

Folklore Sampling is Optimal for Exact Hopsets: Confirming the $\sqrt{n}$ Barrier*

Greg Bodwin Gary Hoppenworth
 University of Michigan University of Michigan
 bodwin@umich.edu garytho@umich.edu

Abstract

For a graph G , a D -diameter-reducing exact hopset is a small set of additional edges H that, when added to G , maintains its graph metric but guarantees that all node pairs have a shortest path in $G \cup H$ using at most D edges. A *shortcut set* is the analogous concept for reachability rather than distances. These objects have been studied since the early '90s, due to applications in parallel, distributed, dynamic, and streaming graph algorithms.

For most of their history, the state-of-the-art construction for either object was a simple folklore algorithm, based on randomly sampling nodes to hit long paths in the graph. However, recent breakthroughs of Kogan and Parter [SODA '22] and Bernstein and Wein [SODA '23] have finally improved over the folklore algorithm for shortcut sets and for $(1 + \varepsilon)$ -approximate hopsets. For either object, it is now known that one can use $O(n)$ hop-edges to reduce diameter to $\tilde{O}(n^{1/3})$, improving over the folklore diameter bound of $\tilde{O}(n^{1/2})$. The only setting in which folklore sampling remains unimproved is for exact hopsets. Can these improvements be continued?

We settle this question negatively by constructing graphs on which any exact hopset of $O(n)$ edges has diameter $\tilde{\Omega}(n^{1/2})$. This improves on the previous lower bound of $\Omega(n^{1/3})$ by Kogan and Parter [FOCS '22]. Using similar ideas, we also polynomially improve the current lower bounds for shortcut sets, constructing graphs on which any shortcut set of $O(n)$ edges reduces diameter to $\tilde{\Omega}(n^{1/4})$. This improves on the previous lower bound of $\Omega(n^{1/6})$ by Huang and Pettie [SIAM J. Disc. Math. '18]. We also extend our constructions to provide lower bounds against $O(p)$ -size exact hopsets and shortcut sets for other values of p ; in particular, we show that folklore sampling is near-optimal for exact hopsets in the entire range of parameters $p \in [1, n^2]$.

*This work was supported by NSF:AF 2153680.

1 Introduction

In graph algorithms, many basic problems ask to compute information about the shortest path distances or reachability relation among node pairs in an input graph. In parallel, distributed, dynamic, or streaming settings, algorithm complexity often scales with the *diameter* of the graph, i.e., the smallest integer d such that every connected node pair has a path of at most d edges. Therefore, a popular strategy to optimize these algorithms is to add a few edges to the input graph in preprocessing, with the goal to reduce diameter while leaving the relevant distance or reachability information unchanged. In the context of reachability, this set of additional edges is called a *shortcut set*.

Definition 1 (Shortcut Sets). *For a directed graph G , a D -diameter reducing shortcut set is a set of additional edges H such that:*

- *Every edge $(u, v) \in H$ is in the transitive closure of G ; that is, there exists a $u \rightsquigarrow v$ path in G .*
- *For every pair of nodes (s, t) in the transitive closure of G , there exists an $s \rightsquigarrow t$ path in $G \cup H$ using at most D edges.*

Shortcut sets were introduced by Thorup [33], after they were used implicitly in prior work. Many algorithmic applications of shortcut sets and their relatives were discovered in the following years [2, 6, 10, 15–20, 24–26, 34], but actual *constructions* of shortcut sets were elusive. For most of their history, essentially the only known construction was the following simple algorithm: randomly sample a set S of $|S| = n^{1/2}$ vertices, and add a shortcut edge between each pair of sampled nodes that lie in the transitive closure of the input graph. To argue correctness: for any nodes s, t in the graph where the shortest path $\pi(s, t)$ has length $\gg \tilde{O}(n^{1/2})$, with high probability we sample nodes u, v in S that respectively hit a prefix and suffix of $\pi(s, t)$ of length $\tilde{O}(n^{1/2})$. Using the added shortcut edge (u, v) , we obtain an $s \rightsquigarrow t$ path of length $\tilde{O}(n^{1/2})$. This analysis gives:

Theorem 1 (Folklore, [34]). *Every n -node graph has a $\tilde{O}(n^{1/2})$ -diameter-reducing shortcut set on $O(n)$ edges.*

Remarkably, despite its simplicity, the diameter bound of $\tilde{O}(n^{1/2})$ achieved by the folklore sampling algorithm remained nearly unimproved for 30 years (log factors were removed in [5], improving the diameter bound to $O(n^{1/2})$). This led researchers to wonder if the bound could be improved in the exponent at all. This was finally answered affirmatively in a recent breakthrough of Kogan and Parter [29]:

Theorem 2 ([29]). *The folklore algorithm is **polynomially suboptimal** for shortcut sets. In particular, every n -node graph has a $\tilde{O}(n^{1/3})$ -diameter-reducing shortcut set on $O(n)$ edges.*

Kogan and Parter proved this theorem via an elegant construction based on sampling vertices *and* sampling from a set of carefully-chosen paths from the input graph. Following this, there are two clear avenues for further progress. First, the new diameter bound of $\tilde{O}(n^{1/3})$ is still not necessarily tight. It was still conceivable to improve diameter as far as $O(n^{1/6})$, at which point we encounter a lower bound construction of Huang and Pettie [22] (improving on a classic construction of Hesse [21]). Second, many algorithms aim to compute exact or approximate *shortest paths* of an input graph, rather than *any path* as in the case of shortcut sets/reachability. These algorithms benefit from shortcut-set-like structures that more strongly reduce the number of edges along (near-)shortest paths in the input graph. Such a structure is called a *hopset*:

Definition 2 (Hopsets). *For a graph G and $\varepsilon \geq 0$, a D -diameter reducing $(1+\varepsilon)$ hopset is a set of additional edges H such that:*

- *Every edge $(u, v) \in H$ has weight $w(u, v) := \text{dist}_G(u, v)$.*
- *For every pair of nodes (s, t) in the transitive closure of G , there exists an $s \rightsquigarrow t$ path $\pi(s, t)$ in $G \cup H$ that uses at most D edges, and which satisfies $w(\pi(s, t)) \leq (1 + \varepsilon) \cdot \text{dist}_G(s, t)$.*

When $\varepsilon = 0$, the path $\pi(s, t)$ is required to be an exact shortest path in $G \cup H$, so we call H an *exact hopset*.

A nice feature of the folklore sampling algorithm is that it extends immediately to hopsets with no real changes. This yields:

Theorem 3 (Folklore). *Every n -node graph has a $\tilde{O}(n^{1/2})$ -diameter-reducing (exact or $(1 + \varepsilon)$) hopset on $O(n)$ edges.*

Thus, the hunt is back on for a hopset construction algorithm that beats folklore sampling. Kogan and Parter partially achieved this goal: they extended their shortcut set construction to also show a new diameter bound of $\tilde{O}(n^{2/5})$ for $(1 + \varepsilon)$ hopsets [29]. Bernstein and Wein [7] then developed a clever extension of the Kogan-Parter construction, further improving the bound for $(1 + \varepsilon)$ hopsets to match the one achieved for shortcut sets:

Theorem 4 ([7, 29]). *The folklore algorithm is **polynomially suboptimal** for $(1 + \varepsilon)$ hopsets. In particular, for all fixed $\varepsilon > 0$, every (possibly directed and weighted) n -node graph has a $\tilde{O}(n^{1/3})$ -diameter-reducing shortcut set on $O(n)$ edges.*

Still, both of these improvements required $\varepsilon > 0$, and so neither extended to *exact* hopsets, which still remained as the last holdout where the folklore algorithm had not been improved. The only progress for exact hopsets came on the lower bounds side, where a separate work of Kogan and Parter [29] showed a diameter lower bound of $\Omega(n^{1/3})$ (see also [9]). Was it possible to translate the recent progress on shortcut sets and $(1 + \varepsilon)$ hopsets to exact hopsets, and finally move past folklore?

1.1 Our Results




Figure 1: State-of-the-art bounds for $O(n)$ -size hopsets and shortcut sets, before and after this paper

Our main results are polynomial improvements on the lower bounds for $O(n)$ -size shortcut sets and hopsets, which we depict in Figure 1. These lower bounds *confirm* that the folklore algorithm for exact hopsets is essentially the right one, showing that its diameter bound is optimal up to $\log n$ factors:

Theorem 5 (New). *The folklore algorithm is **near-optimal** for exact hopsets. In particular, there are n -node graphs on which any exact hopset on $O(n)$ edges reduces diameter to $\tilde{\Omega}(n^{1/2})$.*

This provides a strong separation between exact and $(1 + \varepsilon)$ hopsets. Our lower bound holds for both directed and undirected input graphs, but it critically uses edge weights and thus does not extend to unweighted graphs as well. Our lower bound in the body of the paper (Theorem 7) is actually a bit more general than the one stated in Theorem 5 above. Although it is popular to focus on hopsets of size $O(n)$, one can also ask about hopsets of size $O(p)$ for any parameter $p \in [1, n^2]$. The folklore sampling algorithm extends to any p , by adjusting the size of the sampled vertex set to $p^{1/2}$. Our generalized lower bound establishes that the diameter bound from folklore sampling is near-optimal in the entire range of parameters.

We now turn to shortcut sets. Our new lower bound is the following:

Theorem 6 (New). *There are n -node directed input graphs on which any shortcut set on $O(n)$ edges reduces diameter to $\tilde{\Omega}(n^{1/4})$.*

This improves over the previous lower bound of $\Omega(n^{1/6})$ by Huang and Pettie [22], but a polynomial gap to the upper bound of $\tilde{O}(n^{1/3})$ still remains [29]. It is an interesting open question is to narrow this gap further. We note that every $(1 + \varepsilon)$ hopset is also a shortcut set, and so this lower bound extends to $(1 + \varepsilon)$ hopsets as well.¹ In the body of the paper (Theorem 8) we prove a more general theorem giving improved lower bounds against $O(p)$ size hopsets, but this time the parameter range in which our extended theorem is nontrivial is only $p \in [1, n^{5/4}]$.

1.2 Other Related Work and Open Questions

This work is only concerned with the *existential* bounds that can be achieved for shortcut sets and hopsets. Some prior work in the area has also focused on constructions that are efficient in the appropriate computational model. This was perhaps most famously achieved by Fineman [15], whose breakthrough algorithm for parallel reachability was centered around a new shortcut set construction. His construction reduced diameter to $\tilde{O}(n^{2/3})$ using $O(n)$ edges; this is a worse diameter bound than the one achieved by folklore sampling, but crucially its work-efficiency was much better than folklore. This was later improved by Jambulapati, Liu, and Sidford [24], who achieved roughly the diameter bound from folklore sampling with work-efficiency comparable to Fineman [15]. Relatedly, another work of Kogan and Parter [27] gave a construction improving the (centralized) construction time of their shortcut set algorithm.

This work focuses on hopsets for *weighted* graphs, but hopsets can be studied for unweighted graphs as well. Specifically for $(1 + \varepsilon)$ hopsets in undirected unweighted graphs, it is known that far better diameter bounds are achievable. In particular, following preliminary constructions in [11, 26, 32], constructions of Huang and Pettie [23] and Elkin and Neiman [14] showed that one can reduce diameter to $n^{o(1)}$ using $O(n)$ hop edges. These papers actually provide a more fine-grained tradeoff between hopset size, ε , and diameter bound, which is shown to be essentially tight in [1]. Hopsets for undirected unweighted graphs with larger stretch values were studied in [4], and a unification of the various hopset constructions in this setting was developed in [31].

This work focuses on shortcut/hopset bounds in the setting where the edge budget is $O(n)$, or more generally some function of n . These objects are also sometimes studied in a related setting where the edge budget is $O(m)$, where m is the number of edges in the input graph. Lower bounds in this setting were achieved in [21, 22, 30]; most recently, Lu, Williams, Wein, and Xu showed graphs where any $O(m)$ -size shortcut set reduces diameter to $\Omega(n^{1/8})$. On the upper bounds side, it is easy to get upper bounds as functions of both m and n – for example, folklore sampling with $|S| = m^{1/2}$ sampled nodes yields a diameter bound of $\tilde{O}(n/m^{1/2})$, and the construction by Kogan and Parter [29] implies a bound of $\tilde{O}(n^{2/3}/m^{1/3})$ for $O(m)$ -size shortcut sets. However, it is an interesting open problem to construct $O(m)$ -size shortcut/hopsets with nontrivial diameter upper bounds that depend only on n . By “nontrivial,” we mean that one can always assume without loss of generality that the input graph is connected, and thus $m = \Omega(n)$, and so a construction of $O(n)$ -edge shortcut/hopsets is always valid in the $O(m)$ -setting. A nontrivial construction is one that beats that bound.

¹Note that, since the shortcut set lower bound is only for directed graphs, the lower bound only extends to $(1 + \varepsilon)$ hopsets for *directed* input graphs. For $(1 + \varepsilon)$ hopsets in undirected (but possibly weighted) input graphs, Elkin, Gitlitz, and Neiman proved that much better diameter can be achieved [13].

Open Question 1. *Prove that, for any m , every n -node, m -edge directed graph has an $O(m)$ -edge shortcut/hopset that reduces diameter to $O(n^c)$, where c is a constant strictly less than the one that can currently be achieved for $O(n)$ -edge shortcut/hopsets.*

Finally, we remark that the setting of exact hopsets in unweighted graphs seems to be unexplored. Although unpublished to our knowledge, one can obtain a lower bound by applying a standard analysis in [9, 22, 28] to the n -node undirected unweighted distance preserver lower bound graphs constructed by Coppersmith and Elkin [12]. This would imply that any $O(n)$ -size exact hopset on the Coppersmith-Elkin graphs would reduce diameter to $\Omega(n^{1/5})$. Our new hopset lower bound can also be interpreted as an improved lower bound for this setting:²

Corollary 1. *There are n -node undirected unweighted input graphs on which any exact hopset on $O(n)$ edges reduces diameter to $\tilde{\Omega}(n^{1/4})$.*

But on the upper bounds side, folklore sampling remains the best known algorithm, and it only reduces diameter to $\tilde{O}(n^{1/2})$. It would be interesting to narrow this gap, and in particular to confirm or refute whether folklore sampling is near-optimal.

