Complexity Theory
23-26 July 2018
University of Oxford

Abstracts

Josh Alman (MIT)

Title: Limits on All Known (and Some Unknown) Approaches to Matrix Multiplication

Abstract: We study the known techniques for designing Matrix Multiplication (MM) algorithms. The two main approaches are the Laser method of Strassen, and the Group theoretic approach of Cohn and Umans. We define a generalization based on zeroing outs which subsumes these two approaches, which we call the Solar method, and an even more general method based on monomial degenerations, which we call the Galactic method.

We then design a suite of techniques for proving lower bounds on the value of omega, the exponent of MM, which can be achieved by algorithms using many tensors T and the Galactic method. Some of our techniques exploit 'local' properties of T, like finding a sub-tensor of T which is so 'weak' that T itself couldn't be used to achieve a good bound on omega, while others exploit 'global' properties, like T being a monomial degeneration of the structural tensor of a group algebra. Our main result is that there is a universal constant $c > 2$ such that a large class of tensors generalizing the Coppersmith-Winograd tensor CW_q cannot be used within the Galactic method to show a bound on omega better than c, for any q.

In this talk, I'll begin by giving a high-level overview of the algorithmic techniques involved in the best known algorithms for MM, and then I will tell you about our lower bounds. No prior knowledge of MM algorithms will be assumed.

Joint work with Virginia Vassilevska Williams to appear in FOCS 2018.

Marco Carmosino (UCSD)

Title: Hardness Amplification for Non-Commutative Arithmetic Circuits

Abstract: We show that proving mildly super-linear lower bounds on non-commutative arithmetic circuits implies exponential lower bounds on non-commutative circuits. That is, non-commutative circuit complexity is a threshold phenomenon: an apparently weak lower bound actually suffices to show the strongest lower bounds we could desire. This is part of a recent line of inquiry into why arithmetic circuit complexity, despite being a heavily restricted version of Boolean complexity, still cannot prove super-linear lower bounds on general devices. One can view our work as positive news (it suffices to prove weak lower bounds to get strong ones) or negative news (it is as hard to prove weak lower bounds as it is to prove strong ones).

Arkadev Chattopadhyay (TIFR)

Title: A Short List of Equalities Induces Large Sign Rank

Abstract: We exhibit a natural function F on n variables, that can be computed by just a linear sized decision list of 'Equalities', but whose sign rank is exp(n^{1/4}). This yields the following two new unconditional complexity class separations.

(a) Boolean circuit complexity: The function F can be computed by linear sized depth-two threshold formulas when the weights of the threshold gates are unrestricted (THR of THR), but any THR of MAJ circuit (the weights of the bottom threshold gates are polynomially bounded in n) computing F requires size exp(n^{1/4}). This provides the first separation between the boolean circuit complexity classes THR of MAJ and THR of THR. While Amano and Maruoka (2005) and Hansen and Podolskii (2010) emphasized that super-polynomial separations between the two classes remained a basic open problem since the seminal work of Goldmann, Hastad and Razborov (1992), our separation is in fact exponential.

(b) Communication complexity: The function F (under the natural partition of the inputs) lies in the communication complexity class P^MA. Since F has large sign rank, this implies P^MA is not contained in UPP, strongly resolving a recent open problem posed by Goos, Pitassi and Watson (2016). Further, our work highlights for the first time the class 'decision lists of exact thresholds' as a common frontier for making progress on longstanding open problems in Threshold circuits and communication complexity.

This is joint work with Nikhil Mande.

Lance Fortnow (Georgia Tech)

Title: Some Observations on the Raz-Tal Oracle Separating BQP from PH

Abstract: Ron Raz and Avishay Tal recently gave a relativized world where BQP is not contained in the polynomial-time hierarchy. We explore some extensions to that result.

We give simple proofs of the following:
- There exists a relativized world where P = NP different from BQP.
- Under that original Raz-Tal oracle, PH is infinite.

