Synergies in Combinatorics and Theoretical Computer Science

Timetable

	Monday	Tuesday	Wednesday	Thursday	Friday
9am	Richard Montgomery	Max Hopkins	Venkatesan Guruswami	Ruixiang Zhang	Amey Bhangale
10am	Coffee	Coffee	Coffee	Coffee	Coffee
10.30am	Nikhil Bansal	Tim Hsieh	Zander Kelley	Dimitrii Zakharov	Yang Liu
11.30pm	Lunch	Lunch	Lunch	Lunch	Lunch
1.30pm	István Tomon	Rachel Greenfeld	Free	Oliver Janzer	Boris Bukh
2.30pm	Coffee	Coffee	Free	Coffee	Coffee
3pm	Kewen Wu	James Leng	Free	Luke Postle	Jacques Verstraëte
4pm		Problem Session	Free		

Titles and Abstracts

Speaker: Nikhil Bansal

Title: On Beck–Fiala and Komlós Conjectures

Abstract: A conjecture of Komlós states that the discrepancy of any collection of unit vectors is $O(1)$, i.e., for any matrix A with unit columns, there is a vector x with $-1, 1$ entries such that $|Ax|_\infty = O(1)$. The related Beck–Fiala conjecture states that any set system with maximum degree k has discrepancy $O(k^{1/2})$.

I will describe an $O((\log n)^{1/4})$ bound for the Komlós problem, improving upon an $O((\log n)^{1/2})$ bound due to Banaszczyk. Time permitting, we will see how these ideas can be used to resolve the Beck-Fiala conjecture for $k \geq (\log n)^2$.

Speaker: Amey Bhangale

Title: CSPs and Inverse Theorems for k-wise Correlations of Functions.

Abstract: Constraint satisfaction problems (CSPs) provide a rich source of algebraic and analytic questions at the interface of combinatorics, optimization, and

computational complexity. A landmark theorem of Bulatov and Zhuk classifies the complexity of every finite constraint satisfaction problem, showing that each is either solvable in polynomial time or NP-complete.

This dichotomy raises a natural quantitative question: for an NP-complete problem, what is the optimal approximation threshold? More precisely, for each constraint predicate P, one seeks the largest constant $\alpha(P)$ for which there is a polynomial-time algorithm that, on every satisfiable instance, finds an assignment satisfying at least an $\alpha(P)$ fraction of the constraints, while achieving $(\alpha(P) + \epsilon)$ for any $\epsilon > 0$ is NP-hard. Beyond a few examples, such as Håstad's celebrated 7/8 threshold for Max 3SAT, little is known about these optimal constants.

I will present our work on a systematic study of these approximation thresholds. The main ingredients are new inverse theorems for k-wise correlations of functions. I will also discuss applications of these inverse theorems to problems in additive combinatorics.

This talk will be based on joint work with Subhash Khot, Yang P. Liu, and Dor Minzer.

Speaker: Boris Bukh

Title: Faster random walk via infrequent steering

Abstract: Random walks on graphs can mix slowly. To speed it up, imagine that at each step instead of choosing the neighbor at random, there is a small probability $\epsilon > 0$ that we can choose it. We show that in this case, at least for graphs of bounded degree, there is a way to steer the walk so that we visit every vertex in $n^{1+o(1)}$ many steps. The key to this result is a way, due to Tung and Ueltzen, to decompose arbitrary graphs into small-diameter pieces.

Joint work with Quentin Dubroff.

Speaker: Rachel Greenfeld

Title: Between translational tilings and configurations of low pattern complexity

Abstract: I will discuss an interesting connection between these structures, and how it can be leveraged to study Nivat's conjecture in symbolic dynamics.

