# ( $1+\epsilon$ )-Approximate $f$-Sensitive Distance Oracles 

Shiri Chechik * Sarel Cohen ${ }^{\dagger} \quad$ Amos Fiat ${ }^{\ddagger} \quad$ Haim Kaplan ${ }^{\text {§ }}$


#### Abstract

An $f$-Sensitive Distance Oracle with stretch $\alpha$ preprocesses a graph $G(V, E)$ and produces a small data structure that is used to answer subsequent queries. A query is a triple consisting of a set $F \subset E$ of at most $f$ edges, and vertices $s$ and $t$. The oracle answers a query ( $F, s, t$ ) by returning a value $\tilde{d}$ which is equal to the length of some path between $s$ and $t$ in the graph $G \backslash F$ (the graph obtained from $G$ by discarding all edges in $F$ ). Moreover, $\tilde{d}$ is at most $\alpha$ times the length of the shortest path between $s$ and $t$ in $G \backslash F$. The oracle can also construct a path between $s$ and $t$ in $G \backslash F$ of length $\tilde{d}$. To the best of our knowledge we give the first nontrivial $f$-sensitive distance oracle with fast query time and small stretch capable of handling multiple edge failures. Specifically, for any $f=o\left(\frac{\log n}{\log \log n}\right)$ and a fixed $\epsilon>0$ our oracle answers queries $(F, s, t)$ in time $\widetilde{O}(1)$ with $(1+\epsilon)$ stretch using a data structure of size $n^{2+o(1)}$. For comparison, the naïve alternative requires $m^{f} n^{2}$ space for sublinear query time.


## 1 Introduction.

Dealing with failures is an essential part of modern computing. Built in processes that deal with failures are an essential part of many computing environments, from massive storage devices, large scale parallel computation, and communication networks.

In this paper we study how to answer queries
*Tel Aviv University, Israel. E-mail: shiri.chechik@gmail.com. This research was supported by the Israel Science Foundation grant No. 1528/15, the Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11) and the Blavatnik Fund.
${ }^{\dagger}$ Tel Aviv University, Israel. E-mail: sarelcoh@post.tau.ac.il. This research was supported by the Israel Science Foundation grant No. 1528/15, Israel Science Foundation grant No. 1841/14, the Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11), the Blavatnik Fund and GIF grant No. 1161.
$\ddagger$ Tel Aviv University, Israel. E-mail: fiat@tau.ac.il. This research was supported by Israel Science Foundation grant No. 1841/14, the Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11) and the Blavatnik Fund.
§Tel Aviv University, Israel. E-mail: haimk@post.tau.ac.il. This research was supported by Israel Science Foundation grant No. 1841/14, the Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11), the Blavatnik Fund and GIF grant No. 1161.
in the presence of failures in a network, where one assumes apriori a bound on the number of simultaneous failures. Such a bound is especially relevant in the context of independent failures where the probability of multiple failures decreases exponentially with the number of failures. In particular, we consider distance and shortest path queries in the presence of edge failures for undirected weighted graphs.

Our goal is to preprocess a graph $G=(V, E)$ and produce some (hopefully small) data structure used to answer subsequent queries of the form $(F, s, t)$, $F \subset E,|F| \leq f$. The answer to such a query is an approximation to the length of the shortest path between $s$ and $t$ in the graph $G \backslash F$ (the graph obtained from $G$ by discarding all edges in $F$ ). Moreover, we want to be able to answer queries that ask to compute such an approximate shortest path. We call such a data structure an $f$-sensitive distance oracle.

There are several parameters to be considered with respect to suggested $f$-sensitive distance oracles.

- The size of the data structure.
- The time needed to answer distance queries $(F, s, t)$.
- The approximation ratio (also known as the stretch), how close are our estimated distances to the actual distances?
- The time required for the preprocessing phase.

The $f$-sensitive distance oracle is said to have a stretch of $\alpha$ if the estimated distance returned by the data structure is at most $\alpha$ times the actual distance. I.e., for every set $F \subset E$ such that $|F| \leq f$ and for every pair of vertices $s$ and $t$, the $f$-sensitive distance oracle returns an estimated distance $\widehat{\mathrm{d}}_{G \backslash F}(s, t)$, that satisfies $\mathrm{d}_{G \backslash F}(s, t) \leq \widehat{\mathrm{d}}_{G \backslash F}(s, t) \leq \alpha \mathrm{d}_{G \backslash F}(s, t)$ (where $\mathrm{d}_{G \backslash F}(s, t)$ is the distance from $s$ to $t$ in the graph $G \backslash F)$.

The naïve solution to the exact problem is to store for every vertex $v$ and for every set $F$ with $|F| \leq$ $f$, a shortest path tree rooted by $v$ and a distance vector giving the distances from $v$ to all other vertices. Thus, the total space required by the naïve solution is $O\left(m^{f} n^{2}\right)$. Using the naïve approach, a query $(F, s, t)$ returns the exact solution and requires $\widetilde{O}(f)$ time (to
find the appropriate distance vector). To print out the path one also needs to follow the predecessor links in the appropriate shortest path tree. This can be done in time proportional to the number of edges of the shortest path.

Another naïve solution requires linear space and also returns the exact solution. This solution does not preprocess the graph, and upon query ( $F, s, t$ ) run Dijkstra's algorithm from $s$ in the graph $G \backslash F$. This is space efficient, but query time is high.

In section 2 we present a simple exact solution with query time $O\left(f^{2}\right)$ and space $O\left(n^{2} \cdot n^{f}\right)$. The primary difficulty that we address is how to replace the $n^{f}$ factor in this space bound, by roughly $\log ^{f} n \cdot 1 / \epsilon^{f}$. We achieve this at the cost of returning an approximate shortest path rather than an exact shortest path (with stretch $1+\epsilon$ ), and slightly increasing the query time. The basic idea that we use to reduce the size of the data structure is to consider sets of edges (called segments) rather than singleton edges. More specifically, instead of considering the deletion of each edge individually, we consider the deletion of the entire segment containing the edge. When done appropriately, the deletion of segments rather than individual edges, maintains approximate shortest paths.

