

Simpler and Higher Lower Bounds for Shortcut Sets

Virginia Vassilevska Williams*

Yinzhan Xu[†]

Zixuan Xu[‡]

Abstract

We study the well-known shortcut set problem: how much can one decrease the diameter of a directed graph by adding a small set of shortcuts from the transitive closure of the graph.

We provide a variety of lower bounds. First, we vastly simplify the recent construction of Bodwin and Hoppenworth [FOCS 2023] which showed an $\tilde{\Omega}(n^{1/4})$ lower bound for the diameter of a directed unweighted n -node graph after adding $O(n)$ shortcut edges. We highlight that our simplification completely removes the use of the convex sets by Bárány and Larman [Math. Ann. 1998] used in all previous lower bound constructions. Our simplification also removes the need for randomness and further removes some log factors. It allows us to generalize the construction to higher dimensions, which in turn can be used to show the following results:

- There is an $\Omega(n^{1/5})$ lower bound for the diameter of the graph after adding $O(m)$ shortcuts, where m denotes the number of edges in the input graph.
- For all $\varepsilon > 0$, there exists a $\delta > 0$ such that there are n -vertex $O(n)$ -edge graphs G where adding any shortcut set of size $O(n^{2-\varepsilon})$ keeps the diameter of G at $\Omega(n^\delta)$. This improves the sparsity of the constructed graph compared to a known similar result by Hesse [SODA 2003].
- For any integer $d \geq 2$, there exists a graph $G = (V, E)$ on n vertices and $S \subseteq V$ with $|S| = \tilde{\Theta}(n^{3/(d+3)})$, such that when adding $O(n)$ or $O(m)$ shortcuts, the sourcewise diameter (the largest distance from some vertex in S to some reachable vertex in the graph) is $\tilde{\Omega}(|S|^{1/3})$. This initiates the study of sourcewise diameter in the setting of the shortcut set problem; previously, the study of the sourcewise variant is popular in a wide variety of related problems such as spanners and distance preservers. Complementing this lower bound result, we also provide an upper bound: we show that, we can reduce the sourcewise diameter to $\tilde{O}(\sqrt{|S|})$ by adding $O(n)$ shortcut edges.

1 Introduction

In a variety of computational models, the complexity of shortest paths problems crucially depends on the *diameter* of the underlying graph: the largest shortest paths distance between reachable pairs. In dynamic graph data structures, for instance, the update time and query time can both depend on the diameter. In parallel algorithms, e.g. in the PRAM model, the span of many graph problems is a function of the diameter.

With these applications in mind, Thorup [Tho92] introduced the notion of a *shortcut set* of a graph $G = (V, E)$: a set of edges H from the *transitive closure* of G such that $(V, E \cup H)$ has small diameter.

By definition, shortcut sets preserve all reachability information in G , yet can drastically reduce the diameter, thus significantly lowering the complexity of graph computations in many settings (e.g. [ASZ20, BGW20, CFR20, Fin20, HKN14, JLS19, KS21, KS97]).

Two main settings are of interest in the shortcut set literature: How much can one drop the diameter of an n -node, m -edge directed unweighted graph by adding a shortcut set of (1) $O(n)$ edges? (2) $O(m)$ edges?

Both settings preserve the sparsity of the graph, within a constant factor. Until recently, the only known shortcut set construction in either setting was as follows (see [UY91]). Given any n -node graph and any D , sample a set S of $O(n \log n/D)$ vertices uniformly at random, and add all $\tilde{O}(n^2/D^2)$ edges¹ from the transitive closure between pairs of vertices in S . It is easy to show that this achieves a diameter of $O(D)$. Setting D to $\sqrt{n}$ shows that $\tilde{O}(n)$ shortcuts are sufficient to achieve diameter $O(\sqrt{n})$, and setting D to $n/\sqrt{m}$ shows that $\tilde{O}(m)$ shortcuts are sufficient to drop the diameter to $O(n/\sqrt{m})$. The polylogs in the $\tilde{O}$ were removed by [BRR10].

*Massachusetts Institute of Technology. virgi@mit.edu. Supported by NSF Grants CCF-2129139 and CCF-2330048 and BSF Grant 2020356.

[†]Massachusetts Institute of Technology. xyzhan@mit.edu. Partially supported by NSF Grants CCF-2129139 and CCF-2330048 and BSF Grant 2020356.

[‡]Massachusetts Institute of Technology. zixuanxu@mit.edu.

¹ $\tilde{O}(t(n))$ is $O(t(n)\text{poly log } n)$ and $\tilde{\Omega}(t(n))$ is $\Omega(t(n)/\text{poly log}(n))$.

In a recent breakthrough, Kogan and Parter [KP22b] improved the above folklore construction. They showed that for any D , to drop the diameter to $\leq D$, it is sufficient to add a shortcut set of size $\tilde{O}(n^2/D^3)$ if $D < n^{1/3}$ and $\tilde{O}((n/D)^{3/2})$ if $D \geq n^{1/3}$.

In the two most commonly studied scenarios, we get that

- with $\tilde{O}(n)$ shortcuts one can achieve diameter $O(n^{1/3})$, and
- with $\tilde{O}(m)$ shortcuts one can achieve diameter $O((n^2/m)^{1/3})$.

Several papers achieved lower bound constructions. Improving upon Hesse [Hes03], Huang and Pettie [HP18] constructed a family of n -node, m -edge graphs in which one cannot drop the diameter under $\Omega(n^{1/6})$ using $O(n)$ shortcuts, and under $\Omega(n^{1/11})$ using $O(m)$ shortcuts. Lu, Vassilevska Williams, Wein, and Xu [LVWX22] improved the second lower bound, showing that with $O(m)$ shortcuts one cannot drop the diameter below $\Omega(n^{1/8})$. Most recently, Bodwin and Hoppenworth [BH23] improved the first lower bound as well, showing that with $O(n)$ shortcuts one cannot drop the diameter below $\tilde{\Omega}(n^{1/4})$.

Bodwin and Hoppenworth's construction is in fact very sparse: the graphs have n nodes and $O(n)$ edges. However, for technical reasons² their construction does not imply an improvement for the $O(m)$ -shortcut setting, where the $\Omega(n^{1/8})$ lower bound remains the best known.

Thus, using $O(n)$ shortcuts one can achieve diameter $O(n^{1/3})$ but cannot go below diameter $\tilde{\Omega}(n^{1/4})$. Using $O(m)$ shortcuts one can achieve diameter $O((n^2/m)^{1/3})$ and cannot go below diameter $\Omega(n^{1/8})$. The gap in the second case above is particularly big, especially for small m . As m becomes closer and closer to n , the upper bound exponent becomes closer and closer to $1/3$, whereas the lower bound exponent is only $1/8$. For the $O(n)$ -shortcut setting the exponents $1/3$ and $1/4$ are much closer.

Thus a pressing question is:

Question: Can we reduce the gap between $\Omega(n^{1/8})$ and $O((n^2/m)^{1/3})$ for the diameter using $O(m)$ shortcuts?

1.1 Our Results The main contribution of our paper is a simplification and generalization of the Bodwin-Hoppenworth [BH23] lower bound construction. In particular, our new lower bound construction is fully deterministic and unlike all previous lower bound constructions no longer relies on a result of [BL98] on the number of integer points in a d -dimensional ball around the origin. Due to the simplicity of our construction, we are able to generalize it and obtain several new results.

First, we recover the same lower bound as [BH23] in the $O(n)$ -shortcut setting; in fact we additionally remove some log factors.

THEOREM 1.1. *For any $C > 0$, there exist infinitely many n -vertex directed unweighted graph $G = (V, E)$ such that for shortcut set H of size $|H| \leq Cn$, the graphs $G \cup H$ must have diameter $\Omega(n^{1/4})$.*

The construction of Bodwin and Hoppenworth [BH23] gives a diameter lower bound of $\Omega(n^{1/4}/\log^{9/4} n)$. Our construction gets rid off the $\log^{9/4} n$ loss. Similar to [BH23], our construction also extends to the setting where we are allowed to add p shortcut edges for some other values of p . For $p \leq n$, our bound is the same as theirs, up to $\tilde{O}(1)$ factors; for $p \gg n$, our bound is actually better than theirs, improving from $\tilde{\Omega}(n^{5/4}/p)$ to $\Omega(n/p^{3/4})$. This was achieved by removing some small non-tightness in their analysis.

