Harvard Theory of Computing Seminar

About: The Harvard Theory of Computing seminar is the primary seminar series of the Harvard Theory of Computation Group.

Time and place: The meetings happen roughly every other **Wednesday** from **3:45-5pm.** (Typically there are refreshments before the seminar from 3:15-3:45, outside the seminar room. The talks end at 4:45 with 4:45-5:00pm reserved for Q&A.) The meetings in Fall 2024 will be held in the Harvard SEC, room 3.301-3.303 (tentative)

Mailing list: To join the theory seminar mailing list, please request access to the Theory of Computing Google Group or email Emma at eduncan@seas.harvard.edu.

Calendar: Here's a link to add our google calendar, or you can view it here.

Fall 2024 Organizers: Anurag Anshu, Sitan Chen, Allison Choat, Rebekah Keely, Aaron (Louie) Putterman, Jake Ruotolo, Emma Duncan

Upcoming Talks:

- 09/13/2024: Special seminar: Avi Wigderson (more details)
- 09/18/2024: Postdoc day (more details)
- 09/25/2024: Jason Li (more details)
- 10/09/2024: Siqi Liu (more details)
- 10/17/2024: Special seminar: Claudia Goldin (more details)
- 10/23/2024: Sinho Chewi (more details)
- 11/06/2024: Joan Bruna (more details)
- 11/20/2024: Abhradeep Thakurta (more details)
- 12/04/2024: Jinyoung Park (more details)
- 12/11/2024: Toniann Pitassi (more details)

Past Talks:

Spring 2024 (info):

- 01/24/2024: Postdoc day (more details)
- 02/07/2024: Pravesh Kothari (more details)
- 02/21/2024: Jane Lange (more details)
- 03/06/2024: Raghu Meka (more details)
- 04/03/2024: Yael Tauman Kalai (more details)
- 04/17/2024: Mohsen Ghaffari (more details)
- 04/24/2024: Sanjeev Khanna (more details)

Older seminars: Here's a partial list.

September 13, 2024: Avi Wigderson

Co-hosted by the CS Lecture Series and the Theory of Computing Seminar Series

Time: 3:45 - 5:00 p.m.

Location: Winokur Family Hall 1.321

Title: **The Value of Errors in Proofs**

Speaker: **Avi Wigderson (Institute for Advanced Study, Princeton University)**

Abstract:

A few years ago, a group of theoretical computer scientists posted a paper on the Arxiv with the strange-looking title "MIP* = RE", surprising and impacting not only complexity theory but also some areas of math and physics. Specifically, it resolved, in the negative, the "Connes' embedding conjecture" in the area of von-Neumann algebras, and the "Tsirelson problem" in quantum information theory. It further connects Turing's seminal 1936 paper which defined algorithms, to Einstein's 1935 paper with Podolsky and Rosen which challenged quantum mechanics. You can find the paper here: https://arxiv.org/abs/2001.04383.

As it happens, both acronyms MIP* and RE represent proof systems, of a very different nature. To explain them, we'll take a meandering journey through the classical and modern definitions of *proof*. I hope to explain how the methodology of computational complexity theory, especially modeling and classification (of both problems and proofs) by algorithmic efficiency, naturally leads to the generation of new such notions and results (and more acronyms, like NP). A special focus will be on notions of proof which allow interaction, randomness, and errors, and their surprising power and magical properties.

September 18, 2024: Postdoc Day

Time: 3:45 - 5:00 p.m.

Location: SEC 2.122 + 2.123

Title: **Postdocs Day**

Speakers: **Lunjia Hu**, **Han Shao**, **Avi Karchmer**, **Nikhil Vyas**

Abstracts:

Lunjia Hu: **Calibration Error for Decision Making**

Where predictions gain value by assisting decision making, calibration allows predictions to be reliably interpreted as probabilities by the decision makers. We propose a decision-theoretic calibration error, the Calibration Decision Loss (CDL), defined as the maximum improvement in decision payoff obtained by calibrating the predictions, where the maximum is over all payoff-bounded decision tasks. Vanishing CDL guarantees the payoff loss from miscalibration vanishes simultaneously for all downstream decision tasks. We show a separation between CDL and existing calibration error metrics, including the most well-studied metric Expected Calibration Error (ECE). Our main technical contribution is a new algorithm for online calibration that achieves O(sqrt(T)log T) expected CDL, bypassing the Ω(T^{0.528}) lower bound for ECE by Qiao and Valiant (2021).

