of $O(n^{2/3}/\epsilon^{4/3})$ on the query complexity for this problem. Our algorithm is simple, but the analysis requires great care in bounding the variance on the number of violating triangles in a sample. When ϵ is a slowly decreasing function of n (rather than a constant, as is standard), we prove a lower bound of matching dependence on n of $\Omega(n^{2/3})$, ruling out any property testers with $o(n^{2/3})$ query complexity unless their dependence on $1/\epsilon$ is super-polynomial. Next, we turn to tree metrics and ultrametrics. While there were known upper and lower bounds, we considerably improve these bounds showing essentially tight bounds of $\tilde{O}(1/\epsilon)$ on the sample complexity. We also show a lower bound of $\Omega(1/\epsilon^{4/3})$ on the query complexity. Our upper bounds are derived by doing a more careful analysis of a natural, simple algorithm. For the lower bounds, we construct distributions on NO instances, where it is hard to find a witness showing that these are not ultrametrics.

Yiqiao Bao
University of Pennsylvania

yiqiaob@seas.upenn.edu

Sampath Kannan
University of Pennsylvania
University of California, Berkeley
kannan@seas.upenn.edu

Erik Waingarten
Simons Institute
ewaingar@seas.upenn.edu

CP4

Near-Optimal-Time Quantum Algorithms for Approximate Pattern Matching

Approximate Pattern Matching is among the most fundamental string-processing tasks. Given a text T of length n, a pattern P of length m, and a threshold k, the task is to identify the fragments of T that are at *distance* at most k to P. We consider the two most common distances: Hamming distance (the number of mismatches or character substitutions) in *Pattern Matching with Mismatches* and edit distance (the minimum number of character insertions, deletions, and substitutions) in *Pattern Matching with Edits*. We present quantum algorithms with a time complexity of $\tilde{O}(\sqrt{nk}+\sqrt{n/m}\cdot k^2)$ for Pattern Matching with Mismatches and $\hat{O}(\sqrt{nk}+\sqrt{n/m}\cdot k^{3.5})$ for Pattern Matching with Edits; both algorithms use $\hat{O}(\sqrt{nk})$ quantum queries. The number queries is near-optimal for $k = o(m)$, and the running times are near-optimal for $k \ll m^{1/3}$ and $k \ll m^{1/6}$, respectively. As a major technical contribution, we give a faster algorithm to solve a system of b substring equations of the form $S[x..x') = S[y..y')$. The goal is to construct a generic solution string whose characters are equal only when necessary. While this is known to be possible in $O(n+b)$ time [GKRRW16], we show $\tilde{O}(b^2)$ classical time is sufficient to obtain a grammar-like $\tilde{O}(b)$-size representation of S.

Tomasz Kociumaka
INSAIT
Max Planck Institute for Informatics
tomasz.kociumaka@insait.ai

Jakob Nogler
ETH Zurich
jnogler@ethz.ch

Philip Wellnitz
Max Planck Institute for Informatics,
Saarland Informatics Campus, Germany
wellnitz@nii.ac.jp

CP4

Tight Sampling Bounds for Eigenvalue Approximation

We consider the problem of estimating the spectrum of a symmetric bounded entry (not necessarily PSD) matrix via entrywise sampling. This problem was introduced by [Bhattacharjee, Dexter, Drineas, Musco, Ray '22], where it was shown that one can obtain an ϵn additive approximation to all eigenvalues of A by sampling a principal submatrix of dimension $\frac{\mathrm{poly}(\log n)}{\epsilon^3}$. We improve their analysis by showing that it suffices to sample a principal submatrix of dimension $\tilde{O}(\frac{1}{\epsilon^2})$ (with no dependence on n). This matches known lower bounds and therefore resolves the sample complexity of this problem up to $\log \frac{1}{\epsilon}$ factors. Using similar techniques, we give a tight $\tilde{O}(\frac{1}{\epsilon^2})$ bound for obtaining an additive $\epsilon\|A\|_F$ approximation to the spectrum of A via squared row-norm sampling, improving on the previous best $\tilde{O}(\frac{1}{\epsilon^8})$ bound. We also address the problem of approximating the top eigenvector for a bounded entry, PSD matrix A. In particular, we show that sampling $O(\frac{1}{\epsilon})$ columns of A suffices to produce a unit vector u with $u^T Au \geq \lambda_1(A) - \epsilon n$. This matches what one could achieve via the sampling bound of [Musco, Musco'17] for the special case of approximating the top eigenvector, but does not require adaptivity.

William J. Swartworth
University of California, Los Angeles
wswartworth@gmail.com

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

CP4

Lower Bounds for Convexity Testing

We consider the problem of testing whether an unknown and arbitrary set $S \subseteq^n$ (given as a black-box membership oracle) is convex, versus ε-far from every convex set, under the standard Gaussian distribution. The current state-of-the-art testing algorithms for this problem make $2^{\tilde{O}(\sqrt{n})\cdot \mathrm{poly}(1/\varepsilon)}$ non-adaptive queries, both for the standard testing problem and for tolerant testing. We give the first lower bounds for convexity testing in the black-box query model: * We show that any one-sided tester (which may be adaptive) must use at least $n^{\Omega(1)}$ queries in order to test to some constant accuracy $\varepsilon > 0$. * We show that any non-adaptive *tolerant* tester (which may make two-sided errors) must use at least $2^{\Omega(n^{1/4})}$ queries to distinguish sets that are ε_1-close to convex versus ε_2-far from convex, for some absolute constants $0 < \varepsilon_1 < \varepsilon_2$. Finally, we also show that for any constant $c > 0$, any non-adaptive tester (which may make two-sided errors) must use at least $n^{1/4-c}$ queries in order to test to some constant accuracy $\varepsilon > 0$.

Xi Chen
Columbia University
xc2198@columbia.edu

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

Shivam Nadimpalli, Rocco A. Servedio
Columbia University
shivamn@mit.edu, rocco@cs.columbia.edu

Erik Waingarten
Simons Institute
ewaingar@seas.upenn.edu

CP4

Relative-Error Monotonicity Testing

The standard model of Boolean function property testing is not well suited for testing sparse functions which have few satisfying assignments, since every such function is close (in the usual Hamming distance metric) to the constant-0 function. In this work we propose and investigate a

new model for property testing of Boolean functions, called relative-error testing, which provides a natural framework for testing sparse functions. This new model defines the distance between two functions $f, g : \{0,1\}^n \to \{0,1\}$ to be

$$rel - dist(f,g) := \frac{|f^{-1}(1) \triangle g^{-1}(1)|}{|f^{-1}(1)|}.$$

To compensate for this, algorithms in the new model have access both to a black-box oracle for the function f being tested and to a source of independent uniform satisfying assignments of f. In this paper we first give a few general results about the relative-error testing model; then, we give a detailed study of algorithms and lower bounds for relative-error testing of monotone Boolean functions. We give upper and lower bounds which are parameterized by $N = |f^{-1}(1)|$, the sparsity of the function f being tested. Our results show that there are interesting differences between relative-error monotonicity testing of sparse Boolean functions, and monotonicity testing in the standard model. These results motivate further study of the testability of Boolean function properties in the relative-error model.

Tianqi Yang, Xi Chen
Columbia University
tianqi@cs.columbia.edu, xichen@cs.columbia.edu

Anindya De
University of Pennsylvania
de.anindya@gmail.com

Yizhi Huang, Yuhao Li, Shivam Nadimpalli, Rocco A. Servedio
Columbia University
yizhi@cs.columbia.edu, yuhaoli@cs.columbia.edu, shivamn@mit.edu, rocco@cs.columbia.edu

CP5

Batched K-Mer Lookup on the Spectral Burrows-Wheeler Transform

Since their emergence some two decades ago, indexes based on the Burrows-Wheeler transform (BWT) have been intensely studied and today find wide use in genomics, where they form the basis of software tools for read alignment and k-mer lookup—routine tasks in modern data-intensive bioinformatics pipelines. BWT-based indexes reduce an existential query for a pattern P of length m to a sequence of up to m pairs of rank queries on a sequence derived from the underlying indexed data. In general these rank queries exhibit poor locality of memory reference, with each pair causing one or two cache misses, something that has become generally accepted as a limitation of these indexes. However, in the above mentioned applications a typical experimental run will search for 100s of millions—even billions—of patterns using the index. In this paper we show that, taken across such a large set of patterns, rank queries do exhibit locality of memory reference and that this can be exploited by reorganising the order in which rank queries are issued. We show this leads to significant performance gains—in particular, k-mer lookup queries can be answered several times faster when a batch of patterns is treated holistically.

Simon J. Puglisi, Jarno Alanko, Elena Biagi
University of Helsinki
simon.j.puglisi@gmail.com, jarno.alanko@helsinki.fi, elena.biagi@helsinki.fi

Joel Mackenzie
University of Queensland
joel.mackenzie@uq.edu.au

CP5

Graph Neural Networks As Ordering Heuristics for Parallel Graph Coloring

The graph coloring problem asks for an assignment of the minimum number of distinct colors to vertices in an undirected graph with the constraint that no pair of adjacent vertices share the same color. The problem is a thoroughly studied NP-hard combinatorial problem with several real-world applications. As such, a number of greedy heuristics have been suggested that strike a good balance between coloring quality, execution time, and also parallel scalability. In this work, we introduce a graph neural network (GNN) based ordering heuristic and demonstrate that it outperforms existing greedy ordering heuristics both on quality and performance. Previous results have demonstrated that GNNs can produce high-quality colorings but at the expense of excessive running time. The current paper is the first that brings the execution time down to compete with existing greedy heuristics. Our GNN model is trained using both supervised and unsupervised techniques. The experimental results show that a 2-layer GNN model can achieve execution times between the largest degree first (LF) and smallest degree last (SL) ordering heuristics while outperforming both on coloring quality. Increasing the number of layers improves the coloring quality further, and it is only at four layers that SL becomes faster than the GNN. Finally, our GNN-based coloring heuristic achieves superior scaling in the parallel setting compared to both SL and LF.

Kenneth Langedal, Fredrik Manne
University of Bergen
kenneth.langedal@uib.no, fredrik.manne@uib.no

CP5

Discrete Transforms of Quantized Persistence Diagrams

Topological data analysis leverages topological features to analyze datasets, with applications in diverse fields like medical sciences and biology. A key tool of this theory is the persistence diagram, which encodes topological information but poses challenges for integration into standard machine learning pipelines. We introduce Qupid (QUantized Persistence and Integral transforms of Diagrams), a novel and simple method for vectorizing persistence diagrams. First, Qupid uses a binning procedure to turn persistence diagrams into finite measures on a grid and then applies discrete transforms to these measures. Key features are the choice of log-scaled grids that emphasize information contained near the diagonal in persistence diagrams, combined with the use of discrete transforms to enhance and efficiently encode the obtained topological information. We conduct an in-depth experimental analysis of Qupid, showing that the simplicity of our method results in very low computational costs while preserving highly competitive performances compared to state-of-the-art methods across numerous classification tasks on both synthetic and real-world datasets. Finally, we provide experimental evidence that our method is robust to a decrease in the grid resolution used.

Vadim Lebovici
Universite Sorbonne Paris Nord
lebovici@math.univ-paris13.fr

Matteo Palo
Department of Mathematics, ETH Zurich, Switzerland
mapalo@student.ethz.ch

Olympio Hacquard
KUIAS, Kyoto University, Japan
hacquard.olympio.47i@st.kyoto-u.ac.jp

Michael E. Van Huffel
Department of Mathematics, ETH Zurich, Switzerland
michavan@student.ethz.ch

CP5

Scalable Multilevel and Memetic Signed Graph Clustering

In this study, we address the complex issue of graph clustering in signed graphs, which are characterized by positive and negative weighted edges representing attraction and repulsion among nodes, respectively. The primary objective is to efficiently partition the graph into clusters, ensuring that nodes within a cluster are closely linked by positive edges while minimizing negative edge connections between them. To tackle this challenge, we first develop a scalable multi-level algorithm based on label propagation and local search. Then we develop a memetic algorithm that incorporates a multilevel strategy. This approach meticulously combines elements of evolutionary algorithms with local refinement techniques, aiming to explore the search space more effectively than repeated executions. Our experimental analysis reveals that our new algorithms significantly outperforms existing state-of-the-art algorithms.

Felix Hausberger, Marcelo Fonseca Faraj
Heidelberg University
hausberger@stud.uni-heidelberg.de,
marcelofaraj@informatik.uni-heidelberg.de

Christian Schulz
Heidelberg University, Germany
christian.schulz@informatik.uni-heidelberg.de

CP6

Lift-and-Project Integrality Gaps for Santa Claus

This paper is devoted to the study of the MaxMinDegree Arborescence (MMDA) problem in layered directed graphs of depth $\ell \leq O(\log n/\log\log n)$, which is an important special case of the Santa Claus problem. Obtaining a polylogarithmic approximation for MMDA in polynomial time is of high interest as it is a necessary condition to improve upon the well-known 2-approximation for makespan scheduling on unrelated machines by Lenstra, Shmoys, and Tardos [FOCS'87]. The only way we have to solve the MMDA problem within a polylogarithmic factor is via an elegant recursive rounding of the $(\ell-1)^{th}$ level of the Sherali-Adams hierarchy, which needs time $n^{O(\ell)}$ to solve. However, it remains plausible that one could obtain a polylogarithmic approximation in polynomial time by using the same rounding with only 1 round of the Sherali-Adams hierarchy. As a main result, we rule out this possibility by constructing an MMDA instance of depth 3 for which an integrality gap of $n^{\Omega(1)}$ survives 1 round of the Sherali-Adams hierarchy. This result is tight since it is known that after only 2 rounds the gap is at most polylogarithmic on depth-3 graphs. Second, we show that our instance can be "lifted' via a simple trick to MMDA instances of any depth $\ell \in \Omega(1) \cap o(\log n/\log\log n)$ (the whole range of interest), for which we conjecture that an integrality gap of $n^{\Omega(1/\ell)}$ survives $\Omega(\ell)$ rounds of Sherali-Adams.

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

CP6

The Submodular Santa Claus Problem

We consider the problem of allocating indivisible resources to players so as to maximize the minimum total value any player receives. This problem is sometimes dubbed the Santa Claus problem and its different variants have been subject to extensive research towards approximation algorithms over the past two decades. In the case where each player has a potentially different additive valuation function, Chakrabarty, Chuzhoy, and Khanna [FOCS'09] gave an $O(n^{\epsilon})$-approximation algorithm with polynomial running time for any constant $\epsilon > 0$ and a polylogarithmic approximation algorithm in quasi-polynomial time. We show that the same can be achieved for monotone submodular valuation functions, improving over the previously best algorithm due to Goemans, Harvey, Iwata, and Mirrokni [SODA'09], which has an approximation ratio of more than $\sqrt{n}$. Our result builds up on a sophisticated LP relaxation, which has a recursive block structure that allows us to solve it despite having exponentially many variables and constraints.

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

Sarah Morell
Technische Universität Berlin
morell@math.tu-berlin.de

Lars Rohwedder
Maastricht University
lars.rohwedder@gmail.com

CP6

A Tight $(3/2+\varepsilon)$-Approximation Algorithm for Demand Strip Packing

We consider the Demand Strip Packing problem (DSP), in which we are given a set of jobs, each specified by a processing time and a demand. The task is to schedule all jobs such that they are finished before some deadline D while minimizing the peak demand, i.e., the maximum total demand of tasks executed at any point in time. DSP is closely related to the Strip Packing problem (SP), in which we are given a set of axis-aligned rectangles that must be packed into a strip of fixed width while minimizing the maximum height. DSP and SP are known to be NP-hard to approximate to within a factor below $\frac{3}{2}$. To achieve the essentially best possible approximation guarantee, we prove a structural result. Any instance admits a solution with peak demand at most $\left(\frac{3}{2}+\varepsilon\right)Opt$ satisfying one of two properties. Either (i) the solution leaves a gap for a job with demand Opt and processing time $\mathcal{O}(\varepsilon D)$ or (ii) all jobs with demand greater than $\frac{Opt}{2}$ appear sorted by demand in immediate succession. We then provide two efficient algorithms that find a solution with maximum demand at most $\left(\frac{3}{2}+\varepsilon\right)Opt$ in the respective case. A central observation, which sets our approach apart from previous ones for DSP, is that the properties (i) and (ii) need not be efficiently decidable: We can simply run both algorithms

and use whichever solution is the better one.

Franziska Eberle
Mathematics Department
London School of Economics and Political Science
f.eberle@tu-berlin.de

Felix Hommelsheim
Universität Bremen
fhommels@uni-bremen.de

Malin Rau
Chalmers University of Technology
malin.rau@chalmers.se

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

CP6

Approximating Unrelated Machine Weighted Completion Time Using Iterative Rounding and Computer Assisted Proofs

We revisit the unrelated machine scheduling problem with the weighted completion time objective. It is known that independent rounding achieves a 1.5 approximation for the problem, and many prior algorithms improve upon this ratio by leveraging strong negative correlation schemes. On each machine i, these schemes introduce strong negative correlation between events that some pairs of jobs are assigned to i, while maintaining non-positive correlation for all pairs. Our algorithm deviates from this methodology by relaxing the pairwise non-positive correlation requirement. On each machine i, we identify many groups of jobs. For a job j and a group B not containing j, we only enforce non-positive correlation between j and the group as a whole, allowing j to be positively-correlated with individual jobs in B. This relaxation suffices to maintain the 1.5-approximation, while enabling us to obtain a much stronger negative correlation within groups using an iterative rounding procedure: at most one job from each group is scheduled on i. We prove that the algorithm achieves a $(1.36+\epsilon)$-approximation, improving upon the previous best approximation ratio of 1.4 due to Harris. While the improvement may not be substantial, the significance of our contribution lies in the relaxed non-positive correlation condition and the iterative rounding framework.

Shi Li
University at Buffalo
shili@nju.edu.cn

CP6

A Subexponential Time Algorithm for Makespan Scheduling of Unit Jobs with Precedence Constraints

In a classical scheduling problem, we are given a set of n jobs of unit length with precedence constraints, and the goal is to find a schedule of these jobs on m identical machines that minimizes the makespan. In standard 3-field notation, it is denoted as $Pm|\text{prec}, p_j = 1|C_{\max}$. For $m = 2$ machines, the problem can be solved in polynomial time. Settling the complexity for any constant $m \geq 3$ is a long-standing open question in the field, asked by Lenstra and Rinnooy Kan [OR 1978] in the late 70s and prominently featured in the textbook of Garey and Johnson. Since then, the problem has been thoroughly investigated, but nontrivial solutions had been found only in special cases or relaxed settings. For example, despite the possibility of the problem being polynomially solvable in the exact setting, just the existence of an approximation-scheme is widely regarded as a major open problem (see the survey of Bansal [MAPS 2017]), but so far, only superpolynomial approximations are known. In this paper, we make the first progress on the exact complexity of $Pm|\text{prec}, p_j = 1|C_{\max}$. We present an algorithm that runs in $2^{O(\sqrt{n}\log n)}$ time for $m = O(1)$. Before our work, only a $2^{O(n)}$ time exact algorithm was known by Held and Karp [ACM 1961].

Jesper Nederlof
Utrecht University
j.nederlof@uu.nl

Céline Swennenhuis
Utrecht University
Post NL
cswennenhuis@gmail.com

Karol Wegrzycki
Saarland University and Max Planck Institute for Informatics
wegrzycki@cs.uni-saarland.de

CP7

Nearly Optimal Dynamic Set Cover: Breaking the Quadratic-in-f Time Barrier

The dynamic set cover problem has been extensively studied since the pioneering works of [BHI, ICALP'15] and [GKKP17, STOC'17]. The input is a set system on a fixed collection of sets and a dynamic universe of elements, where each element appears in a most f sets and the cost of each set lies in the range $[1/C, 1]$; the ultimate goal is to maintain a set cover under insertions and deletions of elements, with optimal bounds on both the approximation factor and the update time. Previous work in the low-frequency regime ($f = O(\log n)$) achieved a deterministic $(1+\epsilon)f$-approximation with $O(f^2/\epsilon^3 + f/\epsilon^2 \log C)$ amortized update time [BHNW, SODA'21] and a randomized f-approximation with $O(f^2)$ expected amortized update time for the unweighted case [AS, ESA'21]. For high frequencies ($f = \Omega(\log n)$), [SU, STOC'23] achieved $(1+\epsilon)\ln n$ approximation with $O(f \log n)$ update time. We break the $\Omega(f^2)$ update time barrier with two results: 1. A randomized $(1+\epsilon)f$-approximation with $O(f/\epsilon^2 \log^* f + f/\epsilon^3 + f/\epsilon^2 \log C)$ expected amortized update time against adaptive adversaries. 2. A deterministic $(1+\epsilon)f$-approximation with $O(f \log f/\epsilon + f/\epsilon^3 + f/\epsilon^2 \log C)$ amortized update time. Assuming element updates are specified explicitly, our randomized algorithm achieves near-optimal bounds, exceeding the time needed to specify an element by only a $\log^* f$ factor.

Anton Bukov, Shay Solomon, Tianyi Zhang
Tel Aviv University
bukov.anton@gmail.com, solo.shay@gmail.com, tianyiz21@tauex.tau.ac.il

CP7

Fully Dynamic Algorithms for Graph Spanners Via Low-Diameter Router Decomposition

A t-spanner of an undirected n-vertex graph G is a sparse subgraph of G that preserves all pairwise distances between its vertices to within multiplicative factor t, also

called the stretch. It is well known that any n-vertex graph admits a $(2k-1)$-spanner with $O(n^{1+1/k})$ edges, and that this tradeoff is optimal under the Erdos Girth Conjecture. In this paper we investigate the problem of efficiently maintaining spanners in the fully dynamic setting with an adaptive adversary. One of our main results is a deterministic algorithm, that, for any $512 \leq k \leq (\log n)^{1/49}$ and $1/k \leq \delta \leq 1/400$, maintains a spanner H of a fully dynamic graph with stretch $\text{poly}(k) \cdot 2^{O(1/\delta^6)}$ and size $|E(H)| \leq O(n^{1+O(1/k)})$, with worst-case update time $n^{O(\delta)}$ and recourse $n^{O(1/k)}$. Our algorithm relies on a new technical tool that we develop, called low-diameter router decomposition. We design a deterministic algorithm that maintains a decomposition of a fully dynamic graph into edge-disjoint clusters with bounded vertex overlap, where each cluster C is a bounded-diameter router, so any reasonable multicommodity demand can be routed along short paths and with low congestion in C.

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

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

CP7

Tree-Packing Revisited: Faster Fully Dynamic Min-Cut and Arboricity

A tree-packing is a collection of spanning trees of a graph. It has been a useful tool for computing the minimum cut in static, dynamic, and distributed settings. In particular, [Tho'07] used them to obtain his dynamic min-cut algorithm with $\tilde{O}(\lambda^{14.5}\sqrt{n})$ worst-case update time. We re-examine this relationship, showing that we need to maintain fewer spanning trees for such a result; we show that we only need to pack $\Theta(\lambda^3 \log m)$ greedy trees to guarantee a 1-respecting cut or a trivial cut in some contracted graph. Based on this structural result, we then provide a deterministic algorithm for fully dynamic exact min-cut, that has $\tilde{O}(\lambda^{5.5}\sqrt{n})$ worst-case update time, for min-cut value bounded by λ. In particular, this also leads to an algorithm for general fully dynamic exact min-cut with $\tilde{O}(m^{1-1/12})$ amortized update time, improving upon $\tilde{O}(m^{1-1/31})$ [GHN+'23]. We also give the first fully dynamic algorithm that maintains a $(1+\varepsilon)$-approximation of the fractional arboricity – which is strictly harder than the integral arboricity. Our algorithm is deterministic and has $O(\alpha \log^6 m/\varepsilon^4)$ amortized update time, for arboricity at most α. We extend these results to a Monte Carlo algorithm with $O(\text{poly}(\log m, \varepsilon^{-1}))$ amortized update time against an adaptive adversary. Our algorithms work on multi-graphs as well.

Tijn De Vos
Department of Computer Science, University of Salzburg
tdevos@cs.sbg.ac.at

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

CP7

Fully Dynamic Approximate Minimum Cut in Sub-polynomial Time Per Operation

Dynamically maintaining the minimum cut in a graph G under edge insertions and deletion is a fundamental problem in dynamic graph algorithms for which no conditional lower bound on the time per operation exists. In an n-node graph the best known $(1+o(1))$-approximate algorithm takes $\tilde{O}(\sqrt{n})$ update time [Thorup2007]. If the minimum cut is guaranteed to be $(\log n)^{o(1)}$, a deterministic exact algorithm with $n^{o(1)}$ update time exists [JST2024]. We present the first fully dynamic algorithm for $(1+o(1))$-approximate minimum cut with $n^{o(1)}$ update time. Our main technical contribution is to show that it suffices to consider small-volume cuts in suitably contracted graphs.

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

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

CP7

Settling the Pass Complexity of Approximate Matchings in Dynamic Graph Streams

A semi-streaming algorithm in dynamic graph streams processes any n-vertex graph by making one or multiple passes over a stream of insertions and deletions to edges of the graph in $O(n \text{ polylog}(n))$ space. Semi-streaming algorithms for dynamic streams were first obtained in the seminal work of Ahn, Guha, and McGregor in 2012, alongside the introduction of the graph sketching technique, which remains the de facto way of designing algorithms in this model. We settle the pass complexity of approximating maximum matchings in dynamic streams via semi-streaming algorithms by improving the state-of-the-art in both upper and lower bounds: (i) We present a randomized sketching based semi-streaming algorithm for $O(1)$-approximation of maximum matching in dynamic streams using $O(\log \log n)$ passes. The approximation ratio of this algorithm can be improved to $(1+\epsilon)$ for any fixed $\epsilon > 0$ even on weighted graphs using standard techniques. This exponentially improves upon several $O(\log n)$ pass algorithms developed for this problem since the introduction of the model. (ii) We also prove that any semi-streaming algorithm (not only sketching based) for $O(1)$-approximation of maximum matching in dynamic streams requires $\Omega(\log \log n)$ passes. This presents the first multi-pass lower bound for this problem, which is already optimal, settling a longstanding open question in this area.

Janani Sundaresan
University of Waterloo
jsundaresan@uwaterloo.ca

Sepehr Assadi
Rutgers University, U.S.
sepehr@assadi.info

Soheil Behnezhad
Northeastern University
s.behnezhad@northeastern.edu

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

CP8

A Quantum Speed-Up for Approximating the Top Eigenvectors of a Matrix

Finding a good approximation of the top eigenvector of a given $d \times d$ matrix A is a basic and important computational problem, with many applications. We give two different quantum algorithms that, given query access to the entries of a Hermitian matrix A and assuming a constant eigenvalue gap, output a classical description of a good approximation of the top eigenvector: one with time complexity $\tilde{\mathcal{O}}(d^{1.75})$ and one with time complexity $d^{1.5+o(1)}$ (the first one has a slightly better dependence on the ℓ_2-error of the approximating vector than the second, and uses different techniques of independent interest). Both of our quantum algorithms provide a polynomial speed-up over the best-possible classical algorithm. Our quantum algorithms run a version of the classical power method that is robust to certain benign kinds of errors, where we implement each matrix-vector multiplication with small and well-behaved error on a quantum computer, in different ways for the two algorithms. Our first algorithm estimates the matrix-vector product one entry at a time, using a new "Gaussian phase estimation' procedure. Our second algorithm uses block-encoding techniques to compute the matrix-vector product as a quantum state, from which we obtain a classical description by a new time-efficient unbiased pure-state tomography procedure.

Yanlin Chen
Academia Sinica
yanlin@cwi.nl

Andras Gilyen
Renyi Institute Budapest
gilyen@renyi.hu

Ronald de Wolf
CWI and University of Amsterdam
rdewolf@cwi.nl

CP8

Triply Efficient Shadow Tomography

Given copies of a quantum state ρ, a shadow tomography protocol aims to learn all expectation values from a fixed set of observables, to within a given precision ϵ. We say that a shadow tomography protocol is *triply efficient* if it is sample- and time-efficient, and only employs measurements that entangle a constant number of copies of ρ at a time. Here we describe a framework for two-copy shadow tomography that uses an initial round of Bell measurements to reduce to a fractional coloring problem in an induced subgraph of G with bounded clique number. This coloring problem can be addressed using techniques from graph theory known as *chi-boundedness*. Using this framework we give the first triply efficient shadow tomography scheme for the set of local fermionic observables, which arise in a broad class of interacting fermionic systems in physics and chemistry. We also give a triply efficient scheme for the set of all n-qubit Pauli observables. Our protocols for these tasks use two-copy measurements, which is necessary: sample-efficient schemes are provably impossible using only single-copy measurements. Finally, we give a shadow tomography protocol that compresses an n-qubit quantum state into a poly(n)-sized classical representation, from which one can extract the expected value of any of the 4^n Pauli observables in poly(n) time, up to a small constant error.

Robbie King
Caltech
robbieking1000@gmail.com

David Gosset
IQC Waterloo
dngosset@gmail.com

Robin Kothari
Google Research
robin@robinkothari.com

Ryan Babbush
Google
babbush@google.com

CP8

On Estimating the Trace of Quantum State Powers

We investigate the computational complexity of estimating the trace of quantum state powers $\mathrm{tr}(\rho^q)$ for an n-qubit mixed quantum state ρ, given its state-preparation circuit of size poly(n). This quantity is closely related to the Tsallis entropy $\mathrm{S}_q(\rho) = \frac{1-\mathrm{tr}(\rho^q)}{q-1}$, where $q = 1$ corresponds to the von Neumann entropy. For any non-integer $q \geq 1 + \Omega(1)$, we provide a quantum estimator for $\mathrm{S}_q(\rho)$ with time complexity poly(n), exponentially improving the prior best results of exp(n). Our speedup is achieved by introducing efficiently computable uniform approximations of positive power functions into quantum singular value transformation. Our quantum algorithm reveals a sharp phase transition between the case of $q = 1$ and constant $q > 1$ in the computational complexity of the Quantum q-Tsallis Entropy Difference Problem (TsallisQED$_q$), particularly deciding whether the difference $\mathrm{S}_q(\rho_0) - \mathrm{S}_q(\rho_1)$ is at least 0.001 or at most -0.001: - For any $1+\Omega(1) \leq q \leq 2$, TsallisQED$_q$ is BQP-complete. - For any $1 \leq q \leq 1 + \frac{1}{n-1}$, TsallisQED$_q$ is QSZK-hard. The hardness results are derived from reductions based on new inequalities for the quantum q-Jensen-(Shannon-)Tsallis divergence with $1 \leq q \leq 2$, which are of independent interest.

Yupan Liu
Nagoya University
yupan.liu@gmail.com

Qisheng Wang
University of Edinburgh
qishengwang1994@gmail.com

CP8

Polynomial-Time Classical Simulation of Noisy IQP Circuits with Constant Depth

Sampling from the output distributions of quantum computations comprising only commuting gates, known as instantaneous quantum polynomial (IQP) computations, is believed to be intractable for classical computers, and hence this task has become a leading candidate for testing the capabilities of quantum devices. Here we demonstrate that for an arbitrary IQP circuit undergoing dephasing or de-

polarizing noise, whose depth is greater than a critical $O(1)$ threshold, the output distribution can be efficiently sampled by a classical computer. Unlike other simulation algorithms for quantum supremacy tasks, we do not require assumptions on the circuit's architecture, on anti-concentration properties, nor do we require $\Omega(\log(n))$ circuit depth. We take advantage of the fact that IQP circuits have deep sections of diagonal gates, which allows the noise to build up predictably and induce a large-scale breakdown of entanglement within the circuit. Our results suggest that quantum supremacy experiments based on IQP circuits may be more susceptible to classical simulation than previously thought. Furthermore, we show that the critical depth threshold of our algorithm is tight, and below this threshold there are noisy IQP circuits which are hard to sample from. Thus we demonstrate that noisy IQP circuits exhibit a phase transition in the computational complexity of sampling, as circuit depth is increased.

Joel Rajakumar, James Watson, Yi-Kai Liu
University of Maryland
joeldraj4@gmail.com, 05watson.j@gmail.com, yikailiu@umd.edu

CP8

Quartic Quantum Speedups for Planted Inference

We describe a quantum algorithm for the Planted Noisy kXOR problem (also known as sparse Learning Parity with Noise) that achieves a nearly quartic (4th power) speedup over the best known classical algorithm while also only using logarithmically many qubits. Our work generalizes and simplifies prior work of Hastings, by building on his quantum algorithm for the Tensor Principal Component Analysis (PCA) problem. We achieve our quantum speedup using a general framework based on the Kikuchi Method (recovering the quartic speedup for Tensor PCA), and we anticipate it will yield similar speedups for further planted inference problems. These speedups rely on the fact that planted inference problems naturally instantiate the Guided Sparse Hamiltonian problem. Since the Planted Noisy kXOR problem has been used as a component of certain cryptographic constructions, our work suggests that some of these are susceptible to super-quadratic quantum attacks.

Alexander Schmidhuber
Massachusetts Institute of Technology
alexsc@mit.edu

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

Robin Kothari
Google Research
robin@robinkothari.com

Ryan Babbush
Google
babbush@google.com

CP9

Spiderdan: Matching Augmentation in Demand-Aware Networks

Graph augmentation is a fundamental and well-studied problem that arises in network optimization. We consider a new variant of this model motivated by reconfigurable communication networks. In this variant, we consider a given physical network and the measured communication demands between the nodes. Our goal is to augment the given physical network with a matching, so that the shortest path lengths in the augmented network, weighted with the demands, are minimal. We prove that this problem is NP-hard, even if the physical network is a cycle. We then use results from demand-aware network design to provide a constant-factor approximation algorithm for adding a matching in case that only a few nodes in the network cause almost all the communication. For general real-world communication patterns, we design and evaluate a series of heuristics that can deal with arbitrary graphs as the underlying network structure. Our algorithms are validated experimentally using real-world traces (from e.g., Facebook) of data centers.

Aleksander Figiel, Darya Melnyk, André Nichterlein, Arash Pourdamghani, Stefan Schmid
TU Berlin
a.figiel@tu-berlin.de, melnyk@tu-berlin.de, andre.nichterlein@tu-berlin.de, pourdamghani@tu-berlin.de, schmiste@gmail.com

CP9

Linear Assignment on Tile-Centric Accelerators: Redesigning Hungarian Algorithm on IPUs

The Hungarian algorithm is a popular solution for the linear assignment problem that finds correspondences between sets of items. Despite its popularity, this algorithm suffers from significant efficiency shortcomings which hinders its application to large datasets. To overcome this challenge, we design a parallel version of the algorithm for novel tile-centric accelerators which consist of thousands of small processing cores equipped with memory that is local to each core. This allows them to overcome the memory latency and bandwidth limits of CPUs and GPUs, but, in turn, data and work must be carefully distributed among the many cores. We select the tile-centric Intelligence Processing Unit and discuss its capabilities and the challenges for algorithm design. We then present HunIPU, which outperforms the best GPU algorithm running 6x faster on synthetic datasets and up to 32x on real-world datasets for graph alignment.

Cheng Huang
Aarhus University
cheng@cs.au.dk

Alexander Mathiasen
Independent Researcher
alexm@graphcore.ai

Josef Dean
Graphcore
josefd@graphcore.ai

Johannes Langguth
Simula Research Laboratory
langguth@simula.no

Davide Mottin
Aarhus University
davide@cs.au.dk

Ira Assent
Aarhus University, Denmark
Department of Computer Science

ira@cs.au.dk

CP9

Exploring the Landscape of Distributed Graph Sketching

Recent work has initiated the study of dense graph processing using graph sketching methods, which drastically reduce space costs by lossily compressing information about the input graph. In this paper, we explore the strange and surprising performance landscape of sketching algorithms. We highlight both their surprising advantages for processing dense graphs that were previously prohibitively expensive to study, as well as the current limitations of the technique. Most notably, we show how sketching can avoid bottlenecks that limit conventional graph processing methods. Single-machine streaming graph processing systems are typically bottlenecked by CPU performance, and distributed graph processing systems are typically bottlenecked by network latency. We present Landscape, a distributed graph-stream processing system that uses linear sketching to distribute the CPU work of computing graph properties to distributed workers with no need for worker-to-worker communication. As a result, it overcomes the CPU and network bottlenecks that limit other systems. In fact Landscape achieves a stream ingestion rate one-fourth that of maximum sustained RAM bandwidth. Additionally, we prove that for any sequence of graph updates and queries Landscape consumes at most a constant factor more network bandwidth than is required to receive the input stream. We show that this system can ingest up to 332 million stream updates per second on a graph with 2^{17} vertices.

David Tench
Lawrence Berkeley National Lab
dtench@pm.me

Evan West
Stony Brook University
etwest@cs.stonybrook.edu

Kenny Zhang
MIT
kzzhang@csail.mit.edu

Michael A. Bender, Daniel Delayo
Stony Brook University
bender@cs.stonybrook.edu, ddelayo@cs.stonybrook.edu

Martin Farach-Colton
Rutgers University
martin@farach-colton.com

Gilvir Gill
Stony Brook University
gigill@cs.stonybrook.edu

Tyler Seip
MongoDB
tylerjseip@gmail.com

Victor Zhang
Meta Platforms
xskt.victor@gmail.com

CP9

The Constrained Layer Tree Problem and Applications to Solar Farm Cabling

Motivated by the cabling of solar farms, we study the problem Constrained Layer Tree. At its core, it asks whether there exists a tree that connects a set of sources (the leaves) to one sink (the root) such that certain capacity constraints at the inner nodes are satisfied. Our main algorithmic contribution is a dynamic program with various optimizations for Constrained Layer Tree. It outperforms the previously used MILP by multiple orders of magnitude. Moreover, our experiments show that the somewhat abstract problem Constrained Layer Tree is actually the core of the cabling problem in solar farms, i.e., the feasible solution produced by our dynamic program can be used to bootstrap an MILP that can then find good solutions for the original cabling problem efficiently.