Related Work: The problem of $f$-sensitive distance oracles is closely related to dynamic shortest paths algorithms. In dynamic algorithm as opposed to the $f$ sensitive model, the failures/updates are not given as a batch but rather as an adversarial sequence of update and query operations. The problem of dynamic shortest path was extensively studied since the early 80 's (e.g. $[13,20,21,22,3,9,24,25,23,5,4,16,15,18,17,1])$ and for many problems we have dynamic algorithms with bounds close to optimal. The situation with the $f$-sensitive model is perhaps surprisingly very different and many key problems in this model are far from being fully understood.

The first distance sensitive oracle was in the context of directed graphs [10]. It maintained exact distances and was capable of handling a single edge failure. The space requirement of this oracle is $O\left(n^{2} \log n\right)$ and its query time is $O(\log n)$.

This was later generalized to handle a single vertex or edge failure in [11]. Demetrescu et al. [11] presented an exact 1 -sensitive distance oracle of size $O\left(n^{2} \log n\right)$, $O(1)$ query time and $\widetilde{O}\left(m n^{2}\right)$ preprocessing time.

Later, in two consecutive papers, Karger and Bernstein improved the preprocessing time (while keeping the space and query time unchanged), first in to $O\left(n^{2} \sqrt{m}\right)$ in [6] and then to $\widetilde{O}(m n)$ in [7].

Afterwards, by a quite involved construction, Duan and Pettie [12] considered the case of two failures (ver-
tices or edges) with exact distances. The size of their oracle is $O\left(n^{2} \log ^{3} n\right)$, the query time is $O(\log n)$ and the construction time is polynomial. They claim: "It may yet be possible to find distance oracles capable of handling any fixed number of failures. However, the sheer complexity of our algorithm suggests that moving beyond dual-failures will require a fundamentally different approach to the problem". Further details on their dual-failure data-structure appear in Duan's PhD Thesis [27].

Khanna and Baswana [19] considered approximate 1-sensitive distance oracles for unweighted graphs. More precisely, they presented a data structure of size $O\left(k^{5} n^{1+1 / k} \frac{\log ^{3} n}{\epsilon^{4}}\right),(2 k-1)(1+\epsilon)$ stretch and $O(k)$ query time.

The problem of 1 -sensitive distance oracles was also studied with a special focus on improving the preprocessing time at the cost of a polynomial query time (see $[26,14]$ ).

The case of handling more than two edge failures (while keeping query time small) was also considered in [8] - albeit the stretch was $\Omega(f)$ which is quite large. Let $W$ denote the weight of the heaviest edge in the graph assuming the smallest edge weight is 1 . For any integer parameter $k$, [8] gave an $f$-sensitive distance oracle of size $O\left(f k n^{1+1 / k} \log (n W)\right),(8 k-2)(f+1)$ stretch and $O\left(f \cdot \log ^{2} n \cdot \log \log ^{2} n\right)$ query time. Note that even with $k=1$ and for a constant number of failures, the stretch is quite large, and therefore irrelevant in many settings.
Our Results: To the best of our knowledge, we present the first non trivial $f$-sensitive distance oracle with small stretch and fast query time capable of handling multiple edge failures for weighted undirected graphs. We are unaware of any previous construction that deals with more than 2 failures and gives stretch arbitrarily close to one (See Tables 1, 2). We deal with an arbitrary (predetermined) upper bound on the number of failures, and arbitrarily small (constant) stretch. More precisely, our main result is the following theorem.

Theorem 1.1. Let $G=(V, E)$ be a weighted undirected graph, with $|V|=n$, and $|E|=m$, where all edges have positive weight in the range $[1, W]$. Let $D=n \cdot W$ (which is an upper bound on the diameter of the graph), and let $\epsilon$ be a parameter such that $0<\epsilon<1$. One can construct in $O\left(\left(f n^{5}+n^{4} \log D+n^{3} \log D \log \log D\right)(\log D / \epsilon)^{f}\right)$ time an f-sensitive distance oracle for $G$ with $(1+\epsilon)$ stretch, $O\left(n^{2}(\log D / \epsilon)^{f} \cdot f \log D\right)$ space, and $O\left(f^{5} \log D\right)$ query time.

In addition, we show that one can replace the $\log ^{f} D$

| Model | $f-\operatorname{Max} \#$ <br> failures | Approximation | Data Structure Size | Query time | Preprocessing Time | Ref. |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Directed Edge Failure | 1 | 1 | $O\left(n^{2} \log n\right)$ | $O(\log n)$ | $O\left(m n^{2} \log n+n^{3} \log ^{2} n\right)$ | [10] |
| Directed Edge/vertex Failure | 1 | 1 | $O\left(n^{2} \log n\right)$ | $O(1)$ | $O\left(m n^{2}\right)$ | [11] |
| Directed Edge/vertex Failure | 1 | 1 | $O\left(n^{2} \log n\right)$ | $O(1)$ | $\widetilde{O}(m n)$ | [6], [7] |
| Directed Edge/vertex Failure | 2 | 1 | $O\left(n^{2} \log ^{3} n\right)$ | $O(\log n)$ | poly $(n)$ | [12] |
| Undirected Unweighted Edge/vertex Failure | 1 | $\begin{aligned} & (2 k-1)(1+\epsilon) \\ & k \geq 1, \epsilon>0 \end{aligned}$ | $O\left(k^{5} n^{1+1 / k} \log ^{3} n / \epsilon^{4}\right)$ | $O(k)$ | $O\left(k m n^{1+1 / k}\right)$ | [19] |
| Undirected Weighted Edge Failure | $f \in Z$ | $\begin{aligned} & (8 k-2)(f+1) \\ & k \geq 1 \end{aligned}$ | $O\left(f k n^{1+1 / k} \log (n W)\right)$ | $O\left(f \cdot \log ^{2} n \cdot \log \log ^{2} n\right)$ | poly $(n)$ | [8] |

Table 1: Previous $f$-sensitive Distance Oracles. Note that if the number of failures $f>2$ then the stretch is $(8 k-2)(f+1) \geq 24$.

