

On Diameter Approximation in Directed Graphs

Amir Abboud   

Weizmann Institute of Science, Rehovot, Israel

Mina Dalirrooyfard   

Massachusetts Institute of Technology, Cambridge, MA, USA

Ray Li   

University of California Berkeley, CA, USA

Virginia Vassilevska Williams   

Massachusetts Institute of Technology, Cambridge, MA, USA

Abstract

Computing the diameter of a graph, i.e. the largest distance, is a fundamental problem that is central in fine-grained complexity. In undirected graphs, the Strong Exponential Time Hypothesis (SETH) yields a lower bound on the time vs. approximation trade-off that is quite close to the upper bounds.

In *directed* graphs, however, where only some of the upper bounds apply, much larger gaps remain. Since $d(u, v)$ may not be the same as $d(v, u)$, there are multiple ways to define the problem, the two most natural being the *(one-way) diameter* ($\max_{(u,v)} d(u, v)$) and the *roundtrip diameter* ($\max_{u,v} d(u, v) + d(v, u)$). In this paper we make progress on the outstanding open question for each of them.

- We design the first algorithm for diameter in sparse directed graphs to achieve $n^{1.5-\varepsilon}$ time with an approximation factor better than 2. The new upper bound trade-off makes the directed case appear more similar to the undirected case. Notably, this is the first algorithm for diameter in sparse graphs that benefits from fast matrix multiplication.
- We design new hardness reductions separating roundtrip diameter from directed and undirected diameter. In particular, a 1.5-approximation in subquadratic time would refute the All-Nodes k -Cycle hypothesis, and any $(2 - \varepsilon)$ -approximation would imply a breakthrough algorithm for approximate ℓ_∞ -Closest-Pair. Notably, these are the first conditional lower bounds for diameter that are not based on SETH.

2012 ACM Subject Classification Theory of computation $\rightarrow$ Problems, reductions and completeness; Theory of computation $\rightarrow$ Graph algorithms analysis

Keywords and phrases Diameter, Directed Graphs, Approximation Algorithms, Fine-grained complexity

Digital Object Identifier 10.4230/LIPIcs.ESA.2023.2

Related Version *Full Version:* <https://arxiv.org/abs/2307.07583>

Funding *Amir Abboud:* This project has received funding from the European Research Council (ERC) under the European Union's Horizon Europe research and innovation programme (grant agreement No 101078482). Additionally, Amir Abboud is supported by an Alon scholarship and a research grant from the Center for New Scientists at the Weizmann Institute of Science.

Mina Dalirrooyfard: Partially supported by an Akamai Fellowship.

Ray Li: Supported by the NSF Mathematical Sciences Postdoctoral Research Fellowships Program under Grant DMS-2203067, and a UC Berkeley Initiative for Computational Transformation award.

Virginia Vassilevska Williams: Partially supported by the National Science Foundation Grant CCF-2129139.

Acknowledgements We would like to thank Piotr Indyk, Karthik C.S., and the participants of the Fine-Grained Approximation Algorithms & Complexity Workshop (FG-APX 2019) at Bertinoro 2019 for many helpful discussions.



© Amir Abboud, Mina Dalirrooyfard, Ray Li, and Virginia Vassilevska Williams; licensed under Creative Commons License CC-BY 4.0

31st Annual European Symposium on Algorithms (ESA 2023).

Editors: Inge Li Gørtz, Martin Farach-Colton, Simon J. Puglisi, and Grzegorz Herman; Article No. 2; pp. 2:1–2:17

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

The diameter of the graph is the largest shortest paths distance. A very well-studied parameter with many practical applications (e.g. [23, 36, 45, 15]), its computation and approximation are also among the most interesting problems in Fine-Grained Complexity (FGC). Much effort has gone into understanding the approximation vs. running time tradeoff for this problem (see the survey [43] and the progress after it [14, 13, 34, 35, 28, 25]).

Throughout this introduction we will consider n -vertex and m -edge graphs that, for simplicity, are *unweighted* and *sparse* with $m = n^{1+o(1)}$ edges¹. The diameter is easily computable in $\tilde{O}(mn) = n^{2+o(1)}$ time² by computing All-Pairs Shortest Paths (APSP). One of the first and simplest results in FGC [41, 46] is that any $O(n^{2-\varepsilon})$ time algorithm for $\varepsilon > 0$ for the exact computation of the diameter would refute the well-established Strong Exponential Time Hypothesis (SETH) [30, 18]. Substantial progress has been achieved in the last several years [41, 19, 14, 13, 34, 35, 28, 25], culminating in an approximation/running time lower bound tradeoff based on SETH, showing that even for undirected sparse graphs, for every $k \geq 2$, there is no $2 - 1/k - \delta$ -approximation algorithm running in $\tilde{O}(n^{1+1/(k-1)-\varepsilon})$ time for some $\delta, \varepsilon > 0$.

In terms of upper bounds, the following three algorithms work for both undirected and directed graphs:

1. compute APSP and take the maximum distance, giving an exact answer in $\tilde{O}(n^2)$ time,
2. compute single-source shortest paths from/to an arbitrary node and return the largest distance found, giving a 2-approximation in $\tilde{O}(n)$ time, and
3. an algorithm by [41, 19] giving a 3/2-approximation in $\tilde{O}(n^{1.5})$ time.


For undirected graphs, there are some additional algorithms, given by Cairo, Grossi and Rizzi [17] that qualitatively (but not quantitatively) match the tradeoff suggested by the lower bounds: for every $k \geq 1$ they obtain an $\tilde{O}(n^{1+1/(k+1)})$ time, almost- $(2 - 1/2^k)$ approximation algorithm, meaning that there is also a small constant additive error.

The upper and lower bound tradeoffs for undirected graphs are depicted in Figure 1 ; a gap remains (depicted as white space) because the two trade-offs have different rates. In directed graphs, however, the gap is significantly larger because an upper bound trade-off is missing (the lower bound tradeoff follows immediately because it is a harder problem). One could envision for instance, that the conditional lower bounds for directed diameter could be strengthened to show that if one wants a $(2 - \varepsilon)$ -approximation algorithm, then it must take at least $n^{1.5-o(1)}$ time. Since the work of [17], the main open question (also asked by [43]) for diameter algorithms in directed graphs has been:

Why are there only three approximation algorithms for directed diameter, but undirected diameter has an infinite approximation scheme? Is directed diameter truly harder, or can one devise further approximation algorithms for it?

¹ Notably, however, our algorithmic results hold for general graphs, and our hardness results hold even for very sparse graphs.

² The notation $\tilde{O}(f(n))$ denotes $O(f(n) \text{ poly log}(f(n)))$.



■ **Figure 1** Undirected diameter algorithms and hardness.

Directed is Closer to Undirected

Our first result is that one *can* devise algorithms for directed diameter with truly faster running times than $n^{1.5}$, and approximation ratios between $3/2$ and 2 . It turns out that the directed case has an upper bound tradeoff as well, albeit with a worse rate than in the undirected case. Conceptually, this brings undirected and directed diameter closer together. See Figure 2 for our new algorithms.

