

Approximate Distance Oracles Subject to Multiple Vertex Failures *

Ran Duan[†] Yong Gu[‡] Hanlin Ren[§]

Abstract

Given an undirected graph $G = (V, E)$ of n vertices and m edges with weights in $[1, W]$, we construct vertex sensitive distance oracles (VSDO), which are data structures that preprocess the graph, and answer the following kind of queries: Given a source vertex u , a target vertex v , and a batch of d failed vertices D , output (an approximation of) the distance between u and v in $G - D$ (that is, the graph G with vertices in D removed). An oracle has stretch α if it always holds that $\delta_{G-D}(u, v) \leq \tilde{\delta}(u, v, D) \leq \alpha \cdot \delta_{G-D}(u, v)$, where $\delta_{G-D}(u, v)$ is the actual distance between u and v in $G - D$, and $\tilde{\delta}(u, v, D)$ is the distance reported by the oracle.

In this paper we construct efficient VSDOs for any number d of failures. For any constant $c \geq 1$, we propose two oracles:

- The first oracle has size $n^{2+1/c}(\log n/\epsilon)^{O(d)} \cdot \log W$, answers a query in $\text{poly}(\log n, d^c, \log \log W, \epsilon^{-1})$ time, and has stretch $1 + \epsilon$, for any constant $\epsilon > 0$.
- The second oracle has size $n^{2+1/c} \text{poly}(\log(nW), d)$, answers a query in $\text{poly}(\log n, d^c, \log \log W)$ time, and has stretch $\text{poly}(\log n, d)$.

Both of these oracles can be preprocessed in time polynomial in their space complexity. These results are the first approximate distance oracles of poly-logarithmic query time for any constant number of vertex failures in general undirected graphs. Previously there are $(1 + \epsilon)$ -approximate d -edge sensitive distance oracles [Chechik et al. 2017] answering distance queries when d edges fail, which have size $O(n^2(\log n/\epsilon)^d \cdot d \log W)$ and query time $\text{poly}(\log n, d, \log \log W)$.

1 Introduction

Real-life networks are prone to failures. Usually, there can be several failed nodes or links, but the graph topology will not deviate too much from the underlying failure-free graph. A typical problem is to find the shortest path between two nodes in a network that avoids a specific set of failed nodes or links. This motivates the d -failure model, in which we should preprocess a graph, such that upon a small number (d) of failures, we can “recover” from these failures quickly.

In their pioneering work, Demetrescu and Thorup [27] designed a data structure that can maintain all-pairs shortest paths under one edge failure. In other words, for each triple (u, v, f) where u, v are vertices and f is a failed edge, the data structure can output the length of the shortest path from u to v that does not go through f , in $O(\log n)$ query time. A subsequent work [28] extends the structure to also handle one vertex failure, and improves the query time to $O(1)$. The one-failure case is studied extensively in literature [23, 7, 8, 33, 64, 36, 5, 9, 37, 19, 52].

People also tried to find structures handling multiple failures. For undirected graphs, we can answer connectivity queries under d edge failures¹ [51, 31, 32] and d vertex failures [31, 32] in $\text{poly}(d, \log n)$ time. Chechik et al. [21] designed a data structure that maintains $O(d)$ -approximate shortest paths under d edge failures in an undirected graph, and Bilò et al. [10] improved the approximation ratio to $2d + 1$. For any $\epsilon > 0$, Chechik et al. [20] designed a data structure that $(1 + \epsilon)$ -approximates shortest paths under d edge failures in an undirected graph, with space complexity $O(n^2(\log n/\epsilon)^d \cdot d \log W)$ and query time $\text{poly}(\log n, d, \log \log W)$, where W is the ratio of the largest edge weight to the smallest edge weight. More related work can be found in Section 1.1.

However, despite much effort, it was not known if one can maintain (approximate) shortest paths under multiple *vertex* failures. This problem was addressed as an open problem in [5, 21, 20], and also in Chechik’s PhD thesis [18].

In this paper we build efficient data structures that answer approximate distance queries under multiple vertex failures for general undirected graphs, answering the above question in the affirmative. A *vertex-sensitive distance oracle* (VSDO) for a weighted undirected graph $G = (V, E)$ is a data structure that given a set of failed vertices $D \subseteq V$ and $u, v \in V \setminus D$, outputs (an estimate of) the length of the shortest path from u to v that avoids all vertices in D . We assume a known upper bound d on the number of failures, i.e. for any query

*This work has been supported in part by the Zhongguancun Haihua Institute for Frontier Information Technology.

[†]Institute for Interdisciplinary Information Sciences, Tsinghua University.

[‡]Institute for Interdisciplinary Information Sciences, Tsinghua University.

[§]Institute for Interdisciplinary Information Sciences, Tsinghua University.

¹We can also use dynamic connectivity structures with poly-logarithmic worst case update time [41, 35, 63] to handle d edge failures.

(u, v, D) , we always have $|D| \leq d$. We will be concerned with the following parameters of a VSDO:

- Space complexity, i.e. the amount of space that the data structure occupies.
- Query time, i.e. the time needed to answer one query (u, v, D) .
- Approximation ratio, a.k.a. stretch: A VSDO has stretch α if it always holds that $\delta_{G-D}(u, v) \leq \tilde{\delta}(u, v, D) \leq \alpha \cdot \delta_{G-D}(u, v)$, where $\delta_{G-D}(u, v)$ is the actual distance between u and v in $G - D$ (i.e. G with D disabled), and $\tilde{\delta}(u, v, D)$ is the output of the VSDO.

We will not be particularly interested in the preprocessing time of VSDOs; nevertheless, all VSDOs in this paper can be preprocessed in time polynomial in their space complexity.

In this paper, n and m denote the number of vertices and edges respectively. Let W be the ratio of the largest edge weight to the smallest edge weight. W.l.o.g. we can assume that edge weights are real numbers in $[1, W]$.

1.1 More Related Work.

Sensitivity oracles. For the case of two vertex failures, Duan and Pettie [30] showed that exact distances in a directed weighted graph can be queried in $O(\log n)$ time, with an oracle of size $O(n^2 \log^3 n)$, and Choudhary [22] designed an oracle of $O(n)$ size that handles single source reachability queries in directed graphs in $O(1)$ time.

The general problem of d failures has also received attention on *planar* graphs: Borradaile et al. [15] constructed a data structure that maintains connectivity under d vertex failures, and Charalampopoulos et al. [16] designed a data structure that answers exact distance queries under d vertex failures.

In a recent breakthrough, van den Brand and Saranurak [62] gave an oracle that handles an arbitrary number d of edge failures in *directed* graphs. Their oracle can answer reachability queries in $O(d^\omega)$ time, and exact distance queries in $n^{2-\Omega(1)}$ time (for small integer weights), where $\omega < 2.3728639$ is the matrix-multiplication exponent [24, 57, 65, 44].

We summarize the sensitivity connectivity/distance oracles in the full version (Table 2 of Appendix C) of this paper.

Fault-tolerant structures. A related concept is *fault-tolerant (FT) spanners*: a subgraph G' of G is a d -FT spanner if, after removing any d vertices, the remaining parts of G' is a spanner of the remaining parts of G . It might be *a priori* surprising that sparse FT spanners exist, but Chechik et al. [17] gave the first

construction of d -FT $(2k - 1)$ -spanners with $O(d^2 k^{d+1} \cdot n^{1+1/k} \log^{1-1/k} n)$ edges. Subsequent papers [29, 12, 14] improved the number of edges to $O(n^{1+1/k} d^{1-1/k})$, which is optimal assuming the girth conjecture of Erdős [34].

Besides FT spanners, there are many other kinds of fault-tolerant structures, e.g. [11, 45, 48, 49, 10, 13, 47, 50]. We refer the reader to the excellent survey of [46].

Dynamic shortest path. There are dynamic all-pairs shortest path structures handling vertex updates. Thorup [59] gave a fully dynamic all-pairs shortest paths structure with worst-case update time $\tilde{O}(n^{2.75})$, and Abraham et al. [1] gave a randomized worst-case update time bound $\tilde{O}(n^{2+2/3})$. Recently, Brand and Nanongkai [61] gave a $(1 + \epsilon)$ -approximate algorithm for maintaining APSP under edge insertions and deletions with worst-case update time $\tilde{O}(n^{1.863}/\epsilon^2)$ for directed graphs. Using it we can construct a $\tilde{O}(d\epsilon^{-2}n^{1.863})$ query time oracle for vertex failures in undirected graphs by a simple reduction. Other fully or partial dynamic shortest path structures include [2, 6, 26, 40, 38, 39, 43, 55, 54, 58, 56].

1.2 Our results. We provide the first constructions of approximate VSDOs for general undirected graphs with poly-logarithmic query time. Our main results are as follows:²

THEOREM 1.1. *For any constants $c \geq 1$ and $\epsilon > 0$, we can construct VSDOs for undirected graphs with:*

- space complexity $n^{2+1/c} \log W \cdot (\epsilon^{-1} \log n)^{O(d)}$, query time $\tilde{O}(d^{2c+6} \epsilon^{-1} \log \log W)$ and stretch $1 + \epsilon$;*
- space complexity $\tilde{O}(n^{2+1/c} d^3 \log(nW))$, query time $\tilde{O}(d^{2c+9} \log \log(nW))$ and stretch $O(d^{c+2} \log^6 n)$.*

Each oracle can be preprocessed in time polynomial in their space complexity.³ Our constructions also allow an actual approximate shortest path to be retrieved in an additional time of $O(\ell)$, where ℓ is the number of edges in the reported path.

Using existing structures, we need either $n^{\Omega(d)}$ space or $\Omega(n)$ query time.⁴ Thus our results are the first of its kind.

² \tilde{O} hides $\text{poly}(\log n)$ factors.

³See the full version (Table 3 of Appendix C) for precise time bounds.

⁴We can use the d -fault tolerant spanner [17, 29, 12, 14] with the brute-force query algorithm, build n^{d-2} two-failure distance oracles [30], use the dynamic shortest path algorithms [61], or use the oracle [62] which also works for directed graphs. But none of these solutions provide both $n^{o(d)}$ space and $o(n)$ query time.

1.3 A brief overview. In this section, we briefly introduce the ideas needed to construct the desired VSDOs.

The edge-sensitive distance oracle of [20]. Our first VSDO depends on [20] which handles d edge failures. Therefore we briefly describe their oracle first. It may be helpful to think of their query algorithm as a recursive one.

Given $u, v \in V$ and a set D of d edge failures, let P_{ans} be the shortest u - v path in $G - D$, which we are searching for. The oracle first partitions the shortest path P from u to v in G (which may go through failures) into $\tilde{O}(\epsilon^{-1} \log W)$ short segments. Consider a segment X that contains some failed edges. If P_{ans} does not go through X , then we can “preprocess” the graph $G - X$ and search for P_{ans} in $G - X$. Otherwise, if P_{ans} goes through some vertex $x \in X$, then we can pick an *arbitrary* vertex $w \in X$ such that there are no failed edges between x and w , and *pretend* that P_{ans} passes through w . That is, we recursively find the shortest paths in $G - D$ from u to w and from w to v and concatenate them. It is easy to see that this brings an additive error of at most $2|X|$ to our solution, where $|X|$ is the length of X .

Thus, we want to find a small set of intermediate vertices, which we denote as H , with the following property: For every vertex x and failure f , if x has distance at most $|X|$ to f , then there is some $w \in H$ such that x also has distance at most $|X|$ to w in $G - D$. As it turns out that the query time is polynomial in $|H|$, the size of H should be small.

There is a natural choice of H : we simply let it be the set of vertices incident to some failed edges. It is easy to see that $|H| \leq 2d$, thus the query algorithm runs in time polynomial of d . The above property is also true: given any vertex x and a nearby failure f , we can walk along the path from x to f until we meet a failed edge, then the vertex w we stop at is both in H and close to x . We can control the total additive error (i.e. the sum of $2|X|$ ’s over the “recursion”) to be at most $\epsilon \cdot |P_{\text{ans}}|$, by partitioning each path into *sufficiently short* segments.

Note that, for the sake of intuition, we have omitted some important details, such as how to “preprocess” $G - X$ (by a decision tree structure) and how to implement the query algorithm (non-recursively).

The “high-degree” obstacle. The obvious difficulty of handling vertex failures is the presence of failed vertices with very high degrees. If every failed vertex has degree $\leq \Delta$, we can simply simulate an edge-failure distance oracle [21, 20] and delete at most $d \cdot \Delta$ edges from it. Equivalently, we can define the set of intermediate vertices H as those active vertices adjacent to some failure, then $|H| \leq d \cdot \Delta$ and we run the above

query algorithm. However the techniques of [21, 20] do not seem to work for high-degree vertex failures. For example, techniques in [21] only guarantee a stretch of $\geq \Delta$, and techniques in [20] require $\text{poly}(\Delta)$ query time, therefore both are unsatisfactory when $\Delta = \Omega(n)$.

By the construction of $(2k - 1)$ -stretch spanners with $O(n^{1+1/k})$ edges [4], we can construct a $(2 \log n - 1)$ -stretch spanner with $O(n)$ edges. We note that the query algorithm works even if every failed vertex has a small degree in the spanner (rather than in the whole graph): We can define H to be the set of vertices adjacent to some failed vertex *in the spanner*. If P_{ans} goes through some vertex x that has distance $|X|$ to a failed vertex f , the distance between x and f in the spanner is $O(|X| \log n)$, and there must be some $w \in H$ that has distance $O(|X| \log n)$ to x in $G - D$. By partitioning the paths into shorter segments, we can still control the additive error, i.e. the sum of $O(|X| \log n)$ over the “recursion”, to be less than $\epsilon \cdot |P_{\text{ans}}|$.

High-degree hierarchy: A first attempt.

Given the “high-degree” obstacle, it is natural to see whether the “high-degree hierarchy” of [31] may help us. Plugging the spanners⁵ into the hierarchy of [31], we obtain a structure as follows. The vertices are partitioned into $p = O(\log n)$ levels; let U_i be the set of vertices with level $\geq i$. So we have a sequence of vertex sets $V = U_1 \supseteq U_2 \supseteq \dots \supseteq U_p \supseteq U_{p+1} = \emptyset$, and the i -th level is the set $U_i \setminus U_{i+1}$.⁶ For every i , let G_i be the induced subgraph of $V \setminus U_{i+1}$. We do not have a complete spanner for G_i ; we can only afford to build a “subset-spanner” that preserves the distances in G_i , among vertices in $U_i \setminus U_{i+1}$ (instead of $V \setminus U_{i+1}$). The structure guarantees that every failed vertex in the subset-spanner of any level has low degrees.