Thomas Bläsius, Max Göttlicher, Sascha Gritzbach, Wendy Yi
Karlsruhe Institute of Technology
thomas.blaesius@kit.edu, max.goettlicher@kit.edu, sascha.gritzbach@gmx.de, wendy.yi@kit.edu

CP10

Constraint Satisfaction Problems with Advice

We initiate the study of algorithms for constraint satisfaction problems with ML oracle advice. We introduce two models of advice and then design approximation algorithms for Max-Cut, Max 2-Lin, and Max 3-Lin in these models. In particular, we show the following. 1. For Max-Cut and Max 2-Lin, we design an algorithm that yields near-optimal solutions when the average degree is larger than a threshold degree, which only depends on the amount of advice and is independent of the instance size. We also give an algorithm for nearly satisfiable Max 3-Lin instances with quantitatively similar guarantees. 2. Further, we provide impossibility results for algorithms in these models. In particular, under standard complexity assumptions, we show that Max 3-Lin is still $1/2 + \eta$ hard to approximate given access to advice, when there are no assumptions on the instance degree distribution. Additionally, we also show that Max 4-Lin is $1/2 + \eta$ hard to approximate even when the average degree of the instance is linear in the number of variables.

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

Konstantin Makarychev
Northwestern University
konstantin@northwestern.edu

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

CP10

Approximately Counting Knapsack Solutions in Subquadratic Time

We revisit the classic Knapsack problem, which asks to count the Boolean points $(x_1, x_2, \ldots, x_n) \in \{0,1\}^n$ in a

given half-space $\sum_{i=1}^{n} W_i x_i \leq T$. This P-complete problem is known to admit $(1 \pm \varepsilon)$-approximation. Before this work, [Dyer, STOC 2003]'s $\widetilde{O}(n^{2.5} + n^2\varepsilon^{-2})$-time randomized approximation scheme remains the fastest known in the natural regime of $\varepsilon > 1/\operatorname{poly}\log n$. In this paper, we give a randomized $(1 \pm \varepsilon)$-approximation algorithm for Knapsack in $\widetilde{O}(n^{1.5}\varepsilon^{-2})$ time (in the standard word-RAM model), achieving the first sub-quadratic dependence on n. Such sub-quadratic running time is rare in the approximate counting literature in general, as a large class of algorithms naturally faces a quadratic-time barrier.

Weiming Feng
Institute for Theoretical Studies
ETH Zürich
weiming.feng@eth-its.ethz.ch

Ce Jin
EECS
MIT
cejin@mit.edu

CP10

Lipschitz Continuous Algorithms for Covering Problems

Combinatorial algorithms are widely used for decision-making and knowledge discovery, and it is important to ensure that their output remains stable even when subjected to small perturbations in the input. Failure to do so can lead to several problems, including costly decisions, reduced user trust, potential security concerns, and lack of replicability. Unfortunately, many fundamental combinatorial algorithms are vulnerable to small input perturbations. To address the impact of input perturbations on algorithms for weighted graph problems, Kumabe and Yoshida (FOCS'23) recently introduced the concept of Lipschitz continuity of algorithms. This work explores this approach and designs Lipschitz continuous algorithms for covering problems, such as the minimum vertex cover, set cover, and feedback vertex set problems. Our algorithm for the feedback vertex set problem is based on linear programming, and in the rounding process, we develop and use a technique called cycle sparsification, which may be of independent interest.

Soh Kumabe
The University of Tokyo
sohkuma0213@gmail.com

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

CP10

Balancing Notions of Equity: Trade-Offs Between Fair Portfolio Sizes and Achievable Guarantees

Motivated by fairness concerns, we study existence and computation of portfolios defined as: given an optimization problem with feasible solutions D, a class $\mathbf{C}$ of fairness objective functions, a set $X \subseteq D$ of feasible solutions is an α-approximate portfolio if for each objective $f \in \mathbf{C}$, there is an α-approximation for f in X. We study the trade-off between the size $|X|$ of the portfolio and its approximation factor α for various combinatorial problems, such as scheduling, covering, and facility location, and choices of $\mathbf{C}$ as top-k, ordered and symmetric monotonic norms. Our results include: (i) an α-approximate portfolio of size $O\left(\frac{\log d}{\log(\alpha/4)}\right)$ for ordered norms and lower bounds of size $\Omega\left(\frac{\log d}{\log\alpha + \log\log d}\right)$ for the problem of scheduling identical jobs on d unidentical machines, (ii) $O(\log n)$-approximate $O(\log n)$-sized portfolios for facility location for symmetric monotonic norms, and (iii) $\log^{O(r^2)} d$-size $O(1)$-approximate portfolios for ordered norms and $O(\log d)$-approximate for symmetric monotonic norms for covering polyhedra with a constant r number of constraints. The latter result uses our novel OrderAndCount framework that obtains an exponential improvement in portfolio sizes compared to the current state-of-the-art, which may be of independent interest.

Swati Gupta
Massachusetts Institute of Technology
swatig@mit.edu

Jai Moondra, Mohit Singh
Georgia Institute of Technology
jmoondra3@gatech.edu, mohit.singh@isye.gatech.edu

CP10

Approximating Traveling Salesman Problems Using a Bridge Lemma

We give improved approximations for two metric Traveling Salesman Problem (TSP) variants. In Ordered TSP (OTSP) we are given a linear ordering on a subset of nodes $o_1, ..., o_k$. The TSP solution must have that o_{i+1} is visited at some point after o_i for each $1 <= i < k$. This is the special case of Precedence-Constrained TSP (PTSP) in which the precedence constraints are given by a single chain on a subset of nodes. In k-Person TSP Path (k-TSPP), we are given pairs of nodes $(s_1, t_1), ..., (s_k, t_k)$. The goal is to find an s_i-t_i path with minimum total cost such that every node is visited by at least one path. We obtain a $3/2 + e^{-1} < 1.878$ approximation for OTSP, the first improvement over a trivial alpha+1 approximation where alpha is the current best TSP approximation. We also obtain a $1 + 2 * e^{-1/2} < 2.214$ approximation for k-TSPP, the first improvement over a trivial 3-approximation. These algorithms both use an adaptation of the Bridge Lemma that was initially used to obtain improved Steiner Tree approximations [Byrka et al., 2013]. Roughly speaking, our variant states that the cost of a cheapest forest rooted at a given set of terminal nodes will decrease by a substantial amount if we randomly sample a set of non-terminal nodes to also become terminals such provided each non-terminal has a constant probability of being sampled. We believe this view of the Bridge Lemma will find further use for improved vehicle routing approximations beyond this paper.

Tobias Mömke
University of Augsburg, Germany
moemke@informatik.uni-augsburg.de

Martin Böhm
University of Wroclaw
boehm@cs.uni.wroc.pl

Zachary Friggstad
University of Alberta
zacharyf@ualberta.ca

Joachim Spoerhase

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

CP10
Min-CSPs on Complete Instances

Given a fixed arity $k \geq 2$, Min-k-CSP on complete instances involves a set of n variables V and one nontrivial constraint for every k-subset of variables (so there are $\binom{n}{k}$ constraints). The goal is to find an assignment that minimizes unsatisfied constraints. Unlike Max-k-CSP that admits a PTAS on dense or expanding instances, the approximability of Min-k-CSP is less understood. For some CSPs like Min-k-SAT, there's an approximation-preserving reduction from general to dense instances, making complete instances unique for potential new techniques. This paper initiates a study of Min-k-CSPs on complete instances. We present an $O(1)$-approximation algorithm for Min-2-SAT on complete instances, the minimization version of Max-2-SAT. Since $O(1)$-approximation on dense or expanding instances refutes the Unique Games Conjecture, it shows a strict separation between complete and dense/expanding instances. Then we study the decision versions of CSPs, aiming to satisfy all constraints; which is necessary for any nontrivial approximation. Our second main result is a quasi-polynomial time algorithm for every Boolean k-CSP on complete instances, including k-SAT. We provide additional algorithmic and hardness results for CSPs with larger alphabets, characterizing (arity, alphabet size) pairs that admit a quasi-polynomial time algorithm on complete instances.

Aditya Anand, Euiwoong Lee
University of Michigan
adanand@umich.edu, euiwoong@umich.edu

Amatya Sharma
University of Michigan, Ann Arbor
amatya@umich.edu

CP11
New Combinatorial Insights for Monotone Apportionment

The apportionment problem constitutes a fundamental problem in democratic societies: How to distribute a fixed number of seats among a set of states in proportion to the states' populations? In this paper, we connect the design of monotone apportionment methods to classic problems from discrete geometry and combinatorial optimization and explore the extent to which randomization can enhance proportionality. We first focus on the well-studied family of stationary divisor methods, which satisfy the strong population monotonicity property, and show that this family produces only a slightly superlinear number of different outputs as a function of the number of states. While our upper and lower bounds leave a small gap, we show that closing this gap would solve a long-standing open problem from discrete geometry, known as the complexity of k-levels in line arrangements. As we show that randomizing over divisor methods can only partially overcome their violation of the quota axiom, we propose a relaxed version of divisor methods in which the total number of seats may slightly deviate from the house size. By randomizing over these methods, we can simultaneously satisfy population monotonicity, quota, and ex-ante proportionality. Finally, we turn our attention to house-monotone and quota-compliant methods. We provide a polyhedral characterization based on network flows, which implies a simple description of all ex-ante proportional randomized methods satisfying these axioms.

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

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

Ulrike Schmidt-Kraepelin
TU Eindhoven
u.schmidtkraepelin@gmail.com

Alexandros Tsigonias-Dimitriadis
European Central Bank
alexandrostsigdim@gmail.com

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

CP11
Prophet Secretary and Matching: the Significance of the Largest Item

The prophet secretary problem is a combination of the prophet inequality and the secretary problem, where elements are drawn from known independent distributions and arrive in uniformly random order. In this work, we design 1) a 0.688-competitive algorithm, that breaks the 0.675 barrier of blind strategies (Correa, Saona, Ziliotto, 2021), and 2) a 0.641-competitive algorithm for the prophet secretary matching problem, that breaks the $1 - 1/e \approx 0.632$ barrier for the first time. Our second result also applies to the query-commit model of weighted stochastic matching and improves the state-of-the-art ratio (Derakhshan and Farhadi, 2023).

Ziyun Chen
University of Washington
ziyuncc@cs.washington.edu

Zhiyi Huang, Dongchen Li
The University of Hong Kong
zhiyi@cs.hku.hk, dongchen.li@connect.hku.hk

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

CP11
New Prophet Inequalities Via Poissonization and Sharding

This work introduces sharding and Poissonization as a unified framework for analyzing prophet inequalities. Sharding involves splitting a random variable into several independent random variables, shards, that collectively mimic the original variable's behavior. We combine this with Poissonization, where these shards are modeled using a Poisson distribution. Despite the simplicity of our framework, we improve the competitive ratio analysis of a dozen well studied prophet inequalities in the literature, some of which have been studied for decades. This includes the Top-1-of-k prophet inequality, prophet secretary inequality, and semi-online prophet inequality, among others. This

approach not only refines the constants but also offers a more intuitive and streamlined analysis for many prophet inequalities in the literature. Furthermore, it simplifies proofs of several known results and may be of independent interest for other variants of the prophet inequality, such as order-selection.

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

CP11

Prophet Inequalities: Competing with the Top ℓ Items Is Easy

We explore a prophet inequality problem, where the values of a sequence of items are drawn i.i.d. from some distribution, and an online decision maker must select one item irrevocably. We establish that the worst-case competitive ratio between the expected optimal performance of the online decision maker compared to that of a prophet who uses the average of the top ℓ items is exactly ℓ/c_ℓ, where c_ℓ is the solution to an integral equation. This quantity ℓ/c_ℓ is larger than $1-e^{-\ell}$. This implies that the bound converges exponentially fast to 1 as ℓ grows. In particular for $\ell = 2$, $2/c_2 \approx 0.966$ which is much closer to 1 than the classical bound of 0.745 for $\ell = 1$. Additionally, we prove asymptotic lower bounds for the competitive ratio of a more general scenario, where the decision maker is permitted to select k items. This subsumes the k multi-unit i.i.d. prophet problem and provides the current best asymptotic guarantees, as well as enables broader understanding in the more general framework. Finally, we prove a tight asymptotic competitive ratio when only static threshold policies are allowed.

Mathieu Molina, Nicolas Gast, Patrick Loiseau
Inria
mathieu.molina@inria.fr, nicolas.gast@inria.fr,
patrick.loiseau@inria.fr

Vianney Perchet
ENSAE - Criteo
vianney.perchet@normalesup.org

CP11

An Elementary Predictor Obtaining $2\sqrt{T}+1$ Distance to Calibration

Blasiok et al. [2023] proposed distance to calibration as a natural measure of calibration error that unlike expected calibration error (ECE) is continuous. Recently, Qiao and Zheng [2024] (COLT 2024) gave a non-constructive argument establishing the existence of a randomized online predictor that can obtain $O(\sqrt{T})$ distance to calibration in expectation in the adversarial setting, which is known to be impossible for ECE. They leave as an open problem finding an explicit, efficient, deterministic algorithm. We resolve this problem and give an extremely simple, efficient, deterministic algorithm that obtains distance to calibration error at most $2\sqrt{T}+1$.

Eshwar Ram Arunachaleswaran
Indian Institute of Science
eshwarram.arunachaleswaran@gmail.com

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

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

Mirah Shi
University of Pennsylvania
mirahshi@seas.upenn.edu

CP11

Designing Automated Market Makers for Combinatorial Securities: A Geometric Viewpoint

Designing automated market makers (AMMs) for prediction markets on combinatorial securities over large outcome spaces poses significant computational challenges. Prior research has primarily focused on combinatorial prediction markets within specific set systems (e.g., intervals, permutations). In this paper, we establish a comprehensive framework for designing AMMs on arbitrary set systems, by building a novel connection between the design of AMMs for combinatorial prediction markets and the range query problem in computational geometry. We present a unified framework for both analyzing the computational complexity and designing efficient AMMs. We first demonstrate the equivalence between price queries and trade updates under the popular combinatorial logarithmic market scoring rule market and the range query and range update problem. Building on this equivalence, we construct sublinear time algorithms when the VC dimension of the set system is bounded and show the non-existence of such algorithms for unbounded VC dimension cases. We then extend this approach to AMMs for combinatorial prediction markets with quadratic and power scoring rules. Finally, we show that the multi-resolution market design can be naturally integrated into the partition-tree scheme. Additionally, we introduce the combinatorial swap operation problem for automated market makers in decentralized finance and show that it can be efficiently reduced to range update problems.

Prommy Sultana Hossain
George Mason University
phossai@gmu.edu

Xintong Wang
Rutgers University
xintong.wang@rutgers.edu

Fang-Yi Yu
University of Michigan
fangyiyu@gmu.edu

CP12

Unbreakable Decomposition in Close-to-Linear Time

Unbreakable decomposition, introduced by [Cygan-Lokshtanov-Pilipczuk-Pilipczuk-Saurabh, SICOMP 2019; Cygan-Komosa-Lokshtanov-Pilipczuk-Pilipczuk-Saurabh-Wahlstrm, TALG 2020], has proven to be one of the most powerful tools for parameterized graph cut problems in recent years. Unfortunately, all known constructions require at least $\Omega_k\left(mn^2\right)$ time, given an undirected graph with n vertices, m edges, and cut-size parameter k. In this work, we show the first close-to-linear-time parameterized algorithm that computes an unbreakable decomposition. More precisely, for an undirected graph G and any $0 < \epsilon \leq 1$, our algorithm runs in time $2^{O(\frac{k}{\epsilon}\log\frac{k}{\epsilon})}m^{1+\epsilon}$ and

computes an $(O(k/\epsilon), k)$-unbreakable decomposition of G, where each bag has adhesion at most $O(k/\epsilon)$. This immediately opens up possibilities for obtaining close-to-linear time algorithms for numerous problems whose only known solution is based on unbreakable decomposition.

Aditya Anand, Euiwoong Lee
University of Michigan
adanand@umich.edu, euiwoong@umich.edu

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

Yaowei Long
Tsinghua University
yaoweil@umich.edu

Thatchaphol Saranurak
University of Michigan
thsa@umich.edu

CP12

Packing Short Cycles

Cycle packing is a fundamental problem in optimization, graph theory, and algorithms. Motivated by recent advancements in finding vertex-disjoint paths between a specified set of vertices that either minimize the total length of the paths [Bjrklund, Husfeldt, ICALP 2014; Mari, Mukherjee, Pilipczuk, and Sankowski, SODA 2024] or request the paths to be shortest [Lochet, SODA 2021], we consider the following cycle packing problems: Min-Sum Cycle Packing and Shortest Cycle Packing. In Min-Sum Cycle Packing, we try to find, in a weighted undirected graph, k vertex-disjoint cycles of minimum total weight. Our first main result is an algorithm that, for any fixed k, solves the problem in polynomial time. We complement this result by establishing the W[1]-hardness of Min-Sum Cycle Packing parameterized by k. The same results hold for the version of the problem where the task is to find k edge disjoint cycles. Our second main result concerns Shortest Cycle Packing, which is a special case of Shortest Cycle Packing that asks to find a packing of k shortest cycles in a graph. We prove this problem to be fixed-parameter tractable (FPT) when parameterized by k on weighted planar graphs. We also obtain a polynomial kernel for the edge-disjoint variant of the problem on planar graphs. Deciding whether Min-Sum Cycle Packing is FPT on planar graphs and whether Shortest Cycle Packing is FPT on general graphs remain challenging open questions.

Matthias Bentert
University of Bergen
matthias.bentert@uib.no

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

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

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

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

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

M. S. Ramanujan
University of Warwick
r.maadapuzhi-sridharan@warwick.ac.uk

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

Kirill Simonov
TU Wien
kirillsimonov@gmail.com

CP12

The Primal Pathwidth Seth

Motivated by the importance of dynamic programming (DP) in parameterized complexity, we consider fine-grained questions, such as the following: (i) can Dominating Set be solved in time $(3-\epsilon)^{pw}n^{O(1)}$? (where pw is the pathwidth) (ii) can Coloring be solved in time $pw^{(1-\epsilon)pw}n^{O(1)}$? (iii) can a short reconfiguration between two size-k independent sets be found in time $n^{(1-\epsilon)k}$? Such questions are well-studied: for some the answer is No under the SETH, while for others coarse-grained lower bounds are known under the ETH. Even though questions such as the above seem "morally equivalent" as they all ask if a simple DP can be improved, the problems concerned have varying complexities, ranging from single-exponential FPT to XNLP-complete. This paper's main contribution is to show that, despite their varying complexities, these questions are not just morally equivalent, but in fact they are the same question in disguise. We achieve this by putting forth a natural complexity assumption which we call the Primal Pathwidth-Strong Exponential Time Hypothesis (pw-SETH) and which states that 3-SAT cannot be solved in time $(2-\epsilon)^{pw}n^{O(1)}$, for any $\epsilon > 0$, where pw is the pathwidth of the primal graph of the input CNF formula. We then show that numerous fine-grained questions in parameterized complexity, including the ones above, are equivalent to the pw-SETH, and hence to each other.

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

CP12

Fixed-Parameter Tractability of Hedge Cut

In the Hedge Cut problem, the edges of a graph are partitioned into groups called hedges, and the question is what is the minimum number of hedges to delete to disconnect the graph. Ghaffari, Karger, and Panigrahi [SODA 2017] showed that Hedge Cut can be solved in quasipolynomial-time, raising the hope for a polynomial time algorithm. Jaffke, Lima, Masarik, Pilipczuk, and Souza [SODA 2023] complemented this result by showing that assuming the Exponential Time Hypothesis (ETH), no polynomial-time algorithm exists. In this paper, we show that Hedge Cut is fixed-parameter tractable param-

eterized by the solution size ℓ by providing an algorithm with running time $\binom{O(\log n)+\ell}{\ell} \cdot m^{O(1)}$, which can be upper bounded by $c^\ell \cdot (n+m)^{O(1)}$ for any constant $c > 1$. This running time captures at the same time the fact that the problem is quasipolynomial-time solvable, and that it is fixed-parameter tractable parameterized by ℓ. We further generalize this algorithm to an algorithm with running time $\binom{O(k \log n)+\ell}{\ell} \cdot n^{O(k)} \cdot m^{O(1)}$ for Hedge k-Cut.

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

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

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

Daniel Lokshtanov
UCSB
daniello@ucsb.edu

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

CP12

Parameterized Approximation for Capacitated d-Hitting Set with Hard Capacities

In the Capacitated d-Hitting Set problem input is a universe U equipped with a capacity function $\mathsf{cap} : U \to \mathbb{N}$, and a collection $\mathcal{A}$ of subsets of U, each of size at most d. The task is to find a minimum size subset S of U and an assignment $\phi : \mathcal{A} \to S$ such that, for every set $A \in \mathcal{A}$ we have $\phi(A) \in A$ and for every $x \in U$ we have $|\phi^{-1}(x)| \leq \mathsf{cap}(x)$. In Weighted Capacitated d-Hitting Set each element of U has a positive integer weight and the goal is to find a minimum weight solution. In this paper we initiate the study of parameterized (approximation) algorithms for Capacitated d-Hitting Set. Capacitated d-Hitting Set is a well studied problem and is known to admit a d-approximation algorithm and no $(d-\epsilon)$-approximation under UGC for any $\epsilon > 0$. Further, unweighted Capacitated d-Hitting Set for $d \geq 3$ is W[1]-hard parameterized by solution size. Our main result is a parameterized approximation algorithm that runs in time $\left(\frac{k}{\epsilon}\right)^k 2^{k^{O(kd)}} (|U| + |\mathcal{A}|)^{O(1)}$ and either concludes that there is no solution of size at most k or outputs a solution S of size at most $4/3 \cdot k$ and weight at most $2+\epsilon$ times the minimum weight of a solution whose size is at most k. We also complement our algorithmic results with hardness results.

Vaishali Surianarayanan
UC Santa Barbara
vaishalisurianarayanan@gmail.com

Daniel Lokshtanov
UCSB
daniello@ucsb.edu

Abhishek Sahu
National Institute of Science Education and Research, HBNI
asahuiitkgp@gmail.com

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

Jie Xue
New York University Shanghai
jiexue@nyu.edu

CP12

Crossing Number in Slightly Superexponential Time

In the Crossing Number problem, the input consists of a graph G and integer k. The task is to determine whether there exists a drawing of G with crossing number at most k, and to output such a drawing if it exists. Grohe [STOC 2001, JCSS 2004] gave an algorithm for Crossing Number with running time $f(k)n^2$ where $f(k) = 2^{2^{2^{2^{?(k)}}}}$. He conjectured that there exists an algorithm with running time $2^{O(k)}n$. Kawarabayashi and Reed [STOC 2007] outlined an algorithm with running time $f(k)n$ where $f(k) = 2^{2^{2^{?(k)}}}$. Combining the main combinatorial lemma by Kawarabayashi and Reed with the recent algorithm for Crossing Number parameterized treewidth plus k by de Verdiere and Magnard [ESA 2021] would yield a running time of $f(k)n$ where $f(k) = 2^{O(k^4 \log k)}$. This still falls far away from the dependency on k in the conjecture by Grohe. Furthermore, critical details of the proof of the correctness of the algorithm of Kawarabayashi and Reed, and, in particular, of the aforementioned combinatorial lemma, have never been published. In this work, we give an algorithm with running time $2^{O(k \log k)}n$. Thus, our algorithm resolves Grohe's 23-year old conjecture up to a logarithmic factor in k in the exponent.

Jie Xue
New York University Shanghai
jiexue@nyu.edu

CP13

Engineering Optimal Parallel Task Scheduling

Matthew Akram
Karlsruhe Institute of Technology
mazfh85246@gmail.com

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

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

Dominik Schreiber
Karlsruhe Institute of Technology
mail@dominikschrieber.de

CP13

Another L Makes It Better? Lagrange Meets LLL and May Improve BKZ Pre-Processing

We present a new variant of the LLL lattice reduction algorithm, inspired by Lagrange notion of pair-wise reduction,

called L4. Similar to LLL, our algorithm is polynomial in the dimension of the input lattice, as well as in $\log M$, where M is an upper-bound on the norm of the longest vector of the input basis. We experimentally compared the norm of the first basis vector obtained with LLL and L4 up to dimension 200. On average we obtain vectors that are up to 16% shorter. We also used our algorithm as a pre-processing step for the BKZ lattice reduction algorithm with blocksize 24. In practice, up to dimension 140, this allows us to reduce the norm of the shortest basis vector on average by 3%, while the runtime does not significantly increases. In 10% of our tests, the whole process was even faster.

Sebastien Balny, Claire Delaplace, Gilles Dequen
MIS Laboratory, Université de Picardie Jules Verne, Amiens,
sebastien.balny@u-picardie.fr, claire.delaplace@u-picardie.fr, gilles.dequen@u-picardie.fr

CP13

Constructions, Bounds, and Algorithms for Peaceable Queens

The Peaceable Queens Problem asks to determine the maximum number $a(n)$ such that there is a placement of $a(n)$ white queens and $a(n)$ black queens on an $n \times n$ chessboard so that no queen can capture any queen of the opposite color. In this paper, we consider the peaceable queens problem and its variant on the toroidal board. For the regular board, we show that $a(n) \leq 0.1716n^2$, for all sufficiently large n. This improves on the bound $a(n) \leq 0.25n^2$ of [van Bommel and MacEachern, Armies of chess queens, Math Intelligencer,40(2):1015, 2018]. For the toroidal board, we provide new upper and lower bounds. Somewhat surprisingly, our bounds show that there is a sharp contrast in behavior between the odd torus and the even torus. Our lower bounds are given by explicit constructions. For the upper bounds, we formulate the problem as a non-linear optimization problem with at most 100 variables, regardless of the size of the board. We solve our non-linear program exactly using modern optimization software. We also provide a local search algorithm and a software implementation which converges very rapidly to solutions which appear optimal. Our algorithm is sufficiently robust that it works on both the regular and toroidal boards. For example, for the regular board, the algorithm quickly finds the so-called Ainley construction. Thus, our work provides some further evidence that the Ainley construction is indeed optimal.

Katie Clinch
SCHOOL OF COMPUTER SCIENCE AND ENGINEERING
UNSW
k.clinch@unsw.edu.au

Matthew Drescher
Department of Electrical and Computer Engineering
University of California, Davis
knavely@gmail.com

Tony Huynh
Université Libre de Bruxelles
tony.bourbaki@gmail.com

Abdallah Saffidine
SCHOOL OF COMPUTER SCIENCE AND ENGINEERING
UNSW
abdallah.saffidine@gmail.com

CP13

A Greedy Algorithm for Low-Crossing Partitions for General Set Systems

Simplicial partitions are a fundamental structure in computational geometry, as they form the basis of optimal data structures for range searching and several related problems. Current algorithms are built on very specific spatial partitioning tools tailored for certain geometric cases. This severely limits their applicability to general set systems. In this work, we propose a simple greedy heuristic for constructing simplicial partitions of any set system. We present a thorough empirical evaluation of its behavior on a variety of geometric and non-geometric set systems, showing that it performs well on most instances.

Monika Csikos
Université Paris Cité
csikos@irif.fr

Alexandre Louvet, Nabil Mustafa
Université Sorbonne Paris Nord
alexandre.louvet@mailo.fr, nabil.mustafa@univ-paris13.fr

CP13

HyperSteiner: Computing Heuristic Hyperbolic Steiner Minimal Trees

We propose HyperSteiner – an efficient heuristic algorithm for computing Steiner minimal trees in the hyperbolic space. HyperSteiner extends the Euclidean Smith-Lee-Liebman algorithm, which is grounded in a divide-and-conquer approach involving the Delaunay triangulation. The central idea is rephrasing Steiner tree problems with three terminals as a system of equations in the Klein-Beltrami model. Motivated by the fact that hyperbolic geometry is well-suited for representing hierarchies, we explore applications to hierarchy discovery in data. Results show that HyperSteiner infers more realistic hierarchies than the Minimum Spanning Tree and is more scalable to large datasets than Neighbor Joining.

Aniss A. Medbouhi
KTH Royal Institute of Technology
medbouhi@kth.se

Alejandro García-Castellanos
University of Amsterdam
a.garciacastellanos@uva.nl

Giovanni Luca Marchetti, Danica Kragic
KTH Royal Institute of Technology
glma@kth.se, dani@kth.se

Erik Johannes Bekkers
University of Amsterdam
e.j.bekkers@uva.nl

CP14

Parallel and Distributed Expander Decomposition: Simple, Fast, and Near-Optimal

Expander decompositions have become one of the central frameworks in the design of fast algorithms. For an undi-

rected graph $G = (V, E)$, a near-optimal ϕ-expander decomposition is a partition $V_1, V_2, \ldots, V_k$ of the vertex set V where each subgraph $G[V_i]$ is a ϕ-expander, and only an $\widetilde{O}(\phi)$-fraction of the edges cross between partition sets. In this article, we give the first near-optimal *parallel* algorithm to compute ϕ-expander decompositions in near-linear work $\widetilde{O}(m/\phi^2)$ and near-constant span $\widetilde{O}(1/\phi^4)$. Our algorithm is very simple and likely practical. Our algorithm can also be implemented in the distributed Congest model in $\tilde{O}(1/\phi^4)$ rounds. Our results surpass the theoretical guarantees of the current state-of-the-art parallel algorithms [Chang-Saranurak PODC'19, Chang-Saranurak FOCS'20], while being the first to ensure that only an $\widetilde{O}(\phi)$ fraction of edges cross between partition sets. In contrast, previous algorithms [Chang-Saranurak PODC'19, Chang-Saranurak FOCS'20] admit at least an $O(\phi^{1/3})$ fraction of crossing edges, a polynomial loss in quality inherent to their random-walk-based techniques. Our algorithm, instead, leverages flow-based techniques and extends the popular sequential algorithm presented in [Saranurak-Wang SODA'19].

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

Thatchaphol Saranurak
University of Michigan
thsa@umich.edu

CP14

Quasilinear-Time Eccentricities Computation, and More, on Median Graphs

Computing the diameter, and more generally, all eccentricities of an undirected graph is an important problem in algorithmic graph theory and the challenge is to identify graph classes for which their computation can be achieved in subquadratic time. Using a new recursive scheme based on the structural properties of median graphs, we provide a quasilinear-time algorithm to determine all eccentricities for this well-known family of graphs. The gist of our technique is to identify the balanced and unbalanced parts of the Θ-class decomposition of median graphs, which are then processed using different recursive schemes. The exact running time of our algorithm is in $O(n \log^4 n)$. This outcome not only answers a question asked by Bnteau et al. (2020) but also greatly improves the recent combinatorial algorithm of Berg et al. (2022) for the same problem, running in time $O(n^{1.6408} \log^{O(1)} n)$. As our second main contribution, we further propose a distance oracle for median graphs with both poly-logarithmic size and query time. More specifically, we present a combinatorial algorithm which computes for any median graph G, in quasilinear time $O(n \log^4(n))$, vertex-labels of size $O(\log^3(n))$ such that any distance of G can be retrieved in time $O(\log^4(n))$ thanks to these labels.

Pierre Bergé
Université Clermont Auvergne, France
pierre.berge@uca.fr

Ducoffe Guillaume
University of Bucharest, Romania
guillaume.ducoffe@ici.ro

Habib Michel
IRIF, Université Paris Cité, France
habib@irif.fr

CP14

A Cut-Matching Game for Constant-Hop Expanders

This paper extends and generalizes the well-known cut-matching game framework and provides a novel cut-strategy that produces constant-hop expanders. Constant-hop expanders are a significant strengthening of regular expanders with the additional guarantee that any demand can be (obliviously) routed along constant-hop flow-paths - in contrast to the $\Omega(\log n)$-hop paths in expanders. Cut-matching games for expanders are key tools for obtaining linear-time approximation algorithms for many hard problems, including finding (balanced or approximately-largest) sparse cuts, certifying the expansion of a graph by embedding an (explicit) expander, as well as computing expander decompositions, hierarchical cut decompositions, oblivious routings, multi-cuts, and multi-commodity flows. The cut-matching game of this paper is crucial in extending this versatile and powerful machinery to constant-hop and length-constrained expanders and has been already been extensively used. For example, as a key ingredient in several recent breakthroughs, including, computing constant-approximate k-commodity (min-cost) flows in $(m + k)^{1+\epsilon}$ time as well as the optimal constant-approximate deterministic worst-case fully-dynamic APSP-distance oracle - in all applications the constant-approximation factor directly traces to and crucially relies on the expanders from a cut-matching game guaranteeing constant-hop routing paths.

Bernhard Haeupler
Sofia University "St. Kliment Ohridski"
ETH Zurich
bernhard.haeupler@inf.ethz.ch

Jonas Huebotter
ETH Zurich
jonas.huebotter@inf.ethz.ch

Mohsen Ghaffari
MIT
ghaffari@mit.edu

CP14

Deterministic Online Bipartite Edge Coloring

We study online bipartite edge coloring, with nodes on one side of the graph revealed sequentially. The trivial greedy algorithm is $(2 - o(1))$-competitive, which is optimal for graphs of low maximum degree, $\Delta = O(\log n)$ [BNMN IPL'92]. Numerous online edge-coloring algorithms outperforming the greedy algorithm in various settings were designed over the years (e.g., [AGKM FOCS'03, BMM SODA'10, CPW FOCS'19, BGW SODA'21, KLSST STOC'22, BSVW STOC'24]), all crucially relying on randomization. A commonly-held belief, first stated by [BNMN IPL'92], is that randomization is necessary to outperform greedy. We refute this belief. We present a deterministic algorithm that beats greedy for sufficiently large $\Delta = \Omega(\log n)$, and in particular has competitive ratio $\frac{e}{e-1} + o(1)$ for all $\Delta = \omega(\log n)$. We obtain our result via a new simple randomized algorithm that works against adaptive adversaries (as opposed to oblivious adversaries

assumed by prior work). This implies the existence of a similarly-competitive deterministic algorithm [BDBKTW STOC'90]. A key ingredient in our algorithm (and the reason for its competitive ratio) is the use of contention resolution schemes of [FV FOCS'06]. This is the first use of contention resolution schemes, which are randomized algorithms for randomized inputs, that yields a deterministic algorithm for deterministic settings.

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

Ola Svensson, Radu Vintan
EPFL
ola.svensson@epfl.ch, radu.vintan@epfl.ch

David Wajc
Technion
david.wajc@gmail.com

CP14

Eulerian Graph Sparsification by Effective Resistance Decomposition

We provide an algorithm that, given an n-vertex m-edge Eulerian graph with polynomially bounded weights, computes an $\breve{O}(n\log^2 n\cdot\varepsilon^{-2})$-edge ε-approximate Eulerian sparsifier with high probability in $\breve{O}(m\log^3 n)$ time (where $\breve{O}(\cdot)$ hides polyloglog(n) factors). Due to a reduction from [Peng-Song, STOC '22], this yields an $\breve{O}(m\log^3 n + n\log^6 n)$-time algorithm for solving n-vertex m-edge Eulerian Laplacian systems with polynomially-bounded weights with high probability, improving upon the previous state-of-the-art runtime of $\Omega(m\log^8 n + n\log^{23} n)$. We also give a polynomial-time algorithm that computes $O(\min(n\log n\cdot\varepsilon^{-2}+n\log^{5/3} n\cdot\varepsilon^{-4/3}, n\log^{3/2} n\cdot\varepsilon^{-2}))$-edge sparsifiers, improving the best such sparsity bound of $O(n\log^2 n\cdot\varepsilon^{-2} + n\log^{8/3} n\cdot\varepsilon^{-4/3})$ [Sachdeva-Thudi-Zhao, ICALP '24]. In contrast to prior Eulerian graph sparsification algorithms which used either short cycle or expander decompositions, our algorithms use a simple efficient effective resistance decomposition scheme we introduce. Our algorithms apply a natural sampling scheme and electrical routing (to achieve degree balance) to such decompositions. Our analysis leverages new asymmetric variance bounds specialized to Eulerian Laplacians and tools from discrepancy theory.

Arun Jambulapati
University of Washington
jmblpati@gmail.com

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

Aaron Sidford
Stanford University
sidford@stanford.edu

Kevin Tian
Microsoft Rsearch
kjtian@cs.utexas.edu

Yibin Zhao
University of Toronto
ybzhao@cs.toronto.edu

CP15

A Reduction from Multi-Parameter to Single-Parameter Bayesian Contract Design

The main result of this paper is an almost approximation-preserving polynomial-time reduction from the most general multi-parameter Bayesian contract design (BCD) to single-parameter BCD. That is, for any multi-parameter BCD instance I^M, we construct a single-parameter instance I^S such that any β-approximate contract (resp. menu of contracts) of I^S can in turn be converted to a $(\beta-\epsilon)$-approximate contract (resp. menu of contracts) of I^M. The reduction is in time polynomial in the input size and $\log(\frac{1}{\epsilon})$; moreover, when $\beta = 1$ (i.e., the given single-parameter solution is exactly optimal), the dependence on $\frac{1}{\epsilon}$ can be removed, leading to a polynomial-time exact reduction. This efficient reduction is somewhat surprising because in the closely related problem of Bayesian mechanism design, a polynomial-time reduction from multi-parameter to single-parameter setting is believed to not exist. Our result demonstrates the intrinsic difficulty of addressing moral hazard in Bayesian contract design, regardless of being single-parameter or multi-parameter.

Matteo Castiglioni
Politecnico di Milano
matteo.castiglioni@polimi.it

Junjie Chen, Minming Li
City University of Hong Kong
junjchen9-c@my.cityu.edu.hk, minming.li@cityu.edu.hk

Haifeng Xu
The University of Chicago
haifengxu@uchicago.edu

Song Zuo
Google Research
szuo@google.com

CP15

Hiring for An Uncertain Task: Joint Design of Information and Contracts

We initiate the computational problem of jointly designing information and contracts. We consider three possible classes of contracts with decreasing flexibility and increasing simplicity: ambiguous contracts, menus of explicit contracts and explicit single contract. Ambiguous contracts allow the principal to conceal the applied payment schemes through a contract that depends on the unknown state of nature, while explicit contracts reveal the contract prior to the agents decision. Our results show a trade-off between the simplicity of the contracts and the computational complexity of the joint design. Indeed, we show that an approximately-optimal mechanism with ambiguous contracts can be computed in polynomial time. However, they are convoluted mechanisms and not well-suited for some real-world scenarios. Conversely, explicit menus of contracts and single contracts are simpler mechanisms, but they cannot be computed efficiently. In particular, we show that computing the optimal mechanism with explicit menus of contracts and single contracts is APX-Hard. We also characterize the structure of optimal mechanisms. Interestingly, direct mechanisms are optimal for both the most flexible ambiguous contracts and the least flexible ex-

plicit single contract, but they are suboptimal for that with menus of contracts. Finally, motivated by our hardness results, we turn our attention to menus of linear contracts and single linear contracts. We show that both problems admit an FPTAS.

