

Monochromatic Triangles, Triangle Listing and APSP

Virginia Vassilevska Williams* and Yinzhan Xu†

CSAIL

MIT

Cambridge, USA

{virgi, xyzhan}@mit.edu

Abstract—All-Pairs Shortest Paths (APSP) is one of the most basic problems in graph algorithms. Given an n -node directed or undirected graph with integer weights in $\{-n^c, \dots, n^c\}$ and no negative cycles, APSP asks to compute the shortest paths distance between every pair of vertices. The fastest known algorithm for APSP runs in $n^3/2^{\Theta(\sqrt{\log n})}$ time [Williams’14], and no truly subcubic time algorithms are known.

One of the main hypotheses in fine-grained complexity is that APSP requires $n^{3-o(1)}$ time. Another famous hypothesis in fine-grained complexity is that the 3SUM problem for n integers (which can be solved in $O(n^2)$ time) requires $n^{2-o(1)}$ time. Although there are no direct reductions between 3SUM and APSP, it is known that they are related: the $(\min, +)$ -convolution problem reduces in a fine-grained way to both, and both fine-grained reduce to the Exact Triangle problem.

In this paper we find more relationships between these two problems and other basic problems. Pătraşcu had shown that under the 3SUM hypothesis the All-Edges Sparse Triangle problem in m -edge graphs requires $m^{4/3-o(1)}$ time. The latter problem asks to determine for every edge e , whether e is in a triangle. It is equivalent to the problem of listing m triangles in an m -edge graph where $m = \tilde{O}(n^{1.5})$, and can be solved in $O(m^{1.41})$ time [Alon et al.’97] with the current matrix multiplication bounds, and in $\tilde{O}(m^{4/3})$ time if $\omega = 2$.

We show that one can reduce Exact Triangle to All-Edges Sparse Triangle, showing that All-Edges Sparse Triangle (and hence Triangle Listing) requires $m^{4/3-o(1)}$ time also assuming the APSP hypothesis. This allows us to provide APSP-hardness for many dynamic problems that were previously known to be hard under the 3SUM hypothesis.

We also consider the All-Edges Monochromatic Triangle problem. Via work of [Lincoln et al.’20], our result on All-Edges Sparse Triangle implies that if the All-Edges Monochromatic Triangle problem has an $O(n^{2.5-\epsilon})$ time algorithm for $\epsilon > 0$, then both the APSP and 3SUM hypotheses are false. The fastest algorithm for All-Edges Monochromatic Triangle runs in $\tilde{O}(n^{(3+\omega)/2})$ time [Vassilevska et al.’06], and our new reduction shows that if $\omega = 2$, this algorithm is best possible, unless 3SUM or APSP can be solved faster. Besides 3SUM, previously the only problems known to be fine-grained reducible to All-Edges Monochromatic Triangle were the seemingly easier problems directed unweighted APSP and Min-Witness Product [Lincoln et al.’20]. Our reduction shows that this problem is much harder. We also connect the problem to other “intermediate” problems, whose runtimes are between $O(n^\omega)$ and $O(n^3)$, such as the Max-Min product problem.

*Supported by an NSF CAREER Award, NSF Grants CCF-1740519 and CCF-1909429, a BSF Grant BSF:2012338, a Google Research Fellowship and a Sloan Research Fellowship.

†Supported by an NSF CAREER Award and NSF Grant CCF-1740519.

Keywords—fine-grained complexity; all-pairs shortest paths; triangle detection;

I. INTRODUCTION

All-Pairs Shortest Paths (APSP) is one of the most fundamental problems in graph algorithms. In one of the most general variants of the problem, one is given an n -node directed or undirected graph with integer weights in $\{-n^c, \dots, n^c\}$ with no negative cycles and one needs to compute the shortest paths distance between every pair of vertices. A central question in graph algorithms is how fast APSP in n node graphs can be solved. The fastest known algorithm runs in $n^3/2^{\Theta(\sqrt{\log n})}$ time [1, 2], and no $O(n^{3-\epsilon})$ time for $\epsilon > 0$, so called *truly subcubic* algorithms are known.

One of the main hypotheses in fine-grained complexity, the APSP hypothesis, is that APSP requires $n^{3-o(1)}$ time¹. Another famous hypothesis is the 3SUM hypothesis that the 3SUM problem for n integers in $\{-n^c, \dots, n^c\}$ requires $n^{2-o(1)}$ time. 3SUM asks if three of the integers sum to 0 and can be solved in $O(n^2)$ time. Although there are no direct reductions between 3SUM and APSP, it is known that they are related, in a sense much more related than the third core problem of fine-grained complexity Orthogonal Vectors (OV). For instance, the $(\min, +)$ -convolution problem is known to be fine-grained reducible to both 3SUM and APSP, so that if APSP is in truly subcubic time, or 3SUM is in truly subquadratic time, then $(\min, +)$ -convolution is in truly subquadratic time [3, 4, 5, 6, 7, 8]. Further, the Exact Triangle problem is known to be harder than both APSP and 3SUM [5, 8], so that if it has a truly subcubic time algorithm, then both the APSP and the 3SUM hypotheses would be false. Meanwhile, it is not known how 3SUM and APSP (or $(\min, +)$ -convolution and Exact Triangle) are related to OV^2 .

In this paper we provide more relationships to other problems that 3SUM and APSP have in common. Pătraşcu [4] showed that 3SUM can be reduced in truly subquadratic

¹All of the hypotheses are for the Word-RAM model of computation with $O(\log n)$ bit words.

²The Min-Weight k -Clique problem can be reduced to both APSP and *moderate-dimension* OV [9], but it is not known whether it can be reduced to OV.

time to the All-Edges Sparse Triangle problem of determining for every edge e in an n -node, $m = \tilde{O}(n^{1.5})$ -edge³ graph, whether e is in a triangle. Pătraşcu’s reduction implies that under the 3SUM hypothesis, All-Edges Sparse Triangle requires $m^{4/3-o(1)}$ time. All-Edges Sparse Triangle is known to be equivalent to the problem of listing m triangles in an m -edge graph where $m = \tilde{O}(n^{1.5})$ [10], and can be solved in $O(m^{1.41})$ time [11] with the current matrix multiplication bounds [12, 13], and in $\tilde{O}(m^{4/3})$ time if the exponent of square matrix multiplication ω is 2.

Our main result is a reduction from Exact Triangle to All-Edges Sparse Triangle, thus showing that All-Edges Sparse Triangle (and hence Triangle Listing) requires $m^{4/3-o(1)}$ time also assuming the APSP hypothesis. This allows us to provide APSP-hardness for many dynamic problems.

We also consider the All-Edges Monochromatic Triangle problem (AE-Mono Δ) [14, 15, 16] in which one is given an n -node graph G with colors on the edges, and one is asked to return for every edge e , whether it appears in a monochromatic triangle in G . The fastest algorithm for AE-Mono Δ runs in $\tilde{O}(n^{(3+\omega)/2})$ time [14, 15]. Via work of [16], our reduction from Exact Triangle to All-Edges Sparse Triangle implies that if AE-Mono Δ has an $O(n^{2.5-\epsilon})$ time algorithm for $\epsilon > 0$, then both the APSP and 3SUM hypotheses are false. This shows that if $\omega = 2$, the known algorithm for AE-Mono Δ is best possible, unless 3SUM and APSP can both be solved faster.

Lincoln et al. [16] showed that 3SUM is in fact fine-grained equivalent to the *Monochromatic Convolution* problem, which is the convolution version of AE-Mono Δ . These two latter problems are very related (e.g. the known algorithms for them are analogous, the only difference being the use of FFT vs Fast Matrix Multiplication), to the extent that one might conjecture that they are fine-grained equivalent. If this bold conjecture were true, then 3SUM would be equivalent to both problems, and thus APSP would reduce to 3SUM. We leave determining the veracity of this conjecture to future work.

Besides 3SUM, previously, the only problems known to be fine-grained reducible to AE-Mono Δ were the seemingly easier problems directed unweighted APSP and Min-Witness Product [16]. Our new reduction shows that this problem is much harder. We also connect the problem to multiple so-called “intermediate” problems, whose runtimes are between $O(n^\omega)$ and $O(n^3)$, such as the Max-Min Product problem that has been studied in relation to All Pairs Bottleneck Paths [17, 18].

A. Our results

Here we give more details about the results summarized above.

³The \tilde{O} notation in this paper hides poly-logarithmic factors.

Reductions from Exact Triangle.: The Exact Triangle problem asks, given an n -node graph with integer edge weights in $\{-n^c, \dots, n^c\}$ for some constant c and a target T , whether there are three vertices p, q, r so that $w(p, q) + w(q, r) + w(r, p) = T$. Exact Triangle is equivalent to the version of Exact Triangle in which the target T is 0 [5]. This is called the Zero Triangle problem (also known as Zero-3-Clique or Love Triangle).

The brute-force algorithm for Exact Triangle runs in $O(n^3)$ time. Meanwhile, as mentioned earlier, an $O(n^{3-\epsilon})$ time algorithm for $\epsilon > 0$ for Exact Triangle would violate both the APSP hypothesis and the 3SUM hypothesis. The *Exact Triangle hypothesis* states that Exact Triangle requires $n^{3-o(1)}$ time in the word-RAM model of computation with $O(\log n)$ bit words. This hypothesis is at least as believable as the hypothesis that at least one of the 3SUM and APSP hypotheses is true.

