Fine-grained Complexity and Algorithm Design: Fall 2015 Final Program Report

Russell Impagliazzo Daniel Marx Ramamohan Paturi
Virginia Vassilevska Williams Richard Ryan Williams

1 Background

Traditionally, theoretical computer science has been divided into "algorithm design" and "computational complexity". Algorithm designers concentrated on developing resource-efficient methods to solve specific problems, whereas complexity theorists considered characteristics of groups of problems that determine their likely difficulty. Over the last approximately twenty years, some inter-related research themes have emerged within theoretical computer science that challenge this dichotomy.

A first trend was a movement within complexity to consider more exact notions of complexity, bringing up issues closer to those in algorithm design. While traditionally, complexity classes qualitatively distinguish between "hard" computational problems and "easy" problems, several new lines of inquiry have considered quantitative distinctions between problems, trying to identify the likely *exact* time complexities of problems. These apply both to problems complete for large classes such as NP and problems within P. One such line of work began with the 3-SUM and related conjectures from computational geometry, showing the difficulty of a variety of problems within P assuming that 3-SUM cannot be solved in sub-quadratic time. A parallel line of work at the exponential-time level (for NP-hard problems) centered around the Exponential Time Hypothesis (ETH) and Strong Exponential Time Hypothesis (SETH). Bridging polynomial and exponential time was part of the goal of parameterized complexity, which makes algorithm design and complexity multi-dimensional by identifying parameters of the instance that determine the instance's computational difficulty. While these areas were differently motivated, over the years, more and more connections were found between them, with hypotheses such as ETH having strong ramifications for both parameterized complexity and the complexity of problems within P.

A second emerging theme was a duality between circuit lower bounds (proving there are no efficient algorithms of a certain type for solving problems) and designing improved algorithms for other problems. One of the first such connections was in the area of *derandomization*, where a problem that is hard for circuits can be used to construct a pseudo-random generator, and thus a deterministic simulation of any probabilistic algorithm. More surprisingly, connections in the other direction were also shown, first for derandomization, and then for circuit satisfiability, where an improved algorithm yields a circuit lower bound as a consequence. This was used by Williams to prove the first lower bound against ACC^0 circuits, an open problem for over twenty years. Circuit lower bound techniques have also been used to design more efficient algorithms, perhaps starting with the classical learning algorithm of Linial, Mansoor and Nisan, but really hitting its stride more recently.

When these ideas are combined, they yield even more surprising results. Potentially, a small improvement in the running time of algorithms for a polynomial time problem could give an improved SAT algorithm, and hence a new lower bound. Or, alternatively, a circuit lower bound could be used to design new algorithms. The past 2-3 years have seen a flurry of activity towards realizing this potential, and this momentum was greatly increased by the Simons Institute program.

2 Goals

More specifically, the goals of the program were:

1. To bring together researchers from traditional complexity theory, traditional algorithm design, and parameterized algorithms and complexity to work together on exploiting the connections between complexity and algorithm design.

2. To strengthen the connections between lower bounds and algorithm design, between algorithms for hard problems and those for polynomial time problems, and between parameterized and traditional complexity and algorithm design.

3. To utilize these connections to develop new algorithms and prove new lower bounds.

As we will now describe, progress on all of these fronts during the program exceeded our already high expectations.

3 Fostering inter-area collaborations

One of our primary goals in the program was to foster collaboration between algorithms and complexity researchers. Here are three illustrations of successful collaborations between different areas.

3.1 Proofs of work

The first story is a collaboration that crossed over several Simons Institute programs, starting with the summer 2015 Cryptography program, continuing into our program in Fall 2015, and continuing into the Pseudorandomness program in Spring 2017.

To maintain consistency in their transaction records in a totally distributed environment, crypto-currencies such as BitCoin rely on the concept of "proofs of work", a way to certify that computational effort has been expended. While this idea was introduced in the early 90's in the theoretical computer science community in work by Dwork and Naor, the proofs of work currently used involve ad hoc, untested assumptions about cryptographic hash functions. Furthermore, since the problems used in proofs of work are artificial constructs, the immense computational resources (energy, time, and physical devices devoted to bitcoin mining) expended on constructing proofs of work is wasted without any benefits to society.

In a series of recent papers, Ball, Rosen, Sabin, and Vasudevan bring a rigorous foundation to proofs of work. Their starting point is "complexity within P", the effort to identify problems within polynomial time where the known algorithms are conjectured to be optimal. This effort was one of the focal points in our program. They give average-case hard versions of these problems, and then show that solutions to these problems can be probabilistically verified in strictly less time than it takes to compute them (under the standard conjectures). This gives the first proofs of work with security provable from well-known computational assumptions. Moreover, because these proofs start from problems of actual algorithmic interest, and are in fact capable of coding a variety of other problems, in sequel work, the authors show that these proofs of work can be made "for a purpose", replacing the computational waste of current crypto-currencies with socially useful computation. (There have been prior attempts to make a proof of work for a purpose, but these have again been ad hoc, and the algorithmic utility limited, e.g., for finding Cunningham chains of prime numbers in Primecoin.)

