Fulkerson 100 program: lightning talks, Thursday July 18

4:15-4:22pm **Zhuan Khye Koh; Centrum Wiskunde & Informatica (CWI)**

Title: A strongly polynomial algorithm for the minimum cost generalized flow problem

Abstract: We give a strongly polynomial algorithm for minimum cost generalized flow, and as a consequence, for all linear programs with at most two variables per inequality. Previously, strongly polynomial algorithms were only known for the primal and dual feasibility problems. Our approach is to show that the path-following interior point method of Allamigeon et al. '22 terminates in a strongly polynomial number of iterations for minimum cost generalized flow. We achieve this by bounding the 'straight line complexity' of the central path, which is the minimum number of pieces required by a piecewise affine curve to multiplicatively approximate the central path.

Based on joint work with Daniel Dadush, Bento Natura, Neil Olver and László Végh.

4:22-4:29pm **Ronen Wdowinski; University of Waterloo**

Title: Bounded degree (hyper)graphs with no transversals

Abstract: For a multi-hypergraph G and a partition P of its vertex set (or edge set), there are various theorems stating that if the size of each class of P is sufficiently large compared to the maximum degree of G, then there exists an independent set (or matching) of G that is a transversal of P. These include Haxell's theorem (2001) on independent transversals, and a theorem of Aharoni, Berger, and Meshulam (2005) on full rainbow matchings. Generalizing previous constructions by Szabó and Tardos (2006) and others, we describe a general iterative method for producing multi-hypergraphs and partitions with no such transversals which are extremal for these theorems. We also describe applications of the method to list (edge) coloring based on color degree. This is based on joint work with Penny Haxell.

4:29-4:36pm **Luze Xu; UC Davis**

Title: Extended formulations for polyhedra with bimodular constraint matrices

Abstract: We are motivated by integer linear programs (ILPs) defined by constraint matrices with bounded determinants. Such matrices generalize the notion of totally-unimodular matrices.When the determinants are bounded by 2, the matrix is called bimodular. Artmann et al. give a polynomial-time algorithm for solving any ILP defined by a bimodular constraint matrix.Complementing this result, Conforti et al. give a compact extended formulation for a particular class of bimodular-constrained ILPs, namely those that model the stable set polytope of a graph with odd cycle packing number 1. We demonstrate that their compact extended formulation can be modified to hold for polyhedra such that (1) the constraint matrix is bimodular, (2) the row-matroid generated by the constraint matrix is cographic and (3) the right-hand side is a linear combination of the columns of the constraint matrix. This generalizes the important special case from Conforti et al. concerning 4-connected graphs with odd cycle transversal number at least four.

4:37-4:44pm **Youngho Yoo; Texas A&M University**

Title: Approximating TSP walks in subcubic graphs

Abstract: We prove that every subcubic graph on n vertices with n_2 vertices of degree 2 has a spanning closed walk of length at most $(5n + n_2)/4 - 1$, confirming a conjecture of Dvorák, Král, and Mohar. This bound is best possible and we characterize the extremal subcubic examples meeting this bound. We also give a quadratic time combinatorial algorithm to find such a TSP walk. In particular, we obtain a (5/4)-approximation algorithm for the Graphic TSP on cubic graphs.

Joint work with Michael Wigal and Xingxing Yu.

4:44-4:51pm **Mik Zlatin; CMU**

Title: The matroid base augmentation problem

Abstract: Imagine that we are given a base $B \subseteq E$ of a matroid M over ground set E. Each element of $E \setminus B$ has a cost, and our goal is to augment B by adding a minimum cost set $F \subseteq E \setminus B$ so that $B + F$ is a full rank set even after any element is deleted.
This problem captures the well-known Tree Augmentation Problem when M is graphic. We characterize the approximability of BAP for other common matroid classes, such as co-graphic, transversal, laminar and binary matroids. We end with a conjecture about BAP on regular matroids which may be of independent interest.

4:51-4:58pm **Noah Weninger; University of Waterloo**

Title: Interdiction of minimum spanning trees and other matroid bases

Abstract: In the minimum spanning tree (MST) interdiction problem, we are given a graph $G = (V, E)$ with edge weights, and want to find some $X \subseteq E$ satisfying a knapsack constraint such that the MST weight in $(V, E \setminus X)$ is maximized. Since MSTs of G are the minimum weight bases in the graphic matroid of G, this problem is a special case of matroid interdiction on a matroid $M = (E, I)$, in which the objective is instead to maximize the minimum basis weight in $M \setminus X$. By reduction from 0-1 knapsack, matroid interdiction is NP-complete, even for uniform matroids. We seek exact algorithms with good performance in practice. We present a dynamic programming upper bound which only requires that a simpler discrete derivative problem can be approximated for the given matroid. To get an exact algorithm, we use branch-and-bound with the entries in this dynamic programming table as the bounds for each node. For graphic matroids, we approximate the simpler problem by solving a sequence of minimum cut problems. For partition matroids, the dynamic program produces exact solutions in pseudopolynomial time. The running time of our algorithm is asymptotically faster than the best known MST interdiction algorithm, up to polylog factors. Furthermore, our algorithm can be adapted for any matroid, and achieves state-of-the-art computational performance: we solved all available MST interdiction instances from the literature, and in many cases reduced the best running time from hours to seconds.