Joint work with Yifan Wu. To appear in FOCS 2024.

Han Shao: **Learning From Strategic Data Sources**

In contrast with standard classification tasks, strategic classification involves agents strategically modifying their features in an effort to receive favorable predictions. For instance, given a classifier determining loan approval based on credit scores, applicants may open or close their credit cards to fool the classifier. The learning goal is to find a classifier robust against strategic manipulations. We study the fundamental mistake bound and sample complexity in the strategic classification.

Avi Karchmer: **Cryptography and Complexity Theory in the Design and Analysis of Machine Learning**

Cryptography and Complexity vs. Machine Learning---generally, the former two are in opposition to the latter. But when life gives you lemons... make lemonade!

- Cryptography can help us design more secure and private ML algorithms
- We may “mine” technical machinery from Complexity theory for faster and more robust ML algorithms

● Both Cryptography and Complexity can help us reason about the ML “real world” (e.g., when is training on text and images more effective than training on just text?)
This "Preview" talk will highlight to the audience various situations in which Cryptography and Complexity theory can positively impact the design and analysis of Machine Learning.

Nikhil Vyas: **Understanding and Improving the Shampoo Optimizer**

Shampoo, a second-order optimization algorithm utilizing a Kronecker product preconditioner, has recently gained attention within the machine learning community. We show how Shampoo can be interpreted as a Kronecker-factored approximation of Newton's method. Building on this perspective, we explore modifications to Shampoo that lead to enhanced performance.
Based on works with coauthors David Brandfonbrener, Lucas Janson, Sham Kakade, Eran Malach, Depen Morwani, Itai Shapira, and Rosie Zhao.

September 25, 2024: Jason Li (Carnegie Mellon University)

Time: 3:45 - 5:00 p.m.

Location: SEC 3.301-3.303

Title: **Minimum Isolating Cuts: A new tool for solving minimum cut problems**

Speaker: **Jason Li (Carnegie Mellon University)**

Jason Li is an assistant professor at CMU working on fundamental graph optimization problems such as minimum cut and maximum flow. His research efforts have led to state-of-the-art algorithms for a wide array of classic problems, including deterministic global minimum cut, minimum k-way cut, Gomory-Hu tree or all-pairs minimum cut, and single-source shortest path in parallel.

Abstract:

Minimum cut problems are among the most well-studied questions in combinatorial optimization. In this talk, I will introduce a simple but powerful new tool for solving minimum cut problems called the minimum isolating cuts. I will show how this tool can be employed to obtain faster algorithms for several fundamental min-cut problems, namely global min-cut, Steiner min-cut, and all-pairs min-cut. For these problems, the new results represent the first improvement in their runtimes in several decades.

These results are in collaboration with Amir Abboud, Robert Krauthgamer, Danupon Nanongkai, Thatchaphol Saranurak, and Ohad Trabelsi.

October 9, 2024: Siqi Liu (DIMACS/IAS)

Time: 3:45 - 5:00 p.m.

Location: SEC 3.301-3.303

Title: **Sparse spectral and coboundary high dimensional expanders**

Speaker: **Siqi Liu (DIMACS/IAS)**

Siqi Liu is a postdoc at the Institute for Advanced Study hosted by Avi Wigderson. She is generally interested in theoretical computer science. Most recently she is interested in graphs and their applications to coding theory, property testing, and complexity.

Abstract:

High dimensional expanders (HDXs) are a class of expanding hypergraphs. Their importance is constantly growing with their roles in the breakthroughs in Markov chains [ALGV19], locally testable codes [DELLM22], and quantum codes [PP22, ABN23]. HDXs generalize the notion of expander graphs but differ from the latter in a couple ways. First, while various definitions of expansion are equivalent over graphs, they are different in hypergraphs, yielding different types of HDXs. Secondly, while sparse expanders are abundant in random graph models, bounded-degree HDXs seem to be rare in random hypergraph models.

So far the sparest random construction of HDXs comes from random geometric graphs (RGGs) [LMSY23]. In this talk we will introduce a deterministic construction of HDXs which can be viewed as a derandomization of the RGG construction. These hypergraphs are both spectral HDXs and coboundary HDXs. They are the sparsest known coboundary HDXs (over any group). In particular, their 1-skeletons are abelian Cayley graphs.