The collaboration between these researchers started in the Cryptography program at Simons. However, things really got started when Manuel Sabin, a graduate student at Berkeley, participated in the Fine-Grained Complexity program. He became aware not only of "complexity within P", but of techniques such as embedding functions in algebraic extensions to make average-case hard versions. He was also aware of a recent result of Williams discovered during the program, giving a way to probabilistically verify proofs of unsatisfiability faster than exhaustive search. This was adapted into the probabilistic procedure to verify proofs of work. Thus, both the general tools of fine-grained complexity and the specific program at the Simons Institute were essential to this project. While still in its early stages, this project has immense potential for impact on crypto-currencies, which themselves have immense potential for economic, social, and technological impact. This work appeared in STOC 2017 [13] and Crypto 2017 [14].

3.2 New algorithms for RNA folding and related problems

The discovery of the first sub-cubic algorithms for RNA folding, Language Edit Distance, and Min-Plus Products (with bounded differences) constitutes the first major algorithmic improvement on these textbook dynamic programming problems, as well as an example of how fine-grained lower bounds can inform algorithm design.

While very different in details and applications, these problems all had similar known algorithms, that combined dynamic programming with divide-and-conquer to get an $O(n^3)$ algorithm. The $O(n^3)$ algorithms for these problems are exercises that one can give an undergraduate algorithms class, and the improvements over this methodical algorithm in the literature were small. For example, RNA Folding is the following problem: given a string s of length n over the alphabet A, U, C, G where A can match U and C can match G but no other matches are possible (or, more generally, for a larger alphabet with pairs of matching letters), find a maximum set of disjoint matches $\{(s[i_1], s[j_1]), \ldots, (s[i_k], s[j_k]\}$ where $i_t < j_t$ for all t and there are no two matches $(s[i_t], s[j_t]), (s[i_r], s[j_r])$ with $i_t < i_r < j_t < j_r$. One can think of the problem as finding a maximum set of balanced parentheses where there are several types of parentheses. There is an easy $O(n^3)$ time dynamic programming algorithm for RNA-folding. The fastest previous algorithm for the problem, due to Venkatachalam et al., ran in $O(n^3/\log^2 n)$ time. The Language Edit Distance (LED) Problem is a natural extension to the Context Free Grammar (CFG) Recognition problem: given a length-n string x and a (constant size) CFG encoding a context free language L, determine whether x is in L. In the LED problem one is also given a string x and a CFG defining L and one is to determine the minimum edit distance between x and any $y \in L$. The LED problem also has an $O(n^3)$ algorithm but nothing much better.

Looking at the best conditional lower bounds for these problems revealed a significant gap. An earlier paper by Amir Abboud, Arturs Backurs and Virginia Vassilevska-Williams on the RNA Folding and LED problems showed that unless k-clique can be solved faster, both problems on an input string of length n require n^ω time where $\omega < 2.38$ is the exponent of square matrix multiplication. Even though the two problems only had essentially cubic time algorithms, no higher lower bound than n^ω seemed possible. This left open the possibility that one might obtain truly subcubic algorithms; and if this could be done without matrix multiplication, it would lead to an improvement in "combinatorial" matrix multiplication. Thus, the lower bounds indicate that matrix multiplication should be key to any improvements.

One approach seemed particularly promising. In the 1970s, Valiant had developed an $O(n^\omega)$ time algorithm for CFG recognition. This algorithm can be viewed as a clever reduction to Boolean matrix multiplication. Valiant's approach, with a bit of modification, also works for LED and RNA-folding, except that instead of Boolean matrix multiplication, the reduction is to $(\min, +)$ matrix product. Unfortunately, the best known algorithms for $(\min, +)$ matrix product on $n \times n$ matrices run in essentially cubic time and obtaining a truly faster algorithm is an important open problem.

Fortunately, however, the matrices arising from Valiant's approach, when applied to RNA-folding and LED, are highly structured. In particular, any two consecutive entries in these matrices (row-wise or column-wise) differ only by 1 in absolute value. Thus one only needs to find a truly subcubic algorithm for $(\min, +)$-product of such *bounded differences* (BD) matrices. Recent work of Chan and Lewenstein had produced a truly subquadratic algorithm for the related bounded monotone $(\min, +)$-convolution problem, and it seemed plausible that similar ideas might apply to the BD $(\min, +)$-product problem.

Co-organizer Virginia Vassilevska-Williams was part of the team proving the conditional lower bound. Simons Fellow Barna Saha had previously generalized Valiant's technique. While at the Simons Institute, Virginia proposed to Uri Zwick and Fabrizio Grandoni to work on developing a faster algorithm for RNA Folding. Independently, Barna Saha proposed the same research problem to Karl Bringmann and Fabrizio. Because of the shared interest and because everyone was at the Simons Institute, Barna, Karl, Fabrizio, Uri[1] and Virginia started meeting to brainstorm ideas, sometimes all five, sometimes in smaller groups. Many ideas were discussed: breaking up the matrices into blocks and only using block representatives, subtracting carefully chosen row and column vectors to scale the matrices, using fast $(\min, +)$-product algorithms for other types of structured matrices, etc. After a month or so, all the puzzle pieces fitted together and a truly subcubic algorithm was developed. The entire research project was completed at the Simons Institute.

