Average-Case Fine-Grained Hardness

Marshall Ball
Columbia University
New York, NY, USA
marshall@cs.columbia.edu

Manuel Sabin
UC Berkeley
Berkeley, CA, USA
msabin@berkeley.edu

Alon Rosen
Efi Arazi School of Computer Science, IDC
Herzilya, Israel
alon.rosen@idc.ac.il

Prashant Nalini Vasudevan
CSAIL, Massachusetts Institute of Technology
Cambridge, MA, USA
prashvas@mit.edu

ABSTRACT

We present functions that can be computed in some fixed polynomial time but are hard *on average* for any algorithm that runs in slightly smaller time, assuming widely-conjectured *worst-case* hardness for problems from the study of fine-grained complexity. Unconditional constructions of such functions are known from before (Goldmann et al., IPL '94), but these have been canonical functions that have not found further use, while our functions are closely related to well-studied problems and have considerable algebraic structure.

Based on the average-case hardness and structural properties of our functions, we outline the construction of a *Proof of Work* scheme and discuss possible approaches to constructing fine-grained One-Way Functions. We also show how our reductions make conjectures regarding the worst-case hardness of the problems we reduce from (and consequently the Strong Exponential Time Hypothesis) *heuristically falsifiable* in a sense similar to that of (Naor, CRYPTO '03).

We prove our hardness results in each case by showing fine-grained reductions from solving one of three problems – namely, Orthogonal Vectors (OV), 3SUM, and All-Pairs Shortest Paths (APSP) – in the worst case to computing our function correctly on a uniformly random input. The conjectured hardness of OV and 3SUM then gives us functions that require $n^{2-o(1)}$ time to compute on average, and that of APSP gives us a function that requires $n^{3-o(1)}$ time. Using the same techniques we also obtain a conditional average-case time hierarchy of functions.

CCS CONCEPTS

• **Theory of computation** → **Problems, reductions and completeness; Cryptographic primitives; Complexity classes; Complexity theory and logic;**

KEYWORDS

Average-Case Complexity, Fine-Grained Complexity, Cryptography, Heuristic Falsifiability, Proofs of Work, Worst-Case to Average-Case Reduction

ACM Reference format:

Marshall Ball, Alon Rosen, Manuel Sabin, and Prashant Nalini Vasudevan. 2017. Average-Case Fine-Grained Hardness. In *Proceedings of 49th Annual ACM SIGACT Symposium on the Theory of Computing, Montreal, Canada, June 2017 (STOC'17)*, 14 pages.
DOI: 10.1145/3055399.3055466

1 INTRODUCTION

Since the 1970s we have had a notion of what we consider “easy” and what we consider “hard.” Polynomial-time computable has long been synonymous to efficient and easy, while showing a problem NP-complete was to condemn it as intractable. In our recent history, however, this categorization has been called into question: SAT instances, the flagship of NP-complete problems, are solved on the daily [13], while algorithms that run in as little as quadratic time may be prohibitively expensive for some practical problems such as DNA sequencing, due to large input sizes.

Thus, in the “real world,” our notions of easy and hard may not always align with our classical views. The main problem here is our choice of analysis. For SAT, we classify it as “hard” when it often may be more appropriately classified as “easy” because complexity theory typically employs *worst-case* analysis. That is, we may be adhering to an overly-pessimistic metric, when, in practice, the SAT instances we come across may be much more benign. In part to combat this sort of problem, average-case complexity was introduced in [43]. By considering distributions over problem instances, we can at least hope to argue about the performance of heuristic algorithms in practice.

Similarly, the *practical hardness* of a problem with quadratic time complexity is invisible to our typical “coarse-grained” analysis that only distinguishes between polynomial and not polynomial. Within the past decade, the field of fine-grained complexity has quickly developed [55], mapping out (conditional) hardness of natural problems *within P*. By introducing fine-grained reductions, a picture is emerging of a few main islands amongst the web of reductions, giving us an increasingly clearer classification of the relative hardness of fine-grained problems. Through such reductions, the more exact practical hardness of problems, such as DNA sequencing’s quadratic time barrier [10], has been given evidence for.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

STOC'17, Montreal, Canada

© 2017 ACM. 978-1-4503-4528-6/17/06...\$15.00

DOI: 10.1145/3055399.3055466

However, while average-case analysis and fine-grained analysis independently address issues in classical complexity theory, average-case analysis is still coarse-grained and fine-grained analysis is still worst-case. A more complete theory attempting to capture the notion of “complexity” in our world should begin by marrying average-case and fine-grained analysis.

In this paper we do so by providing average-case fine-grained hardness conjectures and show them to follow from widely conjectured worst-case assumptions on well-studied fine-grained problems. Alternatively viewed, we give new routes for the falsifications of these worst-case conjectures.

1.1 Our Results

We present worst-case-to-average-case fine-grained reductions from the three main islands of fine-grained complexity theory. We recall these three problems here to frame our work, and their relevance is discussed in Section 2.

- *Orthogonal Vectors*: The OV problem on vectors of dimension d (denoted OV_d) is to determine, given two sets U, V of n vectors from $\{0, 1\}^{d(n)}$ each, whether there exist $u \in U$ and $v \in V$ such that $\langle u, v \rangle = 0$ (over $\mathbb{Z}$). (If left unspecified, d is to be taken to be $\lceil \log^2 n \rceil$.)
- *3SUM*: The 3SUM problem is to determine whether a given set $S \subset \{-n^3, \dots, n^3\}$ of size n contains three distinct elements a, b, c such that $a + b = c$.
- *All-Pairs Shortest Path*: Given an edge-weighted (undirected or directed) graph on n vertices, the APSP problem is to find the distances between every pair of vertices, where edge weights are in $\{1, \dots, n^c\}$ for some sufficiently large c .

We give a family of polynomials over finite fields corresponding to each of these, called $\mathcal{F}\text{OV}$, $\mathcal{F}\text{3SUM}$, and $\mathcal{F}\text{ZWT}$ respectively, and conjecture these polynomials to be hard to evaluate on uniformly chosen inputs. To support these conjectures we prove worst-case-to-average-case fine-grained reductions from OV, 3SUM, and APSP to their respective families of polynomials (where 3SUM reduces also to $\mathcal{F}\text{ZWT}$). Specifically, we show:

- If OV requires $n^{2-o(1)}$ time to decide in the worst-case, then $\mathcal{F}\text{OV}$ requires $n^{2-o(1)}$ time to evaluate with probability $3/4$ on uniformly chosen inputs.
- If 3SUM requires $n^{2-o(1)}$ time to decide in the worst-case, then $\mathcal{F}\text{3SUM}$ requires $n^{2-o(1)}$ time to evaluate with probability $3/4$ on uniformly chosen inputs.
- If APSP requires $n^{3-o(1)}$ time or 3SUM requires $n^{2-o(1)}$ time to decide in the worst-case, then $\mathcal{F}\text{ZWT}$ requires $n^{3-o(1)}$ time to evaluate with probability $3/4$ on uniformly chosen inputs.

Further, we conjecture a fourth family of polynomials, $\mathcal{F}\text{TC}$, to also be average-case hard and support this with fine-grained reductions from 3SUM, and APSP, and k -SAT. The reduction from k -SAT makes $\mathcal{F}\text{TC}$ hard under the *Strong Exponential Time Hypothesis* (SETH), which states that there is no $\epsilon > 0$ such that k -SAT can be solved in time $\tilde{O}(2^{n(1-\epsilon)})$ for all values of k .

- If either APSP requires $n^{3-o(1)}$ time, 3SUM requires $n^{2-o(1)}$ time, or SETH holds, then $\mathcal{F}\text{TC}$ requires $n^{3-o(1)}$ time to evaluate with probability $3/4$ on uniformly chosen inputs.

We note that SETH implies that OV requires $n^{2-o(1)}$ time to decide in the worst-case and so $\mathcal{F}\text{OV}$ is also hard on average under the stronger assumption of SETH. Thus, $\mathcal{F}\text{OV}$ and $\mathcal{F}\text{TC}$ are our mostly strongly supported average-case hardness results. $\mathcal{F}\text{TC}$ only can become easy if SETH breaks and both 3SUM and APSP become easy, while $\mathcal{F}\text{OV}$, even with a broken SETH, remains hard unless *all* first-order graph problems become easy since [29] shows that all such problems reduce to (moderate-dimension) OV.

Our results crucially rely on the fact that the polynomials in $\mathcal{F}\text{OV}$, $\mathcal{F}\text{3SUM}$, $\mathcal{F}\text{ZWT}$, and $\mathcal{F}\text{TC}$ have degree $\text{polylog}(n)$, which is very low. This extremely low degree enables us to invoke in a *fine-grained way* the classic random self-reducibility of evaluating low-degree polynomials, first used to show the average-case hardness of computing the Permanent when assuming its worst-case hardness [26, 44], or more generally to show local correctability of Reed-Muller codes [30].

Beyond low degree, our polynomials are efficient to evaluate in time that *tightly matches* their conjectured average-case hardness. These two properties are what distinguishes our polynomials from naively representing a decision problem with a multilinear extension, which has “low” degree n and, having exponentially many terms, can take exponential time to compute, both of which are too large for our fine-grained analysis. A key technique we use to bypass the naïve construction, then, is to look at the structure of specific problems from fine-grained complexity to tailor very low-degree efficiently-computable polynomials to them. The matching upper and lower bounds are precisely what allows us to capture the complexity of our problems in the fine-grained setting and open the door for applications.

Extensions. We extend our results to a generalization of OV, called k -OV whose hardness can also be based on the SETH. To this end, we define a corresponding polynomial $\mathcal{F}\text{OV}^k$ which is computable in $\tilde{O}(n^k)$ time in the worst-case. Using the same ideas as in the above worst-case to average-case reductions we show:

- If k -OV requires $n^{k-o(1)}$ time to decide in the worst-case, then $\mathcal{F}\text{OV}^k$ requires $n^{k-o(1)}$ time to evaluate with probability $3/4$ on uniformly chosen inputs.

We note that this yields a tight average-case time hierarchy: an average-case problem computable in time n^k but not much faster for every integer k (these can be extended to rational numbers through standard padding techniques). Unconditional average-case time hierarchies are known – e.g. [31] – but these are based on canonical functions that have not found further use than as a proof of this concept, while our functions are closely related to well-studied problems and have considerable algebraic structure.

In the full version of this paper, we, building on [18], use local list decoding to show that our families of polynomials remain hard even when asking for their successful evaluation on just a $1/\text{polylog}(n)$ fraction of inputs. We extend this to attain a smooth trade-off between the running time and the upper bound on the probability of success of any average-case solver. In subsequent work, [12] extends this to prove a *direct sum theorem* for some of our polynomials.

In the full version, we additionally show that $\mathcal{F}OV$ remains hard to evaluate even over very small field sizes by applying an isolation lemma to OV , which may be of independent interest itself.

Applications. Leveraging the structure of our problems, and using ideas from [54], we construct a *Proof of Work* scheme based on their average-case hardness. In subsequent work, [12] extends our Proofs of Work framework to introduce Proofs of *Useful Work*. That is, Proofs of Work whose completion is useful for solving practical problems. We pose the application of creating fine-grained cryptography as an open problem, and outline a structural barrier to achieving fine-grained one-way functions from our results.

