Synergies of combinatorics and theoretical computer science

Workshop at the EPFL Bernoulli Center, August 19–23, 2024

Monday:

9:00–9:15	welcome remarks
9:15–10:00	Rico Zenklusen: *Random-Assignment Matroid Secretary Without Knowing the Matroid*
10:05–10:45	coffee break
10:45–11:30	Matthew Kwan: *Resolution of the Quadratic Littlewood-Offord problem*
11:35–12:20	Vera Traub: *The Bidirected Cut Relaxation for Steiner Tree has Integrality Gap Smaller than 2*
12:25–14:30	lunch break
14:30–15:15	Sorrachai Yingchareonthawornchai: *How to Search and Sort using Forbidden 0-1 Matrix Theory*
15:20–16:05	Alexey Gordeev: *Combinatorial Nullstellensatz and the Erdős box problem*
16:10–16:45	coffee break
16:45–17:30	Sophie Huiberts: *Short Stories about Linear Programming*
17:35–	welcome aperitif

Tuesday:

9:15–10:00	Nati Linial: *The Rank-Ramsey Problem and the Log-Rank Conjecture*
10:05–10:45	coffee break
10:45–11:30	Lianna Hambardzumyan: *On rank-like parameters of Boolean matrices*
11:35–12:20	Rob Morris: *Geometric conjectures and Ramsey numbers*
12:25–14:30	lunch break
14:30–15:15	Pravesh K Kothari: *Spectral Refutation via Kikuchi Matrices and Applications*
15:20–16:05	Mehtaab Sawhney: *Improved Bounds for Szemerédi's Theorem*
16:10–16:45	coffee break
16:45–17:30	Venkatesan Guruswami: *Combinatorial challenges in coding theory: A sampler*

Wednesday:

9:15–10:00	Benny Sudakov: *SDP, MaxCut, discrepancy and log-rank-conjecture*
10:05–10:45	coffee break
10:45–11:30	Hannaneh Akrami: *Epistemic EFX Allocations Exist for Monotone Valuations*
11:35–12:20	Zixuan Xu: *Essential covers of the hypercube requires many hyperplanes*
12:25–14:30	lunch break
14:30–15:15	Sammy Luo: *A New Polynomial Method in Additive Combinatorics*
15:20–15:40	coffee break
15:40–16:25	Maya Sankar: *On the Generalized Ramsey–Turan Density of Cliques*
18:30–	workshop dinner in Grandvaux (with option to go for a walk in the vineyards beforehand)

Thursday:

9:15–10:00	Daniel Král': *Matroid depth and width parameters*
10:05–10:45	coffee break
10:45–11:30	Nathan Klein: *Ghost Value Augmentation for k-Edge-Connectivity*
11:35–12:20	Jinyoung Park: *Dedekind's problem and beyond*
12:25–14:30	lunch break
14:30–15:15	Matija Bucic: *Robust sublinear expanders*
15:20–16:05	Omar Alrabiah: *Near-Tight Bounds for 3-Query Locally Correctable Binary Linear Codes via Rainbow Cycles*
16:10–16:45	coffee break
16:45–17:30	Oliver Janzer: *Edge-disjoint cycles with the same vertex set*

Friday:

9:15–10:00	Parinya Chalermsook: *Approximation Schemes for Clustering through Scatter Dimension*
10:05–10:45	coffee break
10:45–11:30	Peter Manohar: *New Spectral Techniques in Algorithms, Combinatorics, and Coding Theory: the Kikuchi Matrix Method*
11:35–12:20	Daniel Dadush: *Column Bounds for the Circuit Imbalance Measure*
12:25–	lunch and departure

Abstracts

Rico Zenklusen: *Random-Assignment Matroid Secretary Without Knowing the Matroid*
The Matroid Secretary Problem (MSP) is a well-known online selection problem about selecting a heavy independent set among elements revealing their weights online in random order. The existence of an O(1)-competitive MSP algorithm is a notorious open problem known as the Matroid Secretary Conjecture. Intense research since MSP's inception led to progress and O(1)-competitive algorithms for various special cases and variants. Unfortunately, these algorithms heavily rely on knowing the matroid upfront, which is arguably a very undesirable property for trying to approach the general MSP conjecture.