After getting this breakthrough result, the collaboration continued. The algorithm turns out to be

[1] Although Uri Zwick eventually decided to leave the project, he had many instrumental ideas that the rest of the collaborators are grateful for.

even more general than the original goal for it, and has numerous other applications that are still being investigated. This work appeared in FOCS 2016 [31].

3.3 The polynomial method and algorithm design

Ryan Williams, a co-organizer of the program, has developed a line of work showing how the polynomial method from circuit complexity can be applied to the design of new algorithms for long-studied problems. In circuit complexity, the *polynomial method* is a general approach to proving low-depth circuit lower bounds. One (a) models low-complexity functions *approximately* with low-degree polynomials, then (b) proves that certain simple functions do not have low-degree polynomial approximations. A generation of papers from the late 80s and early 90s proved many circuit complexity lower bounds in this way.

Collaborative work between Josh Alman, Ryan Williams (both in residence at the Simons Institute) and Timothy Chan of U. Waterloo (short-term visitor at Simons, and prominent algorithms and geometry researcher) has led to new algorithmic applications of the polynomial method. At the heart of the new results is the discovery of new polynomial representations for the OR of MAJORITY function. For example, one can compute the OR of s MAJORITY functions on n variables with a so-called *probabilistic* polynomial threshold function (probabilistic PTF), a distribution of polynomials, each polynomial having degree at most $\tilde{O}(n^{1/3})$. (This is in contrast with the known $\sqrt{n}$-degree lower bounds for probabilistic polynomials computing MAJORITY, due to Razborov and Smolensky.) These polynomial constructions have led to new algorithms for computing approximate Closest Pair in Euclidean distance, new circuit satisfiability algorithms, and stronger circuit lower bounds against $NEXP$. This collaboration, which again appeared in FOCS 2016 [9], nicely illustrates how algorithms researchers can learn new tricks from complexity theorists, in this case by studying polynomials.

4 Further research highlights

The program has produced too many and too diverse a set of excellent results to do them all justice in this brief report. In this section, we outline a few of our favorites.

4.1 Edit distance can simulate sub-linear space Turing machines

A line of celebrated recent research had established that, if the Strong Exponential Time Hypothesis holds, then the quadratic-time algorithms for many basic problems comparing strings and sequences cannot be improved to an exponent less than 2. The problems shown hard are basic for areas from computational biology to document analysis to computational geometry, and have been subject to intense algorithmic interest, with relatively little progress over the first algorithms to be discovered. This connection shows that a major breakthrough on algorithm design for these problems would give both improved CNF SAT algorithms and new circuit lower bounds for weak circuit classes.

However, during the program, the breakthrough result of Abboud, Hansen, Vassilevska-Williams and Williams [4] pushed this connection into hyperdrive. They give major quantitative and qualitative improvements over the reductions between SAT and these string and sequence comparison problems. They can reduce satisfiability not just of CNF's but of any circuit model that uses sub-polynomial parallel time or sub-polynomial space to these string problems, in such a way that only polylogarithmic improvements to the quadratic time algorithms are needed to get improved circuit SAT algorithms, and hence lower bounds, for these classes of circuits. This yields two equally fascinating possibilities: either there is almost no improvement possible for the currently known algorithms for these basic problems; or by finding such a minor improvement, we can leap-frog the classes for which we have circuit lower bounds to classes just short of arbitrary circuits. Much exciting work during the program generalized or built on this breakthrough, which strengthened the lower-bound/algorithm connection beyond our most audacious hopes.

4.2 Circuit lower bounds and learning algorithms

Work by Carmosino, Impagliazzo, Kabanets and Kolokolova [35] both strengthened connections between circuit lower bounds and learning algorithms and used lower bounds to design new learning algorithms. The

starting point for this work is the classic result of Linial, Mansour and Nisan, which used the proof of the lower bound for constant depth unbounded fan-in circuits (AC^0) to design a quasi-polynomial time learning algorithm for such circuits over the uniform distribution (utilizing Fourier analysis). While this makes an intuitive connection between lower bounds and learning algorithms, there was no known general technique for deriving a learning algorithm from lower bounds. Indeed, no such learning algorithms were known for several circuit classes for which we did have lower bounds, such as extending constant depth Boolean circuits with prime modular counting gates (ACC^0_p). The new work made such a general connection, showing that any natural proof (in the Razborov-Rudich sense) of a circuit lower bound yields a corresponding learning algorithm for the same class of circuits. Moreover, it applied this general connection to the case of ACC^0_p, yielding the first quasi-polynomial learning algorithm for such circuits. This is a great example of both strengthening the lower-bound/algorithm connection and of using computational complexity as an algorithm design tool. The algorithm itself consists of several repurposed computational complexity tools, combining a natural property of Razborov and Rudich with the correctness proof for the Nisan-Wigderson derandomization technique. This work was done while all collaborators were participating in the program, and won the Best Paper Award at the 2016 Computational Complexity Conference.