Matteo Castiglioni
Politecnico di Milano
matteo.castiglioni@polimi.it

Junjie Chen
City University of Hong Kong
junjchen9-c@my.cityu.edu.hk

CP15

A Multi-Dimensional Online Contention Resolution Scheme for Revenue Maximization

We study multi-buyer multi-item sequential item pricing mechanisms for revenue maximization with the goal of approximating a natural fractional relaxation – the ex ante optimal revenue. We assume that buyers' values are subadditive but make no assumptions on the value distributions. While the optimal revenue, and therefore also the ex ante benchmark, is inapproximable by any simple mechanism in this context, previous work has shown that a weaker benchmark that optimizes over so-called "buy-many" mechanisms can be approximated. Approximations are known, in particular, for settings with either a single buyer or many unit-demand buyers. We extend these results to the much broader setting of many subadditive buyers. We show that the ex ante buy-many revenue can be approximated via sequential item pricings to within an $O(\log^2 m)$ factor, where m is the number of items; a logarithmic dependence on m is also necessary. Our approximation is achieved through the construction of a new multi-dimensional Online Contention Resolution Scheme (OCRS), that provides an online rounding of the optimal ex ante solution. Chawla et. al [2023] previously constructed an OCRS for revenue for unit-demand buyers, but their construction relied heavily on the "almost single dimensional" nature of unit-demand values. Prior to that work, OCRSes have only been studied in the context of social welfare maximization for single-parameter buyers.

Trung Dang
University of Texas at Austin
dddtrung@cs.utexas.edu

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

Dimitrios Christou, Zhiyi Huang
University of Texas at Austin
dimitrios.christou@hotmail.com, zhiyih@cs.utexas.edu

Gregory Kehne
Washington University in St. Louis
University of Texas at Austin
kehne@wustl.edu

Rojin Rezvan
University of Texas at Austin
rojinrezvan@utexas.edu

CP15

Multi-Agent Combinatorial Contracts

Combinatorial contracts are emerging as a key paradigm in algorithmic contract design, paralleling the role of combinatorial auctions in algorithmic mechanism design. In this paper we study natural combinatorial contract settings involving teams of agents, each capable of performing multiple actions. This scenario extends two fundamental special cases: the single-agent combinatorial action model of [Duetting et al. 2021], and the multi-agent binary-action model of [Babaioff et al. 2012, Duetting et al. 2023]. This setting presents fundamentally different challenges compared to the previous special cases, as it lacks key properties that have been crucial for resolving these scenarios. To navigate these challenges, we develop a broad set of novel tools that allow us to establish approximation guarantees for this setting. Our main result is a constant-factor approximation for multi-agent multi-action problems with submodular rewards, given access to value and demand oracles. This result is tight: we show that this problem admits no PTAS (even under binary actions). As a byproduct of our main result, we devise an FPTAS, given value and demand oracles, for single-agent combinatorial action scenarios with general reward functions, which is of independent interest. Finally, we show that for subadditive rewards, perhaps surprisingly, the gap between the optimal welfare and the principal's utility scales logarithmically (rather than linearly) with the size of the action space.

Paul Duetting
Google Research
duetting@google.com

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

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

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

CP15

Majorized Bayesian Persuasion and Fair Selection

We address the fundamental problem of selection under uncertainty by modeling it from the perspective of Bayesian persuasion. In our model, a decision maker with imperfect information always selects the option with the highest expected value. We seek to achieve fairness among the options by revealing additional information to the decision maker and hence influencing its subsequent selection. To measure fairness, we adopt the notion of majorization, aiming at simultaneously approximately maximizing all symmetric, monotone, concave functions over the utilities of the options. As our main result, we design a novel information revelation policy that achieves a logarithmic-approximation to majorization in polynomial time. On the other hand, no policy, regardless of its running time, can achieve a constant-approximation to ma-

jorization. Our work is the first non-trivial majorization result in the Bayesian persuasion literature with multi-dimensional information sets.

Siddhartha Banerjee
Cornell University
sbanerjee@cornell.edu

Kamesh Munagala
Duke University, USA
kamesh@cs.duke.edu

Yiheng Shen
Duke University
yiheng.shen@duke.edu

Kangning Wang
Rutgers University
kn.w@rutgers.edu

CP16

Linear Equations with Monomial Constraints and Decision Problems in Abelian-by-Cyclic Groups

We show that it is undecidable whether a system of linear equations over the Laurent polynomial ring $\mathbb{Z}[X^{\pm}]$ admit solutions where a specified subset of variables take value in the set of monomials $\{X^z \mid z \in \mathbb{Z}\}$. In particular, we construct a finitely presented $\mathbb{Z}[X^{\pm}]$-module, where it is undecidable whether a linear equation $X^{z_1}\boldsymbol{f}_1+\cdots+X^{z_n}\boldsymbol{f}_n = \boldsymbol{f}_0$ has solutions $z_1, \ldots, z_n \in \mathbb{Z}$. This contrasts the decidability of the case $n = 1$, which can be deduced from Noskov's Lemma. We apply this result to settle a number of problems in computational group theory. We show that it is undecidable whether a system of equations has solutions in the wreath product $\mathbb{Z} \wr \mathbb{Z}$, providing a negative answer to an open problem of Kharlampovich, López and Myasnikov (2020). We show that there exists a finitely generated abelian-by-cyclic group in which the problem of solving a single (spherical) quadratic equation is undecidable, answering an open problem of Lysenok and Ushakov (2021). We also construct a finitely generated abelian-by-cyclic group, different to that of Mishchenko and Treier (2017), in which the Knapsack Problem is undecidable. In contrast, we show that the problem of Coset Intersection is decidable in all finitely generated abelian-by-cyclic groups.

Ruiwen Dong
Saarland University
ruiwen.dong@magd.ox.ac.uk

CP16

An Efficient Uniqueness Theorem for Overcomplete Tensor Decomposition

We give a new, constructive uniqueness theorem for tensor decomposition. It applies to order 3 tensors of format $n \times n \times p$ and can prove uniqueness of decomposition for generic tensors up to rank $r = 4n/3$ as soon as $p \geq 4$. One major advantage over Kruskal's uniqueness theorem is that our theorem has an algorithmic proof, and the resulting algorithm is efficient. Like the uniqueness theorem, it applies in the range $n \leq r \leq 4n/3$. As a result, we obtain the first efficient algorithm for overcomplete decomposition of generic tensors of order 3. For instance, prior to this work it was not known how to efficiently decompose generic tensors of format $n \times n \times n$ and rank $r = 1.01n$ (or rank $r \leq (1+\epsilon)n$, for some constant $\epsilon > 0$). Efficient overcomplete decomposition of generic tensors of format $n \times n \times 3$ remains an open problem. Our results are based on the method of commuting extensions pioneered by Strassen for the proof of his $3n/2$ lower bound on tensor rank and border rank. In particular, we rely on an algorithm for the computation of commuting extensions recently proposed in a companion paper, and on the classical diagonalization-based "Jennrich algorithm" for undercomplete tensor decomposition.

Pascal Koiran
LIP, ENS Lyon
pascal.koiran@ens-lyon.fr

CP16

Faster Linear Systems and Matrix Norm Approximation Via Multi-Level Sketched Preconditioning

We present a new class of preconditioned iterative methods for solving linear systems of the form $Ax = b$. Our methods construct a low-rank preconditioner using sparse matrix sketching, which itself is inverted using additional levels of sketching and preconditioning. We prove that this approach leads to convergence bounds that depend on a natural average condition number of A, yielding faster runtimes for a number of fundamental problems: 1. We show how to solve any $n \times n$ linear system that is well-conditioned except for k outlying large singular values in $\tilde{O}(n^{2.065} + k^{\omega})$ time, improving on a result of [Derezinski, Yang, STOC 2024] for all $kn^{0.78}$. 2. We give the first $\tilde{O}(n^2 + d_{\lambda}{}^{\omega})$ time algorithm for solving a regularized linear system $(A + \lambda I)x = b$, where A is positive semidefinite with effective dimension $d_{\lambda} = tr(A(A + \lambda I)^{-1})$. 3. We give faster algorithms for approximating a variety of matrix norms. For example, for the Schatten 1-norm (nuclear norm), we give an algorithm that runs in $\tilde{O}(n^{2.11})$ time, improving on an $\tilde{O}(n^{2.18})$ method of [Musco et al., ITCS 2018]. Interestingly, previous state-of-the-art algorithms for the problems above relied on stochastic iterative methods, like stochastic coordinate and gradient descent. By leveraging tools from matrix sketching, our work takes a completely different approach.

Michal Derezinski
University of Michigan
derezin@umich.edu

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

Jiaming Yang
University of Michigan
jiamyang@umich.edu

CP16

Improving the Leading Constant of Matrix Multiplication

Algebraic matrix multiplication algorithms are designed by bounding the rank of matrix multiplication tensors, and then using a recursive method. However, designing algorithms in this way quickly leads to large constant factors: if one proves that the tensor for multiplying $n \times n$ matrices has rank $\leq t$, then the resulting recurrence shows that $M \times M$ matrices can be multiplied using $O(n^2 \cdot M^{\log_n t})$ operations, where the leading constant scales proportionally to n^2. Even modest increases in n can blow up the leading constant too much to be worth the slight decrease

in the exponent of M. Meanwhile, the asymptotically best algorithms use very large n, such that n^2 is larger than the number of atoms in the visible universe! In this paper, we give new ways to use tensor rank bounds to design matrix multiplication algorithms, which lead to smaller leading constants than the standard recursive method. Our main result shows that, if the tensor for multiplying $n \times n$ matrices has rank $\leq t$, then $M \times M$ matrices can be multiplied using only $n^{O(1/(\log n)^{0.33})} \cdot M^{\log_n t}$ operations. In other words, we improve the leading constant in general from $O(n^2)$ to $n^{O(1/(\log n)^{0.33})} < n^{o(1)}$.

Hantao Yu, Josh Alman
Columbia University
hy2751@columbia.edu, josh@cs.columbia.edu

CP16

More Asymmetry Yields Faster Matrix Multiplication

We present a new improvement on the laser method for designing fast matrix multiplication algorithms. The new method further develops the recent advances by [Duan, Wu, Zhou FOCS 2023] and [Vassilevska Williams, Xu, Xu, Zhou SODA 2024]. Surprisingly the new improvement is achieved by incorporating more asymmetry in the analysis, circumventing a fundamental tool of prior work that requires two of the three dimensions to be treated identically. The method yields a new bound on the square matrix multiplication exponent $\omega < 2.371339$, improved from the previous bound of $\omega < 2.371552$. We also improve the bounds of the exponents for multiplying rectangular matrices of various shapes.

Josh Alman
Columbia University
josh@cs.columbia.edu

Ran Duan
Tsinghua University, Beijing, China
duanran@mail.tsinghua.edu.cn

Virginia Vassilevska Williams
MIT
virgi@mit.edu

Yinzhan Xu
Massachusetts Institute of Technology
xyzhan@mit.edu

Zixuan Xu
MIT
zixuanxu@mit.edu

Renfei Zhou
Tsinghua University
renfeiz@andrew.cmu.edu

CP17

Quasi-Monte Carlo Beyond Hardy-Krause

The classical approaches to numerically integrating a function f are Monte Carlo (MC) and quasi-Monte Carlo (QMC) methods. MC methods use random samples to evaluate f and have error $O(\sigma(f)/\sqrt{n})$, where $\sigma(f)$ is the standard deviation of f. QMC methods are based on evaluating f at explicit low-discrepancy points, and as given by the classical Koksma-Hlawka inequality, they have error $\widetilde{O}(\sigma_{\mathsf{HK}}(f)/n)$, where $\sigma_{\mathsf{HK}}(f)$ is the variation of f in the sense of Hardy and Krause. These two methods have distinctive advantages and shortcomings, and a fundamental question is to find a method that combines the advantages of both. We give a simple randomized algorithm that produces QMC point sets with the following desirable features: (1) It achieves substantially better error than the classical Koksma-Hlawka inequality. In particular, it has error $\widetilde{O}(\sigma_{\mathsf{SO}}(f)/n)$, where $\sigma_{\mathsf{SO}}(f)$ is a new measure of variation that we introduce, which is substantially smaller than the Hardy-Krause variation. (2) The algorithm only requires random samples from the underlying distribution, which makes it as flexible as MC. (3) It automatically achieves the best of both MC and QMC (and the above improvement over Hardy-Krause variation and Koksma-Hlawka inequality) in an optimal way. (4) The algorithm is extremely efficient, with an amortized $\widetilde{O}(1)$ runtime per sample.

Nikhil Bansal
University of Michigan
bansal@gmail.com

Haotian Jiang
University of Chicago
jhtdavid96@gmail.com

CP17

Improved List Size for Folded Reed-Solomon Codes

Folded Reed-Solomon (FRS) codes are variants of Reed-Solomon codes, known for their optimal list decoding radius. We show explicit FRS codes with rate R that can be list decoded up to radius $1-R-\epsilon$ with lists of size $\mathcal{O}(1/\epsilon^2)$. This improves the best known list size among explicit list decoding capacity achieving codes. We also show a more general result that for any $k \geq 1$, there are explicit FRS codes with rate R and distance $1-R$ that can be list decoded arbitrarily close to radius $\frac{k}{k+1}(1-R)$ with lists of size $(k-1)^2+1$. Our results are based on a new and simple combinatorial viewpoint of the intersections between Hamming balls and affine subspaces that recovers previously known parameters. We then use folded Wronskian determinants to carry out an inductive proof that yields sharper bounds.

Shashank Srivastava
TTIC
shashanks@ias.edu

CP17

Tight Streaming Lower Bounds for Deterministic Approximate Counting

We study the streaming complexity of k-counter approximate counting. In the k-counter approximate counting problem, we are given an input string in $[k]^n$, and we are required to approximate the number of each j's ($j \in [k]$) in the string. Typically we require an additive error $\leq \frac{n}{3(k-1)}$ for each $j \in [k]$ respectively, and we are mostly interested in the regime $n \gg k$. We prove a lower bound result that the deterministic and worst-case k-counter approximate counting problem requires $\Omega(k \log(n/k))$ bits of space in the streaming model, while no non-trivial lower bounds were known before. In contrast, trivially counting the number of each $j \in [k]$ uses $O(k \log n)$ bits of space. Our main proof technique is analyzing a novel potential function. Our lower bound for k-counter approximate counting also implies the optimality of some other streaming algorithms.

For example, we show that the celebrated Misra-Gries algorithm for heavy hitters [MG82] has achieved optimal space usage.

Yichuan Wang
Tsinghua University
yichuan-21@mails.tsinghua.edu.cn

CP18

Flip Dynamics for Sampling Colorings: Improving $(11/6-\varepsilon)$ Using A Simple Metric

We present improved bounds for randomly sampling k-colorings of graphs with maximum degree Δ; our results hold without any further assumptions on the graph. The Glauber dynamics is a simple single-site update Markov chain. Jerrum (1995) proved an optimal $O(n \log n)$ mixing time bound for Glauber dynamics whenever $k > 2\Delta$ where Δ is the maximum degree of the input graph. This bound was improved by Vigoda (1999) to $k > (11/6)\Delta$ using a "flip' dynamics which recolors (small) maximal 2-colored components in each step. Vigoda's result was the best known for general graphs for 20 years until Chen et al. (2019) established optimal mixing of the flip dynamics for $k > (11/6-\varepsilon)\Delta$ where $\varepsilon \approx 10^{-5}$. We present the first substantial improvement over these results. We prove an optimal mixing time bound of $O(n \log n)$ for the flip dynamics when $k \geq 1.809\Delta$. Our proof utilizes path coupling with a simple weighted Hamming distance for "unblocked' neighbors.

Charlie A. Carlson
UCSB
charlieannecarlson@ucsb.edu

Eric Vigoda
University of California Santa Barbara
vigoda@ucsb.edu

CP18

A Polylogarithmic Approximation for Directed Steiner Forest in Planar Digraphs

We consider Directed Steiner Forest (DSF), a fundamental problem in network design. The input to DSF is a directed edge-weighted graph $G = (V, E)$ and a collection of vertex pairs $\{(s_i, t_i)\}_{i\in[k]}$. The goal is to find a minimum cost subgraph H of G such that H contains an s_i-t_i path for each $i \in [k]$. DSF is NP-Hard and is known to be hard to approximate to a factor of $\Omega(2^{\log^{1-\epsilon}(n)})$ for any fixed $\epsilon > 0$ [DK'99]. DSF admits approximation ratios of $O(k^{1/2+\epsilon})$ [CEGS'11] and $O(n^{2/3+\epsilon})$ [BBMRY'13]. In this work we show that in planar digraphs, an important and useful class of graphs in both theory and practice, DSF is much more tractable. We obtain an $O(\log^6 k)$-approximation algorithm via the junction tree technique. Our main technical contribution is to prove the existence of a low density junction tree in planar digraphs. To find an approximate junction tree we rely on recent results on rooted directed network design problems [FM'23, CJKZZ'24], in particular, on an LP-based algorithm for the Directed Steiner Tree problem [CJKZZ'24]. Our work and several other recent ones on algorithms for planar digraphs [FM'23, KS'21, CJKZZ'24] are built upon structural insights on planar graph reachability and shortest path separators [Thorup'04].

Chandra Chekuri, Rhea Jain
University of Illinois at Urbana-Champaign
chekuri@illinois.edu, rheaj3@illinois.edu

CP18

Congestion-Approximators from the Bottom Up

We develop a novel algorithm to construct a congestion-approximator with polylogarithmic quality on a capacitated, undirected graph in nearly-linear time. Our approach is the first bottom-up hierarchical construction, in contrast to previous top-down approaches including that of R"acke, Shah, and Taubig (SODA 2014), the only other construction achieving polylogarithmic quality that is implementable in nearly-linear time (Peng, SODA 2016). Similar to R"acke, Shah, and Taubig, our construction at each hierarchical level requires calls to an approximate max-flow/min-cut subroutine. However, the main advantage to our bottom-up approach is that these max-flow calls can be implemented directly without recursion. More precisely, the previously computed levels of the hierarchy can be converted into a pseudo-congestion-approximator, which then translates to a max-flow algorithm that is sufficient for the particular max-flow calls used in the construction of the next hierarchical level. As a result, we obtain the first non-recursive algorithms for congestion-approximator and approximate max-flow that run in nearly-linear time, a conceptual improvement to the aforementioned algorithms that recursively alternate between the two problems.

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

CP18

(Almost) Ruling Out Seth Lower Bounds for All-Pairs Max-Flow

The All-Pairs Max-Flow problem has gained significant popularity in the last two decades, and many results are known regarding its fine-grained complexity. Despite this, wide gaps remain in our understanding of the time complexity for several basic variants of the problem, including for directed or undirected input graphs that are edge- or node-capacitated, and where the capacities are unit or arbitrary. In this paper, we aim to bridge this gap by providing algorithms, conditional lower bounds, and non-reducibility results. Notably, we show that for most problem settings, deterministic reductions based on the Strong Exponential Time Hypothesis (SETH) cannot rule out $O(n^{4-\varepsilon})$ time algorithms for some small constant $\varepsilon > 0$, under a hypothesis called NSETH. To obtain our results for undirected graphs with unit node-capacities (aka All-Pairs Vertex Connectivity), we design a new randomized Las Vegas $O\big(m^{2+o(1)}\big)$ time combinatorial algorithm. This is our main technical result, improving over the recent $O\big(m^{11/5+o(1)}\big)$ time Monte Carlo algorithm [Huang et al., STOC 2023] and matching their $m^{2-o(1)}$ lower bound (up to subpolynomial factors), thus essentially settling the time complexity for this setting of the problem.

Ohad Trabelsi
Toyota Technological Institute at Chicago
ohadt@ttic.edu

CP18

Certificates in P and Subquadratic-Time Computation of Radius, Diameter, and All Eccentricities

in Graphs

In the context of fine-grained complexity, we investigate the notion of certificate enabling faster polynomial-time algorithms. We specifically target radius (minimum eccentricity), diameter (maximum eccentricity), and all-eccentricity computations for which quadratic-time lower bounds are known under plausible conjectures. In each case, we introduce a notion of certificate as a specific set of nodes from which appropriate bounds on all eccentricities can be derived in subquadratic time when this set has sublinear size. The existence of small certificates is a barrier against SETH-based lower bounds for these problems. We indeed prove that for graph classes with small certificates, there exist randomized subquadratic-time algorithms for computing the radius, the diameter, and all eccentricities respectively. Moreover, these notions of certificates are tightly related to algorithms probing the graph through one-to-all distance queries and allow to explain the efficiency of practical radius and diameter algorithms from the literature. Our formalization enables a novel primal-dual analysis of a classical approach for diameter computation that leads to algorithms for radius, diameter and all eccentricities with theoretical guarantees with respect to certain graph parameters. This is complemented by experimental results on various types of real-world graphs showing that these parameters appear to be low in practice. Finally, we obtain refined results for several graph classes.

Feodor F. Dragan
Kent State University
Dept. of Computer Science
dragan@cs.kent.edu

Guillaume Ducoffe
ICI National Institute for Research and Development in Informatics
guillaume.ducoffe@ici.ro

Michel Habib
IRIF
habib@irif.fr

Laurent Viennot
Inria
laurent.viennot@inria.fr

CP19

Forall-Exist Statements in Pseudopolynomial Time

Given a convex set $Q \subseteq R^m$ and an integer matrix $W \in Z^{mn}$, we consider statements of the form $\forall b \in (Q \cap Z^m)$ $\exists x \in Z^n$ s.t. $Wx \leq b$. Such statements can be verified in polynomial time with the algorithm of Kannan and its improvements if n is fixed and Q is a polyhedron. The running time of the best-known algorithms is doubly exponential in n. We provide a pseudopolynomial-time algorithm if m is fixed. Its running time is $(m\Delta)^{O(m^2)}$ where Δ is the largest absolute value of an entry in W. Furthermore it applies to general convex sets Q.

Eleonore Bach
EPFL
eleonore.bach@epfl.ch

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

Thomas Rothvoss
University of Washington
rothvoss@uw.edu

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

CP19

Integer Programs with Nearly Totally Unimodular Matrices: the Cographic Case

It is a notorious open question whether integer programs (IPs), with an integer coefficient matrix M whose subdeterminants are all bounded by a constant Δ in absolute value, can be solved in polynomial time. We answer this question in the affirmative if we further require that, by removing a constant number of rows and columns from M, one obtains a submatrix A that is the transpose of a network matrix. Our approach focuses on the case where A arises from M after removing k rows only, where k is a constant. We achieve our result in two main steps, the first related to the theory of IPs and the second related to graph minor theory. First, we derive a strong proximity result for the case where A is a general totally unimodular matrix: Given an optimal solution of the linear programming relaxation, an optimal solution to the IP can be obtained by finding a constant number of augmentations by circuits of $\begin{bmatrix} A & \mathbf{I} \end{bmatrix}$. Second, for the case where A is transpose of a network matrix, we reformulate the problem as a maximum constrained integer potential problem on a graph G. We observe that if G is 2-connected, then it has no rooted $K_{2,t}$-minor for $t = \Omega(k\Delta)$. We leverage this to obtain a tree-decomposition of G into highly structured graphs for which we can solve the problem locally. This allows us to solve the global problem via dynamic programming.

Manuel Aprile
Università di Padova
manuelf.aprile@gmail.com

Samuel Fiorini
Université Libre de Bruxelles
samuel.fiorini@ulb.be

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

Stefan Kober
Université Libre de Bruxelles
kober.stefan@gmail.com

Michal Seweryn
Charles University Prague
michalsew@gmail.com

Stefan Weltge
Technical University of Munich
weltge@tum.de

Yelena Yuditsky
McGill
yuditskyl@gmail.com

CP19

The Change-of-Measure Method, Block Lewis

Weights, and Approximating Matrix Block Norms

Given a matrix $\mathbf{A} \in \mathbb{R}^{k\times n}$, a partitioning of $[k]$ into groups $S_1, \ldots, S_m$, an outer norm p, and inner norms such that either $p \geq 1$ and $p_1, \ldots, p_m \geq 2$ or $p_1 = \cdots = p_m = p \geq 1/\log n$, we prove that there is a sparse weight vector $\beta \in \mathbb{R}^m$ such that $\sum_{i=1}^m \beta_i \cdot \|\mathbf{A}_{S_i} x\|_{p_i}^p \approx_{1\pm\varepsilon} \sum_{i=1}^m \|\mathbf{A}_{S_i} x\|_{p_i}^p$, where the number of nonzero entries of β is at most $\widetilde{O_{p,p_i}}(\varepsilon^{-2} n^{\max(1,p/2)})$. Additionally, we give efficient algorithms to find the sparse weight vector β in several regimes of p and $p_1, \ldots, p_m$. Our results imply an algorithm for minimizing sums of Euclidean norms in $\widetilde{O}(\varepsilon^{-1}\sqrt{n})$ linear system solves, improving over the previously known $\widetilde{O}(\sqrt{m}\log(1/\varepsilon))$ iteration complexity when $m \gg n$. Our main technical contribution is a generalization of the *change-of-measure* method that Bourgain, Lindenstrauss, and Milman used to obtain the analogous result when every group has size 1. Our generalization allows one to analyze changes of measure beyond those implied by D. Lewis's original construction, including the measure implied by the block Lewis weights and natural approximations of this measure.

Naren S. Manoj, Max Ovsiankin
Toyota Technological Institute Chicago
nsm@ttic.edu, maxov@ttic.edu

CP19

Complexity of Polytope Diameters Via Perfect Matchings

The (monotone) diameter of a polytope is a fundamental parameter with important connections to the efficiency of the simplex method. In 1984 Frieze and Teng proved the first cornerstone result in this direction by establishing that computing the diameter of an input polytope is weakly NP-hard. In a recent breakthrough-paper, Sanit (FOCS 2018) studied the diameter of a special class of graph-based polytopes, known as fractional matching polytopes, and showed that determining their diameters is NP-hard, thus establishing strong NP-hardness of computing the diameter of polytopes. As our first main result, we show that computing the diameter of perfect matching polytopes (of bipartite graphs) is NP-hard, giving an alternative, short proof for the strong NP-hardness of polytope diameters. In our second main result, we give a precise graph-theoretic description of the monotone diameter of perfect matching polytopes and use this description to prove the novel result that computing the monotone diameter of an input polytope is strongly NP-hard. Finally, as a consequence of these results, we solve an open problem posed and reiterated by Sanit; Kafer; and Borgwardt, Grewe, Kafer, Lee and Sanit; by proving the strong ¶-hardness of computing

the so-called circuit diameter of polytopes.

Christian Nöbel, Raphael Steiner
ETH Zurich
christian.noebel@math.ethz.ch, rsteine@ethz.ch

CP19

Testing Approximate Stationarity Concepts for Piecewise Affine Functions

We study the basic computational problem of detecting approximate stationary points for continuous piecewise affine (PA) functions. Our contributions span multiple aspects, including complexity, regularity, and algorithms. Specifically, we show that testing first-order approximate stationarity concepts, as defined by commonly used generalized subdifferentials, is computationally intractable unless P=NP. To facilitate computability, we consider a polynomial-time solvable relaxation by abusing the convex subdifferential sum rule and establish a tight characterization of its exactness. Furthermore, addressing an open issue motivated by the need to terminate the subgradient method in finite time, we introduce the first oracle-polynomial-time algorithm to detect so-called near-approximate stationary points for PA functions. A notable byproduct of our development in regularity is the first necessary and sufficient condition for the validity of an equality-type (Clarke) subdifferential sum rule. Our techniques revolve around two new geometric notions for convex polytopes and may be of independent interest in nonsmooth analysis. Moreover, some corollaries of our work on complexity and algorithms address open questions in the literature. To demonstrate the versatility of our results, we complement our findings with applications to a series of structured piecewise smooth functions, including rho-margin-loss SVM, piecewise affine regression, and nonsmooth neural networks.

Lai Tian
The Chinese University of Hong Kong, Hong Kong
tianlai.cs@gmail.com

Anthony So
The Chinese University of Hong Kong
manchoso@se.cuhk.edu.hk

CP20

Strict Self-Assembly of Discrete Self-Similar Fractals in the Abstract Tile Assembly Model

This paper answers a long-standing open question in tile-assembly theory, namely that it is possible to strictly assemble discrete self-similar fractals (DSSFs) in the abstract Tile-Assembly Model (aTAM). We prove this in 2 separate ways, each taking advantage of a novel set of tools. One of our constructions shows that specializing the notion of a quine, a program which prints its own output, to the language of tile-assembly naturally induces a fractal structure. The other construction introduces self-describing circuits as a means to abstractly represent the information flow through a tile-assembly construction and shows that such circuits may be constructed for a relative of the Sierpinski carpet, and indeed many other DSSFs, through a process of fixed-point iteration. This later result, or more specifically the machinery used in its construction, further enable us to provide a polynomial time procedure for deciding whether any given subset of $\mathbb{Z}^2$ will generate an aTAM producible DSSF. To this end, we also introduce the Tree Pump Theorem, a result analogous to the important Window Movie Lemma, but with requirements on the set of productions rather than on the self-assembling system itself.

Florent Becker
Université dOrléans
florent.becker@univ-orleans.fr

Daniel Hader, Matthew Patitz
University of Arkansas
dhader@uark.edu, patitz@uark.edu

CP20

Fast Static and Dynamic Approximation Algorithms for Geometric Optimization Problems: Piercing, Independent Set, Vertex Cover, and Matching

We develop simple and general techniques to obtain faster (near-linear time) static approximation algorithms, as well as efficient dynamic data structures, for four fundamental geometric optimization problems: minimum piercing set (MPS), maximum independent set (MIS), minimum vertex cover (MVC), and maximum-cardinality matching (MCM).

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

Timothy M. Chan
UIUC
tmc@illinois.edu

CP20

Ptases for Euclidean Tsp with Unit Disk and Unit Square Neighborhoods

The Euclidean Traveling Salesman Problem with Neighborhoods (ETSPN) is a well-studied problem in computational geometry and has a wealth of results. In this problem, given a set of geometric neighborhoods (or regions), the goal is to compute a shortest route that visits at least one point of each neighborhood. The problem is a generalization of the standard Euclidean TSP and hence is also NP-hard, even when the neighborhoods are disjoint unit disks or unit squares in the plane. A longstanding open problem on this topic is the existence of PTASes for ETSPN with unit disk (and unit square) neighborhoods. Prior to this work, the best-known approximation factor for unit disks is 6.75, and PTASes are only known for the special case where the unit disks/squares are of bounded depth, i.e., each point lies in at most a constant number of disks. In this paper, we resolve this open problem by giving the first PTASes for ETSPN with unit disks and unit squares. Our PTASes rely on new insights to the problem together with various classical tools such as Arora's technique and Baker's shifting technique.

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

Sayan Bandyapadhyay
Portland State University
sayanb@pdx.edu

katie clinch
UNSW Sydney
katharine.clinch@gmail.com

Daniel Lokshtanov
UCSB
daniello@ucsb.edu

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

Jie Xue
New York University Shanghai
jiexue@nyu.edu

CP20

Path and Intersections: Characterization of Quasi-metrics in Directed Okamura-Seymour Instances

We study the following distance realization problem. Given a quasi-metric D on a set T of terminals, does there exist a directed Okamura-Seymour graph that realizes D as the (directed) shortest-path distance metric on T? We show that, if we are further given the circular ordering of terminals lying on the boundary, then Monge property is a sufficient and necessary condition. This generalizes previous results for undirected Okamura-Seymour instances. With the circular ordering, we give a greedy algorithm for constructing a directed Okamura-Seymour instance that realizes the input quasi-metric. The algorithm takes the dual perspective concerning flows and routings, and is based on a new way of analyzing graph structures, by viewing graphs as *paths and their intersections*. We believe this new understanding is of independent interest and will prove useful in other problems in graph theory and graph algorithms. We also design an efficient algorithm for finding such a circular ordering that makes D satisfy Monge property, if one exists. Combined with our result above, this gives an efficient algorithm for the distance realization problem.

Yu Chen
University of Pennsylvania
yu.chen@nus.edu.sg

Zihan Tan
Rutgers University
zihantan1993@gmail.com

CP20

Flipping Non-Crossing Spanning Trees

For a set P of n points in general position in the plane, the flip graph $F(P)$ has a vertex for each non-crossing spanning tree on P and an edge between any two spanning trees that can be transformed into each other by one edge flip. The diameter $diam(F(P))$ of this graph is subject of intensive study. For points in general position, it is between $3n/2-5$ and $2n-4$, with no improvement for 25 years. For points in convex position, it lies between $3n/2-5$ and $\approx 1.95n$, where the lower bound was conjectured to be tight up to an additive constant and the upper bound is a recent breakthrough improvement over several bounds of the form $2n - o(n)$. In this work, we provide new upper and lower bounds on $diam(F(P))$, mainly focusing on points in convex position. We show $14n/9 - O(1) \leq diam(F(P)) \leq 5n/3 - 3$, by this disproving the conjectured upper bound of $3n/2$ for con-

vex position, and substantially improving both the long-standing lower bound for general position and the recent new upper bound for convex position. We complement these by showing that if one of T, T' has at most two boundary edges, then $dist(T, T') \leq 3d/2 < 3n/2$, where $d = |T - T'|$ is the number of edges in one tree that are not in the other. To prove both the upper and the lower bound, we introduce a new powerful tool, namely, acyclic subsets in an associated conflict graph, which might be of independent interest.

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

Håvard Bjerkevik
University at Albany
hbjerkevik@albany.edu

Linda Kleist
Universität Potsdam
kleist@cs.uni-potsdam.de

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

CP21

On Optimal Testing of Linearity

Linearity testing has been a focal problem in property testing of functions. We combine different known techniques and observations about Linearity testing in order to resolve two recent versions of this task. First, we focus on the online-manipulation-resilient model introduced by Kalemaj, Raskhodnikova and Varma (Theory of Computing 2023). In this model, up to t data entries are adversarially manipulated after each query is answered. Ben-Eliezer, Kelman, Meir, and Raskhodnikova (ITCS 2024) showed an asymptotically optimal Linearity tester that is resilient to t manipulations per query, but fails if t is too large. We simplify their analysis for the regime of small t, and for larger values of t we instead use sample-based testers, as defined by Goldreich and Ron (ACM Transactions on Computation Theory 2016). We complement our result by showing that when t is *very* large, any reasonable property, and in particular Linearity, cannot be tested at all. Second, we consider Linearity over the reals with proximity parameter ε. Fleming and Yoshida (ITCS 2020) gave a tester using $O(1/\varepsilon \cdot log(1/\varepsilon))$ queries. We simplify their algorithms, showing an optimal tester that only uses $O(1/\varepsilon)$ queries. This modification works for the low-degree testers presented in Arora, Bhattacharyya, Fleming, Kelman, and Yoshida (SODA 2023) too, resulting in optimal testers for degree-d polynomials, for any constant d.

Vipul Arora
Department of Computer Science, School of Computing,
National University of Singapore
vipul@comp.nus.edu.sg

Esty Kelman
Boston University, and
Massachusetts Institute of Technology
ekelman@mit.edu

Uri Meir
Blavatnik School of Computer Science,
Tel-Aviv University
urimeir.cs@gmail.com

CP21

Simple Sublinear Algorithms for (Delta + 1) Vertex Coloring Via Asymmetric Palette Sparsification

The palette sparsification theorem (PST) of Assadi, Chen, and Khanna (SODA 2019) states that in every graph G with maximum degree ?, sampling a list of O(log n) colors from 1, . . . , ? + 1 for every vertex independently and uniformly, with high probability, allows for finding a (? + 1) vertex coloring of G by coloring each vertex only from its sampled list. PST leads to a host of sublinear algorithms for (?+1) vertex coloring, including in semi-streaming, sublinear time, and MPC models, which are all proven to be nearly optimal. While being a natural and simple-to-state theorem, PST suffers from two drawbacks. Firstly, all its known proofs require technical arguments and secondly, finding the coloring of the graph from the sampled lists in an efficient manner requires a complicated algorithm. We show that a natural weakening of PST addresses both these drawbacks while still leading to sublinear algorithms of similar quality. In particular, we prove an asymmetric palette sparsification theorem (APST) that allows for list sizes of the vertices to have different sizes and only bounds the average size of these lists. The benefit of this weaker requirement is that we can now easily show the graph can be (? + 1) colored from the sampled lists using the standard greedy coloring algorithm. This way, we can recover nearly-optimal bounds for (? + 1) vertex coloring in all the aforementioned models using algorithms that are much simpler to implement and analyze.

Sepehr Assadi
Rutgers University, U.S.
sepehr@assadi.info

Helia Yazdanyar
University of Waterloo
hyazdanyar@uwaterloo.ca

CP21

Sublinear-Time Algorithm for MST-Weight Revisited

For graphs of average degree d, positive integer weights bounded by W, and accuracy parameter $\epsilon > 0$, [Chazelle, Rubinfeld, Trevisan; SICOMP'05] have shown that the weight of the minimum spanning tree can be $(1 + \epsilon)$-approximated in $\tilde{O}(Wd/\epsilon^2)$ expected time. This algorithm is frequently taught in courses on sublinear time algorithms. However, the $\tilde{O}(Wd/\epsilon^2)$-time variant requires an involved analysis, leading to simpler but much slower variations being taught instead. Here we present an alternative that is not only simpler to analyze, but also improves the number of queries, getting closer to the nearly-matching information theoretic lower bound. In addition to estimat-

ing the weight of the MST, our algorithm is also a perfect sampler for sampling uniformly at random an edge of the MST. At the core of our result is the insight that halting Prim's algorithm after an expected $\tilde{O}(d)$ number of steps, then returning the highest weighted edge of the tree, results in sampling an edge of the MST uniformly at random. Via repeated trials and averaging the results, this immediately implies an algorithm for estimating the weight of the MST. Since our algorithm is based on Prim's, it naturally works for non-integer weighted graphs as well.