It is natural to define H as the set of neighbors of failures in the subset-spanners, and $|H|$ will be small. If P_{ans} goes through some vertex x that has distance $|X|$ to a failed vertex f , and x and f are in the *same level*, then we can find an intermediate vertex $w \in H$ that has distance $O(|X| \log n)$ to x in $G - D$, and we are fine. But what if x and f are in different levels? In this case, the x - f path may not be preserved by the “subset-spanner”, thus not captured by H . In [31, Section 4], the authors used ad hoc structures to preserve connectivity between different levels; it appears difficult to extend these structures to also handle $((1 + \epsilon)$ -approximate)

⁵The reason that we need to plug in a spanner, rather than the original graph, is that we can only plug in a *sparse* graph into the high-degree hierarchy.

⁶In the hierarchy structure of Section 2.2, each U_{i+1} is not necessarily a subset of U_i ; this issue is not essential, so for simplicity, in the brief overview we will assume each U_{i+1} is indeed a subset of U_i .

distances.

Our ideas. It is inconvenient that the spanner at level i only preserves distances inside $U_i \setminus U_{i+1}$. Therefore, our first idea is to “extend” the spanners to also preserve distances at lower levels: the spanner at level i should preserve distances between any pair of vertices (x, y) , where $x \in U_i \setminus U_{i+1}$ and $y \in V \setminus U_{i+1}$. Note that we still only guarantee that every vertex failure has small degrees in the *original* spanners; they may have large degrees in the extended spanners.

We implement the spanners by *tree covers*, and there is a natural way to “extend” them. The extended tree cover consists of a collection of trees whose union is a spanner that preserves distances between $U_i \setminus U_{i+1}$ and $V \setminus U_{i+1}$. Moreover, each tree is a shortest path tree rooted in $U_i \setminus U_{i+1}$ (the highest level of G_i). See Section 2.1 for more details.

Recall that in the query algorithm, we have a non-failure vertex x that is close to a failure f , and we want to find an intermediate vertex $w \in H$ that is close to x in $G - D$. Suppose that x is at a higher level than f . If we walk from f (at a lower level) to x (at a higher level), it seems that our first step should go to the parent of f in some tree. Actually, this intuition can be rigorously proved! See the proof of Lemma 3.2. Therefore, if H consists of the neighbors of every failure (in the original spanners) and the parents of every failure in each tree (in the extended tree covers), then we can deal with every (x, f) such that the level of x is at least that of f . Every failure is only in $\tilde{O}(1)$ trees, thus $|H|$ is indeed small.

We need to adapt the query algorithm to ensure that f never has a higher level than x . Let P be the shortest u - v path in the original graph, and we partition P into short segments. Consider a segment X that contains failures, and let i be the highest level of any failure in X . If P_{ans} does not contain any vertex in X with level at least i , then we can “preprocess” the graph $G - (X \cap U_i)$ and search for P_{ans} in this subgraph. Otherwise P_{ans} goes through some $x \in (X \cap U_i)$, and by definition, the level of x cannot be smaller than the level of any failure in X . Therefore, we can find some intermediate vertex $w \in H$ close to x , “pretend” that P_{ans} goes through w , and continue.

The above discussion implies a data structure with space complexity roughly n^3 . To reduce the space complexity by a factor of $n^{1-o(1)}$, we prove a structural theorem (Theorem 4.1) for shortest paths under vertex failures, which allows us to compress such paths. (The corresponding theorem [20, Theorem 3.1] does not hold for vertex failures.) Curiously, the proof of this theorem also relies on Lemma 3.2.

On oracle (b). Although oracle (b) has a larger stretch compared to oracle (a), we think it is also of interest, since it is the first oracle that handles $\omega(\log n)$ failures in polynomial space and $\text{poly}(\log n)$ query time, within a reasonable stretch.⁷ Note that setting $\epsilon = \omega(1)$ (e.g. $\epsilon = \log n$) in oracle (a) does *not* improve its space complexity to $n^2 \log^{o(d)} n$, so oracle (b) is *not* a direct corollary of oracle (a).

1.4 Notation. In this paper, $\log x = \log_2 x$, $\ln x = \log_e x$. For a set S and an integer k , $|S|$ is the cardinality of S , and we denote $\binom{S}{k} = \{S' \subseteq S : |S'| = k\}$, and $\binom{S}{\leq k}, \binom{S}{\geq k}$ are defined analogously. For two sets X and Y , define their Cartesian product as $X \times Y = \{(x, y) : x \in X, y \in Y\}$. We use \circ as the concatenation operator for paths or sequences. For paths P_1, P_2 , if u is the last vertex in P_1 and v is the first vertex in P_2 , then $P_1 \circ P_2$ is well-defined if $u = v$ or (u, v) is an edge in G .

For a graph H and $u, v \in V(H)$, $w_H(u, v)$ denotes the length of the edge between u and v ($w_H(u, v) = +\infty$ if such an edge does not exist), $\delta_H(u, v)$ denotes the length of the shortest path in H from u to v and $\pi_H(u, v)$ denotes the corresponding shortest path. If $S \subseteq V(H)$ is a subset of vertices, then $\delta_H(u, S) = \min\{\delta_H(u, v) : v \in S\}$. ($\delta_H(u, \emptyset) = +\infty$.) We omit the subscript H if $H = G$ is the input graph. We define $H[S]$ as the subgraph induced by S , and $H - S = H[V(H) \setminus S]$. We use nW as an upper bound of the diameter of any (connected) subgraph of G . We assume that the shortest path between every pair of vertices in any subgraph is unique (see Section 3.4 of [25]).

For a path P and $u, v \in P$, define $P[u, v]$ as the portion from u to v in P , and sometimes this notation emphasizes the *direction* from u to v . Let $(u = x_0, x_1, \dots, x_{\ell-1}, x_\ell = v)$ denote the path $P[u, v]$, then we define $P(u, v) = P[x_1, v]$, $P[u, v) = P[u, x_{\ell-1}]$ and $P(u, v) = P[x_1, x_{\ell-1}]$. Define $|P|$ as the length of path P . For a tree T rooted at r and a vertex $x \in V$, define the depth of x , denoted by $\text{dep}_T(x)$, as the (weighted) distance from x to r in T .

In this paper, D denotes the set of $\leq d$ failed vertices. For convenience, we always assume $n \geq 3$ and $d \geq 2$.

Note that we also define some more notations at the end of Section 2.2, which is relevant to the “high-degree hierarchy”.

2 Source-Restricted Tree Covers in High-Degree Hierarchy

Our VSDO is based on a variant of the *high-degree hierarchy* of [31], which we equip with the *source-restricted*

⁷It seems that even $O(\sqrt{n})$ stretch was open before this result.

tree covers of [60, 53] to approximately preserve distances.

2.1 Source-restricted tree covers. Let $G = (V, E)$ be an undirected graph. A *tree cover* of G is, informally, a set of trees such that every vertex $v \in V$ is in a small number of trees, and for every two vertices $u, v \in V$, there is a tree that approximately preserves their distance $\delta(u, v)$. In this paper, we relax the second condition, requiring it to hold only for every $u \in S, v \in V$, where S is some subset of V . Following terminologies of [53], we call such tree covers *source-restricted*.

Throughout this paper, $k = \ln n$.⁸ We define *source-restricted tree cover* as follows.

DEFINITION 2.1. *Given $S \subseteq V$, an S -restricted tree cover is a set of rooted trees $\{T(w) : w \in S\}$, such that the following hold.*

- a) *For every $w \in S$, there is exactly one tree $T(w)$ rooted at w , spanning a subset of V (which we denote as $V(T(w))$).*
- b) *For every $u \in S, v \in V$, there is some $w \in S$ such that $u, v \in V(T(w))$, and the distance between u and v in $T(w)$ is at most $(2k - 1)\delta(u, v)$.*
- c) *Every vertex $v \in V$ is in at most $kn^{1/k}(\ln n + 1) \leq 2e \ln^2 n$ trees.*

In [60], Thorup and Zwick constructed approximate distance oracles, and they noticed that their constructions are also good tree covers. A simple modification of their construction (see [53]) yields source-restricted tree covers.

THEOREM 2.1. *Given a graph $G = (V, E)$ and $S \subseteq V$, we can compute in deterministic polynomial time an S -restricted tree cover $\mathcal{T}(S) = \{T(w) : w \in S\}$ such that for any $u \in S, v \in V$, the vertex w in Definition 2.1 b) can be found in $O(k)$ time.*

A proof sketch of Theorem 2.1 can be found in Appendix B of the full version; we also refer the interested reader to [60, 53] for details.

For $S \subseteq V$, we denote $\mathcal{T}(S)$ as the S -restricted tree cover constructed in Theorem 2.1. For $S, R \subseteq V$, we denote $\mathcal{T}_R(S)$ as the $(S \setminus R)$ -restricted tree cover $\mathcal{T}(S \setminus R)$ in $G - R$.

For technical reasons (namely, we want the hierarchy structure in Section 2.2 to have a reasonable size),

⁸Our construction works for any parameter k , but the complexity is proportional to $kn^{1/k}$, so we minimize it by setting $k = \ln n$.




Figure 1: A sample tree in $\mathcal{T}(S)$. Black vertices are in S , green vertices are trunk vertices not in S , and bold edges denote the subtree induced by trunk vertices. We also include a table of degrees and pseudo-degrees of some sample vertices.

we need that the number of “high-degree” vertices in $\mathcal{T}(S)$ is only $o(|S|/d)$, where d is the number of failures. However, here we defined the tree cover $\mathcal{T}(S)$ to span not only S , but maybe some other vertices in V .⁹ So we can only prove degree bounds of the following form: the number of vertices with high degree w.r.t. the “trunk” parts of the tree cover is $o(|S|/d)$. The precise definitions are as follows.

DEFINITION 2.2. *Consider $S \subseteq V, T \in \mathcal{T}(S), v \in V(T)$. We say v is a trunk vertex of T if there are $u, w \in S$ such that v lies on the path from u to w in T . The subtree (subgraph) of T induced by trunk vertices of T is denoted as $\text{Trunk}(T)$. The pseudo-degree of a vertex $v \in V(T)$, denoted as $\text{pdeg}_T(v)$, is the degree of v in $\text{Trunk}(T)$. If v is not a trunk vertex of T , then $\text{pdeg}_T(v) = 0$.*

Note that vertices in $\text{Trunk}(T)$ are not necessarily in S . See Fig. 1 as an example.

The following property will be useful in Section 3: for a vertex v that is not in $\text{Trunk}(T)$, its path in T to any vertex in S must go through its parent. (This is because the root of T is always in S .)

Let $s = 4e \cdot d^{c+1} \ln^2 n + 1$ be a degree threshold, where $c \geq 1$ is any constant. Define $\text{Hi}(\mathcal{T}(S))$ as the set of vertices in V that has pseudo-degree $> s$ in some tree in $\mathcal{T}(S)$. We prove our desired upper bound on $|\text{Hi}(\mathcal{T}(S))|$.

LEMMA 2.1. *For any $S \subseteq V, |\text{Hi}(\mathcal{T}(S))| \leq \frac{|S|}{2d^{c+1}}$.*

Proof. For a tree T , let $\text{Leaf}(T)$ be the set of leaves of T . Then there are at most $\left\lfloor \frac{|\text{Leaf}(T)| - 2}{s - 1} \right\rfloor$ vertices in T that has degree $> s$ [31, Lemma 3.1]. For any $T \in \mathcal{T}(S)$, $\text{Leaf}(\text{Trunk}(T)) \subseteq S$ by definition, thus

$$\sum_{T \in \mathcal{T}(S)} |\text{Leaf}(\text{Trunk}(T))| \leq \sum_{v \in S} |\{T \in \mathcal{T}(S) : v \in T\}|.$$

⁹This corresponds to the informal description of “extending” tree covers in Section 1.3.

Algorithm 1 Path finding algorithm

- 1: $U_1 \leftarrow V$
 - 2: **for** $i \leftarrow 1$ to h **do**
 - 3: If U_i is a leaf then set $p \leftarrow i$ and halt
 - 4: Let W_1, W_2, \dots, W_d be the children of U_i and artificially define $W_0 = \emptyset, W_{d+1} = W_d$.
 - 5: Let $j \in [0, d]$ be minimal such that $D \cap (W_{j+1} \setminus W_j) = \emptyset$
 - 6: If $j = 0$ then set $p \leftarrow i$ and halt
 - 7: Otherwise $U_{i+1} \leftarrow W_j$
-

Since every $v \in S$ appears in $\leq 2e \ln^2 n$ trees in $\mathcal{T}(S)$, we have

$$\sum_{T \in \mathcal{T}(S)} |\text{Leaf}(\text{Trunk}(T))| \leq |S| \cdot 2e \ln^2 n,$$

thus

$$\begin{aligned} |\text{Hi}(\mathcal{T}(S))| &\leq \sum_{T \in \mathcal{T}(S)} \left\lfloor \frac{|\text{Leaf}(\text{Trunk}(T))|}{s-1} \right\rfloor \\ &\leq \frac{|S| \cdot 2e \ln^2 n}{s-1} = \frac{|S|}{2d^{c+1}}. \end{aligned}$$

□

2.2 The high-degree hierarchy. We use a simplified version of the high-degree hierarchy in [31]. Fix a parameter $c \geq 1$, the hierarchy structure is a set of $O(n^{1/c})$ representations of the graph, such that for every set of d failures, we can find some representation in which all failed vertices have low pseudo-degrees in their relevant tree covers.

DEFINITION 2.3. *The hierarchy tree is a rooted tree in which every node¹⁰ corresponds to a subset of V . The root corresponds to V . Each node U ($U \subseteq V$) stores a tree cover $\mathcal{T}(U)$, and each edge (U, W) , where U is the parent of W , stores a tree cover $\mathcal{T}_W(U)$, which is $\mathcal{T}(U \setminus W)$ in $G - W$. The hierarchy tree is constructed as follows. Let U be any node. If $\text{Hi}(\mathcal{T}(U)) = \emptyset$, then U is a leaf; otherwise let W_1, W_2, \dots, W_d be its children, where*

$$(2.1) \quad \begin{aligned} W_1 &= \text{Hi}(\mathcal{T}(U)), \\ W_i &= W_{i-1} \cup \text{Hi}(\mathcal{T}_{W_{i-1}}(U)). \quad (\text{for } 2 \leq i \leq d) \end{aligned}$$

Then we recursively deal with all W_1, W_2, \dots, W_d .

There are two main differences compared with the hierarchy structure in [31].