► **Theorem 1.** *Let $k = 2^{t+2}$ for a nonnegative integer $t \geq 0$. For every $\varepsilon > 0$ (possibly depending on m), there exists a randomized $2 - \frac{1}{k} + \varepsilon$ -approximation algorithm for the diameter of a directed weighted graphs in time $\tilde{O}(m^{1+\alpha}/\varepsilon)$, for*


$$\alpha = \frac{2(\frac{2}{\omega-1})^t - \frac{(\omega-1)^2}{2}}{(\frac{2}{\omega-1})^t(7-\omega) - \frac{\omega^2-1}{2}}.$$

The constant $2 \leq \omega < 2.37286$ in the theorem refers to the fast matrix multiplication exponent [6]. A surprising feature of our algorithms is that we utilize fast matrix multiplication techniques to obtain faster algorithms for a problem in sparse graphs. Prior work on shortest paths has often used fast matrix multiplication to speed-up computations, but to our knowledge, all of this work is for dense graphs (e.g. [7, 44, 47, 24]). Breaking the $n^{1.5}$ bound with a combinatorial algorithm is left as an open problem.

Roundtrip is Harder

One unsatisfactory property of the shortest paths distance measure in directed graphs is that it is not symmetric ($d(u, v) \neq d(v, u)$) and is hence not a metric. Another popular distance measure used in directed graphs that is a metric is the roundtrip measure. Here the *roundtrip distance* $\tilde{d}(u, v)$ between vertices u, v is $d(u, v) + d(v, u)$.

Roundtrip distances were first studied in the distributed computing community in the 1990s [22]. In recent years, powerful techniques were developed to handle the fast computation of sparse roundtrip spanners, and approximations of the *minimum* roundtrip distance, i.e.



■ **Figure 2** Directed diameter algorithms and hardness. All tradeoffs hold for both weighted and unweighted graphs (though citations may differ for weighted vs. unweighted).

the shortest cycle length, the girth, of a directed graph [38, 21, 26, 20]. These techniques give hope for new algorithms for the *maximum* roundtrip distance, the roundtrip diameter of a directed graph.

Only the first *two* algorithms in the list in the beginning of the introduction work for roundtrip diameter: compute an exact answer by computing APSP, and a linear time 2-approximation that runs SSSP from/to an arbitrary node. These two algorithms work for any distance *metric*, and surprisingly there have been no other algorithms developed for roundtrip diameter. The only fine-grained lower bounds for the problem are the ones that follow from the known lower bounds for diameter in undirected graphs, and these cannot explain why there are no known subquadratic time algorithms that achieve a better than 2-approximation.

Are there $O(n^{2-\varepsilon})$ time algorithms for roundtrip diameter in sparse graphs that achieve a $2 - \delta$ -approximation for constants $\varepsilon, \delta > 0$?

This question was considered e.g. by [4] who were able to obtain a hardness result for the related roundtrip radius problem, showing that under a popular hypothesis, such an algorithm for roundtrip radius does not exist. One of the main questions studied at the “Fine-Grained Approximation Algorithms and Complexity Workshop” at Bertinoro in 2019 was to obtain new algorithms or hardness results for roundtrip diameter. Unfortunately, however, no significant progress was made, on either front.

The main approach to obtaining hardness for roundtrip diameter, was to start from the Orthogonal Vectors (OV) problem and reduce it to a gap version of roundtrip diameter, similar to all known reductions to (other kinds of) diameter approximation hardness. Unfortunately, it has been difficult to obtain a reduction from OV to roundtrip diameter that has a larger gap than that for undirected diameter; in Section 4.1 we give some intuition for why this is the case.

In this paper we circumvent the difficulty by giving stronger hardness results for roundtrip diameter starting from *different* problems and hardness hypotheses. We find this intriguing because all previous conditional lower bounds for (all variants of) the diameter problem were

based on SETH. In particular, it gives a new approach for resolving the remaining gaps in the undirected case, where higher SETH-based lower bounds are provably impossible (under the so-called NSETH) [35].

Our first negative result conditionally proves that any $5/3 - \varepsilon$ approximation for roundtrip requires $n^{2-o(1)}$ time; separating it from the undirected and the directed one-way cases where a 1.5-approximation in $\tilde{O}(n^{1.5})$ time is possible. This result is based on a reduction from the so-called All-Nodes k -Cycle problem.

► **Definition 2** (All-Nodes k -Cycle in Directed Graphs). *Given a k partite directed graph $G = (V, E)$, $V = V_1 \cup \dots \cup V_k$, whose edges go only between “adjacent” parts $E \subseteq \bigcup_{i=1}^k V_i \times V_{i+1 \bmod k}$, decide if all nodes $v \in V_1$ are contained in a k -cycle in G .*

This problem can be solved for all k in time $O(nm)$, e.g. by running an APSP algorithm, and in subquadratic $O(m^{2-1/k})$ for any fixed k [8]. Breaking the quadratic barrier for super-constant k has been a longstanding open question; we hypothesize that it is impossible.

► **Hypothesis 3.** *No algorithm can solve the All-Nodes k -Cycle problem in sparse directed graphs for all $k \geq 3$ in $O(n^{2-\delta})$ time, with $\delta > 0$.*

Similar hypotheses have been used in recent works [5, 37, 10, 40]. The main difference is that we require all nodes in V_1 to be in cycles; such variants of hardness assumptions that are obtained by changing a quantifier in the definition of the problem are popular, see e.g. [4, 16, 1].

► **Theorem 4.** *Under Hypothesis 3, for all $\varepsilon, \delta > 0$, no algorithm can $5/3 - \varepsilon$ approximate the roundtrip diameter of a sparse directed unweighted graph in $O(n^{2-\delta})$ time.*

We are thus left with a gap between the linear time factor-2 upper bound and the subquadratic factor-5/3 lower bound. A related problem with a similar situation is the problem of computing the eccentricity of all nodes in an undirected graph [4]; there, 5/3 is the right number because one can indeed compute a 5/3-approximation in subquadratic time [19]. Could it be the same here?

Alas, our final result is a reduction from the following classical problem in geometry to roundtrip diameter, establishing a barrier for any better-than-2 approximation in subquadratic time.

► **Definition 5** (Approximate ℓ_∞ Closest-Pair). *Let $\alpha > 1$. The α -approximate ℓ_∞ Closest-Pair (CP) problem is, given n vectors $v_1, \dots, v_n$ of some dimension d in $\mathbb{R}^n$, determine if there exists v_i and v_j with $\|v_i - v_j\|_\infty \leq 1$, or if for all v_i and v_j , $\|v_i - v_j\|_\infty \geq \alpha$.*

Closest-pair problems are well-studied in various metrics; the main question being whether the naive n^2 bound can be broken (when d is assumed to be $n^{o(1)}$). For ℓ_∞ specifically, a simple reduction from OV proves a quadratic lower bound for $(2 - \varepsilon)$ -approximations [31]; but going beyond this factor with current reduction techniques runs into a well-known “triangle-inequality” barrier (see [42, 33]). This leaves a huge gap from the upper bounds that can only achieve $O(\log \log n)$ approximations in subquadratic time [31]. Cell-probe lower bounds for the related nearest-neighbors problem suggest that this log-log bound may be optimal [11]; if indeed constant approximations are impossible in subquadratic time then the following theorem implies a tight lower bound for roundtrip diameter.

► **Theorem 6.** *If for some $\alpha \geq 2, \varepsilon > 0$ there is a $2 - \frac{1}{\alpha} - \varepsilon$ approximation algorithm in time $O(m^{2-\varepsilon})$ for roundtrip diameter in unweighted graphs, then for some $\delta > 0$ there is an α -approximation for ℓ_∞ -Closest-Pair with vectors of dimension $d \leq n^{1-\delta}$ in time $\tilde{O}(n^{2-\delta})$.*

of the diameter in directed unweighted graphs in time $O(m^{1.458})$. In Section 4.1 we give an overview of the hardness reductions. In Section 4.2, we prove a weakening of Theorem 6 that only holds for weighted graphs.

