Algorithms and Complexity aspects of Optimization
Aspects algorithmiques et complexes de l'optimisation
(Org: **Vijay Bhattiprolu** and/et **Joseph Cheriyan** (Waterloo))

NATHAN BENDETTO-PROENÇA, University of Waterloo
An approximation algorithm for the weighted fractional cut-covering problem

A cut in a graph $G = (V, E)$ is a set of edges which has precisely one endpoint in S, for a given subset S of V. The fractional cut-covering number is the optimal value of a linear programming relaxation for the problem of covering each edge by a set of cuts. We define a semidefinite programming relaxation of fractional cut covering whose approximate optimal solutions may be rounded into a fractional cut cover via a randomized algorithm. These results arise from the tight connection between fractional cut covering and the maximum cut problem. By pinpointing the fundamental aspects of this relationship via antiblocker and gauge duality, we are not only able to obtain dual results to the celebrated work of Goemans and Williamson, but also to tie both approximation constants, to obtain new optimality certificates, and to relate both problems to geometric representation of graphs.

This is joint work with Marcel K. de Carli Silva, Cristiane Maria Sato, and Levent Tunçel.

MEHRDAD GHADIRI, Georgia Tech
Bit Complexity of Efficient Optimization

Optimization is a fundamental area in applied math, with applications ranging from economics to machine learning. There has been abundant interest in recent years in improving the asymptotic running times of algorithms for fundamental convex optimization problems. However, most of these works assume infinite precision for arithmetic operations. In this talk, we discuss the bit complexity of fundamental convex optimization problems under fixed-point arithmetic, including linear programming and p-norm regression. This requires analyzing the stability of the underlying inverse maintenance processes and how the errors propagate through the iterations.

HAMED HATAMI, McGill University
A Borsuk-Ulam lower bound for sign-rank

I will present a new topological argument based on the Borsuk-Ulam theorem to prove a lower bound on sign-rank. I will mention applications to communication complexity, the theory of dimension reductions, and learning theory. This talk is based on joint work with Kaave Hosseini and Xiang Meng.

EUIWOONG LEE, University of Michigan
Correlation Clustering with Sherali-Adams

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. Correlation Clustering has been used to model a large number of clustering problems in practice, making it one of the most widely studied clustering formulations. The approximability of Correlation Clustering has been actively investigated, culminating in a 2.06-approximation algorithm in 2015 based on rounding the standard linear program (LP) relaxation. Since it has been known that the standard LP cannot give better than a 2-approximation, it has remained an open question to determine if the approximation factor of 2 can be reached, or even breached. In this work, we answer this question affirmatively by showing that there exists a $(1.994 + \epsilon)$-approximation algorithm based on a strengthened LP relaxation called the Sherali-Adams hierarchy.

SHI LI, University at Buffalo
Online Unrelated-Machine Load Balancing and Generalized Flow with Recourse

I this talk, we discuss the online unrelated-machine load balancing problem with recourse, where the algorithm is allowed to re-assign prior jobs. We give a $(2 + \epsilon)$-competitive algorithm for the problem with $O_\epsilon(\log n)$-amortized recourse per job. This is the first $O(1)$-competitive algorithm for the problem with reasonable recourse, and the competitive ratio nearly matches the long-standing best-known offline approximation guarantee.

We also present an $O(\log \log n / \log \log \log n)$-competitive algorithm for the problem with $O(1)$-amortized recourse, improving upon the previous best $O(\log \log n)$-competitive ratio of Gupta et al., which works only for the special case of the restricted assignment model. The algorithm is based on our recent $O(\log \log m / \log \log \log m)$-online rounding algorithm for the problem without recourse.

Both algorithms are based on producing a fractional solution online first. To do so, we introduce and study the online algorithm for the generalized network flow problem (also known as network flow problem with gains) with recourse. We give an online algorithm for the problem with amortized recourse of $O(1/\epsilon)$ and capacity-violation of 1+epsilon. The $(1 + \epsilon)$-factor improves upon the corresponding $(2 + \epsilon)$-factor of Gupta et al., which only works for the ordinary network flow problem.

The talk is based on two papers: [Krishnaswany-Li-Suriyanarayana,2022] and [Li-Xian,2021].

ALEKSANDR NIKOLOV, University of Toronto
Computing and Using Factorization Norms

A factorization norm of a matrix A is the value of an optimization problem over factorizations $A = BC$ with the objective to minimize the product of some matrix norms of B and C. Usually in this context we think of A as an operator between two non-Euclidean spaces, of B and C as factoring A through Euclidean space, and of the factorization norm as measuring how much the factoring distorts A. Factorization norms and the optimal factoring then allow solving non-Euclidean problems by mapping them to more easily solvable problems in Euclidean space. Originating in functional analysis, factorization norms have found many applications, e.g., to communication complexity, discrepancy theory, and private data analysis. In this talk I will discuss some of these applications, and also how to use tools from optimization to compute some factorization norms.

RAFAEL OLIVEIRA, University of Waterloo
Optimization, Invariant Theory, Computer Science and Math

What do the following problems, from seemingly unrelated areas of mathematics, quantum information theory and computer science, have in common?