I will talk about how one can get an O(1)-competitive algorithm without knowing the matroid for Random-Assignment MSP, where weights are assigned uniformly at random to the elements. This settles an open question raised by both Soto [SIAM Journal on Computing 2013] and Oveis Gharan & Vondrák [Algorithmica 2013], and leads to the first well-known MSP variant with an O(1)-competitive algorithm that does not need to know the matroid upfront. Our approach is based on first approximately learning the rank-density curve of the matroid, which we then exploit algorithmically.

This is joint work with Richard Santiago and Ivan Sergeev.

Matthew Kwan: *Resolution of the Quadratic Littlewood-Offord problem*
Consider a quadratic polynomial $Q(\xi_1, \ldots, \xi_n)$ of a random binary sequence $\xi_1, \ldots, \xi_n$. To what extent can $Q(\xi_1, \ldots, \xi_n)$ concentrate on a single value? This is a quadratic version of the classical Littlewood-Offord problem; it was was popularised by Costello, Tao and Vu in their study of symmetric random matrices, and has since become a rich source of connections between combinatorics, probability and computer science. In this talk we will discuss a new essentially optimal bound for the quadratic Littlewood-Offord problem, as conjectured by Nguyen and Vu. Joint work with Lisa Sauermann.

Vera Traub: *The Bidirected Cut Relaxation for Steiner Tree has Integrality Gap Smaller than 2*
The Steiner tree problem is one of the most prominent network design problems. It asks to connect a given set of terminals in the cheapest possible way. The best-known approximation algorithms for Steiner tree involve enumeration of a (polynomial but) very large number of candidate components and have therefore high running times.

A promising ingredient for the design of fast and accurate approximation algorithms for Steiner tree is the bidirected cut relaxation (BCR). This linear programming relaxation is known to be integral in the spanning tree case [Edmonds'67], that is, when all vertices are terminals. For general instances, however it was not known whether the integrality gap of BCR is better than the integrality gap of the natural undirected relaxation, which is well-known to be exactly 2. We resolve this question by proving that the integrality gap of BCR ist at most 1.9988.

This is joint work with Jarek Byrka and Fabrizio Grandoni.

Sorrachai Yingchareonthawornchai: *How to Search and Sort using Forbidden 0-1 Matrix Theory*
We will present a survey on recent progress in pattern-avoiding access in binary search trees and its application to sorting pattern-avoiding permutations. We will explore the key techniques in the area that allow us to achieve nearly tight amortized analysis through the lens of forbidden 0-1 matrix theory, including the new matrix decomposition method. We will also present new tight extremal bounds on 0-1 matrices forbidding product patterns, which bound the time required to sort a pattern-avoiding permutation.

Alexey Gordeev: *Combinatorial Nullstellensatz and the Erdős box problem*
In the talk, I will show how Lasoń's generalization of Alon's Combinatorial Nullstellensatz can be used to obtain lower bounds on Turán numbers of complete r-partite r-uniform hypergraphs. As an example, I will give a short and simple explicit construction of a hypergraph free of copies of the complete r-partite r-uniform hypegraph with parts of size 2, thereby providing a lower bound for the so-called Erdős box problem. This asymptotically matches best known bounds when $r \leq 4$. The talk is based on my recent work https://doi.org/10.1016/j.disc.2024.114037.

Sophie Huiberts: *Short Stories about Linear Programming*
We will learn how to derive the LP dual and complementary slackness from Newtonian physics. We will also regard the first ever LP to be solved, the diet problem, and learn more about it's social context in the 1940's.