2 Preliminaries

All logs are base e unless otherwise specified. For reals $a \geq 0$, let $[\pm a]$ denote the real interval $[-a, a]$. For a boolean statement φ , let $\mathbf{1}[\varphi]$ be 1 if φ is true and 0 otherwise.

For a vertex v in a graph, let $\deg(v)$ denote its degree. For $r \geq 0$, let $B_r^{\text{in}}(v) = \{u : d(u, v) \leq r\}$ be the in-ball of radius r around v , and let $B_r^{\text{out}}(v) = \{u : d(v, u) \leq r\}$ be the out-ball of radius r around v . For $r \geq 0$, let $B_r^{\text{in}+}(v)$ be $B_r^{\text{in}}(v)$ and their in-neighbors, and let $B_r^{\text{out}+}(v)$ be $B_r^{\text{out}}(v)$ and their out-neighbors.

Throughout, let $\omega \leq 2.3728596$ denote the matrix multiplication constant. We use the following lemma which says that we can multiply sparse matrices quickly.

► **Lemma 7** (see e.g. Theorem 2.5 of [32]). *We can multiply a $a \times b$ and a $b \times a$ matrix, each with at most ac nonzero entries, in time $O(ac \cdot a^{\frac{\omega-1}{2}})$.⁴*

We repeatedly use the following standard fact.

► **Lemma 8.** *Given two sets $B \subset V$ with B of size k and V of size $2m$, a set of $4(m/k) \log m$ uniformly random elements of V contains an element of B with probability at least $1 - \frac{1}{m^2}$.*

Proof. The probability that B is not hit is $(1 - \frac{k}{2m})^{4m/k \log m} \leq e^{-2 \log m} = \frac{1}{m^2}$. ◀


3 7/4-approximation of directed (one-way) diameter

In this section, we prove Theorem 1 in the special case of $t = 0$ and unweighted graphs. That is, we give a 7/4-approximation of the (one-way) diameter of a directed unweighted graph in $O(m^{1.4575})$ time. For the rest of this section, let $\alpha = \frac{\omega+1}{\omega+5} \leq 0.4575$.

Before stating the algorithm and proof, we highlight how our algorithm differs from the undirected algorithm of [17]. At a very high level, all known diameter approximation algorithms compute some pairs of distances, and use the triangle inequality to infer other distances, saving runtime. Approximating diameter in directed graphs is harder than in undirected graphs because distances are not symmetric, so we can only use the triangle inequality “one way.” For example, we always have $d(x, y) + d(y, z) \geq d(x, z)$, but not necessarily $d(x, y) + d(z, y) \geq d(x, z)$. The undirected algorithm [17] crucially uses the triangle inequality “both ways,” so it was not clear whether their algorithm could be adapted to the directed case. We get around this barrier using matrix multiplication together with the triangle inequality to infer distances quickly. We consider the use of matrix multiplication particularly interesting because, previously, matrix multiplication had only been used for diameter in dense graphs, but we leverage it in sparse graphs.

► **Theorem 9.** *Let $\alpha = \frac{\omega+1}{\omega+5}$. There exists a randomized 7/4-approximation algorithm for the diameter of an unweighted directed graph running in $\tilde{O}(m^{1+\alpha})$ time.*

⁴ In [32], this runtime of $O(ac \cdot a^{\frac{\omega-1}{2}})$ is stated only for the case $ac > a^{(\omega+1)/2}$. However, the runtime bound for this case works for other cases as well so the lemma is correct for all matrices.



■ **Figure 4** Steps 5 and 6. If $d(a, b) \geq D$ and Steps 2, 3, and 4 do not accept, with high probability, set $\hat{S}$ hits the $D/7$ out- and in- neighborhoods of a and b at vertices s and s' , respectively, that must have distance at least $5D/7$ by the triangle inequality. Thus, checking all pairs of distances in $S^{out} \times S^{in}$, which can be done quickly with sparse matrix multiplication, distinguishes at Step 6 whether the diameter is at least D or less than $4D/7$.

Proof. It suffices to show that, for any positive integer $D > 0$, there exists an algorithm $\mathcal{A}_D$ running in time $\tilde{O}(m^{1+\alpha})$ that takes as input any graph and accepts if the diameter is at least D , rejects if the diameter is less than $4D/7$, and returns arbitrarily otherwise. Then, we can find the diameter up to a factor of $7/4$ by running binary search with $\mathcal{A}_D$,⁵ which at most adds a factor of $O(\log n)$.

We now describe the algorithm $\mathcal{A}_D$. The last two steps, illustrated in Figure 4 contain the key new ideas.

1. First, we apply a standard trick that replaces the input graph on n vertices and m edges with an $2m$ -vertex graph of max-degree-3 that preserves the diameter: replace each vertex v with a $\deg(v)$ -vertex cycle of weight-0 edges and where the edges to v now connect to distinct vertices of the cycle. From now on, we work with this max-degree-3 graph on $2m$ vertices.
2. Sample $4m^\alpha \log m$ uniformly random vertices and compute each vertex's in- and out-eccentricity. If any such vertex has (in- or out-) eccentricity at least $4D/7$ Accept.
3. For every vertex v , determine if $|B_{D/7}^{out}(v)| \leq m^\alpha$. If such a vertex v exists, determine if any vertex in $B_{D/7}^{out+}(v)$ has eccentricity at least $4D/7$, and Accept if so.
4. For every vertex v , determine if $|B_{D/7}^{in}(v)| \leq m^\alpha$. If such a vertex v exists, determine if any vertex in $B_{D/7}^{in+}(v)$ has eccentricity at least $4D/7$, and Accept if so.
5. Sample $4m^{1-\alpha} \log m$ uniformly random vertices $\hat{S}$. Let $S^{out} = \{s \in \hat{S} : |B_{2D/7}^{out}(s)| \leq m^{1-\alpha}\}$ and $S^{in} = \{s \in \hat{S} : |B_{2D/7}^{in}(s)| \leq m^{1-\alpha}\}$. Compute $B_{2D/7}^{out}(s)$ and $B_{2D/7}^{out+}(s)$ for $s \in S^{out}$, and $B_{2D/7}^{in}(s)$ and $B_{2D/7}^{in+}(s)$ for $s \in S^{in}$.
6. Let $A^{out} \in \mathbb{R}^{S^{out} \times V}$ be the $|S^{out}| \times n$ matrix where $A_{s,v} = \mathbf{1}[v \in B_{2D/7}^{out}(s)]$. Let $A^{in} \in \mathbb{R}^{V \times S^{in}}$ be the $n \times |S^{in}|$ matrix where $A_{v,s}^{in} = \mathbf{1}[v \in B_{2D/7}^{in}(s)]$ if $\lfloor 4D/7 \rfloor = 2\lfloor 2D/7 \rfloor$ and $A_{v,s}^{in} = \mathbf{1}[v \in B_{2D/7}^{in+}(s)]$ otherwise. Compute $A^{out} \cdot A^{in} \in \mathbb{R}^{S^{out} \times S^{in}}$ using sparse matrix multiplication. If the product has any zero entries, Accept, otherwise Reject.