4.3 Subgraph Isomorphism on planar graphs

Another major breakthrough, and another example where conditional lower bounds complimented algorithmic progress, is in the fine-grained parameterized complexity of Subgraph Isomorphism on planar graphs. This can be considered an advance in the optimality program for parameterized algorithms, which aims to obtain a fine-grained understanding of the complexity of parameterized problems. The ultimate goal of this program is to determine the best possible function $f(k)$ that can appear in the running time $f(k) \cdot n^{O(1)}$ of a parameterized algorithm for a given problem. By now, tight or almost tight upper and conditional lower bounds are known for a wide range of natural parameterized problems. An area where the optimality program has been particularly successful is understanding how algorithmic graph problems become easier when restricted to planar or H-minor-free graphs. A so-called "square root phenomenon" has been observed: for most natural parameterized problems on planar graphs, the best possible dependence on k is exponential in the *square root* of k. However, the complexity of Subgraph Isomorphism, which is arguably one of the most basic graph-theoretical problems, was poorly understood on planar graphs: parameterized algorithms that have subexponential dependence on the size of the pattern to be found were known only in very special patterns, such as paths. Fomin, Lokshtanov, Marx, Pilipczuk, Pilipczuk, and Saurabh [44] obtained a breakthrough in the important special case of connected bounded-degree patterns by giving a $2^{\sqrt{k}\text{polylog}k} \cdot n^{O(1)}$ time algorithm for Subgraph Isomorphism on planar graphs. The algorithmic technique developed for this problem is sufficiently robust to allow several extensions, such as generalizations to directed, colored, or weighted versions. As a surprisingly timely coincidence, recent work of Bodlaender and van der Zanden (ICALP 16) shows that the restriction to the special case of connected bounded-degree graphs is unavoidable: if we drop the connectivity requirement or drop the bounded-degree requirement, then (assuming ETH) no $2^{o(k/\log k)} \cdot n^{O(1)}$ time algorithm exists for the problem. Thus, the algorithmic breakthrough has a corresponding breakthrough on conditional lower bounds, showing the exact boundary where progress was possible. The crucial parts of the work were done while three of the authors were participating in the Simons Institute program, and when all authors were present for the program workshop "Satisfiability Lower Bounds and Tight Results for Parameterized and Exponential-Time Algorithms." The paper was presented at IEEE FOCS 2016.

5 Mentoring

We feel the program was particularly successful at involving younger researchers. Each Simons Fellow was assigned a senior mentor whom they met with on a regular basis; many of these relationships became research collaborations. In addition, many graduate students were very active in the program, both the advisees of participants and Berkeley graduate students. A large number of the exciting results proven during the program had one or more student authors.

6 Collaborations with other Simons Institute activities

There was some interesting overlap with the concurrent program, *Economics and Computation*. For example, Christos Papadimitriou organized a very successful seminar on classes within TFNP, in particular, classes relevant to Nash Equilibria and other problems connected with economics. This was very well attended and had a high level of participation from the complexity theorists in our program. Other events of joint interest included the field trip and reverse field trip to highlight joint interests with Silicon Valley researchers. Finally, many program participants found the workshop on Neuroscience, co-sponsored by the Simons Institute, MSRI and the Redwood Center and held during the same semester, very interesting.

One major inter-program success was Aviad Rubinstein's work on approximation of Nash equilibria. Rubinstein is a Ph.D. student at Berkeley working with Papadimitriou, and was a participant in both our program and the concurrent Economics program. His paper "Settling the complexity of computing approximate two-player Nash equilibria" formulates an equivalent of the Strong Exponential Time Hypothesis for the class PPAD (widely used in algorithmic game theory) and proves that, under this hypothesis, there is no polynomial time algorithm for finding approximate Nash equilibria. This paper [86] uses techniques from probabilistically checkable proofs as well as fine-grained complexity and game theory. It won both the Machety Award for Best Student Paper and the overall Best Paper Award at IEEE FOCS 2016, a unique achievement.

7 Impact and subsequent developments

The level of interest in "fine-grained complexity" within the broader TCS community is accelerating, and the program contributed decisively to this trend. For example, three of six plenary lectures at the Highlights of Algorithms conference in Paris in summer 2016 were related to fine-grained complexity and given by program participants. Many theory blogs have listed this as one of their favorite new directions within TCS. We feel the training of young researchers in this area, and the establishment of new collaborations, may be the most lasting consequence of the program.

For example, Pasin Manurangsi, a graduate student at UC, Berkeley, was not an official participant in the program, but attended many of the events and was a student in Impagliazzo's UC Berkeley course which introduced the themes of the program, with participation by many of the other researchers in the program. He applied these ideas to understand the complexity of approximation algorithms for the well-known densest subgraph problem, where the goal is to find a set of vertices of a given size that has as many edges as possible. (Variants of this problem are used to identify communities in social networks and related graphs.) He was able to show, assuming ETH, an almost polynomial hardness of approximation, almost matching known algorithms. This work won the Best Student Paper Award at STOC 2017 [76].