THEOREM 1.2. *For any $p \in [1, n^{4/3}]$, there exist infinitely many n -vertex directed unweighted graphs $G = (V, E)$ such that for shortcut set H of size $|H| \leq p$, the graph $G \cup H$ must have diameter $\Omega(\frac{n}{p^{3/4}})$.*

Next, unlike [BH23], our new construction, coupled with a variant of the potential function argument used by [BH23], allows us to improve upon the $\Omega(n^{1/8})$ lower bound [LVWX22] in the $O(m)$ -shortcut setting, addressing the main question from earlier:

²Their basic construction only shows a lower bound of the diameter for adding $\frac{n}{\alpha \log n}$ shortcut edges for some constant α . For the $O(n)$ -shortcut setting, they have to apply a path subsampling argument which reduces the number of vertices in the graph to n' , so that $\frac{n}{\alpha \log n}$ becomes $c \cdot n'$ for arbitrarily large constant c depending on the parameters for the subsampling. However, the path subsampling argument does not necessarily reduce the number of edges. As a result, they obtain lower bounds for the $O(n)$ -shortcut setting, but not for the $O(m)$ -shortcut setting.

THEOREM 1.3. *For any constant $C > 0$, there exist infinitely many n -vertex m -edge directed unweighted graphs $G = (V, E)$ such that for shortcut set H of size $|H| \leq Cm$, the graph $G \cup H$ must have diameter $\Omega(n^{1/5})$.*

In many works on spanners, emulators, distance preservers, hopsets and so on (e.g. [CE06, Kav17, KV15, CGK13, KP22a]), one often considers restrictions to the objects of study, where instead of addressing all pairs of vertices, only certain pairs are of interest. The two common settings are (a) the *pairwise* setting in which one is given a set $P \subseteq V \times V$ and one only cares about the paths or distances between node pairs in P , and (b) the *sourcewise* setting in which one is given $S \subseteq V$ and one only cares about pairs in $S \times V$.

Motivated by this, we consider the *sourcewise* setting for shortcut sets: given a graph $G = (V, E)$, a set $S \subseteq V$, if we add a set of edges H from the transitive closure of G , how much can we decrease the sourcewise diameter of G , $\max_{(s,v) \in S \times V, \text{dist}(s,v) < \infty} \text{dist}(s, v)$.

In particular, we show that using $O(n)$ shortcuts, the sourcewise diameter can be reduced to $\tilde{O}(|S|^{1/2})$, following the method of Kogan and Parter [KP22b]. Additionally, for any $d \geq 2$, and $|S| = \tilde{\Theta}(n^{3/(d+3)})$, the sourcewise diameter is $\tilde{\Omega}(|S|^{1/3})$ by adding $O(n)$ or $O(m)$ shortcuts (via a randomized construction).

2 Technical Overview

2.1 Previous Works Before giving an overview of our construction, we first summarize previous approaches to proving lower bounds for the shortcut set problem using $O(n)$ shortcuts. Prior to the work of Bodwin and Hoppenworth [BH23], lower bound constructions for the $O(n)$ -shortcut set problem including [Hes03] and [HP18] are all based on constructing a directed graph G on n -vertices with a set of *critical paths* Π such that $|\Pi| = C \cdot n$ where C can be set as an arbitrarily large constant. The critical path set contains paths of length ℓ that are unique paths between their endpoints, and are *pairwise edge-disjoint*. Then one would need to add at least one shortcut for each critical path in order to decrease the diameter to $o(\ell)$. Using this approach, Hesse [Hes03] obtained a lower bound of $\Omega(n^{1/17})$, and Huang and Pettie [HP18] obtained a lower bound of $\Omega(n^{1/6})$.

Recently, Bodwin and Hoppenworth [BH23] improved the lower bound to $\tilde{\Omega}(n^{1/4})$ by relaxing the property where paths in Π are pairwise edge-disjoint to allow pairs of paths to overlap on *polynomially many* edges. In order to better illustrate our improvements, we give more details of their construction. The construction falls under the same geometric construction framework used in earlier works (e.g. [Hes03, CE06, LVWX22, BH22]) for constructing lower bounds for shortcut sets and related problems. In this framework, every vertex of the graph corresponds to a lattice point in $\mathbb{Z}^d$ and every edge corresponds to a d -dimensional vector. The construction of the set of vectors corresponding to the edges crucially relies on building a *convex set* $B_d(r)$ of vectors of length $O(r)$, taken as the vertices of the convex hull of the set of integer points in a ball of radius r centered at the origin. It is known from [BL98] that $|B_d(r)| = \Theta\left(r^{d \cdot \frac{d-1}{d+1}}\right)$ and $B_d(r)$ has the important property of being *convex*, namely for distinct $\vec{b}_1, \vec{b}_2 \in B_d(r)$, we have $\langle \vec{b}_1, \vec{b}_2 \rangle < \langle \vec{b}_1, \vec{b}_1 \rangle$.

The construction in [BH23] builds on the framework for $d = 2$ with some modifications that enabled them to allow polynomial overlaps between the paths in Π . We now give a high-level sketch of their construction: for sufficiently large positive integer $r > 0$, take the convex set $B := B_2(r^3)$ with size $|B| = \Theta(r^2)$ (but they only keep points in the first quadrant). Let $B = \{\vec{b}_1, \dots, \vec{b}_{|B|}\}$ where the vectors are ordered counterclockwise. The graph $G = (V, E)$ on n vertices with critical paths set Π is defined as follows.

- **Vertex Set.** G is a layered graph with $\ell = \tilde{\Theta}(r^2)$ layers and each layer corresponds to the grid $[\tilde{\Theta}(r^3)] \times [\tilde{\Theta}(r^3)]$.³ A vertex v in layer $i \in [\ell]$ is denoted as (i, x_1, x_2) where the first coordinate specifies that v is in layer i and the last two coordinates specifies that v corresponds to grid point (x_1, x_2) .
- **Edge Set.** Define the set of edge vectors using B as

$$W := \{\vec{w}_i := \vec{b}_{i+1} - \vec{b}_i \mid i \in [|B| - 1]\}.$$

All the edges in G are between adjacent layers. For $i \in [\ell]$, define $E_i = \{\vec{w}_{\lambda_i}, (0, 0)\}$ where $\vec{w}_{\lambda_i}$ is uniformly sampled from the set W with $\lambda_i \in [|B| - 1]$ as the set of edge vectors between layer i and layer $i + 1$. For a vertex (i, x_1, x_2) in layer i , connect it to the vertex $(i + 1, x_1 + e_1, x_2 + e_2)$ for all $(e_1, e_2) \in E_i$ if $(x_1 + e_1, x_2 + e_2)$ stays inside the grid.

³For a nonnegative integer N , $[N]$ denotes $\{1, \dots, N\}$

- **Critical Paths II.** Define the set of directional vectors $D := B$ where $D = \{\vec{d}_1, \dots, \vec{d}_{|B|}\}$ and $\vec{d}_i := \vec{b}_i$. Let the starting zone S for the paths in Π be the middle square whose side length is $1/3$ of the grid length. For each vertex $(1, x_1, x_2)$ in the first layer where $(x_1, x_2) \in S$ and for each directional vector $\vec{d} \in D$, they uniquely define a critical path π in Π as follows: π starts from the vertex $(1, x_1, x_2)$ in layer 1 and in layer i , it takes the edge corresponding to the vector $\arg \max_{\vec{e} \in E_i} \langle \vec{e}, \vec{d} \rangle$ (it can be shown that the choice is unique). Note that if the path falls off the grid, then the path terminates.

The resulting graph satisfies the following properties:

1. $|V| = \ell \cdot \tilde{\Theta}(r^3) \cdot \tilde{\Theta}(r^3) = \tilde{\Theta}(r^8)$, $|E| \leq 2|V|$, $|\Pi| = \tilde{\Theta}(n)$.
2. G has $\ell = \tilde{\Theta}(r^2) = \tilde{\Theta}(n^{1/4})$ layers.
3. Each path in Π is the unique path between its endpoints, and has length $\Theta(\ell)$ with constant probability.
4. With high probability, for any path $\sigma : u \rightsquigarrow v$ of length g in G , there are at most ℓ/g paths in Π containing σ as a subpath.

