# Settling SETH vs. Approximate Sparse Directed Unweighted Diameter (up to (NU)NSETH) 

Ray Li*

April 5, 2021


#### Abstract

We prove several tight results on the fine-grained complexity of approximating the diameter of a graph. First, we prove that, for any $\varepsilon>0$, assuming the Strong Exponential Time Hypothesis (SETH), there are no near-linear time $2-\varepsilon$-approximation algorithms for the Diameter of a sparse directed graph, even in unweighted graphs. This result shows that a simple near-linear time 2-approximation algorithm for Diameter is optimal under SETH, answering a question from a survey of Rubinstein and Vassilevska-Williams (SIGACT '19) for the case of directed graphs.

In the same survey, Rubinstein and Vassilevska-Williams also asked if it is possible to show that there are no $2-\varepsilon$ approximation algorithms for Diameter in a directed graph in $O\left(n^{1.499}\right)$ time. We show that, assuming a hypothesis called NSETH, one cannot use a deterministic SETH-based reduction to rule out the existence of such algorithms.

Extending the techniques in these two results, we characterize whether a $2-\varepsilon$ approximation algorithm running in time $O\left(n^{1+\delta}\right)$ for the Diameter of a sparse directed unweighted graph can be ruled out by a deterministic SETH-based reduction for every $\delta \in(0,1)$ and essentially every $\varepsilon \in(0,1)$, assuming NSETH. This settles the SETH-hardness of approximating the diameter of sparse directed unweighted graphs for deterministic reductions, up to NSETH. We make the same characterization for randomized SETH-based reductions, assuming another hypothesis called NUNSETH.


We prove additional hardness and non-reducibility results for undirected graphs.

[^0]
## 1 Introduction

The diameter $D$ of a graph $G=(V, E)$ is the maximum shortest path distance between two vertices. A basic algorithmic question, Diameter, is computing or approximating the diameter of a graph. The diameter is a useful measure of the complexity of a graph or network, so efficient algorithms for Diameter are desirable in practice [WS98, BE05, PRT12, $\mathrm{BCH}^{+} 15$, LWCW16]. In this work, we prove several tight results on the fine-grained complexity of approximating the diameter of a graph.

Diameter algorithms have been studied in dense graphs, sparse graphs, and special classes of graphs with additional structure [FP80, ACIM99, Epp00, CDHP01, CDV02, BBST07, RW13, CLR ${ }^{+}$14, CGS15, CGR16, Dam16, BHM19, Duc18, GKM ${ }^{+}$18, BN19, CDP19, DHV20]. In this work, we focus our attention on sparse graphs, i.e., graphs satisfying $m=n^{1+o(1)}$ where $n$ is the number of vertices and $m$ is the number of edges. We state our results for generic graphs, but our results are the strongest for sparse graphs, so it is helpful to imagine $m=n^{1+o(1)}$ throughout the paper.

In sparse graphs, the fastest exact Diameter algorithms in fact compute the distances between every pair of vertices, solving the All-Pairs-Shortest-Paths (APSP) problem, which takes $\tilde{O}(m n)$ time. Because finding exact algorithms is challenging, it is natural to try to find fast approximation algorithms. By running shortest path from a single vertex $v$ and returning $\tilde{D}$, the distance of $v$ to or from the furthest vertex, we obtain a simple 2-approximation in $\tilde{O}(m)$ time by the triangle inequality. Improving on the approximation ratio, a line of work [ACIM99, CLR $^{+} 14$, RW13] gave a $3 / 2$-approximation algorithm for Diameter in time $\tilde{O}\left(m^{3 / 2}\right)$. Cairo, Grossi and Rizzi [CGR16] generalized this algorithm to give an almost ${ }^{1}-\left(2-\frac{1}{2^{k}}\right)$-approximation in time $\tilde{O}\left(m^{1+\frac{1}{k+1}}\right)$ (for integers $k \geq 2$ ), but only in undirected graphs.

A natural question is whether these known algorithms are optimal. A line of work [RW13, $\mathrm{CLR}^{+} 14, \mathrm{BRS}^{+} 18$, Bon20] based on fine-grained complexity has given some partial answers to these questions (see Figure 1 for results in directed graphs and Figure 2 for results in undirected graphs). These works prove conditional hardness results assuming the Strong Exponential Time Hypothesis (SETH) [IPZ01]. ${ }^{2}$ In a survey on fine-grained complexity of approximation problems, Rubinstein and Vassilevska-Williams [RW19] asked three questions on the fine-grained complexity of approximating Diameter, and we address two of them in this work. First,

Question 1.1 (Open Question 2.2 of [RW19]). Is the simple near-linear time 2-approximation of Diameter optimal, or do there exist near-linear time algorithms giving better than a 2 -approximation?

In directed graphs, as stated above, $\left[\mathrm{CLR}^{+} 14\right]$ gave a $3 / 2$-approximation in $\tilde{O}\left(m^{3 / 2}\right)$, but for running time $O\left(m^{3 / 2-\delta}\right)$, the basic 2-approximation algorithm is the best known algorithm, as the algorithm of Cairo, Grossi, and Rizzi [CGR16] crucially uses that the graph is undirected. Hence, Rubinstein and Vassilevska-Williams also ask,

Question 1.2 (Open Question 2.3 of [RW19]). Is there an $O\left(m^{3 / 2-\delta}\right)$ time ( $2-\varepsilon$ )-approximation algorithm (for $\delta, \varepsilon>0$ ) for Diameter in directed graphs, or can we show that a $(2-\varepsilon)$-approximation for the diameter in sparse directed graphs needs $n^{3 / 2-o(1)}$ time?

Our contributions include answering Question 1.1 for the case of directed graphs (Corollary 1.4), conditionally answering Question 1.2 for SETH-hardness results (Corollary 1.6), and conditionally resolving the SETH-hardness of Diameter in directed unweighted graphs (Corollary 1.8). We

[^1]highlight our results in the context of existing work in Figure 1 (directed graphs) and Figure 2 (undirected graphs).

SETH-hardness results. Previously, the best fined grained lower bounds for near-linear time algorithms for Diameter in directed graphs ruled out $5 / 3-\varepsilon$-approximation algorithms in unweighted graphs $\left[\mathrm{BRS}^{+} 18\right]$ and $7 / 4-\varepsilon$-approximation algorithms in weighted graphs [Bon20] under SETH. We improve both of these results to $2-\varepsilon$, which is optimal.
Theorem 1.3. Let $k \geq 3$ be an integer and $\varepsilon>0$. Assuming SETH, a $2-\frac{1}{k}-\varepsilon$ approximation of the Diameter of a directed graph, weighted or unweighted, needs time $n^{1+1 /(k-1)-o(1)}$.

In particular, assuming SETH, the simple 2-approximation algorithm is optimal among nearlinear time Diameter algorithms for directed graphs, answering Question 1.1 for the case of directed graphs.

Corollary 1.4. Assuming SETH, for all $\varepsilon>0$, there exists a $\delta>0$ such that a $2-\varepsilon$ approximation of Diameter of a directed graph, weighted or unweighted, needs $n^{1+\delta-o(1)}$ time.

Our result improves over the result in $\left[\mathrm{BRS}^{+} 18\right]$ by improving the approximation ratio, and improves over [Bon20] not only in the approximation ratio, but also in that our result rules out algorithms even in unweighted graphs.

Not only does Corollary 1.4 achieve the best possible approximation ratio of $2-\varepsilon$ for near-linear time algorithms, it also obtains a good dependence of the runtime $n^{1+\delta-o(1)}$ on the approximation ratio $2-\varepsilon$ for approximation ratios less than 2. Later (Theorem 1.7), we in fact show that this tradeoff is conditionally optimal for directed unweighted graphs.

Question 1.1 remains open for undirected graphs. That is, it is open whether the simple 2approximation algorithm of the Diameter in undirected graphs is optimal for near-linear time algorithms. For undirected graphs, the best fine-grained lower bounds for near-linear time algorithms rule out a $5 / 3-\varepsilon$-approximation algorithm $\left[\mathrm{BRS}^{+} 18\right]$. However, this lower bound only holds for weighted graphs, and the best lower bound for unweighted graphs only rules out $8 / 5-\varepsilon$ approximation algorithms $\left[\mathrm{BRS}^{+} 18\right]$. For undirected unweighted graphs we improve this lower bound from $8 / 5-\varepsilon$ to $5 / 3-\varepsilon$, and we make the improvement with (we believe) a simpler reduction (see Theorem 7.1).

Concurrent work by Dalirrooyfard and Wein. In concurrent and independent work, Dalirrooyfard and Wein [DW20] obtained Theorem 1.3 with a different proof.

Non-reducibility results. For intermediate runtimes $O\left(n^{1+\delta}\right)$ for $\delta \in(0,1 / 2)$, for both directed and undirected graphs, and for both weighted and unweighted graphs, there are gaps between the best approximation ratios achieved by algorithms and the best approximation ratios ruled out by hardness results, even in light of Theorem 1.3. Therefore, it is natural to wonder whether stronger hardness of approximation results exist. In this section, we conditionally rule out better SETH-based hardness results with non-reducibility results in some settings. To prove such non-reducibility results, we utilize a framework of Carmosi, Gao, Impagliazzo, Mikhailin, Paturi, Schneider $\left[\mathrm{CGI}^{+} 16\right]$, which originally was used to conditionally prove that there is no deterministic reduction from SETH to 3-SUM and All-Pairs-Shortest-Path. Assuming the Nondeterministic Strong Exponential Time Hypothesis (NSETH) ${ }^{3}$ [CGI $\left.{ }^{+} 16\right]$, (1) we answer Question 1.2 for SETHbased hardness results and (2) we characterize which combinations of approximation ratios and

[^2]Diameter in sparse directed unweighted graphs


Diameter in sparse directed weighted graphs


Figure 1: Prior and new upper bounds, lower bounds, and non-reducibility results for Diameter in sparse directed graphs. Blue regions are feasible, red are infeasible assuming SETH. Different shades of blue or red denote different results. In cyan dotted regions, no deterministic SETH reductions can rule out the existence of algorithms, assuming NSETH, and no randomized SETH reductions can rule out the existence of algorithms assuming NUNSETH.

Diameter in sparse undirected unweighted graphs


Diameter in sparse undirected weighted graphs


Figure 2: Prior and new upper bounds, lower bounds, and non-reducibility results for Diameter in sparse undirected graphs. Blue regions are feasible, red are infeasible assuming SETH. Different shades of blue or red denote different results. In cyan dotted regions, no deterministic SETH reductions can rule out the existence of algorithms, assuming NSETH, and no randomized SETH reductions can rule out the existence of algorithms assuming NUNSETH.
runtimes can be ruled out by deterministic SETH-based hardness results for the diameter of directed unweighted graphs.

While there is some evidence towards the veracity of NSETH, namely that a disproof would yield nontrivial circuit lower bounds [CGI $\left.{ }^{+} 16\right]$, there is also some evidence against NSETH [Wil16], namely that a Merlin-Arthur SETH and Arthur-Merlin SETH are false in a strong way. However, NSETH itself seems challenging to disprove [CGI ${ }^{+} 16$, Wil18], so we view our NSETH-based nonreducibility results as showing a natural barrier to proving better SETH hardness results.

These non-reducibility results are proved without constructing algorithms. In Question 1.2, if the requested algorithm existed, then the question would be settled and there would be no hardness result. However, despite not finding such an algorithm, we are still able to rule out SETH-hardness results. Instead of finding an algorithm, we prove our non-reducibility results by constructing nondeterministic algorithms. This technique of using non-deterministic algorithms to obtain NSETHbased non-reducibility results was also done for the Gomory-Hu tree problem [AKT20]. The nondeterministic choices of our algorithm can be deterministically found in polynomial time, so our non-deterministic algorithms also provide potentially interesting certifying algorithms for Diameter. [ABMR11, MMNS11, Kun18].

Given a problem $\Pi$ and a time complexity $T$, the pair $(\Pi, T)$ is $S E T H-h a r d[C G I+16]$ if there is a deterministic Turing reduction from CNFSAT to $\Pi$ such that, if SETH holds, then $\Pi$ needs $T^{1-o(1)}$ time (see Section 3 for a formal definition). In this language, Theorem 1.3 states that a $2-\frac{1}{k}-\varepsilon$ approximation of Diameter with time complexity $n^{1+1 /(k-1)}$ is SETH-hard. Our results prove that certain approximations of Diameter are not SETH hard under NSETH. Under this definition of SETH-hard, proving that $(\Pi, T)$ is not SETH-hard does not rule out randomized reductions, when the Turing reduction in the SETH-hard definition is a probabilistic oracle machine (again see Section 3 for a formal definition). However, if we assume another hypothesis called NonUniform NSETH (NUNSETH) [CGI $\left.{ }^{+} 16\right]^{4}$, we obtain non-reducibility results even for randomized reductions.

In Question 1.2, Rubinstein and Vassilevska-Williams ask if we can show that a $(2-\varepsilon)$ approximation for the Diameter of sparse directed graph needs $n^{1.5-o(1)}$ time. We shows that the answer is no, if NSETH is true and we restrict ourselves to deterministic SETH reductions, or if NUNSETH is true and we restricted ourselves to randomized SETH reductions.

Theorem 1.5. Under NSETH (NUNSETH), all $\delta>0$ and $\varepsilon>0$, a $5 / 3+\varepsilon$ approximation of the Diameter on directed graphs, weighted or unweighted, with time complexity $m^{(1+\delta) 19 / 13}$ is not SETH-hard for deterministic (randomized) reductions.

As $5 / 3<2$ and $19 / 13<3 / 2$, taking $\delta$ and $\varepsilon$ to be sufficiently small constants in Theorem 1.5 immediately gives following corollary, which states our "no" answer to Question 1.2 for SETHhardness results.

Corollary 1.6. Under NSETH (NUNSETH), there exist $\delta>0$ and $\varepsilon>0$ such that $a 2-\varepsilon$ approximation of the Diameter on directed graphs, weighted or unweighted, with time complexity $m^{3 / 2-\delta}$ is not SETH-hard for deterministic (randomized) reductions.

Towards resolving Question 1.2 fully, Corollary 1.6 can be viewed in two possible ways: On one hand, it can be seen as evidence for an algorithm, that there exists a $\tilde{O}\left(m^{3 / 2-\delta}\right)$-time $(2-$ $\varepsilon)$-approximation of Diameter in directed graphs, answering Question 1.2 in full (see [AKT20, AKT20b] for an example of a progression from non-reducibility/nondeterministic algorithms to an

[^3]algorithm). On the other hand, if one believes such an algorithm does not exist, Corollary 1.6 can be seen as evidence that, unless NSETH is disproved, hypotheses other than SETH are necessary to conditionally rule out such algorithms. To our knowledge, all existing fine-grained hardness results for Diameter are based on deterministic reductions from SETH, so Corollary 1.6 indeed shows that, under NSETH, new ideas would be needed to prove such hardness results.

Using a variant of the technique for Theorem 1.5, we prove that the tradeoff between the runtime and approximation ratio in Theorem 1.3 is optimal among deterministic (randomized) SETH-reductions for directed unweighted graphs unless NSETH (NUNSETH) is disproved.

Theorem 1.7. Let $k \geq 2$ be a positive integer and let $\delta$ and $\varepsilon$ be positive reals. Assuming NSETH (NUNSETH), for any $\delta>0$, a $2-\frac{1}{k}+\varepsilon$ approximation of Diameter on directed unweighted graphs with time complexity $m^{1+1 / k+\delta}$ is not SETH-hard for deterministic (randomized) reductions under NSETH.

Theorem 1.7 is quantitatively stronger than Theorem 1.5 as Theorem 1.7 applies to a larger parameter setting of runtimes and approximation ratios, but Theorem 1.7 is weaker in that only rules out SETH-hardness results for unweighted graphs. Theorem 1.7 in fact also holds for a large class of weighted graphs (see Theorem 6.1 for the more general statement), but not arbitrary weighted graphs as in Theorem 1.5. Combining Theorem 1.3 with Theorem 1.7 yields the following complete characterization of the SETH-hardness of Diameter in sparse directed unweighted graphs.

Corollary 1.8. Assuming NSETH (NUNSETH), for any $\delta \in(0,1)$ and for essentially all $\varepsilon \in$ $(0,1),{ }^{5}$ a $2-\varepsilon$ approximation of Diameter of a sparse directed unweighted graph with time complexity $n^{1+\delta}$ is SETH-hard for deterministic (randomized) reductions if and only if $\lfloor 1 / \varepsilon\rfloor \leq\lfloor 1 / \delta\rfloor$.

Corollary 1.8 settles, for (essentially) every $\delta, \varepsilon \in(0,1)$, whether a $2-\varepsilon$ approximation of Diameter with time complexity $n^{1+\delta}$ is SETH-hard for deterministic reductions, under NSETH, and SETH-hard for randomized reductions, under NUNSETH. We can view Theorem 1.7 as establishing a similar dichotomy to Theorem 1.5: either there exist certain algorithms for directed unweighted diameter matching our SETH-hardness results (Theorem 1.3) or, unless NSETH is disproved, hypotheses other than SETH are needed to improve our hardness results for directed unweighted graphs.

We note that Theorem 1.5 (but not Theorem 1.7) depends on a construction of hopsets in [CFR20], for which currently only an extended abstract is available (see Section 3 for the construction that we use and Appendix A for an explanation of how the construction is implicit in [CFR20]). Faster and better constructions of hopsets in directed graphs would extend Theorem 1.5 to a larger setting of $\delta$ and $\varepsilon$, possibly matching the parameter setting of Theorem 1.7 (see Theorem 5.2 for the quantitative improvement implied by better hopset constructions).

The non-reducibility results of Theorem 1.5 and Theorem 1.7 automatically apply to undirected graphs. However, the guarantees implied for undirected unweighted graphs are better than the guarantees implied for undirected weighted graphs. Using a construction of hopsets for undirected graphs, we prove non-reducibility results for in undirected weighted graphs matching those implied by Theorem 1.7 for undirected unweighted graphs (see Theorem 5.6 for the result). Among other things, we show that, assuming NSETH, the lower bounds of $\left[\mathrm{BRS}^{+} 18\right]$ and Theorem 7.1, which show a $5 / 3-\varepsilon$ approximation needs $n^{3 / 2-o(1)}$ time in undirected graphs, are tight in the approximation ratio.

[^4]
### 1.1 Other related work

See [Wil18] for a survey of fine-grained complexity and [RW19] for a survey of fine-grained complexity for approximation problems.

Diameter algorithms in dense graphs have also been studied. In weighted dense graphs, like in sparse graphs, the fastest Diameter algorithms in fact compute the distances between every pair of vertices, solving the All-Pairs-Shortest-Paths (APSP) problem. The fastest APSP algorithms take time $\tilde{O}\left(\min \left(m n, n^{3} / \exp (\Omega(\sqrt{n}))\right)\right)$ [Wil14, Pet04, PR05]. For unweighted dense graphs, Diameter can be solved [CGS15] in time $\tilde{O}\left(\min \left(m n, n^{\omega}\right)\right)$, where $\omega<2.373$ is the constant for matrix multiplication [Wil12, Le14, Sto10]. In undirected, unweighted graphs, $\left[\mathrm{BRS}^{+} 18\right]$ gave a nearly- $3 / 2$ approximation in $\tilde{O}\left(n^{2}\right)$ expected time.

A number of generalizations and extensions of Diameter have been considered and studied, including eccentricities [CLR $\left.{ }^{+} 14, \mathrm{BRS}^{+} 18\right]$, ST-diameter [ $\mathrm{BRS}^{+} 18$, DWVW19], bichromatic diameter [DWVW19], roundtrip diameter [RW19], and min-diameter [AWW16, DWV $\left.{ }^{+} 19\right]$.

### 1.2 Outline of the paper

In Section 2, we outline the main ideas behind the proofs of our main results. In Section 3, we state some technical preliminaries for the proofs of our main results. In Section 4, we present the full proof of our SETH-hardness result, Theorem 1.3. We then move on to the proofs of our nonreducibility results. As the proofs of the non-reducibility results that use hopsets, Theorems 1.5 and 5.6, are simpler, we prove Theorems 1.5 and 5.6 first in Section 5. In Section 6, we prove Theorem 1.7. In Section 7, we prove Theorem 7.1.

We leave the proofs of lemmas stated in the preliminaries, Section 3, to the appendix, as they are technical and not the focus of this paper. In Appendix A, we justify our use of the hopset constructions, first for undirected graphs, and then for directed graphs, where we use the construction in [CFR20]. In Appendix B, we prove some results on fine-grained reductions.