Gryphon Patlin, Jan van den Brand
Georgia Tech
gpatlin3@gatech.edu, vdbrand@gatech.edu

CP21

Testing Identity of Distributions under Kolmogorov Distance in Polylogarithmic Space

Suppose we have a sample from a distribution D and we want to test whether $D = D^*$ for a fixed distribution D^*. Specifically, we want to reject with constant probability, if the distance of D from D^* is $\geq \varepsilon$ in a given metric. In the case of continuous distributions, this has been studied thoroughly in the statistics literature. Namely, for the well-studied Kolmogorov metric a test is known that uses the optimal $O(1/\varepsilon^2)$ samples. However, this test naively uses also space $O(1/\varepsilon^2)$, and previous work improved this to $O(1/\varepsilon)$. In this paper, we show that much less space suffices – we give an algorithm that uses space $O(\log^4 \varepsilon^{-1})$ in the streaming setting while also using an asymptotically optimal number of samples. This is in contrast with the standard total variation distance on discrete distributions for which such space reduction is known to be impossible. Finally, we state 9 related open problems that we hope will spark interest in this and related problems.

Jakub Tetek
INSAIT, Sofia University
j.tetek@gmail.com

Christian J. Lebeda
Inria
christian.j.lebeda@gmail.com

CP21

How to Design a Quantum Streaming Algorithm Without Knowing Anything About Quantum Computing

A series of work [GKK+08, Kal22, KPV24] has shown that asymptotic advantages in space complexity are possible for quantum algorithms over their classical counterparts in the streaming model. We give a simple quantum sketch that encompasses all these results, allowing them to be derived from entirely classical algorithms using our quantum sketch as a black box. The quantum sketch and its proof of correctness are designed to be accessible to a reader with no background in quantum computation, relying on only a small number of self-contained quantum postulates.

John M. Kallaugher, Ojas Parekh
Sandia National Laboratories
jmkall@sandia.gov, odparek@sandia.gov

Nadezhda Voronova
Boston University
voronova@bu.edu

CP22

On the Uniqueness of Bayesian Coarse Correlated Equilibria in Standard First-Price and All-Pay Auctions

We study the Bayesian coarse correlated equilibrium (BCCE) of continuous and discretised first-price and all-pay auctions under the standard symmetric independent private-values model. Our goal is to determine how the canonical Bayes-Nash equilibrium (BNE) of the auction relates to the outcome when all buyers bid following no-regret algorithms. Numerical experiments show that in two buyer first-price auctions the Wasserstein-2 distance of buyers' marginal bid distributions decline as $O(1/n)$ in the discretisation size in instances where the prior distribution is concave, whereas all-pay auctions exhibit similar behaviour without prior dependence. To explain this convergence to a near-equilibrium, we study uniqueness of the BCCE of the continuous auction, resulting in proofs of convergence of deterministic self-play to a near equilibrium outcome in these auctions. In the all-pay auction, we show that independent of the prior distribution there is a unique BCCE with symmetric, differentiable, and increasing bidding strategies, which is equivalent to the unique strict BNE. In the first-price auction, either the prior is strictly concave or the learning algorithm has to be restricted to strictly increasing strategies. Without such strong assumptions, no-regret algorithms can end up in low-price pooling strategies.

Mete Seref Ahunbay, Martin Bichler
Technical University of Munich
mete.ahunbay@cit.tum.de, bichler@cit.tum.de

CP22

Clock Auctions Augmented with Unreliable Advice

We provide the first analysis of (deferred-acceptance) clock auctions in the learning-augmented framework. Clock auctions satisfy a unique list of appealing properties that make them very practical, but no deterministic clock auction with n bidders can achieve a $O(\log^{1-\epsilon} n)$ approximation of the optimal social welfare for constant $\epsilon > 0$, even in very simple settings. This overly pessimistic result heavily depends on the unrealistic assumption that the designer has absolutely no information regarding the bidders' values. The learning-augmented framework instead assumes that the designer is equipped with some (machine-learned) advice regarding the optimal solution, which can provide very useful guidance if it is accurate, but may be highly unreliable. Our main results are learning-augmented clock auctions that use this advice to achieve stronger performance guarantees whenever the advice is accurate (known as consistency), while maintaining worst-case guarantees even if this advice is arbitrarily inaccurate (known as robustness). Our first clock auction achieves the best of both worlds: $(1 + \epsilon)$-consistency for any constant $\epsilon > 0$ and $O(\log n)$ robustness. We then consider a much stronger notion of consistency and provide an auction that achieves

the optimal trade-off between this notion and robustness. Finally, we prove bounds regarding the "cost of smoothness,' i.e., the loss in robustness required to achieve error tolerance.

Vasilis Gkatzelis, Daniel Schoepflin, Xizhi Tan
Drexel University
gkatz@drexel.edu, schoep@drexel.edu, xizhi@drexel.edu

CP22

Tolls for Dynamic Equilibrium Flows

We consider dynamic network flows and study the following question: Which dynamic edge flows can be implemented as tolled dynamic equilibrium flows? We study this question for the "heterogeneous-user' model, where the flow particles are partitioned into populations having different valuations of travel time and money spent. As our main result, we give the first characterization of this type of implementability showing that for single-source single-destination networks and heterogeneous users, a dynamic edge flow is implementable by tolls if and only if the induced subgraph of the edge flow contains no cycle of positive length containing the destination. For the proof of this result we make several technical contributions: We formulate a novel infinite dimensional optimization problem, where the goal is to minimize the weighted travel times with respect to the fixed network loading induced by the given edge flow. Using the recently introduced concept of parameterized network loadings (cf. [Graf, Harks, Schwarz: A Decomposition Theorem for Dynamic Flows, 2024]), we prove existence of optimal solutions, strong duality, and a characterization of special optimal solutions for which an inequality is tight. These results are then all used for the proof of the above mentioned main characterization.

Julian Schwarz
University of Passau
Faculty of Computer Science and Mathematics
julian.schwarz@uni-passau.de

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

Lukas Graf
University of Passau
Faculty of Computer Science and Mathematics
lukas.graf@uni-passau.de

CP22

Approximating Competitive Equilibrium by Nash Welfare

We explore the relationship between two popular concepts in the allocation of divisible items: competitive equilibrium (CE) and allocations that maximize Nash welfare. While these concepts align for agents with homogeneous concave utility functions, they diverge for non-homogeneous utilities. From a computational perspective, maximizing Nash welfare amounts to solving a convex program for any concave utility functions, whereas computing CE becomes PPAD-hard already for separable piecewise linear concave (SPLC) utilities. We introduce the concept of Gale-substitute utility functions showing that any allocation maximizing Nash welfare provides an approximate-CE with surprisingly strong guarantees, where every agent gets at least half the maximum utility they can get at any CE, and is approximately envy-free. Gale-substitutes include utility functions where computing CE is PPAD hard, such as all separable concave utilities and the previously studied non-separable class of Leontief-free utilities. We introduce a broad new class of utility functions called generalized network utilities. This class includes SPLC and Leontief-free utilities, and we show that all such utilities are Gale-substitutes. Conversely, although some agents may get much higher utility at a Nash welfare maximizing allocation than at a CE, we show a price of anarchy type result: for general concave utilities, every CE achieves at least 0.69 fraction of the maximum Nash welfare, and this factor is tight.

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

Yixin Tao
Shanghai University of Finance and Economics
taoyixin@mail.shufe.edu.cn

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

CP22

Platforms for Efficient and Incentive-Aware Collaboration

Collaboration is crucial for reaching collective goals. However, its potential for effectiveness is often undermined by the strategic behavior of individual agents — a fact that is captured by a high Price of Stability (PoS) in recent literature [BHPS21]. Implicit in the traditional PoS analysis is the assumption that agents have full knowledge of how their tasks relate to one another. We offer a new perspective on bringing about efficient collaboration across strategic agents using information design. Inspired by the increasingly important role collaboration plays in machine learning (such as platforms for collaborative federated learning and data cooperatives), we propose a framework in which the platform possesses more information about how the agents' tasks relate to each other than the agents themselves. Our results characterize how and to what degree such platforms can leverage their information advantage and steer strategic agents towards efficient collaboration. Concretely, we consider collaboration networks in which each node represents a task type held by one agent, and each task benefits from contributions made to the task itself and its neighboring tasks. This network structure is known to the agents and the platform, but only the platform knows each agent's real location. We employ private Bayesian persuasion and design two families of persuasive signaling schemes that the platform can use to guarantee a small total workload when agents follow

the signal.

Kunhe Yang, Nika Haghtalab, Mingda Qiao
University of California, Berkeley
kunheyang@berkeley.edu, nika@berkeley.edu, mingda.qiao.cs@gmail.com

CP23

Improved Spectral Density Estimation Via Explicit and Implicit Deflation

We study algorithms for approximating the spectral density of a symmetric matrix A that is accessed through matrix-vector products. By combining an existing Chebyshev polynomial moment matching method with a deflation step that approximately projects off the largest magnitude eigendirections of A before estimating the spectral density, we give an $\epsilon\sigma_\ell(A)$ error approximation in the Wasserstein-1 metric using $O(\ell \log n + 1/\epsilon)$ matrix-vector products, where $\sigma_\ell(A)$ is the ℓ^{th} largest singular value of A. When A exhibits fast singular value decay, this can be much stronger than prior work, which gives error $\epsilon\sigma_1(A)$ using $O(1/\epsilon)$ matrix-vector products. We also show that our bound is nearly tight: $\Omega(\ell + 1/\epsilon)$ matrix-vector products are required to achieve error $\epsilon\sigma_\ell(A)$. We further show that the popular Stochastic Lanczos Quadrature (SLQ) method matches the above bound, even though SLQ itself is parameter-free and performs no explicit deflation. This explains the strong practical performance of SLQ, and motivates a simple variant that achieves an even tighter error bound. Our error bound for SLQ leverages an analysis that views it as an implicit polynomial moment matching method, along with recent results on low-rank approximation with single-vector Krylov methods. We use these results to show that SLQ can perform implicit deflation as part of moment matching.

Rajarshi Bhattacharjee
University of Massachusetts Amherst
rbhattacharj@cs.umass.edu

Rajesh Jayaram
Google Research
rkjayaram@google.com

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

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

Archan Ray
Sloan Kettering Institute, U.S.
talk2archan@gmail.com

CP23

Solving Polynomial Equations Over Finite Fields

We present a randomized algorithm for solving low-degree polynomial equation systems over finite fields faster than exhaustive search. In order to do so, we follow a line of work by Lokshtanov, Paturi, Tamaki, Williams, and Yu (SODA 2017), Bjrklund, Kaski, and Williams (ICALP 2019), and Dinur (SODA 2021). In particular, we generalize Dinur's algorithm for $\mathbb{F}_2$ to all finite fields, in particular the symbolic interpolation of Bjrklund, Kaski, and Williams, and we use an efficient trimmed multipoint evaluation and interpolation procedure for multivariate polynomials over finite fields by Van der Hoeven and Schost (AAECC 2013). The running time of our algorithm matches that of Dinur's algorithm for $\mathbb{F}_2$ and is significantly faster than the one of Lokshtanov et al. for $q > 2$. We complement our results with tight conditional lower bounds that, surprisingly, we were not able to find in the literature. In particular, under the strong exponential time hypothesis, we prove that it is impossible to solve n-variate low-degree polynomial equation systems over $\mathbb{F}_q$ in time $O((q-\varepsilon)^n)$. As a bonus, we show that under the counting version of the strong exponential time hypothesis, it is impossible to compute the number of roots of a single n-variate low-degree polynomial over $\mathbb{F}_q$ in time $O((q-\varepsilon)^n)$; this generalizes a result of Williams (SOSA 2018) from $\mathbb{F}_2$ to all finite fields.

Holger Dell
ITU Copenhagen
dell@em.uni-frankfurt.de

Anselm Haak
Universität Paderborn
anselm.haak@uni-paderborn.de

Melvin Kallmayer
Goethe University Frankfurt
mel.kallmayer@gmail.com

Leo Wennmann
Maastricht University
leo.wennmann@gmail.com

CP23

Near-Optimal Hierarchical Matrix Approximation from Matrix-Vector Products

We describe a randomized algorithm for producing a near-optimal hierarchical off-diagonal low-rank (HODLR) approximation to an $n \times n$ matrix $\vec{A}$, accessible only though matrix-vector products with $\vec{A}$ and $\vec{A}^T$. We prove that, for the rank-k HODLR approximation problem, our method achieves a $(1+\beta)^{\log(n)}$-optimal approximation in expected Frobenius norm using $O(k \log(n)/\beta^3)$ matrix-vector products. In particular, the algorithm obtains a $(1+\varepsilon)$-optimal approximation with $O(k \log^4(n)/\varepsilon^3)$ matrix-vector products, and for any constant c, an n^c-optimal approximation with $O(k \log(n))$ matrix-vector products. We complement the upper bound with a nearly-matching lower bound. Our algorithm can be viewed as a robust version of widely used "peeling" methods for recovering HODLR matrices and is, to the best of our knowledge, the first matrix-vector query algorithm to enjoy theoretical worst-case guarantees for approximation by any hierarchical matrix class.To control the propagation of error between levels of hierarchical approximation, we introduce a new perturbation bound for low-rank approximation, which shows that the widely used Generalized Nystrm method enjoys inherent stability when

implemented with noisy matrix-vector products.We also introduce a novel "randomly perforated" matrix sketching method to further control the error in the peeling algorithm.

Feyza Duman Keles
NYU
fd2135@nyu.edu

Tyler Chen
New York University
tyler.chen@nyu.edu

Diana Halikias
Cornell University
dh736@cornell.edu

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

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

David Persson
EPFL
david.persson@epfl.ch

CP23

On the Decidability of Presburger Arithmetic Expanded with Powers

We prove that for any integers $\alpha, \beta > 1$, the existential fragment of the first-order theory of the structure $\langle \mathbb{Z}; 0, 1, <, +, \alpha^{\mathbb{N}}, \beta^{\mathbb{N}} \rangle$ is decidable (where $\alpha^{\mathbb{N}}$ is the set of positive integer powers of α, and likewise for $\beta^{\mathbb{N}}$). On the other hand, we show by way of hardness that decidability of the existential fragment of the theory of $\langle \mathbb{N}; 0, 1, <, +, x \mapsto \alpha^x, x \mapsto \beta^x \rangle$ for any multiplicatively independent $\alpha, \beta > 1$ would lead to mathematical breakthroughs regarding base-α and base-β expansions of certain transcendental numbers. Finally, modifying the original proof of Hieronymi and Schulz we show that for any multiplicatively independent $\alpha, \beta > 1$, it is undecidable whether a given formula with at most 3 alternating blocks of quantifiers holds in $\langle \mathbb{N}; 0, 1, <, +, \alpha^{\mathbb{N}}, \beta^{\mathbb{N}} \rangle$.

Toghrul Karimov
Max Planck Institute for Software Systems
toghs@mpi-sws.org

Florian Luca
Stellenbosch University
fluca@sun.ac.za

Joris Nieuwveld
Max Planck Institute for Software Systems
jnieuwve@mpi-sws.org

Joël Ouaknine
Max Planck Institute for Software Systems, Germany
joel@mpi-sws.org

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

CP23

Fast Deterministic Chromatic Number under the Asymptotic Rank Conjecture

In this paper we further explore the recently discovered connection by Bjrklund and Kaski [STOC 2024] and Pratt [STOC 2024] between the asymptotic rank conjecture of Strassen [Progr. Math. 1994] and the three-way partitioning problem. We show that under the asymptotic rank conjecture, the chromatic number of an n-vertex graph can be computed deterministically in $O(1.99982^n)$ time, thus giving a conditional answer to a question of Zamir [ICALP 2021], and questioning the optimality of the $2^n \operatorname{poly}(n)$ time algorithm for chromatic number by Bjrklund, Husfeldt, and Koivisto [SICOMP 2009]. Viewed in the other direction, if chromatic number indeed requires deterministic algorithms to run in close to 2^n time, we obtain a sequence of explicit tensors of superlinear rank, falsifying the asymptotic rank conjecture. Our technique is a combination of earlier algorithms for detecting k-colorings for small k and enumerating k-colorable subgraphs, with an extension and derandomisation of Pratt's tensor-based algorithm for balanced three-way partitioning to the unbalanced case.

Andreas Björklund
IT University of Copenhagen
anbjo@itu.dk

Kevin Pratt
New York University
kp2154@nyu.edu

Petteri Kaski
Aalto University
Department of Computer Science
petteri.kaski@aalto.fi

Thore Husfeldt
IT University of Copenhagen
thore@itu.dk

Radu Curticapean
University of Regensburg and IT University of Copenhagen
radu.curticapean@ur.de

CP24

Almost Tight Bounds for Differentially Private Densest Subgraph

We study the Densest Subgraph (DSG) problem under the additional constraint of differential privacy. DSG is a fundamental theoretical question that plays a central role in graph analytics, and so privacy is a natural requirement. All known private algorithms for Densest Subgraph lose constant multiplicative factors, despite the existence of non-private exact algorithms. We show that, perhaps sur-

prisingly, this loss is not necessary: in both the classic differential privacy model and the LEDP model, we give (ϵ, δ)-differentially private algorithms with no multiplicative loss whatsoever. In other words, the loss is purely additive. Moreover, our additive losses match or improve the previous state of the art additive loss when $1/\delta$ is polynomial in n, and are almost tight: in the centralized setting, our additive loss is $O(\log n/\epsilon)$ while there is a known lower bound of $\Omega(\sqrt{\log n/\epsilon})$. We also give a number of extensions. First, we show how to extend our techniques to both the node-weighted and the directed versions of the problem. Second, we give a separate algorithm with pure differential privacy but with worse approximation bounds. And third, we give a new algorithm for privately computing the optimal density which implies a separation between the structural problem of privately computing the densest subgraph and the numeric problem of privately computing the density of the densest subgraph.

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

Satyen Kale
Apple
satyen.kale@gmail.com

Silvio Lattanzi
Google Zurich
silviol@google.com

Sergei Vassilvitskii
Google Research
sergeiv@google.com

CP24

Local Lipschitz Filters for Bounded-Range Functions with Applications to Arbitrary Real-Valued Functions

We study local filters for the Lipschitz property of real-valued functions $f : V \to [0, r]$, where the Lipschitz property is defined with respect to an arbitrary undirected graph $G = (V, E)$. We give nearly optimal local Lipschitz filters both with respect to ℓ_1-distance and ℓ_0-distance. Prior work only considered unbounded-range functions over $[n]^d$. Jha and Raskhodnikova (SICOMP '13) gave an algorithm for such functions with lookup complexity exponential in d, which Awasthi et al. (ACM Trans. Comput. Theory) showed was necessary in this setting. We demonstrate that important applications of local Lipschitz filters can be accomplished with filters for functions whose range is bounded in $[0, r]$. For functions $f : [n]^d \to [0, r]$, we achieve runtime $(d^r \log n)^{O(\log r)}$ for the ℓ_1-respecting filter and $d^{O(r)}\text{polylog}\, n$ for the ℓ_0-respecting filter, thus circumventing the lower bound. Our local filters provide a novel Lipschitz extension that can be implemented locally. Furthermore, we show that our algorithms are nearly optimal in terms of the dependence on r for the domain $\{0, 1\}^d$, an important special case of the domain $[n]^d$. In addition, our lower bound resolves an open question of Awasthi et al., removing one of the conditions necessary for their lower bound for general range. We prove our lower bound via a reduction from distribution-free Lipschitz testing and a new technique for proving hardness for adaptive algorithms.

Jane Lange
MIT
jlange@mit.edu

Ephraim Linder, Sofya Raskhodnikova
Boston University
ejlinder@bu.edu, sofya@bu.edu

Arsen Vasilyan
MIT
vasilyan@mit.edu

CP24

Private Mean Estimation with Person-Level Differential Privacy

We study person-level differentially private (DP) mean estimation in the case where each person holds multiple samples. DP here requires the usual notion of distributional stability when all of a person's datapoints can be modified. Informally, if n people each have m samples from an unknown d-dimensional distribution with bounded k-th moments, we show that

$$n = \tilde{\Theta}\left(\frac{d}{\alpha^2 m} + \frac{d}{\alpha m^{1/2}\varepsilon} + \frac{d}{\alpha^{k/(k-1)} m\varepsilon} + \frac{d}{\varepsilon}\right)$$

people are necessary and sufficient to estimate the mean up to distance α in ℓ_2-norm under ε-differential privacy (and its common relaxations). In the multivariate setting, we give computationally efficient algorithms under approximate DP and computationally inefficient algorithms under pure DP, and our nearly matching lower bounds hold for the most permissive case of approximate DP. Our computationally efficient estimators are based on the standard clip-and-noise framework, but the analysis for our setting requires both new algorithmic techniques and new analyses. In particular, our new bounds on the tails of sums of independent, vector-valued, bounded-moments random variables may be of interest.

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

Sushant Agarwal
Northeastern University
agarwal.sus@northeastern.edu

Gautam Kamath
University of Waterloo
g@csail.mit.edu

Mahbod Majid
Massachusetts Institute of Technology
mahbod@mit.edu

Argyris Mouzakis
University of Waterloo
amouzaki@uwaterloo.ca

Jonathan Ullman
Northeastern University
jullman@ccs.neu.edu

CP24

Improved Differentially Private Continual Observation Using Group Algebra

In this paper, we show a novel connection between mechanisms for continual weighted prefix sum and a concept in representation theory known as the group matrix. To the best of our knowledge, this is the first application of group algebra in the analysis of differentially private algorithms. Using this connection, we analyze factorization norms that give upper and lower bounds for the additive error under general ℓ_p-norms of the matrix mechanism. This allows us to give 1. the first efficient factorization that improves the best-known non-constructive upper bound on the factorization norm by Mathias (1993) for the prefix-sum matrix, and also improves on the previous best-known constructive bound, 2. the first upper bound on the additive error for a large class of weight functions for weighted prefix sum problems; 3. a general improved upper bound on the factorization norms that depend on algebraic properties of the weighted sum matrices and that applies to a more general class of weighting functions than the ones considered in Henzinger, Upadhyay, and Upadhyay (SODA 2024). Using the known connection between these factorization norms and the ℓ_p-error of continual weighted sum, we give an upper bound on the ℓ_p-error for the continual weighted sum problem for $p \geq 2$.

Jalaj Upadhyay
Rutgers University
jalaj.upadhyay@rutgers.edu

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

CP25

A Simple and Combinatorial Approach to Proving Chernoff Bounds and Their Generalizations

The Chernoff bound is one of the most widely used tools in theoretical computer science. It's rare to find a randomized algorithm that doesn't employ a Chernoff bound in its analysis. The standard proofs of Chernoff bounds are beautiful but in some ways not very intuitive. In this paper, I'll show you a different proof that has four features: the proof offers a strong intuition for why Chernoff bounds look the way that they do; the proof is user-friendly and (almost) algebra-free; the proof comes with matching lower bounds, up to constant factors in the exponent; and the proof extends to establish generalizations of Chernoff bounds in other settings. The ultimate goal is that, once you know this proof (and with a bit of practice), you should be able to confidently reason about Chernoff-style bounds in your head, extending them to other settings, and convincing yourself that the bounds you're obtaining are tight (up to constant factors in the exponent).

William Kuszmaul
Massachusetts Institute of Technology
kuszmaul@mit.edu

CP25

Better Gaussian Mechanism Using Correlated Noise

We present a simple variant of the Gaussian mechanism for answering differentially private queries when the sensitivity space has a certain common structure. Our motivating problem is the fundamental task of answering d counting queries under the add/remove neighboring relation. The standard Gaussian mechanism solves this task by adding noise distributed as a Gaussian with variance scaled by d independently to each count. We show that adding a random variable distributed as a Gaussian with variance scaled by $(\sqrt{d}+1)/4$ to all counts allows us to reduce the variance of the independent Gaussian noise samples to scale only with $(d+\sqrt{d})/4$. The total noise added to each counting query follows a Gaussian distribution with standard deviation scaled by $(\sqrt{d}+1)/2$ rather than $\sqrt{d}$. The central idea of our mechanism is simple and the technique is flexible. We show that applying our technique to another problem gives similar improvements over the standard Gaussian mechanism.

Christian J. Lebeda
Inria
christian.j.lebeda@gmail.com

CP25

A Multilinear Johnson-Lindenstrauss Transform

The JohnsonLindenstrauss family of transforms constitutes a key algorithmic tool for reducing the dimensionality of a Euclidean space with low distortion of distances. Rephrased from geometry to linear algebra, one seeks to reduce the dimension of a vector space while approximately preserving inner products. We present a multilinear generalization of this bilinear (inner product) setting that admits both an elementary randomized algorithm as well as a short proof of correctness using Orlicz quasinorms.

Antonis Matakos
Aalto University
antonis.matakos@aalto.fi

Petteri Kaski
Aalto University
Department of Computer Science
petteri.kaski@aalto.fi

Heikki Mannila
Aalto University
heikki.mannila@aalto.fi

CP25

Only Two Shuffles Perform Card-Based Zero-Knowledge Proof for Sudoku of Any Size

Sudoku is a popular pencil puzzle where a player fills in the empty cells with numbers on an $n \times n$ board so that each row, column, and $(\sqrt{n} \times \sqrt{n})$-block must contain all the

numbers from 1 to n; a typical puzzle size is $n = 9$. In 2007, Gradwohl, Naor, Pinkas, and Rothblum proposed a physical zero-knowledge proof protocol for Sudoku using a physical deck of cards; their card-based protocol requires $3n\ell$ shuffles, where ℓ is a security parameter to eliminate the soundness error. Since the invention of this seminal protocol, several soundness-error-free protocols were constructed to reduce the number of required shuffles; the state-of-the-art one was designed in 2023, which uses $7\sqrt{n} - 5$ shuffles. In this paper, we show that only three or two shuffles are sufficient to construct a zero-knowledge proof protocol for Sudoku, no matter how large n is, i.e., we propose two card-based protocols using constant numbers (namely, 3 and 2) of shuffles. Our proposed protocols are simple and efficient enough for people to execute for a 9×9 Sudoku puzzle in reality.

Kodai Tanaka
Tohoku University
kodai.tanaka.r2@dc.tohoku.ac.jp

Shun Sasaki
Ibaraki University
23nm723g@vc.ibaraki.ac.jp

Kazumasa Shinagawa
Ibaraki University
AIST
kazumasa.shinagawa.np92@vc.ibaraki.ac.jp

Takaaki Mizuki
Tohoku University
AIST
mizuki@cc.tohoku.ac.jp

CP25

Ellipsoid Fitting Up to Constant Via Empirical Covariance Estimation

The ellipsoid fitting conjecture of Saunderson, Chandrasekaran, Parrilo and Willsky considers the maximum number n random Gaussian points in $\mathbb{R}^d$, such that with high probability, there exists an origin-symmetric ellipsoid passing through all the points. They conjectured a threshold of $n = (1 - o_d(1)) \cdot d^2/4$, while until recently, known lower bounds on the maximum possible n were of the form $d^2/(\log d)^{O(1)}$. We give a simple proof based on concentration of sample covariance matrices, that with probability $1 - o_d(1)$, it is possible to fit an ellipsoid through d^2/C random Gaussian points. Similar results were also obtained in two independent works by Hsieh, Kothari, Potechin and Xu [ICALP 2023] and by Bandeira, Maillard, Mendelson, and Paquette [arXiv, July 2023].

June Wu
University of Chicago
jqw@uchicago.edu

Madhur Tulsiani
TTI Chicago
madhurt@ttic.edu

CP26

Improved Bounds for Fully Dynamic Matching Via Ordered Ruzsa-Szemeredi Graphs

In a very recent breakthrough, Behnezhad and Ghafari [FOCS'24] developed a novel fully dynamic randomized algorithm for maintaining a $(1 - \epsilon)$-approximation of maximum matching with amortized update time *potentially* much better than the trivial $O(n)$ update time. The runtime of the BG algorithm is parameterized via the following graph theoretical concept: * For any n, define $\mathsf{ORS}(n)$—standing for *Ordered Ruzsa-Szemeredi Graph*—to be the largest number of edge-disjoint matchings $M_1, \ldots, M_t$ of size $\Theta(n)$ in an n-vertex graph such that for every $i \in [t]$, M_i is an induced matching in the subgraph $M_i \cup M_{i+1} \cup \ldots \cup M_t$. Then, for any fixed $\epsilon > 0$, the BG algorithm runs in

$$O\left(\sqrt{n^{1+O(\epsilon)} \cdot \mathsf{ORS}(n)}\right)$$

amortized update time with high probability, even against an adaptive adversary.

Our Result: In this work, we further strengthen the result of Behnezhad and Ghafari and push it to limit to obtain a randomized algorithm with amortized update time of

$$n^{o(1)} \cdot \mathsf{ORS}(n)$$

with high probability, even against an adaptive adversary. In the limit, i.e., *if* current lower bounds for $\mathsf{ORS}(n) = n^{o(1)}$ are almost optimal, our algorithm achieves an $n^{o(1)}$ update time for $(1-\epsilon)$-approximation of maximum matching, almost fully resolving this fundamental question.

Sepehr Assadi
Rutgers University, U.S.
sepehr@assadi.info

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

Peter Kiss
University of Vienna
peter.kiss@warwick.ac.uk

CP26

Matching Composition and Efficient Weight Reduction in Dynamic Matching

We consider the foundational problem of maintaining a $(1 - \varepsilon)$-approximate maximum weight matching (MWM) in an n-node dynamic graph undergoing edge insertions and deletions. We provide a general reduction that reduces the problem on graphs with a weight range of $\text{poly}(n)$ to $\text{poly}(1/\varepsilon)$ at the cost of just an additive $\text{poly}(1/\varepsilon)$ in update time. This improves upon the prior reduction of Gupta-Peng (FOCS 2013) which reduces the problem to a weight range of $\varepsilon^{-O(1/\varepsilon)}$ with a multiplicative cost of $O(\log n)$. When combined with a reduction of Bernstein-Dudeja-Langley (STOC 2021) this yields a reduction from dy-

namic $(1-\varepsilon)$-approximate MWM in bipartite graphs with a weight range of $\text{poly}(n)$ to dynamic $(1-\varepsilon)$-approximate maximum cardinality matching in bipartite graphs at the cost of a multiplicative $\text{poly}(1/\varepsilon)$ in update time, thereby resolving an open problem in [GP'13; BDL'21]. Additionally, we show that our approach is amenable to MWM problems in streaming, shared-memory work-depth, and massively parallel computation models. We also apply our techniques to obtain an efficient dynamic algorithm for rounding weighted fractional matchings in general graphs. Underlying our framework is a new structural result about MWM that we call the "matching composition lemma' and new dynamic matching subroutines that may be of independent interest.

Aaron Bernstein
Rutgers University
bernstei@gmail.com

Jiale Chen
Stanford University
jialec@stanford.edu

Aditi Dudeja
University of Salzburg
aditi.dudeja@plus.ac.at

Zachary Langley
Rutgers University
zach.langley@rutgers.edu

Aaron Sidford, Ta-Wei Tu
Stanford University
sidford@stanford.edu, taweitu@stanford.edu

CP26

New Philosopher Inequalities for Online Bayesian Matching, Via Pivotal Sampling

We study the polynomial-time approximability of the optimal online stochastic bipartite matching algorithm, initiated by Papadimitriou et al. (EC'21). Here, nodes on one side of the graph are given upfront, while at each time t, an online node and its edge weights are drawn from a time-dependent distribution. The optimal algorithm is PSPACE-hard to approximate within some universal constant. We refer to this optimal algorithm, which requires time to think (compute), as a *philosopher*, and refer to polynomial-time online approximations of the above as *philosopher inequalities*. Our main result is a state-of-the-art 0.678-approximate algorithm; we also achieve a 0.685-approximation for a vertex-weighted special case. Building on our algorithms and the recent black-box reduction of Banihashem et al. (SODA'24), we provide poly-time (pricing-based) truthful mechanisms which 0.678-approximate the social welfare of the optimal online allocation for bipartite matching markets. Our online allocation algorithm relies on the classic pivotal sampling algorithm along with careful discarding to obtain strong negative correlations between offline nodes, while matching using the highest-value edges. Consequently, the analysis boils down to examining a weighted sum X of negatively correlated Bernoulli variables, specifically lower bounding its mass below a threshold, $\mathbb{E}[\min(1, X)]$, of possible independent interest.

Tristan Pollner
Stanford University
tristan.pollner@gmail.com

Mark Braverman
Princeton University
mbraverman@gmail.com

Mahsa Derakhshan
Northeastern University
m.derakhshan@northeastern.edu

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

David Wajc
Technion
david.wajc@gmail.com

CP26

Entropy Regularization and Faster Decremental Matching in General Graphs

We provide an algorithm that maintains, against an adaptive adversary, a $(1-\varepsilon)$-approximate maximum matching in n-node m-edge general (not necessarily bipartite) undirected graph undergoing edge deletions with high probability with (amortized) $O(\text{poly}(\varepsilon^{-1}, \log n))$ time per update. We also obtain the same update time for maintaining a fractional approximate weighted matching (and hence an approximation to the value of the maximum weight matching) and an integral approximate weighted matching in dense graphs. Our unweighted result improves upon the prior state-of-the-art which includes a $\text{poly}(\log n) \cdot 2^{O(1/\varepsilon^2)}$ update time [Assadi-Bernstein-Dudeja 2022] and an $O(\sqrt{m}\varepsilon^{-2})$ update time [Gupta-Peng 2013], and our weighted result improves upon the $O(\sqrt{m}\varepsilon^{-O(1/\varepsilon)} \log n)$ update time due to [Gupta-Peng 2013]. To obtain our results, we generalize a recent optimization approach to dynamic algorithms from [Jambulapati-Jin-Sidford-Tian 2022]. We show that repeatedly solving entropy-regularized optimization problems yields a lazy updating scheme for fractional decremental problems with a near-optimal number of updates. To apply this framework we develop optimization methods compatible with it and new dynamic rounding algorithms for the matching polytope.

Jiale Chen, Aaron Sidford, Ta-Wei Tu
Stanford University
jialec@stanford.edu, sidford@stanford.edu, taweitu@stanford.edu

CP26

Online Dependent Rounding Schemes for Bipartite Matchings, with Applications

The objective is to maximize the ¡i¿rounding ratio¡/i¿ of the output matching M, which is the minimum over

all fractional b-matchings $\mathbf{x}$, and edges e, of the ratio $\Pr[e \in M]/x_e$. In analogy with the highly influential offline dependent rounding schemes of Gandhi et al. (FOCS'02, J.ACM'06), we refer to such algorithms as ¡i¿online dependent rounding schemes¡/i¿ (ODRSes). This problem, with additional restrictions on the possible inputs $\mathbf{x}$, has played a key role in recent developments in online computing. We provide the first generic b-matching ODRSes that impose no restrictions on $\mathbf{x}$. Specifically, we provide ODRSes with rounding ratios of 0.652 and 0.646 for b-matchings and simple matchings, respectively. This breaks the natural barrier of $1 - 1/e$, prevalent for online matching problems, and numerous online problems more broadly. Using our ODRSes, we provide a number of algorithms with similar better-than-$(1 - 1/e)$ ratios for several problems in online edge coloring, stochastic optimization, and more. Our techniques, which have already found applications in several follow-up works (Patel and Wajc SODA'24, Blikstad et al. SODA'25, Braverman et al. SODA'25, and Aouad et al. 2024), include periodic use of ¡i¿offline¡/i¿ contention resolution schemes (in online algorithm design), grouping nodes, and a new scaling method which we call ¡i¿group discount and individual markup¡/i¿.

Joseph (Seffi) Naor
Technion
naor@cs.technion.ac.il

Aravind Srinivasan
University of Maryland, College Park
asriniv1@umd.edu

David Wajc
Technion
david.wajc@gmail.com

CP27

Highway Dimension: a Metric View

Realistic metric spaces (such as road/transportation networks) tend to be much more tractable then general metrics. In an attempt to formalize this intuition, Abraham et. al. (SODA 2010, JACM 2016) introduced the notion of highway dimension. A weighted graph G has highway dimension h if for every ball B of radius $\approx 4r$ there is a hitting set of size h hitting all the shortest paths of length $> r$ in B. Unfortunately, this definition fails to incorporate some very natural metric spaces such as the grid graph, and the Euclidean plane. We relax the definition of highway dimension by demanding to hit only approximate shortest paths. In addition to generalizing the original definition, this new definition also incorporates all doubling spaces (in particular the grid graph and the Euclidean plane). We then construct a PTAS for TSP under this new definition (improving a QPTAS w.r.t. the original more restrictive definition of Feldmann et. al. (SICOMP 2018)). Finally, we develop a basic metric toolkit for spaces with small highway dimension by constructing padded decompositions, sparse covers/partitions, and tree covers. An abundance of applications follow.

Andreas E. Feldmann
KAM
Charles University
feldmann.a.e@gmail.com

Arnold Filtser
Bar-Ilan University
Israel
arnold.filtser@biu.ac.il

CP27

Embedding Probability Distributions into Low Dimensional ℓ_1: Tree Ising Models Via Truncated Metrics

Given an arbitrary set of high dimensional points in ℓ_1, there are known negative results that preclude the possibility of always mapping them to a low dimensional ℓ_1 space while preserving distances with small multiplicative distortion. This is in stark contrast with dimension reduction in Euclidean space (ℓ_2) where such mappings are always possible. While the first non-trivial lower bounds for ℓ_1 dimension reduction were established almost 20 years ago, there has been limited progress in understanding what sets of points in ℓ_1 are conducive to a low-dimensional mapping. In this work, we study a new characterization of ℓ_1 metrics that are conducive to dimension reduction in ℓ_1. Our characterization focuses on metrics that are defined by the disagreement of binary variables over a probability distribution – any ℓ_1 metric can be represented in this form. We show that, for configurations of n points in ℓ_1 obtained from tree Ising models, we can reduce dimension to polylog(n) with constant distortion. In doing so, we develop technical tools for embedding truncated metrics which have been studied because of their applications in computer vision, and are objects of independent interest in metric geometry. Among other tools, we show how any ℓ_1 metric can be truncated with $O(1)$ distortion and $O(\log(n))$ blowup in dimension.