Another example where our program helped shape the career of young reseachers is the case of Marvin Künnemann and Stefan Schneider. Both participated as graduate students in the program, and began collaborating with each other and with Schneider's advisor, Mohan Paturi, during the program. The next year, Künnemann joined UCSD as a post-doc to continue this collaboration. There, they investigated a broad class of problems solved by dynamic programming algorithms, including standard problems such as change-making, and longest chain of nested boxes. They showed that, for many such problems, improving over the dynamic programming was equivalent to improving a static version (e.g,, vector domination is equivalent to chain of nested boxes). This is counter-intuitive, since the dynamic programming algorithm seems inherently sequential, whereas the static versions are embarrassingly parallel. This work [72] will be presented at ICALP 2017, and will also form the major part of Schneider's PhD thesis.

It remains to be seen what heights this sub-area will attain in its full maturity, but it seems well on its way to becoming a recognized major research direction within TCS. In particular, we feel this program has made both the algorithmic and computational complexity communities appreciate each other's contributions to a much greater extent, and be more open to collaboration. As we have seen in many of the examples above, this collaboration can lead to major advances in our understanding of fundamental computational questions.

Fine-Grained Complexity and Algorithm Design, Fall 2015

[1] A. ABBOUD and K. BRINGMANN. Barriers for Shaving Log-Factors for LCS and Frechet. In submission.

[2] A. ABBOUD and G. BODWIN. The 4/3 Additive Spanner Exponent is Tight. In *Proceedings of the 48th ACM Symposium on Theory of Computing (STOC)*, 2016.

[3] A. ABBOUD and S. DAHLGAARD. Popular conjectures as a barrier for dynamic planar graph algorithms. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 477–486, 2016.

[4] A. ABBOUD, T. D. HANSEN, V. V. WILLIAMS, and R. WILLIAMS. Simulating Branching Programs with Edit Distance and Friends or: A Polylog Shaved is a Lower Bound Made. In *Proceedings of the 48th ACM Symposium on Theory of Computing (STOC)*, 2016.

[5] I. ABRAHAM, S. CHECHIK, and S. KRINNINGER. Fully dynamic all-pairs shortest paths with worst-case update-time revisited. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2017.

[6] I. ABRAHAM, D. DURFEE, I. KOUTIS, S. KRINNINGER, and R. PENG. On fully dynamic graph sparsifiers. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 335–344, 2016.

[7] D. ACHLIOPTAS and F. ILIOPOULOS. Focused Stochastic Local Search and the Lovász Local Lemma. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2016.

[8] D. ACHLIOPTAS, F. ILIOPOULOS, and N. VLASSIS. Stochastic Control via Entropy Compression. *arXiv preprint arXiv:1607.06494,* 2016.

[9] J. ALMAN, T. CHAN, and R. WILLIAMS. Polynomial Representations of Threshold Functions with Applications. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2016.

[10] S. ARTEMENKO, R. IMPAGLIAZZO, V. KABANETS, and R. SHALTIEL. Pseudorandomness when the odds are against you. In *Proceedings of the 31st IEEE Conference on Computational Complexity (CCC)*, 2016.

[11] P. AUSTRIN, P. KASKI, M. KOIVISTO, and J. NEDERLOF. Dense Subset Sum May Be the Hardest. In *Proceedings of the 33rd International Symposium on Theoretical Aspects of Computer Science (STACS)*, pp. 3:1–13:14, 2016.

[12] P. AUSTRIN, P. KASKI, M. KOIVISTO, and J. NEDERLOF. Sharper upper bounds for unbalanced uniquely decodable code pairs. In *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, pp. 335–339, 2016.

[13] M. BALL, A. ROSEN, M. SABIN, and P. N. VASUDEVAN. Average-Case Fine-Grained Hardness. In *Proceedings of the 49th ACM Symposium on Theory of Computing (STOC)*, 2017.

[14] M. BALL, A. ROSEN, M. SABIN, and P. N. VASUDEVAN. Proofs of Useful Work. To appear in Proceedings of CRYPTO, 2017.

[15] N. BANSAL, D. DADUSH, and S. GARG. An Algorithm for Komlós Conjecture Matching Banaszczyk's Bound. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 788–799, 2016.

[16] N. BANSAL, S. GARG, J. NEDERLOF, and N. VYAS. Faster Space-Efficient Algorithms for Subset Sum, k-Sum and Related Problems. *arXiv preprint arXiv:1612.02788,* 2016.

[17] N. BANSAL, D. REICHMAN, and S. W. UMBOH. LP-Based Robust Algorithms for Noisy Minor-Free and Bounded Treewidth Graphs. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1964–1979, 2017.

[18] N. BANSAL, O. SVENSSON, and A. SRINIVASAN. Lift-and-Round to Improve Weighted Completion Time on Unrelated Machines. In *Proceedings of the 48th ACM Symposium on Theory of Computing (STOC)*, 2016.

[19] B. BARAK, S. B. HOPKINS, J. KELNER, P. KOTHARI, A. MOITRA, and A. POTECHIN. A nearly tight sum-of-squares lower bound for the planted clique problem. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 428–437, 2016.

[20] P. BEAME and V. LIEW. Towards Verifying Nonlinear Integer Arithmetic. Submitted to Computer Aided Verification, 2017.

[21] P. BEAME and C. RASHTCHIAN. Massively-Parallel Similarity Join, Edge-Isoperimetry, and Distance Correlations on the Hypercube. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 289–306, 2017.