This is based on joint work with Yotam Dikstein and Avi Wigderson.

October 17, 2024: Claudia Goldin

Time: 4:30pm- 5:30 pm

Location: SEC LL2.224

Title: **Why Women Won**

Speaker: **Claudia Goldin**

Claudia Goldin is the Henry Lee Professor of Economics at Harvard University and the 2023 Nobel Laureate in Economic Sciences. She was the director of the NBER's Development of the American Economy program from 1989 to 2017 and is currently co-director of the NBER's Gender in the Economy group. Most of her research interprets the present through the lens of the past and explores the origins of issues of current concern. She is the author of many books including: *Understanding the Gender Gap: An Economic History of American Women*, *The Race between Education and Technology,* and *Career & Family: Women's Century-Long Journey Toward Equity*. Goldin was the president of the American Economic Association and of the Economic History Association. She is a member of the National Academy of Sciences and the American Philosophical Society. She is a fellow of the American Academy of Political and Social Science, the American Academy of Arts and Sciences, and the Society of Labor Economists. She received the Nobel Prize in Economics in 2023, the Nemmers Prize in Economics in 2020 and the BBVA Prize in 2019. She received her BA from Cornell University and her PhD from the University of Chicago. She is an outdoor enthusiast, a bird-watcher, and has trained Pika, her Golden Retriever, to several performance scenting titles.

Abstract:

How, when, and why did women in the US obtain legal rights equal to men's regarding the workplace, marriage, family, Social Security, criminal justice, credit markets, and other parts of the economy and society, decades after they gained the right to vote? The story begins with the civil rights movement and the somewhat fortuitous nature of the early and key women's rights legislation. The women's movement formed and pressed for further rights. Of the 155 critical moments in women's rights history, I've compiled from 1905 to 2023, 45% occurred between 1963 and 1973. The greatly increased employment of women, the formation of women's rights associations, the belief that women's votes mattered, and the unstinting efforts of various members of Congress were behind the advances. But women soon became splintered by marital status, employment, region, and religion far more than men. A substantial group of women emerged in the 1970s to oppose various rights for women, just as they did during the suffrage movement. They remain a potent force today.

October 23, 2024: Sinho Chewi (Yale University)

Time: 3:45 - 5:00 pm

Location: SEC 3.301-3.303

Title: **Mean-field Langevin dynamics and two-layer neural networks**

Speaker: **Sinho Chewi**

Abstract:

I will discuss recent progress on the theory of two-layer neural networks via the mean-field Langevin dynamics. In particular, I will discuss two key aspects of the finite-particle approximation to these dynamics, i.e., noisy gradient descent optimization of 2-layer neural networks: (1) how close these dynamics are to the mean-field limit, quantified by propagation of chaos, and (2) the rate of convergence to stationarity, quantified by a new log-Sobolev inequality that is independent of the number of particles. This is based on joint work with Murat A. Erdogdu, Yunbum Kook, Mufan (Bill) Li, Atsushi Nitanda, and Matthew S. Zhang.

November 6, 2024: Joan Bruna Estrach (NYU)

Time: 3:45 - 5:00 pm

Location: SEC 3.301-3.303

Title: **Posterior Sampling with Denoising Oracles via Tilted Transport**

Speaker: Joan Bruna Estrach
Joan Bruna is a Professor of Computer Science, Data Science and Mathematics (aff) at the Courant Institute and the Center of Data Science, New York University. He is also a visiting scholar at the Flatiron Institute. His research interests are in the mathematical aspects of Machine Learning and their applications to computational science. For his research contributions, he has been awarded the Sloan Fellowship, the NSF Career award, and several best paper awards.

Abstract:

Score-based diffusion models have significantly advanced high-dimensional data generation across various domains, by learning a denoising oracle (or score) from datasets, providing an efficient sampling scheme going beyond typical isoperimetric assumptions. From a Bayesian perspective, they offer a realistic modeling of data priors, used for solving inverse problems through posterior sampling. Although many heuristic methods have been developed recently for this purpose, they lack the quantitative guarantees needed in many scientific applications.