⁵ We have to be careful not to lose a small additive factor. Here are the details: Let D^* be the true diameter. Initialize $hi = n, lo = 0$. Repeat until $hi - lo = 1$: let $mid = \lfloor (hi + lo)/2 \rfloor$, run $\mathcal{A}_{mid}$, if accept, set $lo = mid$, else $hi = mid$. One can check that $hi \geq D^* + 1$ and $lo \leq 7D^*/4$ always hold. If we return lo after the loop breaks, the output is always in $[D^*, 7D^*/4]$.

Runtime. Computing a single eccentricity takes time $O(m)$, so Step 2 takes time $\tilde{O}(m^{1+\alpha})$. For Step 3 checking if $|B_{D/7}^{out}(v)| \leq m^\alpha$ takes $O(m^\alpha)$ time for each v via a partial Breadth-First-Search (BFS). Here we use that the max-degree is 3. If $|B_{D/7}^{out}(v)| \leq m^\alpha$, there are at most $3m^\alpha$ eccentricity computations which takes time $O(m^{1+\alpha})$. Step 4 takes time $O(m^{1+\alpha})$ for the same reason. Similarly, we can complete Step 5 by running partial BFS for each $s \in \hat{S}$ until $m^{1-\alpha}$ vertices are visited. This gives S^{out} and S^{in} and also gives $B_{2D/7}^{out}(s)$ and $B_{2D/7}^{out+}(s)$ for $s \in S^{out}$ and $B_{2D/7}^{in}(s)$ and $B_{2D/7}^{in+}(s)$ for $s \in S^{in}$. For Step 6, the runtime is the time to multiplying sparse matrices. Matrix A^{out} has at most $|\hat{S}| \leq 4m^{1-\alpha} \log m$ rows each with at most $\max_{s \in S^{out}} |B_{2D/7}^{out}(s)| \leq m^{1-\alpha}$ entries, and similarly A^{in} has at most $4m^{1-\alpha} \log m$ columns each with at most $\max_{s \in S^{in}} |B_{2D/7}^{in+}(s)| \leq 3m^{1-\alpha}$ entries. The sparse matrix multiplication takes time $\tilde{O}(m^{(2-2\alpha)} \cdot m^{(1-\alpha)\frac{\omega-1}{2}}) = \tilde{O}(m^{1+\alpha})$ by Lemma 7 with $a = m^{1-\alpha}, b = n, c = m^{1-\alpha}$.

If the Diameter is less than $4D/7$, we always reject. Clearly every vertex has eccentricity less than $4D/7$, so we indeed do not accept at Steps 2, 3, and 4. In Step 5, we claim for every $s \in S^{out}, s' \in S^{in}$ there exists v such that $A_{s,v}^{out} = A_{v,s'}^{in} = 1$, so that $(A^{out} \cdot A^{in})_{s,s'} \geq 1$ for all $s \in S^{out}$ and $s' \in S^{in}$ and thus we reject. Fix $s \in S^{out}$ and $s' \in S^{in}$. By the diameter bound, $d(s, s') \leq \lfloor 4D/7 \rfloor$. Let v be the last vertex on the s -to- s' shortest path such that $d(s, v) \leq \lfloor 2D/7 \rfloor$, and, if it exists, let v' be the vertex after v . Clearly $A_{s,v}^{out} = 1$. We show $A_{v,s'}^{in} = 1$ as well. If $v = s'$, then clearly $v \in B_{2D/7}^{in}(s')$ so $A_{v,s'}^{in} = 1$ as desired. Otherwise $d(s, v) = \lfloor 2D/7 \rfloor$. If $\lfloor 4D/7 \rfloor = 2\lfloor 2D/7 \rfloor$, then $d(v, s') \leq d(s, s') - d(s, v) \leq \lfloor 4D/7 \rfloor - \lfloor 2D/7 \rfloor = \lfloor 2D/7 \rfloor$, so $v \in B_{2D/7}^{in}(s')$ and $A_{v,s'}^{in} = 1$, so again $A_{v,s'}^{in} = 1$. If $\lfloor 4D/7 \rfloor = 2\lfloor 2D/7 \rfloor + 1$, then $d(v', s') \leq d(s, s') - d(s, v') \leq \lfloor 4D/7 \rfloor - (\lfloor 2D/7 \rfloor + 1) = \lfloor 2D/7 \rfloor$, so $v' \in B_{2D/7}^{in}(s')$ and thus $v \in B_{2D/7}^{in+}(s')$ and $A_{v,s'}^{in} = 1$, as desired. This covers all cases, so we've shown we reject.

If the Diameter is at least D , we accept with high probability. Let a and b be vertices with $d(a, b) \geq D$.

If $|B_{3D/7}^{out}(a)| > m^{1-\alpha}$, Step 2 computes the eccentricity of some $v \in B_{3D/7}^{out}(a)$ with high probability (by Lemma 8), which is at least $d(v, b) \geq d(a, b) - d(a, v) \geq 4D/7$ by the triangle inequality, so we accept. Similarly, we accept with high probability if $|B_{3D/7}^{in}(b)| > m^{1-\alpha}$. Thus we may assume that $|B_{3D/7}^{out}(a)|, |B_{3D/7}^{in}(b)| \leq m^{1-\alpha}$ for the rest of the proof.

If $|B_{D/7}^{out}(v)| \leq m^\alpha$ for any vertex v , then either (i) $d(v, b) \geq 4D/7$, in which case v has eccentricity at least $4D/7$ and we accept at Step 3, or (ii) $d(v, b) \leq 4D/7$, in which case there is a vertex $u \in B_{D/7}^{out+}(v)$ on the v -to- b path with $d(u, b) \leq 3D/7$ (take the $u \in B_{D/7}^{out+}(v)$ closest to b on the path). Then $d(a, u) \geq 4D/7$ by the triangle inequality and we accept in Step 3 as we perform a BFS from u . Thus we may assume $|B_{D/7}^{out}(v)| > m^\alpha$ for all vertices v . Similarly, because of Step 4, we may assume $|B_{D/7}^{in}(v)| > m^\alpha$ for all vertices v .

In particular, we may assume $|B_{D/7}^{out}(a)| > m^\alpha$ and $|B_{D/7}^{in}(b)| > m^\alpha$. Figure 4 illustrates this last step. Then $\hat{S}$ hits $B_{D/7}^{out}(a)$ with high probability (by Lemma 8), so $B_{D/7}^{out}(a)$ has some $s \in \hat{S}$ with high probability, and similarly $B_{D/7}^{in}(b)$ has some $s' \in \hat{S}$ with high probability. The triangle inequality implies that $B_{2D/7}^{out}(s) \subset B_{3D/7}^{out}(a)$, so $|B_{2D/7}^{out}(s)| \leq |B_{3D/7}^{out}(a)| \leq m^{1-\alpha}$ and thus $s \in S^{out}$. Similarly $s' \in S^{in}$. By the triangle inequality, we have $d(s, s') \geq d(a, b) - d(a, s) - d(s', b) \geq D - D/7 - D/7 = 5D/7$. Then we must have $(A \cdot B)_{s,s'} = 0$, as otherwise there is a v such that $d(s, v) \leq \lfloor 2D/7 \rfloor$ and $d(v, s') \leq 4D/7 - \lfloor 2D/7 \rfloor$, contradicting $d(s, s') \geq 5D/7$. Hence, we accept at step 5, as desired. $\blacktriangleleft$

4 Hardness Reductions for Roundtrip

4.1 Overview

In this paper we prove hardness results for roundtrip diameter that go beyond the 2 vs. 3 barrier. Before presenting the proofs, let us begin with an abstract discussion on why this barrier arises and (at a high level) how we overcome it.