Our main result is a reduction from Exact Triangle to certain unbalanced versions of Triangle Listing, and All-Edges Triangle Listing in sparse graphs which can then easily be reduced to other problems, including All-Edges Sparse Triangle, AE-Mono Δ .

In our unbalanced triangle problems there are five parameters, $\alpha, \beta, \gamma, \rho, t$. One is given a tripartite graph where the three parts have $n^\alpha, n^\beta, n^\gamma$ vertices, respectively. The fourth parameter ρ controls the edge density of the graph. Roughly speaking, each vertex in the graph has an $n^{-\rho}$ fraction of vertices as its neighbors. In the parameterized version of All-Edges Triangle Listing one is asked to list for every edge e , t triangles containing e (or all triangles containing e if there are fewer than t triangles). In our parameterized version of Triangle Listing, one is asked to list t triangles in the graph (or all triangles if there are fewer than t triangle).

The statements of our reductions to the parameterized versions of Triangle Listing and All-Edges Triangle Listing are a bit technical (see Theorem 3.4 in Section III). We will list some consequences.

Corollary 1.1: Any n -node instance of Exact Triangle can be reduced in $\tilde{O}(n^{2.5})$ time to $\tilde{O}(n)$ instances of All-Edges Sparse Triangle on graphs with $O(n^{1.5})$ edges each. Thus, assuming the Exact Triangle hypothesis, there is no $O(m^{4/3-\epsilon})$ time algorithm for All-Edges Sparse Triangle for any $\epsilon > 0$.

Lincoln et al. [16] showed that computing a certain number of independent instances of All-Edges Sparse Triangle is equivalent to AE-Mono Δ .

Theorem 1.2 ([16]): Computing $m^{1/3}$ independent instances of All-Edges Sparse Triangle with m edges each, where the number of vertices in each instance is $m^{2/3}$ is equivalent up to poly-logarithmic factors to computing an AE-Mono Δ instance where the number of vertices is $n = O(m^{2/3})$.

Combining Corollary 1.1 and Theorem 1.2, we get the following conditional lower bound for AE-Mono Δ .

Corollary 1.3: Assuming the Exact Triangle hypothesis, there is no $O(n^{2.5-\epsilon})$ time algorithm for AE-Mono Δ on n -node graphs for any $\epsilon > 0$.

We also achieve the following conditional lower bound for even sparser triangle detection problems, as another corollary of our main reduction.

Corollary 1.4: Let \mathbb{A} be an algorithm for All-Edges Sparse Triangle for n -node graphs where every node has degree at most $d = n^\delta$ for some $0 < \delta \leq 0.5$. Assuming the Exact Triangle hypothesis, \mathbb{A} cannot run in $O(n^{1-\epsilon}d^2)$ time for any $\epsilon > 0$.

Many 3SUM-hard problems are now also APSP-hard.: Kopelowitz, Pettie and Porat [19] considered the SetDisjointness and SetIntersection problems. They showed 3SUM hardness for both SetDisjointness and SetIntersection, which were in turn used to show 3SUM hardness for many dynamic graph problems.

The SetDisjointness problem can be viewed as an All-Edges Sparse Triangle problem in constrained graphs, and the SetIntersection problem can be viewed as a triangle listing problem. Using our reductions from Exact Triangle to All-Edges Sparse Triangle and Triangle Listing, we show that the SetDisjointness and SetIntersection problems are hard under the Exact Triangle hypothesis. Thus, we immediately get APSP-hardness for a variety of problems that were previously known to be 3SUM-hard, including Dynamic Maximum Cardinality Matching, Incremental Maximum Cardinality Matching, d -Failure Connectivity and Triangle Enumeration in graphs with particular arboricity [19]. One example result is the following.

Corollary 1.5: Assume the Exact Triangle hypothesis (or any one of the 3SUM hypothesis or the APSP hypothesis). For any constants $y \in (0, 1/2)$, $x \in (0, 2y]$, there exists a constant $\epsilon > 0$ so that for graphs with n vertices, m edges, arboricity $\alpha = \Theta(n^x) = \Theta(m^y)$, and $t < m\alpha^{1-\epsilon}$ triangles, it requires $\Omega(m\alpha^{1-o(1)})$ time to list all triangles in the graph.

Notably, the SetIntersection problem of Kopelowitz et al. is a more generalized version of the All-Edges Sparse Triangle problem considered by Pătraşcu [4] and later by Abboud and Williams [20]. Therefore, all conditional lower bound results in [4] and [20] that use All-Edges Sparse Triangle as an intermediate problem also hold under the Exact Triangle hypothesis. These results include (but are not limited to) dynamic st reachability, dynamic st shortest paths, dynamic strong connectivity, subgraph connectivity, Langerman’s problem, Pagh’s problem and Erickson’s problem. An example result from All-Edges Sparse Triangle is the following.

Corollary 1.6: Assume the Exact Triangle hypothesis (or any one of the 3SUM hypothesis or the APSP hypothesis). Then there is no fully dynamic algorithm for s - t reachability that can have $O(m^{4/3-\epsilon})$ processing time, $O(m^{a-\epsilon})$ update time, and $O(m^{2/3-a-\epsilon})$ query time for some $\epsilon > 0$ and $1/6 \leq a \leq 1/3$.

Readers interested in reductions from triangle problems to dynamic graphs problems can check [4, 19, 20] for more details.

Reductions for intermediate problems.: The notion of “intermediate” problems were first devised by Lincoln et al. [16] who studied problems whose current best running times are $\tilde{O}(n^{2.5})$ if $\omega = 2$. With the current bounds on rectangular matrix multiplication [21], the running time of these problems vary. The easiest “intermediate” problems seem to be the unweighted directed APSP [22] problem and the Min-Witness Product problem [23], whose best algorithms run in $\tilde{O}(n^{2.5286})$ time using the best bounds on rectangular matrix multiplication [21]. Following them in complexity, there are the Equality Product problem [24, 25] and the Dominance Product problem [26, 27], which are known to have $\tilde{O}(n^{2.6598})$ time algorithms.

The Equality Product of two $n \times n$ integer matrices A and B is the $n \times n$ matrix C with $C[i, j] = |\{k \mid A[i, k] = B[k, j]\}|$. The Boolean version of the problem, \exists Equality Product asks to determine whether $|\{k \mid A[i, k] = B[k, j]\}| > 0$, for each i, j . The Dominance Product of integer matrices A and B is the matrix C with $C[i, j] = |\{k \mid A[i, k] \leq B[k, j]\}|$. The Boolean version of the problem, \exists Dominance Product asks whether $|\{k \mid A[i, k] \leq B[k, j]\}| > 0$, for each i, j . While the regular versions of Equality Product and Dominance Product are known to be equivalent [24, 25], their Boolean versions are not known to be, although \exists Dominance Product can be reduced to $O(\log n)$ instances of \exists Equality Product [24, 25].

There are also several intermediate problems that currently do not have any improvement brought by fast rectangular matrix multiplication, and whose running times are all $\tilde{O}(n^{(3+\omega)/2}) = \tilde{O}(n^{2.6865})$. These problems include for instance, the aforementioned AE-Mono Δ problem, the (\min, \leq) -product problem, the Max-Min Product problem studied in relation to the All-Pairs Bottleneck Paths and All-Pairs Nondecreasing Paths [17, 18, 28], and the related $(\min, =)$ -product which we introduce in this paper.

Lincoln et al. [16] gave reductions from unweighted directed APSP to Max-Min Product and AE-Mono Δ , showing that if there exists a $T(n)$ time algorithm for Max-Min Product or AE-Mono Δ , then one can also obtain an $\tilde{O}(T(n))$ time algorithm for unweighted directed APSP. They also give reductions from Min-Witness Product to Max-Min Product and AE-Mono Δ .

These reductions are not tight when $\omega > 2$, since they are reductions from a seemingly easier problem (for which improvements via rectangular matrix multiplication are known) to a harder problem (for which no improvements via rectangular matrix multiplication are known). For instance, in order to use the reduction to AE-Mono Δ to get a better algorithm for unweighted directed APSP, one would need to obtain a better than $\tilde{O}(n^{2.5286})$ time algorithm for AE-Mono Δ . This doesn’t seem doable with current techniques.

Hence a natural question is whether one can obtain tight reductions between some of the intermediate problems.

As a first result in this direction, we show that AE-Mono Δ is the hardest “intermediate” problem among the ones mentioned above, as all of the problems can be reduced to it. A key step in the reduction is to study a new problem called *All-Edges Monochromatic Equality Triangle* (AE-MonoEq Δ), which can be viewed as a combination of AE-Mono Δ and \exists Equality Product. We delay its formal definition to Section II. As the first step in the reductions, we reduce AE-Mono Δ to AE-MonoEq Δ .

