Workshop
"Analysis in TCS: testing, learning, and complexity"

November 4 - 8, 2024

organized by
**Sergey Bobkov, Polona Durcik, Alexandros Eskenazis, Irina Holmes Fay,
Paata Ivanisvili, Dor Minzer, Joseph Slote, and Alexander Volberg**

Abstracts

Srinivasan Arunachalam (IBM)

Testing and learning phase states and their variants

Abstract: In these talks I will discuss two topics: (i) tolerantly testing and learning stabilizer states and (ii) learning degree-d phases and testing degree-2 phase states.

Francisco Escudero Gutiérrez (CWI, University of Amsterdam)

An approach towards the Aaronson–Ambainis conjecture via Fourier completely bounded polynomials

Abstract: Aaronson and Ambainis (AA) conjectured in 2008 that low-degree polynomials bounded in the infinity norm have an influential variable. If true, it would imply that quantum computers can only achieve major speedups in a certain class of problems. We will discuss a similar conjecture, where the polynomials are Fourier completely bounded. Informally speaking, a polynomial is Fourier completely bounded if it is bounded when it is evaluated on matrix inputs. In particular, being Fourier completely bounded implies being bounded in the infinity norm, so this new conjecture is weaker than AA conjecture. However, this weaker conjecture has the same implications in quantum computing.

We will give an introduction to AA conjecture, explain why the mentioned weaker conjecture preserves the implications in quantum computing and finally prove a particular case of the weaker conjecture.

Based on *Influences of Fourier Completely Bounded Polynomials and Classical Simulation of Quantum Algorithms*, Francisco Escudero Gutiérrez, Chicago Journal of Theoretical Computer Science.

Tom Gur (University of Cambridge)

Higher-order Fourier analysis in quantum complexity theory

Abstract: A recent line of work obtained a framework of worst-case to average-case reductions for quantum algorithms. The machinery uses tools from additive combinatorics, such as the quasi-polynomial Bogolyubov lemma, to exact linear structures and obtain local correction procedures. However, such techniques are limited to linear problems. We will present a new approach that uses higher-order Fourier analysis to extract low-degree structures towards the end of breaking the linearity barrier. The talk is based on ongoing work with Alexander Golovnev and Julia Wolf.

Hamed Hatami (McGill University)

Sparse graph counting and Kelley–Meka bounds for binary systems

Abstract: The graph counting lemma of Chung, Graham, and Wilson (Combinatorica 1988) is a fundamental result in combinatorics, which states that if a large graph is pseudorandom, then the number of copies of any small graph H in G is close to what is expected from a random graph of the same density. However, this result is only nontrivial when G is a dense graph. In this work, we obtain a counting lemma that works in the sparse setting, and it is well-suited for the density increment arguments in additive combinatorics.

In a recent remarkable breakthrough, Kelley and Meka (FOCS 2023) obtained a strong upper bound on the density of sets of integers without nontrivial three-term arithmetic progressions. We combine our counting lemma with other ideas to establish Kelley–Meka type bounds for all linear patterns defined by translation-invariant systems of binary linear forms, *i.e.*, each form depends on exactly two variables. In particular, we obtain strong bounds for the Turan problem in Abelian Cayley graphs, *i.e.* an upper bound on the maximum edge density of an Abelian Cayley graph with a clique of a given size. To prove our results, we employ some of the recent technology developed by Kelley and Meka and also the follow-up work by Kelley, Lovett, and Meka (STOC 2024). This talk is based on a joint work with Esty Kelman, Yuval Filmus, and Kaave Hosseini.

Pooya Hatami (Ohio State University)

Constant-Cost Communication

Abstract: Some of the most extreme examples of the power of randomness in computing come from communication complexity, where shared randomness can allow two parties to solve non-trivial problems with communication cost independent of the input size. The canonical example is the Equality problem, where two parties hold n-bit strings x and y, respectively, and they wish to decide whether $x = y$. While n bits of deterministic communication are necessary, there exists a simple randomized protocol that solves this by communicating only 2 bits. Can we characterize all communication problems that admit such hyperefficient randomized protocols? What are the structural properties of the Boolean matrices corresponding to such problems? We discuss recent progress and open problems. Based on joint works with Yuting Fang, Mika Göös, Lianna Hambardzumyan, Nathaniel Harms, and Hamed Hatami.

Ohad Klein (Hebrew University)

Title: Slicing all edges of an n-cube requires $n^{2/3}$ hyperplanes

Abstract: Consider the n-cube graph in $\mathbb{R}^n$, with vertices $\{0,1\}^n$ and edges connecting vertices with Hamming distance 1. How many hyperplanes are required in order to dissect all edges? This problem has been open since the 70s. We will discuss this and related problems.

Puzzle: Show that n hyperplanes are sufficient, while $\sqrt{n}$ are not enough.

Dmitry Krachun (Princeton University)

An optimal space lower bound for approximating MAX-CUT

Abstract: In this talk I will show how insights going beyond classical KKL/hypercontractivity bounds allow one to perform a tight analysis of a specific multiplayer communication problem called the *Implicit Hidden Partition* (IHP) problem which leads to a tight lower bound for the space complexity of a streaming algorithm which breaks the 2-approximation barrier for the MAX-CUT value. In doing so I will describe how Fourier analysis on the boolean cube can be used as a language to argue about propagation of information in communication problems. Based on joint work with Michael Kapralov.