Nati Linial: *The Rank-Ramsey Problem and the Log-Rank Conjecture*
A graph is called Rank-Ramsey if (i) Its clique number is small, and (ii) The adjacency matrix of its complement has small rank. We initiate a systematic study of such graphs. Our main motivation is that their constructions, as well as proofs of their non-existence, are intimately related to the famous log-rank conjecture from the field of communication complexity. These investigations also open interesting new avenues in Ramsey theory. Joint work with Gal Beniamini and Adi Shraibman arXiv:2405.07337

Lianna Hambardzumyan: *On rank-like parameters of Boolean matrices*
In this talk, I will define two new rank-like measures: blocky rank for Boolean matrices, and spiky rank for real matrices. A matrix is blocky if it is a "blowup" of an identity matrix. The blocky rank of a matrix is the minimum number of blocky matrices that linearly span the matrix. Initially defined by Hambardzumyan, Hatami, and Hatami (2022), blocky rank was primarily motivated by its connection to communication complexity. Later, Avraham and Yehudayoff (2024) studied blocky rank as a complexity measure and found connections to circuit complexity and combinatorics. In this talk, I will review these connections and extend blocky rank to real matrices by introducing spiky rank. I will focus on the connection between spiky rank and matrix rigidity.

Rob Morris: *Geometric conjectures and Ramsey numbers*
In the talk I will state two natural conjectures about probability distributions (or sets of points) on high-dimensional spheres, and explain how these conjectures (if true) could be used to improve the Erdos-Szekeres upper bound for the r-colour Ramsey numbers.

Pravesh K Kothari: *Spectral Refutation via Kikuchi Matrices and Applications*
We will present a method to reduce certain extremal combinatorial problems to establishing the unsatisfiability of k-sparse linear equations mod 2 (aka k-XOR formulas) with a limited amount of randomness via spectral bounds on certain "Kikuchi" graphs associated with the k-XOR formulas. Kikuchi graphs are appropriately chosen induced subgraphs of Cayley graphs on the hypercube (or sometimes variants such as products of hypercubes) generated by the coefficient vectors of the given equations and their spectral norm can be related to certain natural combinatorial properties of the input formula.

We will discuss the following applications of this method:

1. Proving Feige's conjectured hypergraph Moore bound – the optimal trade-off between the size of the smallest linear dependent subset of a system of k-sparse linear equations modulo 2 and the total number of equations. This theorem generalizes the famous irregular Moore bound of Alon, Hoory and Linial (2002) for graphs (equivalently, 2-sparse linear equations mod 2),

2. Proving a cubic lower bound on 3-query locally decodable codes (LDCs) improving on a quadratic lower bound of Kerenedis and de Wolf (2004),

3. Proving an exponential lower bound on linear 3-query locally correctable codes (LCCs). This result establishes a sharp separation between 3-query LCCs and 3-query LDCs that are known to admit a construction with a sub-exponential length. It is also the first result to obtain any super-polynomial lower bound for >2-query local codes.

Time permitting, we may also discuss applications to a strengthening of Szemeredi's theorem that asks for establishing the minimal size of a random subset of integers S such that every dense subset of integers contains a 3-term arithmetic progression with common difference from S.

Mehtaab Sawhney: *Improved Bounds for Szemerédi's Theorem*
We discuss recent improved bounds for Szemerédi's Theorem. The talk will seek to provide a gentle introduction to higher order Fourier analysis and recent quantitative developments. In particular, the talk will provide a high level sketch for how the inverse theorem for the Gowers norm enters the picture and the starting points for the proof of the inverse theorem. Additionally, the talk (time permitting) will discuss how recent work of Leng on equidistribution of nilsequences enters the picture and is used. No background regarding nilsequences will be assumed. Based on joint work with James Leng and Ashwin Sah.

Venkatesan Guruswami: *Combinatorial challenges in coding theory: A sampler*
I'll highlight some combinatorial open problems arising in or inspired by coding theory.

Benny Sudakov: *SDP, MaxCut, discrepancy and log-rank-conjecture*
Semidefinite programming (SDP) is a powerful method used in many important approximation algorithms. In this talk, I discuss a different aspect of SDP and demonstrate how it can be employed to offer concise proofs for several well-known and new estimates related to MaxCut, as well as the discrepancy of graphs and matrices. I also explain how the discrepancy result leads to an improvement in Lovett's best-known upper bound on the log-rank conjecture.