Open Question 2. *Is the folklore algorithm near-optimal for exact hopsets in unweighted graphs? Or, alternately, does every n -node unweighted graph have an $O(n)$ -size exact hopset that reduces diameter to $O(n^{1/2-c})$, for some $c > 0$?*

2 Technical Overview

2.1 Recap of Prior Work

In order to explain the strategy used for our improved lower bounds, it will be helpful to first recall the construction of Huang and Pettie [22] for shortcut set lower bounds, and the construction of Kogan and Parter [28] for hopset lower bounds. The Huang-Pettie shortcut set lower bound is a construction of a directed graph G and a set of paths Π with the following properties:

1. Each path in Π is the unique path in G between its endpoints
2. The paths in Π are pairwise edge-disjoint
3. There are $|\Pi| = cn$ paths, where c is a constant that can be selected as large as we want
4. Subject to the above constraints, we want to make the paths in Π as long as possible. In particular, in [22], every path in Π has the same length $\ell = \Theta(n^{1/6})$.

Let us see why these properties imply a lower bound. Each time we add an edge (u, v) to our hopset H , by uniqueness and edge-disjointness of paths in Π , there can be at most one path $\pi \in \Pi$ where the distance between its endpoints decreases due to (u, v) . Thus, if we build a shortcut set of size $|H| = |\Pi| - 1$, then for at least one path $\pi \in \Pi$ the distance between its endpoints is the same in G as in $G \cup H$. Thus, the final diameter of the graph is at least $\ell = \Theta(n^{1/6})$, giving the lower bound.

The Kogan-Parter exact hopset lower bound is similar, except that each path $\pi \in \Pi$ is only required to be a unique *shortest* path between its endpoints in the weighted graph G . This is a more relaxed constraint than requiring π to be the unique path of any kind, and this additional freedom in the construction lets us improve the path lengths to $\ell = \Theta(n^{1/3})$. Besides that, the argument is identical.

²This corollary is not immediate from the discussion so far: since our shortcut set lower bound is directed, it is not clear that it would imply a lower bound against *undirected* unweighted exact hopsets. This holds specifically because our shortcut set lower bound construction happens to be layered.

2.2 Allowing Paths to Overlap

The change in our construction is a relaxation of item (2); that is, the paths in our constructions are not pairwise edge-disjoint. This has appeared in prior work only in a rather weak form: all of the lower bounds against $O(m)$ -size shortcut sets use paths that may intersect pairwise on a single edge [21, 22, 30]. These constructions begin with a system of paths as above, and then apply a tool called the *alternation product* which introduces path overlap. However, the alternation product is not a useful tool for our purposes, and it does not appear in this paper at all. The alternation product *harms* the path lengths of the construction (relative to its number of nodes), in exchange for also reducing the number of edges m relative to the number of paths in the construction. This is useful for constructing lower bounds against $O(m)$ -size shortcut/hopsets, but is not helpful for our goal of constructing lower bounds against $O(n)$ -size shortcut/hopsets.

In our construction, paths that may intersect pairwise on *polynomially* many edges. This property arises from an entirely different technique, and for an entirely different reason: our goal is to use this overlap to get *improved* path lengths. Let us first observe why we can tolerate some path overlap while maintaining correctness of the lower bound. Suppose our shortcut set has a budget of cn edges, and we construct a graph G and a set of $|\Pi| = 2cn$ paths, where each path has length ℓ and each path is the unique path between its endpoints. Let P be the set of node pairs that are the endpoints of paths in Π , and consider the following potential function over shortcut sets H , which simply sums distances over critical pairs:

$$\phi(H) := \sum_{(s,t) \in P} \text{dist}_{G \cup H}(s, t)$$

Initially, we have $\phi(\emptyset) = |\Pi| \cdot \ell = 2cn\ell$. Then we add edges to our shortcut set one at a time, gradually reducing ϕ . How much could any given shortcut edge $(u, v) \in H$ reduce ϕ ? Clearly it could reduce by $\ell - 1$, in the case where (u, v) are a pair in P , since this reduces $\text{dist}(u, v)$ from ℓ to 1. This is acceptable: if *all* edges reduce ϕ by at most $\ell - 1$, then the final potential would be

$$\phi(H) \geq \phi(0) - (\ell - 1)cn = 2\ell cn - (\ell - 1)cn = (\ell + 1)cn.$$

Thus, over the $|P| = 2cn$ critical pairs, the average distance in $G \cup H$ is $\Theta(\ell)$, and so the lower bound still holds.

So the only overlap constraint we need to enforce in our lower bound is that no shortcut edge can reduce the potential ϕ by more than ℓ . This is a much more forgiving constraint than edge-disjoint paths. For example, for two internal nodes u, v with $\text{dist}_G(u, v) = \ell/2$, we could allow two different paths to coincide on a $u \rightsquigarrow v$ subpath: adding (u, v) to the shortcut set would then reduce ϕ by only $2 \cdot (\ell/2 - 1) = \ell - 2$. In general, for two nodes at distance ℓ/x , we can safely allow x paths to coincide on the subpath between these nodes while maintaining correctness of the lower bound.

2.3 Constructing Overlapping Paths

The previous part explains why we are *allowed* overlapping paths, but it is still not clear how to leverage this freedom into an improved lower bound construction. This is where our technical contribution lies. It is again a bit easier to explain the new idea in the context of shortcut sets, but the intuition is essentially the same in the context of hopsets.

Let us return to the previous lower bound constructions. For the shortcut lower bounds of [22], one constructs an $(\ell + 1)$ -layered directed graph where the nodes in each layer are a copy of a grid within $\mathbb{Z}^2$. The next step is to construct a set of convex vectors C . A key perspective shift in this paper is that we think of the vectors in C as playing two *independent* roles in this previous construction:

- They play the role of *edge vectors*: we include an edge from a node u in layer i to a node v in layer $i + 1$ iff the difference between the grid points u, v is a vector $v - u = \vec{c} \in C$.
- They also play the role of (*objective*) *direction vectors*. The paths $\pi \in \Pi$ are indexed by a node u in the first layer and a vector $\vec{c} \in C$, and we generate $\pi_{u, \vec{c}}$ by using u as its first node, and then iteratively selecting its node in the next layer by adding $\vec{c}$ to the node in the previous layer. (A technical detail is that only choices of $(u, \vec{c})$ are allowed that reach the last layer without the path falling off the side of the

grid.) Notice the argument for path uniqueness: using $\vec{c}$ as an objective direction, due to the convexity of the vectors in C , the edge vector $\vec{c}$ itself is the one that maximizes progress in the direction $\vec{c}$. Thus, no alternate path beginning at u can reach the grid point $u + \ell\vec{c}$ within ℓ steps, since it necessarily makes less progress in the direction of $\vec{c}$, implying path uniqueness.

Our constructions disentangle these two uses of the vector set C : we depart from prior work by explicitly using a separate edge vector set C and direction vector set D . These vector sets crucially do *not* have the same size: instead we will have $|D| \gg |C|$, and this difference allows for a technical optimization of parameters leading to improved lower bounds. Roughly, we can choose D large enough to have $|\Pi| = \Theta(n)$, while also allowing $|C| \ll n$. This smaller size $|C|$ can be achieved using shorter convex vectors, which in turn lets us pack more layers into the construction without worrying about paths falling off the sides of the grid before reaching the final layer.

We use the following generalized process for iteratively generating critical paths. Each path $\pi \in \Pi$ is indexed by a node u in the first layer and a direction vector $\vec{d} \in D$, and at each layer, we greedily select the edge vector $\vec{c} \in C$ that maximizes progress in the objective direction $\vec{d}$. Since $|D| \gg |C|$, by necessity many different objective directions will all select the same edge vector at each layer. This leads to overlapping paths discussed above, but more technical ingredients are still needed to ensure that paths don't overlap too much. We explain these next.

2.4 Symmetry Breaking

There is an important problem with the construction sketched so far. Consider two nearby direction vectors $\vec{d}_1, \vec{d}_2 \in D$, which have the same optimizing edge vector $\vec{c}$. Then for any start node u , the paths $\pi_{u, \vec{d}_1}, \pi_{u, \vec{d}_2}$ will simply select the same edge vector $\vec{c}$ at each layer, and these two paths will entirely coincide. In other words, the “extra paths” bought by using $|D| \gg |C|$ are actually just identical copies of a much smaller set of paths, which is not interesting or useful.

We therefore need to somehow break the symmetry between paths that use nearby direction vectors, getting them to eventually choose different edge vectors at some layer to split apart. This is where our lower bound constructions diverge; we will need to use two different symmetry-breaking strategies for shortcut and hopset lower bounds.

Hopset Lower Bounds and ε -Shifting. In our hopset lower bound construction G , like [28], our vertices can be interpreted as points in $\mathbb{R}^2$. More specifically, they initially form a square grid within the integer lattice $\mathbb{Z}^2$, and the columns of this grid act as layers of G . Our edges initially have the form $e = (v, v + \vec{c})$ for edge vectors $\vec{c} \in C$; edge vectors always have first coordinate 1, so that they go from one layer to the next. Initially, the weight of an edge is the Euclidean distance between its endpoints.

Our symmetry-breaking step is a random operation where for each layer i we choose a random variable ε_i sampled uniformly from the interval $(0, 1)$, and we shift the i^{th} layer upwards so that its nodes are offset ε_i higher than the nodes in the previous layer. The shifts therefore compound across the layers. See Figure 2 for a picture.

Our ε -shifting strategy does not affect the edge set of G , nor does it affect the set of direction vectors in any way, but it does affect the Euclidean distance between nodes in adjacent layers, and hence it changes the edge weights. It achieves symmetry-breaking for roughly the following reason. In our greedy generation of paths, a path with direction vector $\vec{d}$ will use the closest edge vector $\vec{c}$ at each level. If two paths π_1 and π_2 with direction vectors $\vec{d}_j, \vec{d}_k \in D$ intersect at a node v in the i^{th} layer of G , then there will be an interval $(a, b) \subseteq (0, 1)$ such that if ε_{i+1} lands in (a, b) , then π_1, π_2 have different closest edge vectors after shifting. Thus, in this event, the paths π_1 and π_2 split apart at v and never reconverge (this is formalized in Lemma 3). The size of the interval (a, b) , and hence the probability that it gets hit by ε_{i+1} , is proportional to the distance between $\vec{d}_j$ and $\vec{d}_k$. The effect is that paths generated by nearby direction vectors tend to intersect on long subpaths, while paths generated by far apart direction vectors intersect on shorter subpaths or perhaps just a single node, but with high probability all pairs of paths split apart eventually. Paths generated by *the same* direction vector remain parallel, and do not intersect at all.

There is a technical detail remaining: we still need to prove that each critical path is a unique shortest path between its endpoints. In [28], the critical paths correspond to lines in Euclidean space, and since edge




Figure 2: Vertex set of the graph G used for our lower bounds against exact hopsets. Each parameter ε_i is the amount the i^{th} column is shifted upwards in the plane, relative to the previous column; the $\{\varepsilon_i\}$ values are chosen uniformly and independently from the interval $(0, 1)$.

weights correspond to Euclidean distances, the analogous unique shortest paths property follows instantly from the geometry of $\mathbb{R}^2$. Since our critical paths are generated by a more involved process, it is much more technical to prove that they are unique shortest paths. Proposition 3 contains the optimization lemma that needs to hold for our process to generate unique shortest paths, and to push it through, it turns out that we essentially need the *derivative* of edge weights to be proportional to Euclidean distances. We therefore differ again from [28] by squaring all of our edge weights, meaning that our graph metric is ultimately quite different from $\mathbb{R}^2$. See Section 4.4 for additional details.

Shortcut Set Lower Bounds and Edge Vector Subsampling. Shortcut set lower bounds are unweighted, and this makes the technique of ε -sampling essentially useless in this setting, since it only affects edge weights in the construction and it does not change the edge *set*. For shortcut sets, we need an entirely different symmetry-breaking strategy that actually changes the edge set from layer to layer.

Our starting graph is similar to the one used by Huang and Pettie [22], mentioned earlier. Each layer of the graph is an independent copy of a square grid in $\mathbb{Z}^2$. We generate a large convex set $W \subseteq \mathbb{R}^2$; initially, W plays the role of both edge vectors C and direction vectors D . However, for the sake of symmetry-breaking, we do not put edges between *all* nodes in adjacent layers whose difference is a vector in W . Instead, at each layer i we randomly sample exactly two adjacent vectors $\vec{c}_{\lambda_i}, \vec{c}_{\lambda_i+1} \in W$, and we use only these two edge vectors to generate edges to the next layer. This is depicted in Figure 3.

The fact that the sampled vectors are *adjacent*, and hence typically close together, allows for a key optimization in the construction. The rate at which paths drift apart from each other is much slower than in [22], even when they are generated by very different direction vectors. This allows us to apply a carefully-chosen translation of the grid from layer to layer, in order to keep all of the paths contained in the grid. This in turn lets us pack many more layers into the construction while still ensuring that all of our paths stay within the confines of the grid.

As before, paths $\pi \in \Pi$ are generated greedily: for direction vector $\vec{d} \in D$, an associated path π will traverse the sampled edge vector in each layer that maximizes progress in the objective direction $\vec{d}$. For two paths π_1, π_2 with direction vectors $\vec{d}_j, \vec{d}_k \in D$ and $v \in \pi_1 \cap \pi_2$, these paths split at v in the event that the sampled edge vectors $\vec{c}_{\lambda_i}, \vec{c}_{\lambda_i+1} \in W$ lie *between* $\vec{d}_j$ and $\vec{d}_k$ in D . This again leads to behavior where paths generated by nearby direction vectors tend to coincide on long subpaths, while paths generated by far apart direction vectors have smaller intersections, but with high probability all pairs of paths split apart eventually. Paths with the same direction vector again remain parallel. This is formalized in Lemma 10.




Figure 3: Our symmetry-breaking strategy for shortcut set lower bounds starts with a large set of convex vectors, but independently subsamples adjacent pairs of convex vectors to generate the edges between adjacent layers. In this picture, there are 4 edge vectors and 4 layers, but only two of the edge vectors (in blue) are sampled and available between any given pair of adjacent layers. For clarity, we have only drawn the edges leaving one particular node in each layer.

3 Preliminaries

We use the following notations:

- For a path π , we use $|\pi|$ to denote the number of **nodes** in π . This is 1 different from the (unweighted) length of π . In weighted graphs, we write $w(\pi)$ for the sum of edge weights in π .
- We write $\text{dist}_G(s, t)$ for the shortest path distance from node s to node t in graph G (counting edge weights, if G is weighted). We write $\text{hopdist}_G(s, t)$ for the least number of edges contained in any $s \rightsquigarrow t$ shortest path.
- We use $\langle \cdot, \cdot \rangle$ to denote the standard Euclidean inner product.

4 Exact Hopsets

In this section we will prove the following theorem.