| Model | $f-\operatorname{Max} \#$ failures | Approximation | Data Structure Size | Query time | Preprocessing <br> Time | Ref. |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Undirected Unweighted Edge Failure | $f \in Z$ | $\begin{aligned} & 1+\epsilon \\ & \epsilon>0 \end{aligned}$ | $O\left(n^{3}(\log n / \epsilon)^{f}\right)$ | $O\left(f^{4}\right)$ | $O\left(n^{2}(\log n / \epsilon)^{f}(m+n \log n)\right)$ | This Paper |
| Undirected Unweighted Edge Failure | $f \in Z$ | $\begin{gathered} 1+\epsilon \\ \epsilon>0 \end{gathered}$ | $O\left(n^{2}(\log n / \epsilon)^{f} \cdot f \log n\right)$ | $O\left(f^{5} \log n\right)$ | $O\left(\left(f n^{5}\right)(\log n / \epsilon)^{f}\right)$ | This Paper |
| Undirected Weighted Edge Failure | $f \in Z$ | $\begin{aligned} & 1+\epsilon \\ & \epsilon>0 \end{aligned}$ | $O\left(n^{3}(\log D / \epsilon)^{f}\right)$ | $O\left(f^{4}\right)$ | $O\left(n^{2}(\log D / \epsilon)^{f}(m+n \log n)\right)$ | This Paper |
| Undirected Weighted Edge Failure | $f \in Z$ | $\begin{aligned} & 1+\epsilon \\ & \epsilon>0 \end{aligned}$ | $O\left(n^{2}(\log D / \epsilon)^{f} \cdot f \log D\right)$ | $O\left(f^{5} \log D\right)$ | $O\left(\left(f n^{5}+n^{4} \log D+n^{3} \log D \log \log D\right)(\log D / \epsilon)^{f}\right)$ | This Paper |
| Undirected Weighted Edge Failure | $f \in Z$ | $\begin{aligned} & 1+\epsilon \\ & \epsilon>0 \end{aligned}$ | $O\left(n^{3}(\log n / \epsilon)^{f}(\log W / \log n)\right)$ | $O\left(f^{4} \log \log W\right)$ | $O\left(n^{2}(\log n / \epsilon)^{f}(\log W / \log n)(m+n \log n)\right)$ | This Paper |
| Undirected Weighted Edge Failure | $f \in Z$ | $\begin{aligned} & 1+\epsilon \\ & \epsilon>0 \end{aligned}$ | $O\left(n^{2}(\log n / \epsilon)^{f} \cdot f \log W\right)$ | $O\left(f^{5} \log n \log \log W\right)$ | $O\left(f n^{5}(\log n / \epsilon)^{f}(\log W / \log n)\right)$ | This Paper |

Table 2: New $f$-sensitive Distance Oracles. Note that the stretch is now $1+\epsilon$ for arbitrary small $\epsilon>0$. The first two rows describe our results for unweighted graphs, rows 3 and 4 give our results for weighted graphs as a function of $D=n \cdot W$, and the last two rows improve the space bound as a function of the max edge weight $W$.
factor in Theorem 1.1 by (roughly) $\log ^{f} n$, at the cost of an additional $\log \log W$ factor in the query time.

The rest of the paper is organized as follows. In Section 2 we present an $f$-sensitive $(1+\epsilon)$ stretch distance oracle that requires $O\left(n^{3}(\log D / \epsilon)^{f}\right)$ space, and $O\left(f^{4}\right)$ query time, where $D=n \cdot W$ is an upper bound on the diameter of the graph (the edge weights are normalized so that the minimum edge has weight 1 and $W$ is the maximum edge weight). In Section 3 we reduce the space requirements of our data structure to $O\left(n^{2}(\log D / \epsilon)^{f} \cdot f \log D\right)$ and obtain Theorem 1.1. The main idea used in Section 3 is that we do not need to explicitly store the paths in our data-structures. Rather we can represent the paths implicitly as the concatenation of at most $O(f \log D)$ shortest paths in
the original graph $G$. We do this by developing new tools introduced in Section 3 (decomposable paths and expaths).

We further improve the space in Section 4 and replace the $\log ^{f} D$ factor in Theorem 1.1 by (roughly) $\log ^{f} n$ at the cost of an additional $\log \log W$ factor in the query time (see Table 2 for the full details).
1.1 Preliminaries. We use the following notations:

Let $P$ be a path. $|P|$ is defined to be the sum of the weights of its edges.
$\mathrm{d}_{G}(u, v)$ is the distance from $u$ to $v$ in $G$, abbreviated $\mathrm{d}(u, v)$ when $G$ is known from the context. $P_{G}(u, v)$ denotes an arbitrary shortest path from $u$ to $v$ in $G$ and is abbreviated $P(u, v)$ when $G$ is known from the con-
text.
Let $P$ be a path and $x, y$ be two vertices along the path. Then $P[x, y]$ is the subpath of $P$ from $x$ to $y$ (including $x$ and $y$ ).

We now define the path concatenation $(\cdot)$ operator. Let $P_{1}=\left(x_{1}, x_{2}, \ldots, x_{r}\right)$ and $P_{2}=\left(y_{1}, y_{2}, \ldots, y_{t}\right)$ be two paths. Then $P=P_{1} \cdot P_{2}$ is defined as the path $P=\left(x_{1}, x_{2}, \ldots, x_{r}, y_{1}, y_{2}, \ldots, y_{t}\right)$, and it is well defined if either $x_{r}=y_{1}$ or $\left(x_{r}, y_{1}\right) \in E$.

## $2 \widetilde{O}\left(n^{3}\right)$ Approximate Distance Sensitive Oracles.

In this section we present an $f$-sensitive distance oracle for a given graph $G$ with $(1+\epsilon)$ stretch, $O\left(n^{3}(\log D / \epsilon)^{f}\right)$ space, and $O\left(f^{4}\right)$ query time.