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

Spencer Compton, Chirag Pabbaraju
Stanford University
comptons@stanford.edu, cpabbara@stanford.edu

CP27

Bounding ε-Scatter Dimension Via Metric Sparsity

A recent work of Abbasi et al. [FOCS 2023] introduced the notion of ε-scatter dimension of a metric space and showed a general framework for efficient parameterized approximation schemes (so-called EPASes) for a wide range of clustering problems in classes of metric spaces that admit a bound on the ε-scatter dimension. Our main result is such a bound for metrics induced by graphs from any fixed proper minor-closed graph class. The bound is double-exponential in ε^{-1} and the Hadwiger number of the graph class and is accompanied by a nearly tight lower bound that holds even in graph classes of bounded treewidth. On the way to the main result, we introduce metric analogs of well-known graph invariants from the theory of sparsity, including generalized coloring numbers and flatness (aka uniform quasi-wideness), and show bounds for these in-

variants in proper minor-closed graph classes. Finally, we show the power of newly introduced toolbox by showing a coreset for k-Center in any proper minor-closed graph class whose size is polynomial in k (but the exponent of the polynomial depends on the graph class and ε^{-1}).

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

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

CP27

Outlier-robust Mean Estimation near the Breakdown Point via Sum-of-Squares

We revisit the problem of estimating the mean of a high-dimensional distribution in the presence of an ε-fraction of adversarial outliers. When ε is at most some sufficiently small constant, previous works can achieve optimal error rate efficiently [**?**, **?**]. As ε approaches the breakdown point $\frac{1}{2}$, all previous algorithms incur either sub-optimal error rates or exponential running time. In this paper we give a new analysis of the canonical sum-of-squares program introduced in [**?**] and show that this program efficiently achieves optimal error rate for all $\varepsilon \in [0, \frac{1}{2})$. The key ingredient for our results is a new identifiability proof for robust mean estimation that focuses on the overlap between the distributions instead of their statistical distance as in previous works. We capture this proof within the sum-of-squares proof system, thus obtaining efficient algorithms using the sum-of-squares proofs to algorithms paradigm [**?**].

Hongjie Chen, Deepak Narayanan Sridharan, David Steurer
ETH Zurich
hongjie.chen@inf.ethz.ch, dsridharan@inf.ethz.ch, dsteurer@inf.ethz.ch

CP27

The Johnson-Lindenstrauss Lemma for Clustering and Subspace Approximation: From Coresets to Dimension Reduction

We study the effect of Johnson-Lindenstrauss transforms in various projective clustering problems, generalizing results which only applied to center-based clustering [Makarychev-Makarychev-Razenshteyn '19]. We ask the general question: for a Euclidean optimization problem and an accuracy parameter $\epsilon \in (0, 1)$, what is the smallest target dimension $t \in \mathbb{N}$ such that a Johnson-Lindenstrauss transform $\mathbf{\Pi}\colon \mathbb{R}^d \to \mathbb{R}^t$ preserves the cost of the optimal solution up to a $(1+\epsilon)$-factor. We give a new technique which uses coreset constructions to analyze the effect of the Johnson-Lindenstrauss transform. Our technique, in addition applying to center-based clustering, improves on (or is the first to address) other Euclidean optimization problems, including: * For (k, z)-subspace approximation: we show that $t = \tilde{O}(zk^2/\epsilon^3)$ suffices, whereas the prior best bound, of $O(k/\epsilon^2)$, only applied to the case $z = 2$ [CEMMP15]. * For (k, z)-flat approximation: we show $t = \tilde{O}(zk^2/\epsilon^3)$ suffices, completely removing the dependence on n from the prior bound $\tilde{O}(zk^2 \log n/\epsilon^3)$ of [KR15]. * For (k, z)-line approximation: we show $t = O((k \log\log n + z + \log(1/\epsilon))/\epsilon^3)$ suffices, and ours is the first to give any dimension reduction result.

Erik Waingarten
Simons Institute
ewaingar@seas.upenn.edu

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

CP28

A Coarse Erdős-Pósa Theorem

An induced packing of cycles in a graph is a set of vertex-disjoint cycles with no edges between them. We generalise the classic Erdős-Pósa theorem to induced packings of cycles. More specifically, we show that there is a function $f(k) = O(k \log k)$ such that for every integer $k > 0$, every graph G contains either an induced packing of k cycles or a set X of at most $f(k)$ vertices such that the closed neighbourhood of X intersects all cycles in G. Our proof is constructive and yields a polynomial-time algorithm. Furthermore, we show that for every integer $d > 0$, if a graph G does not contain two cycles at distance more than d, then G contains sets $X_1, X_2 \subseteq V(G)$ with $|X_1| \leq 12(d+1)$ and $|X_2| \leq 12$ such that, after removing the ball of radius $2d$ around X_1 or the ball of radius $3d$ around X_2, the resulting graphs are forests. As a corollary, we prove that every graph with no $K_{1,t}$ induced subgraph and no induced packing of k cycles has tree-independence number at most $O(tk \log k)$, and one can construct a corresponding tree-decomposition in polynomial time. This resolves a special case of a conjecture of Dallard et al. (arXiv:2402.11222), and implies that on such graphs, many NP-hard problems are solvable in polynomial time. On the other hand, we show that the class of all graphs with no $K_{1,3}$ induced subgraph and no two cycles at distance more than 2 has unbounded tree-independence number.

Jungho Ahn
Korea Institute for Advanced Study
junghoahn@kias.re.kr

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

Tony Huynh
Université Libre de Bruxelles
tony.bourbaki@gmail.com

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

CP28

Unique-Neighbor Expanders with Better Expan-

sion for Polynomial-Sized Sets

A (d_1, d_2)-biregular bipartite graph $G = (L \cup R, E)$ is called left-(m, δ) unique-neighbor expander iff each subset S of the left vertices with $|S| \leq m$ has at least $\delta d_1 |S|$ unique-neighbors, where unique-neighbors mean vertices with exactly one neighbor in S. We can also define right/two-sided expanders similarly. In this paper, we give the following three strongly explicit constructions of unique-neighbor expanders with better unique-neighbor expansion for polynomial-sized sets, while sufficient expansion for linear-sized sets is also preserved: (1) Two-sided $(n^{1/3-\epsilon}, 1-\epsilon)$ lossless expanders for arbitrary $\epsilon > 0$ and aspect ratio. (2) Left-$(\Omega(n), 1-\epsilon)$ lossless expanders with right-$(n^{1/3-\epsilon}, \delta)$ expansion for some $\delta > 0$. (3) Two-sided-$(\Omega(n), \delta)$ unique-neighbor expanders with two-sided-$(n^{\Omega(1)}, 1/2-\epsilon)$ expansion. The second construction exhibits the first explicit family of one-sided lossless expanders with unique-neighbor expansion for polynomial-sized sets from the other side and constant aspect ratio. The third construction gives two-sided unique-neighbor expanders with additional $(1/2-\epsilon)$ unique-neighbor expansion for two-sided polynomial-sized sets, which approaches the 1/2 requirement in Lin and Hsieh (arXiv:2203.03581).

Yeyuan Chen
University of Michigan, Ann Arbor
yeyuanch@umich.edu

CP28

Weak Coloring Numbers of Minor-Closed Graph Classes

Weak coloring numbers are a family of graph parameters studied extensively in structural and algorithmic graph theory. We study the growth rate of weak coloring numbers of graphs excluding a fixed graph as a minor. Van den Heuvel et al. (European J. of Combinatorics, 2017) showed that for a fixed graph X, the maximum r-th weak coloring number of X-minor-free graphs is polynomial in r. We determine this polynomial up to a factor of $\mathcal{O}(r \log r)$ and up to a factor of $\mathcal{O}(\log r)$ in a family of graphs of bounded treewidth. Moreover, we tie the exponent of the polynomial to a structural property of X, namely, 2-treedepth. Our result can be applied to improve several well-known bounds on weak coloring numbers. For instance, we show that for planar graphs of bounded treewidth, the maximum r-th weak coloring number is in $\mathcal{O}(r^2 \log r)$, which is the best possible.

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

CP28

Planar Graphs in Blowups of Fans

We show that every n-vertex planar graph is contained in the graph obtained from a fan by blowing up each vertex by a complete graph of order $O(\sqrt{n}\log^2 n)$. Equivalently, every n-vertex planar graph G has a set X of $O(\sqrt{n}\log^2 n)$ vertices such that $G-X$ has bandwidth $O(\sqrt{n}\log^2 n)$. This result holds in the more general setting of bounded row treewidth graphs, which includes bounded genus graphs, graphs excluding a fixed apex graph as a minor, and k-planar graphs for fixed k. These results are obtained using two ingredients. The first is a new local sparsification lemma, which shows that every n-vertex planar graph G has a set of $O((n \log n)/D)$ vertices whose removal results in a graph with local density at most D. The second is a generalization of a method of Feige and Rao, that relates bandwidth and local density using volume-preserving Euclidean embeddings.

Pat Morin
Carleton University
morin@scs.carleton.ca

CP28

An Analogue of Reed's Conjecture for Digraphs

Reed in 1998 conjectured that every graph G satisfies $\chi(G) \leq \lceil \frac{\Delta(G)+1+\omega(G)}{2} \rceil$. As a partial result, he proved the existence of $\varepsilon > 0$ for which every graph G satisfies $\chi(G) \leq \lceil (1-\varepsilon)(\Delta(G)+1) + \varepsilon\omega(G) \rceil$. We propose an analogue conjecture for digraphs. Given a digraph D, we denote by $\overrightarrow{\chi}(D)$ the dichromatic number of D, which is the minimum number of colours needed to partition D into acyclic induced subdigraphs. We let $\overleftrightarrow{\omega}(D)$ denote the size of the largest biclique (a set of vertices inducing a complete digraph) of D and $\tilde{\Delta}(D) = \max_{v \in V(D)} \sqrt{d^+(v) \cdot d^-(v)}$. We conjecture that every digraph D satisfies $\overrightarrow{\chi}(D) \leq \lceil \frac{\tilde{\Delta}(D)+1+\overleftrightarrow{\omega}(D)}{2} \rceil$, which if true implies Reed's conjecture. As a partial result, we prove the existence of $\varepsilon > 0$ for which every digraph D satisfies $\overrightarrow{\chi}(D) \leq \lceil (1-\varepsilon)(\tilde{\Delta}(D)+1) + \varepsilon\overleftrightarrow{\omega}(D) \rceil$. This implies both Reed's result and an independent result of Harutyunyan and Mohar for oriented graphs.

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

Lucas Picasarri-Arrieta
National Institute of Informatics
lpicasarr@nii.ac.jp

CP29

Optimal Prefix-Suffix Queries with Applications

We revisit the classic border tree data structure [Gu, Farach, Beigel, SODA 1994] that answers the following

prefix-suffix queries on a string T of length n over an integer alphabet $\Sigma = [0, \sigma)$: for any $i, j \in [0, n)$ return all occurrences of T in $T[0\,..\,i]T[j\,..\,n-1]$. The border tree of T can be constructed in $\mathcal{O}(n)$ time and answers prefix-suffix queries in $\mathcal{O}(\log n + \mathsf{Occ})$ time, where Occ is the number of occurrences of T in $T[0\,..\,i]T[j\,..\,n-1]$. Our contribution here is the following. We present a completely different and remarkably simple data structure that can be constructed in the optimal $\mathcal{O}(n/\log_\sigma n)$ time and supports queries in the optimal $\mathcal{O}(1)$ time. Our result is based on a new structural lemma that lets us encode the output of any query in constant time and space. We also show a new direct application of our result in pattern matching on node-labeled graphs.

Solon P. Pissis
CWI Amsterdam
solon.pissis@cwi.nl

CP29

3sum in Preprocessed Universes: Faster and Simpler

We revisit the 3SUM problem in the preprocessed universes setting. We present an algorithm that, given three sets A, B, C of n integers, preprocesses them in quadratic time, so that given any subsets $A' \subseteq A$, $B' \subseteq B$, $C' \subseteq C$, it can decide if there exist $a \in A'$, $b \in B'$, $c \in C'$ with $a + b = c$ in time $O(n^{1.5} \log n)$. In contrast to both the first subquadratic $\tilde{O}(n^{13/7})$-time algorithm by Chan and Lewenstein (STOC 2015) and the current fastest $\tilde{O}(n^{11/6})$-time algorithm by Chan, Vassilevska Williams, and Xu (STOC 2023), which are based on the Balog–Szemerédi–Gowers theorem from additive combinatorics, our algorithm uses only standard 3SUM-related techniques, namely FFT and linear hashing modulo a prime. It is therefore not only faster but also simpler. Just as the two previous algorithms, ours not only decides if there is a single 3SUM solution but it actually determines for each $c \in C'$ if there is a solution containing it. We also modify the algorithm to still work in the scenario where the set C is unknown at the time of preprocessing. Finally, even though the simplest version of our algorithm is randomized, we show how to make it deterministic losing only polylogarithmic factors in the running time.

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

Shashwat Kasliwal, Pratush Sharma
IIT Delhi
shashwat.iitd.math@gmail.com, pratyushkbs@gmail.com

CP29

Pure Binary Finger Search Trees

We present dynamic binary search trees where each node only stores a value and pointers to its parent and its children. We denote such binary search trees pure binary search trees. Our structure supports finger searches in worst-case $O(\lg d)$ time, where d is the rank difference between the node given by the finger and the node found by the search. Inserting a new node with a successor or predecessor value for a node pointed to by a finger and deleting a node in the tree pointed to by a finger are supported in amortized $O(1)$ time and worst-case $O(\lg n)$ time, where n is the number of nodes in the tree. The temporary working space during the operations is $O(1)$ words. The result is obtained by an alternative representation of the red-black trees by Guibas and Sedgewick [FOCS 1978] that encodes bits of information in the tree structure, generalizing the encoding of 2-3-trees by Brown [IPL 1979], and rearranging the nodes in a red-black tree ("folding" left and right paths) such that the predecessor and successor of a node can always be found in worst-case constant time. The same time bounds can easily be obtained by, say, red-black trees and AVL trees augmented with pointers to the predecessor and successor of each node. The novelty of our result is that we store no extra information than the binary tree structure. The structure can be represented by two pointers per value, i.e., the same representation as a doubly linked list.

Gerth S. Brodal
Department of Computer Science, Aarhus University
gerth@cs.au.dk

Casper Rysgaard
Aarhus University
Department of Computer Science
rysgaard@cs.au.dk

CP29

A Simple Algorithm for Dynamic Carpooling with Recourse

We give an algorithm for the fully-dynamic carpooling problem with recourse: Edges arrive and depart online from a graph G with n nodes according to an adaptive adversary. Our goal is to maintain an orientation H of G that keeps the discrepancy, defined as $\max_{v \in V} |\deg^+_H(v) - \deg^-_H(v)|$, small at all times. We present a simple algorithm and analysis for this problem with recourse based on cycles that simplifies and improves on a result of Gupta et al. [SODA '22].

Yuval Efron, Shyamal Patel, Cliff Stein
Columbia University
ye2210@columbia.edu, shyamalpatelb@gmail.com, cliff@ieor.columbia.edu

CP29

Multi-Dimensional Approximate Counting

In this work, we present a simple and optimal d-dimensional counter with Euclidean relative error σ, i.e., $\mathbb{E}|\hat{x} - x|^2 < \sigma^2 |x|^2$ where $|x| = \sqrt{\sum_{j=1}^d x_j^2}$, with a matching lower bound. We prove that on the one hand there exists a $(\log_2 \log_2 n + O(d \log_2 \sigma^{-1}))$-bit d-dimensional counter with relative error σ. On the other hand, any d-dimensional counter with relative error σ takes at least $(\log_2 \log_2 n + \Omega(d \log_2 \sigma^{-1}))$ bits of space. The upper and lower bounds are proved with ideas that are strikingly simple. The upper bound is constructed with a certain variable-length integer encoding and the lower bound is derived from a straight-

forward volumetric estimation of sphere covering.

Dingyu Wang
University of Michigan
wangdy@umich.edu

CP30

Understanding Memory-Regret Trade-Off for Streaming Stochastic Multi-Armed Bandits

We study the stochastic multi-armed bandit problem in the P-pass streaming model. In this problem, the n arms are present in a stream and at most $m < n$ arms and their statistics can be stored in the memory. We give a complete characterization of the optimal regret in terms of m, n and P. Specifically, we design an algorithm with $\tilde{O}\left((n-m)^{1+\frac{2^P-2}{2^{P+1}-1}} n^{\frac{2-2^{P+1}}{2^{P+1}-1}} T^{\frac{2^P}{2^{P+1}-1}}\right)$ regret and complement it with an $\tilde{\Omega}\left((n-m)^{1+\frac{2^P-2}{2^{P+1}-1}} n^{\frac{2-2^{P+1}}{2^{P+1}-1}} T^{\frac{2^P}{2^{P+1}-1}}\right)$ lower bound when the number of rounds T is sufficiently large. Our results are tight up to a logarithmic factor in n and P.

Yuchen He, Zichun Ye
Shanghai Jiao Tong University
yuchen_he@sjtu.edu.cn, alchemist@sjtu.edu.cn

Chihao Zhang
John Hopcroft Center for Computer Science
Shanghai Jiao Tong University
chihao@sjtu.edu.cn

CP30

Streaming Algorithms Via Local Algorithms for Maximum Directed Cut

We explore the use of local algorithms in the design of streaming algorithms for the Maximum Directed Cut problem. Specifically, building on the local algorithm of Buchbinder, Feldman, Seffi, and Schwartz [SIAM Journal on Computing, 2015] and Censor-Hillel, Levy, and Shachnai [ALGOSENSORS, 2017], we develop streaming algorithms for both adversarially and randomly ordered streams that approximate the value of maximum directed cut in bounded-degree graphs. In n-vertex graphs, for adversarially ordered streams, our algorithm uses $O(n^{1-\Omega(1)})$ (sub-linear) space and for randomly ordered streams, our algorithm uses logarithmic space. Moreover, both algorithms require only one pass over the input stream. With a constant number of passes, we give a logarithmic-space algorithm which works even on graphs with unbounded degree on adversarially ordered streams. Our algorithms achieve any fixed constant approximation factor less than $\frac{1}{2}$. In the single-pass setting, this is tight: known lower bounds show that obtaining any constant approximation factor greater than $\frac{1}{2}$ is impossible without using linear space in adversarially ordered streams and $\Omega(\sqrt{n})$ space in randomly ordered streams, even on bounded degree graphs.

Santhoshini Velusamy
Toyota Technological Institute at Chicago
santhoshinivelusamy@gmail.com

Raghuvansh Saxena
Microsoft Research
raghuvansh.saxena@gmail.com

Noah Singer
Carnegie Mellon University
ngsinger@cs.cmu.edu

Madhu Sudan
Harvard
madhu@cs.harvard.edu

CP30

Universal Perfect Samplers for Incremental Streams

If $G : \mathbb{R}_+ \to \mathbb{R}_+$, the G-moment of a vector $\mathbf{x} \in \mathbb{R}_+^n$ is $G(\mathbf{x}) = \sum_{v\in[n]} G(\mathbf{x}(v))$ and the G-sampling problem is to select an index $v_* \in [n]$ according to its contribution to the G-moment, i.e., such that $\mathbb{P}(v_* = v) = G(\mathbf{x}(v))/G(\mathbf{x})$. In this paper we focus on the exact G-sampling problem, where G is selected from the following class of functions.

$$\mathcal{G} = \left\{ G(z) = c1(z > 0) + \gamma_0 z + \int_0^\infty (1 - e^{-rz})\,\nu(dr) \;\middle|\; c, \gamma_0 \geq 0, \nu \text{ is no} \right.$$

We develop G-samplers for a vector $\mathbf{x} \in \mathbb{R}_+^n$ that is presented as an incremental stream of positive updates. In particular, for any $G \in \mathcal{G}$, we give a very simple G-sampler that uses 2 words of memory and stores at all times a $v_* \in [n]$, such that $\mathbb{P}(v_* = v)$ is exactly $G(\mathbf{x}(v))/G(\mathbf{x})$. In addition, we give a "universal' $\mathcal{G}$-sampler that uses $O(\log n)$ words of memory w.h.p., and given any $G \in \mathcal{G}$ at query time, produces an exact G-sample. Our sampling framework is simple and versatile, and can easily be generalized to sampling from more complex objects like graphs and hypergraphs.

Dingyu Wang
University of Michigan
wangdy@umich.edu

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

CP30

Streaming and Communication Complexity of Load-Balancing Via Matching Contractors

In the load-balancing problem, we have an n-vertex bipartite graph between a set of clients and servers. The goal is to find an assignment of all clients to the servers, while minimizing the maximum number of clients assigned to any given server. Motivated by the streaming complexity of this problem, we study load-balancing in the one-way (two-party) communication model. We show that settling this communication complexity problem is equivalent to a natural sparsification problem for vertex-expansion. We then prove a dual interpretation of this sparsifier, showing that the minimum density of a sparsifier is effectively equivalent to the maximum density for a new extremal graph family, which we call Matching-Contractors; these graphs

are intimately connected to, and also generalize, the well-known Ruza-Szemeredi graphs. Thus, the one-way communication complexity of load-balancing can be reduced to a purely graph theoretic question: what is the maximum density of a matching-contractor on n vertices? Finally, we give a novel combinatorial construction of some-what dense matching-contractors, which implies a strong one-way communication lower bound for load-balancing: any one-way protocol with $\tilde{O}(n)$ communication cannot achieve a better than $n^{\frac{1}{4}-o(1)}$-approximation. Our result also implies the first non-trivial lower bounds for semi-streaming load-balancing in the edge-arrival model, ruling out $n^{\frac{1}{4}-o(1)}$-approximation in a single-pass.

Sepehr Assadi
Rutgers University, U.S.
sepehr@assadi.info

Aaron Bernstein, Zach Langley
Rutgers University
bernstei@gmail.com, zblangley@gmail.com

Lap Chi Lau, Robert Wang
University of Waterloo
lapchi@uwaterloo.ca, r585wang@uwaterloo.ca

CP30

Near-Optimal Relative Error Streaming Quantile Estimation Via Elastic Compactors

Given a stream of elements $x_1, x_2, \ldots, x_n$ and a query x, a relative-error quantile estimation algorithm can estimate the rank of x with respect to the stream, up to a multiplicative $\pm\epsilon \cdot \text{rank}(x)$ error. Notably, this requires the sketch to obtain more precise estimates for the ranks of elements on the tails of the distribution, as compared to the additive $\pm\epsilon n$ error regime. Previously, the best known algorithms for relative error achieved space $\tilde{O}(\epsilon^{-1}\log^{1.5}(\epsilon n))$ (Cormode, Karnin, Liberty, Thaler, Veselỳ, 2021) and $\tilde{O}(\epsilon^{-2}\log(\epsilon n))$ (Zhang, Lin, Xu, Korn, Wang, 2006). In this work, we present a nearly-optimal streaming algorithm for the relative-error quantile estimation problem using $\tilde{O}(\epsilon^{-1}\log(\epsilon n))$ space. To surpass the $\Omega(\epsilon^{-1}\log^{1.5}(\epsilon n))$ barrier of the previous approach, our algorithm crucially relies on a new data structure, called an elastic compactor, which can be dynamically resized over the course of the stream. Interestingly, we design a space allocation scheme which adaptively allocates space to each compactor based on the "hardness" of the input stream. This approach allows us to avoid using the maximal space simultaneously for every compactor and facilitates the improvement in the total space complexity.

Hongxun Wu
EECS
UC Berkeley
wuhx@berkeley.edu

Elena Gribelyuk, Pachara Sawettamalya, Huacheng Yu
Princeton University
eg5539@princeton.edu, ps3122@princeton.edu, yuhch123@gmail.com

CP31

Computing the Second and Third Systoles of a Combinatorial Surface

Given a weighted, undirected graph G cellularly embedded on a topological surface S, we describe algorithms to compute the second shortest and third shortest closed walks of G that are neither homotopically trivial in S nor homotopic to the shortest non-trivial closed walk or to each other. Our algorithms run in $O(n^2 \log n)$ time for the second shortest walk and in $O(n^3)$ time for the third shortest walk. We also show how to reduce the running time for the second shortest homotopically non-trivial closed walk to $O(n \log n)$ when both the genus and the number of boundaries are fixed. Our algorithms rely on a careful analysis of the configurations of the first three shortest homotopically non-trivial curves in S. As an intermediate step, we also describe how to compute a shortest essential arc between one pair of vertices or between all pairs of vertices of a given boundary component of S in $O(n^2)$ time or $O(n^3)$ time, respectively.

Matthijs Ebbens
University of Cologne
ymebbens@gmail.com

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

CP31

A Fast Algorithm for Computing Zigzag Representatives

Zigzag filtrations of simplicial complexes generalize the usual filtrations by allowing simplex deletions in addition to simplex insertions. The barcodes computed from zigzag filtrations encode the evolution of homological features. Although one can locate a particular feature at any index in the filtration using existing algorithms, the resulting representatives may not be compatible with the zigzag: a representative cycle at one index may not map into a representative cycle at its neighbor. For this, one needs to compute compatible representative cycles along each bar in the barcode. Even though it is known that the barcode for a zigzag filtration with m insertions and deletions can be computed in $O(m^\omega)$ time, it is not known how to compute the compatible representatives so efficiently. For a non-zigzag filtration, the classical matrix-based algorithm provides representatives in $O(m^3)$ time, which can be improved to $O(m^\omega)$. However, no known algorithm for zigzag filtrations computes the representatives with the $O(m^3)$ time bound. We present an $O(m^2 n)$ time algorithm for this problem, where $n \leq m$ is the size of the largest complex in the filtration.

Tamal K. Dey
Department of Computer Science
Purdue University
tamaldey@purdue.edu

Tao Hou
Department of Computer Science
University of Oregon
taohou@uoregon.edu

Dmitriy Morozov
Lawrence Berkeley National Laboratory
dmitriy@mrzv.org

CP31

Partitioning a Polygon Into Small Pieces

We study the problem of partitioning a given simple polygon P into a minimum number of connected polygonal pieces, each of bounded size. We describe a general technique for constructing such partitions that works for several notions of 'bounded size,' namely that each piece must be contained in an axis-aligned or arbitrarily rotated unit square or a unit disk, or that each piece has bounded perimeter, straight-line diameter or geodesic diameter. The problems are motivated by practical settings in manufacturing, finite element analysis, collision detection, vehicle routing, shipping and laser capture microdissection. The version where each piece should be contained in an axis-aligned unit square is already known to be NP-hard [Abrahamsen and Stade, FOCS, 2024], and the other versions seem no easier. Our main result is to develop constant-factor approximation algorithms, which means that the number of pieces in the produced partition is at most a constant factor larger than the cardinality of an optimal partition. Existing algorithms [Damian and Pemmaraju, Algorithmica, 2004] do not allow Steiner points, which means that all corners of the produced pieces must also be corners of P. This has the disappointing consequence that a partition often does not exist, whereas our algorithms always produce meaningful partitions.

Mikkel Abrahamsen, Nichlas Rasmussen, Mads Jensen
University of Copenhagen
miab@di.ku.dk, nichlas.rasmussen@gmail.com, mvje@di.ku.dk

CP31

Minimum Convex Hull and Maximum Overlap of Two Convex Polytopes

We study the problem of minimizing the convex hull of two convex polytopes with n vertices in total under translation in d-dimensional space $\mathbb{R}^d$ for any fixed dimension $d \geq 2$. For $d = 2$, we present a deterministic $O(n)$-time algorithm returning a translation minimizing the area of the convex hull, improving upon the previously best $O(n \log n)$-time algorithm. Our algorithm returns the smallest area of convex hulls under translation in the same time, and thus it is optimal. For $d \geq 3$, we present a deterministic algorithm with running time $O(n^{(d+1)/2})$ for odd d and $O(n^{d/2} \log^d n)$ for even d. This improves substantially upon the previously best algorithm by a factor at least $n^{(d-1)/2} \log n$. We also consider the variant that two input polytopes are restricted to remain disjoint, and present a deterministic algorithm with running time $O(n^{d+1})$ for odd d and $O(n^d \log^{d-1} n)$ for even d. This improves substantially upon the previously best algorithm for $d > 3$ by factor $n^{O(d^2)}$. We also study the problem of maximizing the overlap of two convex polytopes under translation in d-dimensional space $\mathbb{R}^d$ for $d \geq 3$. We give an $O(n^{2\lfloor d/2 \rfloor})$-time algorithm, improving substantially upon the previously best algorithm by a factor at least $n^{1-3/d} \log^{d+1} n$.

Mook Kwon Jung, Seokyun Kang, Hee-Kap Ahn
POSTECH
jmg1032@postech.ac.kr, sykang0330@postech.ac.kr, heekap@postech.ac.kr

CP31

An Efficient Regularity Lemma for Semi-Algebraic Hypergraphs

We establish stronger and more efficient regularity theorems for k-uniform hypergraphs $H = (P, E)$, where P is a point set in R^d, and the edge set E is determined by a semi-algebraic relation of bounded description complexity. For $0 < \epsilon \leq 1$, we show that one can construct in expected $O(n \log 1/\epsilon)$ time, an equitable partition $P = U_1 \uplus \ldots \uplus U_K$ into $K = O(1/\epsilon^{d+1+\delta})$ subsets, for any $0 < \delta$, so that all but ϵ-fraction of the k-tuples $U_{i_1}, \ldots, U_{i_k}$ are homogeneous: we have that either $U_{i_1} \times \ldots \times U_{i_k} \subseteq E$ or $(U_{i_1} \times \ldots \times U_{i_k}) \cap E = \emptyset$. The best previously known partition, due to Fox, Pach and Suk requires $\Omega(n^{k-1}/\epsilon^c)$ time and yields $K = 1/\epsilon^c$ parts, where the constant c is not stated explicitly. In contrast to the previous regularity lemmas, our partition of P does not depend on the set E, provided its semi-algebraic description complexity does not exceed a certain constant. As a by-product, we show that in any k-partite k-uniform hypergraph $(P_1 \uplus \ldots \uplus P_k, E)$ of bounded semi-algebraic description complexity in R^d and with $|E| \geq \epsilon \prod_{i=1}^k |P_i|$, one can find, in expected time $O(\sum_{i=1}^k (|P_i| + 1/\epsilon) \log(1/\epsilon))$, subsets $Q_i \subseteq P_i$ of size $|Q_i| \geq |P_i|/\epsilon^{d+1+\delta}$, so that $Q_1 \times \ldots \times Q_k \subseteq E$.

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

CP32

Maximum Span Hypothesis: A Potentially Weaker Assumption Than Gap-ETH for Parameterized Complexity

The Gap Exponential Time Hypothesis rules out FPT algorithms providing (nearly) tight inapproximability results for a host of fundamental problems in parameterized complexity. One of the downsides of working under Gap-ETH is that the assumption is not inherently in the parameterized complexity world, and therefore one of the main research directions is to replace Gap-ETH with weaker assumptions. In this paper, we propose a hypothesis called the Maximum Span Hypothesis (MSH), which roughly asserts that given a collection of n vectors in $\mathbb{F}^{\text{poly}(k) \cdot \log n}$ such that there is a k-dimensional subspace containing $2^{\Omega(k)}$ input vectors, the goal of finding poly(k) input vectors which are contained in some k-dimensional subspace is W[1]-hard. Assuming MSH, we obtain near optimal inapproximability ratio for the k-clique problem and polynomial inapproximability ratio for the 2-CSP problem (on k variables and alphabet size n). Assuming a strengthening of MSH

with additional completeness guarantees, we are able to obtain near optimal inapproximability ratio for the k-biclique problem and some constant inapproximability ratio for the Densest k-subgraph problem. Finally, we prove that Gap-ETH implies a mild version of MSH.

Karthik C. S.
Rutgers University
karthik.cs@rutgers.edu

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

CP32

Hardness of Counting Small Induced Subgraphs: Fourier vs. Sylow

This is a merged presentation of the following two papers: "From Graph Properties to Graph Parameters: Tight Bounds for Counting on Small Subgraphs" and "Counting Small Induced Subgraphs: Hardness via Fourier Analysis" Given a graph property and a number k, the problem #IndSub(Φ) asks to count the induced k-vertex subgraphs satisfying Φ in an input graph G. Since this problem generalizes many counting problems, there has been extensive work on determining the parameterized complexity of #IndSub(Φ) for various Φ. This merged talk presents two papers dealing with this in different ways. The first paper, by Dring, Marx, and Wellnitz, extends previous results by showing that #IndSub(Φ) cannot be solved in $n^{o(k)}$ time for any nontrivial edge-monotone graph property Φ, assuming ETH. To this end, they analyze so-called alternating enumerators by studying the fixed points of Sylow group actions. The second paper, by Curticapean and Neuen, shows the same tight lower bound, but by introducing new techniques. More precisely, they observe that the problem can be translated into the world of Boolean functions. Here, many well-known techniques from Fourier analysis allow them to analyze the alternating enumerator and show further results. For instance, they show new hardness results whenever at most $(1-\varepsilon)^{\binom{k}{2}}$ graphs on k vertices satisfy Φ. They also improve the lower bounds for various other graph properties. Lastly, both papers generalize #IndSub(Φ) to graph parameters and present nontrivial, edge-monotone graph parameters for which #IndSub(Φ) is FPT.

Philip Wellnitz
Max Planck Institute for Informatics,
Saarland Informatics Campus, Germany
wellnitz@nii.ac.jp

Simon Döring
Max Planck Institute for Informatics
Saarbrücken, Germany
sdoering@mpi-inf.mpg.de

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

Radu Curticapean
University of Regensburg and IT University of Copenhagen
radu.curticapean@ur.de

Daniel Neuen
CISPA Helmholtz Center for Information Security
dneuen@mpi-inf.mpg.de

CP32

Parameterizing the quantification of CMSO: model checking on minor-closed graph classes

Given a graph G and a vertex set X, the annotated treewidth $\mathsf{tw}(G, X)$ of X in G is the maximum treewidth of an X-rooted minor of G, i.e., a minor H where the model of each vertex of H contains some vertex of X. That way, $\mathsf{tw}(G, X)$ can be seen as a measure of the contribution of X to the tree-decomposability of G. We introduce the logic CMSO/tw as the fragment of monadic second-order logic on graphs obtained by restricting set quantification to sets of bounded annotated treewidth. We prove the following Algorithmic Meta-Theorem (AMT): for every non-trivial minor-closed graph class, model checking for CMSO/tw formulas can be done in quadratic time. Our proof works for the more general CMSO/tw+dp logic, that is CMSO/tw enhanced by disjoint-path predicates. Our AMT can be seen as an extension of Courcelle's theorem to minor-closed graph classes where the bounded-treewidth condition in the input graph is replaced by the bounded-treewidth quantification in the formulas. Our results yield, as special cases, all known AMTs whose combinatorial restriction is non-trivial minor-closedness.

Ignasi Sau
CNRS, LIRMM, Montpellier
ignasi.sau@lirmm.fr

Giannos Stamoulis
University of Montpellier and the National Center for Scientific Research
giannos.stamoulis@mimuw.edu.pl

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

CP32

Finding irrelevant vertices in linear time on bounded-genus graphs

The irrelevant vertex technique provides a powerful tool for the design of parameterized algorithms for a wide variety of problems on graphs. A common characteristic of these problems, permitting the application of this technique on surface-embedded graphs, is the fact that every graph of large enough treewidth contains a vertex that is irrelevant, in the sense that its removal yields an equivalent instance of the problem. The straightforward application of this technique yields algorithms with running time that is quadratic in the size of the input graph. This running time is due to the fact that it takes linear time to detect one irrelevant vertex and the total number of irrelevant vertices to be detected is linear as well. Using advanced

techniques, sub-quadratic algorithms have been designed for particular problems, even in general graphs. However, designing a general framework for linear-time algorithms has been open, even for the bounded-genus case. In this paper we introduce a general framework that enables finding in linear time an entire set of irrelevant vertices whose removal yields a bounded-treewidth graph, provided that the input graph has bounded genus. Our method is applicable to a wide variety of known graph containment or graph modification problems where the irrelevant vertex technique applies. Examples include the (Induced) Minor Folio problem, the (Induced) Disjoint Paths problem, and the $\mathcal{F}$-Minor-Deletion problem.

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

Stavros Kolliopoulos
Department of Informatics and Telecommunications, National and Kapodistrian University of Athens, Greece
sgk@di.uoa.gr

Giannos Stamoulis
University of Montpellier and the National Center for Scientific Research
giannos.stamoulis@mimuw.edu.pl

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

CP32

Losing Treewidth In The Presence Of Weights

In the Weighted Treewidth-η Deletion problem we are given a node-weighted graph G and we look for a vertex subset X of minimum weight such that the treewidth of $G - X$ is at most η. We show that Weighted Treewidth-η Deletion admits a randomized polynomial-time constant-factor approximation algorithm for every fixed η. Our algorithm also works for the more general Weighted Planar F-M-Deletion problem. This work extends the results for unweighted graphs by [Fomin, Lokshtanov, Misra, Saurabh; FOCS '12] and answers a question posed by [Agrawal, Lokshtanov, Misra, Saurabh, Zehavi; APPROX/RANDOM '18] and [Kim, Lee, Thilikos; APPROX/RANDOM '21]. The presented algorithm is based on a novel technique of random sampling of so-called protrusions.

Michal Wlodarczyk
University of Warsaw
m.wlodarczyk@mimuw.edu.pl

CP33

A Simple Lower Bound for Set Agreement in Dynamic Networks

Given a positive integer k, k-set agreement is the distributed task in which each process $i \in [n]$ in a group of n processing nodes starts with an input value $x_i \in \{0, \ldots, k\}$, and must output a value y_i such that (1) for every $i \in [n]$, y_i is the input value of some process, and (2) $|\{y_i : i \in [n]\}| \leq k$. The case $k = 1$ correspond to (binary) consensus, arguably the most studied problem in distributed computing. While lower bounds for consensus have been obtained for most of the standard distributed computing models, the design of lower bounds for k-set agreement with $k > 1$ is notoriously known to be much more difficult, and remains open for many models. The main techniques for designing lower bounds for k-set agreement with $k > 1$ use tools from algebraic topology. The algebraic topology tools are difficult to manipulate, and require a lot of care for avoiding mistakes. This difficulty increases when the communications are mediated by a network of arbitrary structure. Recently, the KNOWALL model has been specifically designed as a first attempt to understand the LOCAL model through the lens of algebraic topology, and Castaeda et al.(2021) have designed lower bounds for k-set agreement in the KNOWALL model, with applications to dynamic networks. In this work, we re-prove the same lower bound for k-set agreement in the KNOWALL model. This new proof stands out in its simplicity, which makes it accessible to a broader audience.