Hannaneh Akrami: *Epistemic EFX Allocations Exist for Monotone Valuations*
We study the fundamental problem of fairly dividing a set of indivisible items among agents with (general) monotone valuations. The notion of envy-freeness up to any item (EFX) is considered to be the most fascinating fairness concept in this line of work. Unfortunately, despite significant efforts, existence of EFX allocations is a major open problem in fair division, thereby making the study of approximations and relaxations of EFX a natural line of research. Recently, Caragiannis et al. introduced a promising relaxation of EFX, called 'epistemic EFX' (EEFX). We say an allocation to be EEFX if, for every agent, it is possible to shuffle the items in the remaining bundles so that she becomes 'EFX-satisfied'. Caragiannis et al. proved existence and polynomial-time computability of EEFX allocations for additive valuations. A natural question asks what happens when we consider valuations more general than additive?

We address this important open question and answer it affirmatively by establishing the existence of EEFX allocations for an arbitrary number of agents with general monotone valuations. To the best of our knowledge, EEFX is the only known relaxation of EFX (beside EF1) to have such strong existential guarantees. Furthermore, we complement our existential result by proving computational and information-theoretic lower bounds. We prove that for an arbitrary number of (more than one) agents (even) with identical submodular valuations, it is PLS-hard to compute EEFX allocations and it requires exponentially-many value queries to do so.

Joint work with Nidhi Rathi (https://arxiv.org/abs/2405.14463)

Zixuan Xu: *Essential covers of the hypercube requires many hyperplanes*
We prove a new lower bound for the smallest possible size of an essential cover of the n-dimensional hypercube, i.e. the smallest possible size of a collection of hyperplanes that forms a minimal cover and such that furthermore every variable appears with a non-zero coefficient in at least one of the hyperplane equations. We show that such an essential cover must contain at least $\Omega(n^{2/3}/(\log n)^{2/3})$ hyperplanes, improving previous lower bounds of Linial–Radhakrishnan, of Yehuda–Yehudayoff and of Araujo–Balogh–Mattos. This new lower bound also implies new lower bounds for some problems in proof complexity. Based on joint work with Lisa Sauermann.

Sammy Luo: *A New Polynomial Method in Additive Combinatorics*
A central tool in additive combinatorics is the "polynomial method", a family of powerful techniques for studying the existence and size of sets satisfying given properties by encoding them in terms of the zeros of certain polynomials, which can then be analyzed from an algebraic perspective.

In this talk, we introduce a new form of the polynomial method based on what we call "shift operators". We show how to take advantage of the many useful properties of these operators to give new proofs of the core results of both relatively classical and modern versions of the polynomial method—including the Combinatorial Nullstellensatz, the Croot-Lev-Pach method, and Dvir's method of multiplicities—thereby suggesting a potential way of unifying these tools. We also touch on some possible new directions in which the method may be fruitfully applied.

Maya Sankar: *On the Generalized Ramsey–Turan Density of Cliques*
We study the generalized Ramsey–Turán function $\mathrm{RT}(n, K_s, K_t, o(n))$, which is the maximum possible number of copies of K_s in an n-vertex K_t-free graph with independence number sublinear in n. The case $s = 2$ was settled by Erdős, Sós, Bollobás, Hajnal, and Szemerédi in the 1980s. We combinatorially resolve the general case for all $s \geq 3$, showing that the (asymptotic) extremal graphs for this problem have simple (bounded) structures. In particular, it implies that the extremal structures follow a periodic pattern when t is much larger than s. Our results disprove a conjecture of Balogh, Liu, and Sharifzadeh and show that a relaxed version does hold.

Daniel Kráľ: *Matroid depth and width parameters*
Depth and width parameters of graphs, e.g., tree-width, path-width and tree-depth, play a crucial role in algorithmic and structural graph theory. These notions are of fundamental importance in the theory of graph minors, fixed parameter complexity and the theory of combinatorial sparsity. In this talk, we will survey structural and algorithmic results concerning width and depth parameters of matroids. We will view matroids as purely combinatorial objects and discuss their structural properties related to depth and width decompositions. As an algorithmic application, we will present matroid based algorithms that can uncover a hidden Dantzig-Wolfe-like structure of instances of integer programs (if such structure is present).