[22] S. BHATTACHARYA, M. HENZINGER, and D. NAMONGKAI. New deterministic algorithms for full dynamic matching. In *Proceedings of the 48th ACM Symposium on Theory of Computing (STOC)*, 2016.
[23] A. BJÖRKLUND, H. DELL, and T. HUSFELDT. The Parity of Set Systems Under Random Restrictions with Applications to Exponential Time Problems. In *ICALP (1)*, pp. 231–242, 2015. Journal version under review.
[24] A. BJÖRKLUND and P. KASKI. How proofs are prepared at Camelot. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, pp. 391–400, 2016.
[25] H. L. BODLAENDER, J. NEDERLOF, and T. C. VAN DER ZANDEN. Subexponential time algorithms for embedding H-minor free graphs. In *Leibniz International Proceedings in Informatics (LIPIcs)*, vol. 55, 2016.
[26] G. BODWIN and S. KRINNINGER. Fully Dynamic Spanners with Worst-Case Update Time. In *Proceedings of 24th Annual European Symposium on Algorithms (ESA)*, pp. 17:1–17:18, 2016.
[27] I. BONACINA, N. GALESI, and N. THAPEN. Total space in resolution. *SIAM Journal on Computing,* 45, 5, pp. 1894–1909, 2016.
[28] R. BOPPANA, J. HÅSTAD, C. H. LEE, and E. VIOLA. Bounded independence vs. moduli. In *Leibniz International Proceedings in Informatics (LIPIcs)*, vol. 60, pp. 24:1–24:9, 2016.
[29] C. BRAND, H. DELL, and M. ROTH. Fine-grained dichotomies for the Tutte plane and Boolean #CSP. In *Proceedings of the 11th International Symposium on Parameterized and Exact Computation (IPEC)*, 2016.
[30] K. BRINGMANN. A near-linear pseudopolynomial time algorithm for subset sum. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1073–1084, 2017.
[31] K. BRINGMANN, F. GRANDONI, B. SAHA, and V. V. WILLIAMS. Truly Sub-cubic Algorithms for Language Edit Distance and RNA-Folding via Fast Bounded-Difference Min-Plus Product. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 375–384, 2016.
[32] K. BRINGMANN, A. GRØNLUND, and K. G. LARSEN. A Dichotomy for Regular Expression Membership Testing. *arXiv preprint arXiv:1611.00918v2,* 2016.
[33] K. BRINGMANN, L. KOZMA, S. MORAN, and N. NARAYANASWAMY. Hitting Set for hypergraphs of low VC-dimension. In *Proceedings of 24th Annual European Symposium on Algorithms (ESA)*, pp. 23:1–23:18, 2016.
[34] B. BUKH, V. GURUSWAMI, and J. HÅSTAD. An Improved Bound on the Fraction of Correctable Deletions. *IEEE Transactions on Information Theory,* 63, 1, pp. 93-103, Jan 2017.
[35] M. CARMOSINO, R. IMPAGLIAZZO, V. KABANETS, and A. KOLOKOLOVA. Learning Algorithms from Natural Proofs. In *Proceedings of the 31st IEEE Conference on Computational Complexity (CCC)*, 2016. Best paper award.
[36] M. L. CARMOSINO, J. GAO, R. IMPAGLIAZZO, I. MIHAJLIN, R. PATURI, and S. SCHNEIDER. Nondeterministic extensions of the strong exponential time hypothesis and consequences for non-reducibility. In *Proceedings of the 7th ACM Conference on Innovations in Theoretical Computer Science (ITCS)*, pp. 261–270, 2016.
[37] R. CHEN, R. SANTHANAM, and S. SRINIVASAN. Average-Case Lower Bounds and Satisfiability Algorithms for Small Threshold Circuits. ECCC Technical Report TR15-191, 2015.
[38] R. CURTICAPEAN, H. DELL, and D. MARX. Homomorphisms are a good basis for counting small subgraphs. To appear in STOC, 2017.
[39] R. CURTICAPEAN. Parity Separation: A Scientifically Proven Method for Permanent Weight Loss. In *Proceedings of the 43rd International Colloquium on Automata, Languages, and Programming (ICALP)*, pp. 47:1–47:14, 2016.
[40] R. CURTICAPEAN, H. DELL, and T. HUSFELDT. Modular Subgraph Counting. In preparation.
[41] R. CURTICAPEAN and D. MARX. Tight conditional lower bounds for counting perfect matchings on graphs of bounded treewidth, cliquewidth, and genus. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1650–1669, 2016.
[42] A. DRUCKER, J. NEDERLOF, and R. SANTHANAM. A short note on Merlin-Arthur protocols for subset sum. *Information Processing Letters,* 118, pp. 15–16, 2017.
[43] A. DRUCKER, J. NEDERLOF, and R. SANTHANAM. Exponential Time Paradigms Through the Polynomial Time Lens. In *Proceedings of 24th Annual European Symposium on Algorithms (ESA)*, pp. 36:1–36:14, 2016.
[44] F. FOMIN, D. LOKSHTANOV, D. MARX, M. PILIPCZUK, M. PILIPCZUK, and S. SAURABH. Subexponential parameterized algorithms for planar and apex-minor free graphs via low treewidth pattern covering. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 515–524, 2016.
[45] F. FOMIN, S. GASPERS, D. LOKSHTANOV, and S. SAURABH. Exact Algorithms via Monotone Local