Theorem 1.7: If AE-Mono Δ has an $\tilde{O}(n^{(3+\omega)/2-\epsilon})$ time algorithm for $0 \leq \epsilon \leq \frac{\omega}{2} - 1$, then AE-MonoEq Δ has an $\tilde{O}(n^{(3+\omega)/2-\delta})$ time algorithm for $\delta \geq 0$. Moreover, if $\epsilon > 0$ then $\delta > 0$.

If we use $\epsilon = 0$ in Theorem 1.7, we immediately get an $\tilde{O}(n^{(3+\omega)/2})$ time algorithm for AE-MonoEq Δ , showing that AE-MonoEq Δ is indeed another “intermediate” problem. If $\omega > 2$, we can use $\epsilon > 0$ in Theorem 1.7 to get that any slight improvements over the current best algorithm for AE-Mono Δ implies an improved algorithm for AE-MonoEq Δ .

As suggested by the names of the problems, AE-Mono Δ is a special version of AE-MonoEq Δ (See Section II for the formal definition), so that any runtime improvements for AE-MonoEq Δ imply runtime improvements of AE-Mono Δ . Thus, AE-MonoEq Δ and AE-Mono Δ are sub- $n^{(3+\omega)/2}$ -fine-grained equivalent when $\omega > 2$.

Next, we show that one can reduce all the “intermediate” problems mentioned above to AE-MonoEq Δ .

Theorem 1.8: Suppose there exists a $T(n)$ time algorithm for AE-MonoEq Δ , then there exist $\tilde{O}(T(n))$ time algorithms for all of the following:

- unweighted directed APSP,
- Min-Witness Product,
- \exists Equality Product,
- \exists Dominance Product,
- $(\min, =)$ -product,
- Max-Min Product,
- (\min, \leq) -product.

Some reductions in the above theorem were already known. First of all, both Max-Min Product and \exists Dominance Product can be reduced to (\min, \leq) -product [17]. Next, both Min Witness Product and unweighted directed APSP can be reduced to Max-Min Product [16]. As mentioned earlier, \exists Dominance Product reduces to \exists Equality Product, and the latter easily reduces to $(\min, =)$ -product. Thus, to prove Theorem 1.7, it suffices to give reductions from $(\min, =)$ -product and (\min, \leq) -product to AE-MonoEq Δ .

Since AE-MonoEq Δ has an $\tilde{O}(n^{(3+\omega)/2})$ time algorithm, Theorem 1.7 provides alternative algorithms for $(\min, =)$ -product and Max-Min Product. These new algorithms are also potentially simpler as they do not involve dealing with

sparse matrix products, which were the main source of difficulty in the previous $\tilde{O}(n^{(3+\omega)/2})$ time algorithms for the problems

Combining Theorem 1.7 and Theorem 1.8, we obtain that AE-Mono Δ is the hardest “intermediate” problem among all “intermediate” problems considered in [16], in the sense that if there is any improvement of AE-Mono Δ over the $\tilde{O}(n^{(3+\omega)/2})$ running time when $\omega > 2$, there will also be improvements for $(\min, =)$ -product, Max-Min Product and (\min, \leq) -product.

AE-MonoEq Δ can be viewed as many independent instances of a problem called AE-Eq Δ , in which we are given a graph G with edge values, and we are asked to decide for each edge e in the graph, whether it is in a triangle such that at least two of its three edges share the same edge value. Via techniques used in the proof of Theorem 4 in [16], we can show that computing a single instance of AE-MonoEq Δ of size n is equivalent to, up to poly-logarithmic factors, computing a certain number of instances of AE-Eq Δ on graphs with n vertices where the total number of edges across all instances is $\Theta(n^2)$.

Motivated by the simple nature of AE-Eq Δ and its relationship to our AE-MonoEq Δ problem, we consider the monochromatic versions of other intermediate problems. The most interesting ones of these are arguably the monochromatic versions of \exists Equality Product and $(\min, =)$ -product which we call Monochromatic Equality Product and Monochromatic $(\min, =)$ -product.

In the Monochromatic Equality Product problem (MonoEq), we are given a tripartite graph G on vertex parts $I \cup J \cup K$. Each edge e in G has a color $c(e)$. All edges e in $I \times K$ and $J \times K$ has a value $v(e)$. For every (i, j) , we need to decide if there exists k such that $v[i, k] = v[j, k]$ and $c[i, k] = c[j, k] = c[i, j]$. MonoEq can be regarded as a combination of many sparse \exists Equality Product instances, where we are given sparse matrices A and B , and we need to compute their \exists Equality Product result on a small number of entries. MonoEq is a special case of AE-MonoEq Δ , so there exists an $\tilde{O}(n^{(3+\omega)/2})$ time algorithm for it.

The input to the Monochromatic $(\min, =)$ -product problem (MonoMinEq) is the same as the input to MonoEq. For the output, instead of only determining for each (i, j) the existence of k such that $v[i, k] = v[j, k]$ and $c[i, k] = c[j, k] = c[i, j]$, we also have to output the minimum value of such a $v[i, k]$. MonoMinEq can be viewed as combination of many sparse $(\min, =)$ -product instances.

The best known algorithms for \exists Equality Product and $(\min, =)$ -product have different running times, and it is unclear whether they are equivalent; clearly \exists Equality reduces to $(\min, =)$ -product, but a reduction in the other direction would imply an improvement over the known algorithms for $(\min, =)$ -product.

Surprisingly, we are able to show that the monochromatic versions, MonoEq and MonoMinEq, are equivalent up to

poly-logarithmic factors.

Theorem 1.9: If there exists a $T(n)$ time algorithm for Monochromatic Equality Product, then there exists an $\tilde{O}(T(n))$ time algorithm for Monochromatic $(\min, =)$ -product, and vice versa.

Theorem 1.9 also implies an $\tilde{O}(n^{(3+\omega)/2})$ time algorithm for MonoMinEq.

II. PRELIMINARIES

In this section, we recall formal definitions of problems considered in this paper and define notations used in the proofs.

A. Hardness Sources

Definition 2.1 (3SUM): Given n integers from the set $\{-n^c, \dots, n^c\}$ for constant c , determine if three of the integers sum to 0.

Conjecture 2.2 (3SUM hypothesis): In the word-RAM model with $O(\log n)$ bit words, there is no $O(n^{2-\epsilon})$ for $\epsilon > 0$ time algorithm for 3SUM.

Definition 2.3 (APSP): Given an n -node directed graph with integer weights in $\{-n^c, \dots, n^c\}$ and no negative cycles, compute the shortest paths distance between every pair of vertices.

Conjecture 2.4 (APSP hypothesis): In the word-RAM model with $O(\log n)$ bit words, there is no $O(n^{3-\epsilon})$ for $\epsilon > 0$ time algorithm for APSP.

Definition 2.5 (Exact Triangle, Exact Δ): Given an integer T and an n -node graph with integer edge weights in $\{-n^c, \dots, n^c\}$ for some constant c , decide whether there are three vertices p, q, r so that $w(p, q) + w(q, r) + w(r, p) = T$.

Conjecture 2.6 (Exact Triangle hypothesis): In the word-RAM model with $O(\log n)$ bit words, there is no $O(n^{3-\epsilon})$ for $\epsilon > 0$ time algorithm for Exact Triangle.

It is known that either the 3SUM hypothesis or the APSP hypothesis implies the Exact Triangle hypothesis [5, 8].

Definition 2.7 (Zero Triangle, Zero Δ): Given an n -node graph with integer edge weights in $\{-n^c, \dots, n^c\}$ for some constant c , decide whether there are three vertices p, q, r so that $w(p, q) + w(q, r) + w(r, p) = 0$

It is known that Exact Triangle is sub-cubic fine-grained equivalent to Zero Triangle [5].

B. Graph Problems

Definition 2.8 (AE-Sparse Δ): Given an n -node m -edge graph $G = (V, E)$, determine for every edge $e \in E$ whether e is in a triangle.

Definition 2.9 (UnweightedAPSP): Given an n -node directed unweighted graph $G = (V, E)$, compute the shortest paths distance between every pair of vertices.

C. Set Problems

we define two problems investigated in [19]: SetDisjointness and SetIntersection.

Definition 2.10 (SetDisjointness): Given a universe U , a family $\mathcal{F} \subseteq 2^U$ of subsets of U , and q pairs of queries $(S, S') \in \mathcal{F} \times \mathcal{F}$, determine for each query (S, S') whether $S \cap S'$ is empty.

Definition 2.11 (SetIntersection): Given a universe U , a family $\mathcal{F} \subseteq 2^U$ of subsets of U , q pairs of queries $(S, S') \in \mathcal{F} \times \mathcal{F}$ and a number T , output elements of $S \cap S'$ for each query (S, S') . It is allowed to terminate the algorithm once it outputs T elements in total.

D. Matrix Product Problems

Definition 2.12 (MinWitness): Given two $n \times n$ Boolean matrices A, B , compute an $n \times n$ matrix C such that

$$C_{i,j} = \min(\{k : A_{i,k} = B_{k,j} = 1\} \cup \{\infty\}).$$

Definition 2.13 ((max, min)-product): Given two $n \times n$ integer matrices A, B , compute an $n \times n$ matrix C such that

$$C_{i,j} = \max_k \min \{A_{i,k}, B_{k,j}\}.$$

We define a generic (\oplus, \otimes) -product, where \oplus maps a set of integers to an integer, and \otimes maps two integers to a Boolean value.