5-5:07pm **Ndiamé Ndiaye; McGill University**

Title: Proving k-leaf powers have no simple characterization

Abstract: A graph $G = (V, E)$ is a k-leaf power if its vertices can be embedded in the leaves of a tree T such that (u, v) is an edge of G if and only if the distance between u and v in T is at most k. For $k < 5$, it was known that there is a simple characterization of k-leaf powers and it was long conjectured that such a characterization exists for any finite k. We prove that such a characterization cannot exist for k which is at least 5.

5:07-5:14pm **Joseph Poremba; University of British Columbia**

Title: Uncrossable multicommodity flows

Abstract: We investigate uncrossed multiflows, multiflows on planar graphs with the property that the curves identified by their flow paths do not cross in the plane. We demonstrate that given a fractional such multiflow, it can be rounded to an integral multiflow of congestion 2, motivating them as an interesting combinatorial object. We explore the complexity of computing uncrossed multiflows. On the one hand we can prove that determining whether a fractional uncrossed multiflow exists is NP-hard. In contrast, an integral (uncrossed) multiflow can be computed efficiently when the number of faces with demands is bounded.

5:14-5:21pm **Haripriya Pulyassary; Cornell University**

Title: Network flow problems with electric vehicles

Abstract: Electric vehicle (EV) adoption in long-distance logistics faces challenges such as range anxiety and uneven distribution of charging stations. Two pivotal questions emerge: How can EVs be efficiently routed in a charging network considering range limits, charging speeds and prices? And, can the existing charging infrastructure sustain the increasing demand for EVs in long-distance logistics? This talk addresses these questions by introducing an EV network flow model that incorporates range constraints and nonlinear charging rates. We identify conditions under which polynomial-time solutions can be obtained for optimal single EV routing, maximum flow, and minimum-cost flow problems.

This is joint work with Kostas Kollias, Aaron Schild, David Shmoys, and Manxi Wu.

5:22-5:29pm **Rebecca Whitman; UC Berkeley**

Title: Hereditary Nordhaus-Gaddum graphs

Abstract: Nordhaus and Gaddum proved in 1956 that the sum of the chromatic number χ of a graph G and its complement is at most $|G| + 1$. The Nordhaus-Gaddum graphs are the class of graphs satisfying this inequality with equality, and are well-understood. We consider a hereditary generalization: graphs G for which all induced subgraphs H of G satisfy $\chi(H) + \chi(\overline{H}) \geq |H|$. We characterize the forbidden induced subgraphs of this class and find its intersection with a number of common classes, including line graphs. We also discuss χ-boundedness and algorithmic results.

5:29-5:36pm **Rian Neogi, University of Waterloo**

Title: Budget-feasible mechanism design

Abstract: In the setting of budget feasible mechanism design, a buyer wants to purchase items from a set of sellers. Each seller i supplies one item, upon which they incur a cost of c_i. The buyer wants to buy a set of items from the sellers so as to maximize a certain valuation function, subject to a knapsack constraint stating that the total cost of the items bought is at most a certain budget. The true cost c_i is private information that is only known to seller i and not to the buyer. Thus, the buyer must work with the seller's reported costs. The goal is to design a mechanism that is truthful, in the sense that the sellers do not have an incentive to deviate from reporting their true costs, and budget feasible, in the sense that the total payments made to the sellers is within a given budget B.

In this talk, I will speak about our recent results in budget feasible mechanism design. When the valuation function is XOS, we design budget feasible mechanisms that are simpler and better compared to the state of art, both in terms of the approximation factors they obtain and in terms of requiring weaker oracle access to the valuation function. When the valuation function is subadditive, we give the first polynomial time explicitly constructed constant factor approximation to the problem, whereas previously only an existential result was known. We also study the setting in which sellers are allowed to supply multiple items to the buyer, which we dub multidimensional budget feasible mechanism design. In this setting, we find that there cannot be any constant factor approximate budget feasible mechanism with respect to the typical benchmark of the algorithmic optimum. We instead propose alternate benchmarks for this problem and show positive results with respect to these benchmarks.

This talk is based on joint work with Chaitanya Swamy and Kanstantsin Pashkovich.