In this talk, we introduce the tilted transport technique, which leverages the quadratic structure of the log-likelihood in linear inverse problems in combination with the prior denoising oracle to transform the original posterior sampling problem into a new `boosted' posterior that is provably easier to sample from. We quantify the conditions under which this boosted posterior is strongly log-concave, highlighting the dependencies on the condition number of the measurement matrix and the signal-to-noise ratio. The resulting posterior sampling scheme is shown to reach the computational threshold predicted for sampling Ising models [Kunisky'23] with a direct analysis, and is further validated on high-dimensional Gaussian mixture models and scalar field phi-4 models.

Joint work with Jiequn Han (Flatiron Institute).

November 20, 2024: Abhradeep Thakurta

Time: 3:45 - 5:00 pm

Location: SEC 3.301-3.303

Title: Large Scale Private Learning on Data Streams, and the BLTs

Speaker: **Abhradeep Thakurta (Google DeepMind)**

Abhradeep Guha Thakurta is a staff research scientist in Google DeepMind (formerly Google Research - Brain Team). His primary research interest is in the intersection of data privacy and machine learning. He focuses on demonstrating, both theoretically and in practice, that it is possible to design differentially private learning algorithms that can scale to industrial workloads. Prior to Google, Abhradeep was a faculty at UC Santa Cruz. And even before that he worked at Apple and Yahoo Labs as research scientists. He did his Ph.D. from The Pennsylvania State University in 2013, and his postdoctoral appointment was with Stanford University and Microsoft Research Silicon Valley Campus.

Abstract:

In recent years, Differentially Private Follow the Regularized Leader (DP-FTRL) style algorithms have become increasingly popular for training large models over data streams, e.g., in training DP models for Gboard next-word prediction (https://research.google/blog/federated-learning-with-formal-differential-privacy-guarantees/ and https://arxiv.org/abs/2305.18465). At the heart of these algorithms is a correlated noise addition mechanism that factorizes a square lower triangular matrix of all ones (A=BC) and adds noise BZ (with Z ~ N(0,σ^2)^{training steps x dimensions}) to the prefix sum of the gradients observed during training. Prior work required generating and storing the noise matrix in advance, which is prohibitively expensive for large-scale model training.

In this talk, we introduce a new algorithmic construction called the Buffered Linear Toeplitz operator (BLT) that allows generating the noise required for DP in a streaming fashion while requiring a buffer size of approximately log^2(training steps) x dimensions. Furthermore, compared to the class of all Lower Triangular Toeplitz (LTT) factorizations (i.e., B and C being LTT), BLTs are only off by an additive constant in terms of utility under standard error metrics. The talk will draw ideas from rational function approximations and constant recurrence sequences for the construction of BLTs. The talk will also demonstrate the practicality of this work through simulation experiments.

The talk is primarily based on https://arxiv.org/pdf/2404.16706 (to appear in FOCS 2024), and is a joint work with Krishnamurthy (Dj) Dvijotham, Brendan McMahan, Krishna Pillutla, and Thomas Steinke.

December 4, 2024: Jinyoung Park

Time: 3:45 - 5:00 pm

Location: SEC 3.301-3.303

Title: Dedekind's Problem and beyond

Speaker: **Jinyoung Park** **(NYU)**

Jinyoung Park received a bachelor's degree in mathematics education from Seoul National University in 2005 and worked as a secondary school mathematics teacher until 2011. She earned her Ph.D. in mathematics from Rutgers University in 2020, under the supervision of Jeff Kahn. Following her Ph.D., she was a postdoctoral researcher at the Institute for Advanced Study and Stanford University before joining the faculty of the Courant Institute at NYU in 2023. Her research interests include extremal and probabilistic combinatorics. She has received several honors and awards, including the AMS Conant Prize (2025), the Sloan Fellowship (2024), the Dénes König Prize (2024), the Maryam Mirzakhani New Frontiers Prize (2023), and the AWM Dissertation Prize (2022).

Abstract:

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. 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 in 1980's. In this talk, we will discuss recent developments around Dedekind's Problem with connection to the cluster expansion method from statistical physics. Based on joint work with Matthew Jenssen and Alex Malekshahian.

December 11, 2024: Toniann Pitassi

Time: 3:45 - 5:00 pm

Location: SEC 3.301-3.303

Title: Proof Complexity: From Roots to Forest

Speaker: **Toniann Pitassi (Columbia University)**

Toniann Pitassi is the Jeffrey L. and Brenda Bleustein Professor of Computer Science at Columbia University. Pitassi's primary research area is complexity theory and proof complexity, which aim to understand the limits of proofs and computation. Recently her research interests have expanded to include theoretical aspects of machine learning, privacy and fairness, and whatever her graduate students are interested in.