Definition 2.14 ((\oplus, \otimes)-product): Given two $n \times n$ integer matrices A, B , compute an $n \times n$ matrix C such that

$$C_{i,j} = \oplus(\{B_{k,j} : A_{i,k} \otimes B_{k,j}\}_k).$$

We can define operation **NonEmpty** that returns 1 if a set is nonempty, and 0 otherwise. We can therefore define \exists Equality Product as (**NonEmpty**, =)-product, and define \exists Dominance Product as (**NonEmpty**, \leq)-product.

If we define $\min(\emptyset) = \infty$ and $\max(\emptyset) = -\infty$, then $(\min, =)$ -product, (\min, \leq) -product and (\max, \leq) -product all fall into this generic definition without ambiguity.

E. Problems with Colors

Definition 2.15 (AE-Mono Δ): Given an n -node graph $G = (V, E)$, where every edge $e \in E$ has a color $c(e)$. Determine for every edge e , whether it appears in a monochromatic triangle in G .

Definition 2.16 (AE-MonoEq Δ): Given an n -node graph $G = (V, E)$, where each edge $e \in E$ has a color $c(e)$ and a value $v(e)$. Determine for every edge e , whether it appears in a monochromatic triangle in G that at least two of its edges share the same value.

Definition 2.17 (MonoEq): Given a graph $G = (I \cup J \cup K, E)$, where $|I| = |J| = |K| = n$. Each edge e in the graph has a color $c(e)$. Each edge e in $I \times K$ and $J \times K$ has a value $v(e)$. For every (i, j) , decide if there exists k such that $v(i, j) = v(j, k)$ and $c(i, k) = c(j, k) = c(i, j)$.

Definition 2.18 (MonoMinEq): Given a graph $G = (I \cup J \cup K, E)$, where $|I| = |J| = |K| = n$. Each edge e in the

graph has a color $c(e)$. Each edge e in $I \times K$ and $J \times K$ has a value $v(e)$. For every (i, j) , compute the minimum value of $v(i, k)$ where $v(i, k) = v(j, k)$ and $c(i, k) = c(j, k) = c(i, j)$.

Definition 2.19 (Mono. (min, \leq)-product): Given a tripartite graph $G = (I \cup J \cup K, E)$, where $|I| = |J| = |K| = n$. Each edge e in the graph has a color $c(e)$. Each edge e in $I \times K$ and $J \times K$ has a value $v(e)$. For every (i, j) , compute the minimum value of $v(j, k)$ where $v(i, k) \leq v(j, k)$ and $c(i, k) = c(j, k) = c(i, j)$.

F. Notations

For a graph $G = (V, E)$, a node $v \in V$ and $U \subseteq V$, we use $\deg(v, U)$ to denote $|\{u \in U : (v, u) \in E\}|$.

We use $\omega(a, b, c)$ to denote the rectangular matrix multiplication exponent, i.e. the smallest real number z such that the time to multiply an $n^a \times n^b$ matrix by an $n^b \times n^c$ matrix is $O(n^{z+\varepsilon})$ for all $\varepsilon > 0$. In particular, let $\omega = \omega(1, 1, 1)$ be the exponent for square matrix multiplication. It is known that $\omega \in [2, 2.373]$ [12, 13]. The best known bounds for $\omega(a, b, c)$ are in [21].

III. HARDNESS OF SPARSE TRIANGLE LISTING

We define a parameterized version of Triangle Listing. In this version, the graph has three parts of vertices. Each of the three parts can have different sizes, but edge densities between any pair of the parts are the same.

Definition 3.1 ($(\alpha, \beta, \gamma, \rho, t)$ -All-Edges Triangle Listing): Given a tripartite graph G with vertex set $A \cup B \cup C$, such that $|A| = n^\alpha, |B| = n^\beta, |C| = n^\gamma$. Let $X, Y \in \{A, B, C\}$ be two different parts of the graph. For any $v \in X$, it is guaranteed that $\deg(v, Y) \leq O(n^{-\rho}|Y|)$. The $(\alpha, \beta, \gamma, \rho, t)$ -All-Edges Triangle Listing problem asks to list, for each $e \in E \cap (A \times B)$, all triangles containing e if there are fewer than t such triangles or t distinct triangles containing e if there are at least t such triangles.

Definition 3.1 defines the triangle listing problem slightly differently from the usual definition. In previous works (e.g. [10, 29]), the algorithm for triangle listing is required list up to T triangles in the the whole graph, while $(\alpha, \beta, \gamma, \rho, t)$ -All-Edges Triangle Listing asks to list up to t triangles for each edge.

We also define an unbalanced triangle listing problem when we have to list up to T triangles globally.

Definition 3.2 ($(\alpha, \beta, \gamma, \rho, T)$ Triangle Listing): Given a tripartite graph G whose vertex set is $A \cup B \cup C$, such that $|A| = n^\alpha, |B| = n^\beta, |C| = n^\gamma$. Let $X, Y \in \{A, B, C\}$ be two different parts of the graph. For any $v \in X$, $\deg(v, Y) \leq O(n^{-\rho}|Y|)$. The $(\alpha, \beta, \gamma, \rho, T)$ Triangle Listing problem asks to list all triangles in G if there are fewer than T triangles or list T distinct triangles in G if there are at least T triangles.

Triangle Listing and All-Edges Triangle Listing are strongly related problems. For instance, it can be shown

that $(\alpha, \beta, \gamma, \rho, t)$ -All-Edges Triangle Listing can be reduced to, up to polylogarithmic factors, $(\alpha, \beta, \gamma, \rho, n^{\alpha+\beta-\rho}t)$ Triangle Listing, by a reduction similar to the reduction for Theorem 15 in [10]. This reduction means that if we have hardness for $(\alpha, \beta, \gamma, \rho, t)$ -All-Edges Triangle Listing, then we also have hardness for $(\alpha, \beta, \gamma, \rho, T)$ Triangle Listing when $T = n^{\alpha+\beta-\rho}t$. However, this argument won't work when $T < n^{\alpha+\beta-\rho}$, since it would require us to set $t < 1$, which doesn't make sense. Therefore, to circumvent this difficulty, we will directly reduce Exact Triangle to both Triangle Listing and All-Edges Triangle Listing.

The triangle listing problems require us to actually list triangles for some edge e . However, many problems we consider, including All-Edges Sparse Triangle and AE-Mono Δ , only require the algorithms to output whether some triangle containing edge e exists. Thus, in order to reduce to these problems, we define the following version of triangle detection.

Definition 3.3 ($(\alpha, \beta, \gamma, \rho)$ -All-Edges Sparse Triangle):

Given a tripartite graph G whose vertex set is $A \cup B \cup C$, such that $|A| = n^\alpha, |B| = n^\beta, |C| = n^\gamma$. Let $X, Y \in \{A, B, C\}$ be two different parts of the graph. For any $v \in X$, it is guaranteed that $\deg(v, Y) \leq O(n^{-\rho}|Y|)$. The $(\alpha, \beta, \gamma, \rho)$ -All-Edges Sparse Triangle problem asks to determine, for each $e \in E \cap (A \times B)$, whether e is in a triangle or not.

Now we are ready to present the reduction from Exact Triangle to triangle listing problems.

Theorem 3.4: Fix nonnegative constants $0 \leq \alpha, \beta, \gamma \leq 1$ and $\rho < \min\{\alpha, \beta, \gamma\}$. There exists an $\tilde{O}(n^{3-\min\{\alpha, \beta, \gamma\}+\rho})$ time randomized reduction from a Zero Triangle instance with n vertices to $\tilde{O}(n^{3-\alpha-\beta-\gamma+2\rho})$ instances of $(\alpha, \beta, \gamma, \rho, 900n^{\gamma-2\rho} + 1)$ -All-Edges Triangle Listing. Similarly, there is also an $\tilde{O}(n^{3-\min\{\alpha, \beta, \gamma\}+\rho})$ time randomized reduction from a Zero Triangle instance with n vertices to $\tilde{O}(n^{3-\alpha-\beta-\gamma+2\rho})$ instances of $(\alpha, \beta, \gamma, \rho, 8100n^{\alpha+\beta+\gamma-3\rho} + 1)$ Triangle Listing.

Proof:

We will first provide the reduction to All-Edges Triangle Listing. The reduction to Triangle Listing can be obtained via slight modifications.

Step 1:

Fix a Zero Triangle instance G . We can randomly assign its vertices to one of three colors, and only keep edges whose two endpoints have different colors. If we repeat $\Theta(\log n)$ times, a zero triangle in G will remain in at least one of the graphs. Thus it suffices to solve Zero Triangle on a tripartite graph.