Theorem 7. *For any parameter $p \in [1, n^2]$, there exists an n -node weighted undirected graph $G = (V, E, w)$ such that for any exact hopset H of size $|H| \leq p$ where $p \in [1, n^2]$, the graph $G \cup H$ must have hopbound $\Omega\left(\frac{n}{p^{1/2} \log^{1/2} n}\right)$.*

We will prove this via a construction of the following type:

Lemma 1. *For any $p \in [1, n^2]$, there is an infinite family of n -node undirected weighted graphs $G = (V, E, w)$ and sets Π of $|\Pi| = p$ paths in G with the following properties:*

- G has $\ell = \Theta\left(\frac{n}{p^{1/2} \log^{1/2} n}\right)$ layers, and each path in Π starts in the first layer, ends in the last layer, and contains exactly one node in each layer.
- Each path in Π is the unique shortest path between its endpoints in G .

- For any two nodes $u, v \in V$, there are at most $\frac{\ell}{\text{hopdist}_G(u,v)}$ paths in Π that contain both u and v .

Next we show how Lemma 1 implies Theorem 7.

4.1 Proving Theorem 7 using Lemma 1

Fix an n and $p \in [1, n^2]$. Let $G = (V, E, w)$ be the graph in Lemma 1 with associated set Π of $|\Pi| = 2p$ paths in G . Let H be an exact hopset of size $|H| \leq p$. Let $P \subseteq V \times V$ be the set of node pairs that are the endpoints of paths in Π . We define the following potential function over hopsets H , which simply sums hopdistances over critical pairs:

$$\phi(H) := \sum_{(s,t) \in P} \text{hopdist}_{G \cup H}(s, t)$$

Observe that by Lemma 1, we have $\phi(\emptyset) = \sum_{(s,t) \in P} \text{hopdist}_G(s, t) = |\Pi| \cdot (\ell - 1) = 2p(\ell - 1)$. Now fix a pair of nodes $(x, y) \in V \times V$, and let $\Pi' \subseteq \Pi$ be the set of paths $\pi \in \Pi$ such that $x, y \in \pi$. We make the following observations.

- For all $(s, t) \in P$, if the unique shortest $s \rightsquigarrow t$ -path $\pi \in \Pi$ in G is not in Π' , then

$$\text{hopdist}_{G \cup \{(x,y)\}}(s, t) = \text{hopdist}_G(s, t).$$

- For all $(s, t) \in P$, if the unique shortest $s \rightsquigarrow t$ -path $\pi \in \Pi$ in G is in Π' , then

$$\text{hopdist}_G(s, t) - \text{hopdist}_{G \cup \{(x,y)\}}(s, t) \leq \text{hopdist}_G(x, y).$$

Then by Lemma 1, $\phi(\emptyset) - \phi(\{(x, y)\}) \leq |\Pi'| \cdot \text{hopdist}_G(x, y) \leq \ell$. We obtain the following sequence of inequalities:

$$\phi(\emptyset) - \phi(H) \leq \sum_{(x,y) \in H} (\phi(\emptyset) - \phi(\{(x, y)\})) \leq \ell \cdot |H| \leq \ell p.$$

Rearranging, we find that

$$\phi(H) \geq \phi(\emptyset) - \ell p \geq 2p(\ell - 1) - \ell p = (\ell - 2)p.$$

Thus, over the $|P| = 2p$ pairs of path endpoints in P , the average hopdistance in $G \cup H$ is $\Theta(\ell)$, and so there must be a pair $(s, t) \in P$ such that $\text{hopdist}_{G \cup H}(s, t) = \Theta(\ell)$ by the pigeonhole principle.

4.2 Constructing G

Our goal is now to prove Lemma 1. Let n be a sufficiently large positive integer, and let $p \in [n, n^2]$.³ For simplicity of presentation, we will frequently ignore issues related to non-integrality of expressions that arise in our construction; these issues affect our bounds only by lower-order terms. Initially, all edges (u, v) in G will be directed from u to v ; we will convert G into an undirected graph in the final step of our construction.

Vertex Set V .

- Let ℓ be a positive integer parameter of the construction to be specified later. Our graph G will have ℓ layers $L_1, \dots, L_\ell$, and each layer will have n/ℓ nodes, ordered from 1 to n/ℓ . Initially, we will label the j th node in layer L_i with tuple (i, j) . We will interpret the node labeled (i, j) as a point in $\mathbb{R}^2$ with integer coordinates. These n nodes arranged in ℓ layers will be the node set V of graph G .
- We now perform the following random operation on the node labels of V . For each layer L_i , $i \in [1, \ell]$, uniformly sample a random real number in the interval $(0, 1)$ and call it ε_i . Now for each node in layer L_i of G labeled (i, j) , relabel this node with the label

$$\left(i, j + \sum_{k=1}^j \varepsilon_k \right).$$

³We will handle the case where $p \in [1, n]$ later.

Again, we interpret the resulting labels for nodes in V as points in $\mathbb{R}^2$. In a slight abuse of notation, we will treat $v \in V$ as either a node in G or a point in $\mathbb{R}^2$, depending on the context. Less formally: for each layer i , this step shifts the nodes in layer i vertically upwards to be ε_i higher than the previous layer (and thus, these vertical shifts compound across the layers). See Section 2 for intuition on this design choice.

Edge Set E .

- All the edges in E will be between consecutive layers L_i, L_{i+1} of G . We will let E_i denote the set of edges in G between layers L_i and L_{i+1} .
- Just as our nodes in V correspond to points in $\mathbb{R}^2$, we can interpret the edges E in G as vectors in $\mathbb{R}^2$. In particular, for every edge $e = (v_1, v_2) \in E$, we identify e with the corresponding vector $\vec{u}_e := v_2 - v_1$. Note that since all edges in E are between adjacent layers L_i and L_{i+1} , the first coordinate of $\vec{u}_e$ is 1 for all $e \in E$. We will use u_e to denote the 2nd coordinate of $\vec{u}_e$, i.e., for all $e \in E$, we write $\vec{u}_e = (1, u_e)$.
- We begin our construction of E by defining the following set C of vectors:

$$C := \left\{ (1, x) \mid x \in \left[0, \frac{n}{4\ell^2}\right] \right\}.$$

We will refer to the vectors in C as *edge vectors*.

- For each $i \in [1, \ell - 1]$, let

$$C_i := \{\vec{c} + (0, \varepsilon_{i+1})\}_{\vec{c} \in C}.$$

Intuitively: we want the edge vectors in C to point between nodes in adjacent layers, and due to the random vertical shifts between layers applied to the nodes, we need to apply a similar shift to C at each layer to adjust for this.

- For each $v \in L_i$ and edge vector $\vec{c} \in C_i$, if $v + \vec{c} \in V$, then add edge $(v, v + \vec{c})$ to E_i . After adding these edges to E_i , we will have that

$$C_i = \{\vec{u}_e \mid e \in E_i\}.$$

Additionally, note that the case $v + \vec{c} \notin V$ only occurs if $v + \vec{c} = (i + 1, j)$ for some j that is higher than any point in the $(i + 1)$ st layer; that is, $j > \frac{n}{\ell} + \sum_{k=1}^{i+1} \varepsilon_k$.

- For each $e \in E$, if $\vec{u}_e = (1, u_e)$, then we assign edge e the weight $w(e) := u_e^2$.

This completes the construction of our graph $G = (V, E, w)$.

4.3 Direction Vectors, Critical Pairs, and Critical Paths

Our next step is to generate a set of *critical pairs* $P \subseteq V \times V$, as well as a set of *critical paths* Π . Specifically, there will be one critical path $\pi_{s,t} \in \Pi$ going between each critical pair $(s, t) \in P$, and we will show that $\pi_{s,t}$ is the unique shortest (weighted) $s \rightsquigarrow t$ path in G . We will identify our critical pairs and paths by first constructing a set of vectors $D \subseteq \mathbb{R}^2$ that we call *direction vectors*, which we define next.

Direction Vectors D .

- Let $q \in \mathbb{Z}_{\geq 0}$ be a sufficiently large integer parameter to be specified later. The size of q will roughly correspond to the maximum number of edges shared between any two critical paths in G .
- We choose our set of direction vectors D to be ⁴

$$D := \left\{ \left(1, x + \frac{y}{q}\right) \mid x \in \left[1, \frac{n}{4\ell^2} - 1\right] \text{ and } y \in [0, q] \right\}.$$

Note that there are $q + 1$ direction vectors between adjacent vectors $(1, x), (1, x + 1) \in C$ for $x \neq 0$. Additionally, adjacent direction vectors in D differ only by $1/q$ in their second coordinate.

⁴Note that if $(1, x) \in D$, then $x \in \left[1, \frac{n}{4\ell^2}\right]$. However, if $(1, x) \in C$, then $x \in \left[0, \frac{n}{4\ell^2}\right]$. This $+1$ gap between C and D is needed to accommodate the ε -shifting operation used to obtain C_i , and is relevant in the proof of Lemma 3.

Proposition 1. *With probability 1, for every $i \in [1, \ell - 1]$ and every direction vector $\vec{d} = (1, d) \in D$, there is a unique vector $(1, c) \in C_i$ that minimizes $|c - d|$ over all choices of $(1, c) \in C_i$.*

Proof. There are only finitely many choices of $\varepsilon_{i+1} \in (0, 1)$ that result in there being two distinct vectors $(1, c_1), (1, c_2) \in C_i$ such that $|c_1 - d| = |c_2 - d|$. We conclude that the claim holds with probability 1. $\square$

In the following we assume that this event holds, i.e., there is a unique minimizing vector in C_i for all $\vec{d} \in D$. Each of our critical paths π in Π will have an associated direction vector $\vec{d} \in D$, and for all $i \in [1, \ell - 1]$, path π will take an edge vector in C_i that is closest to $\vec{d}$ in the sense of Proposition 1 (see Section 2 for more intuition).

Critical Pairs P and Critical Paths Π .

- We first define a set $S \subseteq L_1$ containing half of the nodes in the first layer L_1 of G :

$$S := \left\{ (1, j + \varepsilon_1) \in L_1 \mid j \in \left[1, \frac{n}{2\ell}\right] \right\}.$$

We will choose our set of demand pairs P so that $P \subseteq S \times L_\ell$. For every node $s \in S$ and direction vector $\vec{d} \in D$, we will create a critical pair $(s, t) \in S \times L_\ell$ and a corresponding critical path $\pi_{s,t}$ to add to P and Π .

- Let $v_1 \in S$, and let $\vec{d} = (1, d) \in D$. The associated path π has start node v_1 . We iteratively grow π , layer-by-layer, as follows. Suppose that currently $\pi = (v_1, \dots, v_i)$, for $i < \ell$, with each $v_i \in L_i$. To determine the next node $v_{i+1} \in L_{i+1}$, let $E_i^{v_i} \subseteq E_i$ be the edges in E_i incident to v_i , and let

$$e_i := \operatorname{argmin}_{e \in E_i^{v_i}} (|u_e - d|).$$

By definition, e_i is an edge whose first node is v_i ; we define $v_{i+1} \in L_{i+1}$ to be the other node in e_i , and we append v_{i+1} to π .

- This completes our construction of P and Π . Note that
 - we will show that the paths generated in this way have distinct endpoints (with high probability), and therefore $|P| = |S||D| \geq \frac{n^2 q}{16\ell^3}$, and
 - every path $\pi_{s,t} \in \Pi$ contains one node in each layer, and therefore its number of nodes is $|\pi_{s,t}| = \ell$.

An important feature for correctness of our construction is that, when we iteratively generate paths, we never reach a point v_i such that $v_i + C_i \not\subseteq L_{i+1}$ (i.e. $v_i + \vec{c} \notin L_{i+1}$ for some $\vec{c} \in C_i$). This follows by straightforward counting, based on the maximum second coordinate used in our edge vectors C and also on our choice of start nodes S as only the “lower half” of the nodes in the first layer. The following proposition expresses this correctness in a particular way, pointing out that for any node v lying on a generated path π , none of the edges from v to the following layer are omitted from the graph due to falling off the top of the grid with a too-high second coordinate.

Proposition 2. *Let $v \in L_i \cap \pi$ for some $i \in [1, \ell - 1]$ and $\pi \in \Pi$. Then $\{\vec{u}_e \mid e \in E_i^v\} = C_i$.*

Proof. Let $v = (i, j)$, and let $(s, t) \in P$ be the endpoints of π . Since $s \in S$, we have $s = (1, s_2) \in \mathbb{R}^2$, where

$$s_2 \leq \frac{n}{2\ell} + \varepsilon_1.$$

Moreover, since for all $e \in E_i$ the corresponding vector $\vec{u}_e = (1, u_e)$ satisfies $\varepsilon_{i+1} \leq u_e \leq \frac{n}{4\ell^2} + \varepsilon_{i+1}$, we have

$$j \leq \frac{n}{2\ell} + (i-1) \cdot \frac{n}{4\ell^2} + \sum_{k=1}^i \varepsilon_k \leq \frac{3n}{4\ell} + \sum_{k=1}^i \varepsilon_k.$$

Then observe that for all $\vec{c} = (1, c) \in C_i$, we have that $v + \vec{c} = (i+1, j+c)$, where

$$j+c \leq \left(\frac{3n}{4\ell} + \sum_{k=1}^i \varepsilon_k \right) + \left(\frac{n}{4\ell^2} + \varepsilon_{i+1} \right) \leq \frac{n}{\ell} + \sum_{k=1}^{i+1} \varepsilon_k.$$

Thus $(i+1, j+c) \in L_{i+1}$, and so we have $(v, v+\vec{c}) \in E_i^v$ for all $\vec{c} \in C_i$. It follows that $\{\vec{u}_e \mid e \in E_i^v\} = C_i$. $\square$

4.4 Critical paths are unique shortest paths

We now verify that graph G and paths Π have the unique shortest path property as stated in Lemma 1.

Lemma 2 (Unique shortest paths). *With probability 1, for every $(s, t) \in P$, path $\pi_{s,t} \in \Pi$ is a unique shortest (weighted) $s \rightsquigarrow t$ -path in G .*

We begin with a technical proposition:

Proposition 3. *Let $b, x_1, \dots, x_k \in \mathbb{R}$. Now consider $\hat{x}_1, \dots, \hat{x}_k$ such that*

- $|\hat{x}_i - b| \leq |x_i - b|$ for all $i \in [1, k]$, and
- $\sum_{i=1}^k x_i = \sum_{i=1}^k \hat{x}_i$.

Then

$$\sum_{i=1}^k x_i^2 \geq \sum_{i=1}^k \hat{x}_i^2,$$

with equality only if $|\hat{x}_i - b| = |x_i - b|$ for all $i \in [1, k]$.

