P vs NP and Complexity Lower Bounds
Abstracts

Eric Allender (Rutgers University, retired)

Title: Progress on the Power of Recognizing Randomness

Abstract: To mark my retirement in 2023, I publicized a list of open questions and offered a bounty for the solution of any of them. You may think of it as a Milli-Millennium Prize; the bounty was 0.1% of the amount offered for the solution of a Millennium Prize Problem. One bounty has been claimed thus far. I'll give an update on the solved problem, and I'll also discuss some related open questions, including a program for possibly separating NP from NL. Trigger warning: This talk will involve Kolmogorov Complexity!

Shuichi Hirahara (National Institute of Informatics)

Title: NP-hardness of the Minimum Circuit Size Problem from Well-Studied Assumptions

Abstract: Whether the Minimum Circuit Size Problem (MCSP) is NP-hard or not is a long-standing open question. Indeed, Levin delayed the publication of his fundamental work on the theory of NP-completeness because he hoped to prove NP-completeness of MCSP.

In this talk, we present the first plausible assumptions under which MCSP is NP-hard. Specifically, we prove that MCSP is NP-hard under deterministic quasi-polynomial-time nonadaptive reductions, assuming (1) subexponentially-secure non-interactive witness indistinguishable proof systems for SAT exist, (2) coNP requires subexponential-size non-deterministic circuits, and (3) P^{NP}/poly does not have circuits of size o(2^n/n).

The talk is based on the joint work with Rahul Ilango.

Christian Ikenmeyer (University of Warwick)

Title: Algebraic metacomplexity and representation theory

Abstract: In the algebraic metacomplexity framework we prove that the decomposition of metapolynomials into their isotypic components can be implemented efficiently, namely with only a quasipolynomial blowup in the circuit size. We use this to resolve an open question posed by Grochow, Kumar, Saks & Saraf (2017). Our result means that many existing algebraic complexity lower bound proofs can be efficiently converted into isotypic lower bound proofs via highest weight metapolynomials, a notion studied in geometric complexity theory. In the context of algebraic natural proofs, it means that without loss of generality algebraic natural proofs can be assumed to be isotypic. Our proof is built on the Poincaré-Birkhoff-Witt theorem for Lie algebras and on classical Gelfand-Tsetlin theory. This is joint work with Maxim van den Berg, Pranjal Dutta, Fulvio Gesmundo, and Vladimir Lysikov.

Valentine Kabanets (Simon Fraser University)

Title: Kolmogorov's approach to P vs NP

Abstract: Shannon's Information Theory allows one to quantify the amount of information (entropy) in a given probability distribution over finite binary strings. In contrast, Kolmogorov defined a notion of algorithmic information present in a given individual finite binary string x, called Kolmogorov complexity of x, and denoted K(x). More precisely, the Kolmogorov complexity of a string x, given a string y, denoted K(x|y), is the size of a smallest computer program that, on input y, prints the string x. Similar to the classical chain rule for Shannon's entropy, there is a chain rule for Kolmogorov complexity (proved by Kolmogorov and Levin in the 1970s), which says, for any two finite binary strings x and y, that K(x,y) is approximately equal to K(x) + K(y|x).

Scaling down to time-bounded Kolmogorov complexity, one can define t-time Kolmogorov complexity of x, given y, denoted K^t (x|y), as the size of a smallest computer program that, on input y, prints x within t computation steps. Does the chain rule hold for time-bounded Kolmogorov complexity? For the upper bound, yes. It is easy to show that K^{2t} (x,y) \leq K^t (x) + K^t(y|x). It is open if the lower bound holds, i.e., whether for some fixed polynomial p, we have K^t (x,y) \geq K^{p(t)} (x) + K^{p(t)} (y|x). The proof of the chain rule for the classical, time-unbounded, Kolmogorov complexity doesn't work (is not time-efficient) in the time-bounded setting of K^t. Kolmogorov is believed to have conjectured that the chain rule for K^t is false. Moreover, Kolmogorov thought that disproving the chain rule for K^t was a good approach towards proving that P is different from NP!

We make a step towards validating Kolmogorov's intuition. We show that a (certain general form of the) chain rule for K^t holds if and only if there is an efficient algorithm for a certain version of approximately computing K^t (x) on given input strings x (and generic efficient derandomization of polynomial-time randomized algorithms is possible). Although we don't know if our version of approximating K^t is actually NP-complete, there is some evidence in the literature that it may well be NP-complete. If so, our result would imply that the chain rule for K^t holds iff P=NP, as predicted by Kolmogorov!

Joint work with Antonina Kolokolova.

Michal Koucky (Charles University)

Title: Recent developments in catalytic computing

Abstracts: Catalytic computing is a model of space bounded computation where in addition to the usual work tape we have an access to a much larger memory that is already full of unrelated data. The data must be preserved upon termination of the computation but during the computation if can be modified arbitrarily. This model was defined in 2014 by Buhrman et al. and serves as a platform for understanding the role of space in computation. Buhrman et al. established that the extra memory provides a surprising computational power. More recently, those ideas inspired surprising algorithms for some well studied problems: the almost log-space algorithm for Tree Evaluation of Cook and Mertz (2024) and in turn the space efficient simulation of time bounded computation of Williams (2025). The model of catalytic computing became a focus of several recent studies that shed light on the role of randomness and non-determinism in this space bounded model of computation. In this talk I will discus some of those recent developments.