Avichai Marmor (Bar-Ilan University)

Permutation mixing via hypercontractivity of symmetric characters

Abstract: The topic of permutation mixing asks, for a given sequence of sets of permutations $B_1, \ldots, B_k \subseteq S_n$, whether the product $x_1 \ldots x_k$ is approximately uniformly distributed, where each x_i is uniformly sampled from B_i. Permutation mixing has been extensively studied over the past few decades, starting with the seminal work of Diaconis and Shahshahani in 1981 on deck shuffling. One major problem in this topic asks to determine the minimal size of a conjugacy-closed set B that ensures it has a mixing-time of k, *i.e.*, that multiplying k uniform elements of B results in mixing. In 2008, Larsen and Shalev [*Invent. math.* 174, 645–687 (2008)] found an almost tight bound on this size, employing novel methods to bound the characters of the symmetric group. These bounds rely heavily on specific properties of the symmetric group and its character formulas.

In this talk, we present a new approach to bounding symmetric characters, drawing on a recently-developed tool from the theory of Boolean functions known as 'hypercontractivity for global functions'. This method improves the bounds obtained by Larsen-Shalev and relies on only a few key properties of symmetric characters—properties that have analogues in various classical groups. This raises the possibility of extending the technique to other groups. After introducing the relevant tools from representation theory, we will outline the proof of the new bound and deduce improved bounds on the size that ensures a set has mixing time k.

The talk is based on joint work with Noam Lifshitz, and no prior knowledge is required.

Dan Mikulincer (University of Washington)

Spectral gaps for measures on the cube via a generalized stochastic localization process

Abstract: Stochastic localization processes are an emerging set of techniques useful in the study of high-dimensional probability. At its core lies the idea that *a priori* complicated measures can be simplified by performing randomized Gaussian tilts. These tilts then allow for the use of various Gaussian comparison inequalities. In many cases, the Gaussian framework is well understood, and this strategy has led to advances in various fields, from convex geometry to complexity theory.

In this talk, motivated by algorithmic questions related to spectral gaps of discrete measures, we will go beyond the Gaussian framework and introduce a variant of the stochastic localization process that allows for non-Gaussian tilts. In the first part of the talk, we will present the original localization process and explain the main natural difficulties that arise when moving beyond the standard setting. The second part of the talk will focus on overcoming these hurdles.

Shivam Nadimpalli (Massachusetts Institute of Technology)

High-dimensional convexity: testing, learning, and complexity

Abstract: I will survey a recent line of work exploring a powerful—but partially understood—analogy between several algorithmic and structural aspects of (i) high-dimensional convex sets over Gaussian space, and (ii) monotone Boolean functions over $\{0,1\}^n$. The talk will be based on a sequence of works with collaborators Anindya De and Rocco Servedio.

Joris Roos (University of Massachusetts Lowell)

Sharp isoperimetric inequalities on the Hamming cube

Abstract: We will discuss an isoperimetric problem for the Hamming cube that goes back to Talagrand and arises from an interpolation between the classical edge- and vertex-isoperimetric problems. In this talk the focus will be on the role of certain two-point inequalities.

Ohad Sheinfeld (Bar-Ilan University)

Improved Covering Results and Intersection Theorems in Symmetric Groups, via Hypercontractivity

Abstract: In this talk, we present two seemingly unrelated results on the symmetric group S_n. A well-known problem asks for a characterization of subsets A of A_n whose square A^2 covers the entire alternating group A_n. We show that if A is a conjugacy class of density at least $\exp(-n^{2/5} - \varepsilon)$, then $A_n = A^2$. This improves upon a seminal result by Larsen and Shalev (*Invent. math.*, 2008), who obtained the same result with $1/4$ in the double exponent. The second problem is an analogue of the Erdős–Sós forbidden intersection problem for families of permutations. A family $\mathcal{F}$ of permutations is called non-$(t-1)$-intersecting if any two permutations in $\mathcal{F}$ do not agree on exactly $(t-1)$ values. We prove that for $t < \frac{\sqrt{n}}{\log(n)}$, the maximal size of a non-$(t-1)$-intersecting family is $(n-t)!$, which improves upon the result of Kupavskii and Zakharov (Adv. in Math., 2024), who proved the same for $t = \tilde{O}(n^{1/3})$.

Both results can be interpreted as bounds on the size of independent sets in certain normal Cayley graphs over symmetric groups. The primary tool we use in both proofs to establish these bounds are the recently obtained hypercontractive inequalities for global functions, developed by Keller, Lifshitz, and Marcus, as well as by Keevash and Lifshitz.

The talk is based on joint work with Nathan Keller and Noam Lifshitz.

Kewen Wu (UC Berkeley)

Quantum State Preparation with Optimal T-Count

Abstract: How many T gates are needed to approximate an arbitrary n-qubit quantum state to within a given precision ϵ? Improving prior work of Low, Kliuchnikov, and Schaeffer, we show that the optimal asymptotic scaling is $\sqrt{2^n \log(1/\epsilon)} + \log(1/\epsilon)$ if we allow ancilla qubits. We also show that this is the optimal T-count for implementing an arbitrary diagonal n-qubit unitary to within error ϵ. We describe applications in which a tensor product of many single-qubit unitaries can be synthesized in parallel for the price of one.