## 2 Techniques

In this section, we sketch the ideas behind our SETH-hardness result (Theorem 1.3) and our nonreducibility results (Theorem 1.7 and Theorem 1.5).

### 2.1 SETH-hardness for $k=5$

First, we highlight the ideas behind our improved lower bound, Theorem 1.3, which shows that a $2-\frac{1}{k}-\varepsilon$ approximation of diameter in directed graphs needs $n^{1+1 /(k-1)-o(1)}$ time. To do so, we sketch the ideas of the proof when $k=5$.

Our construction (for $k=5$ ) is based on a clever lower bound construction for a problem called ST-Diameter, given by Backurs, Roditty, Segal, Vassilevska-Williams, and Wein [BRS ${ }^{+}$18]. This lower bound construction reduces from the 5-Orthogonal-Vectors (5-OV) problem, and use this construction for ST-Diameter to prove that there are no near-linear time (in fact no $O\left(n^{3 / 2-\delta}\right)$ time) $5 / 3-\varepsilon$ approximations for diameter in directed graphs. An important idea in our construction is to instead reduce from a variant of 5 -OV called Single-Set 5 -Orthogonal-Vectors.

5 -OV takes as input 5 sets each of $\tilde{n}$ vectors $A, B, C, D, E \subset\{0,1\}^{r}$ of dimension $r=\Theta(\log \tilde{n})$, and outputs whether there exists $k$ vectors $a \in A, b \in B, c \in C, d \in D, e \in E$ that are orthogonal, i.e., $a[x] \cdots e[x]=0$ for all coordinates $x \in[r]$. Single Set 5 -OV is the same problem, but additionally assumes $A=B=C=D=E$. Assuming SETH, both 5-OV and Single-Set 5-OV need $\tilde{n}^{5-o(1)}$ time [Wil05].


Figure 3: $G_{S T}$ : the ST-Diameter construction from $\left[\mathrm{BRS}^{+} 18\right]$ when $k=5$

From a Single-Set 5 -OV instance $\Phi$ of size $\tilde{n}$, we construct a graph $G$ of size $n \stackrel{\text { def }}{=} \tilde{O}\left(\tilde{n}^{4}\right)$ time such that, if $\Phi$ has a solution, the diameter of $G$ is at least 9 , and otherwise the diameter of $G$ is at most 5 . This shows that distinguishing between graphs of diameter 9 and diameter 5 needs $\tilde{n}^{5-o(1)}=n^{5 / 4-o(1)}$ time assuming SETH, so a $9 / 5-\varepsilon$-approximation of diameter also needs $n^{5 / 4-o(1)}$ time assuming SETH.

To construct the graph $G$, we start with the following construction of a graph $G_{S T}$, which is exactly the construction in $\left[\mathrm{BRS}^{+} 18\right]$ for ST-Diameter. Let the Single-Set 5 -OV instance $\Phi$ have vector sets $A \subset\{0,1\}^{r}$. In $\left[\mathrm{BRS}^{+} 18\right]$, the vertex set of $G_{S T}$ has six vertex subsets $L_{0}, \ldots, L_{5}$, with carefully chosen edges between $L_{i}$ and $L_{i+1}$ for $i=0, \ldots, 4$. Sets $L_{0}$ and $L_{5}$ have one vertex for each element of $A^{4}$, respectively. Sets $L_{1}, \ldots, L_{4}$ have (at most) one vertex for each element of $A^{3} \times[r]^{4}$. The graph $G_{S T}$ has the property that, if the 5-OV instance $\Phi$ has a solution, there exists $u \in L_{0}$ and $v \in L_{5}$ such that $d(u, v) \geq 13$, and otherwise $d(u, v) \leq 5$ for all $u \in L_{0}$ and $v \in L_{5}$. The basic setup of this construction is illustrated in Figure 3.

This construction does not immediately yield a hardness construction for Diameter. Indeed, when the 5-OV instance $\Phi$ has no solution, even though the so-called ST-Diameter between $L_{0}$ and $L_{5}$ is 5 , the Diameter is larger. The distance from a vertex in $L_{0}$ to a vertex in $L_{1}, \ldots, L_{4}$ can be larger than 5 , and the distance between any two vertices in $L_{1}, \ldots, L_{4}$ can also be larger than 5 . In [ $\left.\mathrm{BRS}^{+} 18\right]$, they turn this construction into a hardness construction for Diameter by adding more edges and vertices, but only show a $5 / 3-\varepsilon$ hardness of approximation for near-linear time directed Diameter algorithms.

Our directed Diameter lower bound builds on this construction in a different way. The key step, illustrated in Figure 4 is to (1) add edges from each of $L_{1}, L_{2}, L_{3}, L_{4}$ to $L_{1}$ and (2) add edges from $L_{4}$ to each of $L_{1}, L_{2}, L_{3}, L_{4}$, giving a graph $G^{\prime}$. These edges crucially use the fact that we are reducing from Single-Set orthogonal vectors. These edges have the property that, for any vertex $(a, b, c, d) \in L_{1}$ and any coordinate-tuple $\vec{i} \in[r]^{4}$, the out-neighborhood of $(a, b, c, d)$ is a subset of the out neighborhood of $(a, b, c, \vec{i}) \in L_{1} \cup \cdots \cup L_{4}$. Thus, when $\Phi$ has no solution, any vertex that vertex $(a, b, c, d)$ can reach in 5 steps can also be reached in 5 steps by any vertex $(a, b, c, \vec{i}) \in L_{1} \cup \cdots \cup L_{4}$. Since every vertex of $L_{0}$ is distance 5 from every vertex of $L_{5}$ in graph $G_{S T}$, every vertex in $L_{0} \cup \cdots \cup L_{4}$ is distance 5 from every vertex in $L_{5}$ in the new graph $G^{\prime}$. By a similar argument, any vertex that can reach $(b, c, d, e) \in L_{5}$ in 5 steps can also reach any $(c, d, e, \vec{i}) \in L_{1} \cup \cdots \cup L_{4}$ in 5 steps, so every vertex in $L_{0} \cup \cdots \cup L_{4}$ is distance 5 from every vertex in $L_{1} \cup \cdots \cup L_{5}$.

It remains to ensure that, when $\Phi$ has no solution, any two vertices in $L_{0}$ are at distance 5, and any two vertices in $L_{5}$ are at distance 5 . Because vertices in $L_{0}$ and $L_{5}$ are each identified by elements of $A^{4}$, there is a natural bijection between vertices in $L_{0}$ and vertices in $L_{5}$. In our final construction, we (1) direct the edges of $G^{\prime}$ between $L_{i}$ and $L_{i+1}$ to point towards the vertex in $L_{i+1}$, and (2) contract the pairs of vertices in $L_{0}$ and $L_{5}$ that correspond to the same 4 -tuple $A^{4}$, giving our final Diameter instance $G$. The contraction step also uses that we are reducing from Single-Set orthogonal vectors and not ordinary orthogonal vectors. We already showed that every


Figure 4: $G^{\prime}$ : almost our directed Diameter construction for $k=5$
vertex in $L_{0} \cup \cdots \cup L_{4}$ is distance 5 from every vertex in $L_{1} \cup \cdots \cup L_{5}$, but now that $L_{0}=L_{5}$ (and directing the edges did not disrupt any paths), we have that the Diameter of this final construction $G$ is at most 5 when $\Phi$ has no solution.

With some care, we can also show that the Diameter of $G$ is at least 9 when $\Phi$ has a solution, giving our desired reduction. In particular, we show that, if $a, b, c, d, e$ are orthogonal vectors in $\Phi$, the distance from vertex $(a, b, c, d) \in L_{0}$ to vertex $(b, c, d, e) \in L_{0}$ is at least 9 .

### 2.2 Non-reducibility for undirected unweighted 7/4-Diameter

In this section, we highlight the key ideas of Theorem 1.5 and Theorem 1.7, which show that under (NU)NSETH, a $\alpha$ approximation of the Diameter with time complexity $n^{1+\delta}$ is not SETH-hard, for certain values of $\alpha$ and $\delta$. Define the $D^{\prime} / D$-Diameter problem as the (promise) problem whose input is a graph, and such that an algorithm must always accept when the graph has diameter at least $D^{\prime}$, always reject when the graph has diameter at most $D$, and can accept or reject otherwise. We show (Lemma 3.5) that, to prove for some $\alpha \geq 1$ that an $\alpha$-approximation of Diameter with some time complexity $T$ is not SETH-hard, it suffices to prove that $\alpha D / D$-Diameter with some time complexity $T$ is not SETH-hard for all $D$. To demonstrate how we might prove that some $\alpha D / D$-Diameter is not SETH-hard, we show how to rule out SETH-hardness for distinguishing between undirected unweighted graphs of diameter 7 and 4.

Proposition 2.1. Under NSETH (NUNSETH), for any $\delta>0$, 7/4-Diameter on undirected unweighted graphs with time complexity $m^{4 / 3+\delta}$ is not SETH-hard for deterministic (randomized) reductions.

Following the framework of [CGI ${ }^{+}$16], to prove Proposition 2.1, it (almost, modulo a small technicality described after Lemma 3.1) suffices to prove the following proposition, that 7/4-Diameter on directed unweighted graphs is in ( $\mathrm{N} \cap \operatorname{coN}$ )TIME $\left[m^{4 / 3+o(1)}\right]$.
Proposition 2.2. 7/4-Diameter on undirected unweighted graphs is in ( $\mathrm{N} \cap \operatorname{coN}$ )TIME $\left[m^{4 / 3+o(1)}\right]$.
Proof. By the definition of ( $\mathrm{N} \cap$ coN)TIME, it suffices to (1) verify that a graph of diameter at least 7 indeed has diameter greater than 4 and (2) verify that a graph of diameter at most 4 indeed has diameter less than 7 , each in time $m^{4 / 3+o(1)}$. In a graph of diameter 7 , we can easily verify the graph has diameter at least 7 , and thus greater than 4 , by nondeterministically selecting a vertex and checking with BFS in $O(m)$ time if the eccentricity is at least 7 .

The hard step is verifying that the diameter is less than 7 when the diameter is at most 4 . The key idea is to observe that in a graph of diameter at most 4, one of the following exists.

1. A set $X$ of $\tilde{O}\left(m^{2 / 3}\right)$ vertices such that every vertex is distance 1 from some vertex in $X$.
2. A set $Y$ of $\tilde{O}\left(m^{1 / 3}\right)$ vertices such that every vertex is distance 2 from some vertex in $Y$.

To prove this, if any vertex has $\left|N_{2}(v)\right| \leq m^{1 / 3}$ (here, $N_{r}(v)$ is the set of all vertices at distance at most $r$ from $v$ ), we can take $Y=N_{2}(v)$. Since each vertex is distance at most 4 from $v$, each vertex is distance at most 2 from some vertex in $Y$, as desired. Thus, if $Y$ does not exist, we must have $\left|N_{2}(v)\right| \geq m^{1 / 3}$ for all vertices $v$. In particular, at least $m^{1 / 3}$ edges are incident to $N_{1}(v)$. If we take $X$ to be the vertices incident to $\tilde{O}\left(m^{2 / 3}\right)$ uniformly random edges, we end up with a vertex in each $N_{1}(v)$ with high probability, as desired.

From these two possibilities, we can certify that the diameter is at most 6 by first nondeterministically picking a set $X$ of $\tilde{O}\left(m^{2 / 3}\right)$ vertices and also nondeterministically picking a length 4 path between any two vertices in $X$, and then nondeterministically picking a set $Y$ of $m^{1 / 3}$ vertices. We then run multi-source shortest path from $X$ and single-source shortest path from each element of $Y$. We accept if each vertex has distance 1 to some vertex of $X$ and if all length- 4 paths between pairs of elements in $X$ are valid. We also accept if $Y$ has distance 2 to the rest of the graph and every element of $Y$ has eccentricity 4 . We reject otherwise. This all takes time $\tilde{O}\left(m^{4 / 3}\right)$. Crucially, checking the distances in $X$ only takes time $O\left(|X|^{2}\right)$ rather than $O(|X| \cdot m)$, because we guess the paths between vertices in $X$. If the graph is diameter 4, by above, there either exists $X$ such that we accept or there exists $Y$ such that we accept. On the other hand, one can check that, if we accept, the diameter of the graph is at most 6 , as desired.

In the full proof of Theorems 1.5, 1.7, and 5.6, we need to generalize Proposition 2.2 to (1) arbitrary diameters, (2) directed graphs, (3) weighted graphs, and/or (4) other runtimes and approximation ratios. The first step is the most challenging, and we describe the challenge and our solution here. If, for example, the graph is unweighted and $D$ and $D^{\prime}$ are small then generalizing Proposition 2.2 is straightforward. The problem comes when $D$ and $D^{\prime}$ are large, such as $m^{\Theta(1)}$, in which case the runtime of "guessing the path between pairs in $X$ " could take too long if there are too many vertices on the paths. One solution is to use a hopset, a useful object originally defined by Cohen [Coh00] in the context of parallel algorithms. Formally, a $(\beta, \varepsilon)$-hopset for a graph is a set of edges such that, when added to the graph, any two vertices have a $\beta$-edge path between them whose length is within a multiplicative $(1+\varepsilon)$ factor of the true shortest path. With a hopset, we only need to guess paths of at most $\beta$ vertices, saving on the runtime if $\beta$ is sufficiently small and the hopset is constructed sufficiently quickly, and losing up to an $\varepsilon D$ additive factor in the estimated distances, which we can afford. This works for Theorem 1.5 and Theorem 5.6.

However, in directed graphs, the state-of-the-art constructions of hopsets are either too slow and/or have too large of a hop-bound $\beta$ to be useful for Theorem 1.7. Instead, we leverage the following idea. Suppose we have a directed unweighted graph with a large diameter $D=m^{\Theta(1)}$. In the generalization of Proposition 2.2, we end up taking some set $Y$ formed by the $\tilde{D}$-neighborhood of some vertex $v$ for some $\tilde{D}$ so that $Y$ is distance at most $D-\tilde{D}$ from every vertex. Then, crucially, for every other vertex $u$, by considering the path from $u$ to $v$, there are at least $\varepsilon D$ vertices in $Y$ of distance at most $D-\tilde{D}+\varepsilon D$ from $u$ (this step uses that the graph is unweighted). Thus, we can replace $Y$ by a random subset $Y^{\prime}$ of size $\tilde{O}(|Y| /(\varepsilon D))$, saving a factor of $D$ in the runtime, which we need, and losing only an additive $\varepsilon D$ in the estimated distances, which we can afford.

## 3 Preliminaries

All logs are base $e$ unless otherwise specified. When $a_{1}, a_{2}, \ldots$, is a sequence, the subsequence $a_{i}, \ldots, a_{j}$ with $i>j$ is taken to be an empty sequence. For an integer $a$, let $[a] \stackrel{\text { def }}{=}\{1, \ldots, a\}$.

In graphs, $n$ always denotes the number of vertices, $m$ always denotes the number of edges, and $W_{\max }\left(W_{\min }\right)$ always denotes the maximum (minimum) edge-weight. In a graph $G$, let $d_{G}(u, v)$ be the length of the shortest path from $u$ to $v$. We omit the subscript $G$ when it is clear from the context. We assume all input graphs to Diameter problems are connected/strongly connected, so that $d(u, v)$ always exist. This can be verified in near-linear time and the diameter is not defined otherwise.

Graph notation. In an undirected graph, the eccentricity of a vertex $v$ is defined as $\epsilon(v)=$ $\max _{u \in V} d(u, v)$. In a directed graph, the (out-)in-eccentricity of a vertex $v$ is defined as $\epsilon^{i n}(v)=$ $\max _{u \in V} d(u, v)\left(\epsilon^{\text {out }}(v)=\max _{u \in V} d(v, u)\right)$. In an undirected graph, let $N_{\leq r}(v)=\{u: d(v, u) \leq r\}$ denote the set of vertices of distance at most $r$ from/to $v$. In a directed graph, let $N_{\leq r}^{i n}(v)=$ $\{u: d(u, v) \leq r\}$ denote the set of vertices of distance at most $r$ to $v$. In a directed graph, let $N_{\leq r}^{\text {out }}(v)=\{u, d(v, u) \leq r\}$ denote the set of vertices of distance at most $r$ from $v$. Let $N^{\text {in }}(v)$ denote the in-neighbors of a vertex $v$ and $N^{\text {out }}(v)$ denote out-neighbors of a vertex $v$. Let $E_{\leq r}(v)$ denote the set of edges incident to some vertex of $N_{\leq r}(v)$. Let $E_{\leq r}^{i n}(v)$ denote the set of edges $(a, b)$ (from $a$ to $b$ ) with $b \in N_{\leq r}^{i n}(v)$. Let $E_{\leq r}^{o u t}(v)$ denote the set of edges ( $a, b$ ) with $a \in N_{\leq r}^{o u t}(v)$. For a vertex $v$ in a graph and a subset $X$ of vertices, let $d(v, X) \stackrel{\text { def }}{=} \min _{x \in X} d(v, x)$. Similarly, let $d(X, v) \stackrel{\text { def }}{=} \min _{x \in X} d(x, v)$ (which may be different if the graph is directed.)

SETH. The Strong Exponential Time Hypothesis (SETH) [IPZ01] states that, for every $\varepsilon>0$, there exists a $k$ such that $k$-SAT on $n$ variables cannot be solved in $O\left(2^{(1-\varepsilon) n}\right)$ time (say, on a word-RAM with $O(\log n)$-bit words). For $k \geq 2$, the $k$-Orthogonal Vectors ( $k$-OV) problem asks, given sets $A_{1}, \ldots, A_{k} \subset\{0,1\}^{d}$ of $n$ vectors each, determine whether there exists vectors $a_{i} \in A_{i}$ such that $\prod_{i=1}^{k} a_{i}[x]=0$ for all $x \in[d]$. The Single-Set $k$-Orthogonal Vectors is the $k$-OV problem when $A_{1}=A_{2}=\cdots=A_{k}$. Williams [Wil05] showed that, assuming SETH, for a large enough constant $c, k$-OV and Single-Set $k$-OV need $n^{k-o(1)}$ time when $d=c \log n$.

Promise problem. A promise problem $\Pi$ is a pair of non-intersecting sets, denoted $\left(\Pi_{Y E S}, \Pi_{N O}\right)$, with $\Pi_{Y E S}, \Pi_{N O} \subset\{0,1\}^{*}$. The set $\Pi_{Y E S} \cup \Pi_{N O}$ is called the promise. A Turing reduction from a problem $\Pi_{1}$ to a promise problem $\Pi_{2}$ is an oracle Turing machine $\mathcal{M}^{\Pi_{2}}$ such that the output of $\mathcal{M}^{\Pi_{2}}$ is always correct whenever the outputs of $\Pi_{2}$ are "YES" when the input is in $\Pi_{Y E S}$ and are "NO" when the input is in $\Pi_{N O}$ and are arbitrarily "YES" or "NO" when the input is not in the promise.

Fine grained reductions. Let $\Pi_{1}, \Pi_{2}$ be a decision problem and $T_{1}, T_{2}$ be time bounds. We say that $\left(\Pi_{1}, T_{1}\right)$ fine-grained reduces to $\left(\Pi_{2}, T_{2}\right)$, denoted $\left(\Pi_{1}, T_{1}\right) \leq_{F G R}\left(\Pi_{2}, T_{2}\right)$ if

1. For all $\varepsilon>0$, there exists a $\delta>0$ and a deterministic Turing reduction $M^{\Pi_{2}}$ from $\Pi_{1}$ to $\Pi_{2}$, such that TIME $[\mathcal{M}] \leq T_{1}^{1-\delta}$ and
2. Let $\tilde{Q}(\mathcal{M}, x)$ denote the set of queries made by $\mathcal{M}$ to the oracle on a input $x$ of length $n$. Then the query lengths obey the following time bound:

$$
\sum_{q \in \tilde{Q}(\mathcal{M}, x)}\left(T_{2}(|q|)\right)^{1-\varepsilon} \leq\left(T_{1}(n)\right)^{1-\delta} .
$$

A problem $\Pi$ with time complexity $T$ is SETH-hard if there is a fine-grained reduction from CNFSAT with time complexity $2^{n}$ to $(\Pi, T)$. We note that one can make the same definition
for promise problems and function problems, with one exception for function problems. If $\left(\Pi_{2}, T_{2}\right)$ is a function problem, then we also bound the sizes of the answers given by the $\Pi_{2}$-oracle.

$$
\sum_{q \in \tilde{Q}(\mathcal{M}, x)}\left(\left|\Pi_{2}(q)\right|\right)^{1-\varepsilon} \leq\left(T_{1}(n)\right)^{1-\delta} .
$$

This definition holds even if a function has multiple acceptable outputs, as is the case for approximation problems where the output is acceptable if it is within an $\alpha$-factor of the optimal.

