Linear and Non-Linear Mixed Integer Optimization

Poster Session Abstracts
Tuesday, January 31, 2023

The Maximum Entropy Sampling Problem: Hardness, Solvable Cases and Generating Input
Hessa Al-Thani, University of Michigan

The Maximum-Entropy Sampling problem (MESP) is an NP-Hard problem that aims to find the sample $S \subset N$ that will maximize the log-determinant of a given positive semidefinite matrix C. This problem has an application in the field of environmental design where N is a set of n variates representing information from a sensor network. The goal of the MESP is then to find the most informative subset. We will show that for certain sparse cases of the input matrix there exists a polynomial time algorithm (tridiagonal and spider with a bounded number of legs), while the problem is NP-hard for general arrowhead matrices, we show that a greedy algorithm is optimal for a subclass. We also briefly present open-source R package (MESgenCov v 0.1.0) for fitting a time-series model to multivariate rainwater chemistry data and generating a covariance matrix for use as input for the MESP. This package allows us to make use of readily available concentration data from the NADP/NTN network and use the MESP to find the most informative subset of sensors in the network.

Extending the Primal-Dual Method for Edge Connectivity
Ishan Bansal, School of ORIE, Cornell

Williamson, Goemans, Vazirani and Mihail, in 1995, proposed the primal-dual method for edge augmentation problems where the connectivity requirements are given by the so called uncrossable functions. We address the long-standing open question of extending primal-dual methods to more general classes of functions. We see applications of our results in the recently studied flexible graph connectivity problem where we provide O(1) approximation algorithms. Applications are also seen in the capacitated k-edge connected problem.

Generalized Scaling for the Constrained Maximum-Entropy Sampling Problem
Zhongzhu Chen, University of Michigan

The best techniques for the constrained maximum-entropy sampling problem, a discrete-optimization problem arising in the design of experiments, are via a variety of concave continuous relaxations of the objective function. A standard bound-enhancement technique in this context is scaling. We extend this technique to generalized scaling, we give mathematical results aimed at supporting algorithmic methods for computing optimal generalized scalings, and we give computational results demonstrating the usefulness of generalized scaling on benchmark problem instances.

On the Combinatorial Diameters of Parallel and Series Connections
Weston Grewe, University of Colorado Denver

Combinatorial diameters of polyhedra are a classical topic in linear programming due to its connection to the possibility of an efficient pivot rule for the simplex method. We are interested in the diameters of polyhedra formed from the so-called parallel or series connection of oriented matroids: oriented

matroids are the natural way to connect representable matroid theory with linear programming, and these connections are fundamental operations for the construction of more-complicated matroids from elementary matroid blocks.
We prove that, for polyhedra whose combinatorial diameter satisfies the Hirsch-conjecture bound regardless of the right-hand sides in a standard-form description, the diameters of their parallel or series connections remain small in the Hirsch-conjecture bound. These results are a substantial step toward devising a diameter bound for all polyhedra defined through totally-unimodular matrices, based on Seymour's famous decomposition theorem.

Scarf's Algorithm and Stable Matchings
Chengyue He, Columbia Universiy

Scarf's algorithm gives a pivoting procedure to find a special vertex—a dominating vertex—in down-monotone polytopes. This paper studies the behavior of Scarf's algorithm when employed to find stable matchings. First, it proves that Scarf's algorithm can be
implemented to run in polynomial time, showing the first positive result on its runtime in significant settings. Second, it shows an infinite family of instances where, no matter the pivoting rule and runtime, Scarf's algorithm outputs a matching from an exponentially small subset of all stable matchings, thus showing a structural weakness of the approach.

Improved Approximation Algorithm for Node-Tree Augmentation
Dylan Hyatt-Denesik, Eindhoven University of Technology

Given a graph $G = (V,E)$ with a (edge- or node-)connectivity of $k\in \mathbb{Z}^{+}$ and a set of links $L\subseteq \binom{V}{2}\backslash E$, the problem of (edge- or node-)connectivity augmentation is then to select the cheapest set of links $L'\subseteq L$ that increases the (edge- or node-)connectivity of G from k to $k+1$. This problem is NP-hard for both the node and edge connectivity instances when $k\geq 1$, with edge connectivity being the more well studied of the two. In the edge-connectivity setting, it is well known that the problem can be reduced to $k=2$.