- We simplified the definition of W_1 as in (2.1). This change is not essential, but we feel that it could make the hierarchy tree easier to understand. As a consequence, a node is not necessarily a subset of its parent, which is different from [31] (and Section 1.3).
- More importantly, we store in each node the source-restricted tree covers introduced in Section 2.1. By contrast, [31] only concerns about connectivity, so they used a (somewhat arbitrary) spanning forest instead.

The following lemmas assert that the hierarchy tree “is small, shallow and effectively represents the graph” [31], which are crucial for our data structures.

LEMMA 2.2. (HIERARCHY SIZE AND DEPTH)

Consider the hierarchy tree constructed with high-degree threshold $s = 4e \cdot d^{c+1} \ln^2 n + 1$, then the following hold.

1. *The depth h of the hierarchy tree is at most $\lfloor \frac{1}{c} \log_d n \rfloor$,¹¹ assuming the root has depth 0.*
2. *The number of nodes in the hierarchy tree is at most $O(n^{1/c})$.*

Proof. Let U be a node in the hierarchy tree, W_1, W_2, \dots, W_d be its children (if exist). By Lemma 2.1, we have $|W_1| = |\text{Hi}(\mathcal{T}(U))| \leq \frac{|U|}{2d^{c+1}}$ and $|\text{Hi}(\mathcal{T}_{W_i}(U))| \leq \frac{|U|}{2d^{c+1}}$ for any $0 < i \leq d$. Since $W_{i+1} = W_i \cup \text{Hi}(\mathcal{T}_{W_i}(U))$, it follows that $|W_i| \leq i|U|/2d^{c+1}$ for each i . Therefore $|W_i| \leq |U|/2d^c$ for all $0 < i \leq d$. Any node at the k -th level corresponds to a subset of V with size at most $n/(2d^c)^k$, therefore the depth of the hierarchy tree is at most $h \leq \lfloor \log_{2d^c} n \rfloor \leq \lfloor \frac{1}{c} \log_d n \rfloor$. There are at most $\sum_{i=0}^h d^i \leq 2d^h = O(n^{1/c})$ nodes in the hierarchy tree. □

¹⁰We use “vertex” for nodes in the input graph, and “node” for nodes in the hierarchy tree.

¹¹In subsequent sections we will write h as a shorthand of $O(\log n / \log d)$ in time/space bounds.

LEMMA 2.3. *For any set D of at most d failures, Algorithm 1 finds in $O(hd)$ time a path $U_1(=V), U_2, \dots, U_p$ in the hierarchy tree from root to some node U_p , such that for every tree $T \in \bigcup_{1 \leq i \leq p} \mathcal{T}_{U_{i+1}}(U_i)$ and every $f \in D$, f has pseudo-degree $\leq s$ in T . (Assume $U_{p+1} = \emptyset$.)*

Proof. Algorithm 1 executes at most $O(h)$ iterations since the hierarchy tree has depth at most h . For every node U and vertex $v \in V$, we store the first child W_j of U (or none) that v appears in. It is then easy to implement each iteration in $O(d)$ time. When Algorithm 1 halts at Line 3 or 6, either U_p is a leaf or $D \cap W_1 = \emptyset$, where $W_1 = \text{Hi}(\mathcal{T}(U_p))$ is the first child of U_p . Clearly, in both case we have $D \cap \text{Hi}(\mathcal{T}(U_p)) = \emptyset$.

For any $1 \leq i < p$, since $W_{d+1} = W_d$, Line 5 can always find such j . Let $U_{i+1} = W_j$ be the j -th child of U_i . If $j = d$, then $D \cap (W_{j'+1} \setminus W_{j'}) \neq \emptyset$ for $j' = 0, \dots, d-1$. Since $|D| \leq d$, we have $D \subseteq W_d$, thus $D \cap \text{Hi}(\mathcal{T}_{W_d}(U_i)) = \emptyset$. If $1 \leq j < d$, then we have $D \cap (W_{j+1} \setminus W_j) = \emptyset$. Since $W_{j+1} = W_j \cup \text{Hi}(\mathcal{T}_{W_j}(U_i))$ and $\text{Hi}(\mathcal{T}_{W_j}(U_i)) \cap W_j = \emptyset$, we have $D \cap \text{Hi}(\mathcal{T}_{W_j}(U_i)) = \emptyset$. The lemma follows. \square

In Section 3, Section 4 and Section 5, we always deal with a path $V = U_1, U_2, \dots, U_p$ in the hierarchy tree from the root V to a node U_p (not necessarily a leaf node). Artificially define $U_{p+1} = \emptyset$. In other words:

- We run the preprocessing algorithm for every possible such paths $V = U_1, U_2, \dots, U_p$, and we build a separate data structure for each path. The space complexity is then multiplied by a factor of $O(n^{1/c})$.
- In the query algorithm, given D , we always begin by identifying a path $V = U_1, U_2, \dots, U_p$ using Lemma 2.3, then every failed vertex $f \in D$ has pseudo-degree $\leq s$ in every $\mathcal{T}_{U_{i+1}}(U_i)$ ($1 \leq i \leq p$).

Fix a path $V = U_1, U_2, \dots, U_p$ in the hierarchy tree, and we assume $U_{p+1} = \emptyset$. The following corollary of Theorem 2.1 will be important for us.

COROLLARY 2.1. *Let $x \in U_\ell \setminus U_{\ell+1}$ and $y \in V \setminus U_{\ell+1}$. There is a tree $T \in \mathcal{T}_{U_{\ell+1}}(U_\ell)$ such that the distance between x and y in T is at most $(2k-1)\delta_{G-U_{\ell+1}}(x, y)$. Moreover, the root of such a tree can be found in $O(k)$ time.*

We will denote this tree T as $T_\ell(x, y)$, and denote the path from x to y in T as $\mathcal{P}_\ell(x, y)$.

The set of trees is denoted as $\mathcal{T} = \bigcup_{i=1}^p \mathcal{T}_{U_{i+1}}(U_i)$. The proof of Lemma 2.2 shows that $|U_i| \leq |U_{i-1}|/2d^c$, therefore $|\mathcal{T}| \leq \sum_{i=1}^p |U_i| = O(n)$. In a path $V = U_1, U_2, \dots, U_p$ in the hierarchy tree, since U_{i+1} is not

necessarily a subset of U_i , we define the *level* of a vertex v as:

DEFINITION 2.4. *Fix a path $V = U_1, U_2, \dots, U_p$ in the hierarchy tree, define the level of v to be the largest integer l such that $v \in U_l$, denoted as $l(v)$. Define G_ℓ to be the subgraph of G induced by all vertices with level at most ℓ .*

3 An $(1 + \epsilon)$ -Stretch Oracle with $n^{3+1/c+o(1)}$ Space

In this section we present an oracle with

$$\begin{aligned} \text{space complexity} & n^{3+1/c} \cdot (\epsilon^{-1} \log(nW))^{O(d)}, \\ \text{query complexity} & \text{poly}(\log(nW), \epsilon^{-1}, d), \\ \text{stretch} & 1 + \epsilon, \end{aligned}$$

for any $\epsilon > 0$. In this paper (except Section 5 and Section 4.5) we may assume $d = o\left(\frac{\log n}{\log \log n}\right)$, $W = \text{poly}(n)$, so we can simplify the notation for space complexity to $n^{3+1/c+o(1)}$. We will show how to reduce the space complexity to $n^{2+1/c+o(1)}$ in Section 4.

3.1 Data structure. Let $\epsilon_1 = \epsilon/(2 + \epsilon)$, $\epsilon_2 = \epsilon_1/(2k - 1)$, so $\epsilon_1 < 1$. (Recall $k = \ln n$.) We first define a decomposition of a path P into $O(\log |P|/\epsilon_2)$ segments. This definition has the same spirit as [20, Definition 2.1], but it partitions the *vertices* (excluding u, v), rather than *edges*, into segments.

DEFINITION 3.1. ((ϵ_2 -)SEGMENTS) *Consider a path $P = (u = v_0, v_1, v_2, \dots, v_\ell = v)$. For every $1 \leq i, j < \ell$, we say v_i and v_j are in the same segment if one of the two conditions hold:*

- $|P[u, v_i]|, |P[u, v_j]| \leq |P|/2$ and $\lfloor \log_{1+\epsilon_2} |P[u, v_i]| \rfloor = \lfloor \log_{1+\epsilon_2} |P[u, v_j]| \rfloor$.
- $|P[v_i, v]|, |P[v_j, v]| < |P|/2$ and $\lfloor \log_{1+\epsilon_2} |P[v_i, v]| \rfloor = \lfloor \log_{1+\epsilon_2} |P[v_j, v]| \rfloor$.

It is easy to verify that being in the same segment is indeed an equivalence relation, and each equivalence class is indeed a contiguous segment of the path. (See Fig. 2.) There are $O(\log |P|/\epsilon_2)$ different segments. For a vertex v_i on P , define $\text{seg}(v_i, P)$ as the segment it belongs to. More precisely, $\text{seg}(v_i, P) = P[v_l, v_r]$ where v_l is the leftmost (closest-to- u) vertex in the segment, and v_r is the rightmost (closest-to- v) vertex in the segment. Define $\text{seg}(P)$ as the set of segments on P . Note that u and v do not belong to any segment.

LEMMA 3.1. *Let P be a path from u to v , $x \in P \setminus \{u, v\}$, then $|\text{seg}(x, P)| \leq \epsilon_2 \cdot \min\{|P[u, x]|, |P[x, v]|\}$.*




Figure 2: An illustration of path decomposition, where each rounded rectangle denotes a segment.

Proof. If $|P[u, x]| \leq |P|/2$, $|P[u, x]| \leq |P[x, v]|$ and $\text{seg}(x, P)$ is defined in the first way. Let $i = \lfloor \log_{1+\epsilon_2} |P[u, x]| \rfloor$. For any $y \in \text{seg}(x, P)$, $(1 + \epsilon_2)^i \leq |P[u, y]| < (1 + \epsilon_2)^{i+1}$. Thus $|\text{seg}(x, P)| \leq (1 + \epsilon_2)^{i+1} - (1 + \epsilon_2)^i = \epsilon_2(1 + \epsilon_2)^i \leq \epsilon_2|P[u, x]|$. The case that $|P[x, v]| < |P|/2$ is symmetric. \square

Recall that we build a data structure for every node U_p in the hierarchy tree. Let the path from the root V to U_p be $U_1(= V), U_2 \dots, U_p$ and $U_{p+1} = \emptyset$. The data structure for U_p consists of decision trees $FT(u, v)$ for all pairs of vertices $u, v \in V$, which are constructed as follows:

- Each node¹² $\alpha \in FT(u, v)$ is associated with a set $\text{avoid}(\alpha) \subseteq V$ of vertices that we avoid.
- Denote the root of $FT(u, v)$ as $\text{root}(u, v)$, and let $\text{avoid}(\text{root}(u, v)) = \emptyset$.
- For each node $\alpha \in FT(u, v)$, we store the path $P_\alpha = \pi_{G-\text{avoid}(\alpha)}(u, v)$, i.e. the shortest u - v path in G not passing through $\text{avoid}(\alpha)$. If u, v are not connected in $G - \text{avoid}(\alpha)$, we assume that P_α is a path with length $+\infty$.
- For each node α of depth $< d$ (the root has depth 0), each segment $X \in \text{seg}(P_\alpha)$, and each $1 \leq i \leq p$, we create a child $ch(\alpha, X, i)$ of α , in which

$$\text{avoid}(ch(\alpha, X, i)) = \text{avoid}(\alpha) \cup (X \cap U_i).$$

That is, the path stored in a child of α needs to avoid $\text{avoid}(\alpha)$ and the vertices of U_i in a segment $X \in \text{seg}(P_\alpha)$.

The decision tree has depth d , and each non-leaf node has $O(p \cdot \log |P|/\epsilon_2) = O(h\epsilon^{-1} \log n \log(nW))$ children. (Recall $h = O(\log n / \log d)$ and $\epsilon_2 = \epsilon / ((2 + \epsilon)(2k - 1)) = \epsilon / \Theta(\log n)$.) We store in each node α the path P_α as well as a table of $\text{seg}(v, P_\alpha)$ for each $v \in P_\alpha$, so that we can quickly locate any vertex in P_α . Therefore one decision tree occupies $n \cdot O(h\epsilon^{-1} \log n \log(nW))^d$ space. As there are $O(n^{1/c})$ nodes in the hierarchy tree, and for each node we need to store $O(n^2)$ decision trees, the total space complexity is $n^{3+1/c} \cdot O(h\epsilon^{-1} \log n \log(nW))^d$.

¹²Recall that we use “vertex” for nodes in the input graph, and “node” for nodes in the decision trees and the hierarchy tree.

3.2 Query algorithm. Given u, v and a set of failed vertices D , by Lemma 2.3 we first find a path $U_1(= V), U_2 \dots, U_p$ in the hierarchy tree, and set $U_{p+1} = \emptyset$. Let $\mathcal{T} = \bigcup_{i=1}^p \mathcal{T}_{U_{i+1}}(U_i)$, then by Lemma 2.3, the pseudo-degrees of all $f \in D$ in every tree in \mathcal{T} is at most s .

As in [20], the query algorithm builds an auxiliary graph H , but the definition of H is different from [20]. The query algorithm builds H as Definition 3.2, and outputs $|\pi_H(u, v)|$ as an $(1 + \epsilon)$ -approximation of $|\pi_{G-D}(u, v)|$.¹³

DEFINITION 3.2. (Graph H)

- For a failure $f \in D$ and a tree $T \in \mathcal{T}$, if $f \in V(T)$, then we define the neighbors of f in T as

$$N_T(f) = \{\text{parent}_T(f)\} \cup (\text{children}_T(f) \cap \text{Trunk}(T)).$$

In other words, $N_T(f)$ consists of the parent of f in T , and the set of children of f in T which are trunk vertices. (Note that if $f \in V(T) \setminus \text{Trunk}(T)$, then it is possible that $\text{parent}_T(f) \notin \text{Trunk}(T)$.)

- Define

$$N(f) = \bigcup_{T \in \mathcal{T}} N_T(f).$$

That is, $N(f)$ is the union of $N_T(f)$ ’s over all trees $T \in \mathcal{T}$ such that $f \in T$.

- The vertex set of the auxiliary graph H is

$$V(H) = \left(\{u, v\} \cup \bigcup_{f \in D} N(f) \right) \setminus D.$$

For each $x, y \in V(H)$, the weight of the edge (x, y) in H is equal to $\text{DecTree}(x, y, D)$, as defined in Algorithm 2.

Note that in $V(H)$, vertices except u and v are defined independently from u and v . By Lemma 2.3, for every $f \in D$ and $T \in \mathcal{T}$ containing f , we have $|N_T(f)| \leq s + 1$. Therefore $|V(H)| \leq dp \cdot 2e \ln^2 n \cdot (s + 1) + 2 = O(d^{c+2} h \log^4 n)$.