Nathan Klein: *Ghost Value Augmentation for k-Edge-Connectivity*
We show that every fractionally k-edge-connected weighted graph (i.e. every solution to the canonical k-edge-connectivity linear program) can be rounded to an integral $(k-10)$-edge-connected graph of no greater cost.

This implies that for large constant values of k, fractional k-edge-connectivity and integral k-edge-connectivity are essentially the same. As a byproduct of this result, we show that one can produce a $(k-10)$-edge-connected spanning subgraph (ECSS) of cost no more than the optimal k-ECSS, complementing the existing 2-approximation. In addition, this result implies a $1+O(1/k)$ approximation for the k-edge-connected multi-subgraph problem (k-ECSM), resolving a conjecture of Pritchard from 2011.

Jinyoung Park: *Dedekind's problem and beyond*
The Dedekind's Problem asks the number of monotone Boolean functions, $a(n)$, on n variables. Equivalently, $a(n)$ is the number of antichains in the n-dimensional Boolean lattice $[2]^n$ or the number of elements of the free distributive lattice on n generator. While the exact formula for the Dedekind number $a(n)$ is still unknown, its asymptotic formula has been well-studied. Since any subsets of a middle layer of the Boolean lattice is an antichain, the logarithm of $a(n)$ is trivially bounded below by the size of the middle layer. In the 1960's, Kleitman proved that this trivial lower bound is optimal in the logarithmic scale, and the actual asymptotics was also proved by Korshunov and Sapozhenko in 1980's. In this talk, we will discuss recent developments on some variants of Dedekind's Problem. Based on joint works with Matthew Jenssen, Alex Malekshahian, Michail Sarantis, and Prasad Tetali.

Matija Bucic: *Robust sublinear expanders*
Expander graphs are perhaps one of the most widely useful classes of graphs ever considered. In this talk, we will focus on a fairly weak "sublinear" notion of expander graphs, first introduced by Komlós and Szemerédi around 30 years ago. They have found many remarkable applications ever since and have been developed essentially independently by both the combinatorial and theoretical computer science community. In this talk, we will focus on certain robustness conditions one may impose on sublinear expanders and a number of exciting applications of this idea to a variety of areas including combinatorics, discrete geometry, additive number theory, and coding theory.

Omar Alrabiah: *Near-Tight Bounds for 3-Query Locally Correctable Binary Linear Codes via Rainbow Cycles*
We prove that a binary linear code of block length n that is locally correctable with 3 queries against a fraction $\delta > 0$ of adversarial errors must have dimension at most $O_\delta(\log^2 n * \log\log n)$. This is almost tight in view of quadratic Reed-Muller codes being a 3-query locally correctable code (LCC) with dimension $\Theta(\log^2 n)$. Our result improves, for the binary field case, the $O_\delta(\log^8 n)$ bound obtained in the recent breakthrough of [Kothari and Manohar, 2023] (and the more recent improvement to $O_\delta(\log^4 n)$ announced in [Yankovitz, 2024]).

Previous bounds for 3-query linear LCCs proceed by constructing a 2-query locally decodable code (LDC) from the 3-query linear LCC/LDC and applying the strong bounds known for the former. Our approach is more direct and proceeds by bounding the covering radius of the dual code, borrowing inspiration from [Iceland and Samorodnitsky, 2018]. That is, we show that if $x \to (v_1 * x, v_2 * x, ..., v_n * x)$ is an arbitrary encoding map $F_2^k \to F_2^n$ for the 3-query LCC, then all vectors in F_2^k can be written as a $\tilde{O}_{delta}(\log n)$-sparse linear combination of the v_i's, which immediately implies $k \leq \tilde{O}_\delta((\log n)^2)$. The proof of this fact proceeds by iteratively reducing the size of any arbitrary linear combination of at least $\tilde{\Omega}_\delta(\log n)$ of the v_i's. We achieve this using the recent breakthrough result of [Alon, Bucić, Sauermann, Zakharov, and Zamir, 2023] on the existence of rainbow cycles in properly edge-colored graphs, applied to graphs capturing the linear dependencies underlying the local correction property.