Proof. We will prove the equivalent statement $\sum_{i=1}^k (\hat{x}_i^2 - x_i^2) \leq 0$. Fix an $i \in [1, k]$. First we will show that

$$\hat{x}_i^2 - x_i^2 \leq 2b(\hat{x}_i - x_i).$$

We split our analysis into four cases:

- **Case 1:** $b \leq \hat{x}_i \leq x_i$. In this case, $\hat{x}_i^2 - x_i^2 = (\hat{x}_i + x_i)(\hat{x}_i - x_i) \leq 2b(\hat{x}_i - x_i)$.
- **Case 2:** $x_i \leq \hat{x}_i \leq b$. In this case, $\hat{x}_i^2 - x_i^2 = (\hat{x}_i + x_i)(\hat{x}_i - x_i) \leq 2b(\hat{x}_i - x_i)$.
- **Case 3:** $\hat{x}_i \leq b \leq x_i$. In this case, $\hat{x}_i^2 - x_i^2 = (\hat{x}_i + x_i)(\hat{x}_i - x_i) \leq 2b(\hat{x}_i - x_i)$, since $b - \hat{x}_i \leq x_i - b$.
- **Case 4:** $x_i \leq b \leq \hat{x}_i$. In this case, $\hat{x}_i^2 - x_i^2 = (\hat{x}_i + x_i)(\hat{x}_i - x_i) \leq 2b(\hat{x}_i - x_i)$, since $\hat{x}_i - b \leq b - x_i$.

Then

$$\sum_{i=1}^k (\hat{x}_i^2 - x_i^2) \leq 2b \sum_{i=1}^k (\hat{x}_i - x_i) = 0.$$

This inequality is strict if $|\hat{x}_i - b| < |x_i - b|$ for some $i \in [1, k]$. □

Using Proposition 3, we can now prove Lemma 2.

Proof of Lemma 2. Fix an $(s, t) \in P$, and let $(1, x) \in D$ be the direction vector associated with $\pi_{s,t}$. Let $\hat{x}_1, \dots, \hat{x}_{\ell-1} \in \mathbb{R}$ be real numbers such that the i th edge of $\pi_{s,t}$ has the corresponding vector $(1, \hat{x}_i) \in C_i$ for $i \in [1, \ell-1]$. Now consider an arbitrary $s \rightsquigarrow t$ -path π in G , where $\pi \neq \pi_{s,t}$. Since all edges in G are directed from L_i to L_{i+1} , it follows that π has $\ell-1$ edges and the i th edge of π is in E_i . Let $x_1, \dots, x_{\ell-1} \in \mathbb{R}$ be real numbers such that the i th edge of π has the corresponding vector $(1, x_i) \in C_i$ for $i \in [1, \ell-1]$. Now observe that since π and $\pi_{s,t}$ are both $s \rightsquigarrow t$ -paths, it follows that

$$\sum_{i=1}^{\ell-1} \hat{x}_i = \sum_{i=1}^{\ell-1} x_i.$$

Additionally, by our construction of $\pi_{s,t}$, it follows that

$$|\hat{x}_i - x| \leq |x_i - x|$$

for all $i \in [1, \ell-1]$. In particular, since $\pi \neq \pi_{s,t}$, there must be some $j \in [1, \ell-1]$ such that $\hat{x}_j \neq x_j$, and so by Proposition 1, $|\hat{x}_j - x| < |x_j - x|$ with probability 1. Then by Proposition 3,

$$w(\pi_{s,t}) = \sum_{e \in \pi_{s,t}} w(e) = \sum_{i=1}^{\ell-1} \hat{x}_i^2 < \sum_{i=1}^{\ell-1} x_i^2 = \sum_{e \in \pi} w(e) = w(\pi).$$

Path $\pi_{s,t}$ is a unique shortest $s \rightsquigarrow t$ -path in G , as desired. □

4.5 Critical Paths Intersection Properties

Before finishing our proof of Lemma 1, we will need to establish several properties of the critical paths in Π .

Proposition 4. *Let $\pi_1, \pi_2 \in \Pi$ be two critical paths with the same corresponding direction vector $\vec{d} \in D$. Then $\pi_1 \cap \pi_2 = \emptyset$.*

Proof. Let $v_i^j \in L_i$ denote the i th node of π_j , where $j \in \{1, 2\}$. Note that since π_1 and π_2 share the same direction vector $\vec{d}$, the edges (v_i^1, v_{i+1}^1) and (v_i^2, v_{i+1}^2) have the same corresponding vector $\vec{u}_i \in C_i$ for all $i \in [1, \ell - 1]$ by Proposition 2. By our construction of Π , for each node in the first layer $v \in L_1$, v belongs to at most one path $\pi \in \Pi$ with direction vector $\vec{d}$, so $v_1^1 \neq v_1^2$. Then for all $i \in [1, \ell]$,

$$v_i^1 = v_1^1 + \sum_{i=1}^{i-1} \vec{u}_i \neq v_1^2 + \sum_{i=1}^{i-1} \vec{u}_i = v_i^2. \quad \square$$

Let $\pi_1, \pi_2 \in \Pi$ be two critical paths, and let $v \in V$ be a node in G . We say that paths π_1 and π_2 *split* at v if $v \in \pi_1 \cap \pi_2$ and the node following v in π_1 is distinct from the node following v in π_2 , and we simply say that π_1 and π_2 *split* if there exists some $v \in V$ such that they split at v . Note that since $\pi_1, \pi_2 \in \Pi$ are unique shortest paths in G , paths π_1 and π_2 can split at most once.

Lemma 3. *Fix a node $v \in L_i$, where $i \in [1, \ell - 1]$, and let $\pi_1, \pi_2 \in \Pi$ be critical paths with direction vectors $(1, d_1), (1, d_2) \in D$ such that $v \in \pi_1$ and $v \in \pi_2$. Then paths π_1 and π_2 split at v with probability at least $\min\{|d_1 - d_2|, 1\}$.⁵*

Proof. By Proposition 4, $d_1 \neq d_2$, and assume wlog that $d_1 < d_2$. Let F be the event that the random variable ε_{i+1} was sampled so that

$$(\mathbb{Z} + \varepsilon_{i+1} + 1/2) \cap (d_1, d_2) \neq \emptyset,$$

(where $(d_1, d_2) \subseteq \mathbb{R}$ is the open interval with endpoints d_1 and d_2). Our proof strategy is to show that F implies that π_1, π_2 split at v , and then to show that F occurs with the claimed probability.

F implies that π_1, π_2 split at v . Assume that F occurs. By construction there is a nonnegative integer $c \in \mathbb{Z}$ such that $c + \varepsilon_{i+1} + 1/2$ is in the interval (d_1, d_2) . Since $(d_1, d_2) \subseteq (1, \frac{n}{4\ell^2})$, it follows that vectors $(1, c + \varepsilon_{i+1}), (1, c + \varepsilon_{i+1} + 1) \in \mathbb{R}^2$ are in C_i , because $0 \leq c \leq \frac{n}{4\ell^2} - 1$. More generally, by our choice of sets C and D there are vectors $(1, c_1), (1, c_2) \in C_i$ such that

$$c_1 \leq d_1 \leq c_1 + 1 \quad \text{and} \quad c_2 \leq d_2 \leq c_2 + 1.$$

Now we claim that

$$\operatorname{argmin}_{(1,x) \in C_i} |x - d_1| \neq \operatorname{argmin}_{(1,x) \in C_i} |x - d_2|.$$

To see this, suppose for the sake of contradiction that there is a vector $(1, y) \in C_i$ such that

$$(1, y) = \operatorname{argmin}_{(1,x) \in C_i} |x - d_1| = \operatorname{argmin}_{(1,x) \in C_i} |x - d_2|.$$

Then using our assumption that $d_1 < c + \varepsilon_{i+1} + 1/2 < d_2$, we obtain

$$y - d_1 \leq |y - d_1| \leq \min\{|c_1 - d_1|, |(c_1 + 1) - d_1|\} \leq 1/2 < (c + \varepsilon_{i+1} + 1) - d_1$$

and

$$d_2 - y \leq |y - d_2| \leq \min\{|c_2 - d_2|, |(c_2 + 1) - d_2|\} \leq 1/2 < d_2 - (c + \varepsilon_{i+1}).$$

⁵For the sake of completeness, let us be more precise here about the probability claim being made in this lemma. Consider any two paths π_1, π_2 , indexed by two start nodes and two direction vectors, and consider a node $v \in L_i$. The event that we generate π_1, π_2 in such a way that $v \in \pi_1 \cap \pi_2$ depends only on the random choices of $\varepsilon_1, \dots, \varepsilon_i$. If $v \in \pi_1 \cap \pi_2$, then the event that π_1, π_2 split at v depends only on the random choice of ε_{i+1} . The claim is: conditional on the event that $\{\varepsilon_1, \dots, \varepsilon_i\}$ are selected in such a way that $v \in \pi_1 \cap \pi_2$, the probability that ε_{i+1} is selected such that π_1, π_2 split at v is at least $\min\{|d_1 - d_2|, 1\}$.

Together, these two sequences of inequalities imply that $c + \varepsilon_{i+1} < y < c + \varepsilon_{i+1} + 1$. But this contradicts our assumption that $(1, y) \in C_i$, so we conclude that

$$\operatorname{argmin}_{(1,x) \in C_i} |x - d_1| \neq \operatorname{argmin}_{(1,x) \in C_i} |x - d_2|.$$

By Proposition 2, $\{\vec{u}_e \mid e \in E_i^v\} = C_i$, so we have also shown that

$$\operatorname{argmin}_{\{\vec{u}_e \mid e \in E_i^v\}} |u_e - d_1| \neq \operatorname{argmin}_{\{\vec{u}_e \mid e \in E_i^v\}} |u_e - d_2|.$$

Then π_1 and π_2 must split at v by our construction of the critical paths in Π .

F happens with good probability. Since ε_{i+1} is sampled uniformly at random from the interval $(0, 1)$, it follows that $(\mathbb{Z} + \varepsilon_{i+1} + 1/2) \cap (d_1, d_2) \neq \emptyset$ with probability at least $\min\{d_2 - d_1, 1\}$. $\square$

We will use Lemma 3 to prove the following two lemmas, which capture the key properties of our graph G .

Lemma 4 (Low path overlap). *Let $\pi_1, \pi_2 \in \Pi$ be critical paths with distinct associated direction vectors $(1, d_1), (1, d_2) \in D$. Then:*

- If $|d_1 - d_2| < 1$, then with probability at least $1 - n^{-8}$, we have⁶ $|\pi_1 \cap \pi_2| \leq \frac{8 \log n}{|d_1 - d_2|}$.
- If $|d_1 - d_2| \geq 1$, then $|\pi_1 \cap \pi_2| \leq 1$ (deterministically).

Proof. We begin with the first point; suppose $|d_1 - d_2| < 1$. Suppose we iteratively generate π_1, π_2 one layer at a time. Each time we choose a node v that lies in both π_1 and π_2 , by Lemma 3, π_1 and π_2 split at v with probability at least $|d_1 - d_2|$ (over the random choice of ε_{i+1}). Moreover, since π_1 and π_2 are unique shortest paths in G and G is acyclic, it follows that $\pi_1 \cap \pi_2$ is a contiguous subpath of π_1 and π_2 ; thus, once they split, they can no longer intersect in later layers. The number of nodes in the intersection $|\pi_1 \cap \pi_2|$ is 1 more than the number of consecutive nodes at which π_1, π_2 intersect but do not split. So by the above discussion, we have

$$\begin{aligned} \Pr \left[|\pi_1 \cap \pi_2| > \frac{8 \log n}{|d_1 - d_2|} \right] &\leq (1 - |d_1 - d_2|)^{\frac{8 \log n}{|d_1 - d_2|}} \\ &\leq e^{-|d_1 - d_2| \cdot \frac{8 \log n}{|d_1 - d_2|}} \\ &\leq e^{-8 \log n} \\ &= n^{-8} \end{aligned}$$

For the second point of the lemma: if $|d_1, d_2| \geq 1$, then by Lemma 3, if there is a node $v \in \pi_1 \cap \pi_2$, then π_1 and π_2 split at v with probability 1, and then they can no longer intersect in later layers. So we have $|\pi_1 \cap \pi_2| \leq 1$. $\square$

Since $|\Pi| = p \leq n^2$, we can argue by a union bound that Lemma 4 holds for all $\pi_1, \pi_2 \in \Pi$ simultaneously with probability at least $1 - n^{-4}$. From now on, we will assume that this property holds for our constructed graph G .

Once we specify our construction parameters ℓ and q , the following lemma will immediately imply the third property of G as stated in Lemma 1.

Lemma 5. *Let $x, y \in V$ be distinct nodes in G , and let $z = \text{hopdist}_G(x, y)$. Let*

$$\Pi' := \{\pi \in \Pi \mid x, y \in \pi\}.$$

Then $|\Pi'| \leq \max \left\{ \frac{16q \log n}{z}, 1 \right\}$.

⁶Formally, we consider any two paths π_1, π_2 indexed by two start nodes and direction vectors. When we iteratively generate these paths, the number of nodes in their intersection (possibly 0) depends only on the random choices of $\varepsilon_1, \dots, \varepsilon_\ell$. The probability claim in this lemma is with respect to these random choices.

Proof. Let $\Pi' = \{\pi_1, \dots, \pi_k\}$ and let $(1, d_i) \in D$ be the direction vector associated with π_i for $i \in [1, k]$. Since the paths in Π' all intersect, by Proposition 4 we must have $d_i \neq d_j$ for $i \neq j$. Let $a = \min_{i \in [1, k]} d_i$ and let $b = \max_{i \in [1, k]} d_i$. Then

$$d_b - d_a \geq \frac{k-1}{q},$$

since $|d_i - d_j| \geq 1/q$ for all $(1, d_i), (1, d_j) \in D$ such that $i \neq j$. Thus, by Lemma 4 we must have $\frac{k-1}{q} < 1$, since we have at least two nodes $x, y \in \pi_a \cap \pi_b$. So by Lemma 4,

$$|\pi_a \cap \pi_b| \leq \frac{8 \log n}{d_b - d_a} \leq \frac{8q \log n}{k-1}.$$

Since $x, y \in \pi_a \cap \pi_b$ and π_a and π_b are unique shortest paths in G , it follows that they coincide on their subpaths $\pi_a[x, y] = \pi_b[x, y]$. Moreover, since the hopdistance from x to y in G is z , it follows that $|\pi_a[x, y]| = |\pi_b[x, y]| = z + 1$. Then taken together we have

$$z \leq |\pi_a \cap \pi_b| \leq \frac{8q \log n}{k-1}.$$

Rearranging, we get

$$k \leq \frac{8q \log n}{z} + 1.$$

If $\frac{8q \log n}{z} \geq 1$, then this implies $k \leq \frac{16q \log n}{z}$. Otherwise, if $\frac{8q \log n}{z} < 1$, then this implies that $k \leq 1$ since $k \in \mathbb{Z}$. $\square$

4.6 Finishing the proof of Lemma 1

We note that Theorem 7 is trivial in the parameter regime $p = \Omega(n^2/\log n)$, since its lower bound on hopbound is $\Omega(1)$. So we will assume $p = O(n^2/\log n)$ in the following, with as small of an implicit constant as needed. Let

$$\ell = \frac{n}{2^{10} p^{1/2} \log^{1/2} n} \quad \text{and} \quad q = \frac{\ell}{2^{10} \log n}.$$

We now quickly verify that graph G and associated critical paths Π satisfy the properties of Lemma 1:

- By construction, G has $\ell = \Theta\left(\frac{n}{p^{1/2} \log^{1/2} n}\right)$ layers, and each path in Π travels from the first layer to the last layer.
- Each path $\pi \in \Pi$ is a unique shortest path between its endpoints in G by Lemma 2.
- Since $16q \log n \leq \ell$ and $\text{hopdist}_G(u, v) \leq \ell$, Lemma 5 immediately implies that for all $u, v \in V$, there are at most

$$\max \left\{ \frac{16q \log n}{\text{hopdist}_G(u, v)}, 1 \right\} \leq \frac{\ell}{\text{hopdist}_G(u, v)}$$

paths in Π that contain both u and v .