Randomized fine-grained reductions, denoted $\left(\Pi_{1}, T_{1}\right) \leq_{F G R, r}\left(\Pi_{2}, T_{2}\right)$ are defined exactly as deterministic fine-grained reductions, except that the Turing reduction from $\left(\Pi_{1}, T_{1}\right)$ to $\left(\Pi_{2}, T_{2}\right)$ is a probabilistic machine with some two-sided error bound

$$
\operatorname{Pr}\left[\mathcal{M}^{\Pi_{2}}(x) \in \Pi_{1}(x)\right] \geq \frac{2}{3},
$$

where $\Pi_{1}(x)$ is the set of acceptable outputs for $\Pi_{1}$ on input $x$. Now, a problem $\Pi$ with time complexity $T$ is SETH-hard with a randomized reduction if there is a randomized fine-grained reduction from CNFSAT with time complexity $2^{n}$ to ( $\left.\Pi, T\right)$.

NSETH and non-reducibility. The Nondeterministic Strong Exponential Time Hypothesis (NSETH) [CGI $\left.{ }^{+} 16\right]$ states that for every $\varepsilon>0$, there exists a $k$ such that $k$-TAUT is not in NTIME $\left[2^{n(1-\varepsilon)}\right]$, where $k$-TAUT is the language of all $k$-DNF which are tautologies. The following key result allowed [CGI ${ }^{+} 16$ ] to prove non-reducibility results.

Lemma 3.1 ([CGI ${ }^{+} 16$, Theorem 5.1]). If NSETH holds and $\Pi \in(\mathrm{N} \cap \operatorname{coN}) T I M E[T]$ for some decision or function problem $\Pi$, then $\left(\Pi, T^{1+\delta}\right)$ is not SETH-hard (with a deterministic reduction) for any $\delta>0$.

However, in [CGI ${ }^{+}$16], promise problems are not considered, and we were not able to reprove Lemma 3.1 for promise problems. However, we prove a slightly weaker result, which is still sufficient for our non-reducibility results. To state this result, we need the following definition.

Definition 3.2. Let $(\mathrm{N} \cap \operatorname{coN})^{+}$TIME $[T]$ be the complexity class containing the promise problems $\Pi=\left(\Pi_{Y E S}, \Pi_{N O}\right)$ such that there exists nondeterministic Turing machines $\mathcal{A}_{N}$ and $\mathcal{A}_{c o N}$ each running in time $T$ such that:

1. For all inputs $x \in \Pi_{Y E S}$, there exist nondeterministic choices such that $\mathcal{A}_{N}$ outputs "YES", and for all inputs $x \in \Pi_{N O}$, all nondeterministic choices cause $\mathcal{A}_{N}$ to output "NO".
2. For all inputs $x \in \Pi_{N O}$, there exist nondeterministic choices such that $\mathcal{A}_{c o N}$ outputs "NO", and for all inputs $x \in \Pi_{Y E S}$, all nondeterministic choices cause $\mathcal{A}_{c o N}$ to output "YES".
3. For all $x \notin \Pi_{Y E S} \cup \Pi_{N O}$, either some nondeterministic choices allow $\mathcal{A}_{N}$ to output "YES", or some nondeterministic choices allow $\mathcal{A}_{c o N}$ to output "NO".

For context, we note that satisfying the first two conditions is equivalent to being in ( $\mathrm{N} \cap$ coN)TIME $[T]$. We now state the slightly weaker result that we use, which is proved in Appendix B.

Lemma 3.3 (Lemma 3.1 for promise problems). If NSETH holds and $\Pi \in(\mathrm{N} \cap \operatorname{coN})^{+} \operatorname{TIME}[T]$ for some promise problem $\Pi$, then $\left(\Pi, T^{1+\delta}\right)$ is not SETH-hard for any $\delta>0$.

NUNSETH and non-reducibility for randomized reductions. The Non-Uniform Nondeterministic Strong Exponential Time Hypothesis (NUNSETH) states that, for every $\varepsilon>0$, there exists $k$ such that there are no nondeterministic circuit families of size $O\left(2^{n(1-\varepsilon)}\right)$ recognizing the language $k$-TAUT. The following is a result of [CGI $\left.{ }^{+} 16\right]$ that NUNSETH rules out randomized SETH-hardness results, but adapted to promise problems.

Lemma 3.4 (Lemma 3.9 of [CGI $\left.{ }^{+} 16\right]$ for promise problems). If NUNSETH holds and $\Pi \in(\mathrm{N} \cap$ coN $)^{+}$TIME $[T]$ for some promise problem $\Pi$, then $\left(\Pi, T^{1+\delta}\right)$ is not SETH-hard for randomized reductions for any $\delta>0$.

Diameter approximations. Given positive real numbers $D^{\prime}>D$, the $D^{\prime} / D$-Diameter problem is the promise problem ( $\Pi_{Y E S}, \Pi_{N O}$ ) whose input is a graph, and where $\Pi_{Y E S}$ is the set of graphs with Diameter at least $D^{\prime}$ and $\Pi_{N O}$ is the set of graphs with Diameter at most $D$. For a parameter $\alpha \geq 1$, the $\alpha$-approximate-Diameter problem is the function problem whose input is a graph, and whose output must be a value between $D$ and $\alpha D$, where $D$ is the diameter of the graph.

Our non-reducibility results (e.g. Theorem 1.5 and Theorem 1.7) describe when an $\alpha$-approximation of Diameter is not SETH-hard for some $\alpha$. The following lemma shows that it suffices to prove that the corresponding promise problem $\alpha / 1$-Diameter is not SETH-hard. We note that, even without this lemma, ruling out SETH-hardness for the promise problem is already interesting, as (to our knowledge) all known lower bounds for Diameter [RW13, $\mathrm{CLR}^{+} 14, \mathrm{BRS}^{+} 18$, Bon20] prove SETH-hardness of a corresponding promise problem. We also note that, because of Lemma 3.5, we can sidestep the inconvenience of having to define and work with nondeterministic algorithms for function problems. Lemma 3.5 is proved in Appendix B.

Lemma 3.5. Let $\alpha \geq 1$ and $\beta>0$ be constants, $\rho \in[1, \infty]$, and $T$ be some time complexity. If $\alpha / 1$-Diameter on graphs ${ }^{6}$ satisfying $W_{\max } / W_{\min } \leq \rho$ with time complexity $T$ is not SETH hard for deterministic (randomized) reductions, then $\alpha+\beta$-approximate-Diameter on graphs satisfying $W_{\max } / W_{\min } \leq \rho$ with time complexity $T$ is not SETH-hard for deterministic (randomized) reductions.

Hopsets. In a graph $G=(V, E)$, given a set of edges $E^{\prime}$ on vertices $V$, let $G+E^{\prime}$ denote the graph $G$ with edges $E^{\prime}$ added, i.e. $G+E^{\prime}=\left(V, E \cup E^{\prime}\right)$. In a graph $G$, a $(\beta, \varepsilon)$-hopset [Coh00] is a set of weighted edges $E^{\prime}$ (sometimes called shortcuts) such that, in the graph $G+E^{\prime}$, all shortestpath distances are the same as in $G$, and for any two vertices $v$ and $v^{\prime}$, there exists in $G+E^{\prime}$ a $\beta$-edge path between $v$ and $v^{\prime}$ of length at most $(1+\varepsilon) \cdot d\left(v, v^{\prime}\right)$. In a graph $G$ with diameter $D$, a $(\beta, \varepsilon)$-additive-hopset is a set of weighted edges $E^{\prime}$ such that, in the graph $G+E^{\prime}$, all shortest-path distances are the same as in $G$, and for any two vertices $v$ and $v^{\prime}$, there exists in $G+E^{\prime}$ a $\beta$-edge path between $v$ and $v^{\prime}$ of length at most $d\left(v, v^{\prime}\right)+\varepsilon D$. Note that any $(\beta, \varepsilon)$-hopset is by definition a $(\beta, \varepsilon)$-additive-hopset because all shortest path distances are at most $D$.

We use the following hopset constructions. The first one is a hopset construction for directed graphs. The construction is implicit in [CFR19, CFR20, CFR], and we explain how in Appendix A. We thank Arun Jambulapati for the reference [CFR20].

Lemma 3.6 (Implicit in [CFR19, CFR20, CFR]). Let $\alpha \in(0,1 / 2)$ and $\varepsilon \in(0,1)$. These exists a randomized algorithm with running time $\tilde{O}_{\varepsilon}\left(m^{1+4 \alpha}\right)$ such that, given a directed weighted graph $G$ on $n$ vertices and $m$ edges, computes a set of weighted edges $E^{\prime}$ such that, (1) with probability 1, when $E^{\prime}$ is added to $G$, all shortest path distances stay the same, and (2) with positive probability, $E^{\prime}$ forms a ( $\left.n^{1 / 2-\alpha+o(1)}, \varepsilon\right)$-additive-hopset of $G$.

[^5]The second hopset construction is for undirected graphs.
Lemma 3.7. Let $\delta \in(0,1)$ and $\varepsilon \in(0,1)$. These exists a randomized algorithm with running time $\tilde{O}_{\delta, \varepsilon}\left(m^{1+\delta}\right)$ such that, given an undirected weighted graph $G$ on $n$ vertices and $m$ edges, computes a set of weighted edges $E^{\prime}$ such that, (1) with probability 1, when $E^{\prime}$ is added to $G$, all shortest path distances stay the same, and (2) with positive probability, $E^{\prime}$ forms a $\left(O_{\delta, \varepsilon}(1), \varepsilon\right)$-additive-hopset of $G$.

Lemma 3.7 asks for a hopset construction but with the additional guarantee that, when the randomness is bad, the algorithm fails "safe" according to condition (1). We expect that Lemma 3.7 can be inferred from existing hopset constructions, but we include a proof in Appendix A for completeness.

## 4 Directed Diameter 2- $\varepsilon$ hardness

In this section, we prove Theorem 1.3.
Proof of Theorem 1.3. We show that, given a single-set $k$-OV instance $\Phi$ of size $\tilde{n}$ and dimension $d=\Theta(\log n)$, it is possible to construct a graph $G$ in $\tilde{O}\left(\tilde{n}^{k-1}\right)$ time such that, if $\Phi$ has a solution, then the diameter is at least $2 k-1$, and if $\Phi$ has no solution, then the diameter is at most $k$.

Let $\Phi$ be given by a set $A \subset\{0,1\}^{d}$ of $\tilde{n}$ vectors, where $d=\Theta(\log n)$.
When $a \in\{0,1\}^{d}$ is a vector and $x \in[d]$, we let $a[x]$ denote the $x$ th coordinate $a_{x}$ of $a$. When dealing with elements of $A$, we index coordinates using $a[x]$ rather than $a_{x}$ for clarity and to stay consistent with the notation in [ $\left.\mathrm{BRS}^{+} 18\right]$. We use $\bar{a}$ and $\bar{x}$ to refer to tuples of vectors and indices, respectively. We refer to $\bar{a}$ in $A^{k-1}$ or $A^{k-2}$ as a vector-tuple, and refer to $\bar{x} \in[d]^{k-1}$ as an index-tuple.

We make the following useful definitions.

1. (Property $\left.\left(i, L_{1}\right)\right)$ For $i=1, \ldots, k-1$, an index-tuple $\bar{x} \in[d]^{k-1}$ of $k-1$ indices, and a vector $a \in\{0,1\}^{d}$, we say the pair $(a, \bar{x})$ has property $\left(i, L_{1}\right)$ if $a\left[x_{1}\right]=\cdots=a\left[x_{k-i}\right]=1$.
2. (Property $\left.\left(i, L_{k}\right)\right)$ For $i=2, \ldots, k$, an index-tuple $\bar{x} \in[d]^{k-1}$ of $k-1$ indices, and a vector $a \in\{0,1\}^{d}$, we say the pair ( $a, \bar{x}$ ) has property $\left(i, L_{k-1}\right)$ if $a\left[x_{k+1-i}\right]=\cdots=a\left[x_{k-1}\right]=1$.

Construct a graph $G$ with vertex set $V=L_{0} \cup L_{1} \cup \cdots \cup L_{k-1}$.

1. Create one vertex in $L_{0}$ for each element of $A^{k-1}$.
2. For $2 \leq i \leq k-2$, create one vertex in $L_{i}$ for each element of $(\bar{a}, \bar{x}) \in A^{k-2} \times[d]^{k-1}$.
3. Create one vertex in $L_{1}$ for each element of $(\bar{a}, \bar{x}) \in A^{k-2} \times[d]^{k-1}$ such that $\left(a_{j}, \bar{x}\right)$ has $\operatorname{property}\left(j, L_{1}\right)$ for $1 \leq j \leq k-2$.
4. Create one vertex in $L_{k-1}$ for each element of $(\bar{a}, \bar{x}) \in A^{k-2} \times[d]^{k-1}$ such that ( $a_{j-2}, \bar{x}$ ) has property $\left(j, L_{k-1}\right)$ for $3 \leq j \leq k$.

We use $(\bar{a})_{L_{0}}$ or $(\bar{a}, \bar{x})_{L_{i}}$ to denote the corresponding vertex of $L_{i}$, using the subscript $L_{i}$ for disambiguation. We construct the edges of $G$ as follows.

1. $\left(L_{0} \rightarrow L_{1}\right)$ For all vectors $a_{1}, \ldots, a_{k-1} \in A$ and index-tuples $\bar{x} \in[d]^{k-1}$ such that $\alpha_{1}=$ $\left(a_{1}, \ldots, a_{k-2}, \bar{x}\right)_{L_{1}}$ exists as a vertex in $L_{1}$ and such that $a_{k-1}\left[x_{1}\right]=1$, add an edge from vertex $\left(a_{1}, \ldots, a_{k-1}\right)_{L_{0}} \in L_{0}$ to vertex $\alpha_{1}$.


Figure 5: Here, we illustrate $G$ when $k=5$. For clarity, we use $a, b, c, \ldots$ instead of $a_{1}, a_{2}, a_{3}, \ldots$, and we use $x, y, z, w$ instead of $x_{1}, x_{2}, x_{3}, x_{4}$, and we use $\bar{x} \stackrel{\text { def }}{=}(x, y, z, w)$. The blue cells indicate coordinates that need to equal 1 for the edge to exist. The table on the right indicates all coordinates that need to equal 1 for the path from $(a, b, c, d)$ to $(e, f, g, h)$ to exist.
2. $\left(L_{k-1} \rightarrow L_{0}\right)$ For all vectors $a_{2}, \ldots, a_{k} \in A$ and index-tuples $\bar{x} \in[d]^{k-1}$ such that $\alpha_{k-1}=$ $\left(a_{3}, \ldots, a_{k}, \bar{x}\right)_{L_{k-1}}$ exists as a vertex in $L_{k-1}$ and such that $a_{2}\left[x_{k-1}\right]=1$, add an edge from vertex $\alpha_{k-1} \in L_{k-1}$ to vertex $\left(a_{2}, \ldots, a_{k}\right)_{L_{0}} \in L_{0}$.
3. $\left(L_{i} \rightarrow L_{i+1}\right)$ For any $i=1, \ldots, k-2$, any index-tuple $\bar{x} \in[d]^{k-1}$, and any vector tuples $\bar{a}$ and $\bar{b}$ such that $a_{j}=b_{j}$ for all $j \neq k-1-i$, add an edge from vertex $\alpha_{i}=(\bar{a}, \bar{x})_{L_{i}}$ to vertex $\alpha_{i+1}=(\bar{b}, \bar{x})_{L_{i+1}}$ if both vertices exist.
4. $\left(L_{i} \rightarrow L_{1}\right)$ For every vector-tuple $\bar{a} \in A^{k-2}$ and every two index-tuples $\bar{x}, \bar{x}^{\prime} \in[d]^{k-1}$, and every $1 \leq i \leq k-1$, add a directed edge from $(\bar{a}, \bar{x})_{L_{i}}$ to $\left(\bar{a}, \bar{x}^{\prime}\right)_{L_{1}}$ if both vertices exist.
5. $\left(L_{k-1} \rightarrow L_{i}\right)$ For every vector-tuple $\bar{a} \in A^{k-2}$ and every two index-tuples $\bar{x}, \bar{x}^{\prime} \in[d]^{k-1}$, and every $1 \leq i \leq k-1$, add a directed edge from $(\bar{a}, \bar{x})_{L_{k-1}}$ to $\left(\bar{a}, \bar{x}^{\prime}\right)_{L_{i}}$ if both vertices exist.

We refer to the first three types of edges as vector-changing edges, as the vector-tuple changes, and the index tuple, if present does not. We refer to the latter two types of edges as index-changing edges, as the vector-tuple stays constant, while the index tuple changes.

Runtime. The number of vertices in each $L_{i}$ is at most $O\left(\tilde{n}^{k-1}+\tilde{n}^{k-2} d^{k-1}\right)=\tilde{O}\left(\tilde{n}^{k-1}\right)$. Each edge has at least one endpoint in $L_{1} \cup \cdots \cup L_{k-1}$, and each such vertex has at most $O\left(\tilde{n}+d^{k}\right)$ neighbors. $L_{1} \cup \cdots \cup L_{k-1}$ have at most $O\left(\tilde{n}^{k-2} d^{k-1}\right)$ vertices, so the total number of edges is thus at most $O\left(\tilde{n}^{k-1} d^{2 k-1}\right) \leq \tilde{O}\left(\tilde{n}^{k-1}\right)$. Thus, the runtime to produce $G$ from $\Phi$ is at most $\tilde{O}\left(\tilde{n}^{k-1}\right)$ as desired.
$k$-OV no solution. We first show that if there is no $k$-OV solution, then the diameter is at most $k$. For any tuple $\bar{a}=\left(a_{1}, \ldots, a_{l}\right)$ of at most $k$ vectors, by assumption, there exists some index $x$ such that $a_{1}[x]=a_{2}[x]=\cdots=a_{l}[x]=1$. Let ind $(\bar{a})$ denote one such index $x$.

The crucial claim of this section is the following.
Claim 4.1. For any two vertices in $L_{0}$, there is a length $k$ path from one to the other.

Proof. Let the vertices be $\alpha_{0} \stackrel{\text { def }}{=}\left(a_{1}, \ldots, a_{k-1}\right)_{L_{0}}$ and $\alpha_{k} \stackrel{\text { def }}{=}\left(b_{2}, \ldots, b_{k}\right)_{L_{0}}$. For $i=1, \ldots, k-1$, let $x_{i}=\operatorname{ind}\left(a_{1}, \ldots, a_{k-i}, b_{k-i+1}, \ldots, b_{k}\right)$ and $\bar{x}=\left(x_{1}, \ldots, x_{k-1}\right)$. By definition of $x_{1}, \ldots, x_{k-1}$, for $i=1, \ldots, k-1$, we have $a_{i}\left[x_{1}\right]=\cdots=a_{i}\left[x_{k-i}\right]=1$, so the pair ( $a_{i}, \bar{x}$ ) satisfies property $\left(i, L_{1}\right)$. By definition of $x_{1}, \ldots, x_{k-1}$, for $i=2, \ldots, k$, we have $b_{i}\left[x_{k+1-i}\right]=\cdots=b_{i}\left[x_{k-1}\right]=1$, so the pair $\left(b_{i}, \bar{x}\right)$ satisfies property $\left(i, L_{k-1}\right)$. Thus, $\alpha_{1} \stackrel{\text { def }}{=}\left(a_{1}, \ldots, a_{k-2}, \bar{x}\right)_{L_{1}}$ is a valid vertex in $L_{1}$ and $\alpha_{k-1} \stackrel{\text { def }}{=}\left(b_{3}, \ldots, b_{k}\right)_{L_{k-1}}$ is a valid vertex in $L_{k-1}$, and furthermore there are edges from $\alpha_{0}$ to $\alpha_{1}$ and from $\alpha_{k-1}$ to $\alpha_{k}$.