In 2020, the longstanding approximation barrier of 2 for the case where $k=2$, often called the Cactus Augmentation Problem, was broken by \cite{10.1145/3357713.3384301}, giving a 1.91-approximation. This approximation ratio is found by applying an approximation preserving reduction to (unweighted) node Steiner Trees, and treating this problem as an (unweighted) edge Steiner Tree instance. They then apply an iterative randomized rounding approach for (unweighted) edge Steiner Trees by \cite{10.1145/2432622.2432628}, tailoring the analysis to the particular instances that come about as a result of the reduction to achieve the 1.91-approximation. For the node-connectivity augmentation problem for $k=1$, \cite{nutov20202nodeconnectivity} observed that this problem can be reduced in a similar way to node (unweighted) Steiner Tree instances. They then apply the analysis of \cite{10.1145/3357713.3384301} to this instance, achieving a 1.91-approximation.

A more recent result due to \cite{DBLP:journals/corr/abs-2108-02041} finds an improved approximation ratio for both of these problems by showing that one can modify the iterative randomized rounding technique from \cite{10.1145/2432622.2432628} to apply to the particular (unweighted) node Steiner Tree instances that come from this reduction. To find a better than 1.9-approximation, the core of their analysis then relies on analyzing so-called ``Witness Trees"", achieving a 1.892-approximation. This is the current best approximation for node-connectivity augmentation for $k=1$.

In this work, we show that one can improve the analysis by \cite{DBLP:journals/corr/abs-2108-02041} of these Witness Trees to achieve a 1.8596-approximation, improving the best approximation for node-connectivity augmentation when $k=1$. Furthermore, we provide an example demonstrating that the optimal choice of witness trees in the example yields a bound of $1.84\bar{6}$. Thus, if one wishes to find an approximation factor beating $1.84\bar{6}$ they will need to rely on a completely different analysis.

Stackelberg Vertex Cover

Lennart Kurt Kauther, RWTH Aachen University - Chair of Management Science

In a Stackelberg Network Pricing game, one Player called the Leader controls the prices/weights of some entities (vertices or edges) of the underlying graph. The remaining weights are constant and known to the Leader. The other players, called followers, each have a subgraph on which they independently solve a (combinatorial) optimization problem given the prices defined by the Leader.
The problem then asks for a pricing scheme that maximizes the Leader's revenue which corresponds to the prices of the entities contained in any of the followers' solutions.
I consider the problem Stackelberg Vertex Cover (StackVC) with a single follower, i.e., the Leader controls the prices of some vertices, and the follower aims to find a vertex cover of the graph with minimum weight. Previous research has shown that StackVC is NP-hard on general bipartite graphs. To find the complexity boundary of the problem, I investigated paths – arguably the simplest types of a bipartite graph. I propose a dynamic program to determine optimal prices from the Leader's perspective in linear time. With minor adaptions, this algorithm also extends to trees and forests.

While the exact complexity boundary remains unknown, this remains one of the few positive results in the Stackelberg setting.
My results emphasize the intricacy of the Stackelberg framework by showing that solving problems like StackVC is quite involved even when restricted to acyclic graphs or even paths.

Polyhedral geometry for a structural description of neural networks: linear threshold activations and beyond

Sammy Khalife, Johns Hopkins University

The study of the structure and complexity of certain neural networks, in particular the estimation of their depth and width needed to compute exactly some class of functions, can be analyzed through the lens of polyhedral geometry. Using this approach, one can derive complexity bounds for neural networks with linear threshold activation functions. The class of functions that are representable by such neural networks can be fully characterized, and two hidden layers are necessary and sufficient to represent any function in the class. These insights allow to design an algorithm to solve the empirical risk minimization (ERM) problem to global optimality for these neural networks with a fixed architecture. Furthermore, the ERM problem can also be solved to global optimality with a similar algorithm for a strict superclass of the rectified linear units (ReLU) neural networks.

Polyhedral and Approximation Aspects of Feedback Vertex Set and Pseudoforest Deletion

Shubhang Kulkarni, University of Illinois at Urbana Champaign

A feedback vertex set is a subset of vertices whose deletion makes the graph acyclic. FVS asks for a minimum cost feedback vertex set. A pseudoforest deletion set is a subset of vertices whose deletion ensures that each connected component contains at most one cycle. PFDS asks for a minimum cost pseudoforest deletion set. Both these problems are NP-hard. Approximation algorithms for these problems have relied on the local ratio technique which have subsequently led to primal-dual algorithms via LPs, but these LPs were not known to be solvable efficiently. In this poster, I will present polynomial time solvable LPs that achieve nearly the best possible approximation factor for both FVS and PFDS.

Based on joint work with Karthekeyan Chandrasekaran and Chandra Chekuri.

