Abstracts (sorted by order of talks)

Speaker: Sanjeev Khanna
Title: Sublinear Algorithms for Hierarchical Clustering
Abstract: Hierarchical clustering is a popular unsupervised learning method for organizing data as a rooted tree structure that simultaneously clusters data at multiple levels of granularity. A well-studied recent objective function views the input as a weighted graph with edges indicating similarity between data points, and focuses on finding a tree that minimizes the cost of hierarchical partitioning. The resulting problem is NP-hard, and previous algorithms for approximating this objective require at least linear time/space. In this talk, we will consider algorithms for hierarchical clustering that use sublinear resources (space, time, and communication). Specifically, we will present sublinear algorithms for hierarchical clustering in the streaming model (space), in the query model (time), and in the massively parallel computation (MPC) model (communication). At the core of our algorithmic results is a connection between hierarchical clustering and a suitably relaxed notion of cut sparsifiers of graphs. Our main algorithmic contribution is then to show how cut sparsifiers of this relaxed form can be efficiently constructed in the query model and the MPC model. We complement our algorithmic results by establishing nearly matching lower bounds that rule out algorithms with better performance guarantees in each of these models.
This is joint work with Arpit Agarwal, Huan Li, and Prathamesh Patil.

Speaker: Clifford Stein
Title: Scheduling with Speed Predictions.
Abstract: Algorithms with predictions is a recent framework that has been used to overcome pessimistic worst-case bounds in incomplete information settings. In the context of scheduling, very recent work has leveraged machine-learned predictions to design algorithms that achieve improved approximation ratios in settings where the processing times of the jobs are initially unknown. In this paper, we study the speed-robust scheduling problem where the speeds of the machines, instead of the processing times of the jobs, are unknown and augment this problem with predictions.
In this talk, we give an algorithm that simultaneously achieves, for any x < 1, a 1 + x approximation when the predictions are accurate and a
2+ 2/x approximation when the predictions are not accurate. We also study special cases and evaluate our algorithms performance as a function of the error.
Joint work with Eric Balanski, TingTing Ou and Hao-Ting Wei, all at Columbia.

Speaker: Stefan Weltge
Title: Binary scalar products
Abstract: Let $A,B \subseteq \mathbb{R}^d$ both span $\mathbb{R}^d$ such that $\langle a,b \rangle \in \{0,1\}$ holds for all $a \in A$, $b \in B$. We show that $|A| \cdot |B| \le (d+1)2d$. This allows us to settle a conjecture by Bohn, Faenza, Fiorini, Fisikopoulos, Macchia, and Pashkovich (2015) concerning 2-level polytopes. Such polytopes arise in combinatorial optimization and have the property that for every facet-defining hyperplane H there is a parallel hyperplane H' such that $H \cup H'$ contain all vertices. The authors conjectured that for every d-dimensional 2-level polytope P the product of the number of vertices of P and the number of facets of P is at most $d2^{d+1}$, which we show to be true. This is joint work with Andrey Kupavskii.

Speaker: David Wajc
Title: Dynamic Matching with Better-than-2 Approximation in Polylogarithmic Update Time
Abstract: We present dynamic algorithms with polylogarithmic update time for estimating the size of the maximum matching of a graph undergoing edge insertions and deletions with approximation ratio strictly better than 2. Specifically, we obtain a $(1+1/\sqrt{2}+\varepsilon) \approx (1.707+\varepsilon)$
approximation in bipartite graphs and a $1.973+\varepsilon$ approximation in general graphs. We thus answer in the affirmative the major open question first posed in the influential work of Onak and Rubinfeld (STOC'10) and repeatedly asked in the dynamic graph algorithms literature. Our randomized algorithms also work against an adaptive adversary and guarantee worst-case polylog update time, both w.h.p.
Our algorithms are based on simulating new two-pass streaming matching algorithms in the dynamic setting. Our key new idea is to invoke the recent sublinear-time matching algorithm of Behnezhad (FOCS'21) to efficiently simulate the second pass of our streaming algorithms, while bypassing the well-known vertex-update barrier.