- For all critical pairs $(s, t) \in P$, the hopdistance from s to t in G is $\ell - 1 \gg 16q \log n$. Then by Lemma 5, each of the $|\Pi|$ paths constructed in Section 4.3 have distinct endpoints and thus are distinct. Then

$$|\Pi| = |S||D| \geq \frac{n}{2\ell} \cdot \frac{nq}{8\ell^2} \geq \frac{n^2 q}{16\ell^3} \geq p.$$

We have shown that our *directed* graph G satisfies the properties of Lemma 1 in the regime of $p \in [n, n^2]$. Moreover, our construction still goes through even in the extended regime of $p \in [n/c, n^2]$ for any constant $c > 0$. All that remains is to extend our construction to the entire regime of $p \in [1, n]$ and make G undirected.

Extending the construction to $p \in [1, n]$. We can extend our construction to the regime of $p \in [1, n]$ with a simple modification to G that was previously used in the prior work of [29]. We will sketch the modification here and defer the proof of correctness to Lemma 6 in Appendix A.

Let $G(n, p)$ denote an instance of our originally constructed graph G with input parameters n and $p \in [n, n^2]$. Let n be a sufficiently large integer and let $p \in [1, n]$. Let $G_1 := G(p_1, p)$ where $p_1 = \Theta(p)$ and p_1 divides n . Now for each node v in G_1 , replace v with a directed path $\pi_v = (u_1^v, \dots, u_k^v)$ with $k = n/p_1$ nodes. For all $v \in V$, assign weight 0 to all edges in π_v . For each edge (v_1, v_2) originally in G_1 , add edge $(u_k^{v_1}, u_1^{v_2})$ to the graph. Let G_2 be the resulting graph, and let Π_2 be the updated set of critical paths. This completes the modification.

Lemma 6. *The n -node graph G_2 and the set Π_2 of $|\Pi_2| \geq p$ paths satisfy the properties of Lemma 1.*

Proof. We defer the proof to Appendix A, as it largely follows our earlier analysis. $\square$

Making G undirected. To make G undirected, we use the following standard simple blackbox reduction. Let W be the sum of all edge weights in G , i.e., $W = \sum_{e \in E} w(e)$. For every edge $e \in E$, add $+W$ to the weight $w(e)$ of e , and treat e as an undirected edge. Call the resulting graph G' .

We now argue correctness: in particular, we need to argue that for all $s, t \in V$ such that t is reachable from s in G , π is a shortest weighted (directed) $s \rightsquigarrow t$ -path in G if and only if π is a shortest weighted (undirected) $s \rightsquigarrow t$ -path in G' .

- First, note that for all $s \rightsquigarrow t$ -paths π' in G' , the number of nodes in π' satisfies $|\pi'| \geq |\pi|$ by the construction of G and G' .
- Moreover, if $|\pi'| > |\pi|$, then π' has one more edge than π . Thus, its weighted length in G' satisfies

$$\begin{aligned} w(\pi') &> (|\pi| + 1) \cdot W \\ &= W|\pi| + \sum_{e \in E(G)} w(e) \\ &> w(\pi) \end{aligned}$$

and so π' is not a shortest path. We conclude that if π' is a shortest $s \rightsquigarrow t$ -path in G' , then $|\pi'| = |\pi|$.

- Any $s \rightsquigarrow t$ -path π' in G' with $|\pi|$ edges must use exactly one node in each layer, and thus it respects the original edge directions in G . We conclude that π' is a shortest weighted $s \rightsquigarrow t$ -path in G' if and only if π' is a shortest weighted (directed) $s \rightsquigarrow t$ -path in G .

Lemma 1 is immediate from the above discussion.

5 Shortcut Sets

In this section we will prove the following theorem.

Theorem 8. *For any parameter $p \in [1, n^{5/4}]$, there exists an n -node unweighted directed graph $G = (V, E)$ such that for any shortcut set H of size $|H| \leq p$, the graph $G \cup H$ must have diameter D , where*

$$D = \begin{cases} \Omega\left(\frac{n}{p^{3/4} \log n}\right) & \text{for } p \in [1, n/(\alpha \log n)], \text{ where } \alpha > 0 \text{ is a sufficiently large constant} \\ \Omega\left(\frac{n^{5/4}}{p \log^{9/4} n}\right) & \text{for } p \in [n/(\alpha \log n), n^{5/4}]. \end{cases}$$

In particular, when $p = O(n)$, $G \cup H$ must have diameter $D = \Omega\left(\frac{n^{1/4}}{\log^{9/4} n}\right)$.

We will prove this via a construction of the following type:

Lemma 7. *For any $p \in [1, n/(\alpha \log n)]$, where $\alpha > 0$ is a sufficiently large constant, there is an infinite family of n -node directed unweighted graphs $G = (V, E, w)$ and sets Π of $|\Pi| = p$ paths in G with the following properties:*

- G has $\ell = \Theta\left(\frac{n}{p^{3/4} \log n}\right)$ layers. Each path in Π starts in the first layer, ends in the last layer, and contains exactly one node in each layer.
- Each path in Π is the unique path between its endpoints in G .
- For any two nodes $u, v \in V$, there are at most $\frac{\ell}{\text{hopdist}_G(u,v)}$ paths in Π that contain both u and v .

We quickly verify that Lemma 7 implies Theorem 8.

Proof of Theorem 8 using Lemma 7. Fix an n and $p \in [1, n/(\alpha \log n)]$. Let $G = (V, E, w)$ be the graph in Lemma 7 with associated set Π of $|\Pi| = 2p$ paths in G . Let H be a shortcut set of size $|H| \leq p$. Let $P \subseteq V \times V$ be the set of node pairs that are the endpoints of paths in Π . Since all paths in Π are *unique* paths between their endpoints in G , it follows that

$$\text{diameter}(G \cup H) \geq \max_{(s,t) \in P} \text{dist}_{G \cup H}(s, t) = \max_{(s,t) \in P} \text{hopdist}_{G \cup H}(s, t).$$

Then when $p \in [1, n/(\alpha \log n)]$, we can achieve the bounds in Theorem 8 using the same potential function argument as in Section 4.1.

In order to extend our shortcut set lower bound to $p \in [n/(\alpha \log n), n^{5/4}]$, we will appeal to a more general property of the extremal functions of exact hopsets and shortcut sets, which we defer to Lemma 13 in Appendix B. $\square$

5.1 Constructing the strongly convex vector set $W(q)$

In our construction of the graph G , we will implicitly use the following lemma from [3, 8] that establishes the existence of a dense set of vectors that each extend the farthest in their own direction.⁷

Lemma 8 (Theorem 1 of [3]; Lemma 7 of [8]). *For sufficiently large $r \in \mathbb{Z}_{\geq 0}$, there exists a strongly convex set of integer vectors $W(r) \subseteq \mathbb{Z}^2$ of size $|W(r)| = \Theta(r^{2/3})$, such that*

- for all $\vec{w} \in W(r)$, $\|\vec{w}\| \leq r$,
- every $\vec{w} \in W(r)$ lies in the first quadrant, i.e., both coordinates are positive, and
- for all distinct $\vec{w}_1, \vec{w}_2 \in W(r)$, $\langle \vec{w}_1, \vec{w}_2 \rangle < \langle \vec{w}_1, \vec{w}_1 \rangle$.

In our construction of G , we will use a vector set W from this lemma to help generate edge and direction vectors. We will make use of the following technical property of the vectors in W :

Proposition 5. *Let $W = \{\vec{w}_1, \dots, \vec{w}_q\}$ be a set of vectors as described in Lemma 8, with its vectors ordered counterclockwise. For all $\vec{w}_i, \vec{w}_j, \vec{w}_k \in W$ with $i < j < k$, the following inequalities hold:*

$$\langle \vec{w}_i, \vec{w}_k \rangle < \langle \vec{w}_j, \vec{w}_k \rangle < \langle \vec{w}_k, \vec{w}_k \rangle \quad \text{and} \quad \langle \vec{w}_i, \vec{w}_k \rangle < \langle \vec{w}_i, \vec{w}_j \rangle < \langle \vec{w}_i, \vec{w}_i \rangle.$$

Proof. We will only prove here that $\langle \vec{w}_i, \vec{w}_k \rangle < \langle \vec{w}_j, \vec{w}_k \rangle < \langle \vec{w}_k, \vec{w}_k \rangle$; the other set of inequalities follows from an identical argument. By Lemma 8 we already have $\langle \vec{w}_j, \vec{w}_k \rangle < \langle \vec{w}_k, \vec{w}_k \rangle$, so it remains only to show that $\langle \vec{w}_i, \vec{w}_k \rangle < \langle \vec{w}_j, \vec{w}_k \rangle$.

Let ψ_1 be the inner angle between $\vec{w}_i$ and $\vec{w}_j$ and let ψ_2 be the inner angle between $\vec{w}_j$ and $\vec{w}_k$; thus the inner angle between $\vec{w}_i$ and $\vec{w}_k$ is $\psi_1 + \psi_2$. We first establish a useful inequality:

$$\begin{aligned} \langle \vec{w}_i, \vec{w}_j \rangle &< \langle \vec{w}_j, \vec{w}_j \rangle && \text{by Lemma 8} \\ \|\vec{w}_i\| \|\vec{w}_j\| \cos \psi_1 &< \|\vec{w}_j\|^2 && \text{cosine formula} \\ \|\vec{w}_i\| \|\vec{w}_k\| \cos \psi_1 \cos \psi_2 &< \|\vec{w}_j\| \|\vec{w}_k\| \cos \psi_2 && \text{follows algebraically from previous line} \end{aligned}$$

⁷This is a slightly stronger property than convexity, and hence is sometimes referred to as “strong convexity” in the area.

We will next use the trigonometric identity

$$\begin{aligned}\cos \psi_1 \cos \psi_2 &= \sin \psi_1 \sin \psi_2 + \cos(\psi_1 + \psi_2) \\ &> \cos(\psi_1 + \psi_2)\end{aligned}\quad \text{since } \psi_1, \psi_2 \in (0, \pi/2) \text{ by Lemma 8.}$$

We are now ready to show:

$$\begin{aligned}\langle \vec{w}_i, \vec{w}_k \rangle &= \|\vec{w}_i\| \|\vec{w}_k\| \cos(\psi_1 + \psi_2) && \text{cosine formula} \\ &< \|\vec{w}_i\| \|\vec{w}_k\| \cos \psi_1 \cos \psi_2 && \text{second inequality} \\ &< \|\vec{w}_j\| \|\vec{w}_k\| \cos \psi_2 && \text{first inequality} \\ &< \langle \vec{w}_j, \vec{w}_k \rangle. && \text{cosine formula } \square\end{aligned}$$

5.2 Constructing G

We next construct the graph that will be used for Lemma 7. Let n be a sufficiently large positive integer, and let $p = n/(\alpha_0 \log n)$ for a sufficiently large constant $\alpha_0 > 0$ to be chosen later (we will extend our construction to other choices of $p \in [1, n^{5/4}]$ later). For simplicity of presentation, we will frequently ignore issues related to non-integrality of expressions that arise in our construction; these issues affect our bounds only by lower-order terms. All edges (u, v) in G will be directed from u to v .

Vertex Set V .

- Let r be a positive integer construction parameter to be specified later. Our graph G will have $\ell := \frac{n^{1/3}}{r^{2/3}}$ layers $L_1, \dots, L_\ell$, and each layer will have $n/\ell = n^{2/3}r^{2/3}$ nodes.
- We will label each node in layer L_i , $i \in [1, \ell]$, with a distinct triple in $\{i\} \times [1, n^{1/3}r^{1/3}] \times [1, n^{1/3}r^{1/3}]$. We will interpret the node in L_i labeled (i, j, k) as an integer point $(i, j, k) \in \mathbb{Z}^3$.
- These n nodes arranged in ℓ layers will compose our node set

$$V = [1, \ell] \times [1, n^{1/3}r^{1/3}] \times [1, n^{1/3}r^{1/3}] \subseteq \mathbb{Z}^3$$

of graph G . In a slight abuse of notation, we will treat $v \in V$ as either a node in G or a point in $\mathbb{Z}^3$, depending on the context.

Edge Set E .

- All the edges in E will be between consecutive layers L_i, L_{i+1} of G . We will let E_i denote the set of edges in G between layers L_i and L_{i+1} .
- Just as our nodes in V correspond to points in $\mathbb{R}^3$, we can interpret the edges E in G as vectors in $\mathbb{R}^3$. In particular, for every edge $e = (v_1, v_2) \in E$, we identify e with the corresponding vector $\vec{u}_e := v_2 - v_1$. Note that since all edges in E are between adjacent layers L_i and L_{i+1} , the first coordinate of $\vec{u}_e$ is 1 for all $e \in E$. We will use u_e^i to denote the i th coordinate of $\vec{u}_e$ for $i \in \{2, 3\}$ (i.e. for all $e \in E$, we write $\vec{u}_e = (1, u_e^2, u_e^3)$).
- We begin our construction of E by defining the set of vectors $W := W(r^3) \subseteq \mathbb{Z}^2$, where $W(r^3)$ is the strongly convex set of vectors defined in Lemma 8. Let $W = \{\vec{w}_1, \dots, \vec{w}_q\}$, where vectors $\vec{w}_i$ are ordered counterclockwise and $q := |W| = \Theta(r^2)$. We may assume wlog that $q \leq r^2$ (e.g. by removing vectors from W until this is true).
- Using W , we define our set of *edge vectors* $C \subseteq \mathbb{Z}^2$ as:

$$C := \{\vec{w}_{i+1} - \vec{w}_i \mid i \in [1, q-1]\}.$$

Let $\vec{c}_i := \vec{w}_{i+1} - \vec{w}_i$ denote the i th vector of C .