Using a potential argument, in which the potential is defined as the total distances between the two endpoints of the paths in Π , this construction can give a lower bound of $\tilde{\Omega}(n^{1/4})$ for shortcut sets of size $\leq c_0 \cdot \frac{n}{\log n}$ for some fixed constant c_0 . Using the so-called path subsampling argument, they extended the $\tilde{\Omega}(n^{1/4})$ lower bound to $O(n)$ -sized shortcut sets, which is a substantial improvement compared to the previous best lower bound of $\Omega(n^{1/6})$ given in [HP18].

In the following, we discuss two limitations of [BH23].

The $O(m)$ -shortcut setting. As mentioned earlier, aside from the $O(n)$ -shortcut set problem, another intensely studied setting is the $O(m)$ -shortcut set problem where m is the number of edges in the input graph. An immediate question following from the construction in [BH23] is whether this construction can give an improvement for the $O(m)$ -shortcut setting. In particular, notice that in their construction, the graph has constant, and hence $m = \Theta(n)$, so one may think that in this case a shortcut set of size $O(n)$ is also a shortcut set of size $O(m)$, and thus their construction should directly imply an $\Omega(n^{1/4})$ lower bound in the $O(m)$ -shortcut setting. However, this is not the case.

First we highlight the fact that the construction in [HP18] has $|\Pi| = \Theta(n)$ where the constant can be set arbitrarily large, but in [BH23] the size of Π is $c_0 \cdot \frac{n}{\log n}$ for some fixed constant c_0 that cannot be set arbitrarily large. This means that the construction in [BH23] given above would only give a lower bound for shortcut sets of size $\leq c_0 \cdot \frac{n}{\log n}$, so in order to extend the lower bound to shortcut sets of size $C \cdot n$ for arbitrarily large C , Bodwin and Hoppenworth [BH23] use the so-called path subsampling argument: they consider a smaller graph G' obtained from G by restricting to a random subset of vertices $V' \subseteq V$ and “merging” paths between these vertices with short lengths. Namely, they added the edges between all $u, v \in V'$ such that the distance between u, v is at most $O(\log n)$. Thus G' has fewer vertices compared to G while having roughly the same diameter up to $\log n$ factors. A shortcut H on G' can also be viewed as a shortcut on G as $V(G')$ is a subset of $V(G)$. Furthermore, $n' := |V(G')| < n$, so when $|H| = c_0 \cdot \frac{n}{\log n}$ for some, $|H|$ can be $C \cdot n'$ for arbitrarily large C , if n' (which can be adjusted by the parameters in the subsampling argument) is small enough.

Due to the nature of the operation described above, only the number of vertices in the graph is decreased, and the number of edges does not necessarily decrease. This is the underlying barrier in obtaining lower bounds for $O(m)$ -shortcut sets using the technique in [BH23].

Generalization to higher dimensions A natural follow-up question on the construction given in [BH23] is whether it can be generalized to higher dimensions and give better bounds. Previous constructions in [HP18, LVWX22] can both be easily generalized to higher dimensions and one can optimize the dimension to obtain the best bounds. However, it is unclear how the construction in [BH23] can be generalized to higher dimensions for the following reason: the set of edge vectors W is taken to be the set of differences between adjacent vectors in B ordered counterclockwise. Although it is easy to define the notion of “adjacent vectors” in 2 dimensions, it is unclear how to define an analogous notion in, say, 3 dimensions.

2.2 Our Construction The key observation that led to our construction is the following: notice that in the construction of the critical paths in [BH23], each path has two choices at each layer, which is to take the edge

corresponding to the zero vector $(0,0)$ or to take the edge corresponding to the nonzero vector $\vec{w}_{\lambda_i}$. Thus, a critical path in fact corresponds to a subset of the set of edge vectors W , and the difference between the end point and start point of the path corresponds to sum of the vectors in this subset. The use of the convex set B in the construction of the set of edge vectors W ensures that critical paths are well-defined in the sense that it has a well-defined choice between the two vectors in $E_i = \{\vec{w}_{\lambda_i}, (0,0)\}$ and that the sum of vectors taken in a critical path is the sum of a subset of W whose sum is unique among all subsets, since the subset uniquely maximizes the inner product with the corresponding directional vector. Intuitively, this means that we only need a set of edge vectors W that have lots of “unique subset sums”.

There is in fact a very simple set that satisfies the above described property, which is the set $W := \{1\} \times [r] = \{(1, s) \mid 1 \leq s \leq r\}$ for some integer r . Between layer i and layer $i+1$, we put the set of edge vectors as $E_i = \{\vec{w}, (0,0)\}$ for some $\vec{w} \in W$ and each vector $\vec{w} \in W$ is assigned to a unique E_i , so we have $\ell = r+1$. Then we define the critical paths as the paths that take the nonzero edge vectors of the form $\{1\} \times [t]$ for $t \in [r]$ starting in some starting region in the first layer. Notice that given a starting vertex $(0, x_1, x_2)$, a critical path taking the vectors $\{1\} \times [t]$ will end at the vertex $(r, x_1 + t, x_2 + \sum_{j=1}^t j)$. It is easy to see that the start vertex and end vertex uniquely determine a path since the second coordinate indicates the number t of nonzero vectors taken and the third coordinate indicates that the sum of these t distinct numbers is $\sum_{j=1}^t j$. The unique subset of $[r]$ satisfying these conditions is $[t]$. This construction completely removes the use of the result from [BL98] that bounds the number of vertices of the convex hull over the integer points contained in a d -dimensional ball of radius $r > 0$. In addition, we eliminate the use of randomness in [BH23] and the new construction is easier to analyze. We further remark that our simplified construction also provides a logarithmic improvement for the lower bound of $O(n)$ -shortcut set over [BH23].

In addition, this construction can be easily generalized to higher dimensions. For instance, when each grid is 3-dimensional, we take the set of edge vectors as $W := \{1\} \times \{0\} \times [r] \cup \{0\} \times \{1\} \times [r]$ and let the critical paths correspond to subsets of the form $\{1\} \times \{0\} \times [s] \cup \{0\} \times \{1\} \times [s']$ for $s, s' \in [r]$. It is not hard to see that given the coordinates of the start vertex and the end vertex of a critical path, there is a unique path between the endpoints. Using such generalizations, we are able to obtain improvements in other settings such as the $O(m)$ -shortcut set setting.

In our construction for the $O(m)$ -shortcut set setting, we also use a refinement of the potential function used in [BH23]. In their work, the potential function is defined as the sum of distances in the graph $G \cup H$ between the endpoints over all critical paths⁴, where H is the added shortcut set. Our potential function is similar, but the distances are the distances in $G \cup E' \cup H$, where E' consists of pairs (u, v) where the distance from u to v in G is small. In our construction where each grid is 3-dimensional, adding a shortcut between two vertices with smaller distance will drop the potential function defined in [BH23] more. Thus, intuitively, adding E' to our graph G will make the potential drop be 0 for adding a shortcut between two vertices with very small distance. As a result, we can add more shortcuts before our potential function drops below a constant fraction of its initial value.

We will discuss our constructions in more detail in Section 4 and Section 5.

3 Preliminaries

For nonnegative integer N , we use $[N]$ to denote $\{1, \dots, N\}$, and $\llbracket N \rrbracket$ to denote $\{0, \dots, N-1\}$.

We use $\log$ in this paper to denote $\log$ base 2.

For a graph $G = (V, E)$ and vertices $u, v \in V$, we use $\text{dist}_G(u, v)$ to denote the distance from u to v in G . The subscript G might be dropped if it is clear from context. For a path π , we use s_π to denote its start point and t_π to denote its end point.

4 Base Construction

In this section, we give a simplified deterministic construction of a graph that gives an $\Omega(n^{1/4})$ lower bound in the $O(n)$ -shortcut regime, matching the result in [BH23].

4.1 Construction of $G = (V, E)$ with Path Set P First, we describe our graph construction. Let r be an integer and we may assume without loss of generality that r is a power of 2.

⁴We slightly abuse notation and use $G \cup H$ to denote $(V(G), E(G) \cup H)$.