Speaker: Venkatesan Guruswami

Title: The Richness of CSP Non-Redundancy

Abstract: An instance of a constraint satisfaction problem is called non-redundant if no constraint is implied by the rest; that is, for each constraint, there exists an

assignment that uniquely violates that constraint. The non-redundancy (NRD) of a relation R is the maximum number of constraints in a non-redundant CSP instance over R, as a function of the number n of variables. In recent work, we showed that the extent to which a CSP can be sparsified is governed by the NRD of its associated relation. NRD has also been identified as a key parameter in the kernelization and streaming complexity of CSPs. More broadly, non-redundancy lies at an interesting nexus of logic, extremal combinatorics, and algebra, motivating its detailed study.

In this talk, I will report several results that illuminate the rich landscape of NRD. The NRD of an arity-r relation over a finite domain lies between n and n^r, and when $r = 2$, exhibits a sharp linear-versus-quadratic dichotomy. Prior to recent work, all known examples were consistent with NRD exhibiting only polynomial growth n^c for integer exponents c. We show that the behavior of NRD is substantially richer. In particular, we prove that a simple relation over $\{0, 1, 2\}^3$ has NRD between $n^{1.5}$ and $n^{1.6}$, and that for every rational $s \geq 1$, there exists a relation with NRD $= \Theta(n^s)$. Using connections to high-girth constructions, we show that infinitely many exponents arise as the NRD of relations even for arity 3. Even for Boolean predicates, the NRD does not always have polynomial growth—we exhibit an arity 4 Boolean relation with NRD $n^{3-o(1)}$.

We also revisit Mal'tsev embeddings, the most general sufficient condition currently known for $O(n)$ NRD, and resolve several structural questions about them, using a new tool called Catalan polymorphisms. In particular, we give an example of a relation that admits no Abelian embedding but does embed into a non-Abelian group.

Based on joint works with Joshua Brakensiek, Bart Jansen, Victor Lagerkvist, Aaron (Louie) Putterman, and Magnus Wahlström, appearing in STOC 2025, CP 2026, and FOCS 2026.

Speaker: Max Hopkins

Title: On Codes, Hard Functions, and High Dimensional Expanders

Abstract: Expander graphs - sparse, robustly connected graphs that mimic the structure of the complete graph - form a classical interface between combinatorics, algebra, and TCS. In recent years, the theory of high dimensional expanders (HDX) has deepened this interface, extending the success of expanders to hypergraphs and beyond and a slew of surprising connections across TCS and mathematics.

In this talk, I will introduce local-spectral high dimensional expanders and describe their recent applications to two longstanding problems in complexity and coding theory: hardness amplification and efficient list decoding. In particular, we will see how these objects, which mimic many of the useful properties of independent variables, yield the first fully derandomized XOR lemma and near-optimal local and parallel algorithms for error correction in the high-noise

regime. Finally, I will explain concretely how local-spectral expansion enters these results through a new belief-propagation algorithm on HDX, and how this algorithm may be executed in polylogarithmic time via a new primitive we call (fault-tolerant) strongly explicit routing: the ability to sample short low-congestion routes between arbitrary vertices in time polylogarithmic in the size of the graph (even in the presence of non-responsive nodes).

Based on joint works with Yotam Dikstein, Russell Impagliazzo, and Toniann Pitassi.

Speaker: Tim Hsieh

Title: Explicit Lossless Vertex Expanders

Abstract: In this talk, I will present our construction of explicit lossless vertex expanders. Specifically, we construct d-regular graphs where every small set S of vertices has close to $d|S|$ neighbors. Our construction is based on the tripartite line product between a constant-sized lossless expander and a base graph derived from Ramanujan cubical complexes. In the talk, I will explain the tripartite line product and discuss the role of high-dimensional expanders in the construction.

Based on joint works with Ting-Chun Lin, Alex Lubotzky, Theo McKenzie, Sidhanth Mohanty, Ryan O'Donnell, Pedro Paredes, Assaf Reiner, and Rachel Zhang.

Speaker: Oliver Janzer

Title: Hamiltonicity of regular sublinear expanders

Abstract: We prove that regular sublinear spectral expanders with at least polylogarithmic degree are Hamiltonian. We also prove a similar result for bipartite one-sided expanders. As applications, we obtain highly robust versions of recent important results on the Hamiltonicity of Cayley graphs and Kneser graphs.