- For each layer L_i , $i \in [1, \ell - 1]$, sample a random integer from $[1, q - 1]$ and call it λ_i . We define the set $C_i \subseteq \mathbb{Z}^2$ as

$$C_i := \{(0, 0), \vec{c}_{\lambda_i}\},$$

where $\vec{c}_{\lambda_i}$ is the λ_i th vector of C . Note that C_i contains exactly two vectors, the zero vector and a randomly chosen vector $\vec{c}_{\lambda_i}$ from C . Intuitively: for each layer L_i , we are sampling two adjacent vectors $\vec{w}_{\lambda_i}$ and $\vec{w}_{\lambda_i+1}$ from C and adding $-\vec{w}_{\lambda_i}$ to each of them to obtain C_i . The purpose of adding the normalizing vector $-\vec{w}_{\lambda_i}$ and $\vec{w}_{\lambda_i+1}$ is to reduce the magnitude of the vectors in C_i , as we will formalize in Proposition 6.

- The vectors in set C_i will define the edges in E_i . Specifically, for all $(i, v_1, v_2) \in L_i$ and for all $(c_1, c_2) \in C_i$ such that

$$(i + 1, v_1 + c_1, v_2 + c_2) \in L_{i+1},$$

we add the edge $((i, v_1, v_2), (i + 1, v_1 + c_1, v_2 + c_2))$ to E_i .

This completes the construction of our graph $G = (V, E)$. We now verify that the vectors in C_i have small magnitude in expectation.

Proposition 6. *For all $i \in [1, \ell - 1]$, $\mathbb{E}[\|\vec{c}_{\lambda_i}\|] \leq \frac{2^4 r^3}{q} = \Theta(r)$.*

Proof. Note that the vectors in C correspond to sides of a convex polygon whose vertices are the vectors in W . Since this polygon is contained in a ball of radius r^3 in $\mathbb{R}_{\geq 0}^2$ by Lemma 8, it follows that $\sum_{\vec{c} \in C} \|\vec{c}\| \leq 2\pi r^3$. Note that $|C| = q - 1 \geq q/2$ for sufficiently large q . Then

$$\mathbb{E}[\|\vec{c}_{\lambda_i}\|] = \sum_{\vec{c} \in C} \frac{1}{|C|} \cdot \|\vec{c}\| \leq \frac{2\pi r^3}{|C|} \leq \frac{2^4 r^3}{q}. \quad \square$$

The vectors $\vec{w}_i \in W$ have magnitude roughly r^3 , whereas the vectors in C_i have expected magnitude at most $\Theta(r)$ by Proposition 6. Since each edge in our critical path corresponds to a vector in C_i , ensuring that the vectors in C_i have small magnitude (at least in expectation) will be essential for guaranteeing that the paths $\pi \in \Pi$ are long.

Let us comment here on a discrepancy between this construction and the technical overview. In the technical overview, we stated that we would use the convex vectors W to generate the edge vectors C . Instead, we are using the *difference* between adjacent convex vectors W to generate C . This is an optimization: our plan is to sample *one* edge vector from C , and use it together with the zero vector $(0, 0)$ as the two available vectors between pairs of adjacent layers. This is equivalent to sampling *two adjacent* edge vectors from W , as advertised in the technical overview, and then applying an appropriate translation of the next layer in space. Our strategy lets us use vectors of length $\Theta(r)$, instead of $\Theta(r^3)$, and these shorter vectors ultimately lead to a stronger lower bound.

5.3 Direction Vectors, Critical Pairs, and Critical Paths

Our next step is to generate a set of *critical pairs* $P \subseteq V \times V$, as well as a set of *critical paths* Π . Specifically, there will be one critical path $\pi_{s,t} \in \Pi$ going between each critical pair $(s, t) \in P$, and we will show that $\pi_{s,t}$ is the unique $s \rightsquigarrow t$ path in G .

Direction Vectors D . We choose our set of direction vectors D to be $D := W$, where W is our strongly convex set of $q = \Theta(r^2)$ vectors. We will let $D = \{\vec{d}_1, \dots, \vec{d}_q\}$ be our list of direction vectors, and we will let the i th vector $\vec{d}_i$ of D correspond to the i th vector $\vec{w}_i$ of W , i.e. $\vec{d}_i := \vec{w}_i$, for $i \in [1, q]$. We will simply use the name $\vec{d}_i$ when we wish to emphasize the role of these vectors as direction vectors.

Note that since $D = W$, Proposition 5 also holds for D . That is, if $i, j, k \in [1, q - 1]$ and $i < j < k$, then

$$\langle \vec{d}_i, \vec{d}_k \rangle < \langle \vec{d}_j, \vec{d}_k \rangle < \langle \vec{d}_k, \vec{d}_k \rangle \quad \text{and} \quad \langle \vec{d}_i, \vec{d}_k \rangle < \langle \vec{d}_i, \vec{d}_j \rangle < \langle \vec{d}_i, \vec{d}_i \rangle.$$

Critical Pairs P and Critical Paths Π .

- We first define a set $S \subseteq L_1$, containing a subset of the nodes in the first layer L_1 of G :

$$S := \left\{ (1, j, k) \mid (j, k) \in \left[\frac{1}{3}n^{1/3}r^{1/3}, \frac{2}{3}n^{1/3}r^{1/3} \right] \times \left[\frac{1}{3}n^{1/3}r^{1/3}, \frac{2}{3}n^{1/3}r^{1/3} \right] \right\}.$$

Informally, S is a middle square patch of the nodes in L_1 . The key property of S is that all nodes in S are of distance at least $\frac{1}{3}n^{1/3}r^{2/3}$ from the sides of the square grid corresponding to layer L_1 .

We will choose our set of demand pairs P so that $P \subseteq S \times V$. For every node $s \in S$ and direction vector $\vec{d} \in D$, we will choose a critical pair $(s, t) \in S \times V$ and a corresponding critical path $\pi_{s,t}$ to add to P and Π .

- Let $v_1 \in S$, and let $\vec{d} = (1, d) \in D$. The associated path π has start node v_1 . We iteratively grow π , layer-by-layer, as follows. Suppose that currently $\pi = (v_1, \dots, v_i)$ with each $v_i \in L_i$. To determine the next node $v_{i+1} \in L_{i+1}$, let $E_i^{v_i} \subseteq E_i$ be the edges in E_i incident to v_i , and let $\vec{u}_i = (u_i^1, u_i^2) \in C_i$ be

$$\vec{u}_i := \operatorname{argmax}_{\vec{c} \in C_i} \langle \vec{c}, \vec{d} \rangle.$$

If $(1, u_i^1, u_i^2) = \vec{u}_{e_i}$ for some $e_i \in E_i^{v_i}$, then by definition, e_i is an edge whose first node is v_i ; we define $v_{i+1} \in L_{i+1}$ to be the other node in e_i , and we append v_{i+1} to π . Otherwise, if there is no such edge e_i in $E_i^{v_i}$, then we terminate our construction of path π (i.e. v_i will be the final node in π).

This completes our construction of P and Π . We will show that the paths generated in this way have distinct endpoints (with high probability), and therefore $|P| = |S||D| \geq \frac{n^{2/3}r^{2/3}q}{2^5}$, where $q = \Theta(r^2)$.

An important feature for correctness of our construction is that, when we iteratively generate paths, if we reach a point v_i such that $v_i + C_i \not\subseteq L_{i+1}$ (i.e. $v_i + \vec{c} \notin L_{i+1}$ for some $\vec{c} \in C_i$), then we end our path at v_i . As a consequence, our critical paths in Π may not travel through all ℓ layers of G . However, with nonzero probability, paths in Π travel through a constant fraction of layers, as we prove in the following proposition.

Proposition 7. *Let $\hat{\ell} := \ell \cdot \frac{q}{2^4 r^2} = \Theta(\ell)$. With probability at least $1/2$, for all $\pi \in \Pi$, $|\pi| \geq \hat{\ell}$.*

Proof. Each critical path starts at a node s in S . Each edge of the path corresponds to a vector $(1, c_1, c_2) \in \mathbb{Z}^3$ such that $(c_1, c_2) \in C_i$ for some $i \in [1, \ell - 1]$. Our path ends when we reach the boundary of our vertex set $V = [1, \ell] \times [1, n^{1/3}r^{1/3}] \times [1, n^{1/3}r^{1/3}] \subseteq \mathbb{Z}^3$. We must show that before any such path travels through $\hat{\ell}$ nodes before reaching the boundary. Note that $\hat{\ell} = \ell \cdot \frac{q}{2^4 r^2} \leq \ell$, since $q \leq r^2$.

Let x be the random variable defined as $x := \sum_{i=1}^{\hat{\ell}} \vec{c}_{\lambda_i}$. Then by Proposition 6 and the linearity of expectation,

$$\mathbb{E}[x] \leq \hat{\ell} \cdot \frac{2^4 r^3}{q} = \ell \cdot \frac{r}{2^3} = \frac{1}{2^3} n^{1/3} r^{1/3},$$

where the final equality follows from the fact that $\ell = n^{1/3} r^{-2/3}$. Then by Markov's inequality,

$$\Pr \left[x \leq \frac{1}{4} n^{1/3} r^{1/3} \right] \geq 1/2.$$

Now we claim that if $x \leq \frac{1}{4} n^{1/3} r^{1/3}$, then for all $\pi \in \Pi$, $|\pi| \geq \hat{\ell}$. Let $\pi_{s,t} \in \Pi$ be a critical path for critical pair $(s, t) \in P$. Let $s = (s_1, s_2, s_3) \in \mathbb{Z}^3$ and let $t = (t_1, t_2, t_3) \in \mathbb{Z}^3$. By our construction of critical paths Π , either $t_1 = \ell$ or $t + C_{t_1} \not\subseteq V$. In the first case, $|\pi| = \ell \geq \hat{\ell}$, since path π traveled through all ℓ layers L_i . In the second case, we must have that $\|(t_2, t_3) - (s_2, s_3)\| \geq \frac{1}{3} n^{1/3} r^{1/3}$ by our choice of set S . But since $x \leq \frac{1}{4} n^{1/3} r^{1/3}$, we conclude that $|\pi| \geq \hat{\ell}$. The claim follows. $\square$

We have shown that with nonzero probability, all our paths in Π travel through the first $\hat{\ell}$ layers of G . However, we cannot guarantee that paths in Π travel to layers L_i with $i > \hat{\ell}$. Because of this, we choose to remove all layers L_i , $i > \hat{\ell}$, from G . We replace all critical paths $\pi \in \Pi$ with the truncated subpath of π containing only the first $\hat{\ell}$ nodes of π , and we update our critical pairs $P \subseteq V \times V$ to be the set of all pairs of endpoints of the updated paths in Π .

5.4 Critical paths are unique paths

We now verify that graph G and paths Π have the unique path property as stated in Lemma 7. This will follow straightforwardly from the properties of our set of direction vectors D , particularly Proposition 5.

Lemma 9 (Unique paths). *For every $(s, t) \in P$, path $\pi_{s,t}$ is a unique $s \rightsquigarrow t$ -path in G .*

Proof. Fix a direction vector $\vec{d}_j$ in D . We claim that for all $i \in [1, \ell - 1]$, there is a unique vector $\vec{c} \in C_i$ such that maximizes $\langle \vec{c}, \vec{d}_j \rangle$. Recall that $C_i \subseteq \mathbb{Z}^2$ contains exactly two vectors: the zero vector $(0, 0)$ and the vector $\vec{c}_{\lambda_i} = \vec{w}_{\lambda_i+1} - \vec{w}_{\lambda_i}$.

Now assume that $j \geq \lambda_i + 1$ and observe the following sequence of equivalent inequalities:

$$\begin{aligned} \langle \vec{w}_{\lambda_i}, \vec{w}_j \rangle &< \langle \vec{w}_{\lambda_i+1}, \vec{w}_j \rangle && \text{by Lemma 8 and Proposition 5} \\ \langle \vec{0}, \vec{w}_j \rangle &< \langle \vec{w}_{\lambda_i+1} - \vec{w}_{\lambda_i}, \vec{w}_j \rangle \\ \langle \vec{0}, \vec{d}_j \rangle &< \langle \vec{c}_{\lambda_i+1}, \vec{d}_j \rangle. \end{aligned}$$

When $j < \lambda_i + 1$, an identical argument shows that $\langle \vec{0}, \vec{d}_j \rangle > \langle \vec{c}_{\lambda_i+1}, \vec{d}_j \rangle$. In either case, there is a unique vector $\vec{c} \in C_i$ maximizing $\langle \vec{c}, \vec{d}_j \rangle$ for all $i \in [1, \ell - 1]$.

Now fix a critical pair $(s, t) \in P$ that has $\vec{d}_j \in D$ as its associated direction vector and $\pi_{s,t} = (v_1, \dots, v_k)$ as its critical path (where $s = v_1$ and $t = v_k$). Let $f : \mathbb{R}^3 \mapsto \mathbb{R}^2$ be the function that projects each point in $\mathbb{R}^3$ onto the subspace corresponding to the last two coordinates of $\mathbb{R}^3$, i.e. $f(x, y, z) = (y, z)$ for all $(x, y, z) \in \mathbb{R}^3$.

Let π be an arbitrary $s \rightsquigarrow t$ -path, and note that $|\pi| = |\pi_{s,t}|$ since G is a layered directed graph. Let $\pi = (v'_1, \dots, v'_k)$, where $s = v'_1$ and $t = v'_k$. By our construction of path $\pi_{s,t}$, we must have that for all $i \in [1, k - 1]$,

$$\langle f(v'_{i+1} - v'_i), \vec{d}_j \rangle \leq \langle f(v_{i+1} - v_i), \vec{d}_j \rangle.$$

Now suppose for the sake of contradiction that $\pi \neq \pi_{s,t}$. Then $v'_{i+1} - v'_i \neq v_{i+1} - v_i$ for some $i \in [1, k - 1]$. Then $\langle f(v'_{i+1} - v'_i), \vec{d}_j \rangle < \langle f(v_{i+1} - v_i), \vec{d}_j \rangle$ by the above discussion. But then since π and $\pi_{s,t}$ are both $s \rightsquigarrow t$ -paths,

$$\langle f(t - s), \vec{d}_j \rangle = \sum_{i=1}^{k-1} \langle f(v'_{i+1} - v'_i), \vec{d}_j \rangle < \sum_{i=1}^{k-1} \langle f(v_{i+1} - v_i), \vec{d}_j \rangle = \langle f(t - s), \vec{d}_j \rangle.$$

This is a contradiction, so we conclude that $\pi = \pi_{s,t}$. Then path $\pi_{s,t}$ is a unique $s \rightsquigarrow t$ -path in G . □

5.5 Critical Paths Intersection Properties

Before finishing our proof of Lemma 7, we will need to establish several properties of the critical paths in Π .