For $i=2, \ldots, k-2$, vertex $\alpha_{i} \stackrel{\text { def }}{=}\left(a_{1}, \ldots, a_{k-1-i}, b_{k-i+2}, \ldots, b_{k}, \bar{x}\right)_{L_{i}}$ exists as a vertex in $L_{i}$. For $i=1, \ldots, k-2, \alpha_{i}$ and $\alpha_{i+1}$ have the same index-tuple and have vector tuples that differ only in the $k-1-i$ 'th vector of the vector-tuple. Thus, there is an edge from $\alpha_{i}$ to $\alpha_{i+1}$ for all $i=1, \ldots, k-2$. Hence, $\alpha_{0}, \ldots, \alpha_{k}$ is a length $k$ path from $\alpha_{0}$ to $\alpha_{k}$, as desired.

Claim 4.2. For any vertex $v \in V$ in our graph, there exist vertices $u, w \in L_{0}$ such that $N^{\text {out }}(u) \subset$ $N^{\text {out }}(v)$ and $N^{\text {in }}(w) \subset N^{\text {in }}(v)$.

Proof. If $v \in L_{0}$, simply set $u=w=v$. Otherwise, $v \in L_{1} \cup \cdots \cup L_{k-1}$ has some vector tuple $\left(a_{1}, \ldots, a_{k-2}\right) \in A^{k-2}$ and index tuple $\bar{x} \in[d]^{k-1}$. Let $b \in A$ be an arbitrary vector and let $u=\left(a_{1}, \ldots, a_{k-2}, b\right)_{L_{0}}$ and $w=\left(b, a_{1}, \ldots, a_{k-2}\right)_{L_{0}}$. The only edges out of $L_{0}$ are to $L_{1}$, so all out-neighbors of $u$ are in $L_{1}$ and of the form $\left(a_{1}, \ldots, a_{k-2}, \bar{x}^{\prime}\right)_{L_{1}}$. All such vertices are also outneighbors of $v$ as desired. Similarly, the only edges into $L_{0}$ are from $L_{k-1}$, so all in-neighbors of $w$ are in $L_{k-1}$ and of the form $\left(a_{1}, \ldots, a_{k-2}, \bar{x}^{\prime}\right)_{L_{k-1}}$. All such vertices are also in-neighbors of $v$ as desired.

Let $v, v^{\prime} \in V$. By Claim 4.2, there exists $u, w^{\prime} \in L_{0}$ such that $N^{\text {out }}(u) \subset N^{\text {out }}(v)$ and $N^{\text {in }}\left(w^{\prime}\right) \subset$ $N^{i n}\left(v^{\prime}\right)$. By Claim 4.1, there exists a length $k$ path $u, v_{1}, v_{2}, \ldots, v_{k-1}, w^{\prime}$ from $u$ to $w^{\prime}$. Since $v_{1} \in N^{\text {out }}(u)$, we must have $v_{1} \in N^{\text {out }}(v)$, and since $v_{k-1} \in N^{\text {in }}\left(w^{\prime}\right)$, we must have $v_{k-1} \in N^{\text {in }}\left(v^{\prime}\right)$. Thus, $v, v_{1}, \ldots, v_{k-1}, v^{\prime}$ is a length $k$ path from $v$ to $v^{\prime}$. This holds for any vertices $v$ and $v^{\prime}$, so the diameter is at most $k$, as desired.
$k$-OV solution. We now show that, if there is a $k$-OV solution, then the diameter is at least $2 k-1$. Let $a_{1}, \ldots, a_{k}$ be the $k$ orthogonal vectors, i.e. $a_{1}[x] \cdots a_{k}[x]=0$ for all $x \in[d]$. We claim that the distance from $\alpha_{0}=\left(a_{1}, \ldots, a_{k-1}\right)_{L_{0}} \in L_{0}$ to $\alpha_{k}=\left(a_{2}, \ldots, a_{k}\right)_{L_{0}} \in L_{0}$ is at least $2 k-1$.

Suppose for contradiction there exists a length $\ell \leq 2 k-2$ path $\mathcal{P}$ from $\alpha_{0}$ to $\alpha_{k}$. Let $\beta_{0}=$ $\alpha_{0}, \beta_{1}, \ldots, \beta_{\ell}=\alpha_{k}$ denote the vertices of path $\mathcal{P}$. We note the following two observations about $G$ that follow from the edge definitions.
Fact 4.3. The only edges in graph $G$ that go from vertex subset $L_{i}$ to vertex subset $L_{j}$ for $0 \leq i<$ $j \leq k-1$ are edges from $L_{i}$ to $L_{i+1}$ from $i=0, \ldots, k-1$.
Fact 4.4. In graph $G$, any edge from $L_{i}$ to $L_{j}$ for $k-1 \geq i \geq j \geq 0$ satisfy either $i=k-1$ or $j=1$.

The only edges to $L_{0}$ are those from $L_{k-1}$, and any path from $L_{0}$ to $L_{k-1}$ has length at least $k-1$ by Fact 4.3. Thus, two vertices in path $\mathcal{P}$ in $L_{0}$ must be separated by distance at least $k$. As the path $\beta_{0}, \ldots, \beta_{\ell}$ starts and ends in $L_{0}$ and has length at most $2 k-2$, it follows that the path visits $L_{0}$ only at the first vertex $\beta_{0}$ and the last vertex $\beta_{\ell}$. Thus, $\beta_{1}, \ldots, \beta_{\ell-1}$ lie entirely in $L_{1}, \ldots, L_{k-1}$.

By Fact 4.3, we must have $\beta_{1} \in L_{1}$. Let $r$ be the largest integer such that vertex $\beta_{r}$ is in $L_{1}$. By Fact 4.3, there are at least $k-1$ more vertices on the path $\mathcal{P}$, so $r \leq \ell-(k-1) \leq k-1$. By Fact 4.4, as none of $\beta_{r+1}, \ldots, \beta_{r+k-2}$ are in $L_{1}$, we must have $\beta_{r+i} \in L_{i+1}$ for $i=1, \ldots, k-2$.

Let $s$ denote the largest integer such that the subpath $\beta_{0}, \ldots, \beta_{r}$ has a vertex in subset $L_{s}$, and let $t$ denote the smallest integer such that the subpath $\beta_{r+k-2}, \ldots, \beta_{\ell}$ has a vertex in subset $L_{t}$. By Fact 4.3, the subpath $\beta_{1}, \ldots, \beta_{r-1}$ must stay in $L_{1} \cup \cdots \cup L_{r-1}$. Hence, if $r \geq 2$, then $s \leq r-1$, and $s=r=1$ otherwise. Similarly, the path $\beta_{r+k-1}, \ldots, \beta_{\ell-1}$ has length $\ell-1-(r+k-1) \leq k-r-2$ and thus must stay in $L_{r+1} \cup \cdots \cup L_{k-1}$. Hence, if $r+k-1 \leq \ell-1$, then $t \geq r+1$, and $t=k-1$ otherwise. Thus, if $r \geq 2$ and $r+k-1 \leq \ell-1$, we must have $t-s \geq 2$.

We know $\beta_{1}=\left(a_{1}, \ldots, a_{k-2}, \bar{x}\right)$ for some $\bar{x}$. All vertices in $L_{1}, \ldots, L_{s}$ are of the form $\left(a_{1}^{\prime}, \ldots, a_{k-2}^{\prime}, \bar{x}^{\prime}\right)_{L_{i}}$, and any of the edges between them does not change the vectors $a_{1}^{\prime}, \ldots, a_{k-1-s}^{\prime}: L_{i} \rightarrow L_{1}$ edges do not change any of the vectors, and edges from $L_{i} \rightarrow L_{i+1}$ only changes $a_{k-1-i}^{\prime}$. Hence, we have that $\beta_{r}=\left(a_{1}, \ldots, a_{k-1-s}, a_{k-s}^{\prime}, \ldots, a_{k-2}^{\prime}, \bar{y}\right)_{L_{1}} \in L_{1}$ for some vectors $a_{k-s}^{\prime}, \ldots, a_{k-2}^{\prime}$ and index tuple $\bar{y}$. Because $\beta_{r}$ exist as a vertex in $L_{1}$, the pair $\left(a_{i}, \bar{y}\right)$ has property $\left(i, L_{1}\right)$ for $i \leq k-1-s$.

As the index tuple of $\beta_{r}$ is $\bar{y}$, the path $\beta_{r}, \ldots, \beta_{r+k-2}$ only uses $L_{i} \rightarrow L_{i+1}$ edges, and the index tuple does not change in $L_{i} \rightarrow L_{i+1}$ edges, the index tuple of $\beta_{r+k-2}$ is also $\bar{y}$.

We also have $\beta_{\ell-1}=\left(a_{3}, \ldots, a_{k}, \bar{x}^{\prime}\right)$ for some $\bar{x}^{\prime}$. All vertices in $L_{t}, \ldots, L_{k-1}$ are elements $\left(a_{3}^{\prime}, \ldots, a_{k}^{\prime}, \bar{x}^{\prime}\right)$ of $A^{k-2} \times[d]^{k-1}$, and any of the edges between them does not change the vectors $a_{k+2-t}^{\prime}, \ldots, a_{k}^{\prime}: L_{k-1} \rightarrow L_{i}$ edges do not change any of the vectors, and $L_{i} \rightarrow L_{i+1}$ edges only change $a_{k+1-i}^{\prime}$ (here we indexed vectors starting from $a_{3}^{\prime}$ ). Hence, for some $a_{3}^{\prime \prime}, \ldots, a_{k+1-t}^{\prime \prime}$, we have that $\beta_{r+k-2}$ is the vertex $\left(a_{3}^{\prime \prime}, \ldots, a_{k+1-t}^{\prime \prime}, a_{k+2-t}, \ldots, a_{k}, \bar{y}\right)_{L_{k-1}} \in L_{k-1}$. Because $\beta_{k-2+r}$ exists as a vertex in $L_{k-1}$, the pair $\left(a_{i}, \bar{y}\right)$ has property $\left(i, L_{k-1}\right)$ for $i \geq k+2-t$.

When $1 \leq i \leq k-1-s$, we saw that $\left(a_{i}, \bar{y}\right)$ has property $\left(i, L_{1}\right)$, so $a_{i}\left[y_{j}\right]=1$ for all $j \leq k-i$ and in particular for all $j \leq k-(k-1-s)=s+1$. Similarly, when $k+2-t \leq i \leq k$, we saw that $\left(a_{i}, \bar{y}\right)$ has property $\left(i, L_{k-1}\right)$, so $a_{i}\left[y_{j}\right]=1$ for $j \geq k+1-i$, and in particular for all $j \geq k+1-(k+2-t)=t-1$. If $t-s \geq 2$, then we have that $a_{1}[j]=a_{2}[j]=\cdots=a_{k}[j]=1$ for all $t-1 \leq j \leq s+1$, a contradiction of orthogonality of $a_{1}, \ldots, a_{k}$.

Thus, we must have $t-s \leq 1$. By an earlier argument, if $r>1$ and $r+k-2<\ell-1$, then $t-s \geq 2$. Hence, we must either have $r=1$ or $r+k-2=\ell-1$.

If $r=1$, then $s=1$ and $t=2$, and by above, $\left(a_{i}, \bar{y}\right)$ has property $\left(i, L_{1}\right)$ for $i \leq k-1-s$, and in particular for all $i \leq k-2$. Additionally, $\left(a_{i}, \bar{y}\right)$ has property $\left(i, L_{k-1}\right)$ for $i=k$. Thus, $a_{1}\left[y_{1}\right]=a_{2}\left[y_{1}\right]=\cdots=a_{k-2}\left[y_{1}\right]=1$ and $a_{k}\left[y_{1}\right]=1$, respectively. Furthermore, since $r=1$, we have $\beta_{0}=\left(a_{1}, \ldots, a_{k-1}\right)_{L_{0}}$ and $\beta_{1}=\left(a_{1}, \ldots, a_{k-2}, \bar{y}\right)_{L_{1}}$ are the first two vertices of the path $\mathcal{P}$, so the first edge of the path implies additionally that $a_{k-1}\left[y_{1}\right]=1$. We conclude $a_{1}\left[y_{1}\right]=\cdots=a_{k}\left[y_{1}\right]=1$, contradicting orthogonality.

If $r+k-2=\ell-1$, then $s \leq r-1 \leq \ell-k \leq k-2$ and $t=s+1 \leq k-1$. By above, $\left(a_{i}, \bar{y}\right)$ has property $\left(i, L_{1}\right)$ for $i \leq k-1-s$, and in particular $i=1$. Additionally, $\left(a_{i}, \bar{y}\right)$ has property $\left(i, L_{k-1}\right)$ for $i \geq k+2-t$, and in particular for all $3 \leq i \leq k$. We conclude that $a_{1}\left[y_{k-1}\right]=1$ and $a_{3}\left[y_{k-1}\right]=\cdots=a_{k}\left[y_{k-1}\right]=1$. Additionally, since $r=\ell-1$, we have that $\beta_{\ell-1}=\left(a_{3}, \ldots, a_{k}, \bar{y}\right)_{L_{k-1}}$ and $\beta_{\ell}=\left(a_{2}, \ldots, a_{k}\right)_{L_{0}}$ are the last two vertices of the path $\mathcal{P}$, so the edge between them implies additionally that $a_{2}\left[y_{k-1}\right]=1$. We thus have $a_{1}\left[y_{k-1}\right]=\cdots=a_{k}\left[y_{k-1}\right]=1$, contradicting orthogonality.

We have thus shown that in the three cases that the path $\beta_{0}, \ldots, \beta_{\ell}$ could satisfy, (1) $r>1$ and $r+k-2<\ell-1$, (2) $r=1$, and (3) $r+k-2=\ell-1$, there is a contradiction. This covers all cases, we have found a contradiction, so there cannot exist a length $\ell \leq 2 k-2$ path from $\alpha_{0}$ to $\alpha_{k}$, as desired.

## 5 Non-reducibility via hopsets

### 5.1 Directed weighted graphs

By Lemma 3.5, to prove Theorem 1.5 regarding the non-reducibility of approximate Diameter, it suffices to prove the following theorem on the non-reducibility of the corresponding promise problem.

Theorem 5.1. Let $D$ be a positive real number, $\varepsilon>0$ and $D^{\prime}=(5 / 3+\varepsilon) D$. For all $\delta>0$, under NSETH (NUNSETH), D'D-Diameter on directed weighted graphs with time complexity $m^{19 / 13(1+\delta)}$ is not SETH-hard for deterministic (randomized) reductions.

Theorem 5.1 in fact follows from the following more general theorem, which states that nonreducibility results follow from hopset constructions. Using Theorem 5.2 with the state-of-the-art hopset constructions [CFR20], we obtain Theorem 5.1. Furthermore, faster and better hopset constructions would give even stronger non-reducibility results.

Theorem 5.2. Let $\varepsilon>0$ be fixed and $\beta$ be a real number, possibly depending on $n$ and $m$. Suppose there exists an algorithm with running time $T=T(n, m)$ that, given a directed weighted graph $G$ on $n$ vertices and $m$ edges, computes a set of weighted edges $E^{\prime}$ such that, (1) when $E^{\prime}$ is added to $G$, all shortest path distances stay the same, and (2) $E^{\prime}$ forms a $(\beta, \varepsilon / 2)$-additive hopset of $G$.

Let $D$ be a real number and $k \geq 2$ be a positive integer, and $D^{\prime}=\left(2-\frac{1}{k}+\varepsilon\right) D$. Under NSETH (NUNSETH), for any $\delta>0, D^{\prime} / D$-Diameter on directed weighted graphs with time complexity $\left(m^{1+1 / k} \beta^{1-2 / k}+T\right)^{1+\delta}$ is not SETH-hard for deterministic (randomized) reductions.

By Lemma 3.3 and Lemma 3.4, Theorem 5.2 follows from the following theorem.
Theorem 5.3. Suppose $\varepsilon>0$ and the algorithm of Theorem 5.2 exists. Let $D$ be a real number and $k \geq 2$ be a positive integer, and $D^{\prime}=\left(2-\frac{1}{k}+\varepsilon\right) D$. Then $D^{\prime} / D$-Diameter on directed weighted graphs is in $(\mathrm{N} \cap \operatorname{coN})^{+} \operatorname{TIME}\left[m^{1+1 / k} \beta^{1-2 / k}+T\right]$.

To prove Theorem 5.3, we need a few technical lemmas. The first is a standard lemma about hitting sets.

Lemma 5.4. Let $n, M$ and $K$ be positive integers. Let $A_{1}, \ldots, A_{n}$ denote $\ell$ sets of size at least $K$ over a universe of size $M$. Then there exists a set $X$ of $2 M / K \log n$ elements of $M$ such that $X \cap A_{i} \neq \emptyset$ for all $i=1, \ldots, n$.

Proof. Let $X$ be $2 M / K \log n$ elements chosen uniformly at random with replacement from $U$. The probability $X \cap A_{i}=\emptyset$ is at most $(1-K / M)^{|X|} \leq e^{-|X| \cdot K / M}=n^{-2}$. By the union bound, the probability there exists an $i$ such that $X \cap A_{i}=\emptyset$ is at most $n \cdot n^{2}<1$, so some choice of $X$ yields $X \cap A_{i} \neq \emptyset$ for all $i$.

We apply Lemma 5.4 in the following structural result about graphs with diameter $D$.
Lemma 5.5. Let $G$ be a directed weighted graph of diameter $D$, and let $M_{1}, M_{2}$ be positive real numbers with $M_{1} M_{2} \geq 8 m \log m$, and let $D_{1}, D_{2}$ be positive real numbers with $D_{1}+D_{2} \geq D$. Then $G$ either has a vertex subset $Y$ satisfying $|Y| \leq M_{1}$ and $d(v, Y) \leq D_{1}$ for all vertices $v$, or $G$ has a vertex subset $Z$ satisfying $|Z| \leq M_{2}$ and $d(Z, v) \leq D_{2}$.

Proof. Suppose there exists a vertex $u$ with $\left|E_{\leq D_{2}}^{i n}(u)\right| \leq M_{1} / 2$. Then set $X_{1}$ to be the vertices incident to $E_{\leq D_{2}}^{i n}(u)$. For each vertex $v$, there is a length at most $D$ path $v=v_{0}, v_{1}, \ldots, v_{\ell}=u$ from
$v$ to $u$. If $i$ is the largest index such that $d\left(v_{i}, u\right)>D_{2}$, we have $v_{i} \in X_{1}$, so $d\left(v, v_{i}\right)<D-D_{2}<D_{1}$. Thus, $d\left(v, X_{1}\right)<D_{1}$, as desired.

Now suppose that for all vertices $u$, we have $\left|E_{\leq D_{2}}^{i n}(u)\right| \geq M_{1} / 2$. By Lemma 5.4 on the universe of $m$ edges with $n$ sets $E_{\leq D_{2}}^{i n}(u)$, there exists a set $\bar{F}$ of $2 m /\left(M_{1} / 2\right) \cdot \log n \leq M_{2} / 2$ edges intersecting each $E_{\leq D_{2}}^{i n}(u)$. Setting $X_{2}$ to be the vertices incident to $F$ gives $\left|X_{2}\right| \leq M_{2}$ and $d\left(X_{2}, u\right) \leq D_{2}$ for all vertices $u$.

We now prove Theorem 5.3.
Proof of Theorem 5.3. Let $G$ be the input graph. Recall that $D^{\prime} / D$-Diameter is a promise problem $\left(\Pi_{Y E S}, \Pi_{N O}\right)$ where the instances in $\Pi_{Y E S}$ are graphs of diameter at least $D^{\prime}$, and the instances in $\Pi_{N O}$ are graphs of diameter at most $D$. Let $\mathcal{A}_{N}$ be the nondeterministic algorithm that guesses a vertex, outputs "YES" if the eccentricity is greater than $D$, and "NO" otherwise. For any graph not in $\Pi_{N O}$, i.e., any graph with diameter greater than $D$, there exist nondeterministic choices such that $\mathcal{A}_{N}$ outputs "YES". Furthermore, for graphs in $\Pi_{N O}, \mathcal{A}_{N}$ always outputs "NO". As $\mathcal{A}_{N}$ runs in time $m^{1+o(1)}$, we have that properties 1 and 3 of Definition 3.2 are satisfied.

It now remains to construct a nondeterministic algorithm $\mathcal{A}_{c o N}$ satisfying property 2 of Definition 3.2. Let $\mathcal{A}_{\text {hopset }}$ be the algorithm that is assumed with parameters $\beta$ and $\varepsilon$. Define $\mathcal{A}_{\text {coN }}$ to be the following algorithm.