Vertex set. The graph has $\ell = r + 1$ layers indexed from 0 to r and each layer is the grid $\llbracket 2r \rrbracket \times \llbracket 2r^2 \rrbracket$. In every layer, each vertex corresponds to a point $(x_1, x_2) \in \llbracket 2r \rrbracket \times \llbracket 2r^2 \rrbracket$. We denote a vertex v in layer i as (i, x_1, x_2) where the first coordinate specifies that v is in layer i , and the last two coordinates (x_1, x_2) specifies the grid point v corresponds to.

Edge set. Consider the set of integer vectors $W = \{1\} \times \llbracket r \rrbracket$. We only add edges between adjacent layers as follows: between layer i and layer $i + 1$ for $i = 0, \dots, r - 1$, define the edge vector set $E_i = \{(w_1, w_2), (0, 0)\}$ for some $(w_1, w_2) \in W$. (We will specify how to pick these $(w_1, w_2) \in W$ in details later.) For every vertex (i, x_1, x_2) , we add an edge connecting to the vertex $(i + 1, x_1 + e_1, x_2 + e_2)$ for every $(e_1, e_2) \in E_i$. Note that we will not add an edge if $(x_1 + e_1, x_2 + e_2) \notin \llbracket 2r \rrbracket \times \llbracket 2r^2 \rrbracket$.

Now we describe how to assign the nonzero edge vector $(w_1, w_2) \in W$ to edge vector set E_i for each adjacent layer $i \in \llbracket r \rrbracket$. Let q be a permutation on $\llbracket r \rrbracket$ such that the $(\log r)$ -bit binary representation of q_i is exactly the binary representation of i reversed (for simplicity, we use sets of integers indexed from 0, since integers in $\llbracket r \rrbracket$ can all be represented using $(\log r)$ bits). See Table 1 for an example. Then $E_i = \{(1, q_i), (0, 0)\}$.

i	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
q_i	0	8	4	12	2	10	6	14	1	9	5	13	3	11	7	15

Table 1: Example of the permutation q for $r = 16$.

Critical paths P . For any $(x_1, x_2) \in \llbracket r \rrbracket \times \llbracket r^2 \rrbracket$ and for any $s \in [r]$, we define a critical path as follows. The starting point of the critical path is $(0, x_1, x_2)$. For any layer i , the critical path uses the edge corresponding to $(1, q_i)$ if $q_i < s$; otherwise, it uses the edge corresponding to $(0, 0)$. In particular, the set of nonzero vectors taken by a critical path is of the form $\{1\} \times \llbracket s \rrbracket$ for $s \in [r]$.

Properties of G . By definition of our construction, the graph $G = (V, E)$ with the critical path set P has

$$\begin{aligned} |V| &= \ell \times (2r)(2r^2) = 4r^3(r + 1), \\ |E| &\leq 2(\ell - 1)(2r)(2r^2) = 8r^4, \\ |P| &= r \cdot r^2 \cdot |D| = r^4. \end{aligned}$$

Next we show that all critical paths are unique paths between their endpoints and they have relatively small pairwise overlaps.

LEMMA 4.1. *For every critical path $\pi \in P$, π is the unique path between its endpoints in G .*

Proof. Let $v = (1, x_1, x_2)$ be the start vertex of π and suppose the subset of nonzero edge vectors used by π is exactly the set $\{1\} \times \llbracket s \rrbracket$ where $s \in [r]$. Then we can compute the endpoint of π as $u := (r + 1, y_1, y_2)$ where $y_1 = x_1 + s, y_2 = x_2 + \sum_{i=0}^{s-1} x_2 = x_2 + s(s - 1)/2$.

Suppose for contradiction that there exists another path $\pi' \neq \pi$ from v to u in G . Let S denote the set of nonzero edge vectors in π' , then we know that since $\pi' \neq \pi$ and each adjacent layer corresponds to a unique nonzero edge vector, we must have $S \neq \{1\} \times \llbracket s \rrbracket$. Since all the nonzero edge vectors have the first coordinate equal to 1 and we know the endpoints of π' correspond to the grid point (x_1, x_2) and $(x_1 + s, x_2 + s(s - 1)/2)$, we know that π' must take exactly s nonzero edge vectors. Thus we must have $|S| = s$. Then since the second coordinates of all the edge vectors are exactly the set $\llbracket r \rrbracket$, in order for π' to reach u , the sum of the second coordinates of the s vectors in S must be $s(s - 1)/2$. It is easy to see that the only solution is to take the values $0, \dots, s - 1$ for the second coordinates of the vectors in S and therefore $S = \{1\} \times \llbracket s \rrbracket$. This gives a contradiction to the assumption that $\pi' \neq \pi$. $\square$

LEMMA 4.2. *For any two critical paths $\pi_1, \pi_2 \in P$ corresponding to the same subset of nonzero edge vectors $\{1\} \times \llbracket s \rrbracket$ (and thus having distinct starting vertices), we have that π_1 and π_2 are vertex-disjoint.*

Proof. As π_1, π_2 have different start point, and at each layer, they will pick the same edge vector, they will stay different at each layer. $\square$

LEMMA 4.3. *Given any path $\sigma : u \rightsquigarrow w$ of length g in G from layer a to layer b where $a, b \in \llbracket \ell \rrbracket$ and $b - a = g$, there are at most $8r/g$ critical paths containing σ .*

Proof. Let k be the smallest integer such that $g \geq 2 \cdot \frac{r}{2^k}$, so that $\{a, \dots, b-1\}$ contains an interval $\{i \cdot \frac{r}{2^k}, i \cdot \frac{r}{2^k} + 1, \dots, i \cdot \frac{r}{2^k} + \frac{r}{2^k} - 1\}$ for some $i \in \llbracket 2^k \rrbracket$. Note that by definition of the permutation q ,

$$\{q_{i \cdot \frac{r}{2^k}}, \dots, q_{i \cdot \frac{r}{2^k} + \frac{r}{2^k} - 1}\} = \{i', i' + 2^k, \dots, i' + r - 2^k\},$$

where the k -bit binary representation of i' is the k -bit binary representation of i reversed.

Consider the set of nonzero edge vectors $S = \bigcup_{i=a}^{b-1} E_i \setminus \{(0,0)\}$. By the above discussion, S must contain edge vectors $\{(1, i' + t \cdot 2^k) \mid t \in \llbracket r/2^k \rrbracket\}$.

Suppose the path σ is used by a critical path π with initial point $(0, x_1, x_2)$ and with the set of edge vectors $\{1\} \times \llbracket s \rrbracket$. Now we want to bound the number of critical paths $\pi' \in P$ that can potentially contain σ as a subpath. By Lemma 4.2, it suffices to count the number of $s' \in \llbracket r \rrbracket$ such that some critical path π' passing through the vertex u in layer a associated with the set of nonzero edge vectors $\{1\} \times \llbracket s' \rrbracket$ can contain σ as a subpath.

We show that π' cannot contain σ as a subpath if $|s' - s| > 2^k$. Notice that if $|s' - s| > 2^k$, then S must contain an edge vector $w \in W$ that is in the symmetric difference of $\{1\} \times \llbracket s \rrbracket$ and $\{1\} \times \llbracket s' \rrbracket$ since S contains the set $\{(1, i' + t \cdot 2^k) \mid t \in \llbracket r/2^k \rrbracket\}$. This means that from layer a to layer b , exactly one of π and π' takes the edge vector w , so π and π' must have taken different paths from layer a to layer b .

It follows that the number of critical paths containing σ as a subpath is at most the number of s' such that $|s' - s| \leq 2^k$. By definition of k , we have $k = \lceil \log \frac{2r}{g} \rceil \leq 1 + \log \frac{2r}{g}$. So $2^k \leq 4r/g$ and thus the number of critical paths containing σ as a subpath is at most $8r/g$ as desired. $\square$

We summarize the properties of the graph constructed above as follows.

PROPOSITION 4.1. *There exists a constant $c_0 > 0$, so that there exists infinitely many n -vertex directed unweighted graphs $G = (V, E)$ and path sets P of size $|P| = c_0 n$ with the following properties:*

1. G has $\ell = \Theta(n^{1/4})$ layers.
2. Each path in P starts in layer 0 and ends in layer $\ell - 1$.
3. Each path in P is the unique path between its endpoints in G .
4. For any path σ of length g , there are at most $8\ell/g$ paths in P that contain σ as a subpath.