Proposition 8. *Let $\pi_1, \pi_2 \in \Pi$ be two critical paths with the same corresponding direction vector $\vec{d} \in D$. Then $\pi_1 \cap \pi_2 = \emptyset$.*

Proof. Let $k = \min\{|\pi_1|, |\pi_2|\}$. Let $v_i^j \in L_i$ denote the i th node of π_j , where $j \in \{1, 2\}$ and $i \in [1, k]$. Note that since π_1 and π_2 share the same direction vector $\vec{d}$, edges (v_i^1, v_{i+1}^1) and (v_i^2, v_{i+1}^2) have the same corresponding vector $\vec{u}_i \in C_i$ for all $i \in [1, k - 1]$ by our construction of π_1 and π_2 . Also, for each node $v \in L_1$, v belongs to at most one path $\pi \in \Pi$ with direction vector $\vec{d}$, so $v_1^1 \neq v_1^2$. Then for all $i \in [1, k]$,

$$v_i^1 = v_1^1 + \sum_{i=1}^{i-1} \vec{u}_i \neq v_1^2 + \sum_{i=1}^{i-1} \vec{u}_i = v_i^2. \quad \square$$

Let $\pi_1, \pi_2 \in \Pi$ be two critical paths, and let $v \in V$ be a node in G . We say that paths π_1 and π_2 *split* at v if $v \in \pi_1 \cap \pi_2$ and the node following v in π_1 is distinct from the node following v in π_2 , and we simply say that π_1 and π_2 *split* if there exists some $v \in V$ such that π_1 and π_2 split at v . Note that since $\pi_1, \pi_2 \in \Pi$ are unique paths in G , paths π_1 and π_2 can split at most once.

Lemma 10. Fix a node $v \in L_i$, where $i \in [1, \ell - 1]$, and let $\pi_1, \pi_2 \in \Pi$ be critical paths with direction vectors $\vec{d}_j$ and $\vec{d}_k \in D$, $j, k \in [1, q]$, such that $v \in \pi_1$ and $v \in \pi_2$. Then paths π_1 and π_2 split at v with probability at least $\frac{|j-k|}{q}$.

Proof. Fix a node $v \in L_i$, where $i \in [1, \ell - 1]$, and let $\pi_1, \pi_2 \in \Pi$ be critical paths with direction vectors $\vec{d}_j$ and $\vec{d}_k \in D$, $j, k \in [1, q]$, such that $v \in \pi_1$ and $v \in \pi_2$. By Proposition 8, $j \neq k$, and assume wlog that $j < k$. Let F be the event that the random variable λ_i was sampled so that

$$j \leq \lambda_i < k.$$

Our proof strategy is to show that F implies that π_1, π_2 split at v , and then to show that F occurs with the claimed probability.

F implies that π_1, π_2 split at v . Assume that F occurs. Then $j \leq \lambda_i < \lambda_i + 1 \leq k$. Now observe the following sequence of equivalent inequalities:

$$\begin{aligned} \langle \vec{w}_j, \vec{w}_{\lambda_i} \rangle &> \langle \vec{w}_j, \vec{w}_{\lambda_i+1} \rangle && \text{by Lemma 8 and Proposition 5} \\ \langle \vec{w}_j, \vec{0} \rangle &> \langle \vec{w}_j, \vec{w}_{\lambda_i+1} - \vec{w}_{\lambda_i} \rangle \\ \langle \vec{d}_j, \vec{0} \rangle &> \langle \vec{d}_j, \vec{c}_{\lambda_i} \rangle. \end{aligned}$$

Since $C_i = \{\vec{0}, \vec{c}_{\lambda_i}\}$, by our construction of the critical paths in Π , the above inequality $\langle \vec{d}_j, \vec{0} \rangle > \langle \vec{d}_j, \vec{c}_{\lambda_i} \rangle$ implies that path π_1 takes an edge in E_i corresponding to vector $\vec{0}$. An identical argument will show that $\langle \vec{d}_k, \vec{0} \rangle < \langle \vec{d}_k, \vec{c}_{\lambda_i} \rangle$, so path π_2 takes an edge in E_i corresponding to vector $\vec{c}_{\lambda_i}$. Since paths π_1 and π_2 take different edges in E_i , they must split at v .

F happens with good probability. Random variable λ_i is sampled uniformly from $[1, q - 1]$. Then the event F occurs with probability $\frac{|j-k|}{q-1} \geq \frac{|j-k|}{q}$. $\square$

We will use Lemma 11 to prove the following two lemmas, which capture key properties of our graph G .

Lemma 11. Let $\pi_1, \pi_2 \in \Pi$ be critical paths with associated direction vectors $\vec{d}_j, \vec{d}_k \in D$. Then $|\pi_1 \cap \pi_2| \leq \frac{8q \log n}{|j-k|}$ with probability at least $1 - n^{-8}$.

Proof. If $\pi_1 \cap \pi_2 = \emptyset$, then the claim is immediate, so assume there is a node $v \in \pi_1 \cap \pi_2$. Suppose $v \in L_i$. By Lemma 10, π_1 and π_2 split at v with probability at least $\frac{|j-k|}{q}$. Moreover, conditioning on $v \in \pi_1 \cap \pi_2$, the event that π_1 and π_2 split at v given that $v \in \pi_1$ and $v \in \pi_2$ depends only on our choice of λ_i and is independent of λ_j for $j \neq i$.

Since π_1 and π_2 are unique paths in G , it follows that $\pi_1 \cap \pi_2$ is a contiguous subpath of π_1 and π_2 . The number of nodes in the intersection $|\pi_1 \cap \pi_2|$ is 1 more than the number of consecutive nodes at which π_1, π_2 intersect but do not split. So by the above discussion, we have

$$\begin{aligned} \Pr \left[|\pi_1 \cap \pi_2| > \frac{8q \log n}{|j-k|} \right] &\leq \left(1 - \frac{|j-k|}{q} \right)^{\frac{8q \log n}{|j-k|}} \\ &\leq e^{-\frac{|j-k|}{q} \cdot \frac{8q \log n}{|j-k|}} \\ &\leq e^{-8 \log n} \\ &= n^{-8}. \end{aligned}$$

$\square$

Since $|\Pi| = p \leq n^2$, we can argue by a union bound that Proposition 7 holds and Lemma 4 holds for all $\pi_1, \pi_2 \in \Pi$ simultaneously with probability at least $1 - 1/2 - n^{-4} > 0$. From now on, we will assume that this property holds for our constructed graph G .

Once we specify our construction parameters ℓ and q , the following lemma will immediately imply the third property of G as stated in Lemma 7.

Lemma 12. Let $x, y \in V$ be nodes in G such that the unweighted distance from x to y in G is z , where $0 < z < \infty$. Let $\Pi' \subseteq \Pi$ be the following set of critical paths:

$$\Pi' := \{\pi \in \Pi \mid x, y \in \pi\}.$$

Then $|\Pi'| \leq \max\left\{\frac{16q \log n}{z}, 1\right\}$.

Proof. Let $\Pi' = \{\pi_1, \dots, \pi_k\}$ and let $\vec{d}_{\sigma_i} \in D$ be the direction vector associated with π_i for $i \in [1, k]$. By Proposition 4, $\sigma_i \neq \sigma_j$ for $i \neq j$. Let $a = \min_{i \in [1, k]} \sigma_i$ and let $b = \max_{i \in [1, k]} \sigma_i$. Then $b - a \geq k - 1$, so by Lemma 11,

$$|\pi_a \cap \pi_b| \leq \frac{8q \log n}{|b - a|} \leq \frac{8q \log n}{k - 1}.$$

Additionally, since $x, y \in \pi_a \cap \pi_b$ and π_a and π_b are unique paths, it follows that $\pi_a[x, y] = \pi_b[x, y]$. Moreover, since the unweighted distance from x to y in G is z , it follows that $|\pi_a \cap \pi_b| \geq |\pi_a[x, y]| = |\pi_b[x, y]| = z + 1$. Then taken together we have

$$z \leq |\pi_a \cap \pi_b| \leq \frac{8q \log n}{k - 1}.$$

Rearranging, we get that

$$k \leq \frac{8q \log n}{z} + 1 \leq \max\left\{\frac{16q \log n}{z}, 1\right\}.$$

(We obtain the final inequality from the following observation: if $\frac{8q \log n}{z} \geq 1$, then this implies $k \leq \frac{16q \log n}{z}$. Otherwise, if $\frac{8q \log n}{z} < 1$, then this implies that $k \leq 1$ since $k \in \mathbb{Z}$.) $\square$

5.6 Finishing the proof of Lemma 7

Let

$$r = \frac{n^{1/8}}{2^6 \log^{3/8} n}.$$

Now recall that $q = \Theta(r^2)$, and let $\alpha_1 > 0$ be a constant such that $\alpha_1 r^2 \leq q \leq 2\alpha_1 r^2$. Then

$$\ell = \frac{n^{1/3}}{r^{2/3}} \geq n^{1/4} \log^{1/4} n, \quad \hat{\ell} = \frac{q}{2^7 r^2} \cdot \ell \geq \frac{\alpha_1}{2^7} \cdot \ell, \quad \text{and} \quad q \leq 2\alpha_1 r^2 \leq \frac{\hat{\ell}}{16 \log n}.$$

We now quickly verify that graph G and associated critical paths Π satisfy the properties of Lemma 7:

- By construction, G has $\hat{\ell} = \Theta\left(n^{1/4} \log^{1/4} n\right)$ layers, and each path in Π travels from the first layer to the $\hat{\ell}$ th layer.
- Each path $\pi \in \Pi$ is a unique path between its endpoints in G by Lemma 9.
- Since $16q \log n \leq \ell$ and $\text{hopdist}_G(u, v) \leq \ell$, Lemma 12 immediately implies that for all $u, v \in V$, there are at most

$$\max\left\{\frac{16q \log n}{\text{hopdist}_G(u, v)}, 1\right\} \leq \frac{\ell}{\text{hopdist}_G(u, v)}$$

paths in Π that contain both u and v .

- For all critical pairs $(s, t) \in P$, the unweighted distance from s to t in G is $\ell - 1 \gg 16q \log n$. Then by Lemma 12, each of the $|\Pi|$ paths constructed in Section 5.3 have distinct endpoints and thus are distinct. Then

$$|\Pi| = |S||D| \geq \frac{n^{2/3} r^{2/3} q}{2^5} \geq \frac{\alpha_1 n^{2/3} r^{8/3}}{2^5} \geq \frac{\alpha_1 n}{2^{21} \log n}.$$

Recall that we initially let $p = n/(\alpha_0 \log n)$, for some unspecified constant $\alpha_0 > 0$. Choose α_0 so that $p \leq |\Pi|$.

We have shown that our directed graph G satisfies the properties of Lemma 7 when $p = n/(\alpha_0 \log n)$. All that remains is to extend our construction to the regime $p \in [1, n/(\alpha_0 \log n)]$.

Extending the construction to $p \in [1, n/(\alpha_0 \log n)]$. We can extend our construction to the regime of $p \in [1, n/(\alpha_0 \log n)]$ using the same modification to G that we performed on our exact hopset construction and that was previously used in the prior work of [29]. We will sketch the modification here. The proof of correctness follows from an argument identical to the proof of Lemma 6 in Appendix A.

We use $G(n, p)$ denote an instance of our originally constructed graph G with input parameters n and $p = n/(\alpha_0 \log n)$. Let n be a sufficiently large integer and let $p \in [1, n]$. Let $G_1 := G(p_1, p_1/(\alpha_0 \log p_1))$ where $p = \Theta(p_1/\log p_1)$ and p_1 divides n . Now for each node v in G_1 , replace v with a directed path $\pi_v = (u_1^v, \dots, u_k^v)$ with $k = n/p_1$ nodes. For all $v \in V$, assign weight 0 to all edges in π_v . For each edge (v_1, v_2) originally in G_1 , add edge $(u_k^{v_1}, u_1^{v_2})$ to the graph. Let G_2 be the resulting graph, and let Π_2 be the updated set of critical paths. This completes the modification.

References

- [1] Amir Abboud, Greg Bodwin, and Seth Pettie. A hierarchy of lower bounds for sublinear additive spanners. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 568–576. Society for Industrial and Applied Mathematics, 2017.
- [2] 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*, pages 322–335, 2020.
- [3] Imre Bárány and David G Larman. The convex hull of the integer points in a large ball. *Mathematische Annalen*, 312(1):167–181, 1998.
- [4] Uri Ben-Levy and Merav Parter. New (α, β) spanners and hopsets. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1695–1714. SIAM, 2020.
- [5] Piotr Berman, Sofya Raskhodnikova, and Ge Ruan. Finding sparser directed spanners. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2010)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2010.
- [6] Aaron Bernstein, Maximilian Probst Gutenberg, and Christian Wulff-Nilsen. Near-optimal decremental sssp in dense weighted digraphs. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1112–1122. IEEE, 2020.
- [7] Aaron Bernstein and Nicole Wein. Closing the gap between directed hopsets and shortcut sets. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 163–182. SIAM, 2023.
- [8] Greg Bodwin and Gary Hoppenworth. New additive spanner lower bounds by an unlayered obstacle product. 2022.
- [9] Greg Bodwin, Gary Hoppenworth, and Ohad Trabelsi. Bridge girth: A unifying notion in network design. *arXiv preprint arXiv:2212.11944*, 2022.
- [10] 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*, pages 336–349, 2020.
- [11] Edith Cohen. Polylog-time and near-linear work approximation scheme for undirected shortest paths. *Journal of the ACM (JACM)*, 47(1):132–166, 2000.
- [12] Don Coppersmith and Michael Elkin. Sparse sourcewise and pairwise distance preservers. *SIAM Journal on Discrete Mathematics*, 20(2):463–501, 2006.
- [13] Michael Elkin, Yuval Gitlitz, and Ofer Neiman. Almost shortest paths with near-additive error in weighted graphs. *arXiv preprint arXiv:1907.11422*, 2019.