Pierre Fraigniaud
CNRS
University of Paris Sud
pierre@lri.fr

Minh Hang Nguyen
IRIF, CNRS, Université Paris Cité,France
mhnguyen@irif.fr

Ami Paz
LISN, CNRS and Université Paris-Saclay, France
amipaz.cs@gmail.com

CP33

Trading Prophets: How to Trade Multiple Stocks Optimally

In the (single stock) *trading prophet* problem formulated by Correa et al. [2023], an online algorithm observes a sequence of stock prices and can either buy the stock by paying the current price or sell the currently-held stock for profit. The goal is to maximize overall profit. Correa et al. showed an optimal competitive ratio of $\frac{1}{2}$ for this problem when the stock prices are identically and independently distributed. This work generalizes the model and results of Correa et al. by allowing the algorithm to trade multiple stocks. First, we formulate the (k, ℓ, ℓ')-Trading Prophet Problem, which involves k stocks, allowing the online algorithm to hold up to ℓ stocks, while the offline algorithm can hold at most $\ell' \leq \ell$ stocks. Assuming the prices of different stocks are independent, we show that the optimal competitive ratio is $\min\left\{\frac{1}{2}, \frac{\ell}{k}\right\}$. We further introduce the more general $\mathcal{M}$-Trading Prophet Problem over a matroid $\mathcal{M}$ on the k stocks, where stock prices are possibly correlated but independent over time. The algorithm can only hold an independent subset of stocks. We establish a tight competitive ratio bound of $\frac{1}{1+d}$, where d represents the *density* of the matroid.

Surbhi Rajput
Indian Institute of Technology Delhi
surbhi15121999@gmail.com

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

Rohit Vaish
Indian Institute of Technology Delhi
rvaish@iitd.ac.in

CP33

A Simple Parallel Algorithm with Near-Linear Work for Negative-Weight Single-Source Shortest Path

We give the first parallel algorithm with optimal $\tilde{O}(m)$ work for the classical problem of computing Single-Source Shortest Paths in general graphs with negative-weight edges. In graphs without negative edges, Dijkstra's algorithm solves the Single-Source Shortest Paths (SSSP) problem with optimal $\tilde{O}(m)$ work, but is inherently sequential. A recent breakthrough by [Bernstein, Nanongkai, Wulff-Nilsen; FOCS '22] achieves the same for general graphs. Parallel shortest path algorithms are more difficult and have been intensely studied for decades. Only recently, multiple lines of research culminated in parallel algorithms with optimal work $\tilde{O}(m)$ for various restricted settings, such as approximate or exact algorithms for directed or undirected graphs without negative edges. For general graphs, the best known algorithm by [Ashvinkumar, Bernstein, Cao, Grunau, Haeupler, Jiang, Nanongkai, Su; ESA '24] still requires $m^{1+o(1)}$ work. This paper presents a randomized parallel algorithm for SSSP in general graphs with near-linear work $\tilde{O}(m)$ and state-of-the-art span $n^{1/2+o(1)}$. We follow a novel bottom-up approach leading to a particularly clean and simple algorithm. Our algorithm can be seen as a *near-optimal parallel black-box reduction* from SSSP in general graphs to graphs without negative edges. In contrast to prior works, the reduction in this paper is both parallel and with only polylogarithmic overhead in work and span.

Nick Fischer
INSAIT, University of Sofia
nick.fischer@insait.ai

Bernhard Haeupler
ETH Zurich
bernhard.haeupler@inf.ethz.ch

Rustam Latypov
Aalto University
rustam.latypov@aalto.fi

Antti Roeyskoe, Aurelio L. Sulser
ETH Zurich
antti.roeyskoe@inf.ethz.ch, aurelio.sulser@inf.ethz.ch

CP33

Bidirectional Dijkstra's Algorithm Is Instance-Optimal

While Dijkstra's algorithm has near-optimal time complexity for the problem of finding the shortest *st*-path, in practice, other algorithms are often superior on huge graphs. A prominent such example is the *bidirectional search*, which executes Dijkstra's algorithm from both endpoints in parallel and stops when these executions meet. In this paper, we give a strong theoretical justification for the use of such bidirectional search algorithms. We prove that for weighted multigraphs, both directed and undirected, a careful implementation of bidirectional search is instance-optimal with respect to the number of edges it explores. That is, we prove that no correct algorithm can outperform our implementation of bidirectional search on *any single instance* by more than a constant factor. For unweighted graphs, we show that bidirectional search is instace-optimal up to a factor of $O(\Delta)$ where Δ is the maximum degree of the graph. We also show that this is the best possible.

Bernhard Haeupler
Sofia University "St. Kliment Ohridski"
ETH Zurich
bernhard.haeupler@inf.ethz.ch

Richard Hladík
ETH Zurich
& INSAIT, Sofia University
rihl@uralyx.cz

Václav Rozhon
ETH
vaclavrozhon@gmail.com

Robert Tarjan
Princeton University
ret@princeton.edu

Jakub Tetek
INSAIT, Sofia University
j.tetek@gmail.com

CP33

Efficient Matroid Intersection Via a Batch-Update Auction Algorithm

Given two matroids $\mathcal{M}_1$ and $\mathcal{M}_2$ over the same n-element ground set, the matroid intersection problem is to find a largest common independent set, whose size we denote by r. We present a simple and generic auction algorithm that reduces $(1-\varepsilon)$-approximate matroid intersection to roughly $1/\varepsilon^2$ rounds of the easier problem of finding a maximum-weight basis of a single matroid. Plugging in known primitives for this subproblem, we obtain both simpler and improved algorithms in two models of computation, including: * The first near-linear time/independence-query $(1-\varepsilon)$-approximation algorithm for matroid intersection. Our randomized algorithm uses $\tilde{O}(n/\varepsilon + r/\varepsilon^5)$ independence queries, improving upon the previous $\tilde{O}(n/\varepsilon + r\sqrt{r}/\varepsilon^3)$ bound of Quanrud (2024). * The first sublinear exact parallel algorithms for weighted matroid intersection, using $O(n^{2/3})$ rounds of rank queries or $O(n^{5/6})$ rounds of independence queries. For the unweighted case, our results improve upon the previous $O(n^{3/4})$-round rank-query and $O(n^{7/8})$-round independence-query algorithms of Blikstad (2022).

Joakim Blikstad

KTH Royal Institute of Technology
blikstad@kth.se

Ta-Wei Tu
Stanford University
taweitu@stanford.edu

CP34

Simple Approximation Algorithms for Polyamorous Scheduling

In Polyamorous Scheduling, we are given an edge-weighted graph and must find a periodic schedule of matchings in this graph which minimizes the maximal weighted waiting time between consecutive occurrences of the same edge. This NP-hard problem generalises Bamboo Garden Trimming and is motivated by the need to find schedules of pairwise meetings in a complex social group. We present two different analyses of an approximation algorithm based on the Reduce-Fastest heuristic, from which we obtain first a 6-approximation and then a 5.24-approximation for Polyamorous Scheduling. We also strengthen the extant proof that there is no polynomial-time $(1+\delta)$-approximation algorithm for the Optimisation Polyamorous Scheduling problem for any $\delta < \frac{1}{12}$ unless P = NP to the bipartite case. The decision version of Polyamorous Scheduling has a notion of density, similar to that of Pinwheel Scheduling, where problems with density below the threshold are guaranteed to admit a schedule (cf. the recently proven 5/6 conjecture, Kawamura, STOC 2024). We establish the existence of a similar threshold for Polyamorous Scheduling and give the first non-trivial bounds on the poly density threshold.

Yuriy Biktairov
University of Southern California
biktairo@usc.edu

Leszek Gasieniec
University of Liverpool
lechu@liverpool.ac.uk

Wanchote Po Jiamjitrak
University of Helsinki
wanchote.jiamjitrak@helsinki.fi

N Namrata, Benjamin Smith
University of Liverpool
n.namrata@liverpool.ac.uk, b.m.smith@liverpool.ac.uk

Sebastian Wild
University of Marburg
wild@informatik.uni-marburg.de

CP34

Simple Length-Constrained Minimum Spanning Trees

In the length-constrained minimum spanning tree (MST) problem, we are given an n-node edge-weighted graph G and a length constraint $h \geq 1$. Our goal is to find a spanning tree of G whose diameter is at most h with minimum weight. Prior work of Marathe et al. gave a poly-time algorithm which repeatedly computes maximum cardinality matchings of minimum weight to output a spanning tree whose weight is $O(\log n)$-approximate with diameter $O(\log n) \cdot h$. In this work, we show that a simple random sampling approach recovers the results of Marathe et al.—no computation of min-weight max-matchings needed! Furthermore, the simplicity of our approach allows us to tradeoff between the approximation factor and the loss in diameter: we show that for any $\epsilon \geq 1/\operatorname{poly}(n)$, one can output a spanning tree whose weight is $O(n^{\epsilon}/\epsilon)$-approximate with diameter $O(1/\epsilon) \cdot h$ with high probability in poly-time. This immediately gives the first poly-time $\operatorname{poly}(\log n)$-approximation for length-constrained MST whose loss in diameter is $o(\log n)$.

Ellis Hershkowitz, Richard Huang
Brown University
delhersh@gmail.com, rzhuang351@gmail.com

CP34

Simple Combinatorial Construction of the $k^{o(1)}$-Lower Bound for Approximating the Parameterized k-Clique

In the parameterized k-clique problem, or k-Clique for short, we are given a graph G and a parameter $k \geq 1$. The goal is to decide whether there exist k vertices in G that induce a complete subgraph (i.e., a k-clique). In a breakthrough work of Bingkai Lin [STOC'21], any constant-factor approximation of k-Clique is shown to be W[1]-hard, and subsequently, the inapproximation ratio is improved to $k^{o(1)}$ in the work of Karthik C.S. and Khot [CCC'22], and independently in [Lin, Ren, Sun Wang; ICALP'22] (under the apparently stronger complexity assumption ETH). All the work along this line follows the framework developed by Lin, which starts from the k-vector-sum problem and requires some involved algebraic techniques. This paper presents an alternative framework for proving the W[1]-hardness of the $k^{o(1)}$-FPT-approximation of k-Clique. Using this framework, we obtain a gap-producing self-reduction of k-Clique without any intermediate algebraic problem. More precisely, we reduce from $(k, k-1)$-Gap Clique to (q^k, q^{k-1})-Gap Clique, for any function q depending only on the parameter k, thus implying the $k^{o(1)}$-inapproximability result when q is sufficiently large. Our proof is relatively simple and mostly combinatorial. At the core of our construction is a novel encoding of k-element subset stemmed from the theory of *network coding* and a *(linear) Sidon set* representation of a graph.

Yijia Chen
Shanghai Jiao Tong University
yijia.chen@cs.sjtu.edu.cn

Yi Feng
Shanghai University of Finance and Economics
2021310186@live.sufe.edu.cn

Bundit Laekhanukit
Independent Researcher
lbundit@gmail.com

Yanlin Liu
Ocean University of China

liuyanlin@ouc.edu.cn

CP34

Validating a Ptas for Triangle-Free 2-Matching

A triangle-free (simple) 2-matching is an edge set that has at most 2 edges incident to each vertex and contains no cycle of length 3. For the problem of finding a maximum cardinality triangle-free 2-matching in a given graph, a complicated exact algorithm was proposed by Hartvigsen. Recently, a simple PTAS using local search was presented by Bosch-Calvo, Grandoni, and Ameli, but its validity proof is not easy. In this paper, we show a natural and simple decomposition theorem for triangle-free 2-matchings, which leads to a simpler validity proof of the PTAS for the problem.

Yusuke Kobayashi, Takashi Noguchi
Kyoto University
yusuke@kurims.kyoto-u.ac.jp, tnoguchi@kurims.kyoto-u.ac.jp

CP34

Spectral Sparsification by Deterministic Discrepancy Walk

Spectral sparsification and discrepancy minimization are two well-studied areas that are closely related. Building on recent connections between these two areas, we generalize the "deterministic discrepancy walk' framework by Pesenti and Vladu [SODA 23] for vector discrepancy to matrix discrepancy, and use it to give a simpler proof of the matrix partial coloring theorem of Reis and Rothvoss [SODA 20]. Moreover, we show that this matrix discrepancy framework provides a unified approach for various spectral sparsification problems, from stronger notions including unit-circle approximation and singular-value approximation to weaker notions including graphical spectral sketching and effective resistance sparsification. In all of these applications, our framework produces improved results with a simpler and deterministic analysis.

Lap Chi Lau, Robert Wang
University of Waterloo
lapchi@uwaterloo.ca, r585wang@uwaterloo.ca

Hong Zhou
Fuzhou University
hong.zhou@fzu.edu.cn

CP35

Stronger Adversaries Grow Cheaper Forests: Online Node-Weighted Steiner Problems

We propose an $O(\log k \log n)$-competitive randomized algorithm for online node-weighted Steiner forest. This is essentially optimal and significantly improves over the previous bound of $O(\log^2 k \log n)$ by Hajiaghayi et al. (2017). In fact, our result extends to the more general prize-collecting setting, improving over previous works by a polylogarithmic factor. Our key technical contribution is a randomized online algorithm for set cover and non-metric facility location in a new adversarial model which we call semi-adaptive adversaries. As a by-product of our techniques, we obtain the first deterministic $O(\log |C| \log |F|)$-competitive algorithm for non-metric facility location.

Sander Borst
Centrum Wiskunde & Informatica (CWI)
sborst@mpi-inf.mpg.de

Marek Elias, Moritz Venzin
Bocconi University
marek.elias@unibocconi.it, moritz.venzin@unibocconi.it

CP35

Online Scheduling Via Gradient Descent for Weighted Flow Time Minimization

In this paper, we explore how a natural generalization of Shortest Remaining Processing Time (SRPT) can be a powerful *meta-algorithm* for online scheduling. The meta-algorithm processes jobs to maximally reduce the objective of the corresponding offline scheduling problem of the remaining jobs: minimizing the total weighted completion time of them (the residual optimum). We show that it achieves scalability for minimizing total weighted flow time when the residual optimum exhibits *supermodularity*. Scalability here means it is $O(1)$-competitive with an arbitrarily small speed augmentation advantage over the adversary, representing the best possible outcome achievable for various scheduling problems. Thanks to this finding, our approach does not require the residual optimum to have a closed mathematical form. Consequently, we can obtain the schedule by solving a linear program, which makes our approach readily applicable to a rich body of applications. Furthermore, by establishing a novel connection to *substitute valuations in Walrasian markets*, we show how to achieve supermodularity, thereby obtaining scalable algorithms for various scheduling problems, such as matroid scheduling, generalized network flow, and generalized arbitrary speed-up curves, etc., and this is the *first* non-trivial or scalable algorithm for many of them.

Qingyun Chen, Sungjin Im, Aditya Petety
University of California, Merced
qchen41@ucmerced.edu, sim3@ucmerced.edu, apetety@ucmerced.edu

CP35

Unweighted Layered Graph Traversal: Passing a Crown Via Entropy Maximization

Introduced by Papadimitriou and Yannakakis in 1989, layered graph traversal is a central problem in online algorithms and mobile computing that has been studied for several decades, and which now is essentially resolved in its original formulation. In this paper, we demonstrate that what appears to be an innocuous modification of the problem actually leads to a drastic (exponential) reduction of the competitive ratio. Specifically, we present an algorithm that is $O(\log^2 w)$-competitive for traversing unweighted layered graphs of width w. Our algorithm chooses the agents position simply according to the probability distribution over the current layer that maximizes the sum of entropies of the induced distributions in the preceding lay-

ers.

Romain Cosson
Inria
romain.cosson@inria.fr

Xingjian Bai
MIT
xingjianbai0914@gmail.com

Christian Coester
University of Oxford
christian.coester@cs.ox.ac.uk

CP35

The Power of Proportional Fairness for Non-Clairvoyant Scheduling under Polyhedral Constraints

The Polytope Scheduling Problem (PSP) was introduced by Im, Kulkarni, and Munagala (JACM 2018) as a very general abstraction of resource allocation over time and captures many well-studied problems including classical unrelated machine scheduling, multidimensional scheduling, and broadcast scheduling. In PSP, jobs with different arrival times receive processing rates that are subject to arbitrary packing constraints. An elegant and well-known algorithm for instantaneous rate allocation with good fairness and efficiency properties is the Proportional Fairness algorithm (PF), which was analyzed for PSP by Im et al. We drastically improve the analysis of PF for both the general PSP and several of its important special cases subject to the objective of minimizing the sum of weighted completion times. We reduce the upper bound on the competitive ratio from 128 to 27 for general PSP and to 4 for the prominent class of monotone PSP. For certain heterogeneous machine environments we even close the gap to the lower bound of 2 for non-clairvoyant scheduling. Our analysis also gives the first polynomial-time improvements over the nearly 30-year-old bounds on the competitive ratio of the doubling framework by Hall, Shmoys, and Wein (SODA 1996) for clairvoyant online preemptive scheduling on unrelated machines. Somewhat surprisingly, we achieve this improvement by a non-clairvoyant algorithm, thereby demonstrating that non-clairvoyance is not a (significant) hurdle.

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

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

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

CP35

Putting Off the Catching Up: Online Joint Replenishment Problem with Holding and Backlog Costs

We study an online generalization of the classic Joint Replenishment Problem (JRP) that models the trade-off between ordering costs, holding costs, and backlog costs in supply chain planning systems. A retailer places orders to a supplier for multiple items over time: each request is for some item that the retailer needs in the future, and has an arrival time and a soft deadline. If a request is served before its deadline, the retailer pays a holding cost per unit of the item until the deadline. However, if a request is served after its deadline, the retailer pays a backlog cost per unit. Each service incurs a fixed joint service cost and a fixed item-dependent cost for every item included in a service, irrespective of the units of each item ordered. The goal is to schedule services to satisfy all the online requests while minimizing the sum of the service costs, the holding costs, and the backlog costs. Our general model with holding and backlog costs has not been investigated earlier, and no online algorithms are known even in the make-to-stock version with hard deadlines and non-zero holding costs. We develop a new online algorithm for the general version of online JRP with both holding and backlog costs and establish that it is 30-competitive. Along the way, we develop a 3-competitive algorithm for the single-item case that we build on to get our final result. Our algorithm uses a greedy strategy and its competitiveness is shown using a dual fitting analysis.

Benjamin Moseley, Aidin Niaparast, R. Ravi
Carnegie Mellon University
moseleyb@andrew.cmu.edu, aniapara@andrew.cmu.edu, ravi@andrew.cmu.edu

CP35

Competitive Strategies to Use "Warm Start" Algorithms with Predictions

We consider the problem of learning and using predictions for warm start algorithms with predictions. In this setting, an algorithm is given an instance of a problem, and a prediction of the solution. The runtime of the algorithm is bounded by the distance from the predicted solution to the true solution of the instance. We give competitive guarantees against stronger benchmarks that consider a set of k predictions $\mathbf{P}$. That is, the "optimal offline cost" to solve an instance with respect to $\mathbf{P}$ is the distance from the true solution to the closest member of $\mathbf{P}$. In the distributional setting, we show a simple strategy that incurs cost at most an $O(k)$ factor worse than the optimal offline cost. We then show a way to leverage learnable coarse information, in the form of partitions of the instance space into groups of "similar" instances, that allows us to potentially avoid this $O(k)$ factor. Finally, we consider an online version of the problem, where we compete against offline strategies that are allowed to maintain a moving set of k predictions or "trajectories," and are charged for how much the predictions move. We give an algorithm that does at most $O(k^4 \ln^2 k)$ times as much work as any offline strategy of k trajectories. This algorithm is deterministic (robust to an adaptive adversary), and oblivious to the setting of k. Thus the guarantee holds for all k simultaneously.

Avrim Blum

Toyota Technological Institute at Chicago
avrim@ttic.edu

Vaidehi Srinivas
Northwestern University
vaidehi@u.northwestern.edu

CP36

Faster Two-Dimensional Pattern Matching with k Mismatches

The classical pattern matching asks for locating all occurrences of one string, called the pattern, in another, called the text, where a string is simply a sequence of characters. Due to the potential practical applications, it is desirable to seek approximate occurrences, for example by bounding the number of mismatches. This problem has been extensively studied, and by now we have a good understanding of the best possible time complexity as a function of n (length of the text), m (length of the pattern), and k (number of mismatches). In particular, we know that for $k = \mathcal{O}(\sqrt{m})$, we can achieve quasi-linear time complexity [Gawrychowski and Uznanski, ICALP 2018]. We consider a natural generalisation of the approximate pattern matching problem to two-dimensional strings, which are simply square arrays of characters. In the approximate two-dimensional pattern matching, we are given a pattern of size $m \times m$ and a text of size $n \times n$, and ask for all locations in the text where the pattern matches with at most k mismatches. The asymptotically fastest algorithm for this algorithm works in $\mathcal{O}(kn^2)$ time [Amir and Landau, TCS 1991]. We provide a new insight into two-dimensional periodicity to improve on these 30-years old bounds. Our algorithm works in $\tilde{\mathcal{O}}((m^2 + mk^{5/4})n^2/m^2)$ time, which is $\tilde{\mathcal{O}}(n^2)$ for $k = \mathcal{O}(m^{4/5})$.

Jonas Ellert
ENS Paris
ellert.jonas@gmail.com

Pawel Gawrychowski, Adam Gorkiewicz
University of Wroclaw
gawry@cs.uni.wroc.pl, adamgoorkiewicz@gmail.com

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

CP36

Fast and Simple Sorting Using Partial Information

We consider the problem of sorting n items, given the outcomes of m pre-existing comparisons. We present a simple and natural deterministic algorithm that runs in $O(m + \log T)$ time and does $O(\log T)$ comparisons, where T is the number of total orders consistent with the pre-existing comparisons. Our running time and comparison bounds are best possible up to constant factors, thus resolving a problem that has been studied intensely since 1976 (Fredman, Theoretical Computer Science). The best previous algorithm with a bound of $O(\log T)$ on the number of comparisons has a time bound of $O(n^{2.5})$ and is more complicated. Our algorithm combines three classic algorithms: topological sort, heapsort with the right kind of heap, and efficient search in a sorted list. It outputs the items in sorted order one by one. It can be modified to stop early, thereby solving the important and more general top-k sorting problem: Given k and the outcomes of some pre-existing comparisons, output the smallest k items in sorted order. The modified algorithm solves the top-k sorting problem in minimum time and comparisons, to within constant factors.

Richard Hladík
ETH Zurich
& INSAIT, Sofia University
rihl@uralyx.cz

Bernhard Haeupler
ETH Zurich
bernhard.haeupler@inf.ethz.ch

John Iacono
Université libre de Bruxelles
ulb@johniacono.com

Václav Rozhon
ETH
vaclavrozhon@gmail.com

Robert Tarjan
Princeton University
ret@cs.princeton.edu

Jakub Tetek
INSAIT, Sofia University
j.tetek@gmail.com

CP36

Efficient d-Ary Cuckoo Hashing at High Load Factors by Bubbling Up

A d-ary cuckoo hash table is an open-addressed hash table that stores each key x in one of d random positions $h_1(x), h_2(x), \ldots, h_d(x)$. In the offline setting, where all items are given and keys need only be matched to locations, it is possible to support a load factor of $1 - \epsilon$ while using $d = \lceil \ln \epsilon^{-1} + o(1) \rceil$ hashes. The online setting, where keys are moved as new keys arrive sequentially, has the additional challenge of the time to insert new keys, and it has not been known whether one can use $d = O(\ln \epsilon^{-1})$ hashes to support $\text{poly}(\epsilon^{-1})$ expected-time insertions. In this paper, we introduce *bubble-up cuckoo hashing*, an implementation of d-ary cuckoo hashing that achieves all of the following properties simultaneously: (1) it uses $d = \lceil \ln \epsilon^{-1} + \alpha \rceil$ hash locations per item for an arbitrarily small positive constant α; (2) it achieves expected insertion time $O(\delta^{-1})$ for any insertion taking place at load factor $1 - \delta \leq 1 - \epsilon$; and (3)it achieves expected positive query time $O(1)$, independent of d and ϵ. The first two properties give an essentially optimal value of d without compromising insertion time. The third property is interesting even in the offline setting: it says that, even though negative queries must take time d, positive queries can actually be implemented in $O(1)$ expected time, even when

d is large.

William Kuszmaul
CMU
kuszmaul@cmu.edu

Michael Mitzenmacher
Harvard University
michaelm@eecs.harvard.edu

CP36

Top-K Document Retrieval in Compressed Space

Let $\mathcal{D}$ be a collection of D strings of total length n over an alphabet of size σ. We consider the so-called top-k document retrieval problem: given a short string P and an integer k, list the identifiers of k strings in $\mathcal{D}$ most relevant to P, in decreasing order of relevance. Relevance may be a fixed value associated with the strings where P occurs, or the number of times P occurs in the strings. While RAM-optimal solutions using $O(n \log n)$ bits and $O(|P|/\log_\sigma n+k)$ time exist, solving the problem optimally within space close to $O(n \log \sigma)$ bits is open. We describe a data structure for the top-k document retrieval problem that uses $O(\log \log n)$ bits per symbol on top of any compressed suffix array (CSA) of $\mathcal{D}$, and supports queries in essentially optimal time, in the following sense. Given a CSA using $|\mathrm{CSA}|$ bits of space, that finds the suffix array range of a query string P in time t_{cnt}, and accesses a suffix array entry in time t_{SA}, listing any k pattern occurrences would take time $O(t_{\mathrm{cnt}} + k\, t_{\mathrm{SA}})$. Our top-$k$ data structure uses $|\mathrm{CSA}| + O(n \log \log n)$ bits and reports k most relevant documents that contain P in time $O(t_{\mathrm{cnt}} + k\,(t_{\mathrm{SA}} + \log \log n))$. On every known CSA using $O(n \log \sigma)$ bits, t_{SA} is $\Omega(\log \log n)$ in virtually all cases, thus our time is $O(t_{\mathrm{cnt}} + k\, t_{\mathrm{SA}})$ in most situations.

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

Yakov Nekrich
Michigan Technological University
yakov.nekrich@googlemail.com

CP36

Tight Bounds and Phase Transitions for Incremental and Dynamic Retrieval

Retrieval data structures are data structures that answer key-value queries without paying the space overhead of explicitly storing keys. The problem can be formulated in four settings (static, value-dynamic, incremental, or dynamic), each of which offers different levels of dynamism to the user. In this paper, we establish optimal bounds for the final two settings (incremental and dynamic) in the case of a polynomial universe. Our results complete a line of work that has spanned more than two decades, and also come with a surprise: the incremental setting, which has long been viewed as essentially equivalent to the dynamic one, actually has a phase transition, in which, as the value size v approaches $\log n$, the optimal space redundancy actually begins to shrink, going from roughly $n \log \log n$ (which has long been thought to be optimal) all the way down to $\Theta(n)$ (which is the optimal bound even for the seemingly much-easier value-dynamic setting).

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

William Kuszmaul
CMU
kuszmaul@cmu.edu

Tingqiang Xu, Hangrui Zhou, Renfei Zhou
Tsinghua University
xtq23@mails.tsinghua.edu.cn,
zhouhr23@mails.tsinghua.edu.cn, renfeiz@andrew.cmu.edu

CP36

A Cell Probe Lower Bound for the Predecessor Search Problem in PRAM

We study the predecessor search problem in the classical PRAM model of computation. In this problem, the input is a set of n ℓ-bit integers and the goal is to store the input in a data structure of size $S(n)$ such that given a query value q, the predecessor of q can be found efficiently. We prove a lower bound for this problem in the strongest CRCW PRAM model. A simplified version of the lower bound states that in a K-processor PRAM model with $O(\log n)$-bit registers, the query requires $\Omega(\log_K \log n)$ worst-case time under the realistic setting where the space is near-linear.

Peyman Afshani
Aarhus University
peyman@cs.au.dk

Nodari Sitchinava
University of Hawaii at Manoa
nodari@hawaii.edu

CP37

Fully-Distributed Byzantine Agreement in Sparse Networks

Byzantine agreement is a fundamental problem in fault-tolerant distributed networks that has been studied intensively for the last four decades. Most of these works designed protocols for complete networks. A key goal in Byzantine protocols is to tolerate as many Byzantine nodes as possible. Dwork, Peleg, Pippenger, and Upfal [STOC 1986, SICOMP 1988] were the first to address the Byzantine agreement problem in sparse, bounded degree networks, with a protocol that achieved almost-everywhere agreement among honest nodes. In such networks, all known Byzantine agreement protocols that tolerated a large number of Byzantine nodes were not fully-distributed — in those protocols, nodes are required to have initial knowledge of the entire network topology. Thus, these works raise the fundamental open question of whether one can design Byzantine fully-distributed protocols that tolerate a large number of Byzantine nodes in sparse

networks. The work of Augustine, Pandurangan, and Robinson [PODC 2013] presented the first fully-distributed Byzantine agreement protocol that works in sparse networks, but it tolerated only up to $O(\sqrt{n}/polylogn)$ Byzantine nodes (where n is the network size). We answer the earlier open question by presenting fully-distributed Byzantine agreement protocols, under the full information model, for sparse, bounded degree networks that tolerate significantly more Byzantine nodes — up to $O(n/polylogn)$ of them.

John Augustine
IIT Madras
augustine@cse.iitm.ac.in

Fabien Dufoulon
Lancaster University
f.dufoulon@lancaster.ac.uk

Gopal Pandurangan
University of Houston
gopal@cs.uh.edu

CP37

Partial Synchrony for Free: New Upper Bounds for Byzantine Agreement

Byzantine agreement allows n processes to decide on a common value, in spite of arbitrary failures. In this paper, we introduce Oper, the first generic transformation of deterministic Byzantine agreement algorithms from synchrony to partial synchrony. Oper requires no cryptography, is optimally resilient ($n \geq 3t+1$, where t is the maximum number of failures), and preserves the worst-case per-process bit complexity of the transformed synchronous algorithm. Leveraging Oper, we present the first partially synchronous Byzantine agreement algorithm that (1) achieves optimal $O(n^2)$ bit complexity, (2) requires no cryptography, and (3) is optimally resilient ($n \geq 3t + 1$), thus showing that the Dolev-Reischuk bound is tight even in partial synchrony. Moreover, we adapt Oper for long values and obtain several new partially synchronous algorithms with improved complexity and weaker (or completely absent) cryptographic assumptions. Finally, we demonstrate the broad applicability of the Oper transformation by showcasing its use for randomized synchronous agreement algorithms. Indirectly, Oper contradicts the folklore belief that there is a fundamental gap between synchronous and partially synchronous agreement protocols. In a way, we show that there is no inherent trade-off between the robustness of partially synchronous algorithms on the one hand, and the simplicity/efficiency of synchronous ones on the other hand.

Pierre Civit
EPFL
pierre.civit@epfl.ch

Ayaz Dzulfikar, Seth Gilbert
NUS Singapore
ayaz.dzulfikar@u.nus.edu, gilbert@comp.nus.edu.sg

Rachid Guerraoui, Jovan Komatovic, Manuel Vidigueira
EPFL
rachid.guerraoui@epfl.ch, jovan.komatovic@epfl.ch,
manuel.ribeirovidigueira@epfl.ch

Igor Zablotchi
Mysten Labs
igor@mystenlabs.com

CP37

Parks and Recreation: Color Fault-Tolerant Spanners Made Local

We provide new algorithms for constructing spanners of arbitrarily edge- or vertex-colored graphs, that can endure up to f failures of entire color classes. Recently, Petruschka, Sapir and Tzalik [ITCS'24] gave tight bounds for the (worst-case) size s of such spanners with stretch $(2k-1)$ that are resilient to at most f color faults. Additionally, they showed a greedy-based algorithm for computing spanners of size $\tilde{O}(s)$, running in $\tilde{O}(msf)$ sequential time. Providing faster and/or distributed algorithms was left open. We address this by providing a novel variant of the Baswana-Sen algorithm [RSA'07] in the spirit of Parter's algorithm for vertex fault-tolerant spanners [STOC'22]. In a nutshell, it produces color fault-tolerant spanners of size $\tilde{O}_k(s)$ (hence near-optimal for any fixed k), has optimal locality $O(k)$ (i.e., takes $O(k)$ rounds in the LOCAL model), can be implemented in $O_k(f^{k-1})$ rounds in CONGEST, and takes $\tilde{O}_k(m + sf^{k-1})$ sequential time. To handle the considerably more difficult setting of color faults, our approach differs from [BS07, Par22] by taking a novel edge-centric perspective. Another key technical contribution is in constructing and using collections of short paths that are "colorful at all scales', which we call "parks'. These are intimately connected with the notion of spread set-systems that found use in recent breakthroughs regarding the famous Sunflower Conjecture.

Asaf Petruschka, Merav Parter, Shay Sapir, Elad Tzalik
Weizmann Institute
asaf.petruschka@weizmann.ac.il,
merav.parter@weizmann.ac.il, shay.sapir@weizmann.ac.il,
elad.tzalik@weizmann.ac.il

CP37

On the Locality of Hall's Theorem

The last five years of research on distributed graph algorithms have seen huge leaps of progress: new strong lower bounds have emerged for many central problems and exponential improvements over the state of the art have been achieved for the runtimes of many algorithms. Nevertheless, there are still large gaps between the best known upper and lower bounds for many important problems. We develop a novel algorithm design technique aimed at closing this gap. It ensures a logarithmic runtime by carefully combining local solutions into a globally feasible solution. In essence, each node finds a carefully chosen local solution in $O(\log n)$ rounds and this solution is consistent with the other nodes' solutions without coordination. The local solutions are based on a distributed version of Hall's theorem that may be of independent interest. We showcase our framework by improving on the state of the art for the following fundamental problems: edge coloring, bipartite saturating matchings and hypergraph sinkless orientation. In

particular, we obtain an asymptotically optimal $O(\log n)$-round algorithm for $(3\Delta/2)$-edge coloring in bounded degree graphs. The previously best bound for the problem was $O(\log^4 n)$ rounds, obtained by plugging in the state-of-the-art maximal independent set algorithm from [Ghaffari, Grunau, SODA'23] into the $3\Delta/2$-edge coloring algorithm from [Ghaffari, Kuhn, Maus, Uitto, STOC'18].

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

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

Ananth Narayanan
CISPA Helmholtz Center for Information Security
ananth.narayanan@cispa.de

Florian Schager
TU Graz
florian.schager@tugraz.at

Jara Uitto
Aalto University
jara.uitto@aalto.fi

CP37

Asynchronous 3-Majority Dynamics with Many Opinions

We consider 3-Majority, a probabilistic consensus dynamics on a complete graph with n vertices, each vertex starting with one of k initial opinions. At each discrete time step, a vertex u is chosen uniformly at random. The selected vertex u chooses three neighbors v_1, v_2, v_3 uniformly at random with replacement and takes the majority opinion held by the three, where ties are broken in favor of the opinion of v_3. The main quantity of interest is the consensus time, the number of steps required for all vertices to hold the same opinion. This asynchronous version turns out to be considerably harder to analyze than the synchronous version and so far results have only been obtained for $k = 2$. Even in the synchronous version the results for large k are far from tight. In this paper we prove that the consensus time is $\tilde{\Theta}(\min(nk, n^{1.5}))$ for all k. These are the first bounds for all k that are tight up to a polylogarithmic factor.

Nobutaka Shimizu
Institute of Science Tokyo,
Japan
shimizu.n.ah@m.titech.ac.jp

Colin Cooper
Kings College, London
colin.cooper@kcl.ac.uk

Frederik Mallmann-Trenn
King's College London
frederik.mallmann-trenn@kcl.ac.uk

Tomasz Radzik
Kings College London
tomasz.radzik@kcl.ac.uk

Takeharu Shiraga
Chuo University
shiraga.076@g.chuo-u.ac.jp

CP37

Sublinear-Round Broadcast Without Trusted Setup

Byzantine broadcast is one of the fundamental problems in distributed computing. Many of its practical applications, from multiparty computation to consensus mechanisms for blockchains, require increasingly weaker trust assumptions, as well as scalability for an ever-growing number of users n. This rules out existing solutions which run in a linear number of rounds in n or rely on trusted setup requirements. We propose the first sublinear-round and trustless Byzantine broadcast protocol for dishonest majority. Unlike previous sublinear-round protocols, our protocol assumes neither the existence of a trusted dealer who honestly issues keys and correlated random strings to the parties nor random oracles. Instead, we present a solution whose setup is limited to an unstructured uniform reference string and a plain public key infrastructure (a.k.a. bulletin-board PKI). Our broadcast protocol builds on top of a moderated gradecast protocol which parties can use to reach weak agreement on shared random strings. Using these strings, we can then run in an unbiased fashion a committee-based Byzantine protocol, similar to that of Chan et al. (PKC 2020), which terminates in a sublinear number of rounds. To this end, we propose a novel construction for committee election, which does not rely either on random oracles or on a trusted dealer, and uses NIZKs and time-lock puzzles. Our protocol is resilient against an adaptive adversary who corrupts any constant fraction of parties.

Andreea Alexandru
Duality Technologies
aalexandru@dualitytech.com

Julian Loss
CISPA Helmholtz Center for Information Security
lossjulian@gmail.com

Charalampos Papamanthou
Yale University
charalampos.papamanthou@yale.edu

Giorgos Tsimos
University of Maryland
tsimos@umd.edu

Benedikt Wagner
Ethereum Foundation
benedikt.wagner@ethereum.org

CP38

A Lower Bound for Light Spanners in General Graphs

A recent upper bound by Le and Solomon [STOC '23] has

established that every n-node graph has a $(1+\varepsilon)(2k-1)$-spanner with lightness $O(\varepsilon^{-1}n^{1/k})$. This bound is optimal up to its dependence on ε; the remaining open problem is whether this dependence can be improved or perhaps even removed entirely. We show that the ε-dependence cannot in fact be completely removed. For constant k and for $\varepsilon := \Theta(n^{-\frac{1}{2k-1}})$, we show a lower bound on lightness of

$$\Omega\left(\varepsilon^{-1/k}n^{1/k}\right).$$

For example, this implies that there are graphs for which any 3-spanner has lightness $\Omega(n^{2/3})$, improving on the previous lower bound of $\Omega(n^{1/2})$. An unusual feature of our lower bound is that it is conditional on the girth conjecture with parameter $k-1$ rather than k. We additionally show that this implies certain technical limitations to improving our lower bound further. In particular, under the same conditional, generalizing our lower bound to all ε *or* obtaining an optimal ε-dependence are as hard as proving the girth conjecture for all constant k.