Speaker: Alberto Marchetti Spaccamela
Title: Scheduling DAGs: recent results and challenges
Abstract: The Directed Acyclic Graph (DAGs) is a popular representation to describe the structure of parallel applications and to model the execution of multi-threaded programs that is widely used in cloud computing and in real-time systems.
I will present recent results on DAG scheduling considering different models and focusing on complexity, approximation and discussing open issues motivated by multi-core architectures nowadays widely used for their increased performance over single-core processors.

Speaker: Jakub Tarnawski
Title: Online Edge Coloring via Tree Recurrences and Correlation Decay
Abstract: We give an online algorithm that with high probability computes an edge coloring using (e/(e-1) + o(1))*Delta colors on a graph with maximum degree Delta = omega(log n) under online edge arrivals against oblivious adversaries, making first progress on the conjecture of Bar-Noy, Motwani, and Naor in this general setting. Our algorithm is based on reducing to a matching problem on locally treelike graphs, and then applying a tree recurrences based approach for arguing correlation decay.
This is joint work with Janardhan Kulkarni, Yang P. Liu, Ashwin Sah, and Mehtaab Sawhney.

Speaker: Rasmus Kyng
Title: Maximum Flow and Minimum-Cost Flow in Almost-Linear Time.
Abstract: We give the first almost-linear time algorithm for computing exact maximum flows and minimum-cost flows on directed graphs. By well-known reductions, this implies almost-linear time algorithms for several problems including bipartite matching, optimal transport, and undirected vertex connectivity.
Our algorithm uses a new Interior Point Method (IPM) that builds the optimal flow as a sequence of an almost-linear number of approximate undirected minimum-ratio cycles, each of which is computed and processed very efficiently using a new dynamic data structure.
Our framework extends to give an almost-linear time algorithm for computing flows that minimize general edge-separable convex functions to high accuracy. This gives the first almost-linear time algorithm for several problems including entropy-regularized optimal transport, matrix scaling, p-norm flows, and Isotonic regression.
Joint work with Li Chen, Yang Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva.

Speaker: Robert Weismantel
Title: LP and IP over time
Abstract: A classical result of Ford and Fulkerson states that there exists a max flow over time that has a repeated structure most of the time.
We extend this result to general minimum cost flow over time problems.
We also present a general framework that connects these classical over time versions of flow problems with more general linear and integer optimization problems.
My talk is based on joint work with Miriam Schl"oter and partially with Martin Skutella, Miriam Schlöter and Joe Paat.

Speaker: Mohit Singh
Title: Determinant Maximization : Approximation and Estimation Algorithms
Abstract: In the determinant maximization problem, given a collection of vectors, we aim to pick a subset to maximize the determinant of a natural matrix associated with these vectors. The abstract problem captures problems in multiple areas including machine learning, statistics, convex geometry, Nash social welfare problem from algorithmic game theory and network design problems. We will survey the known results and techniques for the problem. The results vary from arbitrary good approximations to only estimation algorithms. The techniques used in these works vary from geometry of polynomials, sparse solutions to convex programming solutions. In this talk, we will focus on algorithms that adapt the classical matroid intersection algorithms for this problem.

Speaker: Michel Goemans
Title: Brascamp-Lieb Polytopes
Abstract: Brascamp-Lieb polytopes arise from the validity of Brascamp-Lieb inequalities in analysis, and have interesting special cases including linear matroid intersection polytopes. I'll introduce them, discuss some of their properties, and present challenges and some algorithmic results using operator scaling.