¹³The vertex set of H corresponds to the set of “intermediate vertices” (also called H) in Section 1.3.

Algorithm 2 Algorithm DecTree

```

1: function DecTree( $u, v, D$ )
2:    $\alpha \leftarrow \text{root}(u, v)$ 
3:   while  $D \cap P_\alpha \neq \emptyset$  do
4:      $f \leftarrow$  the vertex in  $D \cap P_\alpha$  with the highest
       level, breaking ties arbitrarily
5:      $\alpha \leftarrow \text{ch}(\alpha, \text{seg}(f, P_\alpha), l(f))$   $\triangleright$  Recall  $l(f)$  is
       the level of  $f$ , i.e. the largest  $l$  s.t.  $f \in U_l$ .
6:   return  $|P_\alpha|$ 

```

Consider Algorithm 2. After each iteration, the set $D \cap \text{avoid}(\alpha)$ will contain at least one new vertex (namely f). When $|D \cap \text{avoid}(\alpha)| = d$, we will have that $D \cap P_\alpha = \emptyset$, and the algorithm terminates. Thus the algorithm executes at most d iterations. It is easy to see that each iteration only requires $O(d)$ time, thus the time complexity of Algorithm 2 is $O(d^2)$.


It takes $O(|V(H)|^2 d^2)$ time to build the graph H , and $O(|V(H)|^2)$ time to compute $|\pi_H(u, v)|$. Therefore, the query algorithm runs in $O(|V(H)|^2 d^2) = O(d^{2c+6} h^2 \log^8 n)$ time.

Since finally $D \cap P_\alpha = \emptyset$, we have the following observation:

OBSERVATION 3.1. For all $u, v \in V(H)$, $\text{DecTree}(u, v, D) \geq |\pi_{G-D}(u, v)|$.

For every α and f considered in Algorithm 2, let $\alpha_{\text{next}} = \text{ch}(\alpha, \text{seg}(f, P_\alpha), l(f))$ be the next decision tree node the algorithm considers. If the optimal path $\pi_{G-D}(u, v)$ never intersects the set $\text{avoid}(\alpha_{\text{next}}) \setminus \text{avoid}(\alpha)$ in any iteration, then $\text{DecTree}(u, v, D)$ would indeed return the optimal path $\pi_{G-D}(u, v)$. However, if $\pi_{G-D}(u, v)$ goes through some non-failure vertex $x \in \text{avoid}(\alpha_{\text{next}}) \setminus \text{avoid}(\alpha)$, then x is close to f , and we will show that there is some $w \in V(H)$ close to x , so the optimal path can be approximated by $\pi_{G-D}(u, w) \circ \pi_{G-D}(w, v)$. This is illustrated in the following important lemma. Notice that $\text{avoid}(\alpha_{\text{next}}) \setminus \text{avoid}(\alpha) = \text{seg}(f, P_\alpha) \cap U_{l(f)}$, so in this case, the level of x is no less than the level of f , i.e. $l(x) \geq l(f)$.

LEMMA 3.2. Given a failed vertex f and a non-failure vertex x such that $l(x) \geq l(f)$, if there is a path P in G between x and f which contains no other failed vertices, then there is a vertex $w \in V(H)$ such that $|\pi_{G-D}(x, w)| \leq (2k-1)|P|$.



Proof. Let y be the vertex with the highest level on $P[x, f]$, and suppose $j = l(y)$. Then $l(y) = j \geq l(x) \geq l(f)$. Since there are no vertices on P with level $> j$, P is in $G - U_{j+1}$. Let $T_j(y, f)$ be the tree in the tree cover $\mathcal{T}_{U_{j+1}}(U_j)$, such that the distance between y and f in $T_j(y, f)$ is at most $(2k-1)|\pi_{G-U_{j+1}}(y, f)|$. Let $P' = \mathcal{P}_j(y, f)$ be the path between y and f in $T_j(y, f)$. Let f' be the first failed vertex on $P'[y, f]$, and consider the predecessor w of f' on the path $P'[y, f]$. (That is, $P'[y, w]$ is intact from failures.) Since P' is a path on the tree $T_j(y, f)$, w is either the parent of f' or a child of f' in this tree.

- If w is the parent of f' , then $w \in V(H)$ by the definition of $V(H)$.
- If w is a child of f' , then y is a descendant of w and f' . Since $l(y) = j$, y is a trunk vertex in $T_j(y, f)$. It follows that w , as an ancestor of y , is also a trunk vertex in $T_j(y, f)$, therefore $w \in V(H)$.

Therefore, in either case, we have $w \in V(H)$. Since


$$\begin{aligned}
 |\pi_{G-D}(x, w)| &\leq |P[x, y]| + |P'[y, w]| \\
 &\leq |P[x, y]| + |P'| \\
 &\leq |P[x, y]| + (2k-1)|\pi_{G-U_{j+1}}(y, f)| \\
 &\leq |P[x, y]| + (2k-1)|P[y, f]| \\
 &\leq (2k-1)|P|,
 \end{aligned}$$

the lemma is true. \square

3.3 Proof of correctness. In this section, we show that $|\pi_H(u, v)| \leq (1 + \epsilon)|\pi_{G-D}(u, v)|$, proving the correctness of the query algorithm.

From the algorithm $\text{DecTree}(u, v, D)$, the path we get is the shortest path between u and v in the graph $G - \text{avoid}(\alpha_{\text{last}})$, where α_{last} is the last visited decision tree node of the algorithm. As we discussed before, if the real shortest path $\pi_{G-D}(u, v)$ does not go through any vertex in $\text{avoid}(\alpha_{\text{last}})$, then $\text{DecTree}(u, v, D)$ will return the correct answer. Otherwise, as $\text{avoid}(\alpha_{\text{last}})$ is the union of $\leq d$ sets of the form $\text{seg}(f, P_\alpha) \cap U_i$, $\pi_{G-D}(u, v)$ must go through some vertex x in a set $\text{seg}(f, P_\alpha) \cap U_i$. We can show that such x will be “close” to a vertex w in $V(H)$ (by Lemma 3.2), so we can use the vertices in $V(H)$ as intermediate vertices to obtain an approximate shortest path.

LEMMA 3.3. In the query algorithm $\text{DecTree}(u, v, D)$, let α be a decision tree node it encounters, $f \in D$ be the failed vertex which is selected in Line 4 of Algorithm 2 and $i = l(f)$. (That is, f is the vertex in $D \cap P_\alpha$ with the highest level.) For any non-failure vertex x in $\text{seg}(f, P_\alpha) \cap U_i$, there is a vertex $w \in V(H)$ such that $|\pi_{G-D}(x, w)| \leq \epsilon_1 \min\{|P_\alpha[u, x]|, |P_\alpha[x, v]|\}$.




Proof. Let $f' \in D$ be the failed vertex closest to x on the segment $\text{seg}(f, P_\alpha)$, then there are no failed vertices in $P_\alpha[x, f']$ or $P_\alpha(f', x)$. (W.l.o.g. we assume it is $P_\alpha[x, f']$.) We have $l(f) \geq l(f')$ by Algorithm 2. As $x \in U_i = U_{l(f)}$, we have $l(x) \geq l(f)$, hence $l(x) \geq l(f')$. By Lemma 3.2, there is a vertex $w \in V(H)$ such that $|\pi_{G-D}(x, w)| \leq (2k-1)|P_\alpha[x, f']|$. Since $|P_\alpha[x, f']| \leq \epsilon_2 \min\{|P_\alpha[u, x]|, |P_\alpha[x, v]|\}$ and $\epsilon_2 = \epsilon_1/(2k-1)$, the lemma holds. \square

We show that for $u, v \in V(H)$, if the optimal path $\pi_{G-D}(u, v)$ is *not* found by $\text{DecTree}(u, v, D)$, then we can indeed find some $w \in V(H)$ such that $\pi_{G-D}(u, w) \circ \pi_{G-D}(w, v)$ is a good approximation of $\pi_{G-D}(u, v)$. Moreover, one of $\pi_{G-D}(u, w)$ or $\pi_{G-D}(w, v)$ can be dealt with by Algorithm 2, therefore we only need to “recurse” on the other one.

LEMMA 3.4. *Let $u, v \in V(H)$ and $P = \pi_{G-D}(u, v)$. If $\text{DecTree}(u, v, D) > |P|$, there exist $x \in P \setminus \{u, v\}, y \in \{u, v\}, w \in V(H)$ such that*

- (a) $|\pi_{G-D}(x, y)| \leq \frac{1}{2}|\pi_{G-D}(u, v)|$,
- (b) $|\pi_{G-D}(x, w)| \leq \epsilon_1|\pi_{G-D}(x, y)|$, and
- (c) $\text{DecTree}(y, w, D) \leq |\pi_{G-D}(y, x)| + |\pi_{G-D}(x, w)|$, which is smaller than $|\pi_{G-D}(u, v)|$.



Proof. First we prove that there exists a triple (x, y, w) that satisfies (a) and (b).

Let α be the last decision tree node visited by $\text{DecTree}(u, v, D)$ such that $\text{avoid}(\alpha) \cap P = \emptyset$. Since $\text{DecTree}(u, v, D) > |P|$, the procedure $\text{DecTree}(u, v, D)$ did not terminate at α , i.e. it visited a child α_{next} of α such that $P \cap \text{avoid}(\alpha_{\text{next}}) \neq \emptyset$. Recall that $\text{avoid}(\alpha_{\text{next}}) \setminus \text{avoid}(\alpha) = \text{seg}(f, P_\alpha) \cap U_{l(f)}$ where f is the failure selected by Line 4 of Algorithm 2. Therefore P reaches some vertex $x \in \text{seg}(f, P_\alpha) \cap U_{l(f)}$. Since P_α is the shortest u - v path in $G - \text{avoid}(\alpha)$ and P is some u - v path in $G - \text{avoid}(\alpha)$, we know that $|P_\alpha[u, x]| \leq |P[u, x]|$ and $|P_\alpha[x, v]| \leq |P[x, v]|$.

By Lemma 3.3, there is a vertex $w \in V(H)$ such that

$$\begin{aligned} |\pi_{G-D}(x, w)| &\leq \epsilon_1 \min\{|P_\alpha[u, x]|, |P_\alpha[x, v]|\} \\ &\leq \epsilon_1 \min\{|P[u, x]|, |P[x, v]|\}. \end{aligned}$$

Let y be the endpoint in $\{u, v\}$ that is closer to x , then (x, y, w) satisfies (a) and (b).

Among all triples $x \in P \setminus \{u, v\}, y \in \{u, v\}, w \in V(H)$ satisfying (b), we pick a triple minimizing $|\pi_{G-D}(x, y)|$, and in case of a tie choose a triple minimizing $|\pi_{G-D}(x, w)|$. It is easy to see that (a) is also satisfied. In the following we prove that (c) is satisfied.

We compare the path $P' = \pi_{G-D}(y, x) \circ \pi_{G-D}(x, w)$ between y and w , with the path returned by $\text{DecTree}(y, w, D)$. For the sake of contradiction, suppose $|P'| < \text{DecTree}(y, w, D)$. Let α' be the last decision tree node visited in $\text{DecTree}(y, w, D)$ such that $\text{avoid}(\alpha') \cap P' = \emptyset$. We can also see that P' reaches some vertex $x' \in \text{seg}(f', P_{\alpha'}) \cap U_{l(f')}$, where f' is the failure selected by Line 4 of Algorithm 2. We use Lemma 3.3 again and conclude that there is a vertex $w' \in V(H)$ such that

$$\begin{aligned} |\pi_{G-D}(x', w')| &\leq \epsilon_1 \min\{|P_{\alpha'}[y, x']|, |P_{\alpha'}[x', w]|\} \\ &\leq \epsilon_1 \min\{|P'[y, x']|, |P'[x', w]|\}. \end{aligned}$$

Since $x' \in P' = P'[y, x] \circ P'[x, w]$, there are two cases. (See Fig. 3.)

- If $x' \in P'[y, x]$, then $|\pi_{G-D}(x', w')| \leq \epsilon_1|\pi_{G-D}(x', y)|$, i.e. the triple (x', y, w') also satisfies (b). Since $|\pi_{G-D}(x', y)| < |\pi_{G-D}(x, y)|$, this contradicts our choice of (x, y, w) .
- If $x' \in P'[x, w]$, then $|\pi_{G-D}(x, w')| \leq |P'[x, x']| + |\pi_{G-D}(x', w')| \leq |P'[x, x']| + \epsilon_1|P'[x', w]|$. As $\epsilon_1 < 1$, we have $|\pi_{G-D}(x, w')| < |\pi_{G-D}(x, w)|$, and (x, y, w') also satisfies (b), contradicting our choice of (x, y, w) .

Hence it must be true that $|P'| \geq \text{DecTree}(y, w, D)$.

\square

By these lemmas, we can now prove our desired approximation ratio.

THEOREM 3.1. *For every pair $u, v \in V(H)$, the query algorithm in Section 3.2 returns an $(1 + \epsilon)$ -approximation of $|\pi_{G-D}(u, v)|$.*

Proof. It is easy to see that $|\pi_H(u, v)| \geq |\pi_{G-D}(u, v)|$ for every $u, v \in V(H)$. We prove $|\pi_H(u, v)| \leq (1 + \epsilon)|\pi_{G-D}(u, v)|$ below.




Figure 3: Two cases in the proof of Lemma 3.4.

We sort all pairs of vertices $u, v \in V(H)$ ($u \neq v$) by increasing order of $|\pi_{G-D}(u, v)|$, and prove by induction that $|\pi_H(u, v)| \leq (1 + \epsilon)|\pi_{G-D}(u, v)|$ on this order. For the u, v having the smallest $|\pi_{G-D}(u, v)|$, if $\text{DecTree}(u, v, D) > |\pi_{G-D}(u, v)|$, from Lemma 3.4, there exist $y, w \in V(H)$ so that $|\pi_{G-D}(y, w)| < |\pi_{G-D}(u, v)|$, which is a contradiction. Therefore $|\pi_H(u, v)| = \text{DecTree}(u, v, D) = |\pi_{G-D}(u, v)|$.