We explore several open oracle questions:
- P = NP and BQP = EXP. The challenge here is how to use Raz-Tal to encode into BQP.
- NP in BQP and Sigma-2 not contained in BQP (suggested by Scott Aaronson). This also runs into issues of encoding.
- The weaker statement NP^BQP not contained in BQP^NP. The similar statement for BPP has a relativizable proof of inclusion. We use Raz-Tal to give a candidate language for separation.

Joshua Grochow (University of Colorado)

Title: Complexity in Ideals of Polynomials

Abstract: Given a family of ideals A_n in polynomial rings in a growing number of variables, what can be said about the complexity of families of polynomials f_n where f_n in A_n for all n? This abstract question generalizes a common theme that occurs in algebraic natural proofs for algebraic circuit lower bounds, geometric complexity theory, and algebraic proof complexity. The purpose of

this talk is to discuss how this question arises in these areas. I will also talk about what little is known about this general question, and raise many more questions than I answer.

Yuval Ishai (Technion)

Title: Cryptography and Complexity Theory: Recent Interactions

Abstract: There is a rich history of fruitful interactions between cryptography and complexity theory. The talk will survey some recent interactions and questions in complexity theory that they motivate.

Valentine Kabanets (Simon Fraser University)

Title: The Power of Natural Properties as Oracles

Abstract: We study the power of randomized complexity classes that are given oracle access to a natural property of Razborov and Rudich (*JCSS*, 1997) or its special case, the Minimal Circuit Size Problem (MCSP). We show that in a number of complexity-theoretic results that use the SAT oracle, one can use the MCSP oracle instead. For example, we show that $\mathsf{ZPEXP}^{\mathsf{MCSP}} \not\subseteq \mathsf{P/poly}$, which should be contrasted with the previously known circuit lower bound $\mathsf{ZPEXP}^{\mathsf{NP}} \not\subseteq \mathsf{P/poly}$. We also show that, assuming the existence of Indistinguishability Obfuscators (IO), SAT and MCSP are equivalent in the sense that one has a ZPP algorithm if and only the other one does. We interpret our results as providing some evidence that MCSP may be NP-hard under randomized polynomial-time reductions.

Joint work with Russell Impagliazzo and Ilya Volkovich.

Daniel Kane (UCSD)

Title: Fooling Fourier Shapes

Abstract: We present recent work providing a nearly optimal, explicit pseudorandom generator against linear threshold functions. The basic idea is to fool the Fourier transform of the corresponding linear forms rather than the threshold function. This generalizes further to fooling a class of functions that we call "Fourier shapes". We discuss the generator and some of its applications.

Neeraj Kayal (Microsoft Research)

Title: Proper Learning Algorithms from Lower Bounds for Arithmetic Circuits

Abstract: The proper learning problem for a circuit class C is the following: given oracle access to a function f(x), find a minimal sized circuit in C computing f. This problem, also known as the reconstruction problem for C, is known to be hard even for very simple classes of circuits in both the Boolean and Arithmetic settings.

For arithmetic circuit classes C where we have some understanding in the form of lower bounds, can we exploit this understanding to design proper learning algorithms for C? We give such proper learning algorithms for some subclasses of arithmetic circuits. However, our algorithms are not worst

case but they succeed provided that the unknown circuit T satisfies certain nondegeneracy conditions that are satisfied for example if T is not very large and the field constants in T are chosen randomly.

Based on joint work with Chandan Saha.

Antonina Kolokolova (Memorial University of Newfoundland)

Title: Does Looking Inside a Circuit Help?

Abstract: Celebrated Rice's theorem in computability theory states that any non-trivial semantic property of Turing machines is undecidable. That is, to check if the language accepted by a Turing machine satisfies some property, essentially the only thing one can do is to run the machine.