Joint work with Domagoj Bradac.

Speaker: Zander Kelley

Title: A More Efficient Sifting Lemma and a Stronger 3-Player Communication Lower Bound

Abstract: A central goal in complexity theory is to prove separations between randomized and deterministic models of computation. In communication complexity, a key challenge is to establish lower bounds for multiparty protocols in the number-on-forehead (NOF) model. Recent work by Kelley, Lovett, and

Meka provided a separation for an explicit 3-player function, proving a deterministic lower bound of $\Omega(n^{1/3})$ for a function with an efficient randomized protocol.

The proof of this lower bound involved showing that the 'yes' instances of a function cannot be efficiently covered by small "cylinder intersections" – the basic unit of NOF communication. This analysis hinges on a sifting lemma for bipartite graphs, which guarantees that any graph with a large "grid norm" must contain a smaller, much denser induced subgraph. In this talk, based on joint work with Xin Lyu, I will discuss a new, more efficient version of this sifting lemma. This strengthening of the core technical tool allows us to improve the 3-player lower bound to $\Omega(n^{1/2})$, achieving a stronger separation.

Specifically, we prove a new structural result about small cylinder intersections: that they can be covered by a few reasonably small "slice functions". I will explain how this result can be productively compared with Szemerédi's Triangle Removal Lemma for graphs, which also gives a way to cover small cylinder intersections by something simpler.

Speaker: James Leng

Title: Quasipolynomial Bounds for the ($U^{s+1}[N]$) Inverse Theorem

Abstract: In his 2001 proof of Szemerédi's theorem, Gowers introduced the $U^{s+1}[N]$ norm and established the first quantitative bounds for progressions of arbitrary fixed length. His work raised a fundamental structural question: what can be said about a bounded function whose Gowers norm is large? The qualitative answer was provided by Green, Tao, and Ziegler through the inverse theorem for the $U^{s+1}[N]$ norm: a function with large $U^{s+1}[N]$ norm must correlate with a structured object arising from an s-step nilsequence. However, the quantitative bounds in this theorem were extremely weak. In this talk, I will discuss recent work establishing quasipolynomial bounds for the $U^{s+1}[N]$ inverse theorem. I will explain some of the main ideas behind the proof and how the resulting quantitative structure theorem leads to an improved bound of $N\exp(-(\log\log N)^{c_k})$ for the largest subset of $[N]$ containing no nontrivial k-term arithmetic progression.

This is based on joint work with Ashwin Sah and Mehtaab Sawhney.

Speaker: Yang Liu

Title: Proof of the 3-CSP Inverse Theorem and Applications to Combinatorial Lines

Abstract: Let f, g, h be 1-bounded functions from $S^n \to \mathbb{C}$ for a finite set S satisfying $|E_{(x,y,z)\sim\mu^n}[f(x)g(y)h(z)]| \geq \epsilon$, where μ is a pairwise-connected distribution on $S \times S \times S$. We prove that a random restriction of f must correlate

to a product function, i.e., a function of the form $P(x) = P_1(x_1) \cdots P_n(x_n)$ for x in S^n. The proof goes by introducing an analogue of the U^2 Gowers norm for this setting, which we call the swap norm, and proving an inverse theorem for it. As applications, we give the first reasonable bounds for restricted 3-APs over $\mathbb{F}_p^n$ and for the Density Hales–Jewett problem for combinatorial lines of length 3.

Based on joint work with Amey Bhangale, Subhash Khot, and Dor Minzer.