All previous hardness results for diameter are by reductions from OV (or its generalization to multiple sets). In OV, one is given two sets of vectors of size n and dimension $d = \text{poly log } n$, A and B , and one needs to determine whether there are $a \in A, b \in B$ that are orthogonal. SETH implies that OV requires $n^{2-o(1)}$ time [46]. In a reduction from OV to a problem like diameter, one typically has nodes representing the vectors in A and B , as well as nodes C representing the coordinates, and if there is an orthogonal vector pair a, b , then the corresponding nodes in the diameter graph are far (distance ≥ 3), and otherwise *all* pairs of nodes are close (distance ≤ 2). Going beyond the 2 vs. 3 gap is difficult because each node $a \in A$ must have distance ≤ 2 to each coordinate node in C , regardless of the existence of an orthogonal pair, and then it is automatically at distance $2 + 1$ from any node $b \in B$ because each b has at least one neighbor in C . So even if a, b are orthogonal, the distance will not be more than 3.

The key trick for proving a higher lower bound (say 3 vs. 5) for roundtrip is to have two sets of coordinate nodes, a C^{fwd} set that can be used to go *forward* from A to B , and a C^{bwd} set that can be used to go back. The default roundtrip paths from A/B to each of these two sets will have different forms, and this asymmetry will allow us to overcome the above issue. This is inspired by the difficulty that one faces when trying to make the subquadratic 3/2-approximation algorithms for undirected and directed diameter work for roundtrip.

Unfortunately, there is another (related) issue when reducing from OV. First notice that all nodes within A and within B must always have small distance (or else the diameter would be large). This can be accomplished simply by adding direct edges of weight 1.5 between all pairs (within A and within B); but this creates a dense graph and makes the quadratic lower bound uninteresting. Instead, such reductions typically add auxiliary nodes to simulate the n^2 edges more cheaply, e.g. a star node o that is connected to all of A . But then the node o must have small distance to B , decreasing all distances between A and B .

Overcoming this issue by a similar trick seems impossible. Instead, our two hardness results bypass it in different ways.

The reduction from ℓ_∞ -Closest-Pair starts from a problem that is defined over *one set of vectors* A (not two) which means that the coordinates are “in charge” of connecting all pairs within A . We remark that while OV can also be defined over one set (monochromatic) instead of two (bichromatic) and that it remains SETH hard; that would prevent us from applying the above trick of having a forward and a backward sets of coordinate nodes. Our reduction in Section 4.2 is able to utilize the structure of the metric in order to make both ideas work simultaneously.

The reduction from All-Node k -Cycle relies on a different idea: it uses a construction where only a small set of n pairs $a_i \in A, b_i \in B$ are “interesting” in the sense that we do not care about the distances for other pairs (in order to solve the starting problem). Then the goal becomes to connect all pairs within A and within B by short paths, without decreasing the distance for the (a_i, b_i) pairs. A trick similar to the *bit-gadget* [3, 2] does the job. For the complete reduction see the full version of the paper.

4.2 Weighted Roundtrip $2 - \varepsilon$ hardness from ℓ_∞ -CP

In this section, we highlight the key ideas in Theorem 6 by proving a weaker version, showing the lower bound for weighted graphs. See the full version of the paper for the extension to unweighted graphs.

The main technical lemma is showing that to α -approximate ℓ_∞ -Closest-Pair, it suffices to do so on instances where all vector coordinates are in $[\pm(0.5 + \varepsilon)\alpha]$. Towards this goal, we make the following definition.

► **Definition 10.** *The α -approximate β -bounded ℓ_∞ -Closest-Pair problem is, given n vectors $v_1, \dots, v_n$ of dimension d in $[-\beta, \beta]^d$ determine if there exists v_i and v_j with $\|v_i - v_j\|_\infty \leq 1$, or if for all v_i and v_j , $\|v_i - v_j\|_\infty \geq \alpha$.*

We now prove the main technical lemma.

► **Lemma 11.** *Let $\varepsilon \in (0, 1/2)$ and $\alpha > 1$. If one can solve α -approximate $(0.5 + \varepsilon)\alpha$ -bounded ℓ_∞ -CP on dimension $O(d\varepsilon^{-1} \log n)$ in time T , then one can solve α -approximate ℓ_∞ -CP on dimension d in time $T + O_\varepsilon(dn \log n)$, where in $O_\varepsilon(\cdot)$ we neglect dependencies on ε .*

Proof. Start with an ℓ_∞ instance $\Phi = (v_1, \dots, v_n)$. We show how to construct a bounded ℓ_∞ instance Φ' such that Φ has two vectors with ℓ_∞ distance ≤ 1 if and only if Φ' has two vectors with ℓ_∞ distance ≤ 1 .

First we show we may assume that $v_1, \dots, v_n$ are on domain $[0, \alpha n]$. Suppose that $x \in [d]$. Reindex $v_1, \dots, v_n$ in increasing order of $v_i[x]$ (by sorting). Let $v'_1, \dots, v'_n$ be vectors identical to $v_1, \dots, v_n$ except in coordinate x , where instead

$$v'_i[x] = \sum_{j=0}^{i-1} \min(\alpha, v_{j+1}[x] - v_j[x])$$

for $i = 1, \dots, n$, where the empty sum is 0. We have that $v'_i[x] \leq \alpha n$ for all i , and furthermore $|v'_i[x] - v'_j[x]| \geq \alpha$ if and only if $|v_i[x] - v_j[x]| \geq \alpha$ and also $|v'_i[x] - v'_j[x]| \leq 1$ if and only if $|v_i[x] - v_j[x]| \leq 1$. Hence, the instance given by $v'_1, \dots, v'_n$ is a YES instance if and only if the instance Φ is a YES instance, and is a NO instance if and only if the instance Φ is a NO instance. Repeating this with all other coordinates x gives an instance Φ' such that Φ' is a YES instance if and only if Φ is a YES instance, and Φ' is a NO instance if and only if Φ is a NO instance, and furthermore Φ' has vectors on $[0, \alpha n]$.

Now we show how to construct an ℓ_∞ -CP instance in dimension $O_\varepsilon(d \log n)$ vectors with coordinates in $[\pm(0.5 + \varepsilon)\alpha]$.

► **Lemma 12.** *Let $\varepsilon \in (0, 0.5)$ and $\alpha > 1$. For any real number M , there exists two maps $g : [0, M] \rightarrow [-(0.5 + \varepsilon)\alpha, (0.5 + \varepsilon)\alpha]^{2^{\lceil \varepsilon^{-1} \rceil + 1}}$ and $h : [0, M] \rightarrow [0, M/2]$ such that for all $a, b \in [0, M]$, we have $\min(|a - b|, \alpha) = \min(\|(g(a), h(a)) - (g(b), h(b))\|_\infty, \alpha)$. (here, $(g(\cdot), h(\cdot))$ is a length $2^{\lceil \varepsilon^{-1} \rceil} + 2$ vector.) Furthermore, g and h can be computed in $O_\varepsilon(1)$ time.*

Proof. It suffices to consider when ε^{-1} is an integer. Let $f_z : \mathbb{R} \rightarrow [-(0.5 + \varepsilon)\alpha, (0.5 + \varepsilon)\alpha]$ be the piecewise function

$$f_z(x) = \begin{cases} -(0.5 + \varepsilon)\alpha & \text{if } x \leq z - (0.5 + \varepsilon)\alpha \\ (0.5 + \varepsilon)\alpha & \text{if } x \geq z + (0.5 + \varepsilon)\alpha \\ x - z & \text{otherwise.} \end{cases}$$

2:12 On Diameter Approximation in Directed Graphs

For $a \in [M]$, define $g(a) \in \mathbb{R}^{2\varepsilon^{-1}+1}$ and $h(a) \in \mathbb{R}$ as follows, where we index coordinates by $-\varepsilon^{-1}, \dots, -1, 0, 1, \varepsilon^{-1}$ for convenience

$$\begin{aligned} g(a)_i &= f_{M/2+0.5i\varepsilon\alpha}(a) \text{ for } -\varepsilon^{-1} \leq i \leq \varepsilon^{-1} \\ h(a) &= |a - M/2|. \end{aligned}$$

Clearly g and h have the correct codomain, and they can be computed in $O_\varepsilon(1)$ time. Additionally, note that $f_z(x)$ and $|x - M/2|$ are 1-Lipschitz functions of x for all z , so g is a Lipschitz function and thus $\|g(a) - g(b)\|_\infty \leq |a - b|$.