Speaker: Jens Vygen
Title: "Packing cycles in planar and bounded-genus graphs”
vortragen.
Abstract and Coauthors: See https://arxiv.org/abs/2207.00450

Speaker: Seffi Naor
Title: Online Rounding of Bipartite Matchings
Abstract:Two complementary facets of the online bipartite matching problem are discussed. (1) For numerous online bipartite matching problems, such as edge-weighted matching and matching under two-sided vertex arrivals, state-of-the-art fractional algorithms outperform their randomized integral counterparts. Thus, a natural question is whether we can achieve lossless online rounding of fractional solutions in this setting. Even though lossless online rounding is impossible in general, randomized algorithms do induce fractional algorithms of the same competitive ratio, which by definition are losslessly roundable online. This motivates the addition of constraints that decrease the ``online integrality gap'', thus allowing for lossless online rounding. We characterize a set of non-convex constraints which allow for such lossless online rounding and allow for better competitive ratios than yielded by deterministic algorithms. (2) In a different vein, we study the problem of rounding fractional bipartite matchings in online settings. We assume that a fractional solution is already generated for us online by a black box (via a fractional algorithm, or some machine-learned advice) and provided as part of the input, which we then wish to round. We provide improved bounds on the rounding ratio and discuss several applications.
Based on joint papers with Niv Buchbinder, Aravind Srinivasan, and David Wajc.

Speaker: Alantha Newman
Title: Correlated Rounding for Correlation Clustering
Abstract: Given a complete graph G = (V, E) where each edge is labeled + or -, the correlation clustering problem asks to partition V into clusters to minimize the number of +edges between different clusters plus the number of -edges within the same cluster. The approximability of correlation clustering has been actively investigated [BBC04, CGW05, ACN08], culminating in a 2.06-approximation algorithm [CMSY15], based on rounding the standard LP relaxation. Since the integrality gap for this formulation is 2, it has remained a major open question to determine if the approximation factor of 2 can be reached, or even breached. In this talk, we show how to achieve a factor of 2+eps based on O(1/eps^2) rounds of the Sherali-Adams hierarchy. To round this relaxation, we adapt the correlated rounding originally developed for CSPs [BRS11, GS11, RT12]. To go below this approximation ratio, we go beyond the traditional triangle-based analysis by employing a global charging scheme that amortizes the total cost of the rounding across different triangles. This is joint work with Vincent Cohen-Addad and Euiwoong Lee.

Speaker: Laura Sanità
Title: On the Simplex method for 0/1 polytopes.
Abstract: The Simplex method is one of the most popular algorithms for solving linear programs, but despite decades of study, it is still not known whether there exists a pivot rule that guarantees it will always reach an optimal solution with a polynomial number of steps.
In fact, a polynomial pivot rule is not even known for linear programs over 0/1 polytopes (0/1-LPs), despite the fact that the diameter of a 0/1 polytope is bounded by its dimension.
This talk will focus on the behavior of the Simplex method for 0/1-LPs, and discuss pivot rules that are guaranteed to require only a polynomial number of non-degenerate pivots.
Joint work with: Alexander Black, Jesus De Loera, Sean Kafer.

Speaker: Joakim Blikstad
Title: Nearly Optimal Communication and Query Complexity of Bipartite Matching
Abstract: With a simple application of the cutting planes method, we settle the complexities of the bipartite maximum matching problem (BMM) up to poly-logarithmic factors in five models of computation: the two-party communication, AND query, OR query, XOR query, and quantum edge query models. Our results answer open problems that have been raised repeatedly since at least three decades ago [Hajnal, Maass, and Turan STOC'88; Ivanyos, Klauck, Lee, Santha, and de Wolf FSTTCS'12; Dobzinski, Nisan, and Oren STOC'14; Nisan SODA'21] and tighten the lower bounds shown by Beniamini and Nisan [STOC'21] and Zhang [ICALP'04]. Our communication protocols also work for some generalizations of BMM, such as maximum-cost bipartite b-matching and transshipment, using only $\tilde{O}(|V|)$ bits of communications.
To appear in FOCS'22. Joint work with Jan van den Brand, Yuval Efron, Danupon Nanongkai, and Sagnik Mukhopadhyay.