1. For $\ell=1, \ldots, k-1$, nondeterministically choose vertex subsets $X_{\ell}^{i n}$ and $X_{\ell}^{\text {out }}$ of size $8 m^{1-(\ell / k)} \beta^{-1+2 \ell / k} \log m$.
2. Run single source shortest path from every vertex in $X_{k-1}^{i n}$. If each vertex in $X_{k-1}^{i n}$ has out-eccentricity at most $D$, and if $d\left(v, X_{k-1}^{i n}\right) \leq\left(\frac{k-1}{k}\right) D$ for all vertices $v$, output "NO".
3. Run single source shortest path from every vertex in $X_{k-1}^{o u t}$. If each vertex in $X_{k-1}^{o u t}$ has in-eccentricity at most $D$, and if $d\left(X_{k-1}^{\text {out }}, v\right) \leq\left(\frac{k-1}{k}\right) D$ for all vertices $v$, output "NO".
4. Nondeterministically choose a random seed for the hopset algorithm $\mathcal{A}_{\text {hopset }}$.
5. Run $\mathcal{A}_{\text {hopset }}$ to obtain a set $E^{\prime}$ and a new graph $G^{\prime}=G+E^{\prime}$.
6. For $\ell=1, \ldots, k-2$, run multi-source shortest path from each of $X_{\ell}^{\text {in }}$ and $X_{\ell}^{\text {out }}$.
7. Let $\ell \in\{1, \ldots, k-2\}$ be such that $d\left(v, X_{\ell}^{\text {in }}\right) \leq\left(\frac{\ell}{k}\right) D$ and $d\left(X_{k-1-\ell}^{\text {out }}, v\right) \leq\left(\frac{k-\ell}{k}\right) D$ for all vertices $v$. If no such vertices exist, output "YES".
8. Nondeterministically choose a sequence $s_{x, x^{\prime}}$ of at most $\beta$ vertices for every $x \in X_{\ell}^{i n}$ and $x^{\prime} \in X_{k-1-\ell}^{\text {out }}$. If each $s_{x, x^{\prime}}$ forms a path from $x$ to $x^{\prime}$ of length at most $D(1+\varepsilon / 2)$ in $G^{\prime}$, and $d\left(v, X_{\ell}^{i n}\right) \leq D / 3$ and $d\left(X_{k-1-\ell}^{\text {out }}, v\right) \leq D / 3$ for all vertices $v$ output "NO". Otherwise output "YES".

Runtime. Choose the sets $X_{\ell}^{\text {in }}, X_{\ell}^{\text {out }}$ takes time $\tilde{O}(m)$. Steps 2 and 3 each take $\tilde{O}\left(m^{1+1 / k} \beta^{1-2 / k}\right)$, as each shortest path takes time $\tilde{O}(m)$. Steps 4 and 5 together take time $T$ by assumption on algorithm $\mathcal{A}_{\text {hopset }}$. Step 6 takes time $\tilde{O}(m)$. Step 7 takes time $\tilde{O}\left(\left|X_{\ell}^{\text {in }}\right| \cdot\left|X_{\ell}^{\text {out }}\right| \cdot \beta\right)=\tilde{O}\left(m^{1+1 / k} \beta^{1-2 / k}\right)$.

Thus, the total runtime is $\tilde{O}\left(m^{1+1 / k} \beta^{1-2 / k}+T\right)$.
Correctness. We first show that if the algorithm outputs "NO", the diameter is less than $D^{\prime}$. If we output "NO" at Step 2, then for any vertices $v$ and $v^{\prime}$, there exists a vertex $x \in X_{k-1}^{i n}$ such that $d(v, x) \leq\left(\frac{k-1}{k}\right) D$. As $x \in X_{k-1}^{i n}$ has out-eccentricity at most $D$, we have $d\left(v, v^{\prime}\right) \leq$ $d(v, x)+d\left(x, v^{\prime}\right) \leq\left(\frac{2 k-1}{k}\right) D<D^{\prime}$ by the triangle inequality. If we output "NO" at Step 3, then for any vertices $v$ and $v^{\prime}$, there exists a vertex $x \in X_{k-1}^{\text {out }}$ such that $d\left(x, v^{\prime}\right) \leq\left(\frac{k-1}{k}\right) D$. As $x \in X_{k-1}^{\text {out }}$
has in-eccentricity at most $D$, we have $d\left(v, v^{\prime}\right) \leq d(v, x)+d\left(x, v^{\prime}\right) \leq\left(\frac{2 k-1}{k}\right) D<D^{\prime}$ by the triangle inequality. If we output "NO" at Step 7 , then for any $v, v^{\prime} \in V$, there exists $x$ and $x^{\prime}$ such that $d_{G^{\prime}}(v, x) \leq\left(\frac{\ell}{k}\right) D$, and $d_{G^{\prime}}\left(x^{\prime}, v^{\prime}\right) \leq\left(\frac{k-1-\ell}{k}\right) D$, and we also must have $d_{G^{\prime}}\left(x, x^{\prime}\right) \leq(1+\varepsilon / 2) D$ as $x \in X_{\ell}^{\text {in }}$ and $x^{\prime} \in X_{k-1-\ell}^{\text {out }}$. Then $d_{G^{\prime}}\left(v, v^{\prime}\right) \leq d_{G^{\prime}}(v, x)+d_{G^{\prime}}\left(x, x^{\prime}\right)+d_{G^{\prime}}\left(x^{\prime}, v^{\prime}\right) \leq\left(\frac{2 k-1}{k}+\varepsilon / 2\right) D<$ $D^{\prime}$ by the triangle inequality. As, shortest path distances are the same in $G^{\prime}$ as in $G$, we have $d_{G}\left(v, v^{\prime}\right)<D^{\prime}$ for all $v, v^{\prime} \in V$, as desired.

Now we show that the if the diameter is at most $D$, there exists a sequence of nondeterministic choices such that our algorithm outputs "NO". For $\ell=1, \ldots, k-1$, if there exists $X_{\ell}^{i n}$ exists such that $d\left(v, X_{\ell}^{i n}\right) \leq\left(\frac{\ell}{k}\right) D$, let $X_{\ell}^{i n}$ be that set. Similarly if there exists $X_{\ell}^{\text {out }}$ exists such that $d\left(v, X_{\ell}^{\text {out }}\right) \leq\left(\frac{\ell}{k}\right) D$, let $X_{\ell}^{\text {out }}$ be that set. Call such an $X_{\ell}^{\text {in }}$ or $X_{\ell}^{\text {out }}$ good. Let $X_{\ell}^{\text {in }}$ or $X_{\ell}^{\text {out }}$ be arbitrary if it is not good, and call such an $X_{\ell}^{i n}$ or $X_{\ell}^{o u t}$ bad. Additionally, let the randomness of $\mathcal{A}_{\text {hopset }}$ be such that $E^{\prime}$ is indeed a $(\beta, \varepsilon / 2)$-hopset.

By Lemma 5.5 with $D_{1}=(\ell / k) D, D_{2}=((k-\ell) / k) D, M_{1}=\left|X_{\ell}^{i n}\right|=8 m^{1-\ell / k} \beta^{-1+2 \ell / k} \log m$, and $M_{2}=\left|X_{k-\ell}^{\text {out }}\right|=8 m^{\ell / k} \beta^{1-2 \ell / k} \log m$, we have that either $X_{\ell}^{i n}$ is good or $X_{k-\ell}^{\text {out }}$ is good for all $\ell=1, \ldots, k-1$. Thus, among, $X_{1}^{i n}, \ldots, X_{k-1}^{\text {in }}, X_{1}^{\text {out }}, \ldots, X_{k-1}^{o u t}$, there are at most $k-1$ bad vertex subsets. Hence, among the $k$ sets $\left\{X_{k-1}^{\text {in }}\right\},\left\{X_{k-1}^{o u t}\right\}$, and $\left\{X_{\ell}^{\text {in }}, X_{k-1-\ell}^{\text {out }}\right\}$ for $\ell=1, \ldots, k-2$, one of these sets of vertex subsets has only good vertex subsets. First, if $X_{k-1}^{i n}$ is good, then choosing this $X_{k-1}^{i n}$ as $X_{k-1}^{i n}$ in the algorithm causes step 2 to output "NO": we indeed will see that $d\left(v, X_{k-1}^{i n}\right) \leq\left(\frac{k-1}{k}\right) D$, and all eccentricities are at most $D$ because the diameter is at most $D$. Similarly, if $X_{k-1}^{\text {out }}$ is good, then choosing this $X_{k-1}^{o u t}$ as $X_{k-1}^{o u t}$ in the algorithm causes step 3 to output "NO". Finally, if there exists $X_{\ell}^{i n}$ and $X_{k-1-\ell}^{o u t}$ such that both are good, our algorithm will output "NO" in step 7: since $E^{\prime}$ is a $(\beta, \varepsilon / 2)$-hopset, there there exists an at-most- $\beta$ edge path from $x$ to $x^{\prime}$ in $G^{\prime}$ for any $x, x^{\prime} \in X$ of length at most $d_{G}\left(x, x^{\prime}\right)+(\varepsilon / 2) D \leq(1+\varepsilon / 2) D$, so nondeterministically choosing these paths causes us to output "NO". This covers all possible cases, so we these nondeterministic choices cause us to output "NO", as desired.

Proof of Theorem 5.2. With Theorem 5.3, apply Lemma 3.3 (for deterministic reductions) and apply Lemma 3.4 (for randomized reductions).

Proof of Theorem 5.1. In Theorem 5.2, let $k=3$, and let the proposed hopset algorithm be given by Lemma 3.6 with $\alpha=3 / 26$, which indeed computes a hopset for all directed weighted graphs. It follows that $(5 / 3+\varepsilon) D / D$-Diameter on directed weighted graphs with time complexity $\left(m^{4 / 3+\frac{1 / 2-3 / 26}{3}}+m^{1+4 \cdot 3 / 26}\right)^{1+\delta}=m^{19 / 13(1+\delta)}$ is not SETH-hard. This yields Theorem 5.1.

Proof of Theorem 1.5. Apply Lemma 3.5 with $\rho=\infty, \alpha=\frac{5}{3}+\frac{\varepsilon}{2}, \beta=\frac{\varepsilon}{2}$, and $T=m^{(1+\delta) 19 / 13}$ to Theorem 5.1.

### 5.2 Undirected weighted graphs

Using the same technique, we can prove non-reducibility results for undirected graphs in a larger parameter setting.

Theorem 5.6. Let $k \geq 2$ be a positive integer and $\varepsilon>0$. Under NSETH (NUNSETH), for any $\delta>0$, a $2-\frac{1}{k}+\varepsilon$-approximation of Diameter on undirected weighted graphs with time complexity $m^{1+1 / k+\delta}$ is not SETH-hard for deterministic (randomized) reductions.

Again, by Lemma 3.5, it suffices to prove the following Theorem.

Theorem 5.7. Let $D$ be a positive real number, $k \geq 2$ be a positive integer, $\varepsilon>0$ and $D^{\prime}=$ $\left(2-\frac{1}{k}+\varepsilon\right) D$. Under NSETH (NUNSETH), for any $\delta>0, D^{\prime} / D$-Diameter on undirected weighted graphs with time complexity $m^{1+1 / k+\delta}$ is not SETH-hard for deterministic (randomized) reductions.
Proof. Let $D$ be a real number, $k \geq 2$ be a positive integer, and $D^{\prime}=\left(2-\frac{1}{k}+\varepsilon\right) D$ be as in Theorem 5.6. We note that Theorem 5.2 also holds for undirected graphs, when one assumes a hopset construction for unweighted graphs. Since one can in fact use the same proof, we omit the proof for brevity. By Lemma 3.7, there exists an algorithm running in time $T=\tilde{O}\left(m^{1+1 / k}\right)$, given an undirected weighted graph $G$ on $n$ vertices and $m$ edges, computes a set $E^{\prime}$ of weighted edges such that (1) when $E^{\prime}$ is added to $G$, all shortest path distances stay the same, and (2) $E^{\prime}$ forms a $(\beta, \varepsilon / 2)$-additive hopset of $G$, where $\beta=O_{k, \varepsilon}(1)$. Then, following the proof of Theorem 5.2, we have that for all $\delta>0, D^{\prime} / D$-Diameter on undirected weighted graphs with time complexity $\left(m^{1+1 / k} \beta^{1-2 / k}+T\right)^{1+\delta}$ is not SETH-hard. As $\beta=O_{k, \varepsilon}(1)$ and $T=\tilde{O}\left(m^{1+1 / k}\right)$, for all $\delta>0$, $D^{\prime} / D$-Diameter on undirected weighted graphs with time complexity $m^{1+1 / k+\delta}$ is not SETH-hard, as desired.

Proof of Theorem 5.6. Apply Lemma 3.5 to Theorem 5.7 with $\rho=\infty, \alpha=2-\frac{1}{k}+\frac{\varepsilon}{2}, \beta=\frac{\varepsilon}{2}$, and $T=m^{1+1 / k}$.

## 6 Optimal non-reducibility for directed unweighted graphs

We prove the following theorem, which is a generalization of Theorem 1.7.
Theorem 6.1. Let $k \geq 2$ be a positive integer. Assuming NSETH (NUNSETH), for any $\delta>0$ and $\varepsilon>0$, a $2-\frac{1}{k}+\varepsilon$ approximation of Diameter on directed weighted graphs satisfying $W_{\max } / W_{\min } \leq$ $m^{o(1)}$ with time complexity $m^{1+1 / k+\delta}$ is not SETH-hard for deterministic (randomized) reductions.

As in Section 5, by Lemma 3.5, to prove non-reducibility of the approximation problem, Theorem 6.1, it suffices to prove the non-reducibility of the promise problem.

Theorem 6.2. Let $D$ be a real number, $k \geq 2$ be a positive integer, $\varepsilon>0$ and $D^{\prime}=\left(2-\frac{1}{k}+\varepsilon\right) D$. Under NSETH (NUNSETH), for any $\delta>0, D^{\prime} / D$-Diameter on directed weighted graphs with time complexity $m^{1+1 / k+\delta} \cdot\left(W_{\max } / W_{\min }\right)$ is not SETH-hard for deterministic (randomized) reductions.

By Lemma 3.3 and Lemma 3.4, to prove Theorem 6.2, it suffices to prove the following result.
Theorem 6.3. Let $D$ be a real number, $k$ be a positive integer, $\varepsilon>0$ and $D^{\prime}=\left(2-\frac{1}{k}+\varepsilon\right) D$. $D^{\prime} / D$-Diameter on directed weighted graphs with $W_{\text {min }}>0$ is in $(\mathrm{N} \cap \operatorname{coN})^{+} \operatorname{TIME}\left[\tilde{O}\left(m^{1+1 / k}\right.\right.$. $\left.\left.\left(W_{\max } / W_{\min }\right)\right)\right]$.

We prove Theorem 6.3 even when $D$ can depend on $m$. Since $D$ can depend on $m$, we may without loss of generality assume $W_{\min }=1$, by scaling the edge weights by $1 / W_{\min }$ so that the new parameters $D$ and $D^{\prime}$ are $D / W_{\min }$ and $D^{\prime} / W_{\min }$, respectively, and the quantity $W_{\max } / W_{\min }$ remains unchanged. We now show that $D^{\prime} / D$-Diameter on directed weighted graphs with $W_{\min }=1$ is in $(\mathrm{N} \cap \operatorname{coN})^{+}$TIME $\left[m^{1+1 / k+o(1)} W_{\max }\right]$.

Throughout this section, fix a positive integer $k$ and a positive real number $D$, and a real number $\varepsilon>0$. Let $r=\varepsilon D / 2$. Throughout, $G=(V, E)$ is a graph with minimum edge-weight 1 and maximum edge-weight $W_{\max }$. Suppose we are looking to distinguish between graphs of diameter at most $D$ and graphs of diameter at least $D^{\prime}=\left(2-\frac{1}{k}+\varepsilon\right) D$.

For $\ell=0, \ldots, k$, call a subset $X$ of vertices $(\ell, r, k, D)$-out (or simply $\ell$-out when $r, k, D$ are understood) if $|X| \leq 8 m^{1-\ell / k} W_{\max } \log m$ and call it $\operatorname{good}(\ell, r, k, D)$-out (good $\ell$-out) if additionally
$d(X, v) \leq \frac{\ell}{k} D+r$ for all vertices $v$. For $\ell=0, \ldots, k$, call a subset $X$ of vertices $(\ell, r, k, D)$-in (or simply $\ell$-in when $r, k, D$ are understood) if $|X| \leq 8 m^{1-\ell / k} r^{-1} \log m$ and call it $\operatorname{good}(\ell, r, k, D)$-in $\left(\operatorname{good} \ell\right.$-in) if additionally $d(v, X) \leq \frac{\ell}{k} D$ for all vertices $v$.

Lemma 6.4. For all $\ell=1, \ldots, k-1$, there either exists a set $X$ that is good $\ell$-out or a set $X^{\prime}$ that is good $k-\ell$-in.

Proof. Suppose there exists a vertex $v$ with $\left|E_{\leq \frac{k-\ell}{k} D}^{o u t}(v)\right| \leq m^{1-\ell / k} r$. Let $U$ denote the vertices incident to $E^{*} \stackrel{\text { def }}{=} E_{\leq \frac{k \ell}{k} D}^{o u t}(v)$. For every vertex $u$, let $Z_{u}$ denote the set of vertices $w$ such that $d(w, u) \leq \frac{\ell}{k} D+r$ and such that $w$ is incident to an edge in $E^{*}$.

We claim that each $Z_{u}$ has size at least $r$. consider a shortest path $v=z_{0}, z_{1}, \ldots, z_{s}=u$ from $v$ to $u$. Let $i$ denote the largest index such that $d\left(z_{i}, u\right)>\frac{\ell}{k} D+r$ and $j$ denote the smallest index such that $d\left(v, z_{j}\right)>\frac{k-\ell}{k} D$. We must have $j>i$ or else $d(v, u) \geq d\left(v, z_{j}\right)+d\left(z_{i}, u\right)>D+r$, contradicting diameter $D$. By minimality of $j$, we have $d\left(v, z_{j^{\prime}}\right)<\frac{k-\ell}{k} D$ for $j^{\prime} \leq j-1$. In particular $z_{i+1}, z_{i+2}, \ldots, z_{j-1}$ is in $N_{\leq \frac{k-\ell}{k} D}^{\text {out }}(v)$, so $z_{j^{\prime}-1} z_{j^{\prime}}$ is in $E^{*}$ for $j^{\prime} \leq j$ and thus $z_{j^{\prime}} \in U$ for $j^{\prime} \leq j$. By maximality of $i$, we have $d\left(z_{j^{\prime}}, u\right) \leq \frac{\ell}{k} D+r$ for $j^{\prime}>i$, so $z_{i+1}, \ldots, z_{j}$ are in $Z_{u}$ and $\left|Z_{u}\right| \geq j-i$. We note that

$$
\begin{aligned}
D \geq d(v, u) & =d\left(v, z_{j}\right)+d\left(z_{i}, u\right)-d\left(z_{i}, z_{j}\right) \\
& >\frac{k-\ell}{k} D+\frac{\ell}{k} D+r-d\left(z_{i}, z_{j}\right)=D+r-d\left(z_{i}, d_{j}\right)
\end{aligned}
$$

We conclude that $d\left(z_{i}, z_{j}\right)>r$, so in particular $\left|Z_{u}\right| \geq j-i>r / W_{\max }$.
By Lemma 5.4 with universe as vertex set $U$, which has size at most $M \stackrel{\text { def }}{=} m^{1-\ell / k} r$, with $n$ sets $Z_{u} \subset U$ of size at least $K \stackrel{\text { def }}{=} r / W_{\max }$, we have that there exists a vertex subset $X \subset U$ of size at most $2 M / K \log n=2 m^{1-\ell / k} W_{\max } \log n$ that intersects each $Z_{u}$. This implies that $d(X, u) \leq \frac{\ell}{k} D+r$ for all vertices $u \in V$. We conclude that $X$ is good $\ell$-out.

Now suppose to the contrary that $\left|E_{\leq \frac{k-\ell}{k} D}^{o u t}(v)\right| \geq m^{1-\ell / k} r$ for all vertices $v$. By Lemma 5.4 with universe as the set $E$ of all edges, which has size $m$, with $n$ sets $E_{\leq \frac{k-\ell}{k} D}^{o u t}(v)$ of size at least $m^{1-\ell / k} r$, we have that there exists a edge subset $F$ of size $2 M / K \log n=2 m^{\ell / k} r^{-1} \log n$. Let $X^{\prime}$ be the vertices incident to $F$, so that $\left|X^{\prime}\right| \leq 4 m^{\ell / k} r^{-1} \log m$. Then, $X^{\prime}$ contains a vertex inside each $N_{\leq \frac{k-\ell}{k}}^{\text {out }}(v)$, so $d\left(v, X^{\prime}\right) \leq \frac{k-\ell}{k} D$ for all vertices $v$. We conclude $X^{\prime}$ is $\operatorname{good}(k-\ell)$-in.