- [14] Michael Elkin and Ofer Neiman. Linear-size hopsets with small hopbound, and constant-hopbound hopsets in rnc. In *The 31st ACM Symposium on Parallelism in Algorithms and Architectures*, pages 333–341, 2019.
- [15] Jeremy T Fineman. Nearly work-efficient parallel algorithm for digraph reachability. *SIAM Journal on Computing*, 49(5):STOC18–500, 2019.
- [16] Sebastian Forster and Danupon Nanongkai. A faster distributed single-source shortest paths algorithm. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 686–697. IEEE, 2018.
- [17] Maximilian Probst Gutenberg and Christian Wulff-Nilsen. Decremental sssp in weighted digraphs: Faster and against an adaptive adversary. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2542–2561. SIAM, 2020.
- [18] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Decremental single-source shortest paths on undirected graphs in near-linear total update time. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 146–155. IEEE, 2014.
- [19] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Sublinear-time decremental algorithms for single-source reachability and shortest paths on directed graphs. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 674–683, 2014.
- [20] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Improved algorithms for decremental single-source reachability on directed graphs. In *Automata, Languages, and Programming: 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I 42*, pages 725–736. Springer, 2015.
- [21] William Hesse. Directed graphs requiring large numbers of shortcuts. In *Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 665–669. Society for Industrial and Applied Mathematics, 2003.
- [22] Shang-En Huang and Seth Pettie. Lower Bounds on Sparse Spanners, Emulators, and Diameter-reducing shortcuts. In David Eppstein, editor, *16th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2018)*, volume 101 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 26:1–26:12, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [23] Shang-En Huang and Seth Pettie. Thorup–zwick emulators are universally optimal hopsets. *Information Processing Letters*, 142:9–13, 2019.
- [24] Arun Jambulapati, Yang P Liu, and Aaron Sidford. Parallel reachability in almost linear work and square root depth. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1664–1686. IEEE, 2019.
- [25] 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. SIAM, 2021.
- [26] Philip N Klein and Sairam Subramanian. A randomized parallel algorithm for single-source shortest paths. *Journal of Algorithms*, 25(2):205–220, 1997.
- [27] Shimon Kogan and Merav Parter. Beating Matrix Multiplication for $n^{1/3}$ -Directed Shortcuts. In Miłkołaj Bojańczyk, Emanuela Merelli, and David P. Woodruff, editors, *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022)*, volume 229 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 82:1–82:20, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [28] Shimon Kogan and Merav Parter. Having hope in hops: New spanners, preservers and lower bounds for hopsets. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 766–777. IEEE, 2022.

- [29] Shimon Kogan and Merav Parter. New diameter-reducing shortcuts and directed hopsets: Breaking the barrier. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1326–1341. SIAM, 2022.
- [30] 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 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3311–3331. SIAM, 2022.
- [31] Ofer Neiman and Idan Shabat. A unified framework for hopsets. In *30th Annual European Symposium on Algorithms (ESA 2022)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2022.
- [32] Hanmao Shi and Thomas H Spencer. Time-work tradeoffs of the single-source shortest paths problem. *Journal of algorithms*, 30(1):19–32, 1999.
- [33] Mikkel Thorup. On shortcutting digraphs. In *International Workshop on Graph-Theoretic Concepts in Computer Science*, pages 205–211. Springer, 1992.
- [34] Jeffrey D Ullman and Mihalis Yannakakis. High-probability parallel transitive-closure algorithms. *SIAM Journal on Computing*, 20(1):100–125, 1991.

A Proof of Lemma 6

Let n be a sufficiently large integer, and let $p \in [1, n]$ such that p divides n . Let $p_1 := 2^{30}p$. Recall that to obtain G_2 , we first construct graph $G_1 := G(p_1, p)$, where $G(p_1, p)$ denotes our initial construction of the graph G_1 in Lemma 1 on p_1 nodes with an associated set Π_1 of $|\Pi_1| = p$ paths. We let V_1 denote the nodes and E_1 denote the edges of G_1 . Let P_1 be the set of critical pairs associated with G_1 , and let Π_1 be the corresponding canonical paths. Let ℓ and q be the construction parameters used to construct $G(p_1, p_1)$. Then

$$2^6 p \leq |P_1| \leq 2^{10} p, \quad \ell = \frac{p_1}{2^{10} p^{1/2} \log^{1/2} p_1}, \quad \text{and} \quad q = \frac{p_1}{2^{20} p^{1/2} \log^{3/2} p_1}.$$

Note that for all $\pi \in \Pi_1$, we have that $|\pi| = \ell$. We then modified G_1 by replacing each node $v \in V_1$ with a path $\pi_v = (u_1^v, \dots, u_k^v)$ with $k = n/p_1$ nodes. If an edge $(v_1, v_2) \in E_1$ was originally in G_1 , we replaced it with an edge $(u_k^{v_1}, u_1^{v_2})$. This gave us our final n -node graph $G_2 = (V_2, E_2)$.

Let $P_2 := \{(u_1^s, u_k^t) \mid (s, t) \in P_1\}$, and for all $(u_1^s, u_k^t) \in P_2$, let $\pi_{u_1^s, u_k^t}$ be the path obtained by taking $\pi_{s,t} \in \Pi_1$ and replacing each node $v \in \pi_{s,t}$ with the path π_v . Let $\Pi_2 := \{\pi_{s,t} \mid (s, t) \in P_2\}$. Then it is clear from our construction of G_2 that for all $(s, t) \in P_2$, the path $\pi_{s,t} \in \Pi_2$ is a unique shortest $s \rightsquigarrow t$ -path in G_2 . Additionally, for all $(s, t) \in P_2$, the number of nodes in path $\pi_{s,t} \in \Pi_2$ is at least $|\pi_{s,t}| \geq \ell_2$, where

$$\ell_2 := \ell \cdot \frac{n}{p_1} = \frac{n}{2^{10} p^{1/2} \log^{1/2} p_1} = \Theta\left(\frac{n}{p^{1/2} \log^{1/2} p}\right).$$

We now quickly verify that graph G_2 and associated critical paths Π_2 satisfy the properties of Lemma 1:

- By construction, G_2 has $\ell_2 = \Theta\left(\frac{n}{p^{1/2} \log^{1/2} n}\right)$ layers, and each path in Π_2 travels from the first layer to the last layer.
- Each path $\pi \in \Pi_2$ is a unique shortest path between its endpoints in G_2 . This follows from Lemma 2 and the observation that our path replacement step cannot increase the number of paths between pairs of nodes in G_2 .
- What remains is to show that for any two nodes $u, v \in V_2$, there are at most $\frac{\ell_2}{\text{hopdist}_{G_2}(u,v)}$ paths in Π_2 that contain both u and v . We prove this in two cases:

- **Case 1:** $u, v \in \pi_w$ for some $w \in V_1$. Note that $\text{hopdist}_G(u, v) \leq n/p_1$. By Proposition 4, the number of paths $\pi \in \Pi_1$ such that $w \in \pi$ is at most the number of direction vectors

$$\begin{aligned} |D| &\leq \frac{p_1 q}{4\ell^2} \leq \frac{p^{1/2}}{4 \log^{1/2} p_1} \leq \frac{p^{1/2}}{4 \log^{1/2} p_1} \cdot \frac{n/p_1}{\text{hopdist}_{G_2}(u, v)} \\ &\leq \frac{n}{2^{30} p^{1/2} \log^{1/2} p} \cdot \frac{1}{\text{hopdist}_{G_2}(u, v)} \leq \frac{\ell_2}{\text{hopdist}_{G_2}(u, v)}. \end{aligned}$$

Then by our construction of G_2 , the number of paths $\pi \in \Pi_2$ such that $u, v \in \pi$ is at most $|D| \leq \frac{\ell_2}{\text{hopdist}_{G_2}(u, v)}$.

- **Case 2:** $u \in \pi_{w_1}$ and $v \in \pi_{w_2}$ for distinct $w_1, w_2 \in V_1$. By Lemma 5, the number of paths $\pi \in \Pi_1$ such that $w_1, w_2 \in \pi$ is at most

$$\begin{aligned} \max \left\{ \frac{16q \log n}{\text{hopdist}_{G_1}(w_1, w_2)}, 1 \right\} &\leq \max \left\{ \frac{\ell}{\text{hopdist}_{G_1}(w_1, w_2)}, 1 \right\} \leq \max \left\{ \frac{\ell_2}{\text{hopdist}_{G_2}(u, v)}, 1 \right\} \\ &\leq \frac{\ell_2}{\text{hopdist}_{G_2}(u, v)}. \end{aligned}$$

(The final inequality follows from the fact that $\text{hopdist}_{G_2}(u, v) \leq \ell_2$.)

B Extending our shortcut set lower bound

To extend our shortcut set lower bound so that it holds in the regime of $p > n/(\alpha \log n)$, we will prove a more general statement about the behavior of the extremal function of shortcut sets. We write $\mathcal{S}(n, p)$ for the smallest integer D such that every n -node graph G has a shortcut set H of size $|H| \leq p$ such that $G \cup H$ has diameter at most D .

Lemma 13. *For all positive integers n and $x \in [1, n]$,*

$$\mathcal{S}(n, p/x) \leq x \cdot \mathcal{S}(n, p) \cdot 16 \log n.$$

Lemma 13 essentially states that if we decrease the number of shortcuts allowed in our shortcut set, the extremal function controlling the worst-case size of shortcut sets won't increase by too much. We will use this lemma in the opposite direction to argue that our lower bound of $\mathcal{S}(n, n/(\alpha \log n)) = \Omega\left(\frac{n^{1/4}}{\log^{1/4} n}\right)$ (where $\alpha > 0$ is a sufficiently large constant) that we obtained from our shortcut set construction in Lemma 7 implies lower bounds for shortcut sets with greater than $n/(\alpha \log n)$ shortcuts.

Let $p \in \left[\frac{n}{\alpha \log n}, n^{5/4}\right]$, and let $x = \frac{p}{n/(\alpha \log n)}$. Then by applying our lower bound from Lemma 7 to Lemma 13 we find that

$$\Omega\left(\frac{n^{1/4}}{\log^{1/4} n}\right) \leq \mathcal{S}\left(n, \frac{n}{\alpha \log n}\right) = \mathcal{S}\left(n, \frac{p}{x}\right) \leq x \cdot \mathcal{S}(n, p) \cdot 16 \log n.$$

Rearranging, we find that

$$\mathcal{S}(n, p) = \Omega\left(\frac{n^{5/4}}{p \log^{9/4} n}\right),$$

as claimed in Theorem 8.

We now prove Lemma 13, which will follow from a simple path subsampling argument.

Proof of Lemma 13. Let n be a positive integer, $p \in [1, n^2]$, and $x \in [1, n]$. Let $G = (V, E)$ be a graph on n nodes. We subsample nodes of G to construct a smaller graph $G_1 = (V_1, E_1)$ as follows.

- Independently sample each node V into set V_1 with probability $\frac{1}{2x}$. Then $\mathbb{E}[|V_1|] = \frac{n}{2x}$.
- For all pairs of nodes $u, v \in V$ such that $\text{dist}_G(u, v) \leq 8x \log n$, add directed edge (u, v) to E_1 . This completes the construction of $G_1 = (V_1, E_1)$.

By Markov's inequality, $|V_1| \leq n/x$ with probability at least $\frac{1}{2}$. Assume for now that this does indeed hold, and we have $|V_1| \leq n/x$. Then using the fact that $\mathbf{S}(c_1 n, c_1 p) \leq \mathbf{S}(c_2 n, c_2 p)$ if $c_1 \leq c_2$, we find that there exists a shortcut set H_1 of size $|H_1| \leq p/x$ such that the diameter of $G_1 \cup H_1$ is at most

$$\text{diameter}(G_1 \cup H_1) \leq \mathbf{S}(|V_1|, |H_1|) \leq \mathbf{S}(n/x, p/x) \leq \mathbf{S}(n, p).$$

Now we claim that this implies that $\text{diameter}(G \cup H_1) \leq x \cdot \mathbf{S}(n, p) \cdot 16 \log n$. For every pair of nodes $u, v \in V$ such that v is reachable from u in G , fix a shortest $u \rightsquigarrow v$ -path $\pi_{u,v}$ in G . Then the following statement holds with high probability: for all pairs of nodes $u, v \in V$ such that $\text{dist}_G(u, v) \geq 4x \log n$, path $\pi_{u,v}$ contains a node that was sampled into V_2 , i.e. $\pi_{u,v} \cap V_2 \neq \emptyset$. From now on, we will assume this property holds for our sampled set V_2 .

Consider a pair of nodes $u, v \in V$ such that $4x \log n \leq \text{dist}_G(u, v)$ and v is reachable from u in G , and let π be an $u \rightsquigarrow v$ -path in G . Let u' be the node in $\pi \cap V_1$ closest to u , and let v' be the node in $\pi \cap V_1$ closest to v . By our construction of G_1 and the above discussion, there must be a $u' \rightsquigarrow v'$ -path in G_1 . Moreover, $\text{dist}_{G_1 \cup H_1}(u', v') \leq \text{diameter}(G_1 \cup H_1) \leq \mathbf{S}(n, p)$. Then since $\text{dist}_G(x, y) \leq 8x \log n$ for all edges $(x, y) \in E_1$ in G_1 , it follows that

$$\text{dist}_{G \cup H_1}(u', v') \leq \text{dist}_{G_1 \cup H_1}(u', v') \cdot 8x \log n \leq \mathbf{S}(n, p) \cdot 8x \log n.$$

Then putting it all together,

$$\begin{aligned} \text{dist}_{G \cup H_1}(u, v) &\leq \text{dist}_{G \cup H_1}(u, u') + \text{dist}_{G \cup H_1}(u', v') + \text{dist}_{G \cup H_1}(v', v) \\ &\leq 4x \log n + \text{dist}_{G \cup H_1}(u', v') + 4x \log n \\ &\leq \text{dist}_{G_1 \cup H_1}(u', v') \cdot 8x \log n + 8x \log n \\ &\leq \mathbf{S}(n, p) \cdot 8x \log n + 8x \log n \\ &\leq x \cdot \mathbf{S}(n, p) \cdot 16 \log n. \end{aligned}$$

Finally, to conclude the analysis, we made two assumptions: (1) that $|V_1| \leq n/x$, which occurs with probability $1/2$, and (2) that for all u, v that are sufficiently far apart, we sampled a node on a $u \rightsquigarrow v$ path, which happens with high probability, i.e., $1 - 1/\text{poly}(n)$. By an intersection bound, there is positive probability that both events happen at the same time. So a graph G_1 exists as described, completing the proof. $\square$