4.2 Application: $O(n)$ -Shortcut Lower Bound From Proposition 4.1, we can prove the following lower bound.

PROPOSITION 4.2. *There exists an absolute constant $c_1 > 0$ and infinitely many n -vertex directed unweighted graph $G = (V, E)$ such that for shortcut set H of size $|H| \leq c_1 n$, the graph $G \cup H$ must have diameter $\Omega(n^{1/4})$.*

Proof. Take a graph $G = (V, E)$ from Proposition 4.1 with path set P of size $|P| = 10c_1 n$ for a suitable constant c_1 . Let H be a shortcut set of size $|H| \leq c_1 n$ on G .

We define the following potential function over shortcut sets H as

$$\phi(H) = \sum_{\pi \in P} \text{dist}_{G \cup H}(s_\pi, t_\pi),$$

i.e., $\phi(H)$ is the sum over the distances between pairs of endpoints of the critical paths after adding the shortcut set H . Then by Proposition 4.1 we have

$$\phi(\emptyset) = \sum_{\pi \in P} \text{dist}_G(s_\pi, t_\pi) = 10c_1 n \cdot (\ell - 1).$$

Now for any path $\sigma : u \rightsquigarrow v$ of length g in G , adding the shortcut (u, v) only affects the critical paths π that contains σ as a subpath by Lemma 4.1. For any shortcut set H' and not including the shortcut (u, v) , let $\text{dist}_1(\cdot, \cdot)$

denote the distance function $\text{dist}_{G \cup H'}(\cdot, \cdot)$ and dist_2 denote the distance function $\text{dist}_{G \cup H' \cup \{(u,v)\}}$ for simplicity. Then we have

$$\begin{aligned}
& \text{dist}_1(s_\pi, t_\pi) - \text{dist}_2(s_\pi, t_\pi) \\
& \leq \text{dist}_1(s_\pi, t_\pi) - \min \{ \text{dist}_1(s_\pi, t_\pi), \text{dist}_2(s_\pi, u) + \text{dist}_2(u, v) + \text{dist}_2(v, t_\pi) \} \\
& \leq \max \{ 0, \text{dist}_1(s_\pi, t_\pi) - (\text{dist}_2(s_\pi, u) + \text{dist}_2(u, v) + \text{dist}_2(v, t_\pi)) \} \\
& \leq \max \{ 0, (\text{dist}_1(s_\pi, u) + \text{dist}_1(u, v) + \text{dist}_1(v, t_\pi)) - (\text{dist}_2(s_\pi, u) + \text{dist}_2(u, v) + \text{dist}_2(v, t_\pi)) \} \\
& \leq \text{dist}_1(u, v) - \text{dist}_2(u, v) \\
& \leq g - 1 \leq g.
\end{aligned}$$

By Item 4 in Proposition 4.1, adding the edge (u, v) only affects the distance of $8\ell/g$ pairs of (s_π, t_π) for $\pi \in P$, so we have for any shortcut set H' ,

$$\phi(H') - \phi(H' \cup \{(u, v)\}) \leq g \cdot \frac{8\ell}{g} = 8\ell.$$

Therefore for a shortcut set H of size $|H| \leq c_1 n$, we have

$$\phi(H) \geq \phi(\emptyset) - 8\ell c_1 n \geq 10(\ell - 1)c_1 n - 8\ell c_1 n = (2\ell - 10)c_1 n.$$

Over $|P| = 10c_1 n$ pairs of critical paths, the average distance is $\Theta(\ell)$, which implies that there must exist a pair with distance $\Theta(\ell)$. Then it follows that the diameter of $G \cup H$ is $\Theta(\ell) = \Theta(n^{1/4})$ by Item 1 in Proposition 4.1. $\square$

Similar to [BH23], Proposition 4.1 and Proposition 4.2 can be extended to show that, for every $p \in [1, c_1 n]$, there are n -node graphs for which adding p shortcut edges keeps diameter as $\Omega(\frac{n}{p^{3/4}})$. The basic idea is to replace each vertex in the graph with a path of certain length, so that the number of vertices and number of layers increase, and p remains the same. We will omit its proof for conciseness. For larger p , we can also obtain a lower bound by combining Proposition 4.2 and the path subsampling argument in [BH23]. Directly applying their path subsampling argument, however, will incur the loss of log factors in the diameter lower bound and requires the use of randomization. Thus, we provide a different proof that is similar to the path subsampling argument. Additionally, our bound for superlinear p is slightly better than that in [BH23], as they had a small non-tightness in their analysis.

THEOREM 4.1. *For any $p \in [1, n^{4/3}]$, there exist infinitely many n -vertex directed unweighted graphs $G = (V, E)$ such that for shortcut set H of size $|H| \leq p$, the graph $G \cup H$ must have diameter $\Omega(\frac{n}{p^{3/4}})$.*

Proof. Let $G = (V, E)$ be a graph from Proposition 4.2, such that when adding $p = c_1 N$ shortcut edges, the diameter of G is $\Omega(N^{1/4})$. Also, note that G is a layered graph with $\Theta(N^{1/4})$ layers. We will focus on the $p > c_1 n$ case in the proof.

We then create a graph G' from G as follows. Let $k = N/n$ and let $V' \subseteq V$ be a subset of vertices that keep all vertices in every other k layers so that $|V'| = n$. For every pair of vertices $u, v \in V'$, we add an edge from u to v in G' if and only if u can reach v in G and u and v are in two adjacent layers in V' .

Suppose for contradiction we can add $\leq p$ shortcut edges H so that the diameter of G' becomes $o(\frac{n}{p^{3/4}}) = o(\frac{n}{N^{3/4}})$ (this would imply $N = o(n^{4/3})$, as otherwise, the diameter becomes $o(1) < 1$ which is impossible). Then we can add these same shortcut edges to G . For any pair of vertices $s, t \in V$ where s can reach t , let π be any path from s to t . Then among the first k vertices of π , there must be a vertex s' in V' ; similarly, among the last k vertices of π , there must be a vertex t' in V' . As H reduces the diameter of G' to $o(\frac{n}{N^{3/4}})$, and each edge in G' corresponds to a path of length k in G , the distance from s' to t' in $G \cup H$ must be $o(k \cdot \frac{n}{N^{3/4}}) = o(N^{1/4})$. Consequently, the distance from s to t in $G \cup H$ must be $\leq 2k + o(N^{1/4}) = o(N^{1/4})$ ($k = o(N^{1/4})$ because $N = o(n^{4/3})$). This implies that the diameter of $G \cup H$ is at most $o(N^{1/4})$, contradicting Proposition 4.2. $\square$

Theorem 1.1, which we recall below, is an immediate corollary of Theorem 1.2 by setting $p = O(n)$:

THEOREM 4.2. *For any $C > 0$, there exist infinitely many n -vertex directed unweighted graph $G = (V, E)$ such that for shortcut set H of size $|H| \leq Cn$, the graphs $G \cup H$ must have diameter $\Omega(n^{1/4})$.*

5 Higher Dimensional Construction

5.1 Construction of $G_d = (V_d, E_d)$ with Path Set P_d It is easy to generalize our construction to higher dimensions as follows. Let $d \geq 1$ be an integer and let $G_d = (V_d, E_d)$ be the graph defined as follows. Note that the dimension of the grid in each layer of the graph will be $d + 1$ and that the base construction corresponds to the case where $d = 1$.

Vertex set V_d The graph has $\ell = d \cdot r + 1$ layers and each layer is the grid $\llbracket 2d \cdot r \rrbracket^d \times \llbracket 2d \cdot r^2 \rrbracket$. We denote a vertex v in layer i as $(i, \vec{x})$ where i specifies the layer and $\vec{x} \in \llbracket 2d \cdot r \rrbracket^d \times \llbracket 2d \cdot r^2 \rrbracket$ is a grid point.