Finally, we note that these reductions set up a win-win scenario: either the worst-case conjectures are true and we have average-case fine-grained hard problems *or* we can show them easy and thus break our worst-case conjectures, allowing a breakthrough in fine-grained complexity theory. We explore this notion further by considering the ideas introduced in [48] of falsifiable assumptions to show that our results make the OV , $3SUM$, and $APSP$ conjectures falsifiable in a practical sense. Specifically, in Section 6.2 we show that *empirically* evaluating our polynomial for OV faster than we conjecture possible would give strong heuristic evidence that SAT has faster worst-case algorithms – i.e. that the $SETH$ is false. In this sense, we discuss how our results allow for the *heuristic falsifiability* of conjectures.

1.2 Related Work

Recently, and independently of our work, a sequence of papers [14, 33, 54] has also observed that a number of problems from the fine-grained world can be captured by the evaluation of efficiently computable low-degree polynomials. The focus of these papers has been on using the algebraic structure and low-degree of the polynomials to *delegate computation* of fine-grained problems in quickly verifiable ways. The papers were not, however, concerned with the hardness aspects of complexity and made no average-case claims or guarantees.

A key discovery, then, achieved here and independently in [14, 33, 54], is the utility of looking at the structure of specific computational problems to tailor *very* low-degree polynomials that are efficiently computable to them. From the algorithmic perspective, [14, 33, 54] find utility for delegation of computation, while, from the hardness perspective, we connect it to average-case complexity and applications thereof.

This gives a very rich framework that our results are applicable to, as our worst-case to average-case reductions can be adapted in a straightforward way to work for any problem appropriately expressible as a low-degree polynomial. Many other low-degree polynomials for interesting practical problems are found in [14, 33, 54], with [54] independently discovering our $\mathcal{F}OV$ polynomial and [14] independently finding a polynomial similar to our $\mathcal{F}3SUM$. The work of [33], in fact, identifies a natural class of “locally characterizable sets” that contains problems admitting low-degree polynomials akin to the ones considered here. Many of the above polynomials can fit into the framework of applications of average-case fine-grained hardness, such as the Proofs of Work we introduce.

In subsequent work, [12] extends our Proofs of Work framework to obtain Proofs of *Useful Work*, which can be viewed as combining

the delegation of computation from [14, 33, 54] with the hardness results we achieve here. This allows for the ability to delegate arbitrary instances of practical problems yet guarantee that the response time will require some (tunable) fixed polynomial amount of time (even for easy instances). That is, Proofs of Work whose completion is useful for solving practical problems. Further shown in [12] is that the evaluation of our polynomials (and thus the Proofs of Work) are *non-batchable* over multiple instances. Namely, they prove a *direct sum theorem* for some of our problems.

Lastly, [34] recently shows that fine-grained problems related to DNA sequencing are actually *easy* when given a batch of correlated instances. While these correlated instances are not typically what we consider in average-case complexity, they are distributional notions of inputs on fine-grained problems and so seems to be the closest existing work to average-case fine-grained complexity. The techniques used, however, are very different and focused on attaining *easiness* for specific problems with respect to specific distributions, whereas we focus on attaining hardness and applications of hardness and do so within the emerging low-degree polynomial framework. However, [34] can also be used to make claims similar to our notion of *heuristic falsifiability*, suggesting that whenever the intersection of average-case complexity and fine-grained complexity is considered it may immediately bear interesting fruit for this notion. We discuss this further in Section 6.2.

1.3 Organization

For reference, our paper’s results are discussed as follows:

- Worst-case-to-average-case fine-grained reductions to the problem of evaluating of our families of polynomials $\mathcal{F}OV$, $\mathcal{F}3SUM$, $\mathcal{F}ZWT$, and $\mathcal{F}TC$ in Section 3.1, Appendix C, Section 3.2, and Section 3.3, respectively.
- An infinite average-case time hierarchy arising from the assumption of $SETH$ and two other hierarchies from the assumption of the k -SUM conjecture as discussed in Section 4.
- Definitions of fine-grained cryptography and one-way functions, a barrier to achieving them from our results, and a possible way to bypass this barrier all in Section 5.
- Constructions of Proofs of Work guaranteed by worst-case hardness assumptions in Section 6.1.
- Applications of our average-case results and Proofs of Work to the heuristic falsifiability of worst-case assumptions, including $SETH$, in Section 6.2.
- Open problems in average-case fine-grained hardness, fine-grained cryptography, and heuristically falsifying assumptions in practice in Section 7.

2 WORST-CASE CONJECTURES

We consider a few well-studied problems in fine-grained complexity and conjectures about their worst-case hardness that we use to support our average-case hardness conjectures. For a more comprehensive survey of fine-grained complexity, connections between problems, and formal definitions of concepts like fine-grained reductions, see [55].

(All our discussion will be in the Word RAM model of computation with $O(\log(n))$ -bit words. When we speak of randomized algorithms in a worst-case setting, we mean algorithms that, for

every input, output the correct answer with probability at least $2/3$. And unless specified otherwise, all algorithms and conjectures about algorithms are randomized throughout the paper.)

2.1 Main Islands of Fine-Grained Complexity

First we recall the problems of OV, 3SUM, and APSP defined in Section 1.1. These problems currently remain the three key problems of fine-grained complexity as they partition what is known in the field. That is, there are no known reductions between them, but they reduce to many other problems and, thus, give us the basis for what we generally call hardness within P [55]. This foundation is more formally given, after extensive attempts to find improved algorithms for them, through the following popular hardness conjectures:

- **OV Conjecture:** For any $d = \omega(\log n)$, any algorithm for OV_d requires $n^{2-o(1)}$ time.
- **3SUM Conjecture:** Any algorithm for 3SUM requires $n^{2-o(1)}$ time.
- **APSP Conjecture:** Any algorithm for APSP requires $n^{3-o(1)}$ time.

These conjectures are not only important because they help stratify P, but the truth or falsity of each of them has many ramifications to practical problems.

It has been shown that if the the OV conjecture is true, then many string processing problems, hugely relevant to DNA sequencing and data comparison, also have hardness bounds (typically sub-quadratic) [3, 6, 10, 17]. On the other hand, if a sub-quadratic algorithm for OV is found, Williams [53] gave a reduction to show we would achieve improved algorithms for SAT; more specifically, the well-known Strong Exponential Time Hypothesis (SETH) would break. Thus, as stated in Section 1.1, SETH implies the OV conjecture.

(We note that the results in this paper still go through under a slightly weaker variant of the OV conjecture: for all $\epsilon > 0$, there is no $O(n^{2-\epsilon} \text{poly}(d))$ algorithm for OV_d . This problem, “Moderate Dimension” OV, was shown to be hard for the class of all first-order graph problems in [29].)

Similarly, if the 3SUM conjecture is true, problems in computational geometry [28] and exact weighted subgraph problems [5] are hard. Further, [5] shows that if the 3SUM conjecture is false, we get improved algorithms for many of the same graph problems.

Finally, the APSP conjecture’s truth would give lower bounds for many problems in dense graphs [56] and for dynamic problems [51]. [56] also shows that the conjecture being false gives better algorithms for the dense graph problems.

We note that while it is common to make these conjectures only for deterministic algorithms, we see these as structural beliefs about the problems – that brute force is essentially necessary as there is no structure to algorithmically exploit – and so there is no reason to believe that allowing randomness will allow a significant speed-up. Indeed, these conjecture are often made against randomized expected running time machines as in [55] and randomized one-sided error versions of SETH have been made in [24] and [19]. Further, [22] conjectures and argues for a SETH under randomized two-sided error machines (as we apply to all of our conjectures for the use of worst-case to average-case reductions).

Besides these three main islands of fine-grained complexity theory, a fourth seems to be emerging based on the k -CLIQUE problem [1]. With the current best algorithm solving the problem in $n^{\omega k/3}$ time [49], where ω is the matrix multiplication constant, there has been recent work showing that conjecturing this to be optimal leads to interesting hardness results for other important problems such as parsing languages and RNA folding [2, 9, 11, 16]. To explore delegating computation, [14, 33, 54] all introduce different families of polynomials to express the k -CLIQUE problem, yet none yield analysis akin to those above. The polynomials either yield average-case hardness via our techniques but cannot be computed efficiently enough to give matching upper bounds [33, 54], or they have too large of a degree for our worst-case to average-case reductions [14]. We leave the open problem of finding a family of polynomials to represent k -CLIQUE that both are computable in time $n^{\omega k/3}$ and have degree $n^{o(1)}$.

2.2 Auxiliary Problems

For these practical connections, any reduction to or from the main island problems has interesting consequences (see [55] for a more comprehensive treatment). In our results we achieve reductions *from* these problems and, to help facilitate that, we recall two more problems.

- **Zero-Weight Triangle:** Given an edge-weighted graph on n vertices, the ZWT problem is to decide whether there exists a triangle with edge weights w_1, w_2, w_3 such that $w_1 + w_2 = -w_3$, where edge weights are in $\{-n^c, \dots, n^c\}$ for some sufficiently large c .
- **Triangle-Collection:** Given an graph on n vertices and a partition C of the vertices into colors, the Triangle-Collection problem is to decide whether for each triple of three colors $a, b, c \in C$, there exists vertices x, y, z in the graph that form a triangle and $x \in a, y \in b$, and $z \in c$. That is, each triplet of colors is ‘collected’ by some triangle.

We will follow the approach in [20] of, at times, using ZWT as a proxy for both 3SUM and APSP. That is, both reduce in a fine-grained way to ZWT: APSP reduces to (Negative Weight Triangle [56] and then to) ZWT [52], and 3SUM has a randomized (which suits our purposes) reduction to ZWT in [50, 52]. Thus reducing *from* ZWT reduces from both 3SUM and APSP simultaneously. It then follows that if *either* the 3SUM or APSP conjecture are true, then ZWT requires $n^{3-o(1)}$ time.

Similarly, the Triangle-Collection problem is introduced in [7] as a way to base hardness on the believable conjecture that *at least one of* the SETH, 3SUM, or APSP conjectures are true. To do this, they give fine-grained reductions from *all three* of k -SAT, 3SUM, and APSP so that if any of their conjectures are true, then Triangle-Collection requires $n^{3-o(1)}$ time.

In general, it is better to reduce from problems furthest down a chain of reductions, as assuming those problems to be hard will then be the weakest assumption required – for example, assuming Triangle-Collection requires $n^{3-o(1)}$ time is a *weaker* assumption than assuming that at least one of k -SAT, 3SUM, or APSP are hard. It is an interesting direction to base average-case fine-grained hardness on increasingly weaker assumptions.

For this reason, it would be desirable to reduce from some very practical DNA sequencing problems (e.g. EDIT-DISTANCE and LCS) that are reduced to *from* OV (and thus *k*-SAT). Further, there is mounting evidence that, regardless of the status of *k*-SAT's complexity, these DNA sequencing problems are in fact *very* likely to be hard [4, 29]. We remark, however, that there is a barrier to representing these problems with low-degree polynomials [1]. Namely, representing them with low-degree polynomials would allow for small speedups – i.e. by using the polynomial method [21] – but such speedups (of just shaving some logarithmic factors off of the runtime) have been show to imply new breakthroughs in circuit lower bounds [4].

3 AVERAGE-CASE FINE-GRAINED HARDNESS

We now define the notion of average-case complexity that we shall use and describe the technique we use for our worst-case to average-case reductions. Then, we describe the problems we conjecture to be hard on average and show reductions from the worst-case problems described in Section 2 in support of these conjectures.

Definition 3.1. A family of functions $\mathcal{F} = \{f_n\}$ is *computable in time t on average* if there is an algorithm that runs in $t(n)$ time on the domain of f_n and, for all large enough n , computes f_n correctly with probability at least $3/4$ over the uniform distribution of inputs in its domain.