Is there is a similar phenomenon in the complexity theory world? Following a conjecture posed by Barak et al. in their paper on impossibility of obfuscation, we ask whether working with a description of a Boolean circuit gives any computational advantage for deciding properties of a Boolean function, as opposed to a black-box (oracle) access.

We show that for read-once branching programs there is indeed such an advantage. For general Boolean circuits, we make a step towards resolving this question by showing that if properties of a certain type are easier to decide given a circuit description then there is a non-trivial algorithm for the CircuitSAT problem.

Joint work with Russell Impagliazzo, Valentine Kabanets, Pierre McKenzie and Shadab Romani.

Michael Koucky (Charles University)

Title: Lower Bounds for Combinatorial Algorithms for Boolean Matrix Multiplication

Abstract: We propose models of combinatorial algorithms for the Boolean Matrix Multiplication (BMM), and prove lower bounds on computing BMM in these models. First, we give a relatively relaxed combinatorial model which is an extension of the model by Angluin (1976), and we prove that the time required by any algorithm for the BMM is at least $\Omega(n^3 / 2^{O(\sqrt{ \log n })})$. Subsequently, we propose a more general model capable of simulating the "Four Russians Algorithm". We prove a lower bound of $\Omega(n^{7/3} / 2^{O(\sqrt{ \log n })})$ for the BMM under this model. We use a special class of graphs, called (r,t)-graphs, originally discovered by Rusza and Szemeredi (1978), along with randomization, to construct matrices that are hard instances for our combinatorial models.

Joint work with Debarati Das and Mike Saks.

Jan Krajicek (Charles University)

Title: The Nature of Proof Complexity

Abstract: I shall discuss the structure of contemporary proof complexity, how it splits into a few levels, and how the fundamental problems fit in. I shall also speculate about possible approaches to the problems, illustrating this by a few statements.

Andrea Lincoln (MIT)

Title: Tight Hardness for Shortest Cycles and Paths in Sparse Graphs

Abstract: Fine-grained reductions have established equivalences between many core problems with $\tilde{O}(n^3)$-time algorithms on n-node weighted graphs, such as Shortest Cycle, All-Pairs Shortest Paths (APSP), Radius, Replacement Paths, Second Shortest Paths, and so on. These problems also have $\tilde{O}(mn)$-time algorithms on m-edge n-node weighted graphs, and such algorithms have wider applicability. Are these mn bounds optimal when $m \ll n^2$?

Starting from the hypothesis that the minimum weight $(2\ell+1)$-Clique problem in edge weighted graphs requires $n^{2\ell+1-o(1)}$ time, we prove that for all sparsities of the form $m = \Theta(n^{1+1/\ell})$, there is no $O(n^2 + mn^{1-\epsilon})$ time algorithm for $\epsilon > 0$ for *any* of the below problems

- Minimum Weight $(2\ell+1)$-Cycle in a directed weighted graph,
- Shortest Cycle in a directed weighted graph,
- APSP in a directed or undirected weighted graph,
- Radius (or Eccentricities) in a directed or undirected weighted graph,
- Wiener index of a directed or undirected weighted graph,
- Replacement Paths in a directed weighted graph,
- Second Shortest Path in a directed weighted graph,
- Betweenness Centrality of a given node in a directed weighted graph.

That is, we prove hardness for a variety of sparse graph problems from the hardness of a dense graph problem. Our results also lead to new conditional lower bounds from several related hypothesis for unweighted sparse graph problems including k-cycle, shortest cycle, Radius, Wiener index and APSP.

Joint work with Virginia Vassilevska Williams and Ryan Williams.

Ryan O'Donnell (CMU)

Title: Fooling Polytopes

Abstract: We give an explicit pseudorandom generator with seed length poly(logm,1/δ) · logn that δ-fools intersections of m linear thresholds over {0,1}^n. (The previous best seed length had linear dependence on m.) In particular, this gives a deterministic quasipolynomial-time algorithm for approximately counting the number of 0/1 points in an n dimensional polytope with poly(n) facets.