Edge set E_d Consider the set of integer vectors $W \subseteq \{0, 1\}^d \times \llbracket r \rrbracket$, where the first d coordinates of every vector in W contains exactly a single one and $d - 1$ zeros. Let q be the same permutation on $\llbracket r \rrbracket$ as the one defined in Section 4. For each $k = 1, \dots, d$ and $i \in \llbracket r \rrbracket$, let $\vec{w}_{id+k-1}$ be the vector $(0, \dots, 0, 1, 0, \dots, 0, q_i)$ where the 1 is on the k -th coordinate. Note that $\vec{w}_0, \dots, \vec{w}_{dr-1}$ is a permutation of W .

We only add edges between adjacent layers as follows: between layer i and $i + 1$ for $i \in \llbracket dr \rrbracket$, define the edge vector set $E_i = \{\vec{w}_i, \vec{0}\}$ where $\vec{0}$ is the $(d + 1)$ -dimensional all zero vector. For every vertex $(i, \vec{x})$, we add edge connecting it to the vertex $(i + 1, \vec{x} + \vec{e})$ for $\vec{e} \in E_i$ if $\vec{x} + \vec{e} \in \llbracket 2d \cdot r \rrbracket^d \times \llbracket 2d \cdot r^2 \rrbracket$.

Critical Paths P_d A critical path is defined by a start vertex $(0, \vec{x})$ where $\vec{x} \in \llbracket r \rrbracket^d \times \llbracket r^2 \rrbracket$ and a set of vectors $\bigcup_{i=1}^d \{0\}^{i-1} \times \{1\} \times \{0\}^{d-i} \times \llbracket s_i \rrbracket$ for $s_1, \dots, s_d \in \llbracket r \rrbracket$. For any layer i , the critical path uses the edge corresponding to the nonzero vector w_i if w_i belongs to the set; otherwise, it uses the edge corresponding to $\vec{0}$.

Properties of G_d By definition, the graph $G_d = (V_d, E_d)$ has

$$\begin{aligned} |V_d| &= (\ell - 1) \times (2dr)^d \times (2dr^2) = \Theta(r^{d+3}), \\ |E_d| &\leq 2|V_d| = \Theta(r^{d+3}), \\ |P_d| &= r^d \cdot r^2 \cdot r^d = \Theta(r^{2d+2}). \end{aligned}$$

It is not hard to check that G_d also satisfies the following properties.

PROPOSITION 5.1. (PROPERTIES OF G_d) *The graph $G_d = (V_d, E_d)$ with critical path set P_d satisfies the following:*

1. *For every critical path $\pi \in P_d$, π is the unique path between its endpoints in G .*
2. *For any two critical paths $\pi_1, \pi_2 \in P_d$ associated with the same $s_1, \dots, s_d$ (and thus have distinct starting vertices), we have that π_1 and π_2 are vertex-disjoint.*
3. *Given any path $\sigma : u \rightsquigarrow v$ of length g in G_d from layer a to layer b where $a, b \in \llbracket \ell \rrbracket$ and $b - a = g$, there are at most $O((r/g)^d)$ critical paths containing σ .*

Proof.

1. Suppose the start vertex for π is $u = (0, x_1, \dots, x_d, y)$, and $s_1, \dots, s_d$ are the associated values for π . Then the end point v of π can be computed as

$$\left(\ell - 1, x_1 + s_1, x_2 + s_2, \dots, x_d + s_d, y + \sum_{i=1}^d \frac{s_i(s_i - 1)}{2} \right).$$

Thus, for any path travelling from u to v , it must use s_i nonzero vectors from $\{0\}^{i-1} \times \{1\} \times \{0\}^{d-i-1} \times \llbracket r \rrbracket$ for every $i \in [d]$. Furthermore, the sum of the last coordinate of these vectors must equal $\sum_{i=1}^d \frac{s_i(s_i - 1)}{2}$. As a result, for each i , it must use $\{0\}^{i-1} \times \{1\} \times \{0\}^{d-i-1} \times \llbracket s_i \rrbracket$, because any other choices of vectors would make the sum of the last coordinate larger than desired. Hence, there is a unique path from u to v .

2. As π_1, π_2 have different start vertices, and at each layer, they will pick the same direction vector because they share the same $s_1, \dots, s_d$, so they will stay different at each layer.
3. Suppose the path σ is used by a critical path π with initial point $(0, \vec{x})$ and with associated values $s_1, \dots, s_d$. Now we want to bound the number of critical paths $\pi' \in P$ that can potentially contain σ as a subpath. By Item 2, it suffices to count the number of $s'_1, \dots, s'_d \in \llbracket r \rrbracket$ such that some critical path π' passing through the vertex u in layer a associated with $s'_1, \dots, s'_d$ can contain σ as a subpath.

For every $k \in [d]$, $E_a \setminus \{\vec{0}\}, E_{a+1} \setminus \{\vec{0}\}, \dots, E_{b-1} \setminus \{\vec{0}\}$ contain $\{\{0\}^{k-1} \times \{1\} \times \{0\}^{d-k} \times \{q_i\} \mid \alpha \leq i < \beta\}$ for some α, β where $\beta - \alpha = \Theta(g/d) = \Theta(g)$. By the same proof as Lemma 4.3, if $|s'_k - s_k| > \Omega(r/g)$, then π' cannot contain σ as a subpath. As this holds for every $k \in [d]$, we get that the number of π' that can contain σ as a subpath is $O((r/g)^d)$.

□

5.2 Application I: $O(m)$ -Shortcut Lower Bound For our $O(m)$ -shortcut lower bound, we will use $d = 2$ in the graphs constructed above. Let us summarize the properties.

COROLLARY 5.1. *For infinitely many n , there exists an n -vertex directed unweighted graphs $G = (V, E)$ and path sets P , where $|E| \leq 2n$ and $|P| = \Theta(n^{6/5})$, and with the following properties:*

1. G has $\ell = \Theta(n^{1/5})$ layers.
2. Each path in P starts in layer 0 and ends in layer $\ell - 1$.
3. Each path in P is the unique path between its endpoints in G .
4. For any path σ of length g , there are at most $O((\ell/g)^2)$ paths in P that contain σ .

From Corollary 5.1, we can prove the following lower bound.

THEOREM 5.1. *For any constant $C > 0$, there exist infinitely many n -vertex m -edge directed unweighted graphs $G = (V, E)$ such that for shortcut set H of size $|H| \leq Cm$, the graph $G \cup H$ must have diameter $\Omega(n^{1/5})$.*

Proof. Take any graph $G = (V, E)$ with the associated critical paths P from Corollary 5.1.

Let c be a parameter to be chosen later. Let E' a set of edges that contains (u, v) to E' if $\text{dist}_G(u, v) \leq c$.

Let $\Phi(H)$ be a potential function defined as

$$\Phi(H) = \sum_{\pi \in P} \text{dist}_{G \cup E' \cup H}(s_\pi, t_\pi).$$

Initially,

$$\Phi(\emptyset) = |P| \cdot \frac{\ell}{c} = \Theta\left(\frac{n^{7/5}}{c}\right).$$

Now, suppose we have a shortcut set H' , and we add an edge (u, v) to H' , where u and v are g layers apart and u can reach v in G . For any $\pi \in P$, if π does not have a subpath from u to v , then adding (u, v) will not affect the distance from s_π to t_π , by Item 3 in Corollary 5.1.

If any critical path π contains a subpath from u to v , then there exists a unique path σ from u to v in G , because otherwise, the critical path π will not be the unique path from s_π to t_π . By Item 4 in Corollary 5.1, the number of critical paths that contain σ is $O((\ell/g)^2)$, which implies that the number of critical paths that contain a subpath from u to v is $O((\ell/g)^2)$.