For broader definitions that are more useful when one is concerned with whole classes of problems rather than a handful of specific ones, and for extensive discussions of the merits of the same, we refer the reader to Bogdanov and Trevisan's survey [15].

To achieve average-case hardness for our fine-grained problems, our main technique will be to “express” these problems as low-degree polynomials and then use the random self-reducibility of evaluating these polynomials to attain average-case hard problems.

We now recall the classic random self-reducibility of evaluating low-degree polynomials, first used to show the average-case hardness of computing the Permanent when assuming its worst-case hardness [26, 44]. We can more generally view this as the local correctability of Reed-Muller codes first shown by Gemmell and Sudan [30] and get better error rates using techniques from this perspective. We repeat the proof in Appendix A in order to accurately assess the running time of the algorithm involved.

LEMMA 3.2. Consider positive integers N, D , and p , and an $\varepsilon \in (0, 1/3)$ such that $D > 9$, p is prime and $p > 12D$. Suppose that for some polynomial $f : \mathbb{F}_p^N \rightarrow \mathbb{F}_p$ of degree D , there is an algorithm A running in time t such that A is an average-case solver. That is,

$$\Pr_{x \leftarrow \mathbb{F}_p^N} [A(x) = f(x)] \geq 1 - \varepsilon$$

Then, there is a randomized algorithm B that runs in time at most $O(ND^2 \log^2 p + D^3 + tD)$ such that B is a probabilistic worst-case solver. That is, for any $x \in \mathbb{F}_p^N$:

$$\Pr [B(x) = f(x)] \geq \frac{2}{3}$$

REMARK 1. The range of ε being $(0, 1/3)$ is arbitrary to some extent. It could be any constant smaller than $1/2$ at the cost of p having to be slightly larger.

REMARK 2. An important thing to note here is how B 's runtime depends on f 's degree D . Assuming tD is the high-order term in the runtime, B runs in time $O(tD)$. So if we want our reductions to have low overhead, we will need D to be rather small. For our fine-grained purposes, we need to be careful in what we consider “low” and we will see that we always have D polylogarithmic in N .

We now introduce three families of polynomials that we conjecture average-case hard to evaluate and then give evidence for this by reducing to them from the worst-case problems OV, ZWT, and Triangle-Collection, respectively, and then applying the random self-reducibility of low-degree polynomials as just described. A fourth family of polynomials arising from 3SUM can be seen in Appendix C. The landscape of these reductions is seen in Figure 1.

Figure 1: Arrows represent (fine-grained) reductions and dashed means they're randomized. Thus a dashed self-loop is a worst-case to average-case self-reduction. Our work introduces $\mathcal{F}OV$, $\mathcal{F}3SUM$, $\mathcal{F}ZWT$, and $\mathcal{F}TC$ and the reductions involving them. See Appendix C for C3SUM and $\mathcal{F}3SUM$.

3.1 Orthogonal Vectors

For any n , let $p(n)$ be the smallest prime number larger than n^2 , and $d(n) = \lceil \log^2 n \rceil$ (for brevity, we shall write just p and d). We define polynomials $fOV_n : \mathbb{F}_p^{2nd} \rightarrow \mathbb{F}_p$ over $2nd$ variables. We view these variables as representing the input to OV – we separate the variables into two matrices $U, V \in \mathbb{F}_p^{n \times d}$. The polynomial fOV_n is then defined as follows:

$$fOV_n(U, V) = \sum_{i, j \in [n]} \prod_{\ell \in [d]} (1 - u_{i\ell} v_{j\ell})$$

A similar polynomial was used independently by Williams [54] to construct coMA proof systems for OV with efficient verifiers. Given an OV instance $(U, V) \in \{0, 1\}^{2nd}$, $fOV_n(U, V)$ counts the number of pairs of orthogonal vectors in it – for each pair $i, j \in [n]$, the corresponding summand is 1 if $\langle u_i, v_j \rangle = 0$, and 0 otherwise (there is no modular wrap-around of the sum as $p > n^2$). Also, fOV_n has degree at most $2d$, which is rather low.

Define the family of polynomials $\mathcal{FOV} = \{fOV_n\}$. We show a worst-case to average-case reduction from OV to $\mathcal{FOV}$ that, given an algorithm that computes fOV_n well on average, decides OV on instances of length n without much overhead. This is stated as the following theorem.

THEOREM 3.1. *If $\mathcal{FOV}$ can be computed in $O(n^{1+\alpha})$ time on average for some $\alpha > 0$, then OV can be decided in $\tilde{O}(n^{1+\alpha})$ time in the worst case.*

PROOF. Suppose there were an algorithm A that ran in $O(n^{1+\alpha})$ time and computed fOV_n correctly on more than a $3/4$ fraction of inputs for all large enough n .

In order to be able to use such an average-case algorithm, however, one has to be able to write down inputs to run it on. These inputs to fOV_n are in $\mathbb{F}_p^{2nd}$, and so to work with them it is necessary to know $p = p(n)$, the smallest prime number larger than n^2 . Further, p would have to be computable from n rather efficiently for a reduction that uses A to be efficient. As the following lemma states, this turns out to be possible to do.

LEMMA 3.3 (IMPLIED BY [42]). *The smallest prime number greater than m can be computed deterministically in $\tilde{O}(m^{1/2+\alpha})$ time for any $\alpha > 0$.*

We will then use A and this lemma to decide OV as follows. Given an input $(U, V) \in \{0, 1\}^{2nd}$, first compute $p = p(n)$ – this can be done in $\tilde{O}(n^{1+\alpha})$ time by Lemma 3.3. Once p is known, A can be used along with Lemma 3.2 to compute $fOV_n(U, V)$ in $O(n(2d)^2 \log^2 p + (2d)^3 + 2dn^{1+\alpha}) = \tilde{O}(n^{1+\alpha})$ time, and this immediately indicates membership in OV as observed above. $\square$

COROLLARY 3.4. *If OV requires $n^{2-o(1)}$ time to decide, $\mathcal{FOV}$ requires $n^{2-o(1)}$ time to compute on average.*

Note that our result is then *tight* under the OV conjecture in the sense that our polynomial is computable in $\tilde{O}(n^2)$ time, but in no less (even on average) assuming sub-quadratic hardness of OV. That is, we demonstrate a problem that is quadratic-computable but sub-quadratic-hard on average. It should also be noted that our results can be adapted to the Moderate Dimension OV problem (as mentioned in Section 2) and thus an appropriately parametrized variant of $\mathcal{FOV}$ is average-case hard for the class of all first-order graph problems as defined in [29].

3.2 3SUM and All-Pairs Shortest Path

Recall from Section 2 that both 3SUM and APSP have fine-grained reductions to ZWT, and so we restrict our attention to ZWT. We now show a family of polynomials that can count Zero Weight Triangles.

For any n , let $p(n)$ denote the smallest prime number larger than n^3 and let $d = \lceil \log(2(2n^c + 1)) \rceil + 3$ (c being the constant from the definition of ZWT). We define the polynomial $fZWT_n : \mathbb{F}_p^{n^2 d} \rightarrow \mathbb{F}_p$ as taking in a set E of $n^2 d$ variables where we split them into n^2 sets, w_{ij} , of d variables each for all $i, j \in [n]$:

$$fZWT_n(E) = \sum_{i,j,k \in [n]} \prod_{\ell \in [d]} \left(1 - \left(s_\ell(w_{ij}, w_{jk}) - s_\ell(\bar{w}_{ik}, 0 \dots 01) \right)^2 \right)$$

where $s_\ell : \mathbb{F}_p^{2d} \rightarrow \mathbb{F}_p$ is the polynomial such that if $x, y \in \{0, 1\}^d$, then $s_\ell(x, y)$ equals the ℓ^{th} bit of $(x + y)$ as long as x and y represent numbers in $[-n^c, n^c]$. Such polynomials exist, have degree at most $2d$, and are computable in $O(d \log^2 p)$ time – see Appendix B. Further, $\bar{w}_{ik}$ represents the set of linear polynomials that toggle all the bits in a boolean valued w_{ij} ; so $s_\ell(\bar{w}_{ik}, 0 \dots 01)$ effectively takes the one's complement of w_{ij} and then adds 1, which is exactly the two's complement of w_{ij} .

Now, considering a graph on n vertices with edges weighted from $[-n^c, \dots, n^c]$. We use this polynomial to count zero weight triangles in it: For an edge-weight between nodes i and j we decompose the value to its bit representation in two's complement notation and now have d boolean inputs for w_{ij} . If an edge does not exist between an i and j , we similarly put the bit decomposition of the value $2n^c + 1$ into w_{ij} (note that $i = j$ is possible and we consider there to *not* be an edge for this). Conceptually, we now have weights w_{ij} corresponding to a complete graph on n vertices with the non-edges added at weight $2n^c + 1$. Note that each triangle in it is zero weight if and only if it was a zero weight triangle in the original graph. Thus, collecting these all together we have boolean input $E \in \{0, 1\}^{n^2 d}$. This reduction certainly takes sub-cubic time.

Then, given the binary representation of a ZWT instance, the ℓ^{th} term in the product above checks whether the ℓ^{th} bit of the sum of w_{ij} and w_{jk} equals that of the negation of w_{ik} . If all d bits are equal, then, and only then, the summand is 1, otherwise it is 0. So the sum counts the number of triples of distinct (i, j) , (j, k) , and (i, k) such that $w_{ij} + w_{jk} = -w_{ik}$. Also, the degree of $fZWT_n$ is at most $4d^3 = O(\log^3 n)$.

Define the family of polynomials $\mathcal{FZWT} = \{fZWT_n\}$. The following theorem can be proved identically to Theorem 3.1.

THEOREM 3.2. *If $\mathcal{FZWT}$ can be computed in $O(n^{1.5+\alpha})$ time on average for some $\alpha > 0$, then ZWT can be decided in $\tilde{O}(n^{1.5+\alpha})$ time in the worst case.*

COROLLARY 3.5. *If ZWT requires $n^{3-o(1)}$ time to decide, $\mathcal{FZWT}$ requires $n^{3-o(1)}$ time to compute on average.*

Thus, assuming the ZWT conjecture, using the fact that $fZWT_n$ has n^3 terms and each s_ℓ is computable in $O(d)$ time, we again achieve tightness where $\mathcal{FZWT}$ is cubic-computable but sub-cubic-hard. It is also worth noting that the following corollary frames our result in the more familiar problems of 3SUM and APSP.

COROLLARY 3.6. *If either 3SUM requires $n^{2-o(1)}$ time or APSP requires $n^{3-o(1)}$ time, then $\mathcal{FZWT}$ takes $n^{3-o(1)}$ time to compute on average.*

3.3 SETH, 3SUM, and All-Pairs Shortest Path

We now give our most encompassing worst-case-to-average-case result. Recall from Section 2 that if *any* of k -SAT, 3SUM, or APSP are hard then the Triangle-Collection problem is also hard [7], thus so would be any polynomial based on it. We can hence focus our attention on Triangle-Collection. More specifically, we will look at a restricted version of the problem called Triangle-Collection* shown to be equivalent to Triangle-Collection in [1, 7], whose extra structure we will use to construct low-degree polynomials.