Ryan O'Donnell (Carnegie Mellon University)

Title: No exponential quantum speedup for SIS∞ anymore

Abstract: In 2021, Chen, Liu, and Zhandry presented an efficient quantum algorithm for the average-case ℓ∞-Short Integer Solution (SIS∞) problem, in a parameter range outside the normal range of cryptographic interest, but still with no known efficient classical algorithm. This was particularly exciting since SIS∞ is a simple problem without structure, and their algorithmic techniques were different from those used in prior exponential quantum speedups. We present efficient classical algorithms for all of the SIS∞ and (more general) Constrained Integer Solution problems studied in their paper, showing there is no exponential quantum speedup anymore.

Igor Oliveira (University of Warwick)

Title: Meta-Mathematics of P vs NP

Abstract: We survey results on the formalization and independence of mathematical statements related to major open problems in computational complexity theory. Our primary focus is on recent findings concerning the (un)provability of complexity bounds within theories of bounded

arithmetic. This includes the techniques employed and related open problems, such as the (non)existence of a feasible proof that P = NP.

Toniann Pitassi (University of Toronto)

Title: Data Structure Lower Bounds, Local PRGs, and Range Avoidance: New connections and lower bounds

Pavel Pudlák (Czech Academy of Sciences)

Title: NP vs. coNP and hard tautologies

Abstract: $NP\neq coNP$ is equivalent to the nonexistence of a propositional proof system in which all tautologies have proofs whose lengths are bounded by a polynomial. A plausible stronger hypothesis is that there is no optimal proof system, which means that for every proof system P there exists a proof system Q and a sequence of tautologies $A_1,A_2,...$ such that A_is have proofs of polynomial length in Q, but no such proofs exist in P. An old result of Krajíček and Pudlák states that sequences of such tautologies can be constructed using finite consistency statements. Stronger and more specific conjectures have also been proposed. I will talk about some recent developments in this area.

Ben Rossman (Duke University)

Title: The Mystery of Negations

Abstract: Monotone models let us prove strong lower bounds—famously, we know mon-P ≠ mon-NP and have exponential monotone circuit lower bounds for natural problems like Clique. Yet these successes tell us little about P vs NP. The culprit is a mismatch between our techniques and the distributions they target: the strongest monotone lower bounds work under *weight-separable* distributions, where 0-instances have larger Hamming weight than 1-instances. By contrast, under any *slice* distribution (all inputs of the same weight), Berkowitz ’82 showed that monotone and non-monotone complexity coincide up to a polynomial factor for basic measures such as circuit size and formula size.

This talk—an homage to the eponymous chapter in Jukna’s *Boolean Function Complexity*—explores the open question of whether Berkowitz’s theorem extends to *product* distributions (independent input bits; for graphs, G(n,p)) and why this matters for our understanding of P vs NP. I will outline techniques that yield monotone formula lower bounds for the k-Cycle problem in G(n,1/n) and explain a technical barrier that currently prevents analogous bounds for k-Clique (and Planted Clique) in G(n,p).

Shubhangi Saraf (University of Toronto)

Title: The complexity of factors of polynomials

Abstract: I will talk about a recent result showing that algebraic formulas and constant-depth circuits are closed under taking factors. In other words, the complexity of factors of polynomials computable by algebraic formulas or constant depth algebraic circuits is not much more than the complexity of the original polynomial itself.

This 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.

We will see applications of this result to deterministic algorithms for factoring, hardness/randomness

tradeoffs, as well as GCD computation of polynomials.

This talk is based on joint works with Somnath Bhattacharjee, Mrinal Kumar, Shanthanu Rai, Varun Ramanathan and Ramprasad Saptharishi.

Iddo Tzameret (Imperial College London)

Title: The algebraic approach to proof complexity

Abstract: Proof complexity is one of the central approaches to the fundamental hardness problems in complexity theory. In recent years, significant efforts have been made to bridge the gap between algebraic and proof complexity through a relatively transparent reduction from algebraic circuit-size lower bounds to proof-size lower bounds. In this talk, I will discuss state-of-the-art lower bounds in proof complexity that leverage the algebraic circuit-based approach, establishing it as a new tool that also draws on ideas from existing techniques---such as feasible interpolation, random restrictions, width-size tradeoffs, and lifting. I will also highlight some imminent open problems and potential challenges in this direction.

Leslie Valiant (Harvard University)

Title: A combinatorial quantity in counting complexity

Ryan Williams (MIT)

Title: Simulating Time With Square-Root Space

Abstract: We show that for all functions t(n) >= n, every multitape Turing machine running in time t can be simulated using only O(sqrt{t log t}) space. This is a substantial improvement over Hopcroft, Paul, and Valiant's simulation of time t in O(t/log t) space from 50 years ago [FOCS 1975, JACM 1977]. Among other results, our simulation implies that bounded fan-in circuits of size s can be evaluated on any input in only sqrt{s} poly(log s) space, and that there are explicit problems solvable in O(n) space which require at least n^2/poly(log n) time on every multitape Turing machine, thereby making a little progress on the P versus PSPACE problem. Our simulation reduces the problem of simulating time-bounded multitape Turing machines to a series of implicitly-defined Tree Evaluation instances with nice parameters, leveraging the remarkable space-efficient algorithm for Tree Evaluation recently found by Cook and Mertz [STOC 2024].

Henry Yuen (Columbia University)

Title: The power and limits of constant-time quantum computation

David Zuckerman (University of Texas at Austin)

Title: Randomness extractors, pseudorandom generators, and lower bounds

Abstract: We will survey connections between randomness extractors, pseudorandom generators, and lower bounds.