Suppose G is a tripartite graph with vertex parts A, B, C . First, we split vertex parts A, B, C to parts of size $n^\alpha, n^\beta, n^\gamma$ respectively. We could enumerate all $n^{3-\alpha-\beta-\gamma}$ triples of parts, and detect zero triangle within each triple of parts. In the remainder of the reduction, it suffices to reduce each individual unbalanced zero triangle instance of vertex set sizes $|A| = n^\alpha, |B| = n^\beta, |C| = n^\gamma$ to $(\alpha, \beta, \gamma, \rho, t)$ -All-

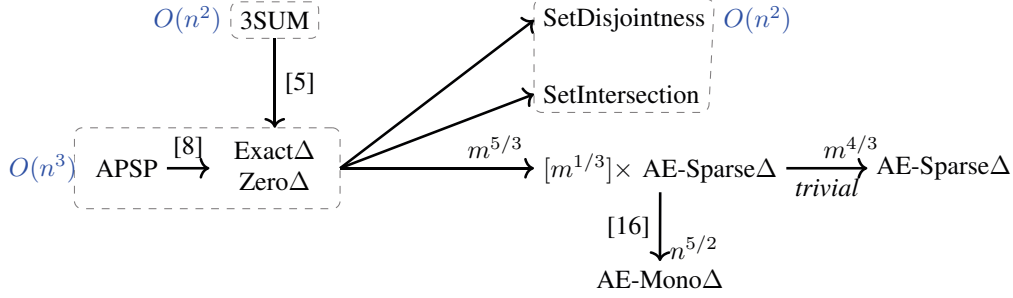


Figure 1. The main reduction in Section III is from Exact Δ /Zero Δ to parameterized versions of triangle detection and triangle listing. This reduction implies hardness for SetDisjointness, SetIntersection, AE-Sparse Δ and AE-Mono Δ .

Edges Triangle Listing instances.

Step 2:

We assume the edge weights $w(\cdot, \cdot)$ in the Zero Triangle instance are integers whose absolute values are bounded by n^k for some constant $k \geq 1$. We pick an arbitrary prime p that is between $100n^k$ and $Dn^k \log n$ for large enough constant D . By the Prime Number Theorem, a random integer in this range is a prime with probability $\Omega(1/\log n)$, so it takes $\tilde{O}(1)$ time to find such a prime by randomly picking integers in this range and test its primality. After we determine p , we can regard all the weights of the graph as elements in \mathbb{F}_p by taking the weight of every edge modulo p . Since the weight of each triangle is in $[-3n^k, 3n^k]$ while $p \geq 100n^k$, the set of zero triangles with respect to the new weights stays the same.

Step 3:

Let $x \in \mathbb{F}_p$ be a random element from \mathbb{F}_p , and let $y_v \in \mathbb{F}_p$ be random elements from \mathbb{F}_p for every $v \in A \cup B \cup C$. For every edge $e = (a, b) \in E \cap (A \times B)$, we set its new weight to $w'(a, b) = xw(e) - y_b + y_a$; for every edge $e = (a, c) \in E \cap (A \times C)$, we set its new weight to $w'(a, c) = xw(e) - y_a + y_c$; for every edge $e = (b, c) \in E \cap (B \times C)$, we set its new weight to $w'(b, c) = xw(e) - y_c + y_b$. Let the graph with the new weights be G' .

Clearly, as long as $x \neq 0$, the set of zero triangles with weights $w'(\cdot, \cdot)$ is exactly the same as the set of zero triangles with weights $w(\cdot, \cdot)$. Thus, false positives occur with probability $\frac{1}{p} \leq \frac{1}{100n^k} \leq 0.01$.

Step 4:

We split \mathbb{F}_p into n^ρ contiguous ranges L_1, \dots, L_{n^ρ} , so that each range has size between $\lfloor p/n^\rho \rfloor$ and $\lceil p/n^\rho \rceil$. We enumerate every triple of i, j, k such that $0 \in L_i + L_j + L_k$. For every pair of i, j , the size of $L_i + L_j$ is $O(p/n^\rho)$. In order for $0 \in L_i + L_j + L_k$, we need $L_k \cap -(L_i + L_j) \neq \emptyset$. Since each L_k has size $\Theta(p/n^\rho)$, there can be at most $O(1)$ ranges L_k that intersect with $-(L_i + L_j)$. Thus, the total number of such triples is $O(n^{2\rho})$.

For each triple (i, j, k) , we consider a subset of edges

$E_{i,j,k}$ defined as

$$\begin{aligned} & \{e \in E \cap (A \times C) : w'(e) \in L_i\} \\ & \cup \{e \in E \cap (B \times C) : w'(e) \in L_j\} \\ & \cup \{e \in E \cap (A \times B) : w'(e) \in L_k\}. \end{aligned}$$

Let $G_{i,j,k} = (A \cup B \cup C, E_{i,j,k})$ be the subgraph of G' with edge set $E_{i,j,k}$. Clearly, if graph G' has a zero triangle, one of $G_{i,j,k}$ will have a zero triangle, and vice versa.

Now we change $G_{i,j,k}$ so that the degree of every vertex is bounded. For each $v \in A$, if $\deg(v, |B|) > 100|B|n^{-\rho} + 200$ or $\deg(v, |C|) > 100|C|n^{-\rho} + 200$, we remove the vertex v from graph $G_{i,j,k}$. We similarly handle vertices in parts B and C that have large degrees.

Finally, we use an algorithm for $(\alpha, \beta, \gamma, \rho, t)$ -All-Edges Triangle Listing for $t = 900n^{c-2\rho} + 1$ on graph $G_{i,j,k}$ to list up to $900n^{c-2\rho} + 1$ triangle for each edge $e \in E_{i,j,k} \cap (A \times B)$. If any of these listed triangles is a zero triangle in the original graph G , we return YES to the Zero Triangle instance.

Step 5:

We repeat the previous steps $100 \log n$ times. If no zero triangle is found in any of these $100 \log n$ tries, we return NO to the Zero Triangle instance.

Analysis

It should be clear why Step 1 through Step 3 works via the in-text explanations. Now we prove why Step 4 and Step 5 work.

Claim 3.5: Fix any zero triangle (a, b, c) in $G_{i,j,k}$ where $a \in A, b \in B, c \in C$. With probability at least 0.94, none of a, b, c will be removed in Step 4 due to having a large degree.

Proof: First consider $\deg(a, B)$. For any $b' \in B \setminus \{b\}$, $w'(a, b') = x \cdot w(a, b') - y_{b'} + y_a$. Conditioned on the fact that $(a, b), (b, c), (c, a) \in E_{i,j,k}$, $y_{b'}$ is a completely new random variable. Therefore, $\Pr[(x, b') \in E_{i,j,k}] = \Pr[w(a, b') - y_{b'} + y_a \in L_k] = \frac{|L_k|}{p} \leq \frac{\lceil pn^{-\rho} \rceil}{p} \leq n^{-\rho} + \frac{1}{p}$. Therefore, the

expected value of $\deg(a, B)$ can be written as

$$\begin{aligned} \mathbb{E} [\deg(a, B)] &= 1 + \mathbb{E} \left[\sum_{b' \neq b} [(a, b') \in E_{i,j,k}] \right] \\ &\leq 1 + |B|n^{-\rho} + \frac{|B|}{p} \leq 2 + |B|n^{-\rho}. \end{aligned}$$

Thus, by Markov's inequality,

$$\Pr [\deg(a, B) > 200 + 100|B|n^{-\rho}] \leq 0.01.$$

We can apply the same argument to $\deg(a, C)$, $\deg(b, A)$, $\deg(b, C)$, $\deg(c, A)$ and $\deg(c, B)$ and take a union bound. Thus, with probability at least 0.94, all of these degrees will be small enough so that none of a, b, c are removed. ■

We also need to show that listing $900n^{\gamma-2\rho} + 1$ triangles for each edge will be enough, i.e. there are not too many false positives for each edge $e \in E_{i,j,k} \cap (A \times B)$.

Claim 3.6: Fix any zero triangle (a, b, c) in $G_{i,j,k}$ where $a \in A, b \in B, c \in C$. The number of vertices c' such that

- 1) $(a, c'), (b, c') \in E_{i,j,k}$, and
- 2) $w(a, b) + w(b, c') + w(c', a) \neq 0$ (Recall that $w(\cdot, \cdot)$ is viewed as elements in \mathbb{F}_p , so all operations are modulo p).

is at most $900n^{\gamma-2\rho}$ with probability at least 0.99.

Proof: Let c' be any $c' \in C$ such that $w(a, b) + w(b, c') + w(c', a) \neq 0$. If $(a, c'), (b, c') \in E_{i,j,k}$, then $w'(a, c'), w'(a, c) \in L_i$ and $w'(b, c'), w'(b, c) \in L_j$. Since L_i and L_j are both ranges of size at most $\lceil pn^{-\rho} \rceil$, we must have

$$\begin{cases} w'(a, c) - w'(a, c') \in [-pn^{-\rho}, pn^{-\rho}] \\ w'(b, c) - w'(b, c') \in [-pn^{-\rho}, pn^{-\rho}]. \end{cases}$$

We can expand out the definition of w' to get

$$\begin{cases} x \cdot (w(a, c) - w(a, c')) + (y_c - y_{c'}) \in [-pn^{-\rho}, pn^{-\rho}] \\ x \cdot (w(b, c) - w(b, c')) + (y_c - y_{c'}) \in [-pn^{-\rho}, pn^{-\rho}]. \end{cases} \quad (1)$$

Each of the two values in Equation (1) is clearly uniformly at random. To show these two values are independent, we consider the sum of these two values, which is $x \cdot (w(a, c) + w(b, c) - w(a, c') - w(b, c'))$. Since $w(a, b) + w(b, c') + w(c', a) \neq 0$, while $w(a, b) + w(b, c) + w(c, a) = 0$, we have $w(a, c) + w(b, c) \neq w(a, c') + w(b, c')$. Thus, the sum of the two values in Equation (1) is the product of x with a nonzero value. Thus, the sum of these two values is also a uniformly random variable. Conditioned on the sum, $x \cdot (w(a, c) - w(a, c')) + (y_c - y_{c'})$ is also uniformly at random, since the sum does not contain the $y_{c'}$ term. Thus the two values in Equation (1) are independent. Therefore, the probability that Equation (1) happens is at most $(2n^{-\rho} + \frac{1}{p})^2 \leq 9n^{-2\rho}$.