Search. In *Proceedings of the 48th ACM Symposium on Theory of Computing (STOC)*, pp. 764–775, 2016.
[46] N. GALESI. Space of refuting random 3CNFS. 2016. In submission.
[47] J. GAO, R. IMPAGLIAZZO, A. KOLOKOLOVA, and R. WILLIAMS. Completeness for First-Order Properties on Sparse Structures with Algorithmic Applications. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 2162–2181, 2017.
[48] A. GOLOVNEV, A. S. KULIKOV, A. V. SMAL, and S. TAMAKI. Circuit size lower bounds and #SAT upper bounds through a general framework: new bounds and limitations. In *Proceedings of the 41st International Symposium on Mathematical Foundations of Computer Science (MFCS)*, pp. 45:1–45:16, 2016.
[49] W. T. GOWERS and E. VIOLA. The multiparty communication complexity of interleaved group products. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 289–294, 2016.
[50] E. HARAMATY, C. H. LEE, and E. VIOLA. Bounded independence plus noise fools products. *Electronic Colloquium on Computational Complexity (ECCC),* TR16-169, 2016.
[51] P. HARSHA and S. SRINIVASAN. Robust Multiplication-based Tests for Reed-Muller Codes. In *Foundations of Software Technology and Theoretical Computer Science*, 2016.
[52] J. HÅSTAD. On the average-case of small-depth circuits. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 79–88, 2016.
[53] M. HENZINGER, A. LINCOLN, S. NEUMANN, and V. WILLIAMS. Conditional Hardness for Sensitivity Problems. In *Proceedings of the 8th ACM Conference on Innovations in Theoretical Computer Science (ITCS)*, 2017.
[54] M. HENZINGER, G. GORANCI, and M. THORUP. Incremental Edge Connectivity in polylogarithmic amortized time. In preparation.
[55] M. HENZINGER, S. KRINNINGER, and D. NANONGKAI. An Almost-Tight Distributed Algorithm for Computing Single-Source Shortest Paths. In *Proceedings of the 48th ACM Symposium on Theory of Computing (STOC)*, 2016.
[56] M. HENZINGER, S. RAO, and D. WANG. Local Flow Partitioning for Faster Edge Connectivity. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1919–1938, 2017.
[57] T. HUSFELDT. Computing Graph Distances Parameterized by Treewidth and Diameter. In *Leibniz International Proceedings in Informatics (LIPIcs)*, vol. 63, 2017.
[58] R. IMPAGLIAZZO, R. JAISWAL, V. KABANETS, B. M. KAPRON, V. KING, and S. TESSARO. Simultaneous security and reliability amplification for a general channel model. 2016. In submission.
[59] R. IMPAGLIAZZO and V. KABANETS. Fourier concentration from shrinkage. In *Proceedings of the 29th IEEE Conference on Computational Complexity (CCC)*, pp. 321–332, 2014. Revised journal version, December 2015.
[60] R. IMPAGLIAZZO, V. KABANETS, A. KOLOKOLOVA, P. MCKENZIE, and S. ROMANI. Does looking inside a circuit help? In preparation.
[61] B. M. P. JANSEN and A. PIETERSE. Optimal sparsification for some binary CSPs using low-degree polynomials. In *Proceedings of the 41st International Symposium on Mathematical Foundations of Computer Science (MFCS)*, vol. 58, pp. 71:1–71:14, 2016.
[62] B. M. P. JANSEN and M. PILIPCZUK. Approximation and kernelization for chordal vertex deletion. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2017.
[63] R. JAYARAM and B. SAHA. Language Edit Distance Approximation via Amnesic Dynamic Programming. Under submission.
[64] J. JEONG, E. J. KIM, and S.-I. OUM. Efficient and constructive algorithms for rank-width and branch-width. Manuscript, 2016.
[65] J. JEONG, E. J. KIM, and S.-I. OUM. Constructive algorithm for path-width of matroids. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1695-1704, 2016.
[66] D. M. KANE and R. WILLIAMS. Super-Linear Gate and Super-Quadratic Wire Lower Bounds for Depth-Two and Depth-Three Threshold Circuits. In *Proceedings of the 48th ACM Symposium on Theory of Computing (STOC)*, pp. 633–643, 2016.
[67] K. KANGAS, T. HANKALA, T. NIINIMÄKI, and M. KOIVISTO. Counting linear extensions of sparse posets. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 603–609, 2016.
[68] M. KARPPA, P. KASKI, and J. KOHONEN. A faster subquadratic algorithm for finding outlier correlations. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2016.

[69] M. KARPPA, P. KASKI, J. KOHONEN, and P. Ó. CATHÁIN. Explicit correlation amplifiers for finding outlier correlations in deterministic subquadratic time. In *Proceedings of 24th Annual European Symposium on Algorithms (ESA), Leibniz International Proceedings in Informatics (LIPIcs)*, vol. 57, pp. 52:1–17, 2016.