Abstract:

One common view of proofs is an excessively formal set of rules that one must figure out how to apply in order to prove or derive something that is obvious in the first place. In this talk I will present proof complexity from a personal perspective, and will argue that proof systems are natural, and essential for understanding algorithms and the limits of computation. I will provide examples showing that proof systems hide behind many fundamental problems in mathematics and computer science, and how the intrinsic difficulty of the underlying problem is closely mirrored by the inherent complexity of their associated proofs. We will then see how Cook and Reckhow's seminal paper from 1979 gave rise to proof complexity as a discipline, and revealed the fundamental connection between proof complexity and the P versus NP problem. I will highlight some of the major discoveries that have been made, with a focus on new results and connections that have been made with other areas in the last 10 years. These new connections have enabled breakthrough discoveries in several areas, including: algorithms for solving distributional learning problems, connections with total NP search problems (TFNP), approximation algorithms, circuit lower bounds, cryptography (secret sharing schemes), and error correcting codes.

Harvard Theory of Computing Seminar (Spring ‘24)

About: The Harvard Theory of Computing seminar is the primary seminar series of the Harvard Theory of Computation Group.

Time and place: The meetings happen roughly every other **Wednesday** from **3:45-5pm.** (Typically there are refreshments before the seminar from 3:15-3:45, outside the seminar room. The talks end at 4:45 with 4:45-5:00pm reserved for Q&A.) The meetings in Spring 2024 will be held in the Harvard SEC, room 3.301-3.303.

Mailing list: To join the theory seminar mailing list, please email Allison at achoat@g.harvard.edu.

Calendar: Here’s a link to add our google calendar, or you can view it here.

Organizers: Sumaira Ahammed, Anurag Anshu, Allison Choat, Rebekah Keely, Aaron (Louie) Putterman, Jake Ruotolo, Madhu Sudan.

January 24, 2024: Postdoc Day

Time: 3:45 - 5:00 p.m.

Location: SEC 3.301-3.303

Title: **Postdocs Day**

Speakers: **Vikrant Singhal**, **Tomer Ezra**, **Eran Malach**

Abstracts:

February 7, 2024: Pravesh Kothari

Time: 3:45 - 5:00 p.m.

Location: SEC 3.301-3.303

Title: **Efficiently Finding Planted Cliques in Semirandom Graphs**

Speaker: **Pravesh Kothari**

Abstract: In this talk, I will present a new polynomial time algorithm for recovering planted cliques in the semi-random graph model of Feige and Kilian from 2001. In the Feige-Kilian model, a graph G on vertex set V is chosen by a combination of benign random and adaptive "adversarial" choices and can be thought of as a "robust" counterpart of the well-studied "fully-random" planted clique problem:
1) choose a planted clique on a set of k vertices S,
2) "benign random choice": include each edge between S and V\S independently at random,
3) "adaptive worst-case choice": deleting any set of edges between S and V\S and adding any set of edges with both endpoints in V\S.

The previous best algorithms for this model succeed if the planted clique has a size of at least n^{2/3} in a graph with n vertices. Our algorithms work for planted-clique sizes approaching n^{1/2} -- the information-theoretic threshold in the semi-random model and a conjectured computational threshold even in the easier fully-random model. This result comes close to resolving an open question of Feige (2019) and Steinhardt (2017). Our algorithm is based on high constant degree sum-of-squares relaxation of the clique SDP and the analysis relies on a new connection between finding planted cliques and efficient certificates of upper bounds on clique number of unbalanced bicliques in bipartite random graphs.

Based on joint work with Rares Buhai and David Steurer.

February 21, 2024: Jane Lange

Time: 3:45 - 5:00 p.m.