Transit Network Design and Transportation Equity

Juan Carlos Martinez Mori, Cornell University

Public infrastructure, and in particular public transit, can be enjoyed by multiple users simultaneously. However, the value any particular user gets from public transit depends on its design: some might experience better service than others. In this poster, I motivate questions I am interested in around network design for public transit, particularly from the perspective of transportation equity. I pose the problem as a budgeted public resource allocation problem with self-interested agents, and highlight trade-offs between fairness criteria, stability, and tractability. The goal of this research direction is to understand the interplay between fairness and stability with the goal of advancing transportation justice: ensuring sufficient transportation access for those who rely on public transit the most.

Diameter Constrained Minimum Spanning Graphs

Jared Miller, Northeastern University

The diameter constrained minimum spanning graph problem is an instance of bicriteria network design, in which each edge in a graph has an associated cost and distance value. A subset of the edges must be selected with minimal total cost while keeping the diameter (measured in distance) below a threshold. This problem is known to be NP-hard and nonapproximable. Branch-and-bound formulations have previously been developed for the case where the spanning graph must additionally be a tree and all edges have a uniform distance. In this work we consider the case where edge distances may be non-uniform, and we lift the tree requirement to allow for minimal cost spanning graphs. We develop a branch-and-bound MILP to find a global optimum spanning graph, perform extensions to Steiner graphs, and extend to cases where the network must additionally satisfy commodity flow constraints.

Exact and heuristic solution approaches for the D-optimality problem

Gabriel Ponte, Federal University of Rio de Janeiro

The D-Optimality problem is a fundamental and challenging integer nonlinear optimization problem coming from statistics. Among a finite set of design points (i.e., specifications of experiments), and a budget for the number of experiments to carry out, it seeks to choose the optimal set of design points, so as to minimize the generalized variance of least-squares parameter estimates for a linear model based on the chosen design points. We investigate local-search heuristics for this problem as well as a branch-and-bound algorithm based on its continuous relaxation. Finally, convex optimization duality theory is applied in the development of valid inequalities for the problem that can reduce the range of possible integer values for the variables.

Memory-efficient algorithms for Max-k-Cut and correlation clustering

Nimita Shinde, ICERM

Max-k-Cut and correlation clustering are fundamental graph partitioning problems. For a graph G=(V,E) with n vertices, the methods with the best approximation guarantees for Max-k-Cut and the Max-Agree variant of correlation clustering involve solving SDPs with O(n^2) constraints and variables. Large-scale instances of SDPs, thus, present a memory bottleneck. In this work, we develop simple polynomial-time Gaussian sampling-based algorithms for these two problems that use O(n+|E|) memory and nearly achieve the best existing approximation guarantees. For dense graphs arriving in a stream, we eliminate the dependence on |E| in the storage complexity at the cost of a slightly worse approximation ratio by combining our approach with sparsification.

Stabilization of Capacitated Matching Games

Lucy Verberk, Eindhoven University of Technology

An edge-weighted, vertex-capacitated graph G is called stable if the value of a maximum-weight capacity-matching equals the value of a maximum-weight fractional capacity-matching. Stable graphs play a key role in characterizing the existence of stable solutions for popular combinatorial games that involve the structure of matchings in graphs, such as network bargaining games (NBG) and cooperative matching games (CMG). In particular, in unit-capacity graphs there is an equivalence between the stability of a graph, existence of a stable solution for NBG, and existence of a stable solution for CMG. We show that this equivalence does not completely extend to the capacitated case: we provide a graph that attains a stable solution for the CMG but that is not stable.

The vertex-stabilizer problem asks to compute a minimum number of players to block (i.e., vertices of G to remove) in order to ensure stability for such games (i.e., stability of the graph). The problem has been shown to be solvable in polynomial-time, for unit-capacity graphs. In contrast, we show that it becomes NP-hard and even hard to approximate, for graphs with arbitrary capacities.

Computation and Certification of Ramsey-Type Numbers

William Wesley, University of California, Davis

Given a linear equation $\E$ and positive integer k, the k-color \textbf{Rado number} $R_k(\E)$ is the smallest number n such that every k-coloring of $\{1,2,\dots,n\}$ contains a monochromatic solution to $\E$. For positive integers r and s, the classical \textbf{Ramsey number} is the smallest n such that every edge 2-coloring of the edges of K_n contains either a clique of size r in the first color or a clique of size s in the second color. Both of these numbers are extremely difficult to compute in general. We study how to compute these numbers and certify their values using \emph{Boolean formulas} and \emph{algebraic geometry}.