- *Triangle-Collection**: Given an undirected, tripartite, colored graph G with n colors and $m = n \log^2 n + 2n \log^4 n$ nodes and with partitions A, B, C of the form:
  - A contains $n \log^2 n$ nodes $a_{\ell,i}$ where $i \in [n], \ell \in [\log^2 n]$ and $a_{\ell,i}$ is colored with color i .
  - B (respectively C) contains $n \log^4 n$ nodes $b_{\ell,i,x}$ (respectively $c_{\ell,i,x}$) where $i \in [n], \ell \in [\log^2 n], x \in [\log^2 n]$ and $b_{\ell,i,x}$ (respectively $c_{\ell,i,x}$) is colored with color i .
  - For each node $a_{\ell,i}$ and colors $j, k \in [n]$, there is exactly one edge from A to B of the form $(a_{\ell,i}, b_{\ell,j,x})$ and exactly one edge from A to C of the form $(a_{\ell,i}, c_{\ell,k,y})$, for some $x, y \in [\log^2 n]$.
  - A node $b_{\ell,j,x}$ can only be connected to nodes of the form $c_{\ell',k,y}$ in C . (There no edges across disparate ℓ 's.)
 For all triples of distinct colors i, j, k , is there a triangle (u, v, w) in G where u has color i , v has color j , and w has color k ?

We now give a polynomial whose evaluation would allow us to decide Triangle-Collection*. For any n , let $p(n)$ denote the smallest prime number larger than n^3 . We define the polynomial $fTC_n : \mathbb{F}_p^m \rightarrow \mathbb{F}_p$ as taking in a set E of $m = (n \log^2 n + 2n \log^4 n)^2$ variables (corresponding to entries in the adjacency matrix of an input graph to the above problem):

$$fTC_n(E) = \sum_{\substack{1 \leq i \\ i < j \\ j < k \\ k \leq n}} \prod_{\substack{\ell, x, y \in [\log^2 n] \\ \pi \in S_{\{i, j, k\}}}} (1 - \delta(\pi(i), \pi(j), \pi(k), \ell, x, y))$$

where for a set X , S_X denotes the set of permutations on X , and δ is defined as follows:

$$\delta(i, j, k, \ell, x, y) = e_{a_{\ell, \pi(i)}, b_{\ell, \pi(j), x}} e_{a_{\ell, \pi(i)}, c_{\ell, \pi(k), y}} e_{b_{\ell, \pi(j), x}, c_{\ell, \pi(k), y}}$$

Consider a tripartite graph as defined above with adjacency matrix E . For each triple of colors, $(i, j, k) \in [n]^3$, if there is a corresponding triangle in the graph then it zeroes out that particular term, otherwise it will evaluate to one. Thus, fTC_n counts the number of colors not collected by a triangle – i.e. the number of violations to being a YES instance – and so, for boolean E , $fTC_n(E) = 0$ if and only if E corresponds to a YES instance of Triangle-Collection*. Moreover, the degree of fTC_n is at most $18 \log^6 n$.

Define the family of polynomials $\mathcal{F}TC = \{fTC_n\}$. The following theorem can be proved identically to Theorem 3.1.

THEOREM 3.3. *If $\mathcal{F}TC$ can be computed in $O(n^{1.5+\alpha})$ time on average for some $\alpha > 0$, then TC can be decided in $\tilde{O}(n^{1.5+\alpha})$ time in the worst case.*

COROLLARY 3.7. *If TC requires $n^{3-o(1)}$ time to decide, $\mathcal{F}TC$ requires $n^{3-o(1)}$ time to compute on average.*

Thus, fTC_n only having n^3 many summands with each being computable in $\text{polylog}(n)$ time, it is easily seen that we again achieve tightness where $\mathcal{F}TC$ is cubic-computable but sub-cubic-hard. More recognizably we attain the following.

COROLLARY 3.8. *If either SETH holds, 3SUM takes $n^{2-o(1)}$ time, or APSP takes $n^{3-o(1)}$ time, then $\mathcal{F}TC$ takes $n^{3-o(1)}$ time to compute on average.*

Note that this does not subsume the hardness of $\mathcal{F}ZWT$ as, even if SETH fails and 3SUM and APSP become easy, the ZWT problem may still be hard and yield hardness for $\mathcal{F}ZWT$.

4 AN AVERAGE-CASE TIME HIERARCHY

In this section we present an infinite collection of generalizations of $\mathcal{F}OV$ that we conjecture form an average-case time hierarchy. That is, for every rational number k the collection contains a function $\mathcal{F}OV^k$ such that $\mathcal{F}OV^k$ is computable in $\tilde{O}(n^k)$ time, but (we conjecture) requires $n^{k-o(1)}$ time to compute even on average. This conjecture is supported by SETH. We describe these generalizations below and indicate how this follows from SETH for integer values of $k \geq 2$ and note that this can be extended to all rational numbers using standard padding techniques.

- *k-Orthogonal Vectors*: For an integer $k \geq 2$, the k -OV problem on vectors of dimension d is to determine, given k sets $(U_1, \dots, U_k)$ of n vectors from $\{0, 1\}^{d(n)}$ each, whether there exist $u_i \in U_i$ for each i such that over $\mathbb{Z}$,

$$\sum_{\ell \in [d(n)]} u_{1\ell} \cdots u_{k\ell} = 0$$

(As with OV, if left unspecified, d is to be taken to be $\lceil \log^2 n \rceil$.)

Similar to how it implies the hardness of OV, (the randomized version of) SETH also implies that for any integer $k \geq 2$, any randomized algorithm for k -OV requires $n^{k-o(1)}$ time – the proof is a natural generalization of that for OV shown in [55]. We next take the same approach we did for OV and define for any integer $k \geq 2$ a family of polynomials $\mathcal{F}OV^k = \{fOV_n^k\}$, where with p being the smallest prime number larger than n^k and $d = \lceil \log^2(n) \rceil$, $fOV_n^k : \mathbb{F}_p^{knd} \rightarrow \mathbb{F}_p$ is defined as:

$$fOV_n^k(U_1, \dots, U_k) = \sum_{u_1 \in U_1, \dots, u_k \in U_k} \prod_{\ell \in [d]} (1 - u_{1\ell} \cdots u_{k\ell})$$

Exactly as fOV_n does, when fOV_n^k 's input is a k -OV instance from $\{0, 1\}^{knd}$, $fOV_n^k(U_1, \dots, U_k)$ counts the number of sets of “orthogonal” vectors in it. Note that the degree of fOV_n^k is at most kd . And also that by simply evaluating each summand and adding them up, the polynomial can be evaluated in $\tilde{O}(n^k)$ time. The following theorem can again be proven in a manner identical to Theorem 3.1.

THEOREM 4.1. *For any integer $k \geq 2$, if $\mathcal{F}OV^k$ can be computed in $O(n^{k/2+\alpha})$ time on average for some $\alpha > 0$, then k -OV can be decided in $\tilde{O}(n^{k/2+\alpha})$ time in the worst case.*

COROLLARY 4.1. *Suppose for every $k \geq 2$, k -OV requires $n^{k-o(1)}$ time to decide. Then for every such k , $\mathcal{F}OV^k$ requires $n^{k-o(1)}$ to compute on average but can be computed in the worst case in $\tilde{O}(n^k)$ time.*

Thus, we can attain the hierarchy from an infinite number of conjectures, one for each k , but, as noted earlier, the entire hierarchy is also implied by the single assumption SETH. We note that for k -OV (and all other problems) we could naively express k -OV with a polynomial via a multilinear extension, yet this polynomial, as discussed in Section 1.2, would have exponentially many terms and degree n . Already the degree is too high for our purposes but not

by much: we may not be too disappointed with $n^{(k-1)-o(1)}$ average-case hardness that degree n would still afford us. The main problem then is that the naïve polynomial may take exponential time to compute and so the upper bound is very far from the lower bound. The tightness of our hierarchy is a key feature then in capturing the hardness of our problems as well as for use in applications such as in Section 6.1 and in [12].

REMARK 3. We can also attain two semi-tight hierarchies from generalizations of the 3SUM problem. That is, if we assume the k -SUM conjecture proposed in [5], we get hardness for two infinite hierarchies, with one based on generalizing $\mathcal{F}$ ZWT and one from generalizing a polynomial introduced in Appendix C, $\mathcal{F}$ 3SUM (the proper generalizations can be based on problems found in [5]). These hierarchies, however, are loose, in that the k -SUM conjecture gives us $(n^{\lceil k/2 \rceil - o(1)})$ hardness at the k^{th} level but, to our knowledge, our generalized polynomials are only $\tilde{O}(n^{k-1})$ computable (as they have n^{k-1} many terms).

5 TOWARDS FINE-GRAINED CRYPTOGRAPHY

Average-case hardness has frequently found use in cryptography, where it is in fact a necessity – it is almost always required that a cryptographic object be hard to defeat on average for it to be useful. In this light, it is a natural question to ask whether the worst-case to average-case reductions presented here, along with the conjectured hardness of the worst-case problems, can be used to do cryptography.

Of course, since these problems are actually computable in polynomial time, one cannot expect to use them to construct standard cryptographic objects, which have to be secure against all polynomial time adversaries. We hence consider *fine-grained* cryptography, where we only ask for security against adversaries that run in some fixed polynomial time.

A major reason for interest in such notions is that a form of cryptography might be realizable even in Impagliazzo's Pessiland (or Heuristica or even Algorithmica) [39]. That is, even if we live in a world where One-Way Functions or "coarse-grained" average-case hardness do not exist, it may still be possible to construct fine-grained cryptographic objects and salvage practical cryptography. Further, even if we *do* have standard cryptography, it may be the case that fine-grained cryptography can use weaker assumptions as it only needs to be secure against moderately powerful adversaries. Though not done so very extensively and not done from the sort fine-grained complexity in [55], such notions have indeed been considered several times before – see for instance [23, 36, 45, 46].

In this section we define a notion of fine-grained cryptography in Section 5.1, show a structural barrier to achieving it in Section 5.2, and outline as an open problem a possible approach to circumventing this barrier in Section 5.3. While this outline has not yet yielded fine-grained cryptography, we use its techniques in Section 6.1 to obtain Proofs of Work. Our framework has since been extended to achieve Proofs of Useful Work [12], which we also mention in Section 6.1

5.1 Fine-Grained One-Way Functions

To illustrate the kinds of objects, approaches, and difficulties fine-grained cryptography might involve, we consider the case of One-Way Functions (OWFs). A fine-grained OWF would capture the same concept as a standard OWF – easy to compute yet hard to invert – but with a more fine-grained interpretation of "easy" and "hard".

Definition 5.1. We say a function $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ is (t, ϵ) -one-way if it can be computed in $O(t(n)^{1-\delta})$ time for some $\delta > 0$, but for any $\delta' > 0$, any $O(t(n)^{1-\delta'})$ -time algorithm A , and all sufficiently large n ,

$$\Pr_{x \leftarrow \{0, 1\}^n} [A(f(x)) \in f^{-1}(f(x))] \leq \epsilon(n, \delta')$$

One approach to constructing such OWFs (that has perhaps been part of folklore for decades) is as follows. Take one of our hard-on-average to evaluate polynomials – e.g. $f\text{OV}_n$ – and suppose there was an input-output sampling algorithm $S \equiv (S_1, S_2)$ that runs in sub-quadratic time in n such that, on uniform input r , $S_1(r)$ is distributed uniformly over the appropriate domain, and $S_2(r) = f\text{OV}_n(S_1(r))$. By our results it can be seen that S_1 is $(n^2, 3/4)$ -one-way if we assume the OV conjecture (strengthened to assume hardness for all sufficiently large input sizes).