Fix some $u, v \in V(H)$, assume that for all pairs $u', v' \in V(H)$ such that $|\pi_{G-D}(u', v')| < |\pi_{G-D}(u, v)|$, it is true that $|\pi_H(u', v')| \leq (1 + \epsilon)|\pi_{G-D}(u', v')|$. If $\text{DecTree}(u, v, D) = |\pi_{G-D}(u, v)|$ then $|\pi_H(u, v)| \leq (1 + \epsilon)|\pi_{G-D}(u, v)|$ follows trivially. Otherwise we use Lemma 3.4 to obtain a triple (x, y, w) , where $x \in \pi_{G-D}(u, v) \setminus \{u, v\}$, $y \in \{u, v\}$, and $w \in V(H)$. We assume w.l.o.g. $y = u$, then $|\pi_{G-D}(w, x)| \leq \epsilon_1 |\pi_{G-D}(u, x)|$. Since $\epsilon_1 < 1$, $|\pi_{G-D}(w, v)| \leq |\pi_{G-D}(w, x)| + |\pi_{G-D}(x, v)| \leq \epsilon_1 |\pi_{G-D}(u, x)| + |\pi_{G-D}(x, v)| < |\pi_{G-D}(u, v)|$, thus $|\pi_H(w, v)| \leq (1 + \epsilon)|\pi_{G-D}(w, v)|$ by induction hypothesis. We have:

$$\begin{aligned}
& |\pi_H(u, v)| \\
& \leq \text{DecTree}(u, w, D) + |\pi_H(w, v)| \\
& \leq |\pi_{G-D}(u, x)| + |\pi_{G-D}(x, w)| + (1 + \epsilon)|\pi_{G-D}(w, v)| \\
& \leq |\pi_{G-D}(u, x)| + \epsilon_1 |\pi_{G-D}(u, x)| \\
& \quad + (1 + \epsilon)(\epsilon_1 |\pi_{G-D}(u, x)| + |\pi_{G-D}(x, v)|) \\
& \leq (1 + \epsilon_1 + (1 + \epsilon)\epsilon_1) |\pi_{G-D}(u, x)| + (1 + \epsilon)|\pi_{G-D}(x, v)| \\
& \leq (1 + \epsilon)|\pi_{G-D}(u, v)|.
\end{aligned}$$

□

We conclude that there is a VSDO with

$$\begin{aligned}
& \text{space complexity} \quad n^{3+1/c} \cdot O\left(\frac{\epsilon^{-1} \log^2 n \log(nW)}{\log d}\right)^d, \\
& \text{query complexity} \quad O\left(\frac{d^{2c+6} \log^{10} n}{\log^2 d}\right), \\
& \text{stretch} \quad 1 + \epsilon.
\end{aligned}$$

In Section 4, we will improve the space complexity to $n^{2+1/c+o(1)}$ when $d = o\left(\frac{\log n}{\log \log n}\right)$ and $W = \text{poly}(n)$, while increasing the query time slightly.

4 An $n^{2+1/c+o(1)}$ -Space $(1 + \epsilon)$ -Stretch oracle

In this section, we discuss the modifications needed to reduce the space complexity to $n^{2+1/c+o(1)}$. Here we set $\epsilon_3 = \frac{\epsilon}{2|V(H)|}$, and $\epsilon_4 = \epsilon_3/(4k - 2)$, where $V(H)$ is

the same as in Section 3 (and Lemma 3.2 still holds), while $E(H)$ are recomputed in this section. We assume that ϵ is small enough, in particular that $\epsilon < 1$ and $\epsilon_4 < \sqrt{2} - 1$.

4.1 A structural theorem. Similar to [20], the main idea is, instead of storing the paths P_α as-is in every node α of $FT(u, v)$, we store an implicit representation of these paths. If the representation has size $\text{poly}(\log(nW), \epsilon^{-1})$ instead of $\Omega(n)$, then our data structure has space complexity $n^{2+1/c+o(1)}$.

In [20], the authors defined *k-decomposable paths*, which are paths that can be represented as the concatenation of at most $k + 1$ shortest paths in G , interleaved with at most k edges. They relied on the fact (Theorem 2 of [3]) that any k -edge-failure shortest path is a k -decomposable path in G , therefore has a succinct representation. Unfortunately, the analogue of this statement in [20] in case of vertex failures does not hold. Even if we only remove *one* vertex (i.e. $|D| = 1$), a shortest path in $G - D$ might not be a k -decomposable path for $k = o(n)$.¹⁴

In this section, we prove a structural theorem similar to the above fact used in [20]. Before we proceed, we need some definitions.

From Lemma 3.4 we can see that for any $u, v \in V(H)$, if the path $\pi_{G-D}(u, v)$ is ϵ_1 -far away from $V(H)$ in the following sense, then $\text{DecTree}(u, v, D)$ indeed finds the distance between u and v in $G - D$:

DEFINITION 4.1. *We say that a path P from u to v is ϵ -far away from $V(H)$ if there are no vertices $x \in P \setminus \{u, v\}, w \in V(H)$ such that $|\pi_{G-D}(x, w)| \leq \epsilon \cdot \min\{|P[u, x]|, |P[x, v]|\}$. (See Fig. 4.)*

Instead of considering all d -failure shortest paths, we only study the ones which are ϵ_3 -far away from $V(H)$. We will use the concept of k -expath as in [20] and re-define it as ϵ_4 -segment expath. Also, instead of considering the concatenation of at most $k + 1$ shortest paths in the original graph G , every segment here is a

¹⁴Consider an unweighted graph $G = (V, E_1 \cup E_2)$ where $V = \{v_i : 0 \leq i \leq n\}$, $E_1 = \{(v_0, v_i) : 1 \leq i \leq n\}$ and $E_2 = \{(v_i, v_{i+1}) : 1 \leq i < n\}$. Then $\pi_{G-\{v_0\}}(v_1, v_n)$ is not a $0.1n$ -decomposable path.





Figure 4: If P is far away from $V(H)$, it means that a certain “diamond”-shaped area does not contain vertices $w \in V(H)$.

shortest path in some G_i . (Recall that G_i is the induced subgraph of G on all vertices of level $\leq i$.)

DEFINITION 4.2. A path P in G is an ϵ -segment expath if the following holds. If we partition P into ϵ -segments as in Definition 3.1, then for every segment $P[x, y]$, there is some $1 \leq i \leq p$ such that $P[x, y]$ is a shortest path in G_i .

The following structural theorem for shortest paths ϵ -far away from $V(H)$ will be crucial to us. Interestingly, it is a consequence of Lemma 3.2.

THEOREM 4.1. For $u, v \in V(H)$, if $\pi_{G-D}(u, v)$ is ϵ_3 -far away from $V(H)$, then it is an ϵ_4 -segment expath.



Proof. Let $P = \pi_{G-D}(u, v)$ and $P[x, y]$ be an ϵ_4 -segment of P . W.l.o.g. assume that x and y are in the first half of P , and x is closer to u than y . By Lemma 3.1, $|P[x, y]| \leq \epsilon_4 |P[u, x]|$. (Recall that $P[x, y]$ is an ϵ_4 -segment.) Consider the vertex z with the highest level on $P[x, y]$, and let its level be $i = l(z)$. Then $P[x, y]$ is a path in G_i . If it is not the shortest path $\pi_{G_i}(x, y)$, then $\pi_{G_i}(x, y)$ must go through some failed vertex in D . (Since otherwise we can find a path in $G - D$ shorter than $\pi_{G-D}(u, v)$.) Let $P' = \pi_{G_i}(x, y)$ and f be the failed vertex on P' closest to x .

Since f is in the graph G_i , we have $l(f) \leq i = l(z)$. There is a path $P'' = P[z, x] \circ P'[x, f]$ connecting z and f that does not go through other failed vertices. By Lemma 3.2, there is a vertex $w \in V(H)$ such that $|\pi_{G-D}(z, w)| \leq (2k - 1)|P''|$. We have

$$\begin{aligned} |\pi_{G-D}(z, w)| &\leq (2k - 1)(|P[z, x]| + |P'[x, f]|) \\ &\leq 2(2k - 1)|P[x, y]| \\ &\leq 2(2k - 1)\epsilon_4|P[u, x]| \\ &\leq \epsilon_3|P[u, z]|, \end{aligned}$$

which contradicts that $\pi_{G-D}(u, v)$ is ϵ_3 -far away from $V(H)$. Therefore, $P[x, y]$ is a shortest path in G_i . \square

4.2 New data structure. We generalize the concept of ϵ_4 -segment expath to ϵ_4 -expath by adding more flexibility.

DEFINITION 4.3. Let $B = \lceil \log_{1+\epsilon_4}(nW) \rceil$. An ϵ_4 -expath P from u to v in G is a path which is a concatenation of subpaths P_0, \dots, P_{2B+1} interleaved with at most $2B + 3$ edges¹⁵, such that the following hold.

- For every $P_k = P[u_k, v_k]$ ($0 \leq k \leq 2B + 1$), P_k is either empty, or a shortest path in G_i for some level $1 \leq i \leq p$.
- If $k < B + 1$, then $|P[u, v_k]| \leq (1 + \epsilon_4)^k$; if $k \geq B + 1$, then $|P[u_k, v]| \leq (1 + \epsilon_4)^{2B+1-k}$.

LEMMA 4.1. An ϵ_4 -segment expath P from u to v is an ϵ_4 -expath.

Proof. Let $j = \lfloor \log_{1+\epsilon_4}(|P|/2) \rfloor + 1$, since $1 + \epsilon_4 < 2$, we have $j < B + 1$. Let P_1, \dots, P_j be the ϵ_4 -segments (possibly empty) in the first half of P such that for every $1 \leq k \leq j$ and $x \in P_k = P[u_k, v_k]$, $\lfloor \log_{1+\epsilon_4} |P[u, x]| \rfloor = k - 1$. Then $|P[u, v_k]| \leq (1 + \epsilon_4)^k$, which satisfies the definition of ϵ_4 -expath. The second half of P is symmetric. \square

Recall that our data structure in Section 3 consists of $O(n^2)$ decision trees, one for each pair $u, v \in V$. Each decision tree node α stores a path P_α , a subset $\text{avoid}(\alpha)$ of V , and the links to its children. The query algorithm builds an auxiliary graph H on the vertex set $V(H)$ defined in Definition 3.2, and uses Algorithm 2 to determine the edge weights in H . At last we output $|\pi_H(u, v)|$ as the approximation of $|\pi_{G-D}(u, v)|$. Our improved data structure also fits into this high-level description, but there are some small changes:

- For every $1 \leq i \leq p$, we also store the shortest path distance matrix of G_i .
- We use ϵ_4 in the definition of segments.
- In every node $\alpha \in FT(u, v)$, we store the shortest ϵ_4 -expath (instead of the general shortest path) from u to v in $G - \text{avoid}(\alpha)$, still denoted as P_α . To save space, for every subpath $P_k = [u_k, v_k]$ which is a shortest path in some G_i , we only need to store a triple (u_k, v_k, i) .
- To check whether f is in a path P_α , for every subpath $P_k = [u_k, v_k]$ which is a shortest path in some G_i , we check whether $|\pi_{G_i}(u_k, f)| + |\pi_{G_i}(f, v_k)| = |\pi_{G_i}(u_k, v_k)|$. By the uniqueness assumption of

¹⁵That is, the concatenation of $e_0, P_0, e_1, P_1, \dots, P_{2B+1}, e_{2B+2}$ where each e_i is either empty or an edge.

shortest paths (see [28]), this method can locate a vertex in P_α .


We now prove the correctness of this data structure, i.e. $|\pi_H(u, v)|$ is always an $(1 + \epsilon)$ -approximation of $|\pi_{G-D}(u, v)|$.

First, it is easy to check that Lemma 3.3 holds for parameter $(2k - 1)\epsilon_4 = \epsilon_3/2$, as follows.

REMINDER OF LEMMA 3.3. *In the query algorithm DecTree(u, v, D), let α be a decision tree node it encounters, $f \in D$ be the failed vertex which is selected in Line 4 of Algorithm 2 and $i = l(f)$. (That is, f is the vertex in $D \cap P_\alpha$ with the highest level.) For any non-failure vertex x in $\text{seg}(f, P_\alpha) \cap U_i$, there is a vertex $w \in V(H)$ such that $|\pi_{G-D}(x, w)| \leq (\epsilon_3/2) \min\{|P_\alpha[u, x]|, |P_\alpha[x, v]|\}$.*

Recall that Lemma 3.4 shows that, in the data structure in Section 3, any shortest path ϵ_1 -far away from $V(H)$ can be found by DecTree. We show that this is also true in the new data structure, where “ ϵ_1 -far away” is changed to “ ϵ_3 -far away”.

LEMMA 4.2. *Let $u, v \in V(H)$, and $P = \pi_{G-D}(u, v)$. If $\text{DecTree}(u, v, D) > |P|$, then P is not ϵ_3 -far away from $V(H)$.*



Proof. For the sake of contradiction, assume P is ϵ_3 -far away from $V(H)$. By Theorem 4.1, P is an ϵ_4 -segment expath.


Let α be the last decision tree node visited by DecTree(u, v, D) such that $\text{avoid}(\alpha) \cap P = \emptyset$. Since $\text{DecTree}(u, v, D) > |P|$, P reaches some vertex $x \in \text{avoid}(\alpha_{\text{next}}) \setminus \text{avoid}(\alpha)$, where α_{next} is the next decision tree node visited by DecTree(u, v, D) after α . Recall that $\text{avoid}(\alpha_{\text{next}}) \setminus \text{avoid}(\alpha) = \text{seg}(f, P_\alpha) \cap U_{l(f)}$, where f is the failed vertex chosen in Line 4 of Algorithm 2. By Lemma 3.3, there is a vertex $w \in V(H)$ such that $|\pi_{G-D}(x, w)| \leq (\epsilon_3/2) \min\{|P_\alpha[u, x]|, |P_\alpha[x, v]|\}$.

As P_α is the shortest ϵ_4 -expath from u to v in $G - \text{avoid}(\alpha)$, and P is some such path, we have $|P| \geq |P_\alpha|$. We will prove $|P_\alpha[u, x]| \leq 2|P[u, x]|$ and $|P_\alpha[x, v]| \leq 2|P[x, v]|$, then it will follow that $|\pi_{G-D}(x, w)| \leq \epsilon_3 \min\{|P[u, x]|, |P[x, v]|\}$, contradicting that P is ϵ_3 -far away from $V(H)$. We only prove $|P_\alpha[u, x]| \leq 2|P[u, x]|$, and the case that $|P_\alpha[x, v]| \leq 2|P[x, v]|$ is symmetric.

Suppose $|P[u, x]| < |P_\alpha[u, x]|/2$, we claim that the path $P[u, x] \circ P_\alpha[x, v]$ is a valid ϵ_4 -expath. Since


$|P[u, x]| < |P_\alpha|/2 \leq |P|/2$, x is closer to u than to v in P . Suppose P_α is composed of subpaths $P_0^\alpha, \dots, P_{2B+1}^\alpha$ interleaved with $\leq 2B + 3$ edges, and P is composed of segments P_1, \dots, P_ℓ . (Every P_j^α and P_j is a shortest path in some G_i .) Recall from the proof of Lemma 4.1 that, if x is in the first half of P , and $x \in P_k$, then $\lfloor \log_{1+\epsilon_4} |P[u, x]| \rfloor = k - 1$.