Lemma 6.5. Let $G$ be a graph with a good $(k-\ell)$-in set $X$ and a good $(\ell-1)$-out set $X^{\prime}$. Suppose further that $d\left(x, x^{\prime}\right) \leq D$ for all $x \in X$ and $x^{\prime} \in X^{\prime}$. Then $G$ has diameter less than $D^{\prime}$.

Proof. Let $v$ and $v^{\prime}$ be vertices in $G$. As $X$ is good $(k-\ell)$-in, we have $d(v, X) \leq \frac{k-\ell}{k} D$. As $X^{\prime}$ is $\operatorname{good}(\ell-1)$-out, we have $d\left(X^{\prime}, v^{\prime}\right) \leq \frac{\ell-1}{k} D+r$. Thus, there exists $x \in X$ and $x^{\prime} \in X^{\prime}$ such that $d(v, x) \leq \frac{k-\ell}{k} D$ and $d\left(x^{\prime}, v^{\prime}\right) \leq \frac{\ell-1}{k} D+r$. By the triangle inequality, we have

$$
\begin{aligned}
d\left(v, v^{\prime}\right) \leq d(v, x)+d\left(x, x^{\prime}\right)+d\left(x^{\prime}, v^{\prime}\right) & \leq \frac{k-\ell}{k} D+D+\frac{\ell-1}{k} D+r \\
& =\left(2-\frac{1}{k}\right) D+r<D^{\prime}
\end{aligned}
$$

This holds for all $v, v^{\prime}$, so the diameter is at most $\left(2-\frac{1}{k}+\varepsilon\right) D$.

Lemma 6.6. Let $G$ be a graph with a good $(k-1)$-out set $X$. Suppose further that $d(v, x) \leq D$ for all vertices $x \in X$ and all vertices $v \in V$. Then $G$ has diameter less than $D^{\prime}$.

Proof. For any vertices $v, v^{\prime} \in V$, there exists a vertex $x \in X$ such that $d\left(x, v^{\prime}\right) \leq \frac{k-1}{k} D+r$. Then by the triangle inequality, $d\left(v, v^{\prime}\right) \leq d(v, x)+d\left(x, v^{\prime}\right) \leq D+\frac{k-1}{k} D+r<D^{\prime}$. This holds for all $v, v^{\prime}$, so the diameter is less than $D^{\prime}$.

Lemma 6.7. For any $\ell=1, \ldots, k-1$, given a $\ell$-out ( $\ell$-in) set $X$ it is possible to check in time $\tilde{O}(m)$ whether $X$ is good $\ell$-out ( $\ell$-in).

Proof. Run multi-source shortest path from $X$ to compute all $d(X, v)$ or all $d(v, X)$.
We now present our algorithm.
Proof of Theorem 6.3. Let $G$ be the input graph. Again recall that $D^{\prime} / D$-Diameter is a promise problem $\left(\Pi_{Y E S}, \Pi_{N O}\right)$ where the instances in $\Pi_{Y E S}$ are graphs of diameter at least $D^{\prime}$, and the instances in $\Pi_{N O}$ are graphs of diameter at most $D$. Let $\mathcal{A}_{N}$ be the nondeterministic algorithm that guesses a vertex, outputs "YES" if the eccentricity is greater than $D$, and "NO" otherwise. For any graph not in $\Pi_{N O}$, i.e., any graph with diameter greater than $D$, there exist nondeterministic choices such that $\mathcal{A}_{N}$ outputs "YES". Furthermore, for graphs in $\Pi_{N O}, \mathcal{A}_{N}$ always outputs "NO". As $\mathcal{A}_{N}$ runs in time $m^{1+o(1)}$, we have that properties 1 and 3 of Definition 3.2 for $(\mathrm{N} \cap \operatorname{coN})^{+} \mathrm{TIME}\left[m^{1+1 / k+o(1)} W_{\max }\right]$ are satisfied.

It now remains to construct a nondeterministic algorithm $\mathcal{A}_{\text {coN }}$ satisfying property 2 of Definition 3.2 for $(\mathrm{N} \cap \operatorname{coN})^{+}$TIME $\left[m^{1+1 / k+o(1)} W_{\max }\right]$. Define $\mathcal{A}_{c o N}$ to be the following algorithm.

1. For each $\ell=0, \ldots, k$, nondeterministically choose a set $X_{\ell}^{\text {out }}$ and $X_{\ell}^{\text {in }}$, such that $X_{\ell}^{\text {out }}$ is $\ell$-out and $X_{\ell}^{i n}$ is $\ell$-in.
2. For each $\ell=0, \ldots, k$, check if $X_{\ell}^{\text {out }}$ is good $\ell$-out, and check if $X_{\ell}^{\text {in }}$ is good $\ell$-in.
3. If $X_{k-1}^{\text {out }}$ is good $(k-1)$-out, then run shortest path from each element of $X_{k-1}^{o u t}$. If each vertex in $X_{k-1}^{\text {out }}$ has in-eccentricity at most $D$, output "NO".
4. Let $\ell \in\{0, \ldots, k-1\}$ be such that $X_{\ell-1}^{\text {out }}$ is $\operatorname{good}(\ell-1)$-out and $X_{k-\ell}^{\text {in }}$ is $\operatorname{good}(k-\ell)$-in. If no such $\ell$ exists, then output "YES".
5. For each $x \in X_{\ell-1}^{\text {out }}$ and each $x^{\prime} \in X_{k-\ell}^{i n}$, nondeterministically choose up to $D$ vertices that form a path from $x$ to $x^{\prime}$.
6. If each path has length at most $D$, output "NO". Otherwise, output "YES".

Runtime. The first step takes time $\tilde{O}(m)$ to choose the $X_{\ell}^{\text {out }}$. The second step takes time $\tilde{O}(m)$, as it takes $\tilde{O}(m)$ time to check if $X_{\ell}^{\text {out }}$ is good $\ell$-out or $X_{\ell}^{\text {in }}$ is good $\ell$-in by Lemma 6.7. The third step takes time $\tilde{O}\left(m \cdot\left|X_{k-1}^{\text {out }}\right|\right) \leq \tilde{O}\left(m^{1+1 / k} W_{\max }\right)$. The fourth step takes constant time (we assume $k$ is constant). The fifth step takes time $O\left(\left|X_{\ell-1}^{\text {out }}\right| \cdot\left|X_{k-\ell}^{\text {in }}\right| \cdot D\right)=\tilde{O}\left(m^{1-(\ell-1) / k} \cdot m^{\ell / k} /(\varepsilon D) \cdot D\right)=$ $\tilde{O}_{\varepsilon}\left(m^{1+1 / k}\right)$. The total running time is thus $\tilde{O}_{\varepsilon}\left(m^{1+1 / k} W_{\max }\right)$

Correctness. We first show that, if we output "NO", the diameter must be less than $D^{\prime}$. Suppose we output "NO" at step 3. Then each $x \in X_{k-1}^{\text {out }}$ satisfies $d(v, x) \leq D$ for all $v \in V$. Then, by Lemma 6.6, the diameter is less than $D^{\prime}$.

Now suppose we output "NO" at step 4, and let $\ell$ be the parameter chosen in Step 4. Then we know $X_{\ell-1}^{\text {out }}$ is $\operatorname{good}(\ell-1)$-out and $X_{k-\ell}^{\text {in }}$ is good $(k-\ell)$-in. Since we output "NO", we also have
that $d\left(x, x^{\prime}\right) \leq D$ for each $x \in X_{\ell-1}^{\text {out }}$ and each $x^{\prime} \in X_{k-\ell}^{\text {in }}$. Setting $X=X_{k-\ell}^{\text {in }}$ and $X^{\prime}=X_{\ell-1}^{\text {out }}$ in Lemma 6.5, we have that the diameter is less than $\left(2-\frac{1}{k}+\varepsilon\right) D=D^{\prime}$, as desired.

We now show that, if the diameter is at most $D$, there exists some nondeterministic choices such that we output "NO". For each $\ell=0, \ldots, k$, if there exists a good $\ell$-in set, let $X_{\ell}^{\text {in }}$ be that set, and if there exists a good $\ell$-out set, let $X_{\ell}^{\text {out }}$ be that set. By Lemma 6.4, for each pair of sets $\left(X_{1}^{\text {out }}, X_{k-1}^{i n}\right), \ldots,\left(X_{k-1}^{\text {out }}, X_{1}^{i n}\right)$, at least one must satisfy the corresponding goodness property. As the set of all vertices is good 0 -out by definition, we have $X_{0}^{\text {out }}$ is good 0 -out. Hence, among, $X_{0}^{\text {out }}, X_{1}^{\text {out }}, \ldots, X_{k-1}^{\text {out }}, X_{1}^{\text {in }}, \ldots, X_{k-1}^{\text {in }}$, there are at least $k$ sets that satisfy their corresponding goodness properties and at most $k-1$ sets that do not.

If $X_{k-1}^{\text {out }}$ is $\operatorname{good}(k-1)$-out, then we will output "NO" at step 3, as the diameter is at most $D$. Furthermore, if, for any $\ell, X_{k-\ell}^{\text {out }}$ is good $(k-\ell)$-out and $X_{\ell-1}^{i n}$ is good $(\ell-1)$-in, then we output "NO" because we choose the corresponding parameter $\ell$ in Step 4, and the path of length $D$ between every pair of vertices $x \in X_{k-\ell}^{o u t}$ and $x^{\prime} \in X_{\ell-1}^{i n}$ exists. Thus, to not output "NO", one set from each of $\left\{X_{k-1}^{\text {out }}\right\},\left\{X_{k-2}^{\text {out }}, X_{1}^{\text {in }}\right\}, \ldots,\left\{X_{0}^{\text {out }}, X_{1}^{\text {in }}\right\}$ must not satisfying the corresponding goodness properties, so $k$ sets cannot satisfy their goodness properties, which is a contradiction of the previous paragraph. Thus, there exist nondeterministic choices such that we output "NO", as desired.

Proof of Theorem 6.2. With Theorem 6.3, apply Lemma 3.3 (for deterministic reductions) and apply Lemma 3.4 (for randomized reductions).

Proof of Theorem 1.7 and Theorem 6.1. Apply Lemma 3.5 with $\rho=m^{o(1)}, \alpha=2-\frac{1}{k}+\frac{\varepsilon}{2}, \beta=\frac{\varepsilon}{2}$, and $T=m^{1+1 / k+\delta}$ to Theorem 6.2.

Proof of Corollary 1.8. First suppose $\lfloor 1 / \varepsilon\rfloor \leq\lfloor 1 / \delta\rfloor$. Let $k=\lfloor 1 / \varepsilon\rfloor+1$. Then $2-\varepsilon<2-\frac{1}{k}$ so there exists a positive $\varepsilon^{\prime}>0$ such that $2-\varepsilon=2-\frac{1}{k}-\varepsilon^{\prime}$. Furthermore, $k-1 \leq\lfloor 1 / \delta\rfloor \leq 1 / \delta$, so $1+\delta \leq 1+1 /(k-1)$. In Theorem 1.3, it is shown that a $2-\frac{1}{k}-\varepsilon^{\prime}=2-\varepsilon$ approximation of the Diameter with time complexity $n^{1+1 /(k-1)}$ is SETH-hard. Hence, since decreasing time complexity preserves SETH-hardness, a $2-\varepsilon$-approximation of the Diameter with time complexity $n^{1+\delta}$ is SETH-hard, as desired.

Now suppose $\lfloor 1 / \varepsilon\rfloor>\lfloor 1 / \delta\rfloor$ and $1 / \varepsilon \neq\lfloor 1 / \delta\rfloor+1$. Let $k=\lfloor 1 / \delta\rfloor+1$. Since $1 / \varepsilon \geq\lfloor 1 / \varepsilon\rfloor \geq$ $\lfloor 1 / \delta\rfloor+1=k$, and $1 / \varepsilon \neq\lfloor 1 / \delta\rfloor+1$, we have $1 / \varepsilon>k$. Thus, there exists an $\varepsilon^{\prime}>0$ such that $2-\frac{1}{k}+\varepsilon^{\prime}=2-\varepsilon$. Additionally, $1 / k<\delta$, so there exists $\delta^{\prime}$ such that $1+1 / k+\delta^{\prime}=1+\delta$. Theorem 1.7 proves that, under NSETH (NUNSETH), a $2-\frac{1}{k}+\varepsilon^{\prime}$ approximation of Diameter in directed unweighted graphs with time complexity $m^{1+1 / k+\delta^{\prime}}$ is not SETH-hard for deterministic (randomized) reductions, so a $2-\varepsilon$ approximation of Diameter in directed unweighted graphs with time complexity $m^{1+\delta}$ is not SETH-hard for deterministic (randomized) reductions, as desired.

## 7 Lower bound for undirected unweighted graphs

In this section, we prove the following result, implying there is no $5 / 3-\varepsilon$ approximation of the diameter of an undirected unweighted graph in near-linear time.

Theorem 7.1. Assuming SETH, for all $\varepsilon>0 a\left(\frac{5}{3}-\varepsilon\right)$-approximation of Diameter in unweighted, undirected graphs on $n$ vertices needs $n^{3 / 2-o(1)}$ time.

Proof. Start with a Single-Set 3 -OV instance $\Phi$ given by a set $A \subset\{0,1\}^{d}$ with $|A|=\tilde{n}$ and $d=c \log n_{N O V}$. We may add the all-1s vector to $A$ without loss of generality, as this does not change whether there is an OV solution or not. We construct a graph with $\tilde{O}\left(\tilde{n}^{2}\right)$ vertices and edges


Figure 6: 5 vs. 3 diameter instance. The coordinates $a[i]$ along the edges must be 1 for the edge to exist.
from the 3 -OV instance such that (1) if $\Phi$ has no solution, any two vertices are at distance 3 , and (2) if $\Phi$ has a solution, then there exists two vertices at distance 5 . Any $(5 / 3-\varepsilon)$-approximation for Diameter distinguishes between graphs of diameter 3 and 5 . Since solving $\Phi$ needs $\tilde{n}^{3-o(1)}$ time under SETH, a $5 / 3-\varepsilon$ approximation of diameter needs $n^{3 / 2-o(1)}$ time under SETH.

Construction of the graph The graph $G$ is illustrated in Figure 6 and constructed as follows. The vertex set $S \cup X$ is defined on

$$
\begin{aligned}
S & =A^{2}, \\
X & =\left\{(a, i, j) \in A \times[d]^{2}: a[i]=a[j]=1\right\}
\end{aligned}
$$

Throughout, we identify tuples $(a, b) \in A^{2}$ and $(a, i, j) \in A \times[d]^{2}$ with vertices of $G$, Throughout we denote vertices in $S$ and $X$ by $(a, b)_{S}$ and $(a, i, j)_{X}$, respectively. The (undirected unweighted) edges are all of the following.

- Edge between $(a, b)_{S}$ and $(a, i, j)_{X}$ if $(a, i, j)_{X}$ exists and $b[i]=1$.
- Edge between $(a, b)_{S}$ and $(a, i, j)_{X}$ if $(a, i, j)_{X}$ exists and $b[j]=1$.
- Edge between $(a, i, j)_{X}$ and $(b, i, j)_{X}$ always if vertices exist.
- Edge between $(a, i, j)_{X}$ and $\left(a, i^{\prime}, j^{\prime}\right)_{X}$ always if vertices exist.

Note that each vertex of $S$ has $O\left(d^{2}\right)$ neighbors, each vertex of $X$ has $O(\tilde{n})$ neighbors, and each vertex of $Y$ has $O(\tilde{n})$ neighbors. The total number of edges and vertices is thus $O\left(\tilde{n}^{2} d^{2}\right)=\tilde{O}\left(\tilde{n}^{2}\right)$. We now show that this construction has diameter 3 when $\Phi$ has no solution and diameter at least 5 when $\Phi$ has a solution.

3-OV no solution Assume that the 3-OV instance has no solution, so that no three (or two) vectors are orthogonal. We show that any pair of vertices have distance at most 3 , by casework on which of $S, X$ the two vertices are in.

- Both vertices are in $S$ : Let the vertices be $(a, b)_{S}$ and $(c, d)_{S}$. As there is no 3-OV solution, there exists indices $i$ and $j$ in $[d]$ such that $a[i]=b[i]=c[i]$ and $a[j]=c[j]=d[j]$. Then $(a, b)_{S}-(a, i, j)_{X}-(c, i, j)_{X}-(c, d)_{S}$ is a valid path.
- One vertex is in $S$ and the other vertex is in $X$ : Let the vertices be $(a, b)_{S}$ and $(c, i, j)_{X}$ : As there is no 3-OV solution, there exists an index $i^{\prime}$ such that $a\left[i^{\prime}\right]=b\left[i^{\prime}\right]=c\left[i^{\prime}\right]=1$. Then $(a, b)_{S}-\left(a, i^{\prime}, i^{\prime}\right)_{X}-\left(c, i^{\prime}, i^{\prime}\right)_{X}-(c, i, j)_{X}$ is a valid path.
- Both vertices are in $X$ : Let the vertices be $(a, i, j)_{X}$ and $\left(c, i^{\prime}, j^{\prime}\right)_{X}$. As there is no 3OV solution, there exists an index $i^{\prime \prime}$ such that $a\left[i^{\prime \prime}\right]=c\left[i^{\prime \prime}\right]$. Then $(a, i, j)_{X}-\left(a, i^{\prime \prime}, i^{\prime \prime}\right)_{X}-$ $\left(c, i^{\prime \prime}, i^{\prime \prime}\right)_{Y}-\left(c, i^{\prime}, j^{\prime}\right)_{X}$ is a valid path.

3-OV has solution Now assume that the 3-OV instance has a solution. That is, assume there exists $a, b, c \in A$ such that $a[i] \cdot b[i] \cdot c[i]=0$ for all $i$. We show there are no paths of length at most 4 from $(a, b)_{S}$ to $(c, b)_{S}$. Any such path must use an $X-X$ edge or else the first entry $a$ of the vertex's tuple does not change. Because of this, the path cannot revisit the set $S$, as it would otherwise need at least 5 edges.

First, suppose that the path does not use an index-changing edge, namely an edge of the form $(a, i, j)_{X}-\left(a, i^{\prime}, j^{\prime}\right)_{X}$. The path must be $(a, b)_{S}-(a, i, j)_{X}-(c, i, j)_{X}-(c, b)_{S}$ or $(a, b)_{S}-(a, i, j)_{X}-$ $(v, i, j)_{X}-(c, i, j)_{X}-(c, b)_{S}$ for some indices $i, j \in[d]$ and some vector $v$. In either case, the existence of vertex $(a, i, j)_{X}$ requires that $a[i]=a[j]=1$, and the existence of vertex $(c, i, j)_{X}$ requires that $c[i]=c[j]=1$. The first edge requires that at least one of the coordinates $b[i]$ or $b[j]$ is 1 . Thus, at least one of $a[i]=b[i]=c[i]=1$ or $a[j]=b[j]=c[j]=1$ holds, contradicting orthogonality of $a, b, c$.

Now suppose the path uses an index-changing edge. The index-changing edge must either be the second edge or the third edge. These cases are symmetric to each other so it suffices to consider only one. If the index-changing edge is the second edge, the path must be $(a, b)_{S}-(a, i, j)_{X}-$ $\left(a, i^{\prime}, j^{\prime}\right)_{X}-\left(c, i^{\prime}, j^{\prime}\right)_{X}-(c, b)$ for some indices $i, j, i^{\prime}, j^{\prime} \in[d]$. The existence of the vertex $\left(a, i^{\prime}, j^{\prime}\right)_{X}$ implies that $a\left[i^{\prime}\right]=a\left[j^{\prime}\right]=1$, and the existence of the vertex $\left(c, i^{\prime}, j^{\prime}\right)_{X}$ implies that $c\left[i^{\prime}\right]=c\left[j^{\prime}\right]=1$. The last edge implies that $b\left[i^{\prime}\right]=1$ or $b\left[j^{\prime}\right]=1$. Thus, at least one of $a\left[i^{\prime}\right]=b\left[i^{\prime}\right]=c\left[i^{\prime}\right]=1$ or $a\left[j^{\prime}\right]=b\left[j^{\prime}\right]=c\left[j^{\prime}\right]=1$ holds, contradicting orthogonality of $a, b, c$.

This shows that $(a, b)_{S}$ and $(c, b)_{S}$ are at distance at least 5 , completing the proof.