Greg Bodwin, Jeremy Flics
University of Michigan
bodwin@umich.edu, jflics@umich.edu

CP38

Spanners in Planar Domains Via Steiner Spanners and Non-Steiner Tree Covers

We study spanners in planar domains, including polygonal domains, polyhedral terrain, and planar metrics. Previous work showed that for any constant $\varepsilon \in (0,1)$, one could construct a $(2+\varepsilon)$-spanner with $O(n \log n)$ edges (Abam et al., 2019), and there is a lower bound of $\Omega(n^2)$ edges for any $(2-\varepsilon)$-spanner (Abam et al., 2015). The main open question is whether linearly many edges suffice and whether the stretch can be reduced to 2. We resolve both problems by showing that for stretch 2, one needs $\Omega(n \log n)$ edges, and for stretch $2+\varepsilon$ for any fixed $\varepsilon \in (0,1)$, $O(n)$ edges are sufficient. Our lower bound is the first super-linear lower bound for stretch 2. En route to achieve our result, we introduce the problem of constructing non-Steiner tree covers for metrics: given a tree and a set of terminals in the tree, our goal is to construct a collection of a small number of dominating trees such that for every pair of terminals, at least one tree in the collection preserves their distance up to a small stretch factor. Here, we identify an unexpected threshold phenomenon around 2 where a sharp transition from n trees to $\Theta(\log n)$ trees and then to $O(1)$ trees happens. Finally, we study $(1+\varepsilon)$-spanners in planar domains using Steiner points. Here, we construct a $(1+\varepsilon)$-spanner where the number of edges depends almost linearly on ε.

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

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

Balázs Keszegh
Rényi institute
keszegh@renyi.hu

Andrey Kupavskii
CNRS
kupavskii@ya.ru

Alexandre Louvet
Université Sorbonne Paris Nord
alexandre.louvet@mailo.fr

Dömötör Pálvölgyi
ELTE
domotor.palvolgyi@ttk.elte.hu

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

CP38

Improved Online Reachability Preservers

A *reachability preserver* is a basic kind of graph sparsifier, which preserves the reachability relation of an n-node directed input graph G among a set of given demand pairs P of size $|P| = p$. We give constructions of sparse reachability preservers in the online setting, where G is given on input, the demand pairs $(s,t) \in P$ arrive one at a time, and we must irrevocably add edges to a preserver H to ensure reachability for the pair (s,t) before we can see the next demand pair. Our main results are: – There is a construction that guarantees a maximum preserver size of

$$|E(H)| \leq O\left(n^{0.72}p^{0.56} + n^{0.6}p^{0.7} + n\right).$$

– Given a promise that the demand pairs will satisfy $P \subseteq S \times V$ for some vertex set S of size $|S| =: \sigma$, there is a construction that guarantees a maximum preserver size of

$$|E(H)| \leq O\left((np\sigma)^{1/2} + n\right).$$

Our techniques also give a small polynomial improvement in the current upper bounds for *offline* reachability preservers, and our results extend to an even stronger model in which we must commit to a path for all possible reachable pairs in G before any demand pairs have been received. As an application, we improve the competitive ratio for Online Unweighted Directed Steiner Forest to $O(n^{3/5+\varepsilon})$, improving on the previous bound of $O(n^{2/3+\varepsilon})$ [Grigorescu, Lin, Quanrud APPROX-RANDOM '21].

Greg Bodwin, Tuong Le
University of Michigan
bodwin@umich.edu, tuongle@umich.edu

CP38

Having Hope in Missing Spanners: New Distance Preservers and Light Hopsets

An r-missing spanner for a graph G is a sparse subgraph $H \subseteq G$ satisfying that for any u, v pair there is a (possibly approximate) u-v shortest path P in G such that $|P \setminus H| \leq r$. That is, H misses at most r edges from every u-v (approximate) shortest path. [Kogan and Parter, FOCS '22] introduced the notion of missing spanners as

an intermediate step for translating hopset constructions into spanners and distance preservers. In this work, we provide new constructions of missing spanners that lead to improved distance preservers. We also present a reduction in the reverse direction to that of [KP'22] by translating a special class of distance preservers into hopsets. Our results provide new links between graph augmentation and reduction structures, which for many years have been studied in isolation. We believe that these connections should have further combinatorial and algorithmic applications.

Shimon Kogan, Merav Parter
Weizmann Institute
shimon.kogan@weizmann.ac.il,
merav.parter@weizmann.ac.il

CP38

Subquadratic Algorithms in Minor-Free Digraphs: (weighted) Distance Oracles, Decremental Reachability, and More

Le and Wulff-Nilsen [SODA '24] initiated a systematic study of VC set systems to unweighted K_h-minor-free directed graphs. We extend their results in the following ways: • We present the first application of VC set systems for real-weighted minor-free digraphs to build the first exact subquadratic-space distance oracle with $O(\log n)$ query time. Prior work using VC set systems only applied in unweighted and integer-weighted digraphs. • We describe a unified system for analyzing the VC dimension of balls and the LP set system (based on Li–Parter [STOC '19]) of Le–Wulff-Nilsen [SODA '24] using pseudodimension. This improves the VC dimension bound of the LP set system in directed graphs to $h - 1$. • We present the first application of these set systems in a dynamic setting. Specifically, we construct decremental reachability oracles with subquadratic total update time and constant query time. Prior to this work, it was not known if this was possible to construct oracles with subquadratic total update time and polylogarithmic query time, even in planar digraphs. • We describe subquadratic time algorithms for unweighted digraphs including (1) constructions of exact distance oracles, (2) computation of vertex eccentricities and Wiener index. The main innovation in obtaining these results is the use of dynamic string data structures.

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

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

CP39

Tree Independence Number IV. Even-Hole-Free Graphs

We prove that the tree independence number of every even-hole-free graph is at most polylogarithmic in its number of vertices. More explicitly, we prove that there exists a constant $c > 0$ such that for every integer $n > 1$ every n-vertex even-hole-free graph has a tree decomposition where each bag has stability (independence) number at most $clog^{10}(n)$. This implies that the Maximum Weight Independent Set problem, as well as several other natural algorithmic problems that are known to be NP-hard in general, can be solved in quasi-polynomial time if the input graph is even-hole-free.

Maria Chudnovsky
Princeton University
mchudnov@math.princeton.edu

Peter Gartland
University of California, Santa Barbara
petergartland@ucsb.edu

Sepehr Hajebi
University of Waterloo
shajebi@uwaterloo.ca

Daniel Lokshtanov
UCSB
daniello@ucsb.edu

Sophie Spirkl
University of Waterloo
sophie.spirkl@uwaterloo.ca

CP39

New Separations and Reductions for Directed Hopsets and Preservers

We study hopsets, shortcut sets, and distance preservers in n-node, m-edge graphs, and show improved bounds in various settings for these problems. Our first set of results is about directed hopsets and shortcut sets. We show an $\Omega(n^{1/2})$ bound for $O(n)$-size approximate hopsets with any given finite stretch, in graphs with arbitrary aspect ratio. This establishes a separation between this setting and $O(n)$-size approximate hopsets for graphs with polynomial aspect ratio. Furthermore, we show an $\Omega(n^{2/7})$ bound for $O(m)$-size exact hopsets in unweighted, undirected graphs and an $\Omega(n^{2/9})$ bound for $O(m)$-size shortcut set, improving the previous bounds. Our second set of results is about exact and approximate distance preservers. We show an $\Omega(n^{2/3}p^{2/3})$ lower bound for approximate distance preservers on n-node graphs with p demand pairs, for directed graphs with arbitrary aspect ratio. This establishes a separation between this problem and reachability preservers. Additionally, we also show an $\widetilde{O}(n^{5/6}p^{2/3} + n)$ upper bound for exact distance preservers for unweighted directed graphs, separating exact distance preservers from consistent tiebreaking schemes in directed, unweighted graphs. Finally, we show a reduction from directed exact distance preservers to undirected exact distance preservers.

Yinzhan Xu
Massachusetts Institute of Technology
xyzhan@mit.edu

Gary Hoppenworth
University of Michigan
garytho@umich.edu

Zixuan Xu
MIT
zixuanxu@mit.edu

CP39

A Topological Proof Of The HellNeetril Dichotomy

We provide a new proof of a theorem of Hell and Neetril [J. Comb. Theory B, 48(1):92110, 1990] using tools from topological combinatorics based on ideas of Lovsz [J. Comb. Theory, Ser. A, 25(3):319324, 1978]. The HellNeetril Theorem provides a dichotomy of the graph homomorphism problem. It states that deciding whether there is a graph homomorphism from a given graph to a fixed graph H is in P if H is bipartite (or contains a self-loop), and is NP-complete otherwise. In our proof we combine topological combinatorics with the algebraic approach to constraint satisfaction problem.

Sebastian Meyer
TU Dresden
sebastian.meyer2@tu-dresden.de

Jakub Opršal
University of Birmingham
UK
j.oprsal@bham.ac.uk

CP39

A Refutation of the Pach-Tardos Conjecture for 0-1 Matrices

The theory of forbidden 0-1 matrices generalizes Turan-style (bipartite) subgraph avoidance, Davenport-Schinzel theory, and Zarankiewicz-type problems, and has been influential in many areas, such as discrete and computational geometry, the analysis of self-adjusting data structures, and the development of the graph parameter twin width. The foremost open problem in this area is to resolve the Pach-Tardos conjecture from 2005, which states that if a forbidden pattern $P \in \{0,1\}^{k\times l}$ is acyclic, meaning it is the bipartite incidence matrix of a forest, then $\mathrm{Ex}(P,n) = O(n\log^{C_P} n)$, where $\mathrm{Ex}(P,n)$ is the maximum number of 1s in a P-free $n \times n$ 0-1 matrix and C_P is a constant depending only on P. This conjecture has been confirmed on many small patterns, specifically all P with weight at most 5, and all but two with weight 6. The main result of this paper is a clean refutation of the Pach-Tardos conjecture. Specifically, we prove that $\mathrm{Ex}(S_0,n), \mathrm{Ex}(S_1,n) \geq n2^{\Omega(\sqrt{\log n})}$, where S_0, S_1 are the outstanding weight-6 patterns.

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

Gábor Tardos
Renyi Institute
tardosgabor@gmail.com

CP39

Recognizing Sumsets is NP-Complete

Sumsets are central objects in additive combinatorics. In 2007, Granville asked whether one can efficiently recognize whether a given set S is a sumset, i.e. whether there is a set A such that $A + A = S$. Granville suggested an algorithm that takes exponential time in the size of the given set, but can we do polynomial or even linear time? This basic computational question is indirectly asking a fundamental structural question: do the special characteristics of sumsets allow them to be efficiently recognizable? In this paper, we answer this question negatively by proving that the problem is NP-complete. Specifically, our results hold for integer sets and over any finite field. Assuming the Exponential Time Hypothesis, our lower bound becomes $2^{\Omega(n^{1/4})}$.

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

Nick Fischer
INSAIT, University of Sofia
nick.fischer@insait.ai

Ron Safier, Nathan Wallheimer
Weizmann Institute of Science
ron.safier@weizmann.ac.il,
nathan.wallheimer@weizmann.ac.il

CP40

Exact Thresholds for Noisy Non-Adaptive Group Testing

In recent years, the mathematical limits and algorithmic bounds for probabilistic group testing have become increasingly well-understood, with exact asymptotic thresholds now being known in general scaling regimes for the noiseless setting. In the noisy setting where each test outcome is flipped with constant probability, there have been similar developments, but the overall understanding has lagged significantly behind the noiseless setting. In this paper, we substantially narrow this gap by deriving exact asymptotic thresholds for the noisy setting under two widely-studied random test designs: i.i.d. Bernoulli and near-constant tests-per-item. These thresholds are established by combining components of an existing information-theoretic threshold decoder with a novel analysis of maximum-likelihood decoding (upper bounds), and deriving a novel set of impossibility results by analyzing certain failure events for optimal maximum-likelihood decoding (lower bounds).

Junren Chen
The University of Hong Kong
chenjr58@connect.hku.hk

Jonathan Scarlett
National University of Singapore

scarlett@comp.nus.edu.sg

CP40

Average-Case Hardness of Parity Problems: Orthogonal Vectors, K-Sum and More

Mina Dalirrooyfard
Morgan Stanley
minad@mit.edu

Andrea Lincoln
Boston University
andrea2@bu.edu

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

Virginia Vassilevska Williams
MIT
virgi@mit.edu

CP40

Sumsets, 3SUM, Subset Sum: Now for Real!

We study a broad class of algorithmic problems with an "additive flavor" such as computing sumsets, 3SUM, Subset Sum and geometric pattern matching. Our starting point is that these problems can often be solved efficiently for *integers*, owed to the rich available tool set including bit-tricks, linear hashing, and the Fast Fourier Transform. However, for *real numbers* these tools are not available, leading to significant gaps in the best-known running times for integer inputs versus for real inputs. In this work our goal is to close this gap. As our key contribution we design a new technique for computing *real* sumsets. It is based on a surprising blend of algebraic ideas (like *Prony's method* and *coprime factorizations*) with combinatorial tricks. We then apply our new algorithm to the aforementioned problems and successfully obtain, in all cases, equally fast algorithms for real inputs.

Nick Fischer
INSAIT, University of Sofia
nick.fischer@insait.ai

CP40

New Applications of 3SUM-Counting in Fine-Grained Complexity and Pattern Matching

The 3SUM problem is one of the cornerstones of fine-grained complexity. Its study has led to countless lower bounds, but as has been sporadically observed beforeand as we will demonstrate againinsights on 3SUM can also lead to *algorithmic* applications. The starting point of our work is that we spend a lot of technical effort to develop new algorithms for 3SUM-type problems such as approximate 3SUM-counting, small-doubling 3SUM-counting, and a deterministic subquadratic-time algorithm for the celebrated Balog-Szemerdi-Gowers theorem from additive combinatorics. All of these are relevant in their own right and may prove useful for future research on 3SUM. But perhaps even more excitingly, as consequences of these tools, we derive diverse new results in fine-grained complexity and pattern matching algorithms, answering open questions from many unrelated research areas. Specifically, our results imply (1) a derandomization of the recent "short cycle removal" technique, (2) a derandomization of the fine-grained equivalence of 3SUM and #3SUM, (3) a faster deterministic algorithm for approximate Text-to-Pattern Hamming Distances, as well as (4) faster algorithms for the interesting k-Mismatch Constellation problem.

Nick Fischer
INSAIT, University of Sofia
nick.fischer@insait.ai

Ce Jin
EECS
MIT
cejin@mit.edu

Yinzhan Xu
Massachusetts Institute of Technology
xyzhan@mit.edu

CP40

Beating Bellman's Algorithm for Subset Sum

Bellman's algorithm for Subset Sum is one of the earliest and simplest examples of dynamic programming, dating back to 1957. For a given set of n integers X and a target t, it computes the set of subset sums $\mathcal{S}(X,t)$ (i.e., the set of integers $s \in [0 \dots t]$ for which there is a subset of X summing to s) in time $O(|\mathcal{S}(X,t)| \cdot n)$. Since then, it has been an important question whether Bellman's seminal algorithm can be improved. This question is addressed in many recent works. And yet, while some algorithms improve upon Bellman's algorithm in specific parameter regimes, such as Bringmann's $\tilde{O}(t+n)$-time algorithm [SODA '17] and Bringmann and Nakos' $\tilde{O}(|\mathcal{S}(X,t)|^{4/3})$-time algorithm [STOC '20], none of the known algorithms beats Bellman's algorithm in all regimes. In particular, it remained open whether Subset Sum is in time $\tilde{O}(|\mathcal{S}(\mathcal{X},\sqcup)| \cdot \setminus^{\infty-\epsilon})$ (for some $\epsilon > 0$). In this work we positively resolve this question and design an algorithm that outperforms Bellman's algorithm in all regimes. Our algorithm runs in time $\widetilde{O}(|\mathcal{S}(X,t)| \cdot \sqrt{n})$, thus improving the time complexity by a factor of nearly $\sqrt{n}$. Our key innovation is the use of a result from additive combinatorics, which has not been applied in an algorithmic context before and which we believe to be of further independent interest for algorithm design.

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

Nick Fischer
INSAIT, University of Sofia
nick.fischer@insait.ai

Vasileios Nakos
National and Kapodistrian University of Athens and

Archimede
s / Athena RC
vasilisnak@di.uoa.gr

CP41

Finding Longer Cycles Via Shortest Colourful Cycle

We consider the parameterised k, e-Long Cycle problem, in which you are given an n-vertex undirected graph G, a specified edge e in G, and a positive integer k, and are asked to decide if the graph G has a simple cycle through e of length at least k. We show how to solve the problem in $1.731^k \operatorname{poly}(n)$ time, improving over the $2^k \operatorname{poly}(n)$ time algorithm by [Fomin et al., TALG 2024], but not the more recent $1.657^k \operatorname{poly}(n)$ time algorithm by [Eiben, Koana, and Wahlstrm, SODA 2024]. When the graph is bipartite, we can solve the problem in $2^{k/2} \operatorname{poly}(n)$ time, matching the fastest known algorithm for finding a cycle of length exactly k in an undirected bipartite graph [Bjrklund et al., JCSS 2017].

Andreas Björklund, Thore Husfeldt
IT University of Copenhagen
anbjo@itu.dk, thore@itu.dk

CP41

Connectivity Carcass of a Vertex Subset in a Graph - Both Odd and Even Case

Let $G = (V, E)$ be an undirected unweighted multi-graph and $S \subseteq V$ be a subset of vertices called the Steiner set. A set of edges with the least cardinality whose removal disconnects S, that is, there is no path between at least one pair of vertices from S, is called a Steiner mincut for S or simply an S-mincut. Connectivity Carcass is a compact data structure storing all S-mincuts in G introduced by Dinitz and Vainshtein in an extended abstract in the Proceedings of ACM STOC 1994. The complete proof of various results of this data structure for the simpler case when the value of S-mincut is odd appeared in an article in SICOMP in the year 2000. Over the last couple of decades, there have been attempts towards the proof for the case when the value of S-mincut is even, but none of them met a logical end. We present the following results. 1. We present the first complete, self-contained, and peer-reviewed exposition of the connectivity carcass which covers both even and odd cases of value of S-mincut. 2. We derive the results using an alternate and much simpler approach. In particular, we derive the results using submodularity of cuts and its generalizations – a well-known property of graphs expressed using a simple inequality. 3. We also show how the connectivity carcass can be helpful in efficiently answering some basic queries related to S-mincuts using some additional insights.

Surender Baswana
IIT Kanpur
sbaswana@gmail.com

Abhyuday Pandey
Indian Institute of Technology, Kanpur
pandey.abhyuday07@gmail.com

CP41

Simpler Optimal Sorting from a Directed Acyclic Graph

Given are a ground set X, some partial order P over X, and some oracle O_L that specifies a linear order L extending P. A *query* to O_L receives $x, x' \in X$ and outputs whether $x <_L x'$ or vice versa. If we denote by $e(P)$ the number of linear orders that extend P, then it follows from basic information theory that $\log e(P)$ is a worst-case lower bound on the number of queries needed to output the sorted order of X. Haeupler, Hladk, Iacono, Rozhon, Tarjan, and Tětek (SODA'25) propose to assume as input a directed acyclic graph, G, with m edges and $n = |X|$ vertices. Denote by P_G the partial order induced by G. Their algorithmic performance is measured in running time and the number of queries used, where they use $\Theta(m+n+\log e(P_G))$ time and $\Theta(\log e(P_G))$ queries to output X in its sorted order. Their analysis uses sophisticated counting arguments, entropy, and recursively defined sets defined over the run of their algorithm. We do away with sophistication. We show that when the input is a directed acyclic graph then the problem admits a simple solution using $\Theta(m + n + \log e(P_G))$ time and $\Theta(\log e(P_G))$ queries. Especially our proofs are much simpler as we avoid the usage of advanced charging arguments, and instead rely upon two observations.

Ivor Van Der Hoog, Eva Rotenberg
Technical University of Denmark
idjva@dtu.dk, eva@rotenberg.dk

Daniel P. Rutschmann
DTU Compute, Technical University of Denmark
daru@dtu.dk

CP41

Connectivity Certificate Against Bounded-Degree Faults: Simpler, Better and Supporting Vertex Faults

An f-edge (or vertex) connectivity certificate is a sparse subgraph that maintains connectivity under the failure of at most f edges (or vertices). It is well known that any n-vertex graph admits an f-edge (or vertex) connectivity certificate with $\Theta(fn)$ edges (Nagamochi and Ibaraki, Algorithmica 1992). A recent work by (Bodwin, Haeupler and Parter, SODA 2024) introduced a new and considerably stronger variant of connectivity certificates that can preserve connectivity under any failing set of edges with bounded degree. For every n-vertex graph $G = (V, E)$ and a degree threshold f, an f-Edge-Faulty-Degree (EFD) certificate is a subgraph $H \subseteq G$ with the following guarantee: For any subset $F \subseteq E$ with $\deg(F) \leq f$ and every pair $u, v \in V$, u and v are connected in $H - F$ iff they are connected in $G - F$. For example, a In their work, [BHP'24] presented an expander-based approach (e.g., using the tools of expander decomposition and expander routing) for computing f-EFD certificates with $O(fn \cdot poly(\log n))$ edges. They also gave a lower bound of $\Omega(fn \cdot \log_f n)$. In this work, we settle the optimal existential size bounds for f-EFD certificates (up to constant factors), and also ex-

tend it to support vertex failures with bounded degrees. Specifically, we show that when $n > f/2$, any n-vertex graph admits an f-EFD certificate with $O(fn \cdot \log(n/f))$ edges.

Elad Tzalik, Merav Parter
Weizmann Institute
elad.tzalik@weizmann.ac.il,
merav.parter@weizmann.ac.il

CP41

A Simplified Parameterized Algorithm for Directed Feedback Vertex Set

no text

Ziliang Xiong
Linkoping University
ziliang.xiong@liu.se

Mingyu Xiao
University of Electronic Science and Technology, China
myxiao@gmail.com

CP42

Coresets for Constrained Clustering: General Assignment Constraints and Improved Size Bounds

Designing small-sized *coresets*, which approximately preserve the costs of the solutions for large datasets, has been an important research direction for the past decade. We consider coreset construction for a variety of general constrained clustering problems. We introduce a general class of assignment constraints, including capacity constraints on cluster centers, and assignment structure constraints for data points (modeled by a convex body $\mathcal{B}$). We give coresets for clustering problems with such general assignment constraints that significantly generalize and improve known results. Notable implications include the first ε-coreset for capacitated and fair k-Median with m outliers in Euclidean spaces whose size is $\tilde{O}(m + k^2\varepsilon^{-4})$, generalizing and improving upon the prior bounds in [Braverman et al., FOCS' 22; Huang et al., ICLR' 23] (for capacitated k-Median, the coreset size bound obtained in [Braverman et al., FOCS' 22] is $\tilde{O}(k^3\epsilon^{-6})$, and for k-Median with m outliers, the coreset size bound obtained in [Huang et al., ICLR' 23] is $\tilde{O}(m + k^3\epsilon^{-5})$), and the first ϵ-coreset of size $\text{poly}(k\epsilon^{-1})$ for fault-tolerant clustering for various types of metric spaces.

Lingxiao Huang, Jian Li
IIIS, Tsinghua University
huanglingxiao1990@126.com, lapordge@gmail.com

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

Xuan Wu
Nanyang Technological University
wu3412790@gmail.com

CP42

Inapproximability of Maximum Diameter Clustering for Few Clusters

In the Max-k-Diameter problem, we are given a set of points in a metric space, and the goal is to partition the input points into k parts such that the maximum pairwise distance between points in the same part of the partition is minimized. The approximability of the Max-k-Diameter problem was studied in the eighties, culminating in the work of Feder and Greene [STOC'88], wherein they showed it is NP-hard to approximate within a factor better than 2 in the ℓ_1 and ℓ_∞ metrics, and NP-hard to approximate within a factor better than 1.969 in the Euclidean metric. This complements the celebrated 2 factor polynomial time approximation algorithm for the problem in general metrics (Gonzalez [TCS'85]; Hochbaum and Shmoys [JACM'86]). Over the last couple of decades, there has been increased interest from the algorithmic community to study the approximability of various clustering objectives when the number of clusters is fixed. In this setting, the framework of coresets has yielded PTAS for most popular clustering objectives, including k-means, k-median, k-center, k-minsum, and so on. In this paper, rather surprisingly, we prove that even when $k = 3$, the Max-k-Diameter problem is NP-hard to approximate within a factor of 1.5 in the ℓ_1-metric (and Hamming metric) and NP-hard to approximate within a factor of 1.304 in the Euclidean metric.

Ashwin Padaki
Columbia University
apadaki@seas.upenn.edu

Henry Fleischmann
Carnegie Mellon University
hfleischmann3@gmail.com

Kyrylo Karlov
Charles University
kirill.karlov1@gmail.com

Karthik C. S.
Rutgers University
karthik.cs@rutgers.edu

Stepan ZHARKOV
Columbia University
styopa@cs.columbia.edu

CP42

A Tight Vc-Dimension Analysis of Clustering Coresets with Applications

We consider coresets for k-median problems, where the goal is to assign points to centers minimizing the sum of distances. Given a point set P, a coreset Ω is a small weighted subset that approximates the cost of P for all candidate solutions up to a $(1 \pm \varepsilon)$ multiplicative factor. In this paper, we give a sharp VC-dimension based analysis for k-median coreset construction. As a consequence, we

obtain improved k-median coreset bounds for planar graph metrics and Frechet metrics.

Chris Schwiegelshohn
Aarhus University
cschwiegelshohn@gmail.com

Vincent Cohen-Addad
Google Research
cohenaddad@google.com

Andrew Draganov
Aarhus University, Denmark
raganovandrew@cs.au.dk

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

David Saulpic
ISTA
david.saulpic@lip6.fr

CP42

Gains-from-Trade in Bilateral Trade with a Broker

We study bilateral trade with a broker, where a buyer and seller interact exclusively through the broker. The broker strategically maximizes her payoff through arbitrage by trading with the buyer and seller at different prices. We study whether the presence of the broker interferes with the mechanism's gains-from-trade (GFT) achieving a constant-factor approximation to the first-best gains-from-trade (FB). We first show that the GFT achieves a 1/36-approximation to the FB even if the broker runs an optimal posted-pricing mechanism under symmetric agents with monotone-hazard-rate distributions. Beyond posted-pricing mechanisms, even if the broker uses an arbitrary incentive-compatible (IC) and individually-rational (IR) mechanism that maximizes her expected profit, we prove that it induces a 1/2-approximation to the first-best GFT when the buyer and seller's distributions are uniform distributions with arbitrary supports. This bound is shown to be tight. We complement such results by proving that if the broker uses an arbitrary profit-maximizing IC and IR mechanism, there exists a family of problem instances under which the approximation factor to the first-best GFT becomes arbitrarily close to zero. We show that this phenomenon persists even if we restrict one of the buyer's or seller's distributions to have a singleton support, or even in the symmetric setting where the buyer and seller have identical distributions.

Suho Shin
University of Maryland
suhoshin@umd.edu

Ilya Hajiaghayi
Takoma Park Middle School
ihajiaghayi@gmail.com

MohammadTaghi Hajiaghayi, Gary Peng
University of Maryland
hajiagha@umd.edu, gpeng1@terpmail.umd.edu

CP42

Efficient Approximation Algorithm for Computing Wasserstein Barycenter under Euclidean Metric

Given a set of probability distributions, in the Wasserstein barycenter problem, one wishes to compute a distribution that minimizes the average Wasserstein distance, or optimal transport cost, from all the input distributions. Wasserstein barycenters preserve common geometric features of the input distributions making them useful in machine learning and data analytics tasks. We present a near-linear time algorithm that computes the Wasserstein barycenter within a relative $(1+\varepsilon)$-approximation in fixed dimensions. Much of the prior work on computing Wasserstein barycenters has focused on the design of such algorithms that compute only an additively approximate Wasserstein barycenter over a collection of distributions. To obtain our results, we first present a dynamic-programming-based algorithm to solve the primal and the dual formulation of the exact Wasserstein barycenter under tree metrics in near-linear time. We then combine our tree-based algorithms with the boosting framework to obtain a $(1+\varepsilon)$-approximation in near-linear time.

Pankaj K. Agarwal
Duke University
pankaj@duke.edu

Sharath Raghvendra
Department of Computer Science
North Carolina State University
skraghve@ncsu.edu

Pouyan Shirzadian
Virginia Tech University
pshirzadian@vt.edu

Keegan Yao
Department of Computer Science
Duke University
keegan.yao@duke.edu ¡keegan.yao@duke.edu¿

CP43

Faster Vizing and Near-Vizing Edge Coloring Algorithms

Vizing's celebrated theorem states that every simple graph with maximum degree Δ admits a $(\Delta+1)$ edge coloring which can be found in $O(m \cdot n)$ time on n-vertex m-edge graphs. After a series of simplifications and variations, this running time was eventually improved by Gabow, Nishizeki, Kariv, Leven, and Terada in 1985 to $O(m\sqrt{n \log n})$ time. This has effectively remained the state-of-the-art (except for a parallel and independent work that will be discussed in the talk). As our main result, we present a novel randomized algorithm that computes a $\Delta + O(\log n)$ coloring of any given simple graph in $O(m \log \Delta)$ expected time; in other words, a near-linear time randomized algorithm for a "near'-Vizing's coloring. As a corollary of this algorithm, we also obtain the following results: * A randomized algorithm for $(\Delta+1)$ edge

coloring in $O(n^2 \log n)$ expected time. This is near-linear in the input size for dense graphs and presents the first polynomial time improvement over the longstanding bounds of Gabow et al. for Vizing's theorem in almost four decades. * A randomized algorithm for $(1+\epsilon)\Delta$ edge coloring in $O(m \log(1/\epsilon))$ expected time for any $\epsilon = \omega(\log n/\Delta)$. The dependence on ϵ exponentially improves upon a series of recent results that obtain algorithms with runtime of $\Omega(m/\epsilon)$ for this problem.

Sepehr Assadi
Rutgers University, U.S.
sepehr@assadi.info

CP43

Even Faster (Delta + 1)-Edge Coloring Via Shorter Multi-Step Vizing Chains

Vizing's Theorem from 1964 states that any n-vertex m-edge graph with maximum degree Δ can be *edge colored* using at most $\Delta+1$ colors. For over 40 years, the state-of-the-art running time for computing such a coloring, obtained independently by Arjomandi [1982] and by Gabow, Nishizeki, Kariv, Leven and Terada [1985], was $\tilde{O}(m\sqrt{n})$. Very recently, this time bound was improved in two independent works, by Bhattacharya, Carmon, Costa, Solomon and Zhang to $\tilde{O}(mn^{1/3})$, and by Assadi to $\tilde{O}(n^2)$. In this paper we present an algorithm that computes such a coloring in $\tilde{O}(mn^{1/4})$ time. Our key technical contribution is a subroutine for extending the coloring to one more edge within time $\tilde{O}(\Delta^2+\sqrt{\Delta n})$. The best previous time bound of any color extension subroutine is either the trivial $O(n)$, dominated by the length of a Vizing chain, or the bound $\tilde{O}(\Delta^6)$ by Bernshteyn [2022], dominated by the length of *multi-step Vizing chains*, which is basically a concatenation of multiple (carefully chosen) Vizing chains. Our color extension subroutine produces significantly shorter multi-step Vizing chains than in previous works, for sufficiently large Δ.

Martin Costa
University of Warwick
martincosta2000@outlook.com

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

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

Tianyi Zhang
ETH Zurich
tianyi.zhang@inf.ethz.ch

CP43

Randomized Greedy Online Edge Coloring Succeeds for Dense and Randomly-Ordered Graphs

Vizing's theorem states that any graph of maximum degree Δ can be properly edge colored with at most $\Delta+1$ colors. In the online setting, it has been a matter of interest to find an algorithm that can properly edge color any graph on n vertices with maximum degree $\Delta = \omega(\log n)$ using at most $(1+o(1))\Delta$ colors. Here we study the nave random greedy algorithm, which simply chooses a legal color uniformly at random for each edge upon arrival. We show that this algorithm can $(1+\epsilon)\Delta$-color the graph for arbitrary ϵ in two contexts: first, if the edges arrive in a uniformly random order, and second, if the edges arrive in an adversarial order but the graph is sufficiently dense, i.e., $n = O(\Delta)$. Prior to this work, the random greedy algorithm was only known to succeed in trees. Our second result is applicable even when the adversary is adaptive, and therefore implies the existence of a deterministic edge coloring algorithm which $(1+\epsilon)\Delta$ edge colors a dense graph. Prior to this, the best known deterministic algorithm for this problem was the simple greedy algorithm which utilized $2\Delta-1$ colors.

Aditi Dudeja
University of Salzburg
aditi.dudeja@plus.ac.at

Rashmika Goswami, Michael Saks
Rutgers University
rg894@scarletmail.rutgers.edu, saks@math.rutgers.edu

CP43

A Sublinear-Time Algorithm for Nearly-Perfect Matchings in Regular Non-Bipartite Graphs

A breakthrough pair of papers by Goel, Kapralov, and Khanna [**?**, **?**] gave the first sublinear-time algorithms for finding large matchings in regular bipartite graphs. In particular, they gave an algorithm based on the idea of randomized depth-first search, that, for any d-regular bipartite graph, finds a perfect matching in $O(n \log n)$ time. (When $d = \omega(\log n)$, this is sublinear in the size of the graph.) We investigate whether large matchings can still be found without the assumption of bipartiteness. On the positive side, we present a randomized algorithm that finds a matching whose size is at least $\lceil \frac{n}{2}(1-\frac{1}{d+1}) - \frac{1}{\log n} \rceil$ on any d-regular graph, and whose expected running time is $O(n \log n)$. As is well known, there exist non-bipartite d-regular graphs whose maximal matching has size only $\frac{n}{2}(1-\frac{1}{d+1})$. Moreover, even when larger matchings do exist, they cannot, in general, be found in sublinear time. Specifically, we prove the lower bound: for all $\varepsilon > 0$, all $d = d(n)$, every algorithm that, given a d-regular graph G, outputs a matching whose size is at least $1-\frac{1-\varepsilon}{d+1}$ times the maximum matching size for G, must have expected running time $\Omega(dn)$ on worst-case inputs. Thus, in a certain precise sense, $1-1/(d+1)$ is the approximation threshold for maximum matching in sublinear time.

Thomas P. Hayes
University of New Mexico, USA
hayest@gmail.com

Varsha Dani
Rochester Institute of Technology

varshadani@gmail.com

CP43

Fully Dynamic $(\Delta + 1)$ Coloring Against Adaptive Adversaries

Over the years, there has been extensive work on fully dynamic algorithms for classic graph problems that admit greedy solutions. Examples include $(\Delta + 1)$ vertex coloring, maximal independent set, and maximal matching. For all three problems, there are randomized algorithms that maintain a valid solution after each edge insertion or deletion to the n-vertex graph by spending polylog(n) time, provided that the adversary is oblivious. However, none of these algorithms work against adaptive adversaries whose updates may depend on the output of the algorithm. In fact, even breaking the trivial bound of $O(n)$ against adaptive adversaries remains open for all three problems. In this paper, we break this linear barrier for the $(\Delta + 1)$ vertex coloring problem. Our algorithm is randomized, and maintains a valid $(\Delta+1)$ vertex coloring after each edge update by spending $\widetilde{O}(n^{8/9})$ time with high probability. To achieve this result, we build on a powerful sparse-dense decomposition of the literature. A major challenge in applying this framework to our setting is that it relies on maintaining a perfect matching of a certain graph. While maintaining a perfect matching (conditionally) requires $n^{1-o(1)}$ time per update, we prove several (of possible independent interest) structural properties of this graph to achieve our sublinear in n update-time.

Omer Wasim, Soheil Behnezhad, Rajmohan Rajaraman
Northeastern University
wasim.o@northeastern.edu,
s.behnezhad@northeastern.edu,
r.rajaraman@northeastern.edu

CP44

Differentiable Approximations for Distance Queries

The widespread use of gradient-based optimization has motivated the adaptation of various classical algorithms into differentiable solvers compatible with learning pipelines. In this paper, we investigate the enhancement of traditional geometric query problems such that the result consists of both the geometric function as well as its gradient. Specifically, we study the fundamental problem of distance queries against a set of points P in $\mathbb{R}^d$, which also underlies various similarity measures for learning algorithms. The main result of this paper is a multiplicative $(1+\varepsilon)$-approximation of the Euclidean distance to P which is differentiable at all points in $\mathbb{R}^d \setminus P$ with asymptotically optimal bounds on the norms of its gradient and Hessian, from a data structure with storage and query time matching state-of-the-art results for approximate nearest-neighbor searching. The approximation is realized as a regularized distance through a partition-of-unity framework, which efficiently blends multiple local approximations, over a suitably defined covering of space, into a smooth global approximation. In order to obtain the local distance approximations in a manner that facilitates blending, we develop a new approximate Voronoi diagram based on a simple point-location data structure, simplifying away both the lifting transformation and ray shooting.

Ahmed Abdelkader, David M. Mount
University of Maryland
ahmadabdolkader@gmail.com, mount@umd.edu

CP44

Facet-Hamiltonicity

We consider *facet-Hamiltonian* cycles of polytopes, defined as cycles in their skeleton such that every facet is visited exactly once. These cycles can be understood as optimal *watchman routes* that guard the facets of a polytope. We consider the existence of such cycles for a variety of polytopes, the facets of which have a natural combinatorial interpretation. In particular, we prove the following results: (1) Every permutahedron has a facet-Hamiltonian cycle. With these cycles we associate what we call *rhombic strips* which encode interleaved Gray codes of the Boolean lattice, one Gray code for each rank. The rhombic strips also have interpretations as simple Venn diagrams. (2) Every generalized associahedron has a facet-Hamiltonian cycle. This generalizes the so-called *rainbow cycles* of Felsner, Kleist, Mtze, and Sering (SIDMA 2020) to associahedra of any finite type. We relate the constructions to the Conway-Coxeter friezes and the bipartite belts of finite type cluster algebras. (3) Graph associahedra of wheels, fans, and complete split graphs have facet-Hamiltonian cycles. We also consider the computational complexity of deciding whether a given polytope has a facet-Hamiltonian cycle and show that the problem is NP-complete, even when restricted to simple 3-dimensional polytopes.

