Complexity Theory Update Meeting

IMSc, Chennai. 21-23 January 2026.

Long talks

1. Jayalal Sarma M N, IITM
   Abstract: The talk will be based on recent developments concerning the Range-Avoid search problem.

2. Meena Mahajan, IMSc
   Title: Proof complexity and Search problems in TFNP and TFPH
   Abstract: The talk will be loosely based on the papers:
   Provably Total Functions in the Polynomial Hierarchy,
   by Noah Fleming, Deniz Imrek, Christophe Marciot (CCC 2025),
   `https://doi.org/10.4230/LIPIcs.CCC.2025.28`, and
   Separations in Proof Complexity and TFNP,
   by Mika Göös, Alexandros Hollender, Siddhartha Jain, Gilbert Maystre, William Pires, Robert Robere, Ran Tao (FOCS 2022), J. ACM 71(4): 26:1-26:45 (2024).
   `https://doi.org/10.1145/3663758`

3. Swagato Sanyal, Sheffield Univ
   The talk will be based on the paper:
   Direct Product Theorems for Randomized Query Complexity,
   by Shalev Ben-David, Eric Blais (FOCS 2025).
   `https://doi.org/10.48550/arXiv.2512.08268`

4. Ramprasad Saptharishi, TIFR
   The talk will be based on the paper:
   Collapsing Catalytic Classes,
   by Michal Koucký, Ian Mertz, Edward Pyne, Sasha Sami (FOCS 2025).
   `https://doi.org/10.48550/arXiv.2504.08444`

5. Rohit Gurjar, IIT Bombay
   The talk will be based on the paper:
   Bipartite Matching is in Catalytic Logspace,
   by Aryan Agarwala, Ian Mertz (FOCS 2025).
   `https://doi.org/10.48550/arXiv.2504.09991`

6. Sumanta Ghosh, ISI Kolkata
   The talk will be based on the paper:
   Improved PIR Schemes using Matching Vectors and Derivatives
   by Fatemeh Ghasemi, Swastik Kopparty, Madhu Sudan (STOC 2025)
   `https://doi.org/10.48550/arXiv.2411.11611`

7. Suryajith Chillara, IIIT Hyderabad.
   The talk will be based on the paper:
   Permanental rank versus determinantal rank of random matrices over finite fields,

by Fatemeh Ghasemi, Gal Gross, Swastik Kopparty (RANDOM 2025).
`https://doi.org/10.48550/arXiv.2512.03221`

8. Anil Shukla, IIT Ropar.
   Title: Proof Complexity of MaxSAT.
   Abstract: Given a CNF formula F, MaxSAT problem asks to find the maximum number of clauses that can be simultaneously satisfiable by any complete assignment to the variables. In proof complexity, we study the hardness in finding the certificate of MaxSAT. That is, if the MaxSAT answer for a CNF formula is k, then proof complexity studies proof systems for efficiently proving the fact that the answer is indeed k. Many MaxSAT proof systems have been recently introduced in the literature. This talk covers a few of them, and is based on the following papers:
   - Resolution for Max-SAT, by María Luisa Bonet, Jordi Levy and Felip Manyà in Artificial Intelligence, Volume 171, Issue 8-9, 2007.
     `https://doi.org/10.1016/j.artint.2007.03.001`
   - MaxSAT Resolution with Inclusion Redundancy, by Ilario Bonacina, Maria Luisa Bonet and Massimo Lauria (SAT 2024).
     `https://doi.org/10.4230/LIPIcs.SAT.2024.7`
   - Redundancy Rules for MaxSAT, by Ilario Bonacina, Maria Luisa Bonet, Sam Buss and Massimo Lauria (SAT 2025).
     `https://doi.org/10.4230/LIPIcs.SAT.2025.7`
   - Polynomial calculus for optimization, by Ilario Bonacina, Maria Luisa Bonet and Jordi Levy in Artificial Intelligence, Volume 337, 2024.
     `https://doi.org/10.1016/j.artint.2024.104208`
9. Jaikumar Radhakrishnan, ICTS Bengaluru
   The talk will be based on the paper:
   Quantum Search with In-Place Queries,
   by Blake Holman, Ronak Ramachandran, Justin Yirka (TQC 2025).
   `https://doi.org/10.48550/arXiv.2504.03620`
10. Nitin Saurabh, IIT Hyderabad
    The talk will be based on the paper:
    A Near-Optimal Polynomial Distance Lemma Over Boolean Slices,
    by Prashanth Amireddy, Amik Raj Behera, Srikanth Srinivasan, Madhu Sudan (ICALP 2025).
    `https://doi.org/10.48550/arXiv.2507.03193`