Speaker: Richard Montgomery

Title: Minimising the reciprocal sum of cycle lengths

Abstract: A central theme in extremal graph theory is to understand the relationship between the density of a graph and the different lengths of the cycles it contains. The cycle-length spectrum, denoted $\mathcal{C}(G)$, is the set of distinct cycle lengths occurring in G. In 1966, Erdős and Hajnal proposed using

$$\sum_{\ell \in \mathcal{C}(G)} \frac{1}{\ell},$$

the sum of the reciprocals of the elements of $\mathcal{C}(G)$ to quantify the richness of the cycle lengths in a graph G. A sequence of increasingly ambitious conjectures culminated in the 1981 conjecture of Erdős that the complete bipartite graph $K_{k,n-k}$ minimises this quantity among all n-vertex graphs with at least $k(n-k)$ edges when n is at least $2k$. I will discuss this conjecture and its proof for all sufficiently large k.

This is joint work with Aleksa Milojević, Alexey Pokrovskiy and Benny Sudakov.

Speaker: Luke Postle

Title: A Proof of Nash-Williams' Conjecture

Abstract: A central open question in extremal design theory is Nash-Williams' Conjecture from 1970, namely, that every triangle-divisible graph on n vertices (for n large enough) with minimum degree at least $0.75n$ has a triangle decomposition. In this talk, we discuss the history of the problem and our recent resolution of this conjecture, as well as other applications in design theory. We also overview the proof, highlighting the new techniques we developed to resolve the fractional version as well as the full conjecture.

Joint work with Michelle Delcourt.

Speaker: István Tomon

Title: The Combinatorics of Factorization Norms

Abstract: I will give a gentle introduction to the ℓ_2-factorization norm of matrices, also known as γ_2-norm or max-norm. While this norm might be familiar to researchers in communication complexity, it may be less well known among combinatorialists. My goal is to convince the audience that this is an interesting quantity to study combinatorially, discussing it from the perspective of extremal and structural graph theory.

Speaker: Jacques Verstraëte

Title: Recent Progress on the Erdős–Rogers Function

Abstract: We let $f_{H,G}(n)$ denote the maximum k such that every n-vertex G-free graph contains an induced F-free subgraph with k vertices. When $F = K_2$ these are determined by Ramsey numbers $r(k, G)$, and the case $F = K_s$ and $G = K_{s+1}$ are the classical Erdős–Rogers functions, denoted $f_{s,s+1}(n)$. In this talk we survey recent progress, including the recent result of Morris, Saharasbudhe and the author showing for each $s \geq 2$ that $f_{s,s+1}(n) = \Theta(\sqrt{n \log n})$ as $n \to \infty$.

Joint work with Rob Morris and Julian Saharasbudhe.

Speaker: Kewen Wu

Title: Locally sampleable distributions: Hamming slices, symmetry, and quantum advantage

Abstract: What distributions can shallow circuits produce if given independent random bits? I will discuss recent progress on this question, focusing on symmetric distributions over the Boolean cube. We will see that even simple-looking targets, such as Hamming slices, will impose strong locality lower bounds. These lower bounds lead to a classification of locally sampleable (uniform) symmetric distributions and several applications in data structure lower bounds and quantum-classical separations.

This is based on joint works with Daniel Grier, Daniel Kane, Jackson Morris, and Anthony Ostuni.

Speaker: Dimitrii Zakharov

Title: Sticky Kakeya estimate in four dimensions

Abstract: The (discretized) Kakeya problem is a question about how much can collections of tubes overlap. A recently developed strategy for proving sharp

bounds for Kakeya and other projection theory questions is to solve the "sticky" case of the problem and then reduce the general problem to that case. This is how Wang and Zahl solved the 3-dimensional case of the Kakeya dimension in 2022–2025. In this work, we prove a sharp sticky Kakeya estimate in $\mathbb{R}^4$. As a byproduct we also obtain a new proof in dimension 3. Ultimately, the argument reduces these questions to different versions of sum-product-type theorems. Joint work with Hong Wang.

Speaker: Ruixiang Zhang

Title: Arithmetic Kakeya and mixed-norm multilinear inequalities

Abstract: I will introduce Katz-Tao's Arithmetic Kakeya Conjecture and connections with mixed-norm multilinear inequalities first studied by Christ. I will also make some comparisons between the mixed-norm inequalities and classical Brascamp-Lieb inequalities.