5.2 Barriers and NSETH

However, as stated, there turn out to be certain barriers to this approach. For instance, if $S(r) = (x, y)$, then r would be a certificate that $f\text{OV}_n(x) = y$ that can be verified in sub-quadratic time. In particular, if $x \in \{0, 1\}^{2nd}$ is a NO instance of OV, then r is a certificate for this that is verifiable in deterministic sub-quadratic time – i.e. that $f\text{OV}_n(x) = 0$.

Further, tracing this back through the reduction of k -SAT to OV (see [55]), this gives us a certificate for NO instances of k -SAT (for any k) that are verifiable in $O(2^{n(1-\epsilon)})$ time for some $\epsilon > 0$ – i.e. short and quickly verifiable certificates for CNF-UNSAT instances. Interestingly, the impossibility of this was recently conjectured and formalized as NSETH (the Non-deterministic Strong Exponential Time Hypothesis) in [20] along with the establishment of its barrier status: its falsification would yield breakthroughs in both circuit complexity and proof complexity.

An alternative view of matters would be that this presents another approach to breaking NSETH. In fact, something weaker would suffice for this purpose – one only needs a sampler that runs in sub-quadratic time and samples $(x, f\text{OV}_n^k(x))$ for some k such that the distribution of x has $\{0, 1\}^{knd}$ in its support.

Note that while a sampler based on OV would, because of its relation to k -SAT, break NSETH, a sampler based on ZWT would only yield quick deterministic certification for APSP instances (the reduction from 3SUM is randomized). So we are left with much weaker "barriers" for $\mathcal{F}$ 3SUM and $\mathcal{F}$ ZWT samplers. Indeed, [20]'s introduction of NSETH was in a direction opposite to ours: while we explore NSETH to argue that OV is unlikely to have small co-nondeterministic complexity because of its relationship to SAT, they introduce it to argue that 3SUM and APSP are unlikely to have a relationship to SAT (as OV does) by showing them to actually have

small co-nondeterministic complexities. Thus they explicitly break both “barriers” for a $\mathcal{F}3\text{SUM}$ or $\mathcal{F}Z\text{WT}$ sampler!

Still, $\mathcal{F}OV$ is much simpler algebraically and so seems to hold more hope for constructing a fine-grained OWF in this manner. We now discuss ways in which we may still salvage a sampler for $\mathcal{F}OV$.

5.3 A Way Around

There are visible ways to skirt this NSETH barrier: Suppose that the sampler S was not perfect – that with some small probability over r it outputs (x, y) such that $fOV_n(x) \neq y$. Immediately, NSETH no longer applies, as now an r such that $S(r) = (x, y)$ may no longer be a sound certificate that $fOV_n(x) = y$. And, as explained below, such a sampler still gives us a fine-grained *distributional* OWF based on the hardness of OV as long as the probability that it is wrong is small enough.

Informally, a distributional OWF is a function f such that $f(x)$ is easy to compute, but for most y 's it is hard to sample a (close to) uniformly random x such that $f(x) = y$. A distributional OWF might not be a OWF itself, but it is known how to derive a OWF from any distributional OWF [40]. And further, this transformation works with quasi-linear overhead using constructions of hash functions from [41].

We claim that S_1 is now a fine-grained distributional OWF. Intuitively, if it were not, then we would, for many x , be able to sub-quadratically sample r almost uniformly from those obeying $S_1(r) = x$. But, since the probability over r is low that S errs, the r we sampled is likely a non-erring one. That is, it is likely that we would have sub-quadratically obtained an r that gives us $S_2(r) = fOV_n(x)$ and thus we likely can sub-quadratically compute $fOV_n(x)$. So if fOV_n is actually hard to compute on average, then such a distributional inverter cannot exist.

While, as mentioned earlier, such an erring sampler would no longer give efficiently verifiable certificates for NO instances of OV , it turns out that it would still yield a “barrier” of a coAM protocol for OV with sub-quadratic verification.

First we note that we get an AM protocol with a sub-quadratic time verifier that takes input (x, y) and proves that $fOV_n(x) = y$, and is complete and sound for *most values* of x . Since the sampler is wrong with only with a very small probability over r , most values of x 's have the property that most values of r satisfying $S_1(r) = x$ also satisfy $S_2(r) = fOV_n(x)$. In the protocol with input (x, y) , the verifier simply asks the prover to prove that for most r 's such that $S_1(r) = x$, it is also the case that $S_2(r) = y$. If S_1 is indeed distributed uniformly, then this comes down to proving a lower bound on the number of r 's such that $S_1(r) = x$ and $S_2(r) = y$, and this can be done in AM using the protocol from [35], in a single round (verifier sends a message and prover responds) and with the verifier running in sub-quadratic time. With a little more analysis and appropriate setting of parameters, this protocol can be shown to work even if S_1 is only close to uniform.

Further, from this protocol one can get a protocol that works for *all values* of x by following the approach in the proof of Lemma 3.2 – given an input (x, y) , the verifier runs the random self-reduction for x and, for each evaluation query to fOV_n that comes up in its course, it asks the prover for the answer along with an AM proof

of its correctness. Thus, done all at once, this gives a single-round coAM protocol for OV with a sub-quadratic time verifier (this is when $y = 0$). This in turns leads to a single-round coAM protocol for $k\text{-SAT}$ with a verifier that runs in $\tilde{O}(2^{n(1-\epsilon)})$ time for some $\epsilon > 0$.

The impossibility of the existence of such a protocol could be conjectured as a sort of $\text{AM}[2]\text{SETH}$. Williams [54] gives a non-interactive MA protocol that comes close to breaking this conjecture, but standard approaches for converting MA protocols to AM protocols [8] seem to incur a prohibitively large overhead in our fine-grained setting.

We next detail the MA protocol from [54] to pose an open problem of efficiently converting this to an AM protocol, thus breaking the barrier to an erring sampler (and thus a fine-grained OWF) with the further hope such a sampler could be “reverse-engineered” from an AM protocol. In Section 6.1, show how this MA protocol is already useful to our framework by constructing Proofs of Work.

5.4 An MA Protocol

We briefly describe the MA protocol for fOV_n achieved in [54], where a more general protocol and framework can be found. Recall that the objective of the prover is to convince the verifier that, for given (x, y) , $fOV_n(x) = y$. We will expand x into $(U, V) \in \mathbb{F}_p^{n \times d} \times \mathbb{F}_p^{n \times d}$, which is the syntax we used when defining fOV_n as:

$$fOV_n(U, V) = \sum_{i,j \in [n]} \prod_{\ell \in [d]} (1 - u_{i\ell} v_{j\ell})$$

For a fixed V , we define the supporting polynomial $fOV_V : \mathbb{F}_p^d \rightarrow \mathbb{F}_p$ as:

$$fOV_V(x_1, \dots, x_d) = \sum_{j \in [n]} \prod_{\ell \in [d]} (1 - x_{\ell} v_{j\ell})$$

Now we may write fOV_n as:

$$fOV_n(U, V) = \sum_{i \in [n]} fOV_V(u_{i1}, \dots, u_{id})$$

Next let $\phi_1, \dots, \phi_d : \mathbb{F}_p \rightarrow \mathbb{F}_p$ be the polynomials of degree $(n-1)$ such that for $i \in [n]$, $\phi_\ell(i) = u_{i\ell}$ (treating i as an element of $\mathbb{F}_p$). Define the polynomial $R_{U,V}(x) = fOV_V(\phi_1(x), \dots, \phi_d(x))$. We can then write:

$$fOV_n(U, V) = \sum_{i \in [n]} R_{U,V}(i)$$

The degree of $R_{U,V}$ is at most $(n-1)d$. If the verifier knew the coefficients of $R_{U,V}$, then it could evaluate it on the points $\{1, \dots, n\}$ in time $\tilde{O}(nd)$ (using fast techniques for batch evaluation on univariate polynomials [27]) and thus compute $fOV_n(U, V)$. This suggests the following protocol. The prover sends over the coefficients of a polynomial R^* that it claims to be $R_{U,V}$.

To verify this claim, the verifier checks whether $R^*(x) = R_{U,V}(x)$ for a random $x \in \mathbb{F}_p$. This can be done because even though the verifier does not know the coefficients of $R_{U,V}$, it can evaluate it at an input x in time $\tilde{O}(nd)$ by first evaluating all the $\phi_\ell(x)$'s (these polynomials can be found using fast interpolation techniques for univariate polynomials [38]), and then evaluating fOV_V using

these values. By the Schwartz-Zippel lemma, since the field is considerably larger than the degree of R^* and $R_{U,V}$, if these are not the same polynomial, the verifier will catch this with high probability. If this check passes, the verifier uses R^* to compute $fOV_n(U, V)$.

We pose the open problem of using the ideas of this protocol to construct an AM protocol that does the same. Further we pose the problem of creating an erring sampler from such a protocol, or more generally constructing a fine-grained OWF. We show in Section 6.1 how to use the existing MA scheme to create Proofs of Work.

6 PROOFS OF WORK AND THE FALSIFIABILITY OF CONJECTURES

In this section we present two other applications of our results. While we focus on $\mathcal{F}OV$, these applications can similarly be made for $\mathcal{F}3SUM$, $\mathcal{F}ZWT$, and $\mathcal{F}TC$ (and $\mathcal{F}OV^k$ from Section 4).

6.1 Proofs of Work

A Proof of Work (PoW) scheme is a means for one party (that we call the *Prover*) to prove to another (called the *Challenger*) that it has expended a certain amount of computational resources. These were introduced by Dwork and Naor [25] as a means to deter spam mail and denial-of-service attacks, and have also found use in cryptocurrencies [47]. Here we formulate a simplified definition of a PoW scheme to illustrate how our results could be used in this context. A PoW scheme consists of three algorithms:

- $\text{Challenge}(1^n)$ takes a difficulty parameter n and produces a challenge c of this difficulty.
- $\text{Solve}(c)$ solves the challenge c and produces a solution s .
- $\text{Verify}(c, s)$ verifies that s is a valid solution to the challenge c .

The idea is that when the challenger wants the prover to work for a certain amount of time (and prove that it has done so), it runs Challenge with the appropriate difficulty parameter to generate a challenge c that it sends over. The prover is then required to produce a solution s to the challenge that the challenger verifies using Verify .

Intuitively, our average-case hardness results for fOV_n says that random inputs, thought of as challenges, must require a certain amount of work for a prover to evaluate fOV_n on. Then, the MA protocol in Section 5.4 gives a quick way for the challenger to verify a solution.

Definition 6.1. The triple of algorithms (Challenge , Solve , Verify) is a Proof of Work (PoW) scheme if the following properties are satisfied:

- Challenge and Verify are computable in time $s(n)$.
- Solve is computable in time $t(n)$ and $\text{Verify}(c, \text{Solve}(c)) = 1$ ($t(n) > s(n)$)
- There are constants $\epsilon, \epsilon' \in [0, 1]$ such that for any A and n satisfying:

$$\Pr_{c \leftarrow \text{Challenge}(1^n)} [\text{Verify}(c, A(c)) = 1] \geq \epsilon$$

it is the case that A takes time at least $t'(n)$ on a ϵ' fraction of challenges of difficulty n . ($t'(n) > s(n)$). Ideally, $t'(n) \approx t(n)$.)

We want Challenge and Verify to be efficiently computable so that the challenger does not have to do too much work in this

process. We want Solve to be more expensive but not by too much so that the prover can indeed produce a valid solution, albeit with more work than that done by the challenger. The lower bound on the difficulty of producing valid solutions ensures that the prover actually has to spend a certain amount of time on this task and that its production of a valid solution is proof that it has spent this time.