The intuition for our construction comes from a simple data structure that can accommodate up to $f$ edge failures and returns the exact shortest path in the graph without the failed edges. Unfortunately, the space requirements are quite large, $O\left(n^{f+3}\right)$. We describe this simple data structure in Appendix A as the paper is selfcontained without it, and we give it to provide some more intuition.
2.1 Improved Data Structure. Similarly to the $O\left(n^{f+3}\right)$ data structure, our $O\left(n^{3}(\log D / \epsilon)^{f}\right)$ data structure consists of rooted trees $F T(u, v)$ for every pair of vertices $u, v \in V$. We further improve upon the space requirement in Section 3.

To define $F T(u, v)$ we first define a general decomposition of a path $P$ into subpaths using a particular subset of the vertices of $P$ which we call netpoints. We use $\epsilon^{\prime}=\epsilon / 3$ in these definitions.

Definition 2.1. (Path Netpoints) Let $P=\left(v_{1}=\right.$ $u, v_{2}, \ldots, v_{k}=v$ ) be a path from $u$ to $v$ in $G$. Let $L$ be the set of all vertices $v_{j}, v_{j+1} \in P$ such that $\left|P\left[u, v_{j}\right]\right|<\left(1+\epsilon^{\prime}\right)^{i} \leq\left|P\left[u, v_{j+1}\right]\right|$ for some integer $i \geq 0$. Define $R$ analogously to be the set of all vertices $v_{j}, v_{j-1} \in P$ such that $\left|P\left[v_{j}, v\right]\right|<\left(1+\epsilon^{\prime}\right)^{i} \leq\left|P\left[v_{j-1}, v\right]\right|$ for some integer $i$. Let netpoints $(P)=L \cup R \cup\{u, v\}$. See Figure 1.

Note that the cardinality of netpoints $(P)$ is $O\left(\log _{1+\epsilon^{\prime}}|P|\right)=O\left(\log _{1+\epsilon}|P|\right)$ and for a small enough $\epsilon, O\left(\log _{1+\epsilon}|P|\right)=O\left(\frac{\log |P|}{\epsilon}\right)=O(\log D / \epsilon)$.

A segment of $P$ is a subpath of $P$ connecting two consecutive netpoints along $P$. We denote by $\operatorname{seg}(P)$ the set of segments of the path $P$. From this definition it follows that the segments in $\operatorname{seg}(P)$ are pairwise disjoint except possibly for their endpoints. For $e \in P$ we define $\operatorname{seg}(e, P)$ to be the segment of $P$ containing $e$, see Figure 2.

The following lemma follows immediately from the definition of netpoints and segments.

Lemma 2.1. Let $P$ be a path from $u$ to $v, e=(x, y)$ be an edge of $P$. Either $\operatorname{seg}(e, P)=\{e\}$ or $|\operatorname{seg}(e, P)| \leq$ $\epsilon^{\prime} \min \{|P[u, x]|,|P[x, v]|\}$.

Proof. Assume with out loss of generality that the vertex $x$ is closer than $y$ to $u$ along $P$ (i.e. $|P[u, x]|<$ $|P[u, y]|)$. Let $i$ be the maximal index such that $\left(1+\epsilon^{\prime}\right)^{i} \leq|P[u, x]|$. Then $\left(1+\epsilon^{\prime}\right)^{i+1}>|P[u, x]|$. If $\left(1+\epsilon^{\prime}\right)^{i+1} \leq|P[u, y]|$ then we get $|P[u, x]|<(1+$ $\left.\epsilon^{\prime}\right)^{i+1} \leq|P[u, y]|$ and by definition both $x$ and $y$ are netpoints of $P$, hence $\operatorname{seg}(e, P)=\{e\}$. Otherwise, $\left(1+\epsilon^{\prime}\right)^{i+1}>|P[u, y]|$ and
$|\operatorname{seg}(e, P)| \leq\left(1+\epsilon^{\prime}\right)^{i+1}-\left(1+\epsilon^{\prime}\right)^{i}=\epsilon^{\prime}\left(1+\epsilon^{\prime}\right)^{i} \leq \epsilon^{\prime}|P[u, x]|$.
Symmetrically, we can show that either $\operatorname{seg}(e, P)=$ $\{e\}$ or $|\operatorname{seg}(e, P)| \leq \epsilon^{\prime}|P[y, v]|$. If $|\operatorname{seg}(e, P)| \leq$ $\epsilon^{\prime}|P[y, v]|$, it follows that $|\operatorname{seg}(e, P)| \leq \epsilon^{\prime}|P[y, v]|<$ $\epsilon^{\prime}|P[x, v]|$ (since $|P[y, v]|<|P[x, v]|$ by assumption). Therefore, either $\operatorname{seg}(e, P)=\{e\}$ or $|\operatorname{seg}(e, P)| \leq$ $\epsilon^{\prime} \min \{|P[u, x]|,|P[x, v]|\}$.

The trees $F T(u, v)$. We associated with each node $\alpha \in F T(u, v)$ a particular graph $G_{\alpha}$, the shortest path $P_{\alpha}:=P_{G_{\alpha}}(u, v)$ which is the shortest path from $u$ to $v$ in the graph $G_{\alpha}$, and a segment $S_{\alpha} \in \operatorname{seg}\left(P_{p(\alpha)}\right)$ where $p(\alpha)$ is the parent of $\alpha$ in $F T(u, v)$. The graphs $G_{\alpha}$ are defined top down as follows.

The graph $G_{r}$ associated with the root $r$ of $F T(u, v)$ is the input graph $G=(V, E)$ itself (we formally define $S_{r}=\emptyset$ ). We introduce a child $\alpha$ of $r$ for each segment $s \in \operatorname{seg}\left(P_{r}\right)$, define $S_{\alpha}=s$, and $G_{\alpha}$ to be the graph obtained from $G_{r}$ by deleting the edges of $s$.

In general, assume that we have defined $F T(u, v)$ up to level $k$. We define level $k+1$ as follows. For each node $\alpha^{\prime}$ of level $k$ for which $P_{\alpha^{\prime}} \neq \emptyset$ we introduce a child $\alpha$, per segment $s \in \operatorname{seg}\left(P_{\alpha^{\prime}}\right)$, define $S_{\alpha}=s$, and $G_{\alpha}$ to be the graph obtained from $G_{\alpha^{\prime}}$ by deleting the edges of $s$. Nodes $\alpha^{\prime}$ of level $k$ for which $P_{\alpha^{\prime}}=\emptyset$ are leaves of $F T(u, v)$.