- Let $x \in P_j^\alpha$, then $j \geq \lfloor \log_{1+\epsilon_4} |P_\alpha[u, x]| \rfloor$. This is because if $j < B + 1$ (recall that $B = \lceil \log_{1+\epsilon_4}(nW) \rceil$ as in Definition 4.3), then $|P_\alpha[u, x]| \leq (1 + \epsilon_4)^j$.
- Let $x \in P_{j'}$, then $j' = \lfloor \log_{1+\epsilon_4} |P[u, x]| \rfloor + 1$. Since $(1 + \epsilon_4)^2 < 2 \leq |P_\alpha[u, x]|/|P[u, x]|$, we have $j' \leq \lfloor \log_{1+\epsilon_4} |P_\alpha[u, x]| \rfloor - 1 \leq j - 1$.



Let $P' = P[u, x] \circ P_\alpha[x, v]$, Consider the following representation of P' as $P'_0, P'_1, \dots, P'_{2B+1}$:

- For $0 \leq i < j'$, $P'_i = P_i$.
- For $i = j'$, $P'_{j'} = P_{j'}[u_{j'}, x]$, where $u_{j'}$ is the endpoint of $P_{j'}$ that lies on $P[u, x]$.
- For $j' < i < j$, $P'_i = \emptyset$.
- For $i = j$, $P'_j = P_j^\alpha[x, v_j^\alpha]$, where v_j^α is the endpoint of P_j^α that lies on $P_\alpha[x, v]$.
- For $j < i \leq 2B + 1$, $P'_i = P_i^\alpha$.



We need to verify that the representation $P'_0, P'_1, \dots, P'_{2B+1}$ satisfies the definition of ϵ_4 -expath. Let u'_i, v'_i be the endpoints of P'_i , i.e. $P'_i = P'[u'_i, v'_i]$, then:

- Case I: $i \leq j'$ (i.e. Items i and ii). In this case, $i < B + 1$, as P'_i lies in the first half of P . Since $|P'[u, v'_i]| = |P[u, v'_i]| \leq (1 + \epsilon_4)^i$, Definition 4.3 is satisfied.
- Case II: $j \leq i < B + 1$. In this case, $|P'[u, v'_i]| = |P[u, x]| + |P_\alpha[x, v'_i]| < |P_\alpha[u, v'_i]| \leq (1 + \epsilon_4)^i$, thus Definition 4.3 is satisfied.
- Case III: $j \geq B + 1$. In this case, $|P'[u'_i, v]| = |P_\alpha[u'_i, v]| \leq (1 + \epsilon_4)^{2B+1-i}$, thus Definition 4.3 is satisfied.

We conclude that P' is a valid ϵ_4 -expath. Since $|P'| < |P_\alpha|$, this contradicts the choice of P_α .

Therefore $|P_\alpha[u, x]| \leq 2|P[u, x]|$, and by symmetry, $|P_\alpha[x, v]| \leq 2|P[x, v]|$. It follows that P is not ϵ_3 -far away from $V(H)$. \square

We prove the following theorem that immediately implies the approximation ratio of the algorithm.

THEOREM 4.2. *For every pair $u, v \in V(H)$, $|\pi_H(u, v)| \leq (1 + \epsilon)|\pi_{G-D}(u, v)|$.*

Proof. For the purpose of the proof, we construct a subgraph H' of H on the same set of vertices (i.e. $V(H)$), but only keep the edges (u, v) where $\pi_{G-D}(u, v)$ is ϵ_3 -far away from $V(H)$. By Lemma 4.2, the weight of every single edge (u, v) in H' is exactly $|\pi_{G-D}(u, v)|$.

We sort all pairs of vertices $u, v \in V(H)$ by nondecreasing order of $|\pi_{G-D}(u, v)|$. For every $u, v \in V(H)$, we define a u - v path in H' inductively in this order, and denote it as $p(u, v)$. The path $p(u, v)$ is defined as follows.

- If $\pi_{G-D}(u, v)$ is ϵ_3 -far away from $V(H)$, $p(u, v)$ consists of a single edge (u, v) .
- If $\pi_{G-D}(u, v)$ is not ϵ_3 -far away from $V(H)$, there exist $x \in \pi_{G-D}(u, v) \setminus \{u, v\}$, $w \in V(H)$ such that $|\pi_{G-D}(x, w)| \leq \epsilon_3 \min\{|\pi_{G-D}(u, x)|, |\pi_{G-D}(x, v)|\}$. Since $\epsilon_3 < 1$, $|\pi_{G-D}(u, w)|$ and $|\pi_{G-D}(w, v)|$ are both smaller than $|\pi_{G-D}(u, v)|$, so $p(u, w)$ and $p(w, v)$ are both well-defined. We concatenate these paths to form $p(u, v)$, i.e. we define $p(u, v) = p(u, w) \circ p(w, v)$.

Let $k(u, v)$ be the number of edges in $p(u, v)$. We prove that for every $u, v \in V(H)$,

$$|\pi_{H'}(u, v)| \leq (1 + \epsilon_3)^{k(u, v)} |\pi_{G-D}(u, v)|.$$

We proceed by induction on $k(u, v)$. When $k(u, v) = 1$, $|\pi_{H'}(u, v)| = |\pi_{G-D}(u, v)|$. Assume this is true for all pairs (u, v) such that $k(u, v) < j$, consider some (u, v) such that $k(u, v) = j$. Let x, w be the vertices selected in the construction of $p(u, v)$, then both $k(u, w)$ and $k(w, v)$ are less than j . As $|\pi_{G-D}(x, w)| \leq (\epsilon_3/2)|\pi_{G-D}(u, v)|$, we have

$$\begin{aligned} & |\pi_{H'}(u, v)| \\ & \leq |\pi_{H'}(u, w)| + |\pi_{H'}(w, v)| \\ & \leq (1 + \epsilon_3)^{j-1} (|\pi_{G-D}(u, w)| + |\pi_{G-D}(w, v)|) \\ & \leq (1 + \epsilon_3)^{j-1} (|\pi_{G-D}(u, v)| + 2|\pi_{G-D}(x, w)|) \\ & \leq (1 + \epsilon_3)^{j-1} (|\pi_{G-D}(u, v)| + 2 \frac{\epsilon_3}{2} |\pi_{G-D}(u, v)|) \\ & \leq (1 + \epsilon_3)^j |\pi_{G-D}(u, v)|. \end{aligned}$$

Thus, for every $u, v \in V(H)$,

$$\begin{aligned} |\pi_H(u, v)| & \leq |\pi_{H'}(u, v)| \\ & \leq \left(1 + \frac{\epsilon}{2|V(H)|}\right)^{|V(H)|} |\pi_{G-D}(u, v)| \\ & \leq e^{\frac{\epsilon}{2}} |\pi_{G-D}(u, v)| \\ & < (1 + \epsilon) |\pi_{G-D}(u, v)|. \quad (\text{since } \epsilon < 1) \end{aligned}$$

\square

Each ϵ_4 -expath can be stored in $O(\epsilon_4^{-1} \log(nW))$ space. Each non-leaf node in the decision tree has $O(h\epsilon_4^{-1} \log(nW))$ children. Thus we have a VSDO of

$$\begin{aligned} \text{space complexity} & \quad n^{2+1/c} \epsilon_4^{-1} \log(nW) \cdot O(h\epsilon_4^{-1} \log(nW))^d, \\ \text{query complexity} & \quad O(d^2 |V(H)|^2 \cdot \epsilon_4^{-1} \log(nW)), \\ \text{stretch} & \quad 1 + \epsilon. \end{aligned}$$

As $\epsilon_4^{-1} = O(|V(H)| \cdot \epsilon^{-1} \log n) = O(d^{c+2} \epsilon^{-1} h \log^5 n)$, the VSDO is of

$$\begin{aligned} \text{space complexity} & \quad n^{2+1/c} \cdot (\epsilon^{-1} d^c \log(nW))^{O(d)}, \\ \text{query complexity} & \quad \tilde{O}(\epsilon^{-1} d^{3c+8} \log W), \\ \text{stretch} & \quad 1 + \epsilon. \end{aligned}$$

We improve both the space complexity and query time in the next subsection.

4.3 An improvement. In Section 4.2, we use ϵ_4 -segments in the decision tree. Therefore, each decision tree node that is not a leaf has $O(\epsilon_4^{-1} \cdot h \log(nW))$ children, and each decision tree node occupies $O(\epsilon_4^{-1} \cdot \log(nW))$ space. As $\epsilon_4^{-1} = \Theta(|V(H)| \cdot \epsilon^{-1} \log n)$, this ϵ_4^{-1} factor may seem too large. In this section, we show that the $|V(H)|$ factor in ϵ_4^{-1} can be shaved.


Let $\epsilon_1 = \epsilon/(2 + \epsilon)$ as in Section 3 and $\epsilon_5 = \epsilon_1/(4k - 2)$. We will use $O(\epsilon_5^{-1} \log(nW))$ space to represent a node in the decision tree $FT(u, v)$. A first attempt would be to store the shortest ϵ_5 -expath in each node α , but we face a technical problem as follows. Suppose $\text{DecTree}(u, v, D)$ does *not* capture the shortest path $P = \pi_{G-D}(u, v)$, then by Lemma 3.4, P is *not* far from $V(H)$. In other words, there are vertices $x \in P$ and $w \in V(H)$ such that $\pi_{G-D}(x, w) \leq \epsilon_1 |P[u, x]|$. (Here we assume w.l.o.g. that x is closer to u .) Let $P_1 = \pi_{G-D}(u, x) \circ \pi_{G-D}(x, w)$, and $P_2 = \pi_{G-D}(w, v)$, we “recursively” find P_1 and P_2 and concatenate them as an approximation of P . The proof of Lemma 3.4 shows that P_1 is far away from $V(H)$, so we may attempt to use Lemma 4.2 to conclude that $|\text{DecTree}(u, w, D)| \leq |P_1|$, and we only need to “recurse” on P_2 . However, Lemma 4.2 relies on Theorem 4.1, which requires P_1 to be a *shortest* path in $G - D$, while $P_1 = \pi_{G-D}(u, x) \circ \pi_{G-D}(x, w)$ is not necessarily the shortest u - w path.

The solution is simple. If P_1 is ϵ_1 -far from $V(H)$, we can use the same proof method of Theorem 4.1, to prove that each segment of $P_1 = \pi_{G-D}(u, x) \circ \pi_{G-D}(x, w)$ is the concatenation of at most two shortest paths in some G_i and G_j . (The original Theorem 4.1 proved that each segment of $\pi_{G-D}(u, v)$ is a shortest path in some G_i .) Therefore, we define *segment bipaths*, in which each segment is the concatenation of two shortest paths in G_i and G_j , rather than one shortest path in G_i as in segment expaths.

DEFINITION 4.4. A path P in G is an ϵ_5 -segment bipath if the following holds. If we partition P into ϵ_5 -segments as in Definition 3.1, for every segment $P[u_k, v_k]$, there exist two levels i, j and a vertex $z \in P[u_k, v_k]$ such that $P[u_k, v_k] = \pi_{G_i}(u_k, z) \circ \pi_{G_j}(z, v_k)$.

The following theorem can be proved by similar arguments as Theorem 4.1.

THEOREM 4.3. For $u, v, w \in V$, let $P = \pi_{G-D}(u, v) \circ \pi_{G-D}(v, w)$. If P is ϵ_1 -far away from $V(H)$, then it is an ϵ_5 -segment bipath.



Proof. (Sketch.) Let $P[x, y]$ be a segment of P . If $v \notin P[x, y]$ then the argument of Theorem 4.1 applies to $P[x, y]$, and there is some $1 \leq i \leq p$ such that $P[x, y] = \pi_{G_i}(x, y)$. If $v \in P[x, y]$, then $P[x, v]$ and $P[v, y]$ are shortest paths in $G - D$ respectively. Let z_1 be the vertex with the highest level in $P[x, v]$, and z_2 be the vertex with the highest level in $P[v, y]$. We proceed with the same argument as in Theorem 4.1, and we can see that $P[x, v]$ is the shortest x - v path in $G_{l(z_1)}$, and $P[v, y]$ is the shortest v - y path in $G_{l(z_2)}$. \square

Similarly we can define ϵ_5 -bipaths:

DEFINITION 4.5. Let $B = \lceil \log_{1+\epsilon_5}(nW) \rceil$. An ϵ_5 -bipath P from u to v in G is a path which is a concatenation of subpaths P_0, \dots, P_{2B+1} interleaved with at most $2B + 3$ edges, such that the following hold.

- For every $P_k = P[u_k, v_k]$ ($0 \leq k \leq 2B + 1$), either P_k is empty, or there exists a vertex $z \in P[u_k, v_k]$ and two levels $1 \leq i, j \leq p$, such that $P[u_k, v_k] = \pi_{G_i}(u_k, z) \circ \pi_{G_j}(z, v_k)$.
- If $k < B + 1$, then $|P[u, v_k]| \leq (1 + \epsilon_5)^k$; if $k \geq B + 1$, then $|P[u_k, v]| \leq (1 + \epsilon_5)^{2B+1-k}$.

We also use ϵ_5 in the definition of segments when constructing decision trees $FT(u, v)$. In each node $\alpha \in FT(u, v)$, we store the shortest ϵ_5 -bipath from u to v as the path P_α . Lemma 3.3 still holds (for parameter $(2k - 1)\epsilon_5 = \epsilon_1/2$).

REMINDER OF LEMMA 3.3. In the query algorithm $\text{DecTree}(u, v, D)$, let α be a decision tree node it encounters, $f \in D$ be the failed vertex which is selected in Line 4 of Algorithm 2 and $i = l(f)$. (That is, f is the vertex in $D \cap P_\alpha$ with the highest level.) For any non-failure vertex x in $\text{seg}(f, P_\alpha) \cap U_i$, there is a vertex $w \in V(H)$ such that $|\pi_{G-D}(x, w)| \leq (\epsilon_1/2) \min\{|P_\alpha[u, x]|, |P_\alpha[x, v]|\}$.

It is easy to verify that the counterparts of Lemma 4.1 and Lemma 4.2 also hold for (segment) bipaths.

LEMMA 4.3. An ϵ_5 -segment bipath P from u to v is an ϵ_5 -bipath.

LEMMA 4.4. (Assume $1 + \epsilon_5 < \sqrt{2}$.) Let $y, w \in V(H)$, $x \in V$, and $P = \pi_{G-D}(y, x) \circ \pi_{G-D}(x, w)$. If $\text{DecTree}(y, w, D) > |P|$, then P is not ϵ_1 -far away from $V(H)$.