Summing over all $c' \in C$, the expected number of c' satisfying conditions 1) and 2) is at most $9n^{\gamma-2\rho}$. Thus, the

probability that the number of such c' exceeds $900n^{\gamma-2\rho}$ is at most 0.01. ■

If an edge (a, b) is in a zero triangle (a, b, c) in the original graph G , then this zero triangle is preserved in one instance $G_{i,j,k}$ before removing any vertices. Then we will report a triangle containing edge (a, b) as long as

- 1) we don't remove any of a, b, c in the vertex removal process in Step 4 (which happens with probability at least 0.94 by Claim 3.5);
- 2) the number of nonzero triangles containing (a, b) in $G_{i,j,k}$ is at most $900n^{\gamma-2\rho}$ (which happens with probability at least 0.99 by Claim 3.6).

Therefore, by union bound, each iteration reports at least one zero triangle with constant probability if the original graph has a zero triangle. Thus, repeating the iterations for $O(\log n)$ time suffices.

Running Time:

We split the n -node graph to $n^{3-\alpha-\beta-\gamma}$ unbalanced graphs in Step 1. For each unbalanced graph, we reduce it to $n^{2\rho}$ instances of $(\alpha, \beta, \gamma, \rho, 900n^{\gamma-2\rho} + 1)$ -All-Edges Triangle Listing. The running time of the reduction is linear (up to poly-logarithmic factors) to the total input size of all the triangle listing instances. Thus, the running time is

$$\begin{aligned} &\tilde{O}(n^{3-\alpha-\beta-\gamma+2\rho} \cdot (n^{\alpha+\beta-\rho} + n^{\beta+\gamma-\rho} + n^{\alpha+\gamma-\rho})) \\ &= \tilde{O}(n^{3-\min\{\alpha, \beta, \gamma\}+\rho}). \end{aligned}$$

The reduction from Zero Triangle to Triangle Listing

The reduction is largely the same as the previous reduction. The only difference in the reduction is that now in Step 4, we use the $(\alpha, \beta, \gamma, \rho, T)$ Triangle Listing algorithm on graph $G_{i,j,k}$ to list up to $T = 8100n^{\alpha+\beta+\gamma-3\rho} + 1$ triangles, and test whether any of these triangles is a zero triangle. The correctness analysis of this reduction requires a more careful analysis of the expected number of triangles in $G_{i,j,k}$. We omit the details and defer the full proof of this part to the full version of this paper. ■

Theorem 3.4 shows hardness for listing triangles in some special graphs. To show the hardness for detecting triangles, we still need a reduction from triangle listing to triangle detection. Theorem 15 in [10] is such a reduction that reduces listing $O(1)$ triangles for each edge to detecting whether each edge is in a triangle; however, that reduction changes the structure of the graph. Specifically, it does not necessarily reduce an $(\alpha, \beta, \gamma, \rho, O(1))$ -All-Edges Triangle Listing instance to $(\alpha, \beta, \gamma, \rho)$ -All-Edges Sparse Triangle instances. Thus, we give a new structure-preserving reduction from triangle listing to all edge triangle detection. The reduction adapts the techniques for finding the witnesses of Boolean matrix multiplication [30, 31].

Proposition 3.7: Let α, β, γ be any positive constants and let $\rho \leq \min\{\alpha, \beta, \gamma\}$. There exists an $\tilde{O}(n^{\alpha+\beta} +$

$n^{\beta+\gamma} + n^{\gamma+\alpha})n^{-\rho}$) time randomized reduction from an $(\alpha, \beta, \gamma, \rho, O(1))$ -All-Edges Triangle Listing instance to $\tilde{O}(1)$ instances of $(\alpha, \beta, \gamma, \rho)$ -All-Edges Sparse Triangle.

We defer the proof of Proposition 3.7 to the full version of this paper.

Using Theorem 3.4 and Proposition 3.7, we obtain the following corollaries. As their proofs are standard, we defer the proofs to the full version.

Corollary 3.8: There exists a reduction from Exact Triangle to $\tilde{O}(n)$ instances of All-Edges Sparse Triangle of $O(n^{1.5})$ edges. Thus, assuming the Exact Triangle hypothesis, there is no $O(m^{4/3-\epsilon})$ time algorithm for All-Edges Sparse Triangle for $\epsilon > 0$.

Corollary 3.9: Assuming the Exact Triangle hypothesis, there is no $O(n^{2.5-\epsilon})$ time algorithm for AE-Mono Δ on n -node graphs for $\epsilon > 0$.

Corollary 3.10: Let \mathbb{A} be an algorithm for All-Edges Sparse Triangle for n -node graphs where every node has degree at most $d = n^\delta$ for some $0 < \delta \leq 0.5$. Assuming the Exact Triangle hypothesis, \mathbb{A} cannot run in $O(n^{1-\epsilon}d^2)$ for $\epsilon > 0$.

Corollary 3.11: For any constant $0 < \theta < 1$, let \mathbb{A} be an algorithm for offline SetDisjointness where $|U| = \Theta(N^{2-2\theta})$, $|\mathcal{F}| = \Theta(N)$, each set in \mathcal{F} has at most $O(N^{1-\theta})$ elements from U , and $q = \Theta(N^{1+\theta})$. Assuming the Exact Triangle hypothesis, \mathbb{A} cannot run in $O(N^{2-\epsilon})$ for $\epsilon > 0$.

Corollary 3.12: For any constants $0 \leq \theta < 1$ and $0 < \delta$, let \mathbb{A} be an algorithm for offline SetIntersection where $|U| = \Theta(N^{1+\delta-\theta})$, $|\mathcal{F}| = \Theta(\sqrt{N^{1+\delta+\theta}})$, each set in \mathcal{F} has at most $O(N^{1-\theta})$ elements from U , $q = \Theta(N^{1+\theta})$, and $T = O(N^{2-\delta})$. Assuming the Exact Triangle hypothesis, \mathbb{A} cannot run in $O(N^{2-\epsilon})$ for $\epsilon > 0$.

IV. REDUCTIONS BETWEEN INTERMEDIATE PROBLEMS

Consider an AE-MonoEq Δ instance $G = (V, E)$. We can copy the vertex set of G three times to parts I, J, K , and then plant the edges of G to $I \times J, J \times K$ and $I \times K$. Thus we only need to solve AE-MonoEq Δ on tripartite graphs. Say we want to report for every edge between parts I and J , whether it is in a monochromatic equality triangle. Depending on where the two edges with the same values are, we can split the problem to three cases shown in Figure 3.

Cases (b) and (c) are symmetric, so it suffices to only consider case (a) and case (b). The following theorem shows that we can solve AE-MonoEq Δ in $\tilde{O}(n^{(3+\omega)/2})$ time. Moreover, when $\omega > 2$, if AE-Mono Δ has a better algorithm, so does AE-MonoEq Δ . However, when $\omega = 2$, we must have $\lambda = 0$ in the theorem, so AE-MonoEq Δ won't necessarily have a better algorithm given a better algorithm for AE-Mono Δ .

Theorem 4.1: Suppose $\omega \geq 2 + \lambda$ for some $\lambda \geq 0$, and suppose AE-Mono Δ has an $O(n^{(3+\omega)/2-\epsilon})$ time algorithm,

then AE-MonoEq Δ has an

$$\tilde{O}\left(n^{(\omega+3)/2 - \frac{\frac{\lambda}{2(\kappa-1)}\epsilon}{\omega+3 + \frac{\lambda}{2(\kappa-1)} - 2\epsilon}}\right)$$

time algorithm, where $\kappa \geq 3$ is a constant depending on ω and λ .

Proof: We first show the algorithm for AE-MonoEq Δ on tripartite graphs $I \cup J \cup K$ with values on edge sets $I \times K$ and $J \times K$ (Case (a) in Figure 3). We define a good triangle (i, j, k) to be a triangle where all its three edges have the same color, and $v(i, k) = v(j, k)$. For each edge $e \in E \cap (I \times J)$, the algorithm needs to report whether it is in a good triangle or not.