## 8 Acknowledgements

The author would like to thank Aviad Rubinstein for many helpful discussions, guidance, encouragement, and feedback on this writeup. The author would like to thank Mary Wootters for helpful discussions and feedback on this writeup. The author would like to thank Joshua Brakensiek for helpful discussions. The author would like to thank Thuy Duong Vuong and Arun Jambulapati for helpful discussions on hopsets, and Arun Jambulapati for the reference [CFR20]. The author would like to thank Nairen Cao, Jeremy T. Fineman, and Katina Russell for helpful discussions on their work [CFR19, CFR20]. The author would like to thank anonymous reviewers for helpful feedback.

## References

[ABMR11] Eyad Alkassar, Sascha Böhme, Kurt Mehlhorn, and Christine Rizkallah. Verification of certifying computations. In International Conference on Computer Aided Verification, pages 67-82, 2011.
[ACIM99] Donald Aingworth, Chandra Chekuri, Piotr Indyk, and Rajeev Motwani. Fast estimation of diameter and shortest paths (without matrix multiplication). SIAM J. Comput., 28(4):1167-1181, 1999.
[AKT20] Amir Abboud, Robert Krauthgamer, and Ohad Trabelsi New algorithms and lower bounds for all-pairs max-flow in undirected graphs. In Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms SODA, 2020, pages 48-61, 2020.
[AKT20b] Amir Abboud, Robert Krauthgamer, and Ohad Trabelsi Subcubic Algorithms for Gomory-Hu Tree in Unweighted Graph. arXiv preprint arXiv:2012.10281.
[AWW16] 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.
[ $\left.\mathrm{BRS}^{+} 18\right]$ 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, 2018.
[BBST07] Boaz Ben-Moshe, Binay Bhattacharya, Qiaosheng Shi, and Arie Tamir. "Efficient algorithms for center problems in cactus networks." Theoretical Computer Science, 378(3): 237-252, 2007.
[BN19] Matthias Bentert, and André Nichterlein. Parameterized complexity of diameter. In International Conference on Algorithms and Complexity, pp. 50-61. Springer, Cham, 2019.
[Bon20] Edouard Bonnet. Inapproximability of Diameter in super-linear time: Beyond the 5/3 ratio. In Symposium on Theoretical Aspects of Computer Science, STACS 2021.
$\left[\mathrm{BCH}^{+} 15\right]$ Michele Borassi, Pierluigi Crescenzi, Michel Habib, Walter A. Kosters, 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. Theor. Comput. Sci., 586:59-80, 2015.
[BE05] Ulrik Brandes and Thomas Erlebach. Network Analysis: Methodological Foundations. Springer-Verlag, 2005.
[BHM19] Karl Bringmann, Thore Husfeldt, and Måns Magnusson. Multivariate Analysis of Orthogonal Range Searching and Graph Distances. Algorithmica, 2020.
[CDHP01] Derek G. Corneil, Feodor F. Dragan, Michel Habib, and Christophe Paul. Diameter determination on restricted graph families. Discrete Applied Mathematics 113(2-3): 143-166, 2001.
[CGR16] 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.
[CFR] Nairen Cao, Jeremy T. Fineman, and Katina Russell. Private communication.
[CFR19] Nairen Cao, Jeremy T. Fineman, and Katina Russell. Efficient Construction of Directed Hopsets and Parallel Approximate Shortest Paths. In Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, pages 336349, 2020.
[CFR20] Nairen Cao, Jeremy T. Fineman, and Katina Russell. Brief Announcement: Improved Work Span Tradeoff for Single Source Reachability and Approximate Shortest Paths. In Proceedings of the 32nd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), July, 2020, pages 511-513.
[CGI ${ }^{+}$16] Marco L. Carmosino, Jiawei Gao, Russell Impagliazzo, Ivan Mihajlin, Ramamohan Paturi, and Stefan Schneider. Nondeterministic extensions of the strong exponential time hypothesis and consequences for non-reducibility. In Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, page 261-270, 2016.
[CLR $\left.{ }^{+} 14\right]$ 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.
[CDV02] Victor Chepoi, Feodor Dragan, and Yann Vaxès. Center and diameter problems in plane triangulations and quadrangulations. In Proc. SODA, pp. 346-355, 2002.
[Coh00] Edith Cohen. Polylog-time and near-linear work approximation scheme for undirected shortest paths. J. ACM, 47(1):132-166, 2000.
[CDP19] David Coudert, Guillaume Ducoffe, and Alexandru Popa. Fully polynomial FPT algorithms for some classes of bounded clique-width graphs. ACM Transactions on Algorithms (TALG), 15(3): 1-57, 2019.
[CGS15] Marek Cygan, Harold N. Gabow, and Piotr Sankowski. Algorithmic applications of baur-strassen's theorem: Shortest cycles, diameter, and matchings. J. ACM, 62(4):28:1-28:30, 2015.
[DWV ${ }^{+}$19] Mina Dalirrooyfard, Virginia Vassilevska Williams, Nikhil Vyas, Nicole Wein, Yinzhan Xu , and Yuancheng Yu. Approximation algorithms for min-distance problems. In 46 th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece, pages 46:1-46:14, 2019.
[DWVW19] Mina Dalirrooyfard, Virginia Vassilevska Williams, Nikhil Vyas, and Nicole Wein. Tight approximation algorithms for bichromatic graph diameter and related problems. In 46 th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece, pages 47:1-47:15, 2019.
[DW20] Mina Dalirrooyfard and Nicole Wein. Tight Conditional Lower Bounds for Approximating Diameter in Directed Graphs. In Symposium on Theory of Computing, STOC 2021, to appear.
[Dam16] Peter, Damaschke. Computing giant graph diameters. In International Workshop on Combinatorial Algorithms, pp. 373-384. Springer, Cham, 2016.
[Duc18] Guillaume Ducoffe, A New Application of Orthogonal Range Searching for Computing Giant Graph Diameters. In 2nd Symposium on Simplicity in Algorithms (SOSA 2019). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.
[DHV20] Guillaume Ducoffe, Michel Habib, and Laurent Viennot. Diameter computation on H-minor free graphs and graphs of bounded (distance) VC-dimension. In Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 19051922. Society for Industrial and Applied Mathematics, 2020.
[Epp00] David Eppstein, Diameter and treewidth in minor-closed graph families. Algorithmica, 27(3-4):275-291, 2000.
[FP80] Arthur M. Farley, and Andrzej Proskurowski. Computation of the center and diameter of outerplanar graphs. Discrete Applied Mathematics 2(3):185-191, 1980. Harvard
$\left[\mathrm{GKM}^{+} 18\right]$ Pawl Gawrychowski, Haim Kaplan, Shay Mozes, Micha Sharir, and Oren Weimann. Voronoi diagrams on planar graphs, and computing the diameter in deterministic $\tilde{O}\left(n^{5 / 3}\right)$ time. In Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 495-514. Society for Industrial and Applied Mathematics, 2018.
[Gol05] Oded Goldreich, On promise problems: A survey. Theoretical computer science. Springer, Berlin, Heidelberg, 254-290, 2006.
[IPZ01] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? J. Comput. Syst. Sci., 63(4):512-530, 2001.
[JLS19] Arun Jambulapati, Yang Liu, and Aaron Sidford. Parallel Reachability in Almost Linear Work and Square Root Depth. 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS). IEEE, 2019.
[Kun18] Marvin Künnemann, On Nondeterministic Derandomization of Freivalds' Algorithm: Consequences, Avenues and Algorithmic Progress. In 26th Annual European Symposium on Algorithms (ESA 2018), 2018.
[Le14] François Le Gall. Powers of tensors and fast matrix multiplication. In International Symposium on Symbolic and Algebraic Computation, ISSAC '14, Kobe, Japan, July 23-25, 2014, pages 296-303, 2014.
[LWCW16] Ting-Chun Lin, Mei-Jin Wu, Wei-Jie Chen, and Bang-Ye Wu. Computing the diameters of huge social networks. In 2016 International Computer Symposium (ICS), pages 6-11. IEEE, 2016.
[MMNS11] Ross M. McConnell, Kurt Mehlhorn, Stefan Näher, and Pascal Schweitzer. Certifying algorithms. Computer Science Review, 5(2):119-161, 2011.
[PRT12] David Peleg, Liam Roditty, and Elad Tal. Distributed algorithms for network diameter and girth. In Automata, Languages, and Programming - 39th International Colloquium, ICALP 2012, Warwick, UK, July 9-13, 2012, Proceedings, Part II, pages 660-672, 2012.
[Pet04] Seth Pettie. A new approach to all-pairs shortest paths on real-weighted graphs. Theor. Comput. Sci., 312(1):47-74, 2004.
[PR05] Seth Pettie and Vijaya Ramachandran. A shortest path algorithm for real-weighted undirected graphs. SIAM J. Comput., 34(6):1398-1431, 2005.
[RW13] Liam Roditty and Virginia Vassilevska Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. In Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013, pages 515-524, 2013.
[RW19] Aviad Rubinstein and Virginia Vassilevska Williams. SETH vs approximation. SIGACT News, 50(4):57-76, 2019.
[Sto10] Andrew Stothers. On the complexity of matrix multiplication. Ph.D. Thesis, U. Edinburgh, 2010.
[WS98] D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. Nature, 393:440-442, 1998.
[Wil05] Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. Theor. Comput. Sci., 348(2-3):357-365, 2005.
[Wil12] Virginia Vassilevska Williams. Multiplying matrices faster than coppersmithwinograd. In Proceedings of the forty-fourth annual ACM symposium on Theory of computing, pages 887-898. ACM, 2012.
[Wil14] Ryan 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, pages 664-673, 2014.
[Wil16] Richard Ryan Williams. Strong ETH breaks with merlin and arthur: Short noninteractive proofs of batch evaluation. In 31st Conference on Computational Complexity, CCC 2016, May 29 to June 1, 2016, Tokyo, Japan, pages 2:1—2:17, 2016.
[Wil18] Virginia Vassilevska Williams. On some fine-grained questions in algorithms and complexity. In Proceedings of the ICM, volume 3. World Scientific, 2018.

## A Hopset results

## A. 1 Undirected hopset

Here, we prove Lemma 3.7. Let $M$ be a positive integer. In a graph $G$, let $\tilde{N}_{M}(v)$ denote the set of vertices visited when Dijkstra is run from a vertex $v$ until $2 M$ edges are visited, and let $\tilde{d}_{M}(v)$ denote the distance from $v$ to the furthest vertex in $\tilde{N}_{M}(v)$. Let $\tilde{E}_{M}(v)$ denote the set of vertices incident to $\tilde{N}_{M}(v)$. By definition, we have $\left|\tilde{E}_{M}(v)\right| \geq M$ (each edge can be visited at most twice).

Lemma A.1. Let $G=(V, E)$ be an undirected weighted graph, and $\underset{\sim}{M}$ be an integer. For each $v$, in time $O(M \log M)$ we can compute $\tilde{N}_{M}(v)$ and $d(v, u)$ for all $u \in \tilde{N}_{M}(v)$.

Proof. Run Dijkstra from $v$ until $2 M$ edges are visited. Then every vertex in $\tilde{N}_{M}(v)$ is visited by definition of $\tilde{N}_{M}(v)$, so we know $d(v, u)$ for all $u \in \tilde{N}_{M}(v)$.

Proof of Lemma 3.7. Any $(\beta, \varepsilon)$-additive-hopset is also a $\left(\beta, \varepsilon^{\prime}\right)$-additive-hopset for $\varepsilon^{\prime}>\varepsilon$, so we may assume without loss of generality that $\varepsilon \leq 1 / 100$. Let $k=\lceil 1 / \delta\rceil$ and assume without loss of generality that $n$ and $m$ are sufficiently large in terms of $k$ (all guarantees suppress dependencies
on $\delta$ ). For $i=1, \ldots, k$, let $M_{i}=m^{(k+1-i) / k}$. Let $D$ denote the diameter of the graph. We use the following algorithm.

1. For $i=1, \ldots, k$, let $S_{i}$ be the vertices incident to $4 m^{i / k} \log m$ uniformly random edges.
2. For $i=1, \ldots, k$, for each $v \in S_{i}$, run the algorithm in Lemma A. 1 for $M=M_{i}$. For each $u \in \tilde{N}_{M_{i}}(v)$, add an edge of weight $d_{G}(v, u)$ from $v$ to $u$ to $E^{\prime}$.

Note that, for $i=1$, we simply end up running Dijkstra from each vertex.
Runtime. For $i=1, \ldots, k$, each Dijkstra takes $O\left(M_{i} \log M_{i}\right)$ time, so the running time for that step is $O\left(m^{i / k} \log m \cdot m^{(k+1-i) / k} \log m\right)=\tilde{O}\left(m^{1+1 / k}\right)$. Thus the total running time is $\tilde{O}\left(m^{1+1 / k}\right)$.

Correctness. First, by Lemma A.1, we know that for each $i=1, \ldots, k$, each $v \in S_{i}$, and each $u \in \tilde{N}_{M_{i}}(v)$, the distance $d(v, u)$ is accurately computed, so the edge added from $v$ to $u$ does not decrease the shortest path from $v$ to $u$ and hence does not decrease any shortest path. Thus, we have the first guarantee that adding $E^{\prime}$ preserves all shortest paths.

Now we show the second guarantee. By same reasoning as in Lemma 5.4, for all $i=1, \ldots, k$, with probability $1-1 / n$, because $4 m^{i / k} \log m \cdot M_{i+1}>2 m \log n$, the edges used to generate $S_{i}$ intersect each of the $n$ sets $\tilde{E}_{M_{i+1}}(v)$ for all vertices $v$. Hence, $S_{i}$ intersects each $\tilde{N}_{M_{i+1}}(v)$. By the union bound, with probability $1-k / n>0$, for all $i=1, \ldots, k$ and all vertices $v$, the set $S_{i}$ intersects each set $\tilde{N}_{M_{i+1}}(v)$. Fix the choice of randomness such that this holds. We now show that $E^{\prime}$ is a $\left(O_{\delta, \varepsilon}(1), 6 \varepsilon\right)$ hopset, which, after reparameterizing $\varepsilon$, shows that we can obtain a $\left(O_{\delta, \varepsilon}(1), \varepsilon\right)$-additive-hopset.

Fix two vertices $v$ and $v^{\prime}$ and a shortest path $v=v_{0}, \ldots, v_{L}=v^{\prime}$ between them. We show how to construct a path in $G+E^{\prime}$ from $v$ to $v^{\prime}$ of length at most $d\left(v, v^{\prime}\right)+\varepsilon D$ using at most $O_{\delta, \varepsilon}(1)$ vertices. We prove the following claim.

Claim A.2. For every $\ell \in\{0, \ldots, L\}$, there exists $\ell^{\prime}>\ell$ and an at-most $k+1$ edge path from $v_{\ell}$ to $v_{\ell^{\prime}}$ such that either (1) $d\left(v_{\ell}, v_{\ell^{\prime}}\right)>\varepsilon^{k} D$ and the path has length at most $(1+4 \varepsilon) d\left(v_{\ell}, v_{\ell^{\prime}}\right)$ or (2) $\ell^{\prime}=L$ and the path has length at most $d\left(v_{\ell}, v_{\ell^{\prime}}\right)+2 \varepsilon D$.

Proof. Fix $\ell$. For each vertex $u$ and $i=2, \ldots, k$, let $f_{i}(u)$ be an arbitrary vertex in $S_{i-1} \cap \tilde{N}_{M_{i}}(v)$. Such a vertex exists by construction of $S_{i-1}$. Furthermore, if $u \in S_{i}$, by the definition of the algorithm, there always exists an edge from $u$ to $f_{i}(u)$. Say a vertex $u \in S_{i}$ is $i$-nonexpanding if $\tilde{d}_{M_{i}}(u) \leq \varepsilon^{i-1} D$ and $i$-expanding otherwise. Note that if $u$ is $i$-nonexpanding, then every vertex $u^{\prime} \in \tilde{N}_{M_{i}}(u)$ satisfies $d\left(u, u^{\prime}\right) \leq \varepsilon^{i-1} D$. On the other hand, if $u$ is $i$-expanding, any vertex $u^{\prime} \notin$ $\tilde{N}_{M_{i}}(u)$ satisfies $d\left(u, u^{\prime}\right)>\varepsilon^{i-1} D$.

Let $u_{k}=v_{\ell}$. For $i=k, \ldots, 2$, let $u_{i-1}=f_{i}\left(u_{i}\right)$. Then $u_{i} \in S_{i}$ for $i=1, \ldots, k$ by definition of $f_{i}$, and there always exists an edge from $u_{i}$ to $u_{i-1}$ for $i=2, \ldots, k$. Clearly $u_{1}$ is 1 -expanding, and the length of all shortest paths is bounded above by the diameter $D$, so $\tilde{d}_{M_{1}}(u) \leq D$. Let $j$ be the largest index such that $u_{j}$ is $j$-expanding. Then $u_{i}$ is $i$-nonexpanding for $j<i \leq k$, so by the triangle inequality

$$
\begin{equation*}
d\left(v_{\ell}, u_{j}\right) \leq \sum_{i=j+1}^{k} d\left(u_{i}, u_{i-1}\right)=\sum_{i=j+1}^{k} d\left(u_{i}, f\left(u_{i}\right)\right) \leq \sum_{i=j+1}^{k} \varepsilon^{i-1} D<2 \varepsilon^{j} D \tag{1}
\end{equation*}
$$

Note this inequality even holds for $j=k$.
First suppose $v_{L} \in \tilde{N}_{M_{j}}\left(u_{j}\right)$. Then the path $u_{k}, \ldots, u_{j}, v_{L}$ is a path of at most $k$ edges from $v_{\ell}=u_{k}$ to $v_{L}$ and has length at most

$$
d\left(v_{\ell}, u_{j}\right)+d\left(u_{j}, v_{L}\right) \leq 2 d\left(v_{\ell}, u_{j}\right)+d\left(v_{\ell}, v_{L}\right)<2 \varepsilon D+d\left(v_{\ell}, v_{L}\right) .
$$

In the first inequality, we used the triangle inequality, and in the last inequality we used (1) and that $j \geq 1$.