5:36-5:43pm **Rhea Jain; University of Illinois, Urbana-Champaign**

Title: LP-based algorithms for two-cost network design

Abstract: In several standard network design problems, given an input graph with fixed edge costs and source/sink pairs of vertices, the goal is to find a low-cost subgraph satisfying some connectivity properties. We consider two-cost network design models, in which edges of the input graph have an associated length along with the fixed cost. One fundamental such problem is the multicommodity buy-at-bulk problem: given demand pairs s_i, t_i with demand d(i), the goal is to route d(i) flow between s_i and t_i while minimizing the sum of the fixed cost (total cost of edges used) and oruting cost (total length of flow paths). Chekuri et al. gave a polylogarithmic approximation via junction trees and posed the question of establishing an upper bound on the integrality gap of a natural LP relaxation.In this talk, we will see that we can obtain a polylogarithmic upper bound on the integrality gap using recent results in hop-constrained oblivious routing. We demonstrate a connection between h-hop-constrained network design problems (in which each demand pair must be connected with a path of at most h edges) and obtain bi-criteria approximation algorithms for these problems as well.

This is joint work with Chandra Chekuri.

5:44-5:51pm **Sharat Ibrahimpur; London School of Economics and Political Science**

Title: Efficient caching with reserves via marking

Abstract: Online caching is among the most fundamental and well-studied problems in the area of online algorithms. Much of traditional caching research studies cache management for a single-user environment. In this work, we introduce and study a model of caching that captures issues that arise in a multi-user environment. In the caching with reserves problem, we have m users sharing a cache that can hold k pages and we are given a reserve requirement k_i for each user i. Page requests arrive online and the goal is to dynamically maintain the cache so that the number of cache misses is minimized while ensuring that the cache holds at least k_i pages belonging to user i at any time.

Unlike the classical setting, the caching with reserves problem is NP-hard even in the offline setting, where the page sequence is known in advance. Here, we give an offline 2-approximation algorithm that is inspired by Belady's farthest-in-the-future rule. In the online setting, we show that a water-filling strategy gives an $O(\log k)$-competitive fractional algorithm, and this can be further turned into a marking-style randomized integral algorithm at a constant factor loss in competitiveness. Our competitive analysis uses a simple potential function to upper bound the cost of our algorithm and complements this with a novel LP dual-fitting based lower bound on the cost of an optimal policy.

This is joint work with Manish Purohit, Zoya Svitkina, Erik Vee, and Joshua Wang.

5:51-5:58pm **Varun Suriyanarayana; Cornell University**

Title: Load balancing with recourse

Abstract: We study the online load balancing problem with unrelated machines from the perspective of recourse. We show that with amortized recourse of O(log n) we can achieve a constant competitive ratio. Along the way we show a constant competitive algorithm for generalized flow with constant amortized recourse.

5:58-6:05pm **CANCELLED: Fatemeh Abbasi; Wroclaw University**

Title: Parameterized approximation for robust clustering in discrete geometric spaces

Abstract: We consider the well-studied (k, z) clustering problem, which generalizes the classic k-median, k-means, and k-center problems and arises in the domains of robust optimization [Anthony, Goyal, Gupta, Nagarajan, Math. Oper. Res. 2010] and in algorithmic fairness [Abbasi, Bhaskara, Venkatasubramanian, 2021 \& Ghadiri, Samadi, Vempala, 2022]. Given a constant $z \geq 1$, the input to (k, z) clustering is a set P of n points in a metric space (M, δ), a weight function $w: P \rightarrow \mathbb{R}_{\geq 0}$ and a positive integer k. Further, each point belongs to one (or more) of the m many different groups $S_1, S_2, \ldots, S_m \subseteq P$. Our goal is to find a set X of k centers such that $\max_{i \in [m]} \sum_p w_p \delta(p, X)^z$ is minimized. Complementing recent work on this problem, we give a comprehensive understanding of the parameterized approximability of the problem in geometric spaces where the parameter is the number k of centers. We prove the following results:

(i) For a universal constant $\eta_0 > 0.0006$, we devise a $3^z(1-\eta_0)$-factor FPT approximation algorithm for (k,z) clustering in *discrete* high-dimensional Euclidean spaces where the set of potential centers is finite. This shows that the lower bound of 3^z for general metrics [Goyal, Jaiswal, Inf. Proc. Letters, 2023] no longer holds when the metric has geometric structure.

(ii) We show that (k,z) clustering in discrete Euclidean spaces is $\left(\sqrt{3/2} - o(1)\right)$-hard to approximate for FPT algorithms, even if we consider the special case k-center in logarithmic dimensions. This rules out a $(1+\epsilon)$-approximation algorithm running in time $f(k,\epsilon)\mathrm{poly}(m,n)$ (also called efficient parameterized approximation scheme or EPAS), giving a striking contrast with the recent EPAS for the *continuous* setting where centers can be placed anywhere in the space [Abbasi et al., FOCS'23].

(iii) However, we obtain an EPAS for (k,z) clustering in discrete Euclidean spaces when the dimension is sublogarithmic (for the discrete problem, earlier work [Abbasi et al., FOCS'23] provides an EPAS only in dimension $o(\log\log n)$). Our EPAS works also for metrics of sub-logarithmic doubling dimension.