Location: SEC 3.301-3.303

Title: **Agnostic Proper Learning of Monotone Functions: Beyond the Black-Box Correction Barrier**

Speaker: **Jane Lange**

Abstract: We give the first agnostic, efficient, proper learning algorithm for monotone Boolean functions. Given 2^Õ(sqrt(n)/ε) uniformly random examples of an unknown function f:{±1}^n→{±1}, our algorithm outputs a hypothesis g:{±1}^n→{±1} that is monotone and (opt+ε)-close to f, where opt is the distance from f to the closest monotone function. The running time of the algorithm (and consequently the size and evaluation time of the hypothesis) is also 2^Õ(sqrt(n)/ε), nearly matching the lower bound of Blais et al (RANDOM '15). We also give an algorithm for estimating up to additive error ε the distance of an unknown function f to monotone using a run-time of 2^Õ(sqrt(n)/ε). Previously, for both of these problems, sample-efficient algorithms were known, but these algorithms were not run-time efficient. Our work thus closes this gap in our knowledge between the run-time and sample complexity. This work builds upon the improper learning algorithm of Bshouty and Tamon (JACM '96) and the proper semiagnostic learning algorithm of Lange, Rubinfeld, and Vasilyan (FOCS '22), which obtains a non-monotone Boolean-valued hypothesis, then "corrects" it to monotone using query-efficient local computation algorithms on graphs. This black-box correction approach can achieve no error better than 2opt+ε information-theoretically; we bypass this barrier by a) augmenting the improper learner with a convex optimization step, and b) learning and correcting a real-valued function before rounding its values to Boolean. Our real-valued correction algorithm solves the "poset sorting'' problem of [LRV22] for functions over general posets with non-Boolean labels.

March 6, 2024: Raghu Meka

Time: 3:45 - 5:00 pm

Location: SEC 3.301-3.303

Title: **New Frontiers in Structure vs Randomness with Applications to Combinatorics, Complexity, Algorithms**

Speaker: **Raghu Meka**

Abstract: In 1936, Erdos and Hajnal asked the following: Suppose you have a set S of integers from {1,2,..., N} that contains at least N / C elements. Then, for large enough N, must S have three equally spaced numbers (i.e., a 3-term arithmetic progression)? In 1946, Behrend showed that C can be at most exp(sqrt(log N)). Since then, the problem has been a cornerstone of the area of additive combinatorics, with the best bound being C = (log N)^(1+c) for some constant c > 0. Recent work obtained an exponential improvement showing that C can be as big as exp((log N)^0.09), thus getting closer to Behrend's construction.

In this talk, I will describe this result and the main ingredient, a new variant of the "structure vs. randomness" paradigm. The latter is an old technique with many applications in complexity theory, algorithm design, and number theory, and the new variant can potentially lead to further progress. I will highlight two such applications:
1. Communication complexity: explicit separations between randomized and deterministic multi-party protocols.
2. Algorithm design: fast combinatorial algorithms for Boolean matrix multiplication, detecting triangles in graphs.

Based on works with Amir Abboud, Nick Fischer, Zander Kelley, Shachar Lovett.

April 3, 2024: Yael Tauman Kalai

Time: 3:45 - 5:00 pm

Location: SEC 3.301-3.303

Title: **Classical Commitments to Quantum States and Application to Succinct Verification of QMA**

Speaker: **Yael Tauman Kalai**

Abstract: We define the notion of a classical commitment scheme to quantum states, which allows a quantum prover to compute a classical commitment to a quantum state, and later open each qubit of the state in either the standard or the Hadamard basis. Our notion is a strengthening of the measurement protocol from Mahadev (STOC 2018). We construct such a commitment scheme from the post-quantum Learning With Errors (LWE) assumption. Our scheme is succinct in the sense that the running time of the verifier in the commitment phase depends only on the security parameter (independent of the size of the committed state), and its running time in the opening phase grows only with the number of qubits that are being opened (and the security parameter). As a corollary we obtain a classical succinct argument system for QMA under the post-quantum LWE assumption. Previously, this was only known assuming post-quantum secure indistinguishability obfuscation. As an additional corollary we obtain a generic way of converting any X/Z quantum PCP into a succinct argument system under the quantum hardness of LWE. This is joint work with Sam Gunn, Anand Natarajan and Agi Villanyi.

April 17, 2024: Mohsen Ghaffari

Time: 3:45 - 5:00 pm

Location: SEC 3.301-3.303

Title: **Parallel Derandomization for Chernoff-like Concentrations**

Speaker: **Mohsen Ghaffari**

Abstract:TBD

April 24, 2024: Sanjeev Khanna

Time: 3:45 - 5:00 pm

Location: SEC 3.301-3.303

Title: **TBD**

Speaker: **Sanjeev Khanna**

Abstract:TBD