Now suppose $v_{L} \notin \tilde{N}_{M_{j}}\left(u_{j}\right)$. Let $\ell^{\prime}$ denote the largest integer such that $v_{\ell^{\prime}-1} \in \tilde{N}_{M_{j}}\left(u_{j}\right)$. Since we assume $v_{L} \notin \tilde{N}_{M_{j}}\left(u_{j}\right)$, we have $\ell^{\prime} \leq L$. By (1), we have $d\left(v_{\ell}, u_{j}\right)=d\left(u_{k}, u_{j}\right)<\varepsilon^{j-1} D$, so $v_{\ell} \in$ $\tilde{N}_{M_{j}}\left(u_{j}\right)$ and thus $\ell^{\prime}>\ell$. Since $u_{j}$ is $j$-expanding, and $v_{\ell^{\prime}+1} \notin \tilde{N}_{M_{j}}\left(u_{j}\right)$, we have $d\left(u_{j}, v_{\ell^{\prime}}\right)>\varepsilon^{j-1} D$. Thus, by the triangle inequality and (1), we have

$$
\begin{equation*}
d\left(v_{\ell}, v_{\ell^{\prime}}\right)>d\left(u_{j}, v_{\ell^{\prime}}\right)-d\left(u_{j}, v_{\ell}\right)>\left(\varepsilon^{j-1}-2 \varepsilon^{j}\right) D . \tag{2}
\end{equation*}
$$

and in particular $d\left(v_{\ell}, v_{\ell^{\prime}}\right)>2 \varepsilon^{k} D$. On the other hand, the path $u_{k}, u_{k-1}, \ldots, u_{j}, v_{\ell^{\prime}-1}, v_{\ell^{\prime}}$ is an at most $k+1$ edge path from $v_{\ell}$ to $v_{\ell^{\prime}}$, and the length is at most

$$
\begin{array}{rlr}
d\left(v_{\ell}, u_{j}\right) & +d\left(u_{j}, v_{\ell^{\prime}-1}\right)+d\left(v_{\ell^{\prime}-1}, v_{\ell^{\prime}}\right) & \\
& \leq 2 d\left(u_{k}, u_{j}\right)+d\left(v_{\ell}, v_{\ell^{\prime}-1}\right)+d\left(v_{\ell^{\prime}-1}, v_{\ell^{\prime}}\right) & \\
& =2 d\left(u_{k}, u_{j}\right)+d\left(v_{\ell}, v_{\ell^{\prime}}\right) & \\
& <2 \varepsilon^{j} D+d\left(v_{\ell}, v_{\ell^{\prime}}\right) & \\
& <\left(1+\frac{2 \varepsilon^{j}}{\varepsilon^{j-1}-2 \varepsilon^{j}}\right) d\left(v_{\ell}, v_{\ell^{\prime}}\right) & \\
& <(1+4 \varepsilon) d\left(v_{\ell}, v_{\ell^{\prime}}\right) & \\
& \text { (by (by }(1))) \\
& (\varepsilon \text { sufficiengle-ineq.) }) \\
&
\end{array}
$$

as desired.
Now we finish the proof. By repeatedly applying Claim A.2, we obtain indices $0=\ell_{0}, \ell_{1}, \ldots, \ell_{p}=$ $\ell$, such that for $0 \leq p^{\prime}<p-1$, there exists an at-most- $k+1$ edge path from $v_{\ell_{p^{\prime}}}$ to $v_{\ell_{p^{\prime}+1}}$ of length at most $(1+\varepsilon) d\left(v_{\ell_{p^{\prime}}}, v_{\ell_{p^{\prime}+1}}\right)$ and an at-most- $k+1$ edge path from $v_{\ell_{p-1}}$ to $v_{\ell_{p}}$ of length at most $d\left(v_{\ell_{p-1}}, v_{\ell_{p}}\right)+\varepsilon D$. Furthermore, we have $p<\varepsilon^{-k}$ because $d\left(v_{\ell_{p^{\prime}}}, v_{\ell_{p^{\prime}+1}}\right) \geq \varepsilon^{k} D$ for all $p^{\prime}<p-1$. In total, this gives a path containing $(k+1) \cdot 2 \varepsilon^{-k} \leq O_{\delta, \varepsilon}(1)$ vertices and of total length at most

$$
\begin{aligned}
\sum_{p^{\prime}=0}^{p-2}(1+4 \varepsilon) d\left(v_{\ell_{p^{\prime}}}, v_{\ell_{p^{\prime}+1}}\right)+\left(d\left(v_{\ell_{p-1}}, v_{\ell_{p}}\right)+2 \varepsilon D\right) & =d\left(v, v^{\prime}\right)+\left(\sum_{p^{\prime}=0}^{p-2} 4 \varepsilon d\left(v_{\ell_{p^{\prime}}}, v_{\ell_{p^{\prime}+1}}\right)\right)+2 \varepsilon D \\
& <d\left(v, v^{\prime}\right)+6 \varepsilon D
\end{aligned}
$$

Thus, $E^{\prime}$ is indeed a $\left(O_{\delta, \varepsilon}(1), 6 \varepsilon\right)$-additive hopset, as desired.

## A. 2 Directed hopset

Lemma 3.6 is implicit in [CFR20] and is confirmed by the authors [CFR]. Here we explain how Lemma 3.6 is implicit in [CFR20]. In [CFR19], building off of [JLS19], a parallel algorithm is given for single-source reachability and single-source approximate shortest paths on directed weighted graphs with $\tilde{O}(m)$ work and $n^{1 / 2+o(1)}$ span. To achieve single-source approximate reachability, they find a set of shortcuts in $\tilde{O}(m)$ time that reduces the diameter to $n^{1 / 2+o(1)}$. To achieve singlesource approximate shortest paths, they use a construction of a $\left(n^{1 / 2+o(1)}, \varepsilon\right)$ hopset on directed weighted graphs in $\tilde{O}_{\varepsilon}(m)$ time. The hopsets are in fact a generalization of the shortcuts.

In [CFR20, Theorem 1.1], the authors generalize this algorithm to obtain a tradeoff between work and span, solving single-source reachability and approximate single-source shortest paths in a directed weighted graph in $\tilde{O}\left(m \rho^{4}\right)$ (in fact, $\tilde{O}\left(m \rho^{2}+n \rho^{4}\right)$ ) work and $n^{1 / 2+o(1)} / \rho$ span. When $\rho=1$, these algorithms are exactly the same as the one in [CFR19]. For larger $\rho$, they follow
the same approach. For single-source reachability, they construct shortcuts to reduce the diameter to $n^{1 / 2+o(1)} / \rho$ in time $\tilde{O}\left(m \rho^{2}\right)$. For single-source approximate shortest path, they construct $\left(n^{1 / 2+o(1)} / \rho, \varepsilon\right)$ hopset on directed weighted graphs in time $\tilde{O}_{\varepsilon}\left(m \rho^{4}\right)$. However, as the paper is only an extended abstract, they only state the shortcut construction [CFR20, Theorem 4.1], and not the full hopset construction. The use of such a hopset construction was confirmed by the authors [CFR].

To obtain the addition guarantee that, with probability 1, the edge additions preserve all shortest paths, we note that the algorithms in [CFR19, CFR20] only adds an edge from $u$ to $v$ of weight $d_{G}(u, v)$, so the edge does not shorten the shortest path from $u$ to $v$ and thus any other vertex. Additionally, since edges are only added, the length of shortest paths never increase.

## B Fine grained reductions

## B. 1 Applying $\left[\mathrm{CGI}^{+} 16\right]$ for promise problems.

Lemma 3.3 follows from the following property, in the same way that [CGI ${ }^{+}$16, Theorem 5.1] follows from [CGI ${ }^{+} 16$, Lemma 3.5]. This property states that fine grained reductions translate ( $\mathrm{N} \cap \mathrm{coN})^{+}$TIME savings for promise problems to ( $\mathrm{N} \cap \operatorname{coN}$ )TIME.

Lemma B. 1 (Analogue of Lemma 3.5 of $\left[\mathrm{CGI}^{+} 16\right]$ for promise problems). Let $\left(\Pi^{\prime}, T^{\prime}\right) \leq_{F G R}$ $(\Pi, T)$, where $\Pi, \Pi^{\prime}$ are promise problems, and suppose $\Pi \in(\mathrm{N} \cap \operatorname{coN})^{+} \mathrm{TIME}\left[T(n)^{1-\varepsilon}\right]$ for some $\varepsilon>0$. Then there exists a $\delta>0$ such that $\Pi^{\prime} \in(\mathrm{N} \cap \operatorname{coN}) \operatorname{TIME}\left[T^{\prime}(n)^{1-\delta}\right]$.

Proof. We proceed as in the proof in [CGI $\left.{ }^{+} 16\right]$. We take $\mathcal{M}^{\Pi}$ to be the deterministic oracle that achieves a fine-grained reduction from $\Pi=\left(\Pi_{Y E S}, \Pi_{N O}\right)$ to $\Pi^{\prime}=\left(\Pi_{Y E S}^{\prime}, \Pi_{N O}^{\prime}\right)$, and construct a nondeterministic machine $\mathcal{M}^{\prime}$ for deciding $\Pi^{\prime}$ to show that $\Pi^{\prime} \in \operatorname{NTIME}\left[T^{\prime}(n)^{1-\delta}\right]$. Let $\mathcal{A}_{N}, \mathcal{A}_{c o N}$ be the algorithms for $\Pi$ in time $T$ given by Definition 3.2. As in [CGI+16], $\mathcal{M}^{\prime}$ guesses a table of queries to $\Pi$ along with their answers. $\mathcal{M}^{\prime}$ then uses the nondeterministic algorithm $\mathcal{A}_{N}$ for $\Pi$ to verify the queries with "YES" answer, and the nondeterministic algorithm $\mathcal{A}_{c o N}$ for $\Pi$ to verify the queries with "NO" answer. We then simulate $\mathcal{M}$, looking up the answers to queries in the table. If a query does not appear in the table, we output "NO".

We now check the correctness of the nondeterministic Turing machine $\mathcal{M}^{\prime}$. If the input to $\Pi^{\prime}$ is in $\Pi_{N O}^{\prime}$, then all nondeterministic choices result in an output of "NO" for the same reason as in [CGI ${ }^{+}$16]: unless the query table is incorrectly labels a query in $\Pi_{Y E S}$ with "NO" or labels a $\Pi_{N O}$ query with "YES", in which case the nondeterministic verifiers for $\Pi$ cause us to output "NO", $\mathcal{M}^{\prime}$ will always output "NO" on an input in $\Pi_{N O}^{\prime}$ by the definition of $\mathcal{M}^{\prime}$.

We now check that for all inputs to $\Pi^{\prime}$ is in $\Pi_{Y E S}^{\prime}$, there exist nondeterministic choices such that $\mathcal{M}^{\prime}$ outputs "YES". We put all queries made by $\mathcal{M}^{\prime}$ in the table, with their corresponding answers. For all queries $q$ not in the promise $\Pi_{Y E S} \cup \Pi_{N O}$, if there exist nondeterministic choices for $\mathcal{A}_{N}$ to output "YES", then give query $q$ answer "YES", and if there exist nondeterministic choices for $\mathcal{A}_{c o N}$ to output "NO", then give query $q$ answer "NO". By property 3 of the definition of $(\mathrm{N} \cap \operatorname{coN})^{+}$TIME, one of these two possibilities must exist for all queries $q$ not in the promise. Now, there exist nondeterministic choices such that $\mathcal{M}^{\prime}$ never outputs "NO" before simulating $\mathcal{M}$, because all queries in the promise $\Pi_{Y E S} \cup \Pi_{N O}$ can be correctly verified by properties 1 and 2 of $(\mathrm{N} \cap \operatorname{coN})^{+}$TIME, and, by construction of the table, all queries not in the promise can also be verified. Thus the simulation of $\mathcal{M}$ succeeds and we correctly output "YES".

As in $\left[\mathrm{CGI}^{+} 16\right]$, the simulation takes time at most $\left(T^{\prime}\right)^{1-\delta}$, and guessing and verifying the query table takes $O\left(\left(T^{\prime}\right)^{1-\delta}\right)$ time. Therefore $\Pi^{\prime} \in \operatorname{NTIME}\left[T^{\prime}(n)^{1-\delta}\right]$.

Similarly, we can design a nondeterministic machine for inputs in $\Pi_{N O}^{\prime}$ to conclude that $\Pi^{\prime} \in$ $(\mathrm{N} \cap \operatorname{coN}) \mathrm{TIME}\left[T^{\prime}(n)^{1-\delta}\right]$.

Proof of Lemma 3.4. The proof is the same as in [CGI $\left.{ }^{+} 16\right]$, except that when we nondeterministically simulate the oracle queries to the promise problem $\Pi$, we simulate as in Lemma B. 1 rather than as in [CGI ${ }^{+} 16$, Lemma C.5].

Suppose for contradiction that $\left(\Pi, T^{1+\delta}\right)$ is SETH-hard so that there is a randomized reduction from CNFSAF with time $2^{n}$ to $\Pi$ with time $T^{1+\delta}$, and let $\mathcal{M}^{\Pi}$ be the randomized oracle machine. Then there exists some $\varepsilon>0$ such that $\operatorname{TIME}[\mathcal{M}] \leq 2^{n(1-\varepsilon)}$ and such that the query lengths satisfy $\sum_{q \in \tilde{Q}(\mathcal{M}, x)} T(|q|) \leq 2^{n(1-\varepsilon)}$ (here we switched $\varepsilon$ and $\delta$ from the definition in Section 3).

Let $m<n^{k}$ be the length in bits of a description of a k-SAT formula on $n$ inputs. By repeating $\mathcal{M}^{\Pi} O(m)$ times and taking the majority answer, we can make the error probability less than $2^{-m}$. Since there are at most $2^{m}$ possible inputs, there is one random tape such that $\mathcal{M}$ has no errors when all the oracle queries to $\Pi$ are correct. Since $\mathcal{M}$ runs in total time $2^{(1-\varepsilon) n}$, this tape will have length at most $O\left(m 2^{(1-\varepsilon) n}\right)$. Then simulate oracle queries nondeterministically, using this nondeterministically chosen random tape, using the same simulation as Lemma B.1. This simulation has time complexity $O\left(m 2^{(1-\varepsilon) n}\right)$, giving a nondeterministic circuit with total size $O\left(m 2^{(1-\varepsilon) n}\right)$.

## B. 2 Reducing approximate Diameter to the promise problem

Let $\rho \in[1, \infty]$, possibly as some function of $m$ and $n$. For $\alpha \geq 1$, let $f_{\alpha, \rho}$ be the function problem $f$ of giving an $\alpha$-approximation of Diameter on graphs satisfying $W_{\max } / W_{\min } \leq \rho$ (all graphs if $\rho=\infty)$. Let $\Pi_{\alpha, \rho}$ be the gap-problem $\alpha / 1$-Diameter on graphs satisfying $W_{\max } / W_{\min } \leq \rho$.

Lemma B.2. For all $\rho \in[1, \infty]$, all constant $\alpha \geq 1$, all constant $\beta>0$, and all time complexities $T$, there is a fine grained reduction from $\left(f_{\alpha+\beta, \rho}, T\right)$ to $\left(\Pi_{\alpha, \rho}, T\right)$.

Proof. Suppose we are given an oracle to $\alpha / 1$-Diameter on graphs with $W_{\max } / W_{\min } \leq \rho$. If we had oracles to $\alpha D / D$-Diameter for all $D$, we can compute $f$ on graphs with $W_{\max } / W_{\min } \leq \rho$ by simply running a binary search on the same graph to accuracy $\beta$, querying $\alpha D / D$-Diameter for different values of $D$, to obtain an interval $[x,(\alpha+\beta) x]$ containing the correct answer. However, we can replace all $\alpha D / D$-Diameter queries with $\alpha / 1$-Diameter queries by simply re-weighting the graph. This preserves the weight condition $W_{\max } / W_{\min } \leq \rho$. This gives a Turing reduction $\mathcal{M}^{\Pi_{\alpha, \rho}}$ running in time $O(\log (1 / \beta))=O_{\beta}(1)$. Let $\tilde{Q}(\mathcal{M}, x)$ denote the set of queries made by $\mathcal{M}$ to the oracle on an input $x$ of length $n$. There are $O_{\beta}(1)$ queries, so, for all $\varepsilon>0$ the query lengths obey the following time bound.

$$
\sum_{q \in \tilde{Q}(\mathcal{M}, x)}(T(|q|))^{1-\varepsilon} \leq O_{\beta}\left(T(n)^{1-\varepsilon}\right)
$$

Hence, this is a fine grained reduction from $\left(f_{\alpha+\beta, \rho}, T\right)$ to $\left(\Pi_{\alpha, \rho}, T\right)$.
We now can prove Lemma 3.5.
Proof of Lemma 3.5. First we consider deterministic reductions. Lemma B. 2 proves that $\left(f_{\alpha+\beta, \rho}, T\right) \leq_{F G R}$ $\left(\Pi_{\alpha, \rho}, T\right)$. By composition of fine-grained reductions ${ }^{7}$, if (CNFSAT, $\left.2^{n}\right) \leq_{F G R}\left(f_{\alpha+\beta, \rho}, T\right)$, then we also have $\left(\operatorname{CNFSAT}, 2^{n}\right) \leq_{F G R}\left(\Pi_{\alpha, \rho}, T\right)$. Hence, taking the contrapositive, we have that if

[^6](CNFSAT, $\left.2^{n}\right) \not Z_{F G R}\left(\Pi_{\alpha, \rho}, T\right)$, then (CNFSAT, $\left.2^{n}\right) \not Z_{F G R}\left(f_{\alpha+\beta, \rho}, T\right)$, which is exactly what Lemma 3.5 asks to prove.

To show that non-SETH-hardness of the promise problem implies non-SETH-hardness of the approximation problem for randomized algorithms, the proof is the same, but we need to show that randomized fine-grained reductions are closed under composition. We prove that these reductions are closed under composition in the scenario that we need.

Lemma B.3. Let $\left(A, T_{A}\right) \leq_{F G R, r}\left(B, T_{B}\right)$ and $\left(B, T_{B}\right) \leq_{F G R, r}\left(C, T_{C}\right)$ where $B$ is an approximate function problem. Then $\left(A, T_{A}\right) \leq_{F G R, r}\left(C, T_{C}\right)$.

Proof. The proof is the same as [CGI ${ }^{+}$16, Lemma C.7], except that we simulate the oracle machines multiple times to amplify the failure probabilities. Specifically, let $\mathcal{M}_{A B}^{B}$ be the (probabilistic) machine that achieves a fine-grained reduction from $\left(A, T_{A}\right)$ to $\left(B, T_{B}\right)$, and let $\mathcal{M}_{B C}^{C}$ be the machine that achieves a fine-grained reduction from $\left(B, T_{B}\right)$ to $\left(C, T_{C}\right)$. We construct a machine $\mathcal{M}_{A C}^{C}$ that achieves a fine-grained reduction from $\left(A, T_{A}\right)$ to $\left(C, T_{C}\right)$.

Let $\mathcal{M}_{A C}^{C}$ on an input to $A$ of length $n_{A}$ simulate $\mathcal{M}_{A B}^{B} 3$ times, and for each oracle call to $B, \mathcal{M}_{A C}^{C}$ simulates $\mathcal{M}_{B C}^{C} O\left(\log T_{A}\left(n_{A}\right)\right)$ times and takes the median output. Problem $B$ has a correct output if it is in some range [OPT, $\alpha \cdot O P T$ ], and by definition of $\mathcal{M}_{B C}^{C}$ at least $2 / 3$ of the outputs to $\mathcal{M}_{B C}^{C}$ lie in that range $[O P T, \alpha \cdot O P T]$ in expectation. Thus, with probability $\frac{1}{100} T_{A}\left(n_{A}\right)^{-2}$ (if the number of repeats of $\mathcal{M}_{B C}^{C}$ is enough), the median output of $\mathcal{M}_{B C}^{C}$ is in the correct range $[O P T, \alpha \cdot O P T]$. Since there are at most $T_{A}\left(n_{A}\right)$ oracle calls to $B$, with probability at least $1-\frac{3}{100} T_{A}\left(n_{A}\right)^{-1}>0.97$, all of the oracle calls to $B$ have a correct output. Furthermore, given that the oracle calls to $B$ have a correct output, one call to $\mathcal{M}_{A B}^{B}$ gives the incorrect output to $A$ with probability at most $\frac{1}{3}$, so the probability that at least two of the three calls to $\mathcal{M}_{A B}^{B}$ give the incorrect output is at most $(1 / 3)^{3}+3 \cdot(2 / 3) \cdot(1 / 3)^{2}<0.26$. Thus, the probability that $\mathcal{M}_{A C}^{C}$ gives an incorrect output is at most $0.26+0.03<\frac{1}{3}$, as desired.

To prove that $\mathcal{M}_{A C}^{C}$ satisfies the required time and query length bounds so that it indeed achieves a fine grained reduction from $\left(A, T_{A}\right)$ to $\left(C, T_{C}\right)$, we follow the same proof as in [CGI $\left.{ }^{+} 16\right]$. The only difference is that the runtime of $\mathcal{M}_{A C}^{C}$ and the number of queries is multiplied by $O\left(\log T_{A}\left(n_{A}\right)\right)$, which has a negligible effect on the overall analysis.

Applying Lemma B. 3 completes the proof of Lemma 3.5 for randomized reductions.


[^0]:    *Department of Computer Science, Stanford University. Email: rayyli@cs.stanford.edu. Research supported by the National Science Foundation (NSF) under Grant No. DGE - 1656518 and by Jacob Fox's Packard Fellowship.

[^1]:    1 "Almost" means the algorithm loses an additive factor that depends on the edge weights.
    ${ }^{2}$ SETH states that, for every $\varepsilon>0$, there exists a positive integer $k$ such that $k$-SAT needs $\Omega\left(2^{(1-\varepsilon) n}\right)$ time.

[^2]:    ${ }^{3}$ NSETH states: for every $\varepsilon>0$, there exists a $k$ such that $k$-TAUT is not in NTIME $\left[2^{n(1-\varepsilon)}\right]$, where $k$-TAUT is the language of all $k$-DNF which are tautologies.

[^3]:    ${ }^{4}$ NUNSETH states that for all $\varepsilon>0$, there exists $k$ such that there are no nondeterministic circuit families of size $O\left(2^{n(1-\varepsilon)}\right)$ recognizing the language $k$-TAUT.

[^4]:    ${ }^{5}$ For every $\delta$, there is only one value of $\varepsilon$ for which the characterization does not apply, namely $1 / \varepsilon=\lfloor 1 / \delta\rfloor+1$.

[^5]:    ${ }^{6}$ If $\rho=\infty$, this is all graphs

[^6]:    ${ }^{7}$ Here, we need [CGI ${ }^{+} 16$, Lemma 3.7] to hold for promise problems, but the proof is the same so we omit it.