For each node $\alpha \in F T(u, v)$ we define $\operatorname{avoid}(\alpha)$ to be the set containing all edges in the segment associated with $\alpha$ and in the segments associated with the ancestors of $\alpha$ in $F T(u, v)$ (equivalently one can define recursively avoid $\left.(\alpha)=\operatorname{avoid}(p(\alpha)) \cup S_{\alpha}\right)$. Then $G_{\alpha}=G \backslash \operatorname{avoid}(\alpha)$. See Figure 3.

Our data structure consists of the tree $F T(u, v)$ for every pair of vertices $u, v$, constructed up to level $f$. In each node $\alpha \in F T(u, v)$ we keep only the path $P_{\alpha}$. We store the edges of $P_{\alpha}$ in a perfect hash table, keeping the index of $\operatorname{seg}\left(e, P_{\alpha}\right)$ with each edge $e$. This index identifies the child $\beta$ of $\alpha$ such that $S_{\beta}=\operatorname{seg}\left(e, P_{\alpha}\right)$.

## Path netpoints and segments.



Figure 1: Path netpoints, the union of both sets of netpoints, from $u$ ( $L$ in text) and from $v$ ( $R$ in text). Although not shown, the same vertex can be chosen both for $L$ and for $R$, and, moreover, the same vertex can be chosen multiple times for $L$ and/or for $R$.


Figure 2: Segments of a path and $\operatorname{seg}(e, P)$.


Figure 3: A tree $F T(u, v)$. The path at the root $r$ is $P_{r}=P_{G}(u, v)$, the shortest path in the graph $G$ between vertices $u$ and $v$. Every node in $F T(u, v)$ has $O\left(\log D / \epsilon^{\prime}\right)$ children. The path $P_{r}$ is split by the vertices of netpoints $\left(P_{r}\right)$ into segments $S_{\alpha}, S_{\beta}$, and $S_{\gamma}$. Associated with each such segment is a child of node $r$, nodes $\alpha, \beta$, and $\gamma$. The path $P_{\alpha}$ is the shortest path in the graph $G_{\alpha}$, where $G_{\alpha}$ is the result of deleting all edges along $S_{\alpha}$ from the graph $G=G_{r}$ (i.e., deleting avoid $(\alpha)$ ). Paths $P_{\beta}$ and $P_{\gamma}$ are defined analogously. The path $P_{\beta}$ is split by netpoints $\left(P_{\beta}\right)$ into segments $P_{\delta}, P_{\eta}$. Segments $P_{\delta}$ and $P_{\eta}$ have associated nodes $\delta$ and $\eta$ - children of node $\beta$. The path $P_{\delta}$ is the shortest path between vertices $u$ and $v$ in the graph $G_{\delta}$, where $G_{\delta}$ is the result of deleting the edges along $S_{\beta}$ and $S_{\delta}$ from $G$. Equivalently, $G_{\delta}$ is the result of deleting the edges along $S_{\delta}$ from $G_{\beta}$. The path $P_{\eta}$ is defined similarly.

Since the number of children of each node of $F T(u, v)$ is $O(\log D / \epsilon)$ and the depth of $F T(u, v)$ is $f$ then $F T(u, v)$ contains $O\left((\log D / \epsilon)^{f}\right)$ nodes. Each node $\alpha$ keeps the path $P_{\alpha}$ in a data structure of size linear in number of edges of $P_{\alpha}$ which is $O(n)$. It follows that the total size of $F T(u, v)$ is $O\left(n(\log D / \epsilon)^{f}\right)$. The total size of the trees $F T(u, v)$ for all pairs $u, v \in V$ is $O\left(n^{3}(\log D / \epsilon)^{f}\right)$. This is the size of the data structure.

Preprocessing. We construct $F T(u, v)$ top down while maintaining $G_{\alpha}$. At each node $\alpha$ we have to run a single source shortest path computation to find $P_{\alpha}$. Subsequently we compute netpoints $\left(P_{\alpha}\right)$ and $\operatorname{seg}\left(P_{\alpha}\right)$ in $O(n)$ time and construct the hash table storing the edges of $P_{\alpha}$ in $O(n)$ expected time. Summing up over all nodes $\alpha$ in all trees $F T(u, v)$ we get that the total preprocessing time is $O\left(n^{2}(\log D / \epsilon)^{f}(m+n \log n)\right)$ expected time.
2.2 Query. Given $F \subset E,|F| \leq f$ (a set of failing edges), and a pair of vertices $s, t \in V$, we approximate
the length of the shortest path from $s$ to $t$ in $G \backslash F$ as follows.
The graph $H$ : We build a small weighted graph $H$ over $V(H)=V(F) \cup\{s, t\}$, i.e. the vertices of $H$ are the endpoints of the edges of $F$, and $s, t^{1}$ For each pair of vertices $u, v \in V(H)$ we decide if $E(H)$ contains $(u, v)$ and assign it a weight $w_{H}(u, v)$ as follows.

We traverse the tree $F T(u, v)$ from its root towards a leaf. When visiting a node $\alpha \in F T(u, v)$, we do one of the following steps.

Case 1: If $P_{\alpha}=\emptyset$ we stop the traversal and don't add an edge between $u, v$ in $H$.

Case 2: If $P_{\alpha} \cap F=\emptyset$ we add an edge $(u, v)$ to $H$ with $w_{H}(u, v)=\left|P_{\alpha}\right|$ and stop the traversal of $F T(u, v)$.

[^0]Case 3: Otherwise, $P_{\alpha}$ must contain a failing edge $e \in$ $F$. We choose arbitrarily an edge $e \in P_{\alpha} \cap F$ and continue the traversal in a subtree of $\alpha$ obtained by following the child $\alpha^{\prime}$ of $\alpha$ such that $S_{\alpha^{\prime}}=$ $\operatorname{seg}\left(e, P_{\alpha}\right)$.