Now, it suffices to show that $\min(\|(g(a), h(a)) - (g(b), h(b))\|_\infty, \alpha) \geq \min(|a - b|, \alpha)$. If a and b are on the same side of $M/2$, then $\|h(a) - h(b)\|_\infty \geq ||a - M/2| - |b - M/2|| = |a - b|$, as desired. Now suppose a and b are on opposite sides of $M/2$, and without loss of generality $a < M/2 < b$. Let $0 \leq i \leq \varepsilon^{-1}$ be the largest integer such that $a \leq M/2 - i\varepsilon\alpha$ ($i = 0$ works so i always exists). If $i = \varepsilon^{-1}$, then $a < M/2 - \alpha$ and

$$\|g(a) - g(b)\|_\infty \geq f_{M/2-0.5\alpha}(b) - f_{M/2-0.5\alpha}(a) \geq 0.5\alpha - (-0.5\alpha) = \alpha \geq \min(|a - b|, \alpha),$$

as desired. Now assume $i < \varepsilon^{-1}$. Let $z = M/2 + (0.5 - i\varepsilon)\alpha$. By maximality of i , we have $a - z \in [-(0.5 + \varepsilon)\alpha, -0.5\alpha]$. We have $g(\cdot)_{\varepsilon^{-1}-2i} = f_z(\cdot)$ by definition of g . By the definition of $f_z(\cdot)$, since $a \in [z - (0.5 + \varepsilon)\alpha, z - 0.5\alpha]$ and $b \geq a$, we have $\min(f_z(b) - f_z(a), \alpha) = \min(b - a, \alpha)$. Thus,

$$\begin{aligned} \min(\|g(a) - g(b)\|_\infty, \alpha) &\geq \min(g(b)_{\varepsilon^{-1}-2i} - g(a)_{\varepsilon^{-1}-2i}, \alpha) \\ &= \min(f_z(b) - f_z(a), \alpha) = \min(b - a, \alpha), \end{aligned}$$

as desired. In either case, we have $\min(\|g(a) - g(b)\|_\infty, \alpha) \geq \min(|a - b|, \alpha)$, so we conclude that $\min(\|g(a) - g(b)\|_\infty, \alpha) = \min(|a - b|, \alpha)$. $\blacktriangleleft$

Iterating Lemma 12 gives the following.

► **Lemma 13.** *Let $\varepsilon \in (0, 1/2)$. There exists a map $g : [0, \alpha n] \rightarrow [\pm(0.5 + \varepsilon)\alpha]^{4\lceil\varepsilon^{-1}\rceil \log n}$ such that for all $a, b \in [0, \alpha n]$, we have $\min(|a - b|, \alpha) = \min(\|g(a) - g(b)\|_\infty, \alpha)$. Furthermore, g can be computed in $O_\varepsilon(\log n)$ time.*


Proof. For $\ell = 1, \dots$, let $M_\ell = \alpha n / 2^{\ell-1}$, and let $g_\ell^* : [M_\ell] \rightarrow [\pm(0.5 + \varepsilon)\alpha]^{2\lceil\varepsilon^{-1}\rceil+1}$ and $h_\ell^* : [M_\ell] \rightarrow [M_{\ell+1}]$ be the functions given by Lemma 12. For $\ell = 0, 1, \dots$, let $g_\ell : [0, \alpha n] \rightarrow [-(0.5 + \varepsilon)\alpha, (0.5 + \varepsilon)\alpha]^{\ell(2\lceil\varepsilon^{-1}\rceil+1)}$ and $h_\ell : [0, \alpha n] \rightarrow [0, \alpha n / 2^\ell]$ be such that $g_0(x) = ()$ is an empty vector, $h_0(x) = x$ is the identity, and for $\ell \geq 1$, $g_\ell(x) = (g_{\ell-1}(x), g_\ell^*(h_{\ell-1}(x)))$ and $h_\ell(x) = h_\ell^*(h_{\ell-1}(x))$. By Lemma 12, we have that

$$\begin{aligned} &\min(\|(g_{\ell-1}(a), h_{\ell-1}(a)) - (g_{\ell-1}(b), h_{\ell-1}(b))\|_\infty, \alpha) \\ &= \min(\|(g_{\ell-1}(a), g_\ell^*(h_{\ell-1}(a)), h_\ell^*(h_{\ell-1}(a))) - (g_{\ell-1}(b), g_\ell^*(h_{\ell-1}(b)), h_\ell^*(h_{\ell-1}(b)))\|_\infty, \alpha) \\ &= \min(\|(g_\ell(a), h_\ell(a)) - (g_\ell(b), h_\ell(b))\|_\infty, \alpha) \end{aligned}$$

for all ℓ . For $\ell = \lceil \log n \rceil$, the vector $g(a) \stackrel{\text{def}}{=} (g_\ell(a), h_\ell(a) - 0.5\alpha)$ has every coordinate in $[\pm(0.5 + \varepsilon)\alpha]$, and by (4.2), we have

$$\begin{aligned} \min(|a - b|, \alpha) &= \min(|g_0(a) - g_0(b)|, \alpha) \\ &= \min(|g_\ell(a) - g_\ell(b)|, \alpha) = \min(\|g(a) - g(b)\|_\infty, \alpha), \end{aligned}$$

as desired. The length of this vector is at most $\lceil \log n \rceil (2\lceil\varepsilon^{-1}\rceil + 1) + 1$, which we bound by $4\lceil\varepsilon^{-1}\rceil \log n$ for simplicity (and pad the corresponding vectors with zeros). $\blacktriangleleft$



■ **Figure 5** The roundtrip diameter instance G for ℓ_∞ -CP hardness.

To finish, let $g : [0, \alpha n] \rightarrow [\pm(0.5 + \varepsilon)\alpha]$ be given by Lemma 13, and let the original ℓ_∞ instance be $v_1, \dots, v_n$. Let the new $(0.5 + \varepsilon)\alpha$ -bounded ℓ_∞ instance be $w_i = (g(v_i[x]))_{x \in [d]}$ of length $4d \lceil \varepsilon^{-1} \rceil \log n$. ◀

We now prove our goal for this section, Theorem 6 for weighted graphs.