While above we use time as the computational resource of interest as we are working in the Word RAM model, these specifications can be in terms of other resources such as space, circuit size, etc. when working in other models. For different applications, there could also be other properties that might be desirable in such schemes such as resistance to parallelism, non-batchability, etc., but we shall ignore these for now.

The MA protocol for fOV_n along with our average-case reductions can be used to construct a PoW scheme based on the hardness of OV as follows.

- $\text{Challenge}(1^n)$ samples a random input $(U, V) \in \mathbb{F}_p^{2nd}$ to fOV_n .
- $\text{Solve}((U, V))$ outputs the coefficients of $R_{U,V}$ (as defined in Section 5.4).
- $\text{Verify}((U, V), R^*)$ checks that $R^*(x) = R_{U,V}(x)$ for a random $x \in \mathbb{F}_p$.

By the arguments in Sections 3.1 and 5.4 regarding the verifier in the MA protocol, both Challenge and Verify above can be computed in $\tilde{O}(n)$ time. Solve can be computed in $\tilde{O}(n^2)$ time either by using the explicit expression for $R_{U,V}$, or by evaluating $R_{U,V}$ on sufficiently many points using the ϕ_ℓ 's and fOV_V and interpolating to find its coefficients.

On the other hand, if any algorithm A produces R^* that passes Verify with probability, say, $5/6$ (over the uniform distribution of (U, V)), then by the soundness of Verify it follows that A is actually producing $R_{U,V}$ with almost the same probability. As noted earlier, given the coefficients of $R_{U,V}$, $fOV_n(U, V)$ can be computed in $\tilde{O}(nd)$ time. So A can be used to compute fOV_n correctly on average with an additive $\tilde{O}(nd)$ overhead.

Corollary 3.4 now implies that if OV is hard for sub-quadratic algorithms, then A cannot be sub-quadratic *and* correct on more than a $3/4$ fraction of inputs. This leaves at least a $\frac{5}{6} - \frac{3}{4} = \frac{1}{12}$ fraction of inputs where A takes $\Omega(n^{2-o(1)})$ time. These numbers can be improved by repetition and by considering the better reductions from the full version of this paper.

Thus, based on the conjectured hardness of OV , the above is a PoW scheme where a party can prove to a challenger that it has run for a certain amount of time where the challenger only has to run for about a square-root of this amount of time.

This framework has since been used to create Proofs of Useful Work [12]. By allowing the ability to delegate arbitrary instances of OV (or other fine-grained problems) while still guaranteeing that the challenges will be hard on average (even for easy instances), [12] achieves Proofs of Work from our framework whose completion is useful for solving practical problems.

REMARK 4. We note that OV is non-batchable due to downward self-reducibility. For instance, let $\ell(n) = O(n^{1-\epsilon})$ for some $\epsilon > 0$. If there is a batch evaluation algorithm that can solve $\ell^2(n)$ instances of size $n/\ell(n)$ in time at most $O(n^{2-\delta})$ for some $\delta > 0$, then OV on an input (U, V) can be solved in sub-quadratic time by partitioning the

U and V into $\ell(n)$ sets of size $n/\ell(n)$ each and running this algorithm on all pairings of these partitions together.

This has recently been extended [12] to show that fOV_n (as well as the above PoW) is non-batchable in the average-case. That is, they show a direct sum theorem for our $\mathcal{F}OV$. That is, multiple instances requires proportionally more time.

6.2 On the Heuristic Falsifiability of Conjectures

The above Proof of Work yields a win-win in the domain of algorithms and complexity. Namely, either we have a PoW or we have the existence of a *provably* sub-quadratic prover that can convince the challenger with sufficient probability which will in turn yield breakthroughs: a randomized sub-quadratic time algorithm for OV and thus a randomized $2^{n(1-\epsilon)}$ time algorithm for k -SAT for some $\epsilon > 0$. In particular, because the hardness of the PoW is over random instances, even a prover that can be *empirically* demonstrated to be sub-quadratic in practice will give *heuristic* evidence that the conjectured hardness of OV or k -SAT is false. In other words, if the above PoW is empirically (after working out the constants hidden by Big-Oh notation) insecure, then (randomized) SETH is heuristically falsified.

This notion of heuristic falsifiability sits directly at the marriage of average-case and fine-grained complexity that we accomplish: Without average-case results, a worst-case conjecture could not be broken without a full proof that an algorithm works on *all* inputs, and, without fine-grained results, a conjecture on large time complexity like SETH could not be feasibly be broken in practice. Thus, it is *precisely* average-case fine-grained hardness that allows us to discuss the heuristic falsifiability of conjectures. While theory and practice can often influence each other indirectly, this marks an interesting connection, akin to the hard-sciences, in which empirical evidence can give concrete and parameterizable theoretical evidence.

REMARK 5. We note that while some may try to claim that SETH is already being falsified in practice - e.g. that we might seem to run in $2^{\sqrt{n}}$ time on practical inputs - there are two main points in which this is different than our heuristic falsifiability. One point is that, from our worst-case to average-case reductions, our notion would be heuristically breaking the worst-case version of SETH and achieve a complexity theoretic claim, as opposed to heuristically breaking an average-case notion of SETH on "nature's distribution" of "practical" inputs, which only says a heuristic claim on how we perform in practice. Secondly, claims of breaking even such an average-case notion of SETH in practice cannot be given too much confidence to since the input sizes must remain very small to be feasible and so not many "data points" can be gathered to see the true shape of the exponential curve. In contrast, our heuristic falsifiability reduces to a quadratic time problem, so that many more "data points" can be gathered for runtimes on much larger input sizes, giving us much more confidence as to if we heuristically break OV's conjecture and thus SETH.

To put this observation in more concrete terms, we consider [48] in which Moni Naor introduced a stronger notion of falsifiable assumptions: Informally, an assumption A is efficiently falsifiable if

there is an efficiently samplable distribution of challenges and an efficient verifier of solutions to these challenges such that, if and only if A is false (we consider the "only if" case), there is an efficient algorithm which causes the verifier to accept with high probability over the challenge distribution.

In our case, any algorithm that solves $\mathcal{F}OV$ in sub-quadratic time falsifies the conjectured hardness of $\mathcal{F}OV$ and thus OV and thus SETH. The sampler simply uniformly draws a set of inputs for fOV_n , and the verifier simply evaluates fOV_n on the instances and compares with the sub-quadratic prover's answer.

Further, the problem underlying the PoW (to output a polynomial for a given fOV_n instance) is falsifiable with the added property that both the sampler *and* the verifier run in sub-quadratic time. Thus to heuristically falsify SETH, a challenger may deploy PoW challenges out into the world and, if they're often prematurely solved, we gather empirical evidence against SETH. Note that this can similarly be done for 3SUM and APSP as their polynomials can also be used for PoWs.

We note that interesting applications of heuristic falsifiability may be inherent to the intersection of average-case complexity and fine-grained complexity: One of the only other works we are aware of that might be considered average-case fine-grained analysis, [34], immediately yields results that notions of heuristic falsifiability can apply to. That is, [34] shows that fine-grained problems related to DNA sequencing are actually *easy* when given a batch of correlated instances; thus, this analysis is average-case over a specific distribution of *correlated instances* for a fine-grained problem. This distribution for which easiness is achieved, however, is motivated to be "realistic" by the correlation of the instances attempting to match current evolutionary theory and how mutations occur within a phylogenetic tree. Thus, if a distribution over correlated instances matches well a theory of evolution *yet* [34]'s algorithm consistently under-performs on real-life data, this may suggest our current theory of evolution is *wrong*. Again, we can (efficiently) test a hypothesis in a concretely parameterizable way and gather evidence against it.

We find it interesting that combining average-case and fine-grained complexity seems to almost immediately bear interesting fruit in the context of heuristic falsifiability. We pose it as an open question to find more ways in which the "scientific method" can be introduced into the highly theoretical field of complexity theory so that conjectures can be *tested* to give concrete parameters for our confidences for them.

7 OPEN PROBLEMS

To our knowledge, our work is the first to attain average-case results in the young field of conditional fine-grained complexity as described in [55], and there are many ripe questions to ask of average-case fine-grained complexity. Amongst several questions regarding average-case algorithms for fine-grained problems, fine-grained cryptography, and the heuristic falsifiability of worst-case assumptions are the following:

- (1) Can other low-degree polynomials be found for other interesting problems in the fine-grained world to match conjectured worst-case lower bounds on them - e.g. can k -CLIQUE reduce

to evaluating a family of polynomials that are both computable in time $n^{\omega k/3}$ and have degree $n^{o(1)}$?

- (2) Are there samplable distributions over which well-studied problems like 3SUM, OV, etc. are hard on average? Towards this, is it possible to reduce any of our polynomial families back to the problems they came from or to problems further down their reduction chain?
- (3) Are there any other worst-case to average-case reduction techniques that might be interesting in the fine-grained world?
- (4) Classically, generic derandomization from worst-case hardness assumptions proceeds in two steps: Achieving average-case hardness from worst-case hardness, and achieving strong pseudorandom generators from average-case hardness. In this paper we achieve the first step in a fine-grained way. Can pseudorandom generators be achieved in a fine-grained way from our average-case hardness?
- (5) Is it possible to construct a single-round coAM protocol for OV (or for $\mathcal{F}$ OV) with a sub-quadratic time verifier (or for k -SAT with a $2^{(1-\epsilon)n}$ -time verifier)?
- (6) Can we construct a fine-grained OWF from worst-case hardness assumptions? Is it possible to realize it from an AM protocol for $\mathcal{F}$ OV as discussed in Section 5.3?
- (7) Standard OWFs are sufficient for secret-key cryptography as they are equivalent to Pseudorandom Generators and Pseudorandom Functions [32, 37]. What similar equivalences hold in the fine-grained world and what fine-grained cryptography can be accomplished from fine-grained OWFs? More generally, what standard cryptographic results translate over to fine-grained cryptography and are there any that hold only in the fine-grained world?
- (8) Can we heuristically falsify any of the three big worst-case fine-grained conjectures (from Section 2)? Are there other ways in which we can develop techniques for practice to influence theory and give concrete and parameterizable theoretical evidence?

A EVALUATING LOW DEGREE POLYNOMIALS

Here we recall and prove Lemma 3.2.

LEMMA 3.2. *Consider positive integers N , D , and p , and an $\epsilon \in (0, 1/3)$ such that $D > 9$, p is prime and $p > 12D$. Suppose that for some polynomial $f : \mathbb{F}_p^N \rightarrow \mathbb{F}_p$ of degree D , there is an algorithm A running in time t such that:*

$$\Pr_{\mathbf{x} \leftarrow \mathbb{F}_p^N} [A(\mathbf{x}) = f(\mathbf{x})] \geq 1 - \epsilon$$

Then there is a randomized algorithm B that runs in time at most $O(ND^2 \log^2 p + D^3 + tD)$ such that for any $\mathbf{x} \in \mathbb{F}_p^N$:

$$\Pr [B(\mathbf{x}) = f(\mathbf{x})] \geq \frac{2}{3}$$

PROOF. The algorithm B works as follows on input $\mathbf{x}$:

- (1) Draw two random points $\mathbf{y}_1, \mathbf{y}_2 \in \mathbb{F}_p^N$ and define the curve $\mathbf{c}(w) = \mathbf{x} + w\mathbf{y}_1 + w^2\mathbf{y}_2$ for $w \in \mathbb{F}_p$.
- (2) For a value of $m (< p)$ to be determined later, compute $(A(\mathbf{c}(1)), \dots, A(\mathbf{c}(m)))$ to get $(z_1, \dots, z_m) \in \mathbb{F}_p$.