Hugo A. Akitaya
Tufts University
hugo_akitaya@uml.edu

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

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

Linda Kleist
Universität Potsdam
kleist@cs.uni-potsdam.de

Robert Lauff
Technische Universität Berlin
lauff@math.tu-berlin.de

CP44

A Discrete Analog of Tuttes Barycentric Embeddings on Surfaces

Tuttes celebrated barycentric embedding theorem describes a natural way to build straight-line embeddings (crossing-free drawings) of a (3-connected) planar graph: map the vertices of the outer face to the vertices of a con-

vex polygon, and ensure that each remaining vertex is in convex position, namely, a barycenter with positive coefficients of its neighbors. Actually computing an embedding then boils down to solving a system of linear equations. A particularly appealing feature of this method is the flexibility given by the choice of the barycentric weights. Generalizations of Tuttes theorem to surfaces of nonpositive curvature are known, but due to their inherently continuous nature, they do not lead to an algorithm. In this paper, we propose a purely discrete analog of Tuttes theorem for surfaces (with or without boundary) of nonpositive curvature, based on the recently introduced notion of reducing triangulations. We prove a Tutte theorem in this setting: every drawing homotopic to an embedding such that each vertex is harmonious (a discrete analog of being in convex position) is a weak embedding (arbitrarily close to an embedding). We also provide a polynomial-time algorithm to make an input drawing harmonious without increasing the length of any edge, in a similar way as a drawing can be put in convex position without increasing the edge lengths.

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

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

Vincent Despré
Université de Lorraine, CNRS, Inria, LORIA, Nancy,
France
vincent.despre@loria.fr

CP44

Fréchet Distance in Subquadratic Time

Let m and n be the numbers of vertices of two polygonal curves in $\mathbb{R}^d$ for any fixed d such that $m \leq n$. Since it was known in 1995 how to compute the Fréchet distance of these two curves in $O(mn \log(mn))$ time, it has been an open problem whether the running time can be reduced to $o(n^2)$ when $m = \Omega(n)$. In the mean time, several well-known quadratic time barriers in computational geometry have been overcome: 3SUM, some 3SUM-hard problems, and the computation of some distances between two polygonal curves, including discrete Fréchet distance, dynamic time warping, and geometric edit distance. It is curious that the quadratic time barrier for Fréchet distance still stands. We present an algorithm to compute the Fréchet distance in $O(mn(\log\log n)^{2+\mu} \log n / \log^{1+\mu} m)$ expected time for some constant $\mu \in (0, 1)$. It is the first algorithm that returns the Fréchet distance in $o(mn)$ time when $m = \Omega(n^{\varepsilon})$ for any fixed $\varepsilon \in (0, 1]$.

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

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

CP44

Relating Interleaving and Frchet Distances Via Ordered Merge Trees

Merge trees are a common topological descriptor for data with a hierarchical component, such as terrains and scalar fields. The interleaving distance, in turn, is a common distance for comparing merge trees. However, the interleaving distance for merge trees is solely based on the hierarchical structure, and disregards any other geometrical or topological properties that might be present in the underlying data. Furthermore, the interleaving distance is NP-hard to compute. We introduce a form of ordered merge trees that can capture intrinsic order present in the data. We further define a natural variant of the interleaving distance, the monotone interleaving distance, which is an order-preserving distance for ordered merge trees. Analogously to the regular interleaving distance for merge trees, we show that the monotone variant has three equivalent definitions in terms of two maps, a single map, or a labelling. Furthermore, we establish a connection between the monotone interleaving distance of ordered merge trees and the Frchet distance of 1D curves. As a result, the monotone interleaving distance between two ordered merge trees can be computed exactly in time near-quadratic in their complexity. The connection between the monotone interleaving distance and the Frchet distance builds a new bridge between the fields of topological data analysis, where interleaving distances are a common tool, and computational geometry, where Frchet distances are studied extensively.

Thijs Beurskens
TU Eindhoven
t.p.j.beurskens@tue.nl

Tim Ophelders
TU Eindhoven
Utrecht University
t.a.e.ophelders@uu.nl

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

Kevin Verbeek
TU Eindhoven
k.a.b.verbeek@tue.nl

CP45

A Parametric Version of the Hilbert Nullstellensatz

Hilbert's Nullstellensatz is a fundamental result in algebraic geometry that gives a necessary and sufficient condition for a finite collection of multivariate polynomials to have a common zero in an algebraically closed field. The associated computational problem HN asks to determine whether a system of polynomials with coefficients in $\mathbb{Q}$ has a common zero over the field of algebraic numbers. In 1996, Koiran showed that HN lies in AM assuming the

Generalised Riemann Hypothesis (GRH). He later generalised this result by showing that the problem DIM, which asks to determine the dimension of the set of solutions of a given polynomial system, also lies in AM subject to GRH. We study the solvability of polynomial equations over arbitrary algebraically closed fields of characteristic zero. We formulate a parametric version of HN, called HNP, in which the input is a system of polynomials with coefficients in a function field $\mathbb{Q}(\mathbf{x})$ and the task is to determine whether the polynomials have a common zero in the algebraic closure $\overline{\mathbb{Q}(\mathbf{x})}$. We observe that Koiran's proof that DIM lies in AM can be interpreted as a randomised polynomial-time reduction of DIM to HNP, followed by an argument that HNP lies in AM. Our main contribution is a self-contained proof that HNP lies in AM that follows the same basic idea as Koiran's argument, namely random instantiation of the parameters, but whose justification is purely algebraic.

Klara Nosan
IRIF, Université Paris Cité
nosan@irif.fr

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

Nikhil Balaji
IIT Delhi
nbalaji@cse.iitd.ac.in

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

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

CP45

Revisiting Tree Canonization using polynomials

Graph Isomorphism (GI) is a fundamental algorithmic problem. Amongst graph classes for which the computational complexity of GI has been resolved, trees are arguably the most fundamental. Tree Isomorphism is complete for deterministic logspace, a tiny subclass of polynomial time, by Lindells result. Over three decades ago, he devised a deterministic logspace algorithm that computes a string which is a canon for the input tree two trees are isomorphic precisely when their canons are identical. Inspired by Miller-Reifs reduction of Tree Isomorphism to Polynomial Identity Testing, we present a new logspace algorithm for tree canonization fundamentally different from Lindells algorithm. Our algorithm computes a univariate polynomial as canon for an input tree, based on the classical Eisensteins criterion for the irreducibility of univariate polynomials. This can be implemented in logspace by invoking the well known Buss et al. algorithm for arithmetic formula evaluation. This algorithm is conceptually very simple, avoiding the delicate case analysis and complex recursion that constitute the core of Lindells algorithm. We illustrate the adaptability of our algorithm by extending it to a couple of other classes of graphs.

V. Arvind
Institute of Mathematical Sciences, Chennai
Chennai Mathematical Institute, Chennai
arvind@imsc.ac.in

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

SALMAN Faris
BITS Pilani, Hyderabad
salmanfaris2048@gmail.com

Asif Khan
Chennai Mathematical Institute, Chennai
asifkhan@cmi.ac.in

CP45

The Quasi-Probability Method and Applications for Trace Reconstruction

In the trace reconstruction problem, one attempts to reconstruct a fixed but unknown string x of length n from a given number of traces $\tilde{x}$ drawn iid from the application of a noisy process (such as the deletion channel) to x. The best known algorithm for the trace reconstruction from the deletion channel is due to Chase, and recovers the input string whp given $\exp(\tilde{O}(n^{1/5}))$ traces [Cha21b]. The main component in Chase's algorithm is a procedure for k-mer estimation, which, for any marker $w \in \{0,1\}^k$ of length k, computes a smoothed distribution of its appearances in the input string x [CGL+23, MS24]. Current k-mer estimation algorithms fail when the deletion probability is above $1/2$, requiring a more complex analysis for Chase's algorithm. Moreover, the only known extension of these approaches beyond the deletion channels is based on numerically estimating high-order differentials of a multi-variate polynomial, making it highly impractical [Rub23]. In this paper, we utilize an approach from the field of quantum error mitigation (the process of using many measurements from noisy quantum computers to simulate a clean quantum computer), called the quasi-probability method [TBG17, PSW22] to construct a simple Monte Carlo method for k-mer estimation which can be easily applied to a much wider variety of channels. Our algorithm is quantum-inspired, but no background in quantum computing is needed to understand this paper.

Ittai Rubinstein
Massachusetts Institute of Technology
ittai.rubinstein@gmail.com

CP45

Experimental Design Using Interlacing Polynomials

We present a unified deterministic approach for experimental design problems using the method of interlacing polynomials. Our framework recovers the best-known approximation guarantees for the well-studied D/A/E-design

problems with simple analysis. Furthermore, we obtain the first non-trivial approximation guarantee for E-design in a new regime. Additionally, our approach provides an optimal approximation guarantee for a generalized ratio objective that generalizes both D-design and A-design.

Lap Chi Lau, Robert Wang
University of Waterloo
lapchi@uwaterloo.ca, r585wang@uwaterloo.ca

Hong Zhou
Fuzhou University
hong.zhou@fzu.edu.cn

CP45

Faster Algorithms for Average-Case Orthogonal Vectors and Closest Pair Problems

We study the average-case version of the Orthogonal Vectors problem, in which one is given as input n vectors from $\{0,1\}^d$ which are chosen randomly so that each coordinate is 1 independently with probability p. Kane and Williams [ITCS 2019] showed how to solve this problem in time $O(n^{2-\delta_p})$ for a constant $\delta_p > 0$ that depends only on p. However, it was previously unclear how to solve the problem faster in the hardest parameter regime where p may depend on d. The best prior algorithm was the best worst-case algorithm by Abboud, Williams and Yu [SODA 2014], which in dimension $d = c \cdot \log n$, solves the problem in time $n^{2-\Omega(1/\log c)}$. In this paper, we give a new algorithm which improves this to $n^{2-\Omega(\log\log c/\log c)}$ in the average case for any parameter p. As in the prior work, our algorithm uses the polynomial method. We make use of a very simple polynomial over the reals, and use a new method to analyze its performance based on computing how its value degrades as the input vectors get farther from orthogonal. To demonstrate the generality of our approach, we also solve the average-case version of the closest pair problem in the same running time.

Josh Alman, Alexandr Andoni, Hengjie Zhang
Columbia University
josh@cs.columbia.edu, andoni@cs.columbia.edu, hz2613@columbia.edu

CP46

Dynamic Independent Set of Disks (and Hypercubes) Made Easier

Maintaining an approximate maximum independent set of a dynamic collection of objects has been studied extensively in the past few years. Recently, Bhore, Nöllenburg, Tóth, and Wulms (SoCG 2024) showed that for (unweighted) disks in the plane, it is possible to maintain $O(1)$-factor approximate solution in polylogarithmic amortized update time. In this work, we provide a *much simpler* dynamic $O(1)$-approximation algorithm, which at the same time improves the number of logarithmic factors in the previous update time bound. Along the way, we also obtain simpler and faster algorithms for dynamic independent set for axis-aligned hypercubes as well as static independent set for fat objects in any constant dimension. All the above results also hold for the minimum piercing set problem for the same classes of objects.

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

Timothy M. Chan
UIUC
tmc@illinois.edu

CP46

A Simple Partially Embedded Planarity Test Based on Vertex-Addition

In the Partially Embedded Planarity problem, we are given a graph G together with a topological drawing of a subgraph H of G. The task is to decide whether the drawing can be extended to a drawing of the whole graph such that no two edges cross. Angelini et al. gave a linear-time algorithm for solving this problem in 2010 (SODA '10). While their paper constitutes a significant result, the algorithm described therein is highly complex: it uses several layers of decompositions according to connectivity of both G and H, its description spans more than 30 pages, and can hardly be considered implementable. We give an independent linear-time algorithm that works along the well-known vertex-addition planarity test by Booth and Lueker. We modify the PC-tree as underlying data structure used for representing all planar drawing possibilities in a natural way to also respect the restrictions given by the prescribed drawing of the subgraph H. The testing algorithm and its proof of correctness only require small adaptations from the comparatively much simpler generic planarity test, of which several implementations exist. If the test succeeds, an embedding can be constructed using the same approaches that are used for the generic planarity test.

Simon D. Fink
TU Wien
sfink@ac.tuwien.ac.at

Ignaz Rutter
University of Passau
rutter@fim.uni-passau.de

Sandhya T P
Stockholm University
thekkumpadan@math.su.se

CP46

On Beating 2^n for the Closest Vector Problem

The Closest Vector Problem (CVP) is a computational problem in lattices that is central to modern cryptography.The study of its fine-grained complexity has gained momentum in the last few years, partly due to the upcoming deployment of lattice-based cryptosystems in practice. In this paper we show positive results for a natural special case of the problem that has hitherto seemed just as hard, namely $(0,1)$-CVP where the lattice vectors are restricted to be sums of subsets of basis vectors (meaning that all coefficients are 0 or 1). All previous hardness results applied to this problem, and none of the previous

algorithmic techniques could benefit from it. We prove the following results, which follow from new reductions from $(0,1)$-CVP to weighted Max-SAT and minimum-weight k-Clique. An $O(1.7299^n)$ time algorithm for exact $(0,1)$-CVP_2 in Euclidean norm, breaking the natural 2^n barrier, as long as the absolute value of all coordinates in the input vectors is $2^{o(n)}$. A computational equivalence between $(0,1)$-CVP_p and Max-p-SAT for all even p (a reduction from Max-p-SAT to $(0,1)$-CVP_p was previously known). The minimum-weight-k-Clique conjecture from fine-grained complexity and its numerous consequences (which include the APSP conjecture) can now be supported by the hardness of a lattice problem, namely $(0,1)$-CVP_2. Similar results also hold for the Shortest Vector Problem.

Rajendra Kumar
Indian Institute of Technology Delhi
rjndr2503@gmail.com

Amir Abboud
Weizmann Institute of Science and INSAIT, Sofia University
amir.abboud@gmail.com

CP46

An Optimal Algorithm for Half-Plane Hitting Set

Given a set P of n points and a set H of n half-planes in the plane, we consider the problem of computing a smallest subset of points such that each half-plane contains at least one point of the subset. The previously best algorithm solves the problem in $O(n^3 \log n)$ time. It is also known that $\Omega(n \log n)$ is a lower bound for the problem under the algebraic decision tree model. In this paper, we present an $O(n \log n)$ time algorithm, which matches the lower bound and thus is optimal. Another virtue of the algorithm is that it is relatively simple.

Gang Liu
Kahlert School of Computing, University of Utah
u0866264@utah.edu

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

CP46

Recursive Lattice Reduction-A Framework for Finding Short Lattice Vectors

We propose a new framework called recursive lattice reduction for finding short non-zero vectors in a lattice or for finding dense sublattices of a lattice. At a high level, the framework works by recursively searching for dense sublattices of dense sublattices (or their duals) with progressively lower rank. Eventually, the procedure encounters a recursive call on a lattice L with relatively low rank k, at which point we simply use a known algorithm to find a shortest non-zero vector in L. We view this new framework as complementary to basis reduction algorithms, which similarly work to reduce an n-dimensional lattice problem with some approximation factor γ to a lower-dimensional exact lattice problem in some lower dimension k. Our framework provides an alternative and arguably simpler perspective, which in particular can be described without explicitly referencing any specific basis or even representation of the lattice. We present a number of specific instantiations of our framework to illustrate its usefulness. Our main concrete result is an efficient reduction that matches the tradeoff between γ, n, and k achieved by the best-known basis reduction algorithms across all parameter regimes. In fact, this reduction also can be used to find dense *sublattices* with any rank ℓ satisfying $\min\{\ell, n-\ell\} \leq n-k+1$, using only an oracle for SVP in k dimensions, which is itself a novel result.

Divesh Aggarwal
National University of Singapore
dcsdiva@nus.edu.sg

Thomas Espitau
PQShield
t.espitau@gmail.com

Spencer Peters
Cornell University
sp2473@cornell.edu

Noah Stephens-Davidowitz
New York University
noahsd@gmail.com

CP47

Faster Approximation Algorithms for Restricted Shortest Paths in Directed Graphs

In the restricted shortest paths problem, we are given a graph G whose edges are assigned two non-negative weights: lengths and delays, a source s, and a delay threshold D. The goal is to find, for each target t, the length of the shortest (s,t)-path whose total delay is at most D. While this problem is known to be NP-hard [Garey and Johnson, 1979] $(1+\epsilon)$-approximate algorithms running in $\tilde{O}(mn)$ time [Goel et al., INFOCOM'01; Lorenz and Raz, Oper. Res. Lett.'01] given more than twenty years ago have remained the state-of-the-art for *directed* graphs. An open problem posed by [Bernstein, SODA'12] — who gave a randomized $m \cdot n^{o(1)}$ time bicriteria $(1+\epsilon, 1+\epsilon)$-approximation algorithm for undirected graphs — asks if there is similarly an $o(mn)$ time approximation scheme for directed graphs. We show two randomized bicriteria $(1+\epsilon, 1+\epsilon)$-approximation algorithms that give an affirmative answer to the problem: one suited to dense graphs, and the other that works better for sparse graphs. On directed graphs with a quasi-polynomial weights aspect ratio, our algorithms run in time $\tilde{O}(n^2)$ and $\tilde{O}(mn^{3/5})$ or better, respectively. More specifically, the algorithm for sparse digraphs runs in time $\tilde{O}(mn^{(3-\alpha)/5})$ for graphs with $n^{1+\alpha}$ edges for any real $\alpha \in [0, 1/2]$.

Vikrant Ashvinkumar
Rutgers
va264@cs.rutgers.edu

Aaron Bernstein
Rutgers University
bernstei@gmail.com

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

CP47

New Approximation Algorithms and Reductions for n-Pairs Shortest Paths and All-Nodes Shortest Cycles

In this paper, we focus on two related problems, the n-Pairs Shortest Paths problem and the All-Nodes Shortest Cycles problem. In the n-PSP problem, given a graph with n nodes, and a set P consisting of at most n pairs of nodes, our objective is to estimate the distances between each pair in P. In the ANSC problem, the objective is to find for each node the shortest cycle that includes that particular node. In both problems, we present new algorithms and reductions that enhance the existing solutions. We present the first reduction for undirected graphs between these problems that assumes an approximation algorithm: We prove that for unweighted graphs, given an algorithm for k-approximation n-PSP, then for any small constant $\epsilon > 0$, one can solve a $k+\epsilon$ approximation to the ANSC in comparable time. For the n-PSP problem, our main result breaks the multiplicative bound of the Thorup-Zwick oracle, and gets a $(\lceil 4k/3 \rceil - 1, \lceil 4k/3 \rceil - 1)$-approximation algorithm in unweighted graphs that runs in $O(kmn^{1/k} + k^2 n^{1+2/k})$ time. For the ANSC problem, we have developed a 4-approximation algorithm that operates in $O(m + n^{1+3/4})$ time by combining two algorithms. The first algorithm finds an almost 2-approximation in $O(n\sqrt{m})$ time. The second algorithm is a construction of a fast 2-multiplicative ANSC spanner a new type of spanner designed to preserve the lengths of the shortest cycles.

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

Itay Hoch
Tel-Aviv University
itayhoch1212@gmail.com

Gur Lifshitz
Tel Aviv University
gurlifshitz@gmail.com

CP47

Faster Single-Source Shortest Paths with Negative Real Weights Via Proper Hop Distance

The textbook algorithm for single-source shortest paths with real-valued edge weights runs in $O(mn)$ time on a graph with m edges and n vertices. A recent breakthrough algorithm by Fineman [Fin24] takes $\tilde{O}(mn^{8/9})$ randomized time. We present an $\tilde{O}(mn^{4/5})$ randomized time algorithm building on ideas from [Fin24].

Yufan Huang, Peter Jin, Kent Quanrud
Purdue University
huan1754@purdue.edu, jin453@purdue.edu, krq@purdue.edu

CP47

Improved Shortest Path Restoration Lemmas for Multiple Edge Failures: Trade-Offs Between Fault-Tolerance and Subpaths

The restoration lemma is a classic result by Afek, Bremler-Barr, Kaplan, Cohen, and Merritt [PODC '01], which describes how the structure of shortest paths in a graph can change when some edges in the graph fail. Their work shows that, after one edge failure, any replacement shortest path avoiding this failing edge can be partitioned into two pre-failure shortest paths. More generally, this implies an additive tradeoff between fault tolerance and subpath count: for any f, k, we can partition any f-edge-failure replacement shortest path into $k+1$ subpaths which are each an $(f-k)$-edge-failure replacement shortest path. This generalized version of the result has found applications in routing, graph algorithms, fault tolerant network design, and more. Our main result improves this to a multiplicative tradeoff between fault tolerance and subpath count. We show that for all f, k, any f-edge-failure replacement path can be partitioned into $O(k)$ subpaths that are each an (f/k)-edge-failure replacement path. We also show an asymptotically matching lower bound. In particular, our results imply that the original restoration lemma is exactly tight in the case $k = 1$, but can be significantly improved for larger k. We also show an extension of this result to weighted input graphs, and we give an efficient algorithms that computes path decompositions satisfying our improved restoration lemmas, which runs in near-linear time for fixed f.

Greg Bodwin, Lily Wang
University of Michigan
bodwin@umich.edu, lilyxy@umich.edu

CP47

All-Hops Shortest Paths

Let $G = (V, E, w)$ be a weighted directed graph without negative cycles. For two vertices $s, t \in V$, we let $d_{\leq h}(s, t)$ be the minimum, according to the weight function w, of a path from s to t that uses at most h edges, or hops. We consider algorithms for computing $d_{\leq h}(s, t)$ for every $1 \leq h \leq n$, where $n = |V|$, in various settings. We consider the single-pair, single-source and all-pairs versions of the problem. We also consider a distance oracle version of the problem in which we are not required to explicitly compute all distances $d_{\leq h}(s, t)$, but rather return each one of these distances upon request. We consider both the case in which the edge weights are arbitrary, and in which they are small integers in the range $\{-M, \ldots, M\}$. For some of our results we obtain matching conditional lower bounds.

Yinzhan Xu
Massachusetts Institute of Technology
xyzhan@mit.edu

Virginia Vassilevska Williams
MIT
virgi@mit.edu

Zoe Xi
Massachusetts Institute of Technology
zoexi@mit.edu

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

CP47

Fine-Grained Optimality of Partially Dynamic Shortest Paths and More

Single Source Shortest Paths (SSSP) is among the most well-studied problems in computer science. In the incremental (resp. decremental) setting, the goal is to maintain distances from a fixed source in a graph undergoing edge insertions (resp. deletions). A long line of research culminated in a near-optimal deterministic $(1+\varepsilon)$-approximate data structure with $m^{1+o(1)}$ total update time over all m updates by Bernstein, Probst Gutenberg and Saranurak [FOCS 2021]. However, there has been remarkably little progress on the exact SSSP problem beyond Even and Shiloach's algorithm [J. ACM 1981] for unweighted graphs. For weighted graphs, there are no exact algorithms beyond recomputing SSSP from scratch in $\widetilde{O}(m^2)$ total update time, even for the simpler Single-Source Single-Target Shortest Path problem (stSP). Despite this lack of progress, known (conditional) lower bounds only rule out algorithms with amortized update time better than $m^{1/2-o(1)}$ in dense graphs. In this paper, we give a tight (conditional) lower bound: any partially dynamic exact stSP algorithm requires $m^{2-o(1)}$ total update time for any sparsity m. We thus resolve the complexity of partially dynamic shortest paths, and separate the hardness of exact and approximate shortest paths, giving evidence as to why no non-trivial exact algorithms have been obtained while fast approximation algorithms are known.

Christopher Ye, Barna Saha
University of California, San Diego
czye@ucsd.edu, christopher.yeeeee@gmail.com

Virginia Vassilevska Williams
MIT
virgi@mit.edu

Yinzhan Xu
Massachusetts Institute of Technology
xyzhan@mit.edu

CP48

Optimal Mixing for Randomly Sampling Edge Colorings on Trees Down to the Max Degree

We address the convergence rate of Markov chains for randomly generating an edge coloring of a given tree. Our focus is on the Glauber dynamics which updates the color at a randomly chosen edge in each step. For a tree T with n vertices and maximum degree Δ, when the number of colors q satisfies $q \geq \Delta + 2$ then we prove that the Glauber dynamics has an optimal relaxation time of $O(n)$, where the relaxation time is the inverse of the spectral gap. This is optimal in the range of q in terms of Δ as Dyer, Goldberg, and Jerrum (2006) showed that the relaxation time is $\Omega(n^3)$ when $q = \Delta + 1$. For the case $q = \Delta + 1$, we show that an alternative Markov chain which updates a pair of neighboring edges has relaxation time $O(n)$. Moreover, for the Δ-regular complete tree we prove $O(n \log^2 n)$ mixing time bounds for the respective Markov chain. Our proofs establish approximate tensorization of variance via a novel inductive approach, where the base case is a tree of height $\ell = O(\Delta^2 \log^2 \Delta)$, which we analyze using a canonical paths argument.

Charlie A. Carlson
UCSB
charlieannecarlson@ucsb.edu

Xiaoyu Chen
Nanjing University
chenxiaoyu233@smail.nju.edu.cn

Weiming Feng
Institute for Theoretical Studies
ETH Zürich
weiming.feng@eth-its.ethz.ch

Eric Vigoda
University of California Santa Barbara
vigoda@ucsb.edu

CP48

Rényi-Infinity Constrained Sampling with D^3 Membership Queries

Uniform sampling over a convex body is a fundamental algorithmic problem, yet the convergence in KL or Rényi divergence of most samplers remains poorly understood. In this work, we propose a constrained proximal sampler, a principled and simple algorithm that possesses elegant convergence guarantees. Leveraging the uniform ergodicity of this sampler, we show that it converges in the Rényi-infinity divergence ($\mathcal{R}_\infty$) with no query complexity overhead when starting from a warm start. This is the strongest of commonly considered performance metrics, implying rates in $\{\mathcal{R}_q, \mathsf{KL}\}$ convergence as special cases. By applying this sampler within an annealing scheme, we propose an algorithm which can approximately sample ε-close to the uniform distribution on convex bodies in $\mathcal{R}_\infty$-divergence with $\widetilde{\mathcal{O}}(d^3 \operatorname{polylog} \frac{1}{\varepsilon})$ query complexity. This improves on all prior results in $\{\mathcal{R}_q, \mathsf{KL}\}$-divergences, without resorting to any algorithmic modifications or post-processing of the sample. It also matches the prior best known complexity in total variation distance.

Yunbum Kook
KAIST
Discrete Mathematics Group, Institute for Basic Science
yb.kook@gatech.edu

Matthew Zhang
University of Toronto
matthew.zhang@mail.utoronto.ca

CP48

FPTAS for Holant Problems with Log-Concave Sig-

natures

For an integer $b \geq 0$, a b-matching in a graph $G = (V, E)$ is a set $S \subseteq E$ such that each vertex $v \in V$ is incident to at most b edges in S. We design a fully polynomial-time approximation scheme (FPTAS) for counting the number of b-matchings in graphs with bounded degrees. Our FPTAS also applies to a broader family of counting problems, namely Holant problems with log-concave signatures. Our algorithm is based on Moitra's linear programming approach (JACM'19). Using a novel construction called the extended coupling tree, we derandomize the coupling designed by Chen and Gu (SODA'24).

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

Zhidan Li, Guoliang Qiu
Shanghai Jiao Tong University
yueenternal@sjtu.edu.cn, guoliang.qiu@sjtu.edu.cn

Chihao Zhang
John Hopcroft Center for Computer Science
Shanghai Jiao Tong University
chihao@sjtu.edu.cn

CP48

Potential Hessian Ascent: The Sherrington-Kirkpatrick Model

We present the first iterative spectral algorithm to find near-optimal solutions for a random quadratic objective over the discrete hypercube, resolving a conjecture of Subag. The algorithm is a randomized "Hessian ascent' in the solid cube, with the objective modified by subtracting an instance-independent potential function. These results lay the groundwork for (possibly) demonstrating low-degree sum-of-squares certificates over high-entropy step distributions for a relaxed version of the Parisi formula.

Juspreet Singh Sandhu
University of California Santa Cruz
jsinghsa@ucsc.edu

David Jekel
University of Copenhagen
daj@math.ku.dk

Jonathan Shi
University of California San Diego
jshi@cs.cornell.edu

CP48

Spectral Independence Beyond Total Influence on Trees and Related Graphs

We study how to establish spectral independence, a key concept in sampling, without relying on total influence bounds, by applying an approximate inverse of the influence matrix. Our method gives constant upper bounds on spectral independence for two well-studied Gibbs distributions known to have unbounded total influences: - The monomer-dimer model on graphs with large girth (including trees). Prior to our work, such results were only known for graphs with constant maximum degrees or infinite regular trees, as shown by Chen, Liu, and Vigoda (STOC '21). - The hardcore model on trees with fugacity $\lambda < e^2$, which significantly improves upon the current threshold $\lambda < 1.3$ proved by Efthymiou, Hayes, tefankovic, and Vigoda (RANDOM '23), and more interestingly, surpasses the long-standing $\lambda_r > e - 1$ lower bound for the reconstruction threshold on trees (by Martin '03). Consequently, we establish optimal $\Omega(n^{-1})$ spectral gaps of the Glauber dynamics for these models on arbitrary trees, regardless of the maximum degree Δ.

Xiaoyu Chen
Nanjing University
chenxiaoyu233@smail.nju.edu.cn

Xiongxin Yang
Northeast Normal University
yangxx500@nenu.edu.cn

Yitong Yin, Xinyuan Zhang
Nanjing University
yinyt@nju.edu.cn, zhangxy@smail.nju.edu.cn

CP48

Mean-Field Potts and Random-Cluster Dynamics from High-Entropy Initializations

A common obstruction to efficient sampling from high-dimensional distributions with Markov chains is the multimodality of the target distribution because they may get trapped far from stationarity. Still, one hopes that this is only a barrier to the mixing of Markov chains from worst-case initializations and can be overcome by choosing high-entropy initializations, e.g., a product or weakly correlated distribution. Ideally, from such initializations, the dynamics would escape from the saddle points separating modes quickly and spread its mass between the dominant modes with the correct probabilities. In this paper, we study convergence from high-entropy initializations for the random-cluster and Potts models on the complete graph—two extensively studied high-dimensional landscapes that pose many complexities like discontinuous phase transitions and asymmetric metastable modes. We study the Chayes–Machta and Swendsen–Wang dynamics for the mean-field random-cluster model and the Glauber dynamics for the Potts model. We sharply characterize the set of product measure initializations from which these Markov chains mix rapidly, even though their mixing times from worst-case initializations are exponentially slow. Our proofs require careful approximations of projections of high-dimensional Markov chains (which are not themselves Markovian) by tractable 1-dimensional random processes, followed by analysis of the latter's escape from saddle points separating stable modes.

Antonio Blanca
Penn State University
blanca@cse.psu.edu

Reza Gheissari

Northwestern University
gheissari@northwestern.edu

Xusheng Zhang
Penn State University
xzz5349@psu.edu

CP49

Low Degree Local Correction Over the Boolean Cube

In this work, we show that the class of multivariate degree-d polynomials mapping $\{0,1\}^n$ to any Abelian group G is locally correctable with $\widetilde{O}_d((\log n)^d)$ queries for up to a fraction of errors approaching half the minimum distance of the underlying code. In particular, this result holds even for polynomials over the reals or the rationals, special cases that were previously not known. Further, we show that they are locally list correctable up to a fraction of errors approaching the minimum distance of the code. These results build on and extend the prior work of Amireddy, Behera, Paraashar, Srinivasan, and Sudan [ABPSS24] (STOC 2024) who considered the case of linear polynomials ($d = 1$) and gave analogous results. Low-degree polynomials over the Boolean cube $\{0,1\}^n$ arise naturally in Boolean circuit complexity and learning theory, and our work furthers the study of their coding-theoretic properties. Extending the results of [ABPSS24] from linear polynomials to higher-degree polynomials involves several new challenges and handling them gives us further insights into properties of low-degree polynomials over the Boolean cube. For local correction, we construct a set of points in the Boolean cube that lie between two exponentially close parallel hyperplanes and is moreover an interpolating set for degree-d polynomials.

Prashanth Amireddy
Harvard University
pamireddy@g.harvard.edu

Amik Raj Behera, Manaswi Paraashar, Srikanth Srinivasan
University of Copenhagen
bamikraj@gmail.com, manaswi.isi@gmail.com, srsr@di.ku.dk

Madhu Sudan
Harvard
madhu@cs.harvard.edu

CP49

Quantum Locally Recoverable Codes

Classical locally recoverable codes, which permit highly efficient recovery from localized errors as well as global recovery from larger errors, provide some of the most useful codes for distributed data storage in practice. In this paper, we initiate the study of quantum locally recoverable codes (qLRCs). After defining quantum local recoverability, we provide an explicit construction of qLRCs based on the classical LRCs of Tamo and Barg (2014), which we show have (1) a close-to-optimal rate-distance tradeoff (near the Singleton bound), (2) an efficient decoder, and (3) permit good spatial locality in a physical implementation. The analysis for both the distance and the efficient decoding of these quantum Tamo-Barg (qTB) codes is significantly more involved than in the classical case. Nevertheless, we obtain close-to-optimal parameters by introducing a "folded" version of these qTB codes, which we analyze using a combination of algebraic techniques. We also present and analyze two additional more basic constructions, namely random qLRCs, and qLRCs from AEL distance amplification. We complement these constructions with Singleton-like bounds that show our qLRC constructions achieve close-to-optimal parameters. We also apply these results to obtain new Singleton-like bounds for quantum LDPC codes. We then show that even the weakest form of a stronger locality property, namely local correctability, is impossible quantumly.

Louis Golowich, Venkatesan Guruswami
UC Berkeley
lgolowich@berkeley.edu, venkatg@berkeley.edu

CP49

More Efficient Approximate k-Wise Independent Permutations from Random Reversible Circuits Via Log-Sobolev Inequalities

We prove that the permutation computed by a reversible circuit with $\widetilde{O}(nk \cdot \log(1/\epsilon))$ random 3-bit gates is ϵ-approximately k-wise independent. Our bound improves on currently known bounds in the regime when the approximation error ϵ is not too small, and is optimal up to logarithmic factors when ϵ is a constant. We obtain our results by analyzing the log-Sobolev constants of appropriate Markov chains rather than their spectral gaps. A corollary of our result concerns the *incompressibility of random reversible circuits* as pointed out by concurrent work of Chen et al. [CHH+24], who showed that a linear-in-k bound for a multiplicative approximation to a k-wise independent permutation implies the linear growth of circuit complexity (a generalization of Shannon's argument).

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

Lucas Gretta, Angelos Pelecanos
UC Berkeley
lucas_gretta@berkeley.edu, apelecan@berkeley.edu

CP49

Improved Explicit Near-Optimal Codes in the High-Noise Regimes

Error-correcting codes are fundamental objects designed to ensure the accurate transmission of data across channels subject to noise or adversarial errors. Depending on the situations, we may need either a uniquely decodable code or a list decodable code, with a favorable tradeoff, along with an efficient decoding algorithm. The well-known lower bounds on the trade-off between rate and distance (or list decoding radius) are the Gilbert-Varshamov bound and the list decoding capacity, both of which are achieved

by random codes. However, achieving these bounds with explicit constructions across various parameter regimes remains a significant challenge. In this talk, I will provide an overview of error-correcting codes and introduce an explicit construction technique called graph concatenated codes. Using this approach, I will present explicit constructions of uniquely decodable and list decodable codes for the high-noise regime, showing improvements in key parameters and more efficient algorithms. These constructions make use of dispersers and a new combinatorial object we refer to as "multi-set dispersers". I will also offer new insights into graph-based codes and their decoding strategies. This is joint work with Xin Li.

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

Songtao Mao
Department of Computer Science,
Johns Hopkins University
smao13@jhu.edu

CP49

Locally Testable Tree Codes

Tree codes (Schulman, STOC 93', IEEE Transactions on Information Theory 96') are codes designed for interactive communication. Encoding in a tree code is done in an online manner: the i-th codeword symbol depends only on the first i message symbols. Codewords should have good tree distance meaning that for any two codewords, starting at the first point of divergence, they should have large Hamming distance. We investigate whether tree codes can be made to be locally testable. That is, can a tester, given oracle access to an alleged codeword w of the tree code, decide whether w is a codeword or far from such, while only reading a sub-linear number of symbols from w. As the main result of this work, we construct, for any $r \geq 3$, a probabilistic tree code that is locally testable using $\tilde{O}(n^{2/r})$ queries. The tester accepts any codeword with probability 1 and rejects strings that are δ_r-far from the code with high probability, where $\delta_r < 1$ degrades with r. Our probabilistic notion of a tree code is a relaxation of the standard notion and allows the encoder to toss random coins. We require that encoded messages are far (in tree distance) from any possible encoding of any other message.

Tamer Mour, Alon Rosen
Bocconi University
tamer.mour@gmail.com, alon.rosen@unibocconi.it

Ron Rothblum
Technion
rothblum@cs.technion.ac.il

CP49

Hermitian Diagonalization in Linear Precision

This work presents an algorithm for Hermitian diagonalization running in near matrix multiplication time requiring only $2\lg(1/\epsilon) + O(\log(n) + \log\log(1/\epsilon))$ bits of precision. Despite the widespread, highly successful use of various algorithms for Hermitian diagonalization in practice, the literature long lacked rigorous guarantees of their performance in finite arithmetic. The recent work of Banks, Garza-Vargas, Kulkarni, and Srivastava (FOCS 2020) changed this by providing an algorithm for diagonalizing any matrix up to backward error, and proving it requires no more than $O(\log^4(n/\epsilon)\log(n))$ bits. This work improves upon their algorithm in the Hermitian setting to dramatically reduce the bit requirement.

Rikhav Shah
University of California, Berkeley
rikhav.shah@berkeley.edu