We will create a new graph G' such that the edges of G' only has colors, instead of having both colors and values. The vertex set of G' will be $I \cup J \cup (K \times \mathcal{V})$, where \mathcal{V} is the set of all values in G (we will create the graph lazily, so we won't actually spend time creating those vertices that end up not having any neighbors). The edge set and the colors on edges between I and J in G' is the same as the edge set between I and J in G . For each edge (i, k) in $E(G)$ with value v and color c , we create an edge between i and (k, v) in G' with the same color c . Similarly, for each edge (j, k) , in $E(G)$ with value v and color c , we create an edge between j and (k, v) in G' with the same color c . We can see that edge (i, j) is in a good triangle in G if and only if it is in a monochromatic triangle in G' . Thus, from now on, we can focus on the All-Edges Monochromatic Triangle problem on the unbalanced graph G' .

Let $D \geq 1$ and $T \geq n$ be some parameters to be fixed later. For each color c , let G'_c be the subgraph of G' consisting of vertices of G' and all edges with color c . For each vertex $(k, v) \in K'$, if its degree in G'_c is at most D , then we can enumerate all pairs of its neighbors and test if these three vertices form a triangle in $\tilde{O}(1)$ time. If a triangle $(i, j, (k, v))$ is found during the enumeration, we record that edge (i, j) is in a good triangle. After the enumeration, we can remove the vertex (k, v) together with all edges incident to it from the graph G'_c . Summing over all c and all (k, v) , these enumerations take $O(n^2D)$ time. We can now assume that, in each G'_c , the degree of every vertex $(k, v) \in K'$ is at least D .

We consider two cases, based on the the remaining size of K' in G'_c . As a high level description, if the size of K' is at least T , we use rectangular matrix multiplication. When the size of K' is less than T , we combine all instances with a small size of K' across every G'_c into a single AE-Mono Δ instance, and use the assumed $O(n^{(3+\omega)/2-\epsilon})$ time algorithm for AE-Mono Δ . We describe and analyze these two cases in more details in the following.

Rectangular matrix multiplication for more unbalanced colors

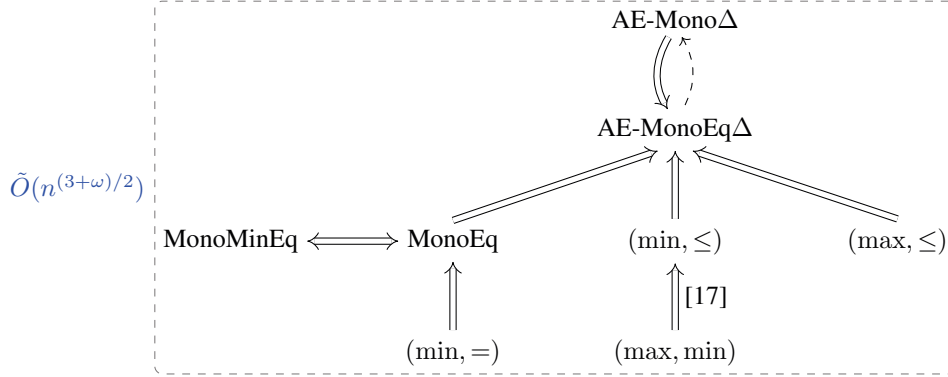


Figure 2. Main reductions in Section IV. Double arrows represent that the running times before and after the reduction are the same up to poly-logarithmic factors. Dashed arrows represent reductions that hold only when $\omega > 2$.




Figure 3. Three cases for AE-MonoEq Δ . The red edges represent edges colors, and the black edges represent edge values.

For a subgraph G'_c , if K' has size at least T , we use rectangular matrix multiplication to determine whether each edge $e \in E(G'_c) \cap (I \times J)$ is in a triangle. To do so, we create an integer matrix X of dimension $|I| \times |K'|$ and an integer matrix Y of dimension $|K'| \times |J|$. Initially all entries in X and Y are zero. If there is an edge between $i \in I$ and $k' \in K'$, we set the (i, k') -th entry of X to 1. Similarly, if there is an edge between $j \in J$ and $k' \in K'$, we set the (k', j) -th entry of Y to 1. It is then not hard to see that edge (i, j) is in a triangle if and only if $(XY)_{i,j} > 0$. Since $|I|, |J| \leq n$, it would take $O(n^{\omega(1, \log_n |K'|, 1)})$ time to multiply X and Y .

In order to analyze rectangular matrix multiplication, we need some bound on $\omega(1, t, 1)$ in the regime when $\omega \geq 2 + \lambda$.

Claim 4.2: There exists a constant κ such that for every $1 \leq t \leq 3$, $\omega(1, t, 1) \leq \omega + (t - 1)(1 - \frac{\lambda}{2(\kappa - 1)})$.

The proof of this claim can be found in the full version of this paper.

Let $t = \log_n T$. If the size of K' in G'_c is n^{q_c} for some $q_c \geq t$, then it takes $O(n^{\omega(1, q_c, 1)})$ time to compute XY . Clearly $1 \leq t \leq q_c \leq 3$, so $n^{\omega(1, q_c, 1)} = O(n^{\omega + (q_c - 1)(1 - \frac{\lambda}{2(\kappa - 1)})})$ by Claim 4.2. Since the degree of every vertex in K' is at least D , and the total number of edges across all G'_c is $O(n^2)$, the sum of the sizes of K' is at most $O(\frac{n^2}{D})$. In other words, $\sum_c n^{q_c} = O(\frac{n^2}{D})$. Therefore, the overall time complexity can be bounded as the following

up to constant factors.

$$\begin{aligned} \sum_c n^{\omega + (q_c - 1)(1 - \frac{\lambda}{2(\kappa - 1)})} &= n^\omega \cdot \sum_c n^{(q_c - 1) / n^{\frac{(q_c - 1)\lambda}{2(\kappa - 1)}}} \\ &\leq n^\omega \cdot \sum_c n^{(q_c - 1) / n^{\frac{(t - 1)\lambda}{2(\kappa - 1)}}} \\ &= \frac{n^{\omega + 1 + \frac{\lambda}{2(\kappa - 1)}}}{DT^{\frac{\lambda}{2(\kappa - 1)}}}. \end{aligned}$$

Therefore, this case takes at most $O(\frac{n^{\omega + 1 + \frac{\lambda}{2(\kappa - 1)}}}{DT^{\frac{\lambda}{2(\kappa - 1)}}})$ time.

All-Edges Monochromatic Triangle for moderately unbalanced colors

In this part, we consider colors c such that the K' part in G'_c has at most T vertices. Since $T \geq n$ we can assume the number of vertices in G'_c is at most $2n + T \leq 3T$. Also, if we sum over the number of edges of G'_c of every c , the total number of edges is at most n^2 .

Note that the graph on each G'_c is an All-Edges Sparse Triangle instance. We can actually combine all these instances to a single All-Edges Monochromatic Triangle instance of size $\tilde{O}(T)$ by a similar reduction of the proof of Theorem 1.2.

The reduction from All-Edges Sparse Triangle to All-Edges Monochromatic Triangle in Theorem 1.2 works as follows [16]. Let H be a multi-graph on $3T$ vertices, initially with no edges. For each All-Edges Sparse Triangle instance on vertex set $[3T]$, we take a random permutation of its vertices, and copy each edge in this instance to H , and color these edges using a color unique to this instance. On expectation, for every $(u, v) \in [3T] \times [3T]$, the multiplicity of (u, v) in H is $\frac{n^2}{(3T)^2} = O(1)$. Since the permutations for graphs G'_c are independent for different colors c , by Chernoff bound, the multiplicity of any (u, v) is $O(\log n)$ with high probability. For every (u, v) in $[3T] \times [3T]$, we arbitrarily label its corresponding $O(\log n)$ edges using numbers from 1 to $O(\log n)$. Then we enumerate $O(\log^3 n)$ triples (i, j, k) of labels. For each triple (i, j, k) , we create a graph on

vertex set $V_1 \sqcup V_2 \sqcup V_3$, where $V_1 = V_2 = V_3 = [3T]$. Between V_1 and V_2 , we add edges in H with label i ; between V_2 and V_3 , we add edges with label j ; between V_3 and V_1 , we add edges with label k . Then a triangle in one All-Edges Sparse Triangle instance corresponds to some monochromatic triangle in one of these $O(\log^3 n)$ instances. This finishes the reduction.

For each of the $O(\log^3 n)$ instances of AE-Mono Δ , we can use the assumed $O(n^{(3+\omega)/2-\epsilon})$ time algorithm for AE-Mono Δ , so we get an $\tilde{O}(T^{(3+\omega)/2-\epsilon})$ time algorithm for this case.

Running Time The three cases have running times $O(n^2 D)$, $O\left(\frac{n^{\omega+1+\frac{\lambda}{2(\kappa-1)}}}{DT^{\frac{\lambda}{2(\kappa-1)}}}\right)$, and $\tilde{O}(T^{(3+\omega)/2-\epsilon})$ respectively.

We can balance them by setting $T = n^{1+\frac{2\epsilon}{\omega+3+\frac{\lambda}{2(\kappa-1)}-2\epsilon}}$ and $D = \frac{T^{(3+\omega)/2-\epsilon}}{n^2}$. The overall running time is

$$\tilde{O}\left(n^{(\omega+3)/2-\frac{\frac{\lambda}{2(\kappa-1)}\epsilon}{\omega+3+\frac{\lambda}{2(\kappa-1)}-2\epsilon}}\right).$$

The case where the values are on edge sets $I \times J$ and $I \times K$ (Case (b) in Figure 3) can be solved similarly with the same running time. ■

Next, we show reductions to AE-MonoEq Δ from many other intermediate problems.