Abstracts: Short talks

1. Speaker: Neha Kuntewar, IIT Madras
   Title: Range Avoidance via Turan-type Bounds in Hypergraphs
   RANDOM 2025, `https://doi.org/10.4230/LIPIcs.APPROX/RANDOM.2025.62`
   Abstract: The range avoidance problem (denoted by AVOID) is as follows: given a Boolean circuit C computing a function from $\{0,1\}^n$ to $\{0,1\}^m$ with $m > n$, find a string that is outside the range of the circuit. The problem has received much attention in complexity

theory during the last couple of years due to its close connections with circuit lower bounds and several important explicit construction problems that are central open questions in the area. Deterministic polynomial time algorithms (even with access to NP oracles) solving this problem are known to imply explicit constructions of various pseudorandom objects like hard Boolean functions, linear codes, PRGs, etc. Given the central nature of the problem, it is natural to study the complexity for special circuit classes. Deterministic polynomial time algorithms are known for AVOID when $m > n$, when each output function depends only on at most two input variables, and for the case of dependency on three bits when $m > n^2/\log n$. On the other hand, for the dependency three case, if we design an algorithm for $m = n + O(n^{2/3})$, it implies explicit construction of rigid matrices, which is another long-standing open problem in the area. In fact, algorithms for solving range avoidance, even when each output function depends on four input variables, imply new circuit lower bounds.

We propose a new approach to solving the range avoidance problem via hypergraphs. We formulate the problem in terms of Turan-type problems in hypergraphs of the following kind: for a fixed k-uniform hypergraph H', what is the maximum number of edges that can exist in a k-uniform hypergraph H which does not have a sub-hypergraph isomorphic to H'? We apply this framework to design new algorithms for special cases of AVOID problem. Observing that even monotone cases of the problem are as hard as general, we use our framework to design polynomial time algorithms for monotone-AVOID when $m > n$, when each function depends on three input variables. We also present generalisations of our algorithm, using appropriately designed hypergraph structures, to solve other special cases of AVOID that might be of independent interest.

2. Speaker: Pratik Shastri, IMSc
   Title: Lower Bounds for Noncommutative Circuits with Low Syntactic Degree
   `https://eccc.weizmann.ac.il/report/2025/133/`, to appear in ITCS 2026.
   Abstract:

3. Speaker: Rohit Narayanan, IIT Gandhinagar
   Title: Monotone Bounded Depth Formula Complexity of Graph Homomorphism Polynomials
   `https://doi.org/10.48550/arXiv.2511.03388`
   Abstract:

4. Speaker: Varun Ramanathan, TIFR
   Title: Deterministic factorization of constant-depth algebraic circuits in subexponential time
   `https://arxiv.org/abs/2504.08063` (FOCS 25), and `https://arxiv.org/abs/2506.23214`
   Abstract:

5. Speaker: Shantanu Rai, TIFR
   Title: Factors of succinctly computable polynomials
   `https://doi.org/10.48550/arXiv.2506.23214`
   Abstract: In this talk, we will see that algebraic formulas and constant-depth circuits are closed under taking factors. In other words, we show that if a multivariate polynomial over a field of characteristic zero has a small constant-depth circuit or formula, then all its factors can be computed by small constant-depth circuits or formulas respectively. The result turns out to be an elementary consequence of a fundamental and surprising result of Furstenberg from the 1960s, which gives a non-iterative description of the power series roots of a bivariate

polynomial. Combined with standard structural ideas in algebraic complexity, we observe that this theorem yields the desired closure results.