► **Theorem 14.** *If for some $\alpha \geq 2, \varepsilon > 0$ there is a $2 - \frac{1}{\alpha} - \varepsilon$ approximation algorithm in time $O(m^{2-\varepsilon})$ for roundtrip diameter in weighted graphs, then for some $\delta > 0$ there is an α -approximation for ℓ_∞ -Closest-Pair with vectors of dimension $d \leq n^{1-\delta}$ in time $\tilde{O}(n^{2-\delta})$.*

Proof. By Lemma 11 it suffices to prove that there exists an $O(n^{2-\delta})$ time algorithm for α -approximate $(0.5 + \varepsilon)\alpha$ -bounded ℓ_∞ -CP for $\varepsilon = (4\alpha)^{-1}$.

Let Φ be the bounded-domain ℓ_∞ -CP instance with vectors $v_1, \dots, v_n \in [\pm(0.5 + \varepsilon)\alpha]^n$. Then construct a graph G (see Figure 5) with vertex set $S \cup X_1 \cup X_2$ where $X_1 = X_2 = [d]$ and $S = [n]$. We identify vertices with the notations i_S, x_{X_1} , and x_{X_2} , for $i \in [n]$ and $x \in [d]$. Draw directed edges

1. from i_S to x_{X_1} , of weight $\alpha + v_i[x]$,
2. from x_{X_1} to i_S , of weight $\alpha - v_i[x]$,
3. from i_S to x_{X_2} , of weight $\alpha - v_i[x]$,
4. from x_{X_2} to i_S , of weight $\alpha + v_i[x]$, and
5. between any two vertices in $X_1 \cup X_2$, of weight α .

Note that all edge weights are nonnegative, and any two vertices in $X_1 \cup X_2$ are roundtrip distance 2α , and any $s \in S$ and $x \in X_1 \cup X_2$ are distance 2α . Suppose Φ has no solution, so that every pair has ℓ_∞ distance α . Then for vertices i_S, j_S , there exists a coordinate x such that $v_i[x] - v_j[x]$ is either $\geq \alpha$ or $\leq -\alpha$. Without loss of generality, we are in the case $v_i[x] - v_j[x] \geq \alpha$. Then the path $i_S \rightarrow x_{X_2} \rightarrow j_S \rightarrow x_{X_1} \rightarrow i_S$ is a roundtrip path of length

$$(\alpha - v_i[x]) + (\alpha + v_j[x]) + (\alpha + v_j[x]) + (\alpha - v_i[x]) = 4\alpha - 2(v_i[x] - v_j[x]) \leq 2\alpha.$$

So when Φ has no solution, the roundtrip diameter is at most 2α .

On the other hand, suppose Φ has a solution i, j such that for all x , $|v_i[x] - v_j[x]| \leq 1$. Then, as every edge has weight at least $(0.5 - \varepsilon)\alpha$,

$$\begin{aligned} d(i_S, j_S) &\geq \min \left(\min_{x \in [d]} (d(i_S, x_{X_1}) + d(x_{X_1}, j_S), d(i_S, x_{X_2}) + d(x_{X_2}, j_S)), 4(0.5 - \varepsilon)\alpha \right) \\ &\geq \min \left(\min_{x \in [d]} (\alpha + v_i[x] + \alpha - v_j[x], \alpha + v_j[x] + \alpha - v_i[x]), 2\alpha - 4\varepsilon\alpha \right) \\ &\geq \min(2\alpha - 1, 2\alpha - 4\alpha\varepsilon) = 2\alpha - 1. \end{aligned}$$

Similarly, we have

$$d(j_S, i_S) \geq 2\alpha - 1,$$

so we have

$$d_{RT}(j_S, i_S) \geq 4\alpha - 2.$$

so in this case the RT-diameter is at least $4\alpha - 2$. A $2 - \alpha^{-1} - \varepsilon$ approximation for RT diameter can distinguish between RT diameter $4\alpha - 2$ and RT-diameter 2α . Thus, a $2 - \alpha - \varepsilon$ approximation for RT diameter solves α -approximate ℓ_∞ -CP. ◀

References

- 1 Amir Abboud, Karl Bringmann, Danny Hermelin, and Dvir Shabtay. Scheduling lower bounds via AND subset sum. *J. Comput. Syst. Sci.*, 127:29–40, 2022. doi:10.1016/j.jcss.2022.01.005.
- 2 Amir Abboud, Keren Censor-Hillel, Seri Khoury, and Ami Paz. Smaller cuts, higher lower bounds. *ACM Trans. Algorithms*, 17(4):30:1–30:40, 2021. doi:10.1145/3469834.
- 3 Amir Abboud, Fabrizio Grandoni, and Virginia Vassilevska Williams. Subcubic equivalences between graph centrality problems, APSP and diameter. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 1681–1697, 2015.
- 4 Amir Abboud, Virginia Vassilevska Williams, and Joshua R. Wang. Approximation and fixed parameter subquadratic algorithms for radius and diameter in sparse graphs. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 377–391, 2016.
- 5 Udit Agarwal and Vijaya Ramachandran. Fine-grained complexity for sparse graphs. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 239–252, 2018.
- 6 Josh Alman and Virginia Vassilevska Williams. A refined laser method and faster matrix multiplication. In Dániel Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 – 13, 2021*, pages 522–539. SIAM, 2021.
- 7 N. Alon, Z. Galil, and O. Margalit. On the exponent of the all pairs shortest path problem. *J. Comput. Syst. Sci.*, 54(2):255–262, 1997.
- 8 N. Alon, R. Yuster, and U. Zwick. Color-coding. *J. ACM*, 42(4):844–856, 1995.
- 9 Bertie Ancona, Keren Censor-Hillel, Mina Dalirrooyfard, Yuval Efron, and Virginia Vassilevska Williams. Distributed distance approximation. In Quentin Bramas, Rotem Oshman, and Paolo Romano, editors, *24th International Conference on Principles of Distributed Systems, OPODIS 2020, December 14-16, 2020, Strasbourg, France (Virtual Conference)*, volume 184 of *LIPIcs*, pages 30:1–30:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.OPODIS.2020.30.
- 10 Bertie Ancona, Monika Henzinger, Liam Roditty, Virginia Vassilevska Williams, and Nicole Wein. Algorithms and hardness for diameter in dynamic graphs. In *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019.
- 11 Alexandr Andoni, Dorian Croitoru, and Mihai Patrascu. Hardness of nearest neighbor under l-infinity. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 424–433, 2008. doi:10.1109/FOCS.2008.89.
- 12 Arturs Backurs, Liam Roditty, Gilad Segal, Virginia Vassilevska Williams, and Nicole Wein. Towards tight approximation bounds for graph diameter and eccentricities. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 267–280. ACM, 2018.