Therefore, adding the edge (u, v) will only affect the distance of $O((\ell/g)^2)$ pairs of (s_π, t_π) for $\pi \in P$. For each such π (for simplicity, we use dist_1 to denote $\text{dist}_{G \cup E' \cup H'}$ and dist_2 to denote $\text{dist}_{G \cup E' \cup H' \cup \{(u, v)\}}$ in the following),

$$\begin{aligned} & \text{dist}_1(s_\pi, t_\pi) - \text{dist}_2(s_\pi, t_\pi) \\ & \leq \text{dist}_1(s_\pi, t_\pi) - \min \{ \text{dist}_1(s_\pi, t_\pi), \text{dist}_2(s_\pi, u) + \text{dist}_2(u, w) + \text{dist}_2(w, t_\pi) \} \\ & \leq \max \{ 0, \text{dist}_1(s_\pi, t_\pi) - (\text{dist}_2(s_\pi, u) + \text{dist}_2(u, w) + \text{dist}_2(w, t_\pi)) \} \\ & \leq \max \{ 0, (\text{dist}_1(s_\pi, u) + \text{dist}_1(u, w) + \text{dist}_1(w, t_\pi)) - (\text{dist}_2(s_\pi, u) + \text{dist}_2(u, w) + \text{dist}_2(w, t_\pi)) \} \\ & \leq \text{dist}_1(u, w) - \text{dist}_2(u, w) \\ & \leq \lceil g/c \rceil - 1. \end{aligned}$$

Therefore, if $g \leq c$, $\Phi(H') - \Phi(H' \cup \{(u, w)\}) = 0$; if $g > c$, $\Phi(H') - \Phi(H' \cup \{(u, w)\}) \leq O((\ell/g)^2) \cdot (\lceil g/c \rceil - 1) = O(\frac{\ell^2}{gc}) = O(\frac{\ell^2}{c^2})$.

Thus, as H contains at most $Cm \leq 2Cn$ edges,

$$\Phi(H) \geq \Phi(\emptyset) - 2Cn \cdot O\left(\frac{\ell^2}{c^2}\right) = \Theta\left(\frac{n^{7/5}}{c}\right) - O\left(\frac{Cn^{7/5}}{c^2}\right).$$

Note that the constant factors hidden in Θ and O above are independent of C and c . Therefore, by setting c big enough, we can ensure that $\Phi(H) \geq \Theta\left(\frac{n^{7/5}}{c}\right)$. By averaging, this implies that there exists $\pi \in P$ where

$$\text{dist}_{G \cup E' \cup H}(s_\pi, t_\pi) \geq \Omega\left(\frac{n^{1/5}}{c}\right).$$

Also, as $\text{dist}_{G \cup H}(s_\pi, t_\pi) \geq \text{dist}_{G \cup E' \cup H}(s_\pi, t_\pi)$, we get that the diameter of $G \cup H$ is $\Omega\left(\frac{n^{1/5}}{c}\right)$. The theorem follows as c is a constant only depending on C . $\square$

5.3 Application II: Near-Quadratic-Shortcut Lower Bound Using the construction of G_d (recall $|V_d| = \Theta(r^{d+3})$, $|P_d| = \Theta(r^{2d+2})$ and other properties in Proposition 5.1), we can show the following Theorem 5.2 that recovers [Hes03, Theorem 4.1]. One small improvement is that the average degree of our construction is a constant, while the average degree of Hesse's construction depends on ε .

THEOREM 5.2. *For all $\varepsilon > 0$, there exists some $\delta > 0$ such that there exist infinitely many n -vertex $O(n)$ -edge graphs G where adding any shortcut set of size $O(n^{2-\varepsilon})$ keeps the diameter of G at $\Omega(n^\delta)$.*

Proof. Let d be a large enough integer so that $(1 - \varepsilon) \cdot \frac{d+3}{d-1} < 1 - \varepsilon/2$. Similar to the proof of Theorem 1.3, we let E' be a set of edges that contain (u, v) if $\text{dist}_{G_d}(u, v) \leq c$, for some parameter c .

Let H be the shortcut set, and let $\Phi(H)$ be a potential function defined as

$$\Phi(H) = \sum_{\pi \in P_d} \text{dist}_{G \cup E' \cup H}(s_\pi, t_\pi).$$

Initially,

$$\Phi(\emptyset) = |P_d| \cdot \frac{\ell}{c} = \Theta\left(\frac{r^{2d+3}}{c}\right).$$

If we add an edge between two vertices that are $g \geq c$ layers apart, the potential function is dropped by at most

$$O\left((r/g)^d \cdot \frac{g}{c}\right) = O\left(\frac{r^d}{g^{d-1}c}\right) \leq O\left(\frac{r^d}{c^d}\right).$$

Thus, if we add $O(r^{d+3} \cdot c^{d-1})$ edges, the potential is still $\Omega\left(\frac{r^{2d+3}}{c}\right)$, so the diameter of the graph is still $\Omega(r/c)$. By setting $c = r^{(1-\varepsilon) \cdot \frac{d+3}{d-1}}$, we conclude that after adding

$$O(r^{d+3} \cdot c^{d-1}) = O(r^{d+3} \cdot r^{(1-\varepsilon) \cdot (d+3)}) = O(n^{2-\varepsilon})$$

edges, the diameter is still

$$\Omega\left(\frac{r}{c}\right) = \Omega\left(\frac{r}{r^{(1-\varepsilon) \cdot \frac{d+3}{d-1}}}\right) = \Omega(r^{\varepsilon/2}) = \Omega(n^{\varepsilon/(2(d+3))}).$$

$\square$

6 Sourcewise $O(n)$ -Shortcut Set

Recall that in the sourcewise $O(n)$ -shortcut set problem, we need to add $O(n)$ shortcut edges to a directed unweighted graph, so that $\max_{(s,v) \in S \times V, \text{dist}(s,v) < \infty} \text{dist}(s, v)$ is minimized. Similar to the $O(n)$ -shortcut set problem, we are interested in how small the quantity could be. We provide both an upper and a lower bound for this problem. The upper bound adapts the techniques of Kogan and Parter [KP22b], and the lower bound uses our higher dimensional construction in Section 5.

6.1 Lower Bound

THEOREM 6.1. *For any integer $d \geq 2$, there exists a graph $G = (V, E)$ on n vertices and $S \subseteq V$ with $|S| = \tilde{\Theta}(n^{3/(d+3)})$, such that when adding $O(n)$ (or $O(m)$) shortcuts, the sourcewise diameter is $\Omega(n^{1/(d+3)})$.*

Proof. In G_d , we independently put each vertex in the first layer in S with probability $q = \frac{\log^2 n}{r^{d+1}}$. Recall that the number of vertices in each layer of G_d is $\Theta(r^{d+2})$, so with high probability, $|S| = \Theta(r^{d+2} \cdot q) = \Theta(r^3 \cdot \log^2 n) = \Theta(n^{3/(d+3)} \cdot \log^2 n)$.

Similar to the proof of Theorem 1.3, let c be a parameter to be chosen later. Let E' be a set of edges that contains (u, v) if $\text{dist}_G(u, v) \leq c$. Additionally, we add (u, v) to E' if u can reach v and v is in the first $\ell/2$ layers.

Let the potential function $\Phi(H)$ for a set of shortcut edges H be defined as

$$\Phi(H) = \sum_{\pi \in P_d, s_\pi \in S} \text{dist}_{G_d \cup E' \cup H}(s_\pi, t_\pi),$$

where we recall that P_d is the set of critical paths in G_d . Because of the way we randomly pick S , the number of critical paths π with $s_\pi \in S$ is $\Theta(|P_d| \cdot q) = \Theta(r^{d+3} \cdot \log^2 n)$ (recall $|P_d| = \Theta(r^{2d+2})$). Therefore, initially, $\Phi(\emptyset) = \Theta(r^{d+4} \cdot \log^2 n/c)$.

Now consider adding an edge (u, v) to H where u can reach v in G . Since if $\text{dist}_{G_d}(u, v) \leq c$, or if v is in the first $\ell/2$ layers, adding it will not change the value of the potential function, we assume that $\text{dist}_{G_d}(u, v) = g > c$ and v is in the second half of layers. First, by Proposition 5.1, the number of critical paths in P using the subpath from u to v is at most $O((r/g)^d)$. Furthermore, for each vertex s in the first layer, the number of critical paths starting from s and at the same time using the subpath from u to v must be using the $s \rightsquigarrow v$ subpath in G_d . As a result, for each fixed s , there are $O((r/\text{dist}_{G_d}(s, v))^d) = O(1)$ critical paths starting from s and using the subpath from u to v . Then by the random sampling of S , with high probability, the number of critical paths π in P using the subpath from u to v with $s_\pi \in S$ is

$$O(\max\{\log n, (r/g)^d \cdot q\}).$$

This implies that, with high probability, the drop of potential caused by adding (u, v) is

$$O(\max\{g/c \cdot \log n, g/c \cdot (r/g)^d \cdot q\}).$$

Because $c \leq g \leq r$, the above can be bounded by

$$O\left(\max\left\{r \log n/c, \frac{r \log^2 n}{c^d}\right\}\right).$$

If we add $O(n)$ edges to H , then the total potential drop can be bounded by

$$O\left(\max\left\{r^{d+4} \log n/c, \frac{r^{d+4} \log^2 n}{c^d}\right\}\right).$$

By taking large enough n and setting c as a large enough constant, we can make it so that the above is less than $\Phi(\emptyset)/2$. Therefore, $\Phi(H) = \Theta(r^{d+4} \log^2 n/c)$ for $H = O(n)$. By averaging, some critical path π with $s_\pi \in S$ will have length $\Omega(r/c)$.