Proof. (Sketches of Lemma 4.3 and Lemma 4.4.) The arguments are essentially the same as Lemmas 4.1 and 4.2, except that each subpath in P and P_α is now a concatenation of two shortest paths in G_i and $G_{i'}$. This does not affect the calculation of lengths of paths in the proofs. In particular, in Lemma 4.2, the representation of $P[u, x] \circ P_\alpha[x, v]$ as ϵ_5 -bipath remains exactly the same, and it is easy to verify the validity of $P[u, x] \circ P_\alpha[x, v]$ as an ϵ_5 -bipath. \square

Recall that the query algorithm builds the graph H on vertex set $V(H)$, adds an edge of weight $\text{DecTree}(x, y, D)$ for each $x, y \in V(H)$, and outputs the value $|\pi_H(u, v)|$. We now prove that the query algorithm has stretch $1 + \epsilon$.

THEOREM 4.4. For every $u, v \in V(H)$, $|\pi_H(u, v)| \leq (1 + \epsilon)|\pi_{G-D}(u, v)|$.

Proof. For all pairs $u, v \in V(H)$, we sort the lengths $|\pi_{G-D}(u, v)|$ in nondecreasing order, and use induction on this order. For each $u, v \in V(H)$, if $\pi_{G-D}(u, v)$ is ϵ_1 -far away from $V(H)$, by Lemma 4.4, $\text{DecTree}(u, v, D) = |\pi_{G-D}(u, v)|$ and we are done. Otherwise let $P = \pi_{G-D}(u, v)$, then there are vertices $x \in P \setminus \{u, v\}$, $y \in \{u, v\}$ and $w \in V(H)$ such that $|P[x, y]| \leq \frac{1}{2}|P|$ and $|\pi_{G-D}(x, w)| \leq \epsilon_1|P[x, y]|$.

Among all such triples (x, y, w) , we choose the triple that minimizes $|P[x, y]|$, and in case of ties choose the

triple that minimizes $|\pi_{G-D}(x, w)|$. W.l.o.g. assume $y = u$. Let $P' = \pi_{G-D}(u, x) \circ \pi_{G-D}(x, w)$, if P' is not ϵ_1 -far away from $V(H)$, then there are vertices $x' \in P' \setminus \{u, w\}$, $y' \in \{u, w\}$ and $w' \in V(H)$ such that $|\pi_{G-D}(x', w')| \leq \epsilon_1 |P'[x', y']|$. The same argument as Lemma 3.4 shows that this is a contradiction to the choice of (x, y, w) (see Fig. 3):

- If $x' \in P'[u, x]$, then the triple (x', y, w') also satisfies that $|\pi_{G-D}(x', w')| \leq \epsilon_1 |P'[x', y]|$, and $|P'[x', y]| < |P[x, y]|$. So we should have chosen the triple (x', y, w') instead of (x, y, w) .
- If $x' \in P'[x, w]$, then $|\pi_{G-D}(x, w')| \leq |P'[x, x']| + |\pi_{G-D}(x', w')| \leq |P'[x, x']| + \epsilon_1 |P'[x', w]| < |\pi_{G-D}(x, w)|$ as $\epsilon_1 < 1$. So we should have chosen the triple (x, y, w') instead of (x, y, w) .

It follows that P' is ϵ_1 -far away from $V(H)$. By Lemma 4.4, we have $\text{DecTree}(u, w, D) \leq |P'| = |\pi_{G-D}(u, x)| + |\pi_{G-D}(x, w)|$. It is easy to see that $|\pi_{G-D}(w, v)| < |\pi_{G-D}(u, v)|$, thus by induction hypothesis $|\pi_H(w, v)| \leq (1 + \epsilon) |\pi_{G-D}(w, v)|$. We have

$$\begin{aligned} & |\pi_H(u, v)| \\ & \leq |\pi_H(u, w)| + |\pi_H(w, v)| \\ & \leq |\pi_{G-D}(u, x)| + |\pi_{G-D}(x, w)| \\ & \quad + (1 + \epsilon)(|\pi_{G-D}(w, x)| + |\pi_{G-D}(x, v)|) \\ & \leq |\pi_{G-D}(u, x)|(1 + \epsilon_1 + (1 + \epsilon)\epsilon_1) + (1 + \epsilon)|\pi_{G-D}(x, v)| \\ & = (1 + \epsilon)|\pi_{G-D}(u, v)|. \end{aligned}$$

□

Since an ϵ_5 -bipath occupies $O(\epsilon_5^{-1} \log(nW))$ space, and each non-leaf node has $O(h\epsilon_5^{-1} \log(nW))$ children, we have a VSDO of

$$\begin{aligned} \text{space complexity} & n^{2+1/c} \epsilon_5^{-1} \log(nW) \cdot O(h\epsilon_5^{-1} \log(nW))^d, \\ \text{query complexity} & O(d^2 |V(H)|^2 \cdot \epsilon_5^{-1} \log(nW)), \\ \text{stretch} & 1 + \epsilon. \end{aligned}$$

As $\epsilon_5^{-1} = O(\epsilon^{-1} \log n)$, the VSDO is of

$$\begin{aligned} \text{space complexity} & n^{2+1/c} \cdot \frac{\log d}{\log n} \cdot O\left(\frac{\epsilon^{-1} \log^2 n \log(nW)}{\log d}\right)^{d+1}, \\ \text{query complexity} & O\left(\frac{\epsilon^{-1} d^{2c+6} \log^{11} n \log(nW)}{\log^2 d}\right), \\ \text{stretch} & 1 + \epsilon. \end{aligned}$$

4.4 Implementation details.

Preprocessing. Given a subgraph G' of G , vertices $s, t \in V$ and $\epsilon' > 0$, we show that the shortest ϵ' -expath from s to t in G' can be computed in polynomial time.

Let $\pi'_{G'}(s, t)$ be the shortest path of the form $\min\{\pi_{G_i}(s, t) \mid 1 \leq i \leq p, \pi_{G_i}(s, t) \subseteq G'\}$ (which may be $+\infty$). First we compute $\pi'_{G'}(s, t)$ for all pairs of $s, t \in V$. Then let $\pi(s, t, j)$ be the shortest s - t path P in G' such that the following hold.

- P is the concatenation of subpaths P_0, \dots, P_j interleaved with $\leq j + 1$ edges. Moreover, denote $P_k = P[u_k, v_k]$, where u_k, v_k are endpoints of P_k and u_k is the closer-to- u one, then $v_j = t$, but there might be an edge between s and u_0 . (That is, P is the concatenation of $e_0, P_0, e_1, P_1, \dots, e_j, P_j$ where each e_i is an edge and each P_i is a subpath.)
- For every $0 \leq k \leq j$, P_k is either empty or a shortest path in G_i for some level $1 \leq i \leq p$.
- For every $0 \leq k \leq j$, $|P[s, v_k]| \leq (1 + \epsilon')^k$.

We use a dynamic programming algorithm to compute $|\pi(s, t, k)|$ for all $k \leq B$. To start with, we artificially define $|\pi(s, t, -1)|$ as:

$$|\pi(s, t, -1)| = \begin{cases} 0 & \text{if } s = t \\ +\infty & \text{if } s \neq t \end{cases}.$$

Given $\{|\pi(s, t, j - 1)|\}$ for all $s, t \in V$, we compute $|\pi(s, t, j)|$ as follows:

$$|\pi(s, t, j)| = \begin{cases} |\tilde{\pi}(s, t, j)| & \text{if } |\tilde{\pi}(s, t, j)| \leq (1 + \epsilon')^j \\ +\infty & \text{otherwise} \end{cases},$$

where

$$(4.2) \quad |\tilde{\pi}(s, t, j)| = \min_{(u, v)} \{|\pi(s, u, j-1)| + w(u, v) + |\pi'(v, t, j)|\}.$$

Then the length of shortest ϵ' -expath is

$$\min_{(u, v)} \{|\pi(s, u, B)| + w(u, v) + |\pi(v, t, B)|\}.$$

Here $w(u, v)$ is the weight of the edge between u and v . If $u = v$ then we assume $w(u, v) = 0$.

We can easily adapt the algorithm to obtain the actual shortest ϵ' -expath.

If we replace the term $\pi'_{G'}(s, t)$ in (4.2) by $\pi''_{G'}(s, t)$, which is defined as the shortest concatenated path of the form $\pi'_{G'}(s, u) \circ \pi'_{G'}(u, t)$, then we can also compute shortest ϵ' -bipaths in polynomial time. Once we have a polynomial-time algorithm for computing the shortest ϵ' -expath or ϵ' -bipath in a subgraph G' , it is easy to see that the whole preprocessing time is polynomial in the space complexity.

Query. An ϵ' -expath from u to v is stored as $O(\epsilon'^{-1} \log(nW))$ triples (x, y, l) , where each triple denotes a subpath $\pi_{G_l}(x, y)$. To check whether a failed vertex f is in an ϵ' -expath P_α , we check every subpath $\pi_{G_l}(x, y)$ whether it contains f by checking whether $\pi_{G_l}(x, f) + \pi_{G_l}(f, y) = \pi_{G_l}(x, y)$. The correctness of this method relies on the uniqueness assumption of shortest paths. If f is in P_α , we can also find the segment it is in, by computing $\lfloor \log_{1+\epsilon'} |P_\alpha[u, f]| \rfloor$ or $\lfloor \log_{1+\epsilon'} |P_\alpha[f, v]| \rfloor$.

If we store the distance matrices of each G_i during preprocessing, then every operation (i.e. checking if $f \in P_\alpha$ and locating $\text{seg}(f, P_\alpha)$) can be done in $O(\epsilon'^{-1} \log(nW))$ time. Therefore the time complexity of Algorithm 2 becomes $O(\epsilon'^{-1} \log(nW) \cdot d^2)$. Similar arguments also apply to ϵ' -bipaths.

Retrieving the actual path. The actual $(1 + \epsilon)$ -approximate shortest path can be efficiently retrieved as follows. (By *retrieving a path efficiently*, we mean finding it in $O(\ell)$ additional time, where ℓ is the number of vertices in the path.)

- For every $1 \leq i \leq p$, we also preprocess the shortest paths of G_i . That is, for every $v \in V(G_i)$, we precompute the incoming shortest path tree rooted at v . Consequently, given any $1 \leq i \leq p$ and $u, v \in V(G_i)$, we can retrieve the path $\pi_{G_i}(u, v)$ efficiently.
- Let $\alpha \in FT(u, v)$ be a decision tree node. Recall that P_α is an ϵ_4 -expath or an ϵ_5 -bipath, therefore a concatenation of $O(\epsilon_4^{-1} \log(nW))$ or $O(\epsilon_5^{-1} \log(nW))$ paths of the form $\pi_{G_i}(x, y)$. Hence, P_α can be retrieved efficiently.
- Let $u, v \in V$ and D be a set of failed vertices. We build the graph H according to Definition 3.2, and find the shortest u - v path in H . Each edge (x, y) in this path corresponds to a path returned by $\text{DecTree}(x, y, D)$, which by Algorithm 2 is P_α for some decision tree node α . The concatenation of these paths P_α for each edge on $\pi_H(u, v)$ forms an $(1 + \epsilon)$ -approximate shortest u - v path in $G - D$. As each P_α can be retrieved efficiently, this path can also be retrieved efficiently.

4.5 A reduction from arbitrary weights to bounded weights. If $W = n^{\omega(1)}$, then we may be unsatisfied with the $\log^d(nW)$ factor in the space complexity of our oracle. We can replace the $\log^d(nW)$ factor by $\log W \log^{d-1} n$ in the space complexity of our data structure, via a reduction from arbitrary weights to bounded weights. This reduction appears in [20, Lemma 4.1] and we notice that it also holds for vertex failures.

LEMMA 4.5. (LEMMA 4.1 OF [20], REPHRASED)
Suppose we have a VSDO for undirected graphs with edge weights in $[1, n^3]$, which occupies S space, needs Q query time and has stretch A . Then we can build a VSDO for undirected graphs with edge weights in $[1, W]$, which occupies $O(S \log W / \log n)$ space, needs $O(Q \log \log W)$ query time and has stretch $(1 + 1/n)A$.

Proof. For every $0 \leq i \leq \frac{\log W}{\log n}$, we build a VSDO \mathcal{O}^i on the graph \tilde{G}^i , which is defined as follows: $V(\tilde{G}^i) = V(G)$

and for each edge (u, v) of weight w in G , if $w \leq n^{i+1}$, then we have an edge (u, v) of weight $\lceil w \cdot n^{-(i-2)} \rceil$ in \tilde{G}^i . Note that the graphs \tilde{G}^i are *monotone* in the sense that, if an edge appears in \tilde{G}^i , then it also appears (albeit with a different weight) in \tilde{G}^{i+1} . Also note that the edge weights in every \tilde{G}^{i+1} is at most n^3 .

Given a query (u, v, D) , we can use binary search to find the smallest integer i such that s and t are connected in $\tilde{G}^i - D$. Then we use the oracle \mathcal{O}^i and \mathcal{O}^{i+1} to compute an A -approximation of the value

$$ans = \min\{\delta_{\tilde{G}^i - D}(u, v) \cdot n^{i-2}, \delta_{\tilde{G}^{i+1} - D}(u, v) \cdot n^{i-1}\}.$$

It remains to prove that $\delta_{G-D}(u, v) \leq ans \leq (1 + 1/n)\delta_{G-D}(u, v)$. That $\delta_{G-D}(u, v) \leq ans$ is trivial. Let \tilde{W} be the largest edge weight in $\pi_{G-D}(u, v)$, and $i_\star = \lfloor \log_n \tilde{W} \rfloor$. Since $\tilde{W} \leq n^{i_\star+1}$, u, v are connected in \tilde{G}^{i_\star} . On the other hand, every edge in G that appears in $\tilde{G}^{i_\star-2}$ has weight at most $n^{i_\star-1}$, thus if u, v are connected in $\tilde{G}^{i_\star-2}$, then $\delta_{G-D}(u, v) \leq (n-1)n^{i_\star-1} < n^{i_\star}$, contradicting the definition of i_\star . Therefore $i_\star - 1 \leq i \leq i_\star$ and $i_\star \in \{i, i+1\}$.