- (3) Run Berlekamp-Welch on $(z_1, \dots, z_m)$. If it succeeds and outputs a polynomial $\hat{g}$, output $\hat{g}(0)$. Otherwise output 0.

To see why the above algorithm works, define the polynomial $g(w) = f(\mathbf{c}(w))$. g is a polynomial of degree $2D$ over the single variable w , with the property that $g(0) = f(\mathbf{x})$. So if we had $(2D+1)$ evaluations of g at different points in $\mathbb{F}_p$, we could retrieve g and compute $f(\mathbf{x})$. While we may not be able to obtain these evaluations directly (since they involve computing f), we do have access to A , which promises to be correct about the value of f on a random point with probability $2/3$.

So we replace the values $g(w) = f(\mathbf{c}(w))$ with $A(\mathbf{c}(w))$, which will hopefully be correct at several points. Now our problem is to retrieve g given a set of its alleged evaluations at m points, some of which may be wrong.

We do this by interpreting $(g(1), g(2), \dots, g(m))$ as the Reed-Solomon encoding of g and running its decoding algorithm on $(A(\mathbf{c}(1)), \dots, A(\mathbf{c}(m)))$, which now corresponds to a corrupt codeword. The Berlekamp-Welch algorithm can do this as long as less than $(m - 2D)/2$ of the m values of $A(\mathbf{c}(w))$ are wrong. We now bound the probability of too many of these values being wrong.

Let Q_w be an indicator variable such that $Q_w = 1$ if and only if $A(\mathbf{c}(w)) \neq f(\mathbf{c}(w))$. Let $Q = \sum_{w=1}^m Q_w$. We note at this point the fact that over the randomness of $\mathbf{y}_1$ and $\mathbf{y}_2$, the distributions of $\mathbf{c}(w)$ and $\mathbf{c}(w')$ for any two distinct $w, w' \in \mathbb{F}_p$ are uniform and independent. This gives us the following statistics:

$$\begin{aligned} \mathbb{E}[Q] &= \epsilon m \\ \text{Var}[Q] &= m\epsilon(1 - \epsilon) \end{aligned}$$

Thus, by the Chebyshev inequality, we have that the probability that more than a $\delta (> \epsilon)$ fraction of $A(\mathbf{c}(w))$'s disagree with $f(\mathbf{c}(w))$ is:

$$\begin{aligned} \Pr [Q > \delta m] &\leq \Pr [|Q - \epsilon m| > (\delta - \epsilon)m] \\ &\leq \frac{\epsilon(1 - \epsilon)}{(\delta - \epsilon)^2 m} \leq \frac{1}{4(\delta - \epsilon)^2 m} \end{aligned}$$

We are interested in $\delta = \frac{m-2D}{2m} = \frac{1}{2} - \frac{D}{m}$. So if we set, for instance, $m = 12D$, the bound on the probability above is at most $1/3$ if $D > 9$ and $\epsilon < 1/3$.

So except with this small probability, the decoding algorithm correctly recovers g as $\hat{g}$, and consequently B computes $f(\mathbf{x})$ correctly.

Generating each $\mathbf{c}(w)$ and running A on it takes $O(ND \log^2 p + t)$ time. The Berlekamp-Welch algorithm then takes $O(m^3)$ time, and the final evaluation of $\hat{g}$ takes $O(D \log^2 p)$. Hence, given A with the above properties, the running time of B is:

$$O(m(ND \log^2 p + t) + m^3 + D \log^2 p) = O(ND^2 \log^2 p + D^3 + tD)$$

□

B POLYNOMIALS COMPUTING SUMS

In this section we write down polynomials that compute the bits of the sum of a pair of numbers that are given to them in bits. Without loss of generality, we will represent numbers in the two's complement form. In both cases where such representations are required (3SUM and ZWT), there are apriori bounds on the sizes of the numbers that come up – these are either numbers in the input or

sums of pairs of these numbers. So we can assume that we always have enough bits to be able to represent these numbers. Under this assumption, addition in the two's complement representation is the same as adding unsigned numbers in the standard place-value representation (and ignoring the final carry). So we will present the polynomials $\{s_\ell\}_{\ell \in [d]}$ that correspond to unsigned addition (and are easier to describe) and these are the polynomials that will be used.

We label the bits of a d -bit number from 1 to d starting from the least significant bit. We then translate the semantics of the ripple-carry adder into polynomials. The polynomial $s_1 : \mathbb{F}_p^{2d} \rightarrow \mathbb{F}_p$ corresponding to the first bit of the sum is:

$$s_1(x, y) = x_1(1 - y_1) + (1 - x_1)y_1$$

The carry from this operation is given by the following polynomial:

$$c_1(x, y) = x_1y_1$$

For every other ℓ , this pair of polynomials can be computed from earlier polynomials and the inputs as follows (hiding the arguments x and y for convenience):

$$\begin{aligned} s_\ell &= (1 - x_\ell)(1 - y_\ell)c_{\ell-1} + (1 - x_\ell)y_\ell(1 - c_{\ell-1}) \\ &\quad + x_\ell(1 - y_\ell)(1 - c_{\ell-1}) + x_\ell y_\ell c_{\ell-1} \\ c_\ell &= x_\ell y_\ell(1 - c_{\ell-1}) + x_\ell(1 - y_\ell)c_{\ell-1} \\ &\quad + (1 - x_\ell)y_\ell c_{\ell-1} + x_\ell y_\ell c_{\ell-1} \end{aligned}$$

It can now be seen that $\deg(s_\ell) = \deg(c_\ell) = \deg(c_{\ell-1}) + 2$. Along with the fact that $\deg(c_1) = 2$, this implies that $\deg(s_\ell) = 2\ell \leq 2d$.

These polynomials can also be computed very easily by evaluating them in order. Given $c_{\ell-1}$, both s_ℓ and c_ℓ take only a constant number of operations to compute. Hence all the s_ℓ 's can be computed in $O(d \log^2 p)$ time.

C CONVOLUTION-3SUM

Here we give another average-case fine-grained hard problem based on the hardness of 3SUM. While Section 3.2 already has a polynomial whose average-case hardness is based on the 3SUM conjecture, we include this one for completeness as it is possible that either is independently hard even if the other is shown to be easy. We first recall the CONVOLUTION-3SUM (C3SUM) problem introduced by [50] where it was shown that 3SUM has a (randomized) fine-grained reduction to C3SUM, thus allowing us to restrict our attention to it.

- C3SUM: Determine whether, when given three n -element arrays, A , B , and C , with entries in $\{-n^3, \dots, n^3\}$, there exist $i, j \in [n]$ such that $A[i] + B[j] = C[i + j]$.

We now define a family of polynomials that can count solutions to a C3SUM instance (when given in binary) and refer the reader to Section 3.2 for all notation and discussion due to its similarity to $\mathcal{F}$ 3SUM.

For any n , let $p(n)$ denote the smallest prime number larger than n^2 and let $d = \lceil 3 \log n \rceil + 3$. We define the polynomial $f3SUM_n : \mathbb{F}_p^{3nd} \rightarrow \mathbb{F}_p$ as taking in sets A , B , and C of nd variables each and where we split each set into n groups of d variables – e.g. $A[i]$ is the i^{th} group of d variables of the nd variables in A .

$$f3SUM_n(A, B, C) = \sum_{i, j \in [n]} \prod_{\ell \in [d]} (1 - (s_\ell(A[i], B[j]) - C[i + j]_\ell)^2)$$

From the same arguments in Section 3.2, we get the following theorem for $\mathcal{F}$ 3SUM = $\{f3SUM_n\}$.

THEOREM C.1. *If $\mathcal{F}$ 3SUM can be computed in $O(n^{1+\alpha})$ time on average for some $\alpha > 0$, then 3SUM can be decided in $\tilde{O}(n^{1+\alpha})$ time in the worst case.*

COROLLARY C.1. *If 3SUM requires $n^{2-o(1)}$ time to decide, $\mathcal{F}$ 3SUM requires $n^{2-o(1)}$ time to compute on average.*

Note that if we did not use C3SUM, we would have had n^3 terms from a more naïve construction from 3SUM and thus a gap between $\mathcal{F}$ 3SUM's computability and its hardness. But, with our current construction having only n^2 terms, we achieve tightness where $\mathcal{F}$ 3SUM is quadratic-computable but sub-quadratic-hard.

ACKNOWLEDGEMENTS

We would like to thank Amir Abboud for suggesting a barrier to achieving low-degree polynomials for EDIT-DISTANCE (Section 2.2) and for pointing us to the Triangle-Collection problem in [7] and a polynomial for it, Shafi Goldwasser for pointing out [54] to us, Tim Roughgarden for clarifications on Remark 5, and Vinod Vaikuntanathan for comments relating to the material in Section 5.3 and Open Question 5. Thanks also to Andrej Bogdanov, Irit Dinur, Ramprasad Satharishi, Eli-Ben Sasson, Amir Shpilka, and Amnon Ta-Shma for useful discussions.

The bulk of this work was performed while the authors were at IDC Herzliya's FACT center and supported by NSF-BSF Cyber Security and Privacy grant #2014/632, ISF grant #1255/12, and by the ERC under the EU's Seventh Framework Programme (FP/2007-2013) ERC Grant Agreement #07952. Marshall Ball is supported in part by the Defense Advanced Research Project Agency (DARPA) and Army Research Office (ARO) under Contract #W911NF-15-C-0236, and NSF grants #CNS-1445424 and #CCF-1423306. Manuel Sabin is also supported by the National Science Foundation Graduate Research Fellowship under Grant #DGE-1106400. Prashant Nalini Vasudevan is also supported by the IBM Thomas J. Watson Research Center (Agreement #4915012803).

REFERENCES

- [1] Amir Abboud. 2017. Personal communication. (2017).
- [2] Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. 2015. If the current clique algorithms are optimal, so is Valiant's parser. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*. IEEE, 98–117.
- [3] Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. 2015. Quadratic-Time Hardness of LCS and other Sequence Similarity Measures. *CoRR* abs/1501.07053 (2015). <http://arxiv.org/abs/1501.07053>
- [4] Amir Abboud, Thomas Dueholm Hansen, Virginia Vassilevska Williams, and Ryan Williams. 2015. Simulating branching programs with edit distance and friends or: A polylog shaved is a lower bound made. *arXiv preprint arXiv:1511.06022* (2015).
- [5] Amir Abboud and Kevin Lewi. 2013. Exact weight subgraphs and the k-sum conjecture. In *International Colloquium on Automata, Languages, and Programming*. Springer Berlin Heidelberg, 1–12.
- [6] Amir Abboud, Virginia Vassilevska Williams, and Oren Weimann. 2014. Consequences of faster alignment of sequences. In *International Colloquium on Automata, Languages, and Programming*. Springer, 39–51.