- perfect matching in bipartite graphs

- word problem for the free skew field

- optimal constant in Brascamp-Lieb inequalities

- one-body quantum marginal problem

- the Paulsen problem

- Horn's problem

- sample complexity of matrix and tensor normal model

As it turns out, these problems are all instances of the moment polytope problem from geometric invariant theory. Moreover, these problems can be cast as (geodesically) convex optimization problems over geodesically convex Riemannian manifolds. In this talk we will discuss these connections and how a recent series of works was able to give efficient algorithms for the problems above (via the unifying view of moment polytopes), as well as mention several open questions.

KOSTYA PASHKOVICH, University of Waterloo
Non-Adaptive Matroid Prophet Inequalities For Minor-Closed Matroid Classes

We consider the problem of matroid prophet inequalities. Kleinberg and Weinberg have constructed a 2-competitive mechanism for all matroids. However, their mechanism recomputes thresholds during its course. In other words, their mechanism is adaptive.

The non-adaptive case is far from resolved. There are known constant-competitive nonadaptive mechanisms for uniform and graphical matroids, but some classes of gammoids do not admit constant-competitive nonadaptive mechanisms.

In this work, we present constant-competitive nonadaptive mechanism for all regular matroids and for further minor-closed families of matroids.

Joint work with Alice Sayutina.

ROBERT ROBERE, McGill University
On the Proof Complexity of Integer Programming Solvers

We discuss recent progress on understanding the complexity of modern integer programming solvers in optimization using tools from propositional proof complexity. In particular, we focus on the recent introduction and study of the so-called "Stabbing Planes" proof system (also known as "Branching Proofs") which very tightly model the execution of such solvers. Both lower bounds and upper bounds on various complexity measures of Stabbing Planes proofs will be discussed, as well as the close relationship between Stabbing Planes proofs and Cutting Planes proofs.

ALEX TUNG, University of Waterloo
Cheeger Inequalities for Vertex Expansion and Reweighted Eigenvalues

The classical Cheeger's inequality relates the edge conductance ϕ of a graph and the second smallest eigenvalue λ_2 of the Laplacian matrix. Recently, Olesker-Taylor and Zanetti discovered a Cheeger-type inequality $\psi^2/\log|V| \lesssim \lambda_2^* \lesssim \psi$ connecting the vertex expansion ψ of a graph $G = (V, E)$ and the maximum reweighted second smallest eigenvalue λ_2^* of the Laplacian matrix.

In this work, we first improve their result to $\psi^2/\log d \lesssim \lambda_2^* \lesssim \psi$ where d is the maximum degree in G, which is optimal up to constant factor. Also, the improved result holds for weighted vertex expansion, answering an open question by Olesker-Taylor and Zanetti. Building on this connection, we then develop a new spectral theory for vertex expansion. We discover that several interesting generalizations of Cheeger inequalities relating edge conductances and eigenvalues have a close analog in relating vertex expansions and reweighted eigenvalues. These include analogs of bipartite Cheeger's inequality, higher order Cheeger's inequality and improved Cheeger's inequality.

Finally, inspired by this connection, we present negative evidence to the $0/1$-polytope edge expansion conjecture by Mihail and Vazirani. We construct $0/1$-polytopes whose graphs have very poor vertex expansion. This implies that the fastest mixing time to the uniform distribution on the vertices of these $0/1$-polytopes is almost linear in the graph size. This does not disprove the conjecture, but this is in contrast with known positive results which proved poly-logarithmic mixing time to the uniform distribution on the vertices of subclasses of $0/1$-polytopes.

(Paper link: https://arxiv.org/abs/2203.06168.)

YIBIN ZHAO, University of Toronto
A Simple and Efficient Parallel Laplacian Solver

A symmetric matrix is called a Laplacian if it has nonpositive off-diagonal entries and zero row sums. Laplacian linear systems naturally appears in various optimization problems, e.g. [Daitch and Spielman, 2008, Madry, 2013, Lee and Sidford, 2014, Dong, Gao, Goranci, Lee, Peng, Sachdeva, and Ye, 2021]. Since the seminal work of Spielman and Teng [2004] on solving Laplacian linear systems in nearly linear time, several algorithms have been designed for the task. Yet, the work of Kyng and Sachdeva [2016] remains the simplest and most practical sequential solver. They presented a solver purely based on random sampling and without graph-theoretic constructions such as low-stretch trees and sparsifiers.

In this work, we extend the result of Kyng and Sachdeva [2016] to a simple parallel Laplacian solver with $O(m \log^3 n \log \log n)$

work and $O(\log^2 n \log \log n)$ depth using the ideas of block Cholesky factorization from Kyng, Lee, Peng, Sachdeva, and Spielman [2016]. Our proof is simple and based on standard matrix concentration inequalities. Compared to the best known parallel Laplacian solvers that achieve polylogarithmic depth due to Lee, Peng, and Spielman [2015], our solver achieves both better depth and, for sparse graphs, better work.

This talk is based on a joint work with Sushant Sachdeva.