[70] E. J. KIM and O. KWON. A Polynomial Kernel for Distance-Hereditary Vertex Deletion. *arXiv preprint arXiv: 1610.07229,* 2016.

[71] P. KOTHARI, R. MEKA, and P. RAGHAVENDRA. Approximating Rectangles by Juntas and Weakly-Exponential Lower Bounds for LP Relaxations of CSPs. *arXiv preprint arXiv:1610.02704,* 2016. To appear in STOC, 2017.

[72] M. KÜNNEMANN, R. PATURI, and S. SCHNEIDER. On the Fine-grained Complexity of One-Dimensional Dynamic Programming. *arXiv preprint arXiv:1703.00941,* 2017. To appear in Proceedings of ICALP, 2017.

[73] K. G. LARSEN and R. WILLIAMS. Faster online matrix-vector multiplication. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 2182–2189, 2017.

[74] A. LINCOLN, V. V. WILLIAMS, J. R. WANG, and R. R. WILLIAMS. Deterministic Time-Space Trade-Offs for k-SUM. In *Proceedings of the 43rd International Colloquium on Automata, Languages, and Programming (ICALP)*, pp. 58:1–58:14, 2016.

[75] D. LOKSHTANOV, R. PATURI, S. TAMAKI, and R. WILLIAMS. Beating Brute Force For Systems of Polynomial Equations Over Finite Fields. 2016. In submission.

[76] P. MANURANGSI. Almost-polynomial ratio ETH-hardness of approximating densest k-subgraph. In *Proceedings of the 49th ACM Symposium on Theory of Computing (STOC)*, 2017. Best Student Paper Award.

[77] D. MARX and M. PILIPCZUK. Subexponential parameterized algorithms for graphs of polynomial growth. *arXiv preprint arXiv:1610.07778v1,* 2016.

[78] D. MEDARAMETLA and A. POTECHIN. Bounds on the norms of uniform low degree graph matrices. In *Proceedings of Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, 2016.

[79] R. MEKA. Explicit resilient functions matching Ajtai-Linial. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2017.

[80] D. MOELLER, R. PATURI, and S. SCHNEIDER. Subquadratic Algorithms for Succinct Stable Matching. In *Proceedings of the 11th International Computer Science Symposium in Russia (CSR)*, 2016.

[81] J. NEDERLOF. Finding Large Set Covers Faster via the Representation Method. In *Proceedings of 24th Annual European Symposium on Algorithms (ESA)*, pp. 69:1-69:15, 2016.

[82] J. PACHOCKI, L. RODITTY, A. SIDFORD, R. TOV, and V. V. WILLIAMS. Approximating Cycles in Directed Graphs: Fast Algorithms for Girth and Roundtrip Spanners. *arXiv preprint arXiv:1611.00721,* 2016.

[83] P. PUDLAK. Incompleteness in the finite domain. *arXiv preprint arXiv:1601.01487v1,* 2016. Submitted to Bulletin of the ASL.

[84] A. RUBINSTEIN. Beyond matroids: Secretary problem and prophet inequality with general constraints. In *Proceedings of the 48th ACM Symposium on Theory of Computing (STOC)*, pp. 324–332, 2016.

[85] A. RUBINSTEIN. Eth-hardness for signaling in symmetric zero-sum games. *CoRR, abs/1510.04991,* 2015.

[86] A. RUBINSTEIN. Settling the complexity of computing approximate two-player Nash equilibria. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 258–265, 2016. Best Paper and Best Student Paper.

[87] T. SAKAI, K. SETO, S. TAMAKI, and J. TERUYAMA. Improved Exact Algorithms for Mildly Sparse Instances of Max SAT. In *Proccedings of the 10th International Symposium on Parameterized and Exact Computation (IPEC)*, vol. 43, pp. 90-101, 2015.

[88] T. SAKAI, K. SETO, S. TAMAKI, and J. TERUYAMA. Bounded Depth Circuits with Weighted Symmetric Gates: Satisfiability, Lower Bounds and Compression. In *Proceedings of the 41st International Symposium on Mathematical Foundations of Computer Science (MFCS)*, pp. 82:1–82:16, 2016.

[89] S. TAMAKI. A satisfiability algorithm for depth 2 threshold circuits with a sub-quadratic number of gates. *Electronic Colloquium on Computational Complexity,* TR16-100, 2016.

[90] M. THORUP. Consistent Hashing with Guaranteed Load Balancing and Hard Capacity Constraints. CoRR abs/ 1608.01350, 2016.

[91] E. VIOLA and A. WIGDERSON. Local Expanders. *Electronic Colloquium on Computational Complexity (ECCC),* Report TR16–129, pp. 1–13, 2016. Journal version in submission.

[92] R. WILLIAMS. Strong ETH Breaks With Merlin and Arthur: Short Non-Interactive Proofs of Batch Evaluation. In *Conference on Computational Complexity (CCC)*, pp. 2:1–2:17, 2016.

[93] R. WILLIAMS. Thinking Algorithmically About Impossibility. In *Leibniz International Proceedings in Informatics (LIPIcs)*, vol. 41, 2015. Invited paper.