We prove in Lemma 2.3 below that this traversal is well defined in the sense that we never reach Case 3 when $\alpha$ is at level $f$. We add an edge to $H$ only if this procedure ends in Case 2. So each edge $(u, v)$ in $H$ corresponds to some path in $G \backslash F$ and $w_{H}(u, v)$ is the length of this path.

Recall that $\operatorname{avoid}(\alpha)$ is the set of edges along the segments $S_{\alpha^{\prime}}$ for all nodes $\alpha^{\prime}$ that are ancestors of $\alpha$ in $F T(u, v)$ (including edges of $S_{\alpha}$ ), $G_{\alpha}$ is the graph $G \backslash \operatorname{avoid}(\alpha)$ and $P_{\alpha}$ is the shortest path between $u$ to $v$ which avoids the edges avoid $(\alpha)$.

Lemma 2.2. During the traversal of $F T(u, v)$ let $\alpha_{d}$ be the node visited at depth $d$, it must be that $\mid \operatorname{avoid}\left(\alpha_{d}\right) \cap$ $F \mid \geq d$.

Proof. By induction on $d$. For $d=0$ this holds vacuously. For $d>1$, we traversed the edge $\left(\alpha_{d-1}, \alpha_{d}\right)$ labeled by $\operatorname{seg}\left(e, P_{\alpha_{d-1}}\right)$ where $e \in F \cap P_{\alpha_{d-1}}$. By definition,

$$
\operatorname{avoid}\left(\alpha_{d}\right)=\operatorname{avoid}\left(\alpha_{d-1}\right) \cup \operatorname{seg}\left(e, P_{\alpha_{d-1}}\right)
$$

Since $e \in F \cap\left(\operatorname{avoid}\left(\alpha_{d}\right) \backslash \operatorname{avoid}\left(\alpha_{d-1}\right)\right)$, we get that

$$
\left|\operatorname{avoid}\left(\alpha_{d}\right) \cap F\right| \geq\left|\operatorname{avoid}\left(\alpha_{d-1}\right) \cap F\right|+1
$$

and the lemma follows by induction.
Lemma 2.3. The query procedure is well-defined in the sense that it never applies Case 3 in a leaf of $\operatorname{FT}(u, v)$.

Proof. Let $\alpha \in F T(u, v)$ be a node that is traversed by the query in which we apply Case 3 . In Case 3 we have $P_{\alpha} \cap F \neq \emptyset$, let $e \in P_{\alpha} \cap F$. Since $e \in P_{\alpha}$ it follows that $e \notin \operatorname{avoid}(\alpha)$, thus $|\operatorname{avoid}(\alpha) \cap F|<|F|=f$. By Lemma 2.2, it must be that the depth of $\alpha$ is less than $f$, and therefore $\alpha$ is not a leaf of $F T(u, v)$.

We use $H$ to answer the query as follows. If $s$ and $t$ are disconnected in $H$, then we answer that $s$ and $t$ are also disconnected in $G \backslash F$. Otherwise, we compute and return $\mathrm{d}_{H}(s, t)$ as our approximation of $\mathrm{d}_{G \backslash F}(s, t)$.
Lemma 2.4. Query time is $O\left(f^{4}\right)$.
Proof. This is the time it takes to build the graph $H$. For each pair $u, v \in V(H)$ (there are $O\left(f^{2}\right)$ such pairs) we traverse the tree $F T(u, v)$ along a root-to-leaf path (the length of the root-to-leaf path is at most $f$ ) and at
each node along the root-to-leaf path we find a segment containing a failing edge. We show below how to find a segment containing a failing edge in each node which we traverse in $O(f)$ time. Therefore, the total time to build the graph $H$ is $O\left(f^{4}\right)$ (that is, $O\left(f^{2}\right)$ trees $F T(u, v)$, multiplied by $O(f)$ nodes on a root-to-leaf path in each tree, multiplied by $O(f)$ time per node to find the appropriate segment). Finding the shortest path from $s$ to $t$ in $H$ takes $\widetilde{O}\left(f^{2}\right)$ time, so the query time is dominated by the time required to build $H$.

To determine if $e \in P$ and to find $\operatorname{seg}(e, P)$ if $e \in P$, we store the edges of $P$ in a perfect hash table, keeping also the index of $\operatorname{seg}(e, P)$ with each edge. This allows to determine if an edge $e$ is in $P$ in $O(1)$ time. To determine if there is an edge $e \in F$ that belongs to $P$ we query the hash table with the edges of $F$ one by one in $O(f)$ total time.
2.3 Correctness. We prove that $\mathrm{d}_{H}(s, t)$ is a good approximation of $\mathrm{d}_{G \backslash F}(s, t)$.

The very high level idea of our analysis is that we distinguish between two cases. The first case is when all failures are "far away" from the shortest path from $s$ to $t$ in $G \backslash F$. In this case, we show that we can find an exact shortest path from $s$ to $t$ by querying the tree $F T(u, v)$ (see Lemma 2.5). Loosely speaking, since all failures are far away from the shortest path, discarding at each level of the search in $F T(u, v)$ a segment containing a failed edge doesn't eliminate any of the edges of $P$. Hence the shortest path survives during the query of $F T(u, v)$. The second case is when there is a "close by" failure to the shortest path from $s$ to $t$ in $G \backslash F$. In this case, we show that we can construct an approximate shortest path that consists only of sub-paths where such subpath is a path between vertices in $V(F) \cup\{s, t\}$ and each sub-path is "far away" from all failures and thus can be found by querying the appropriate tree $F T(u, v)$.

To prove the correctness we use the following definition of the trapezoid of a path.

Definition 2.2. $\left(\operatorname{tr}_{G}(P)\right)$ Let $P$ be an arbitrary path from $u$ to $v$ in $G$. Let $x \in V$ be a vertex and $\ell \geq 0 a$ real number, define

$$
\operatorname{ball}_{G}(x, \ell)=\left\{u \mid \mathrm{d}_{G}(x, u) \leq \ell\right\}
$$

$\operatorname{ball}_{G}(x, \ell)$ is the set of all vertices whose distance in $G$ from $x$ is at most $\ell$. Define the trapezoid of $P$ in $G$ to be $\operatorname{tr}_{G}(P)=\left(\bigcup_{x \in P \backslash\{u, v\}} \operatorname{ball}_{G}\left(x, \epsilon^{\prime} \cdot \min (|P[u, x]|,|P[x, v]|)\right)\right) \backslash\{u, v\}$.