We have $ans \leq \delta_{\tilde{G}^{i_\star} - D}(u, v) \cdot n^{i_\star-2}$. For every edge $e \in E(G)$ with weight w , if e appears in the graph \tilde{G}^{i_\star} , then $(\lceil w \cdot n^{-(i_\star-2)} \rceil) \cdot n^{i_\star-2} \leq w + n^{i_\star-2}$, i.e. every such edge is “overestimated” by an additive error of at most $n^{i_\star-2}$. It follows that $ans \leq \delta_{G-D}(u, v) + (n-1)n^{i_\star-2}$. Since $\delta_{G-D}(u, v) \geq n^{i_\star}$, we have $ans \leq (1 + 1/n)\delta_{G-D}(u, v)$.

As our new oracle computes an A -approximation of ans , its stretch is $(1 + 1/n)A$. \square

Assuming $\epsilon > 2/n$, Lemma 4.5 transforms the VSDO in Section 3 into a VSDO of

$$\begin{aligned} \text{space complexity} & n^{3+1/c} \cdot \frac{\log W}{\log n} \cdot O\left(\frac{\epsilon^{-1} \log^3 n}{\log d}\right), \\ \text{query complexity} & O\left(\frac{d^{2c+6} \log^{10} n \log \log W}{\log^2 d}\right), \\ \text{stretch} & 1 + \epsilon, \end{aligned}$$

and the VSDO in Section 4 into a VSDO of

$$\begin{aligned} \text{space complexity} & n^{2+1/c} \cdot \frac{\log W \log d}{\log^2 n} \cdot O\left(\frac{\epsilon^{-1} \log^3 n}{\log d}\right)^{d+1}, \\ \text{query complexity} & O\left(\frac{\epsilon^{-1} d^{2c+6} \log^{12} n \log \log W}{\log^2 d}\right), \\ \text{stretch} & 1 + \epsilon. \end{aligned}$$

5 A poly(log n, d)-Stretch Oracle

In the full version of this paper, we will also present a VSDO of

$$\begin{aligned} \text{space complexity} & n^{2+1/c} \text{poly}(\log(nW), d), \\ \text{query complexity} & \text{poly}(\log(nW), d), \\ \text{stretch} & \text{poly}(\log n, d). \end{aligned}$$

We omit this result here due to space constraints.

Acknowledgments

We are grateful to anonymous reviewers for helpful comments, bringing [53] to our attention, and pointing out the recent work [42] that allows us to derandomize the oracle in Section 5. We would like to thank Thatchaphol Saranurak for providing an early manuscript of [62], and Zhijun Zhang for helpful comments on a draft version of this paper.

References

- [1] Ittai Abraham, Shiri Chechik, and Sebastian Krinninger. Fully dynamic all-pairs shortest paths with worst-case update-time revisited. In *Proc. 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 440–452, 2017.
- [2] Ittai Abraham, Shiri Chechik, and Kunal Talwar. Fully dynamic all-pairs shortest paths: Breaking the $O(n)$ barrier. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2014*, volume 28 of *LIPICs*, pages 1–16, 2014.
- [3] Yehuda Afek, Anat Bremler-Barr, Haim Kaplan, Edith Cohen, and Michael Merritt. Restoration by path concatenation: fast recovery of MPLS paths. *Distributed Computing*, 15(4):273–283, 2002.
- [4] I. Althöfer, G. Das, D. Dobkin, D. Joseph, and J. Soares. On sparse spanners of weighted graphs. *Discrete and Computational Geometry*, 9:81–100, 1993.
- [5] Surender Baswana and Neelesh Khanna. Approximate shortest paths avoiding a failed vertex: Near optimal data structures for undirected unweighted graphs. *Algorithmica*, 66(1):18–50, 2013.
- [6] Aaron Bernstein. Fully dynamic $(2 + \epsilon)$ approximate all-pairs shortest paths with fast query and close to linear update time. In *Proc. 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 693–702, 2009.
- [7] Aaron Bernstein and David Karger. Improved distance sensitivity oracles via random sampling. In *Proc. 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 34–43, 2008.
- [8] Aaron Bernstein and David Karger. A nearly optimal oracle for avoiding failed vertices and edges. In *Proc. 41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 101–110, 2009.
- [9] Davide Bilò, Keerti Choudhary, Luciano Gualà, Stefano Leucci, Merav Parter, and Guido Proietti. Efficient oracles and routing schemes for replacement paths. In *Proc. 35th Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 96 of *LIPICs*, pages 13:1–13:15, 2018.
- [10] Davide Bilò, Luciano Gualà, Stefano Leucci, and Guido Proietti. Multiple-edge-fault-tolerant approximate shortest-path trees. In *Proc. 33rd Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 47 of *LIPICs*, pages 18:1–18:14, 2016.
- [11] Davide Bilò, Luciano Gualà, Stefano Leucci, and Guido Proietti. Fault-tolerant approximate shortest-path trees. *Algorithmica*, 80(12):3437–3460, 2018.
- [12] Greg Bodwin, Michael Dinitz, Merav Parter, and Virginia Vassilevska Williams. Optimal vertex fault tolerant spanners (for fixed stretch). In *Proc. 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1884–1900, 2018.
- [13] Greg Bodwin, Fabrizio Grandoni, Merav Parter, and Virginia Vassilevska Williams. Preserving distances in very faulty graphs. In *Proc. 44th International Colloquium on Automata, Languages and Programming (ICALP)*, volume 80 of *LIPICs*, pages 73:1–73:14, 2017.
- [14] Greg Bodwin and Shyamal Patel. A trivial yet optimal solution to vertex fault tolerant spanners. In *Proc. 38th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 541–543, 2019.
- [15] Glencora Borradaile, Seth Pettie, and Christian Wulff-Nilsen. Connectivity oracles for planar graphs. In *Proc. 13th Scandinavian Symposium and Workshop on Algorithm Theory (SWAT)*, volume 7357 of *LNCS*, pages 316–327, 2012.
- [16] Panagiotis Charalampopoulos, Shay Mozes, and Benjamin Tebeka. Exact distance oracles for planar graphs with failing vertices. In *Proc. 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2110–2123, 2019.
- [17] S. Chechik, M. Langberg, David Peleg, and L. Roditty. Fault-tolerant spanners for general graphs. In *Proc. 41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 435–444, 2009.
- [18] Shiri Chechik. *Fault-tolerant structures in graphs*. PhD thesis, Weizmann Institute of Science, June 2012.
- [19] Shiri Chechik and Sarel Cohen. Distance sensitivity oracles with subcubic preprocessing time and fast query time. In *Proc. 52nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 1375–1388, 2020.
- [20] Shiri Chechik, Sarel Cohen, Amos Fiat, and Haim Kaplan. $(1 + \epsilon)$ -approximate f -sensitive distance oracles. In *Proc. 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1479–1496, 2017.
- [21] Shiri Chechik, Michael Langberg, David Peleg, and Liam Roditty. f -sensitivity distance oracles and routing schemes. *Algorithmica*, 63(4):861–882, 2012.
- [22] Keerti Choudhary. An optimal dual fault tolerant reachability oracle. In *Proc. 43rd International Colloquium on Automata, Languages and Programming (ICALP)*, volume 55 of *LIPICs*, pages 130:1–130:13, 2016.
- [23] Rezaul Alam Chowdhury and Vijaya Ramachandran. Improved distance oracles for avoiding link-failure. In *Proc. 13th International Symposium on Algorithms and Computation (ISAAC)*, volume 2518 of *LNCS*, pages 523–534, 2002.
- [24] Don Coppersmith and Shmuel Winograd. Matrix multiplication via arithmetic progressions. *Journal of Symbolic Computation*, 9(3):251–280, 1990.

- [25] Camil Demetrescu and Giuseppe F. Italiano. A new approach to dynamic all pairs shortest paths. *Journal of the ACM*, 51(6):968–992, 2004.
- [26] Camil Demetrescu and Giuseppe F. Italiano. Fully dynamic all pairs shortest paths with real edge weights. *Journal of Computer and System Sciences*, 72(5):813–837, 2006.
- [27] Camil Demetrescu and Mikkel Thorup. Oracles for distances avoiding a link-failure. In *Proc. 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 838–843, 2002.
- [28] Camil Demetrescu, Mikkel Thorup, Rezaul Alam Chowdhury, and Vijaya Ramachandran. Oracles for distances avoiding a failed node or link. *SIAM Journal of Computing*, 37(5):1299–1318, 2008.
- [29] Michael Dinitz and Robert Krauthgamer. Fault-tolerant spanners: Better and simpler. In *Proc. 30th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 169–178, 2011.
- [30] Ran Duan and Seth Pettie. Dual-failure distance and connectivity oracles. In *Proc. 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 506–515, 2009.
- [31] Ran Duan and Seth Pettie. Connectivity oracles for failure prone graphs. In *Proc. 42nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 465–474, 2010.
- [32] Ran Duan and Seth Pettie. Connectivity oracles for graphs subject to vertex failures. In *Proc. 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 490–509, 2017.
- [33] Ran Duan and Tianyi Zhang. Improved distance sensitivity oracles via tree partitioning. In *Proc. 15th International Symposium on Algorithms and Data Structures (WADS)*, volume 10389 of *LNCS*, pages 349–360, 2017.
- [34] P. Erdős. Extremal problems in graph theory. In *Proceedings of the Symposium on Theory of Graphs and its Applications*, pages 29–36, 1964.
- [35] David Gibb, Bruce M. Kapron, Valerie King, and Nolan Thorn. Dynamic graph connectivity with improved worst case update time and sublinear space. *CoRR*, abs/1509.06464, 2015.
- [36] Fabrizio Grandoni and Virginia Vassilevska Williams. Improved distance sensitivity oracles via fast single-source replacement paths. In *Proc. 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 748–757, 2012.
- [37] Manoj Gupta and Aditi Singh. Generic single edge fault tolerant exact distance oracle. In *Proc. 45th International Colloquium on Automata, Languages and Programming (ICALP)*, volume 107 of *LIPICs*, pages 72:1–72:15, 2018.
- [38] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Sublinear-time decremental algorithms for single-source reachability and shortest paths on directed graphs. In *Proc. 46th Annual ACM Symposium on Theory of Computing (STOC)*, pages 674–683, 2014.
- [39] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. A subquadratic-time algorithm for decremental single-source shortest paths. In *Proc. 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1053–1072, 2014.
- [40] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Dynamic approximate all-pairs shortest paths: Breaking the $O(mn)$ barrier and derandomization. *SIAM Journal of Computing*, 45(3):947–1006, 2016.
- [41] Bruce M. Kapron, Valerie King, and Ben Mountjoy. Dynamic graph connectivity in polylogarithmic worst case time. In *Proc. 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1131–1142, 2013.
- [42] Karthik C. S. and Merav Parter. Deterministic replacement path covering. *CoRR*, abs/2008.05421, 2020.
- [43] Valerie King. Fully dynamic algorithms for maintaining all-pairs shortest paths and transitive closure in digraphs. In *Proc. 40th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 81–91, 1999.
- [44] François Le Gall. Powers of tensors and fast matrix multiplication. In *Proc. 39th International Symposium on Symbolic and Algebraic Computation (IS-SAC)*, pages 296–303, 2014.
- [45] Merav Parter. Dual failure resilient BFS structure. In *Proc. 2015 ACM Symposium on Principles of Distributed Computing (PODC)*, pages 481–490, 2015.
- [46] Merav Parter. Fault-tolerant logical network structures. *Bulletin of the EATCS*, 118, 2016.
- [47] Merav Parter. Vertex fault tolerant additive spanners. *Distributed Computing*, 30(5):357–372, 2017.
- [48] Merav Parter and David Peleg. Fault tolerant BFS structures: A reinforcement-backup tradeoff. In *Proc. 27th ACM Symposium on Parallelism in Algorithms and Architectures*, pages 264–273, 2015.
- [49] Merav Parter and David Peleg. Sparse fault-tolerant BFS structures. *ACM Transactions on Algorithms*, 13(1):11:1–11:24, October 2016.
- [50] Merav Parter and David Peleg. Fault-tolerant approximate BFS structures. *ACM Transactions on Algorithms*, 14(1):10:1–10:15, January 2018.
- [51] Mihai Pătrașcu and Mikkel Thorup. Planning for fast connectivity updates. In *Proc. 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 263–271, 2007.
- [52] Hanlin Ren. Improved distance sensitivity oracles with subcubic preprocessing time. In *Proc. 28th European Symposium on Algorithms (ESA)*, volume 173 of *LIPICs*, pages 79:1–79:13, 2020.
- [53] Liam Roditty, Mikkel Thorup, and Uri Zwick. Deterministic constructions of approximate distance oracles and spanners. In *Proc. 32nd International Colloquium on Automata, Languages and Programming (ICALP)*, ICALP’05, pages 261–272, 2005.
- [54] Liam Roditty and Uri Zwick. Dynamic approxi-

- mate all-pairs shortest paths in undirected graphs. In *Proc. 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 499–508, 2004.
- [55] Liam Roditty and Uri Zwick. On dynamic shortest paths problems. *Algorithmica*, 61(2):389–401, October 2011.
- [56] Piotr Sankowski. Subquadratic algorithm for dynamic shortest distances. In *Proc. 11th International Computing and Combinatorics Conference (COCOON)*, volume 3595 of *LNCS*, pages 461–470, 2005.
- [57] Andrew James Stothers. *On the complexity of matrix multiplication*. PhD thesis, The University of Edinburgh, 2010.
- [58] Mikkel Thorup. Fully-dynamic all-pairs shortest paths: Faster and allowing negative cycles. In *Proc. 9th Scandinavian Symposium and Workshop on Algorithm Theory (SWAT)*, volume 3111 of *LNCS*, pages 384–396, 2004.
- [59] Mikkel Thorup. Worst-case update times for fully-dynamic all-pairs shortest paths. In *Proc. 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 112–119, 2005.
- [60] Mikkel Thorup and Uri Zwick. Approximate distance oracles. *Journal of the ACM*, 52(1):1–24, 2005.
- [61] Jan van den Brand and Danupon Nanongkai. Dynamic approximate shortest paths and beyond: Subquadratic and worst-case update time. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 436–455. IEEE, 2019.
- [62] Jan van den Brand and Thatchaphol Saranurak. Sensitive distance and reachability oracles for large batch updates. In *Proc. 60th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 424–435, 2019.
- [63] Zhengyu Wang. An improved randomized data structure for dynamic graph connectivity. *CoRR*, abs/1510.04590, 2015.
- [64] Oren Weimann and Raphael Yuster. Replacement paths and distance sensitivity oracles via fast matrix multiplication. *ACM Transactions on Algorithms*, 9(2):14:1–14:13, March 2013.
- [65] Virginia Vassilevska Williams. Multiplying matrices faster than Coppersmith-Winograd. In Howard J. Karloff and Toniann Pitassi, editors, *Proc. 44th Annual ACM Symposium on Theory of Computing (STOC)*, pages 887–898, 2012.