Theorem 4.3: If there is a $T(n)$ time algorithm for AE-MonoEq Δ , then there is an $O(T(n) \log n)$ time algorithm for $(\min, =)$ -product.

The main idea in the proof is parallel binary search. We defer the details to the full version of this paper.

Theorem 4.4: If there is a $T(n)$ time algorithm for AE-MonoEq Δ , then there is an $O(T(n) \log^2 n)$ time algorithm for (\min, \leq) -product.

Theorem 4.4 can be proved by considering the binary representations of the input matrix entries and using the parallel binary search idea. We defer the details to the full version of this paper. Using a similar idea, we can get a reduction to AE-MonoEq Δ from (\max, \leq) product. Even though (\max, \leq) -product looks similar to (\min, \leq) -product, and their best algorithms both run in $\tilde{O}(n^{(3+\omega)})$ time [18], we don't know if they are equivalent.

Proposition 4.5: If there is a $T(n)$ time algorithm for AE-MonoEq Δ , then there is an $O(T(n) \log^2 n)$ time algorithm for (\max, \leq) -product.

Now we consider the Monochromatic Equality Product problem, which can be viewed as Case (a) of AE-MonoEq Δ .

Our proof of Theorem 4.3 only uses Case (a) of AE-MonoEq Δ (see the full version of this paper), so the same proof actually shows a reduction from $(\min, =)$ -product to Monochromatic Equality Product. In fact, we will show that Monochromatic Equality Product is equivalent to Monochromatic $(\min, =)$ -product which is a stronger version of $(\min, =)$ -product.

Theorem 4.6: If there is a $T(n)$ time algorithm for Monochromatic $(\min, =)$ -product then there is an $O(T(n))$ time algorithm for Monochromatic Equality Product. Also, if there is a $T(n)$ time algorithm for Monochromatic Equality Product, then there is an $O(T(n) \log n)$ time algorithm for Monochromatic $(\min, =)$ -product.

We further consider Monochromatic (\min, \leq) -product, which turns out also has an $\tilde{O}(n^{(3+\omega)/2})$ time algorithm.

Proposition 4.7: If there is a $T(n)$ time algorithm for AE-MonoEq Δ , then there is an $O(T(n) \log^2 n)$ time algorithm for Monochromatic (\min, \leq) Product.

We defer the proofs of Theorem 4.6 and Proposition 4.7 to the full version of this paper.

REFERENCES

- [1] R. Williams, "Faster all-pairs shortest paths via circuit complexity," in *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, D. B. Shmoys, Ed. ACM, 2014, pp. 664–673.
- [2] R. R. Williams, "Faster all-pairs shortest paths via circuit complexity," *SIAM J. Comput.*, vol. 47, no. 5, pp. 1965–1985, 2018.
- [3] M. Cygan, M. Mucha, K. Wegrzycki, and M. Włodarczyk, "On problems equivalent to $(\min, +)$ -convolution," *ACM Trans. Algorithms*, vol. 15, no. 1, pp. 14:1–14:25, 2019.
- [4] M. Pătraşcu, "Towards polynomial lower bounds for dynamic problems," in *Proceedings of the forty-second ACM symposium on Theory of computing*, 2010, pp. 603–610.
- [5] V. V. Williams and R. Williams, "Finding, minimizing, and counting weighted subgraphs," *SIAM J. Comput.*, vol. 42, no. 3, pp. 831–854, 2013.
- [6] V. Vassilevska and R. Williams, "Finding, minimizing, and counting weighted subgraphs," in *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, M. Mitzenmacher, Ed. ACM, 2009, pp. 455–464.
- [7] V. V. Williams and R. Williams, "Subcubic equivalences between path, matrix and triangle problems," in *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*. IEEE Computer Society, 2010, pp. 645–654.
- [8] V. V. Williams and R. R. Williams, "Subcubic equivalences between path, matrix, and triangle problems," *J. ACM*, vol. 65, no. 5, pp. 27:1–27:38, 2018.
- [9] A. Abboud, H. Dell, K. Bringmann, and J. Nederlof, "More consequences of falsifying seth and the orthogonal vectors conjecture," in *50th Annual ACM Symposium on Theory of Computing, STOC 2018*. Association for Computing Machinery, Inc, 2018, pp. 445–456.

- [10] L. Duraj, K. Kleiner, A. Polak, and V. Vassilevska Williams, “Equivalences between triangle and range query problems,” in *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2020, pp. 30–47.
- [11] N. Alon, R. Yuster, and U. Zwick, “Finding and counting given length cycles,” *Algorithmica*, vol. 17, no. 3, pp. 209–223, 1997.
- [12] V. Vassilevska Williams, “Multiplying matrices faster than Coppersmith-Winograd,” in *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, H. J. Karloff and T. Pitassi, Eds. ACM, 2012, pp. 887–898.
- [13] F. L. Gall, “Powers of tensors and fast matrix multiplication,” in *International Symposium on Symbolic and Algebraic Computation, ISSAC '14, Kobe, Japan, July 23-25, 2014*, K. Nabeshima, K. Nagasaka, F. Winkler, and Á. Szántó, Eds. ACM, 2014, pp. 296–303.
- [14] V. Vassilevska, R. Williams, and R. Yuster, “Finding heaviest H -subgraphs in real weighted graphs, with applications,” *ACM Trans. Algorithms*, vol. 6, no. 3, pp. 44:1–44:23, 2010.
- [15] —, “Finding the smallest H -subgraph in real weighted graphs and related problems,” in *Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part I*, ser. Lecture Notes in Computer Science, M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, Eds., vol. 4051, 2006, pp. 262–273.
- [16] A. Lincoln, A. Polak, and V. Vassilevska Williams, “Monochromatic triangles, intermediate matrix products, and convolutions,” in *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
- [17] V. Vassilevska, R. Williams, and R. Yuster, “All-pairs bottleneck paths for general graphs in truly sub-cubic time,” in *Proceedings of the 39th Annual ACM Symposium on Theory of Computing, San Diego, California, USA, June 11-13, 2007*, D. S. Johnson and U. Feige, Eds. ACM, 2007, pp. 585–589.
- [18] R. Duan and S. Pettie, “Fast algorithms for (max, min)-matrix multiplication and bottleneck shortest paths,” in *Proceedings of the twentieth annual ACM-SIAM symposium on Discrete algorithms*. SIAM, 2009, pp. 384–391.
- [19] T. Kopelowitz, S. Pettie, and E. Porat, “Higher lower bounds from the 3sum conjecture,” in *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*. SIAM, 2016, pp. 1272–1287.
- [20] A. Abboud and V. V. Williams, “Popular conjectures imply strong lower bounds for dynamic problems,” in *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*. IEEE, 2014, pp. 434–443.
- [21] F. Le Gall and F. Urrutia, “Improved rectangular matrix multiplication using powers of the coppersmith-winograd tensor,” in *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, 2018, pp. 1029–1046.
- [22] U. Zwick, “All pairs shortest paths using bridging sets and rectangular matrix multiplication,” *J. ACM*, vol. 49, no. 3, pp. 289–317, 2002.
- [23] A. Czumaj, M. Kowaluk, and A. Lingas, “Faster algorithms for finding lowest common ancestors in directed acyclic graphs,” *Theor. Comput. Sci.*, vol. 380, no. 1-2, pp. 37–46, 2007.
- [24] V. Vassilevska Williams, “Problem 2 on problem set 2 of CS367, October 15, 2015,” 2015. [Online]. Available: <http://theory.stanford.edu/~virgi/cs367/hw2.pdf>
- [25] K. Labib, P. Uznański, and D. Wolleb-Graf, “Hamming distance completeness,” *Leibniz International Proceedings in Informatics, LIPIcs*, vol. 128, 2019.
- [26] J. Matoušek, “Computing dominances in E^n ,” *Inf. Process. Lett.*, vol. 38, no. 5, pp. 277–278, 1991.
- [27] R. Yuster, “Efficient algorithms on sets of permutations, dominance, and real-weighted APSP,” in *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2009, New York, NY, USA, January 4-6, 2009*, C. Mathieu, Ed. SIAM, 2009, pp. 950–957.
- [28] R. Duan, C. Jin, and H. Wu, “Faster algorithms for all pairs non-decreasing paths problem,” in *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece*, ser. LIPIcs, vol. 132. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019, pp. 48:1–48:13.
- [29] A. Björklund, R. Pagh, V. V. Williams, and U. Zwick, “Listing triangles,” in *International Colloquium on Automata, Languages, and Programming*. Springer, 2014, pp. 223–234.
- [30] N. Alon, Z. Galil, O. Margalit, and M. Naor, “Witnesses for boolean matrix multiplication and for shortest paths,” in *33rd Annual Symposium on Foundations of Computer Science, Pittsburgh, Pennsylvania, USA, 24-27 October 1992*. IEEE Computer Society, 1992, pp. 417–426.
- [31] V. Vassilevska Williams, “Lecture nodes for lecture 8 of CS367, October 15, 2015,” 2015. [Online]. Available: <http://theory.stanford.edu/~virgi/cs367/lecture4-1.pdf>