- 13 Édouard Bonnet. 4 vs 7 sparse undirected unweighted diameter is seth-hard at time $n^{4/3}$. In *Proc. ICALP*, pages 34:1–34:15, 2021.
- 14 Édouard Bonnet. Inapproximability of diameter in super-linear time: Beyond the 5/3 ratio. In *38th International Symposium on Theoretical Aspects of Computer Science, STACS 2021, March 16-19, 2021, Saarbrücken, Germany (Virtual Conference)*, volume 187 of *LIPIcs*, pages 17:1–17:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021.
- 15 Michele Borassi, Pierluigi Crescenzi, Michel Habib, Walter A Koster, Andrea Marino, and Frank W Takes. Fast diameter and radius bfs-based computation in (weakly connected) real-world graphs: With an application to the six degrees of separation games. *Theoretical Computer Science*, 586:59–80, 2015.
- 16 Karl Bringmann and Bhaskar Ray Chaudhury. Polyline simplification has cubic complexity. *J. Comput. Geom.*, 11(2):94–130, 2020. doi:10.20382/jocg.v11i2a5.
- 17 Massimo Cairo, Roberto Grossi, and Romeo Rizzi. New bounds for approximating extremal distances in undirected graphs. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 363–376, 2016.
- 18 Chris Calabro, Russell Impagliazzo, and Ramamohan Paturi. On the exact complexity of evaluating quantified k -cnf. In Venkatesh Raman and Saket Saurabh, editors, *Parameterized and Exact Computation – 5th International Symposium, IPEC 2010, Chennai, India, December 13-15, 2010. Proceedings*, volume 6478 of *Lecture Notes in Computer Science*, pages 50–59. Springer, 2010.
- 19 Shiri Chechik, Daniel H. Larkin, Liam Roditty, Grant Schoenebeck, Robert Endre Tarjan, and Virginia Vassilevska Williams. Better approximation algorithms for the graph diameter. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 1041–1052, 2014.
- 20 Shiri Chechik and Gur Lifshitz. Optimal girth approximation for dense directed graphs. In Daniel Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 – 13, 2021*, pages 290–300. SIAM, 2021.
- 21 Shiri Chechik, Yang P. Liu, Omer Rotem, and Aaron Sidford. Constant girth approximation for directed graphs in subquadratic time. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 1010–1023. ACM, 2020.
- 22 Lenore Cowen and Christopher G. Wagner. Compact roundtrip routing for digraphs. In Robert Endre Tarjan and Tandy J. Warnow, editors, *Proceedings of the Tenth Annual ACM-SIAM Symposium on Discrete Algorithms, 17-19 January 1999, Baltimore, Maryland, USA*, pages 885–886. ACM/SIAM, 1999.
- 23 Pierluigi Crescenzi, Roberto Grossi, Leonardo LANZI, and Andrea Marino. On computing the diameter of real-world directed (weighted) graphs. In Ralf Klasing, editor, *Experimental Algorithms: 11th International Symposium, SEA 2012, Bordeaux, France, June 7-9, 2012. Proceedings*, pages 99–110. Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- 24 Mina Dalirrooyfard and Jenny Kaufmann. Approximation algorithms for min-distance problems in dags. In Nikhil Bansal, Emanuela Merelli, and James Worrell, editors, *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12-16, 2021, Glasgow, Scotland (Virtual Conference)*, volume 198 of *LIPIcs*, pages 60:1–60:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021.
- 25 Mina Dalirrooyfard, Ray Li, and Virginia Vassilevska Williams. Hardness of approximate diameter: Now for undirected graphs. In *Proc. FOCS, FOCS’2021*, pages 1021–1032, 2021.
- 26 Mina Dalirrooyfard and Virginia Vassilevska Williams. Conditionally optimal approximation algorithms for the girth of a directed graph. In Artur Czumaj, Anuj Dawar, and Emanuela Merelli, editors, *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 168 of *LIPIcs*, pages 35:1–35:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020.

- 27 Mina Dalirrooyfard, Virginia Vassilevska Williams, Nikhil Vyas, Nicole Wein, Yinzhan Xu, and Yuancheng Yu. Approximation algorithms for min-distance problems. In *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.
- 28 Mina Dalirrooyfard and Nicole Wein. Tight conditional lower bounds for approximating diameter in directed graphs. In *Proc. STOC, STOC'2021*, pages 1697–1710, 2021.
- 29 Ofer Grossman, Seri Khoury, and Ami Paz. Improved hardness of approximation of diameter in the CONGEST model. In Hagit Attiya, editor, *34th International Symposium on Distributed Computing, DISC 2020, October 12-16, 2020, Virtual Conference*, volume 179 of *LIPICs*, pages 19:1–19:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPICs.DISC.2020.19.
- 30 R. Impagliazzo and R. Paturi. On the complexity of k-sat. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001.
- 31 Piotr Indyk. On approximate nearest neighbors under l-infinity norm. *Journal of Computer and System Sciences*, 63(4):627–638, 2001.
- 32 Haim Kaplan, Micha Sharir, and Elad Verbin. Colored intersection searching via sparse rectangular matrix multiplication. In *Proceedings of the twenty-second annual symposium on Computational geometry*, pages 52–60, 2006.
- 33 Karthik C. S. and Pasin Manurangsi. On closest pair in euclidean metric: Monochromatic is as hard as bichromatic. *Comb.*, 40(4):539–573, 2020. doi:10.1007/s00493-019-4113-1.
- 34 Ray Li. Improved seth-hardness of unweighted diameter. *CoRR*, abs/2008.05106v1, 2020. arXiv:2008.05106.
- 35 Ray Li. Settling seth vs. approximate sparse directed unweighted diameter (up to (nu)nseth). In *Proc. STOC, STOC'2021*, pages 1684–1696, 2021.
- 36 T. C. Lin, M. J. Wu, W. J. Chen, and B. Y. Wu. Computing the diameters of huge social networks. In *2016 International Computer Symposium (ICS)*, pages 6–11, 2016.
- 37 Andrea Lincoln, Virginia Vassilevska Williams, and R. Ryan Williams. Tight hardness for shortest cycles and paths in sparse graphs. In Artur Czumaj, editor, *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 1236–1252. SIAM, 2018.
- 38 Jakub Pachocki, Liam Roditty, Aaron Sidford, Roei Tov, and Virginia Vassilevska Williams. Approximating cycles in directed graphs: Fast algorithms for girth and roundtrip spanners. In Artur Czumaj, editor, *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 1374–1392. SIAM, 2018.
- 39 David Peleg, Liam Roditty, and Elad Tal. Distributed algorithms for network diameter and girth. In Artur Czumaj, Kurt Mehlhorn, Andrew M. Pitts, and Roger Wattenhofer, editors, *Automata, Languages, and Programming – 39th International Colloquium, ICALP 2012, Warwick, UK, July 9-13, 2012, Proceedings, Part II*, volume 7392 of *Lecture Notes in Computer Science*, pages 660–672. Springer, 2012. doi:10.1007/978-3-642-31585-5_58.
- 40 Maximilian Probst Gutenberg, Virginia Vassilevska Williams, and Nicole Wein. New algorithms and hardness for incremental single-source shortest paths in directed graphs. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 153–166, 2020.
- 41 Liam Roditty and Virginia Vassilevska Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. In *Proceedings of the 45th annual ACM symposium on Symposium on theory of computing*, STOC '13, pages 515–524, New York, NY, USA, 2013. ACM. doi:10.1145/2488608.2488673.
- 42 Aviad Rubinfeld. Hardness of approximate nearest neighbor search. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 1260–1268. ACM, 2018. doi:10.1145/3188745.3188916.

- 43 Aviad Rubinfeld and Virginia Vassilevska Williams. Seth vs approximation. *ACM SIGACT News*, 50(4):57–76, 2019.
- 44 R. Seidel. On the all-pairs-shortest-path problem in unweighted undirected graphs. *J. Comput. Syst. Sci.*, 51(3):400–403, 1995.
- 45 Frank W. Takes and Walter A. Kusters. Determining the diameter of small world networks. In *Proceedings of the 20th ACM International Conference on Information and Knowledge Management, CIKM '11*, pages 1191–1196, 2011.
- 46 Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theor. Comput. Sci.*, 348(2-3):357–365, 2005. doi:10.1016/j.tcs.2005.09.023.
- 47 Uri Zwick. All pairs shortest paths using bridging sets and rectangular matrix multiplication. *Journal of the ACM*, 49(3):289–317, 2002. Announced at FOCS'98. doi:10.1145/567112.567114.