Based on joint work with Somnath Bhattacharjee, Mrinal Kumar, Varun Ramanathan, Ramprasad Saptharishi and Shubhangi Saraf.

6. Speaker: Swarnalipa Datta, ISI Kolkata
   Title: Spectral Norm, Economical Sieve, and Linear Invariance Testing of Boolean functions (to appear in STACS 2026)
   Abstract: Given Boolean functions $f, g : \mathbb{F}_2^n \to \{-1, +1\}$, we say they are linearly isomorphic if there exists $A \in GL_n(\mathbb{F}_2)$ such that $f(x) = g(Ax)$ for all x. We study this problem in the tolerant property testing framework under the known–unknown model, where g is given explicitly, and f is accessible only via oracle queries, meaning the algorithm may adaptively request the value of $f(x)$ for inputs $x \in \mathbb{F}_2^n$ of its choice. Given parameters $\epsilon \geq 0$ and $\omega > 0$, the goal is to distinguish whether there exists $A \in GL_n(\mathbb{F}_2)$ such that the normalized Hamming distance between f and $g(Ax)$ is at most ϵ, or whether for every $A \in GL_n(\mathbb{F}_2)$ the distance is at least $\epsilon + \omega$. Our main result is a tolerant tester making $O(m^4/\omega^4)$ queries to f, where m is an upper bound on the spectral norm of g, improving the previous $O(m^{24}/\omega^{24})$ bound of Wimmer and Yoshida. We complement this with a nearly matching lower bound of $\Omega(m^2)$ for constant ω (for example, $\omega = 1/4$), improving the prior $\Omega(\log m)$ lower bound of Grigorescu, Wimmer, and Xie. A key technical ingredient on the algorithmic side is a query-efficient local list corrector. For the lower bound, we give a reduction from communication complexity using a novel subclass of Maiorana–McFarland functions from symmetric-key cryptography.

7. Speaker: Manmatha Roy, ISI Kolkata
   Title: Price of Parsimony: Complexity of Fourier Sparsity Testing
   (NuerIPS 2025 poster) `https://openreview.net/forum?id=7bCPXHq8xV`
   Abstract: A function $f : \mathbb{F}_2^n \longrightarrow \mathbb{R}$ is said to be s-Fourier sparse if its Fourier expansion contains at most s nonzero coefficients. In general, the existence of a sparse representation in the Fourier basis serves as a key enabler for the design of efficient learning algorithms. However, most existing techniques assume prior knowledge of the function's Fourier sparsity, with algorithmic parameters carefully tuned to this value. This motivates the following decision problem: given $s > 0$, determine whether a function is s-Fourier sparse.

   In this work, we study the problem of tolerant testing of Fourier Sparsity for real-valued functions over $\mathbb{F}_2^n$, accessed via oracle queries. The goal is to decide whether a given function is close to being s-Fourier sparse or far from every s-Fourier sparse function. Our algorithm provides an estimator that, given oracle access to the function, estimates its distance to the nearest s-Fourier sparse function with query complexity $\tilde{O}(s)$, for constant accuracy and confidence parameters.

   A key structural ingredient in our analysis is a new spectral concentration result for real-valued functions over $\mathbb{F}_2^n$ when restricted to small-dimensional random affine subspaces. We further complement our upper bound with a matching lower bound of $\Omega(s)$, establishing that our tester is optimal up to logarithmic factors. The lower bound exploits spectral properties of a class of cryptographically hard functions, namely, the Maiorana–McFarland family, in a novel way.