Joint work with Rocco Servedio (Columbia) and Li-Yang Tan (TTI / Stanford).

Igor Oliveira (Oxford)

Title: Hardness Magnification for Natural Problems

Abstract: We show that for several natural problems of interest (including Vertex Cover, Satisfiability, and variants of the Minimum Circuit Size Problem), complexity lower bounds that are barely non-trivial imply super-polynomial or even exponential lower bounds in strong computational models. We term this phenomenon "hardness magnification".

We further explore magnification as an avenue to proving strong lower bounds, and argue that magnification circumvents the "natural proofs" barrier of Razborov and Rudich. Examining some standard proof techniques, we find that they fall just short of proving lower bounds via magnification. As one of our main open problems, we ask whether there are other meta-mathematical barriers to proving lower bounds that rule out approaches combining magnification with known techniques.

Joint work with Rahul Santhanam.

Toniann Pitassi (University of Toronto)

Title: BPP Lifting in Communication Complexity

Abstract: I want to give the main ideas behind the proof, and then briefly discuss some applications and open problems.

Joint work with Mika Goos and Thomas Watson.

Pavel Pudlak (Czech Academy of Sciences)

Title: An Approach to Proving Better Lower Bounds on Bounded Depth Frege Proofs

Abstract: We will present an approach that can potentially lead to a lower bound $2^{n^{1/3d}}$ on depth-*d* Frege proofs of the Tseitin tautologies on random 3-regular graphs. The main idea is to define random restrictions in such a way that they produce, from a Tseitin tautology on a random 3-regular graph, again a Tseitin tautology on a random 3-regular graph on a smaller set of vertices.

This is a work in progress and there are gaps that we have to fill in order to obtain a proof. However, some parts that we already have may be of independent interest, in particular, the form of the switching lemma that we want to use.

Joint work with Navid Talebanfard and Neil Thapen.

Alexander Razborov (University of Chicago)

Title: Grand Challenges in Complexity Theory through the Lens of Proof Theory

Abstract: Given our current inability to even formulate a coherent program towards solving grand challenges in computational complexity (such as P vs. NP), it becomes increasingly important to at least understand this state of affairs; such attempts are often called "barriers". The proof-theoretic approach tries to rule out the existence of solutions within a class of methods that is both rigorously defined and reasonably large, in the hope that this will give at least some insight into the nature of the difficulties. It turns out (and this is where meta-complexity enters the picture) that this approach becomes more successful when the class of methods inherently involves, explicitly or implicitly, the same concepts of feasible computations it intends to study. One framework where this hope has materialized is called Natural Proofs, and the parallel framework in which the corresponding problems are largely open is that of (propositional) proof complexity.

The main purpose of this talk is to give an accessible introduction to this kind of problem, in the hope to attract attention of the new generation of researchers to them.

Ben Rossman (University of Toronto)

Title: Sharper Bounds and Faster #SAT for Regular AC^0 Formulas
Abstract: We say that a boolean function f is "k-critical" if Pr[f under a p-random restriction has decision-tree depth >= t] is at most (pk)^t for every p and t. Hastad's switching lemma states every width-w CNF is O(w)-critical. By an alternative analysis (which bounds the *expected* length of the Razborov encoding of a p-random restriction), we show that every size-m CNF is O(log m)-critical. An extension of this analysis allows us to show that every regular AC^0 formula (in which gates of equal depth have the same fan-in) of size s and depth d+1 is O((1/d)log s)^d-critical. This directly leads to a sharper LMN theorem and faster #SAT algorithms for regular AC^0 formulas, strengthening the best known bounds for AC^0 circuits. As a further corollary, we extend from o(log n / log log n) to o(log n) the number of quantifier alternations for which the QBF-SAT algorithm of Santhanam and Williams (SODA '14) beats exhaustive search.