Speaker: Daniel Dadush
Title: Interior point methods are not worse than Simplex
Abstract: Whereas interior point methods provide polynomial-time linear programming algorithms, the running time bounds depend on bit-complexity or condition measures that can be unbounded in the problem dimension. This is in contrast with the classical simplex method that always admits an exponential bound. We introduce a new polynomial-time path-following interior point method where the number of iterations also admits a combinatorial upper bound O(2^n n^{1.5} log n) for an n-variable linear program in standard form. This complements previous work by Allamigeon, Benchimol, Gaubert, and Joswig (SIAGA 2018) that exhibited a family of instances where any path-following method must take exponentially many iterations.
The number of iterations of our algorithm is at most O(n^{1.5} log n) times the number of segments of any piecewise linear curve in the wide neighborhood of the central path. In particular, it matches the number of iterations of any path following interior point method up to this polynomial factor.
The overall exponential upper bound derives from studying the ‘max central path’, a piecewise-linear curve with the number of pieces bounded by the total length of 2n shadow vertex simplex paths.
This is joint work with Xavier Allamigeon (INRIA / Ecole Polytechnique), Georg Loho (U. Twente), Bento Natura (LSE), Laszlo Vegh (LSE).

Speaker: Danupon Nanongkai
Title: Negative-Weight Single-Source Shortest Paths in Near-linear Time
Abstract: We present a randomized algorithm that computes single-source shortest paths (SSSP) in O(m log^8(n) log W) time when edge weights are integral and can be negative and are >=-W. This essentially resolves the classic negative-weight SSSP problem. In contrast to all recent developments that rely on sophisticated continuous optimization methods and dynamic algorithms, our algorithm is simple: it requires only a simple graph decomposition and elementary combinatorial tools.

Speaker: András Sebő
Title: Duality Ratios for Hitting and Coloring Rectangles
Abstract : According to a celebrated conjecture of Wegner's (1965), the hitting number of axis-parallel rectangles in the plane is less than twice their packing number. No constant bound has been established for the hitting/packing quotient, but an example of Jelínek [in Latex: Jel\'{\i}nek] shows that the two of Wegner's conjecture cannot be improved. Ahlswede and Karapetyan bound this quotient as a linear function of the "aspect ratio" (the larger side of a rectangle divided by the smaller one), measuring how close a rectangle is to a square.
First, we improve Ahlswede and Karapetyan's bound from linear to logarithmic dependence on the aspect ratios. We then raise accordingly modified Wegner type conjectures for squares and unit squares, and also when these are not necessariy axis-parallel. The small constant bounds for the latter require more sophisticated geometric methods, and are still far from the lower bounds presented by the best examples. Similar problems and results can be stated for the chromatic number of rectangles and of squares. Some of the most frustrating, simply stated open challenges will also be shared with the audience.
This is joint work with Marco Caoduro

Speaker: Shay Moran

Title: A Combinatorial Characterization of Minimax in 0/1 Games

Abstract: We will discuss a generalization of the celebrated Minimax Theorem (von Neumann, 1928) for binary zero-sum games.

A simple game which fails to satisfy Minimax is Ephraim Kishon's “Jewish Poker” (see http://www.ephraimkishon.de/en/my_favorite_stories.htm).

In this game, each player picks a number and the larger number wins.

The payoff matrix in this game is *infinite triangular*.

We show this is the only obstruction:

if a game does not contain triangular submatrices of unbounded sizes then the Minimax Theorem holds.

This generalizes von Neumann's Minimax Theorem by removing requirements of finiteness or compactness.

Speaker: Leen Stougie

Title: The Online Traveling Salesman Problem with Predictions, including a general prediction error.

Abstract: In online optimization input arrives over time or one-by-one and an algorithm needs to make decisions without knowledge on future requests. The performance of online algorithms is typically assessed by worst-case competitive analysis.

The assumption in online optimization of not having any prior knowledge about future requests seems overly pessimistic. In particular, in the realm of machine-learning methods and data-driven applications, one may expect to have access to predictions about future requests. However, simply trusting such predictions might lead to very poor solutions, as these predictions come with no quality guarantee. A recent vibrant line of research aims at incorporating such error-prone predictions into online algorithms, to go beyond worst-case barriers. The goal are learning-augmented algorithms with performance that is close to that of an optimal offline algorithm when given accurate predictions (called consistency) and, at the same time, never being (much) worse than that of a best known algorithm without access to predictions (called robustness). Further, the performance of an algorithm shall degrade in a controlled way with increasing prediction error.