The next lemma gives a sufficient condition for having an edge $(u, v)$ in $H$.

Lemma 2.5. Let $u, v \in V(H)$, and let $P$ be an arbitrary
path from $u$ to $v$ in $G \backslash F$. If

$$
\operatorname{tr}_{G \backslash F}(P) \cap V(H)=\emptyset
$$

Proof. Let $\alpha_{0}(=r), \alpha_{1}, \ldots, \alpha_{k}, k \leq f$, be the nodes of $F T(u, v)$ that we traversed during a query. We prove below that $P$ is contained in $G_{\alpha_{i}}$ for every $1 \leq i \leq k$. From this follows that the query concluded the traversal of $F T(u, v)$ in Case 2 (recall the cases of the procedure that decides if to add $(u, v)$ to $H$ in Section 2.2), added $(u, v)$ to $H$, and set $w_{H}(u, v)=\left|P_{\alpha_{k}}\right|=\mathrm{d}_{G_{\alpha_{k}}}(u, v)$. Since $P$ is a path between $u$ and $v$ in $G_{\alpha_{k}}$, and $P_{\alpha_{k}}$ is the shortest path between $u$ and $v$ in $G_{\alpha_{k}}$, we obtain that $w_{H}(u, v) \leq|P|$.

The crux of the proof is to show that $P$ is contained in $G_{\alpha_{i}}$ for $i=0, \ldots, k$. We prove this by induction on $i$. The base case is obvious since $G_{r}=G$. To simplify the notation which we use to establish the induction step, let $\alpha=\alpha_{i}$ and $\alpha^{\prime}=\alpha_{i+1}$. We assume that $P$ is contained in $G_{\alpha}$ and we need to prove that $P$ is contained in $G_{\alpha^{\prime}}$. The only way $P$ can disappear from $G_{\alpha^{\prime}}$ is if there is a segment of $P_{\alpha}$ that contains an edge $e_{F} \in F$ and an edge $e_{P} \in P$. (Note that $e_{F} \neq e_{P}$ since $P \in G \backslash F$.)

Assume by contradiction that such a segment $s \in$ $\operatorname{seg}\left(P_{\alpha}\right)$ exists and assume without loss of generality that there isn't any other edge of $F$ or $P$ between $e_{F}$ and $e_{P}$ in $s$, and thus $P_{\alpha}[x, y]$ is a path in $G \backslash F$. Let $x$ be the endpoint of $e_{P}$ which is closer to $e_{F}$ along $P_{\alpha}$, and let $y$ be the endpoint of $e_{F}$ which is closer to $e_{P}$ along $P_{\alpha}$. Note that $x, y \notin\{u, v\}$. See Figure 4.

The vertex $x$ appears along $P$ and partitions $P$ into two subpaths $P[u, x]$ and $P[x, v]$. Similarly, the vertex $x$ also appears along $P_{\alpha}$ and splits it into two subpaths $P_{\alpha}[u, x]$ and $P_{\alpha}[x, v]$. Since $P_{\alpha}$ is the shortest path from $u$ to $v$ in $G_{\alpha}$ and $P$ is a path in $G_{\alpha}$ it follows that $\left|P_{\alpha}[u, x]\right| \leq|P[u, x]|$ and $\left|P_{\alpha}[x, v]\right| \leq|P[x, v]|$ which implies that

$$
\begin{equation*}
\min \left\{\left|P_{\alpha}[u, x]\right|,\left|P_{\alpha}[x, v]\right|\right\} \leq \min \{|P[u, x]|,|P[x, v]|\} \tag{2.1}
\end{equation*}
$$

Since $x \in V(P), y \in V(H)$, and $\operatorname{tr}_{G \backslash F}(P) \cap V(H)=$ $\emptyset$ by the assumption of the lemma, it follows that $y \notin \operatorname{tr}_{G \backslash F}(P)$ and so (see Figure 5)

$$
\begin{equation*}
\mathrm{d}_{G \backslash F}(x, y)>\epsilon^{\prime} \min \{|P[u, x]|,|P[x, v]|\} . \tag{2.2}
\end{equation*}
$$

Combining Equations (2.1) and (2.2) and since $\left|P_{\alpha}[x, y]\right| \geq \mathrm{d}_{G \backslash F}(x, y)$ (since we assumed $P_{\alpha}[x, y]$ is a path in $G \backslash F)$ we get that

$$
\begin{aligned}
& \quad\left|\operatorname{seg}\left(e_{F}, P_{\alpha}\right)\right| \geq\left|P_{\alpha}[x, y]\right| \geq \mathrm{d}_{G \backslash F}(x, y)> \\
& \epsilon^{\prime} \min \{|P[u, x]|,|P[x, v]|\} \geq \epsilon^{\prime} \min \left\{\left|P_{\alpha}[u, x]\right|,\left|P_{\alpha}[x, v]\right|\right\}
\end{aligned}
$$

Since $\operatorname{seg}\left(e_{F}, P_{\alpha}\right)$ contains both $e_{F}$, and $e_{P}$ and $e_{F} \neq e_{P}$, it cannot be that $\operatorname{seg}\left(e_{F}, P_{\alpha}\right)=\left\{e_{F}\right\}$. Therefore we obtain a contradiction to Lemma 2.1, since neither $\operatorname{seg}\left(e_{F}, P_{\alpha}\right)=\left\{e_{F}\right\}$ nor $\left|\operatorname{seg}\left(e_{F}, P_{\alpha}\right)\right| \leq$ $\epsilon^{\prime} \min \left\{\left|P_{\alpha}[u, x]\right|,\left|P_{\alpha}[x, v]\right|\right\}$.