The same proof works for the $O(m)$ -shortcut case because G_d has $m = \Theta(n)$. $\square$

6.2 Upper Bound We can assume without loss of generality that the input graph G is a directed acyclic graph (DAG) by the following trick used in prior works (e.g., see [Ras10]): for each strongly connected component, we add a star connecting every vertex from and to the same vertex, which would reduce the diameter of each strongly connected component to ≤ 2 using only $O(n)$ shortcuts in total. Then we only need to keep one representative vertex in each strongly connected component, and the resulting graph would become a DAG.

LEMMA 6.1. *Given a graph $G = (V, E)$ and a subset $S \subseteq V$, one can reduce the sourcewise diameter to $\tilde{O}(\sqrt{|S|})$ using $O(n)$ shortcuts.*

Proof Sketch. The proof of this lemma is analogous to the shortcut set construction in [KP22b], so we only provide a proof sketch.

Let $D = \sqrt{|S|}$ be our desired diameter. First, given a DAG G , let G^* be its transitive closure. It is known that (e.g., see [GIL⁺21]) for every ℓ , we can decompose the vertices of any DAG to ℓ chains (a chain is a path), and $2n/\ell$ antichains (an antichain is a set of vertices that do not have edges between each other). Thus, we can decompose G^* into $\ell := \frac{16n}{D}$ chains $\mathcal{P} = \{P_1, \dots, P_\ell\}$ and $2n/\ell$ antichains. Then for each P_i , we use known method [Ras10] that decreases the diameter of it to ≤ 2 by adding $O(|P_i| \log |P_i|)$ shortcut edges. Next, let $\mathcal{P}'$ be a set of chains, and we add P_i to $\mathcal{P}'$ for every $i \in [\ell]$ with probability $\Theta(\log n/D)$. Finally, for every $s \in S$ and every $P_i \in \mathcal{P}'$, we add an edge (s, u) to the shortcut set where u is the first vertex on P_i that is reachable from s (this step is different from [KP22b], since here we add edges from vertices in S , while in [KP22b], they add edges from a random subset of vertices).

It is not difficult to see that, with high probability, the number of shortcut edges added is $\tilde{O}(n + \frac{n}{D^2} \cdot |S|) = \tilde{O}(n)$. The correctness follows by the analysis in [KP22b].

To decrease the number of shortcut edges added to $O(n)$, we perform the following two modifications:

- When we add $O(|P_i| \log |P_i|)$ shortcut edges to P_i to make its diameter ≤ 2 , we instead only add $O(|P_i|)$ edges to make its diameter $\leq \tilde{O}(1)$. This can be done by applying the same method to the chain formed by the $\frac{|P_i|}{\log |P_i|}$ evenly spaced-out vertices on the original chain P_i .
- Slightly increase D to $\tilde{O}(\sqrt{|S|})$ so that the $\tilde{O}(\frac{n}{D^2} \cdot |S|)$ part of the edge bound becomes $O(n)$.

□

References

- [ASZ20] Alexandr Andoni, Clifford Stein, and Peilin Zhong. Parallel approximate undirected shortest paths via low hop emulators. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 322–335, 2020.
- [BGW20] Aaron Bernstein, Maximilian Probst Gutenberg, and Christian Wulff-Nilsen. Near-optimal decremental SSSP in dense weighted digraphs. In *Proceedings of the 61st IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1112–1122, 2020.
- [BH22] Greg Bodwin and Gary Hoppenworth. New additive spanner lower bounds by an unlayered obstacle product. In *Proceedings of the 63rd IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 778–788, 2022.
- [BH23] Greg Bodwin and Gary Hoppenworth. Folklore sampling is optimal for exact hopsets: Confirming the $\sqrt{n}$ barrier. In *Proceedings of the 64th IEEE Symposium on Foundations of Computer Science (FOCS)*, page to appear, 2023.
- [BL98] Imre Bárány and David G. Larman. The convex hull of the integer points in a large ball. *Math. Ann.*, 312(1):167–181, 1998.
- [BRR10] Piotr Berman, Sofya Raskhodnikova, and Ge Ruan. Finding sparser directed spanners. In *Proceedings of the 30th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, volume 8, pages 424–435, 2010.
- [CE06] Don Coppersmith and Michael Elkin. Sparse sourcewise and pairwise distance preservers. *SIAM J. Discret. Math.*, 20(2):463–501, February 2006.
- [CFR20] 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)*, pages 336–349, 2020.
- [CGK13] Marek Cygan, Fabrizio Grandoni, and Telikepalli Kavitha. On pairwise spanners. In *Proceedings of the 30th International Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 20, pages 209–220, 2013.
- [Fin20] Jeremy T. Fineman. Nearly work-efficient parallel algorithm for digraph reachability. *SIAM J. Comput.*, 49(5), 2020.
- [GIL⁺21] Fabrizio Grandoni, Giuseppe F. Italiano, Aleksander Lukaszewicz, Nikos Parotsidis, and Przemyslaw Uznanski. All-pairs LCA in dags: Breaking through the $O(n^{2.5})$ barrier. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 273–289, 2021.
- [Hes03] William Hesse. Directed graphs requiring large numbers of shortcuts. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, page 665–669, 2003.

- [HKN14] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Decremental single-source shortest paths on undirected graphs in near-linear total update time. In *Proceedings of the 55th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 146–155, 2014.
- [HP18] Shang-En Huang and Seth Pettie. Lower bounds on sparse spanners, emulators, and diameter-reducing shortcuts. In *Proceedings of the 16th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT)*, 2018.
- [JLS19] Arun Jambulapati, Yang P. Liu, and Aaron Sidford. Parallel reachability in almost linear work and square root depth. In *Proceedings of the 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1664–1686, 2019.
- [Kav17] Telikepalli Kavitha. New pairwise spanners. *Theory Comput. Syst.*, 61(4):1011–1036, 2017.
- [KP22a] Shimon Kogan and Merav Parter. Having hope in hops: New spanners, preservers and lower bounds for hopsets. In *Proceedings of the 63rd IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 766–777, 2022.
- [KP22b] Shimon Kogan and Merav Parter. New diameter-reducing shortcuts and directed hopsets: Breaking the barrier. In *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1326–1341, 2022.
- [KS97] Philip N. Klein and Sairam Subramanian. A randomized parallel algorithm for single-source shortest paths. *J. Algorithms*, 25(2):205–220, 1997.
- [KS21] Adam Karczmarz and Piotr Sankowski. A deterministic parallel APSP algorithm and its applications. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 255–272, 2021.
- [KV15] Telikepalli Kavitha and Nithin M. Varma. Small stretch pairwise spanners and approximate d-preservers. *SIAM J. Discret. Math.*, 29(4):2239–2254, 2015.
- [LVWX22] Kevin Lu, Virginia Vassilevska Williams, Nicole Wein, and Zixuan Xu. Better lower bounds for shortcut sets and additive spanners via an improved alternation product. In *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3311–3331, 2022.
- [Ras10] Sofya Raskhodnikova. Transitive-closure spanners: A survey. In *Property Testing: Current Research and Surveys*, pages 167–196. Springer Berlin Heidelberg, 2010.
- [Tho92] Mikkil Thorup. On shortcutting digraphs. In *Proceedings of the 18th International Workshop on Graph-Theoretic Concepts in Computer Science (WG)*, page 205–211, 1992.
- [UY91] Jeffrey D. Ullman and Mihalis Yannakakis. High-probability parallel transitive-closure algorithms. *SIAM J. Comput.*, 20(1):100–125, 1991.