In this lecture I will show results within this framework for the online traveling salesman problem (OLTSP). I will show the typical ingredients that such analysis requires and the typical statements about performance that one may expect to see. One of the ingredients is a proper definition of a prediction error. For the OLTSP we have devised such an error measure, the basis of which is broadly applicable.

This is joint work with Giulia Bernardini (UTrieste), Alexander Lindmayer (UBremen), Alberto Marchetti Spaccamela (LSURoma), Nicole Megow (UBremen), Michelle Sweering (CWI).

Speaker: Uriel Feige

Title:On maximin fair allocation of indivisible items to three agents

Abstract: We consider the allocation of indivisible goods to n agents with additive valuation functions, in a setting with no money. A natural fairness requirement is that every agent will get a bundle that she values at least as high as her maximin share (MMS), namely, the value she could guarantee to herself if she were to partition the items into n bundles of her choice, and receive the least valuable of these bundles.

Unfortunately, for three or more agents, there are allocation instances in which no MMS allocation exists.

For the case of three agents, we present sufficient conditions for an MMS allocation to exist. We also consider ρ_3, the largest approximation ratio for which there always is an allocation giving each of the three agents at least a ρ_3 fraction of her MMS. We show that $11/12 \le \rho_3 \le 39/40$, and discuss the prospects of pinning down the exact value of ρ_3.

Based on joint works with Moshe Babaioff, Alexey Norkin, Ariel Sapir, Laliv Tauber.

Speaker: Thomas Rothvoss

Title: Approximate Carathéodory bounds via Discrepancy Theory

Abstract: The approximate Carath\'eodory problem in general form is as follows: Given two symmetric convex bodies $P,Q \subseteq \mathbb{R}^m$, a parameter k and $\mathbf{z} \in \textrm{conv}(X)$ with $X \subseteq P$, find $\mathbf{v}_1,\ldots,\mathbf{v}_k \in X$ so that $\|\mathbf{z} - \frac{1}{k}\sum_{i=1}^k

\mathbf{v}_i\|_Q$ is minimized. Maurey showed that if both P and Q coincide with the $\| \cdot \|_p$-ball, then an error of $O(\sqrt{p/k})$ is possible. We prove a reduction to the vector balancing constant from discrepancy theory which for most cases can provide tight bounds for general P and Q. For the case where P and Q are both $\| \cdot \|_p$-balls we prove an upper bound of $\sqrt{ \frac{\min\{ p, \log (\frac{2m}{k}) \}}{k}}$. Interestingly, this bound cannot be obtained taking independent random samples; instead we use the Lovett-Meka random walk. We also prove an extension to the more general case where P and Q are $\|\cdot \|_p$ and $\| \cdot \|_q$-balls with $2 \leq p \leq q \leq \infty$.
This is joint work with Victor Reis.

Speaker: Hang Zhou
Title: A PTAS for Capacitated Vehicle Routing on Trees
Abstract: We give a polynomial time approximation scheme (PTAS) for the unit demand capacitated vehicle routing problem (CVRP) on trees, for the entire range of the tour capacity. The result extends to the splittable CVRP. This is joint work with Claire Mathieu.

Speaker: Vera Traub
Title: Breaching the 2-Approximation Barrier for the Forest Augmentation Problem
Abstract: The goal of network design is to construct cheap networks that satisfy certain connectivity requirements. A celebrated result by Jain [Combinatorica, 2001] provides a 2-approximation algorithm for a wide class of these problems. However, even for many basic special cases nothing better is known.
In this talk we present the first better-than-2 approximation algorithm for one of the most basic problems in this family, the Forest Augmentation Problem: given an undirected unweighted graph (that w.l.o.g. is a forest) and a collection of extra edges (links), compute a minimum cardinality subset of links whose addition to the graph makes it 2-edge-connected.
This is joint work with Fabrizio Grandoni and Afrouz Jabal Ameli.