- [7] Amir Abboud, Virginia Vassilevska Williams, and Huacheng Yu. 2015. Matching Triangles and Basing Hardness on an Extremely Popular Conjecture. Manuscript: <https://dl.dropboxusercontent.com/u/14999836/publications/MatchTriangles.pdf>. (2015).
- [8] László Babai. 1985. Trading Group Theory for Randomness. In *Proceedings of the 17th Annual ACM Symposium on Theory of Computing, May 6-8, 1985, Providence, Rhode Island, USA*, Robert Sedgwick (Ed.). ACM, 421–429. DOI: <https://doi.org/10.1145/22145.22192>
- [9] Arturs Backurs, Nishanth Dikkala, and Christos Tzamos. 2016. Tight Hardness Results for Maximum Weight Rectangles. *arXiv preprint arXiv:1602.05837* (2016).
- [10] Arturs Backurs and Piotr Indyk. 2014. Edit Distance Cannot Be Computed in Strongly Subquadratic Time (unless SETH is false). CoRR abs/1412.0348 (2014). <http://arxiv.org/abs/1412.0348>
- [11] Arturs Backurs and Christos Tzamos. 2016. Improving Viterbi is Hard: Better Runtimes Imply Faster Clique Algorithms. *arXiv preprint arXiv:1607.04229* (2016).
- [12] Marshall Ball, Alon Rosen, Manuel Sabin, and Prashant Nalini Vasudevan. 2017. Proofs of Useful Work. In submission. (2017).
- [13] Armin Biere, Marijn Heule, and Hans van Maaren. 2009. *Handbook of satisfiability*. Vol. 185. ios press. <http://gauss.eces.uc.edu/Courses/c626/notes/history.pdf>
- [14] Andreas Björklund and Petteri Kaski. 2016. How proofs are prepared at Camelot. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*. ACM, 391–400.
- [15] Andrej Bogdanov and Luca Trevisan. 2006. Average-Case Complexity. *Foundations and Trends in Theoretical Computer Science* 2, 1 (2006). DOI: <https://doi.org/10.1561/04000000004>
- [16] Karl Bringmann, Allan Grönlund, and Kasper Green Larsen. 2016. A Dichotomy for Regular Expression Membership Testing. *arXiv preprint arXiv:1611.00918* (2016).
- [17] Karl Bringmann and Marvin Künnemann. 2015. Quadratic conditional lower bounds for string problems and dynamic time warping. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*. IEEE, 79–97.
- [18] Jin-yi Cai, Aduri Pavan, and D. Sivakumar. 1999. On the Hardness of Permanent. In *STACS 99, 16th Annual Symposium on Theoretical Aspects of Computer Science, Trier, Germany, March 4-6, 1999, Proceedings (Lecture Notes in Computer Science)*, Christoph Meinel and Sophie Tison (Eds.), Vol. 1563. Springer, 90–99. DOI: https://doi.org/10.1007/3-540-49116-3_8
- [19] Chris Calabro, Russell Impagliazzo, Valentine Kabanets, and Ramamohan Paturi. 2003. The Complexity of Unique k-SAT: An Isolation Lemma for k-CNFs. In *18th Annual IEEE Conference on Computational Complexity (Complexity 2003), 7-10 July 2003, Aarhus, Denmark*. 135. DOI: <https://doi.org/10.1109/CCC.2003.1214416>
- [20] Marco L. Carmosino, Jiawei Gao, Russell Impagliazzo, Ivan Mihajlin, Ramamohan Paturi, and Stefan Schneider. 2016. Nondeterministic Extensions of the Strong Exponential Time Hypothesis and Consequences for Non-reducibility. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, Cambridge, MA, USA, January 14-16, 2016*. 261–270. DOI: <https://doi.org/10.1145/2840728.2840746>
- [21] Timothy M Chan and Ryan Williams. 2016. Deterministic apsp, orthogonal vectors, and more: Quickly derandomizing razborov-smolensky. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, 1246–1255.
- [22] M. Cygan, F.V. Fomin, L. Kowalik, D. Lokshantov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. 2015. *Parameterized Algorithms*. Springer International Publishing. 469–470 pages. <https://books.google.co.il/books?id=Frg0CgAAQBAJ>
- [23] Akshay Degwekar, Vinod Vaikuntanathan, and Prashant Nalini Vasudevan. 2016. Fine-Grained Cryptography. In *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part III*. 533–562. DOI: https://doi.org/10.1007/978-3-662-53015-3_19
- [24] Holger Dell, Thore Husfeldt, and Martin Wahlén. 2010. Exponential time complexity of the permanent and the Tutte polynomial. In *International Colloquium on Automata, Languages, and Programming*. Springer, 426–437.
- [25] Cynthia Dwork and Moni Naor. 1992. Pricing via Processing or Combatting Junk Mail. In *Advances in Cryptology - CRYPTO '92, 12th Annual International Cryptology Conference, Santa Barbara, California, USA, August 16-20, 1992, Proceedings*. 139–147. DOI: https://doi.org/10.1007/3-540-48071-4_10
- [26] J Feigenbaum and L Fortnow. 1990. On the Random-Self-Reducibility of Complete Sets. *University of Chicago Technical Report* (1990), 90–22.
- [27] Charles M. Fiduccia. 1972. Polynomial Evaluation via the Division Algorithm: The Fast Fourier Transform Revisited. In *Proceedings of the 4th Annual ACM Symposium on Theory of Computing, May 1-3, 1972, Denver, Colorado, USA*, Patrick C. Fischer, H. Paul Zeiger, Jeffrey D. Ullman, and Arnold L. Rosenberg (Eds.). ACM, 88–93. DOI: <https://doi.org/10.1145/800152.804900>
- [28] Anka Gajentaan and Mark H Overmars. 1995. On a class of $O(n^2)$ problems in computational geometry. *Computational geometry* 5, 3 (1995), 165–185.
- [29] Jiawei Gao and Russell Impagliazzo. 2016. Orthogonal Vectors is hard for first-order properties on sparse graphs. *Electronic Colloquium on Computational Complexity (ECCC)* 23 (2016), 53. <http://eccc.hpi-web.de/report/2016/053>
- [30] Peter Gemmell and Madhu Sudan. 1992. Highly resilient correctors for polynomials. *Information processing letters* 43, 4 (1992), 169–174.
- [31] Mikael Goldmann, Per Grape, and Johan Håstad. 1994. On Average Time Hierarchies. *Inf. Process. Lett.* 49, 1 (1994), 15–20. DOI: [https://doi.org/10.1016/0020-0190\(94\)90048-5](https://doi.org/10.1016/0020-0190(94)90048-5)
- [32] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. 1984. How to Construct Random Functions (Extended Abstract). In *25th Annual Symposium on Foundations of Computer Science, West Palm Beach, Florida, USA, 24-26 October 1984*. 464–479. DOI: <https://doi.org/10.1109/SFCS.1984.715949>
- [33] Oded Goldreich and Guy Rothblum. February 2017. Simple Doubly-Efficient Interactive Proof Systems for Locally-Characterizable Sets. *Electronic Colloquium on Computational Complexity Report TR17-018*. (February 2017).
- [34] Shafi Goldwasser and Dhiraj Holden. 2016. On the Fine Grained Complexity of Polynomial Time Problems Given Correlated Instances. In *Innovations in Theoretical Computer Science (ITCS)*.
- [35] Shafi Goldwasser and Michael Sipser. 1986. Private coins versus public coins in interactive proof systems. In *Proceedings of the eighteenth annual ACM symposium on Theory of computing*. ACM, 59–68.
- [36] Johan Håstad. 1987. One-way permutations in NC^0 . *Inform. Process. Lett.* 26, 3 (1987), 153–155.
- [37] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. 1999. A Pseudorandom Generator from any One-way Function. *SIAM J. Comput.* 28, 4 (1999), 1364–1396. DOI: <https://doi.org/10.1137/S0097539793244708>
- [38] Ellis Horowitz. 1972. A fast method for interpolation using preconditioning. *Inform. Process. Lett.* 1, 4 (1972), 157–163.
- [39] R. Impagliazzo. 1995. A Personal View of Average-case Complexity. In *Proceedings of the 10th Annual Structure in Complexity Theory Conference (SCT'95) (SCT '95)*. IEEE Computer Society, Washington, DC, USA, 134–. <http://dl.acm.org/citation.cfm?id=829497.829786>
- [40] Russell Impagliazzo and Michael Luby. 1989. One-way Functions are Essential for Complexity Based Cryptography (Extended Abstract). In *30th Annual Symposium on Foundations of Computer Science, Research Triangle Park, North Carolina, USA, 30 October - 1 November 1989*. 230–235. DOI: <https://doi.org/10.1109/SFCS.1989.63483>
- [41] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. 2008. Cryptography with constant computational overhead. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*. ACM, 433–442.
- [42] Jeffrey C Lagarias and Andrew M. Odlyzko. 1987. Computing $\pi(x)$: An analytic method. *Journal of Algorithms* 8, 2 (1987), 173–191.
- [43] Leonid A. Levin. 1986. Average Case Complete Problems. *SIAM J. Comput.* 15, 1 (1986), 285–286. DOI: <https://doi.org/10.1137/0215020>
- [44] Richard Lipton. 1991. New directions in testing. *Distributed Computing and Cryptography* 2 (1991), 191–202.
- [45] Ueli M Maurer. 1992. Conditionally-perfect secrecy and a provably-secure randomized cipher. *Journal of Cryptology* 5, 1 (1992), 53–66.
- [46] Ralph C. Merkle. 1978. Secure Communications Over Insecure Channels. *Commun. ACM* 21, 4 (1978), 294–299. DOI: <https://doi.org/10.1145/359460.359473>
- [47] Satoshi Nakamoto. 2008. Bitcoin: A peer-to-peer electronic cash system. (2008).
- [48] Moni Naor. 2003. On cryptographic assumptions and challenges. In *Annual International Cryptology Conference*. Springer, 96–109.
- [49] Jaroslav Nešetřil and Svatopluk Poljak. 1985. On the complexity of the subgraph problem. *Commentationes Mathematicae Universitatis Carolinae* 26, 2 (1985), 415–419.
- [50] Mihai Pătraşcu. 2010. Towards Polynomial Lower Bounds for Dynamic Problems. In *Proceedings of the Forty-second ACM Symposium on Theory of Computing (STOC '10)*. ACM, New York, NY, USA, 603–610. DOI: <https://doi.org/10.1145/1806689.1806772>
- [51] Liam Roditty and Uri Zwick. 2004. On dynamic shortest paths problems. In *European Symposium on Algorithms*. Springer, 580–591.
- [52] Virginia Vassilevska and Ryan Williams. 2009. Finding, minimizing, and counting weighted subgraphs. In *In Proceedings of the Forty-First Annual ACM Symposium on the Theory of Computing*. 455–464.
- [53] Ryan Williams. 2005. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theor. Comput. Sci.* 348, 2-3 (2005), 357–365. DOI: <https://doi.org/10.1016/j.tcs.2005.09.023>
- [54] Ryan Williams. 2016. Strong ETH Breaks With Merlin and Arthur: Short Non-Interactive Proofs of Batch Evaluation. In *31st Conference on Computational Complexity, CCC 2016, May 29 to June 1, 2016, Tokyo, Japan*. 2:1–2:17. DOI: <https://doi.org/10.4230/LIPIcs.CCC.2016.2>
- [55] Virginia Vassilevska Williams. 2015. Hardness of easy problems: Basing hardness on popular conjectures such as the strong exponential time hypothesis. In *Proc. International Symposium on Parameterized and Exact Computation*. 16–28. <http://theory.stanford.edu/~virgi/ipcc-survey.pdf>
- [56] Virginia Vassilevska Williams and Ryan Williams. 2010. Subcubic Equivalences between Path, Matrix and Triangle Problems. *2013 IEEE 54th Annual Symposium on Foundations of Computer Science* 00, undefined (2010), 645–654. DOI: <https://doi.org/doi.ieeeecomputersociety.org/10.1109/FOCS.2010.67>