Srikanth Srinivasan (IIT, Bombay)

Title: A Near-Optimal Depth-Hierarchy Theorem for Small-Depth Multilinear Circuits

Abstract: We study the size blow-up that is necessary to convert an algebraic circuit of product-depth D+1 to one of product-depth D in the multilinear setting.

We show that for every positive D = o(log n/log log n), there is an explicit multilinear polynomial P^D on n variables that can be computed by a multilinear formula of product-depth D+1 and size O(n), but not by any multilinear circuit of product-depth D and size less than exp(n^{Omega(1/D)}). This result is tight up to the constant implicit in the double exponent for all D= o(log n/log log n).

This strengthens a result of Raz and Yehudayoff (Computational Complexity 2009) who prove a quasipolynomial separation for constant-depth multilinear circuits, and a result of Kayal, Nair and Saha (STACS 2016) who give an exponential separation in the case D = 1.

Our separating examples may be viewed as algebraic analogues of variants of the Graph Reachability problem studied by Chen, Oliveira, Servedio and Tan (STOC 2016), who used them to prove lower bounds for constant-depth \emph{Boolean} circuits.

This work is joint with Suryajith Chillara (IIT Bombay, CSE), Christian Engels (IIT Bombay, CSE), and Nutan Limaye (IIT Bombay, CSE).

Avishay Tal (Stanford)

Title: Oracle Separation of BQP and the Polynomial Hierarchy

Abstract: We present an oracle, A, relative to which BQP^A is not contained in PH^A.

Following the approach of Aaronson [STOC, 2010], our oracle separation is obtained by finding a distribution D over Boolean strings of length N such that:
(1) There exists a quantum algorithm that runs in time polylog(N) and distinguishes between D and the uniform distribution over Boolean strings of length N.

(2) No AC0 circuit of quasi-polynomial size can distinguish between D and the uniform distribution with advantage better than polylog(N)/sqrt(N).

Ryan Williams (MIT)

Title: Lower Bounds by Algorithm Design: A Progress Report

Abstract: In 2010, the author proposed a program for proving lower bounds in circuit complexity, via faster algorithms for circuit satisfiability and related problems. This talk will give an overview of how the program works, report on the successes of this program so far, and outline open frontiers that have yet to be resolved.

Virginia Vassilevska Williams (MIT)

Title: Towards Tight Approximation Bounds for Graph Diameter and Eccentricities

Abstract: Among the most important graph parameters is the Diameter, the largest distance between any two vertices. There are no known very efficient algorithms for computing the Diameter exactly. Thus, much research has been devoted to how fast this parameter can be *approximated*. Chechik et al. [SODA 2014] showed that the diameter can be approximated within a multiplicative factor of 3/2 in $\tilde{O}(m^{3/2})$ time. Roditty and Vassilevska W. [STOC 13] showed that unless the Strong Exponential Time Hypothesis (SETH) fails, no $O(n^{2-\varepsilon})$ time algorithm for $\varepsilon>0$ can achieve an approximation factor better than 3/2 in sparse graphs. Thus the above algorithm is essentially optimal for sparse graphs for approximation factors less than 3/2. It was, however, completely plausible that a 3/2-approximation is possible in linear time.

In this work we conditionally rule out such a possibility by showing that unless SETH fails no $O(m^{3/2-\varepsilon})$ time algorithm can achieve an approximation factor better than 5/3. We provide similarly tight results for other distance-related parameters, also providing new algorithms.

To establish the our lower bounds we study the *S*-*T* Diameter problem: Given a graph and two subsets *S* and *T* of vertices, output the largest distance between a vertex in *S* and a vertex in *T*. We give new algorithms and show tight lower bounds that serve as a starting point for all other hardness results. In this talk I will explain the lower bound constructions for *S*-*T* Diameter and will outline how one can extend them for Diameter as well.

Joint work with Arturs Backurs, Nicole Wein, Liam Roditty and Gilad Segal.