Oliver Janzer: *Edge-disjoint cycles with the same vertex set*
In 1975, Erdős asked for the maximum number of edges that an n-vertex graph can have if it does not contain two edge-disjoint cycles on the same vertex set. It is known that Turan-type results can be used to prove an upper bound of $n^{3/2+o(1)}$. However, this approach cannot give an upper bound better than $n^{3/2}$.

We show that, for any fixed k, every n-vertex graph with at least $n\,\text{polylog}(n)$ edges contains k pairwise edge-disjoint cycles with the same vertex set, resolving this old problem in a strong form up to a polylogarithmic factor. The well-known construction of Pyber, Rodl and Szemeredi of graphs without 4-regular subgraphs shows that there are n-vertex graphs with $\Omega(n \log\log n)$ edges which do not contain two cycles with the same vertex set, so the polylogarithmic term in our result cannot be completely removed.

Our proof combines a variety of techniques including sublinear expanders, absorption and a novel tool for regularisation, which is of independent interest. Among other applications, this tool can be used to regularise an expander while still preserving certain key expansion properties.

Joint work with Debsoumya Chakraborti, Abhishek Methuku and Richard Montgomery.

Parinya Chalermsook: *Approximation Schemes for Clustering through Scatter Dimension*
We give a clean and simple approximation scheme for center-based clustering problems. Our algorithm (despite being oblivious to the input structures and objective functions) handles almost all known clustering objectives as well as multiple metric spaces. Our analysis relies on a new extremal property of metric spaces, that we call scatter dimension.

Peter Manohar: *New Spectral Techniques in Algorithms, Combinatorics, and Coding Theory: the Kikuchi Matrix Method*
In this talk, we present a new method to solve algorithmic and combinatorial problems by (1) reducing them to bounding the maximum, over x in $\{0,1\}^n$, of homogeneous degree-q multilinear polynomials, and then (2) bounding the maximum value attained by these polynomials by analyzing the spectral properties of appropriately chosen induced subgraphs of Cayley graphs on the hypercube (and related variants) called "Kikuchi matrices".

We will present the following applications of this method.

(1) Designing algorithms for refuting/solving semirandom and smoothed instances of constraint satisfaction problems;

(2) Proving Feige's conjectured hypergraph Moore bound on the extremal girth vs. density trade-off for hypergraphs;

(3) Proving a cubic lower bound for 3-query locally decodable codes and an exponential lower bound for 3-query locally correctable codes.

Daniel Dadush: *Column Bounds for the Circuit Imbalance Measure*
For a real matrix $A \in \mathbb{R}^{d\times n}$ with non-collinear columns, we show that $n \leq O(d^4\kappa_A)$ where κ_A is the *circuit imbalance measure* of A. The circuit imbalance measure κ is a real analogue of Δ-modularity for integer matrices, satisfying $\kappa_A \leq \Delta_A$ for integer A, which has numerous applications in the context of linear programming.

Similar to the strategy of Geelen, Nelson and Walsh (2021) for proving column bounds for Δ-modular matrices, we reduce to upper bounding the number of elements of a simple real (or complex) representable matroid excluding a line minor. As our main technical contribution, we show that any simple rank d complex representable matroid which excludes a line of length l has at most $O(d^4 l)$ elements. This complements the tight bound of $(l-3)\binom{d}{2}+d$, $l \geq 4$, of Geelen, Nelson and Walsh which holds when the rank d is at least doubly exponential in l.

Our proof of the above relies on an improvement of a Sylvester-Gallai type theorem of Dvir, Saraf and Wigderson (2014). Refining their design matrix technique, we show that for any full-dimensional set of n points in $\mathbb{C}^d$ there always exists a point that lies on at least $(1-\frac{4}{d})n$ many distinct lines (the constant 4 is improved from 12). The excluded minor bound follows by inductively applying this result to find good elements to contract in the matroid, where the improved constant reduces the dependence on d from d^{12} to d^4. Interestingly, by relying geometric techniques, our proof avoids the use of any difficult matroid machinery.