When the condition of Lemma 2.5 does not hold (i.e., $\operatorname{tr}_{G \backslash F}(P) \cap V(H) \neq \emptyset$ ) then the following lemma shows how to find a vertex $z$ in $\operatorname{tr}_{G \backslash F}(P) \cap V(H)$ such that one of the distances from $u$ to $z$ or from $z$ to $v$ can be approximated using Lemma 2.5, and the other path is shorter than $P$, and hence can be approximated recursively. See Figure 7 where $x$ is $u$, there is an edge from $u$ to $z$ in $H$ according to Lemma 2.5, and the approximate length of the dashed path from $z$ to $v$ is found recursively. The proof of Theorem 2.1 below specifies this process, Lemmata 2.5 above and 2.6 below are used in the proof of Theorem 2.1.

Lemma 2.6. Let $u, v \in V(H)$, and let $P=P_{G \backslash F}(u, v)$. If $\operatorname{tr}_{G \backslash F}(P) \cap V(H) \neq \emptyset$ then there exist vertices $x, y$, and $z$ such that (see Figure 6):

$$
\begin{aligned}
& \text { 1. } x \in\{u, v\}, y \in P, z \in \operatorname{tr}_{G \backslash F}(P) \cap V(H) \text {. } \\
& \text { 2. }|P[x, y]| \leq \frac{1}{2}|P[u, v]| \text {. } \\
& \text { 3. } \mathrm{d}_{G \backslash F}(y, z) \leq \epsilon^{\prime} \mathrm{d}_{G \backslash F}(x, y) \text {. } \\
& \text { 4. Let } P^{\prime}=P_{G \backslash F}(x, y) \cdot P_{G \backslash F}(y, z) \text {. Then } \operatorname{tr}_{G \backslash F}\left(P^{\prime}\right) \cap \\
& V(H)=\emptyset \text {. }
\end{aligned}
$$

Proof. Since $\operatorname{tr}_{G \backslash F}(P) \cap V(H) \neq \emptyset$, there exists $z \in \operatorname{tr}_{G \backslash F}(P) \cap V(H)$. By the definition of $\operatorname{tr}_{G \backslash F}(P)$ there exists $y \in P$ such that $z \in$ $\operatorname{ball}_{G \backslash F}\left(y, \epsilon^{\prime} \min \{|P[u, y]|,|P[y, v]|)\right\}$. We then choose $x \in\{u, v\}$ such that $|P[x, y]|=\min \{|P[u, y]|,|P[y, v]|\}$ and obtain that $x, y, z$ satisfy Conditions 1,2 and 3 .

Now among all triples of vertices $x, y, z$ that satisfy Conditions 1,2 , and 3 we choose the triple $x, y, z$ that minimizes $\mathrm{d}_{G \backslash F}(x, y)$, breaking ties in favor of the triple that minimizes $\mathrm{d}_{G \backslash F}(y, z)$. We show that these $x, y, z$ satisfy Condition 4.

Assume by contradiction that $\operatorname{tr}_{G \backslash F}\left(P^{\prime}\right) \cap V(H) \neq$ $\emptyset$. Then there are vertices $z^{\prime} \in \operatorname{tr}_{G \backslash F}\left(P^{\prime}\right) \cap$ $V(H)$ and $y^{\prime} \in P^{\prime}$ such that $\mathrm{d}_{G \backslash F}\left(y^{\prime}, z^{\prime}\right) \leq$ $\epsilon^{\prime} \min \left\{\left|P^{\prime}\left[x, y^{\prime}\right]\right|,\left|P^{\prime}\left[y^{\prime}, z\right]\right|\right\} . \quad$ As $P^{\prime}=P_{G \backslash F}(x, y)$. $P_{G \backslash F}(y, z)$ and $y^{\prime} \in P^{\prime}$ then either $y^{\prime} \in P_{G \backslash F}(x, y)$ or $y^{\prime} \in P_{G \backslash F}(y, z)$.

If $y^{\prime} \in P_{G \backslash F}(x, y) \backslash\{y\}$ then $\left|P\left[x, y^{\prime}\right]\right|<|P[x, y]|$ and this contradicts the choice of $x, y, z$. (We should have preferred $x, y^{\prime}, z^{\prime}$, see the upper part of Figure 6)

If $y^{\prime} \in P_{G \backslash F}(y, z)$ (including the case $y^{\prime}=y$ ), then


Figure 4: The path $P_{\alpha}$ and the edges $e_{P}$ and $e_{F}$ in the proof of Lemma 2.5.


Figure 5: Proof of Lemma 2.5: Since $y \notin \operatorname{tr}_{G \backslash F}(P)$ we have $\mathrm{d}_{G \backslash F}(x, y)>\epsilon^{\prime} \min \{|P[u, x]|,|P[x, v]|\}$.
$\mathrm{d}_{G \backslash F}\left(y^{\prime}, z^{\prime}\right) \leq \epsilon^{\prime} \mathrm{d}_{G \backslash F}\left(y^{\prime}, z\right)$ and therefore
$\mathrm{d}_{G \backslash F}\left(y, z^{\prime}\right) \leq$
$\mathrm{d}_{G \backslash F}\left(y, y^{\prime}\right)+\mathrm{d}_{G \backslash F}\left(y^{\prime}, z^{\prime}\right)=$
$\left(\mathrm{d}_{G \backslash F}(y, z)-\mathrm{d}_{G \backslash F}\left(z, y^{\prime}\right)\right)+\mathrm{d}_{G \backslash F}\left(y^{\prime}, z^{\prime}\right) \leq$
$\mathrm{d}_{G \backslash F}(y, z)-\mathrm{d}_{G \backslash F}\left(z, y^{\prime}\right)+\epsilon^{\prime} \mathrm{d}_{G \backslash F}\left(y^{\prime}, z\right)<$
$\mathrm{d}_{G \backslash F}(y, z)$. (assuming $\epsilon^{\prime}<1$ )
This also contradicts the choice of $x, y, z$. (We should have preferred $x, y, z^{\prime}$ since $\mathrm{d}_{G \backslash F}\left(y, z^{\prime}\right)<\mathrm{d}_{G \backslash F}(y, z)$. See the lower part of Figure 6).

We are now ready to establish the correctness of our query algorithm.

Theorem 2.1. (Correctness) Our algorithm determines if $s$ and $t$ are connected in $G \backslash F$ and if they are connected we have that (see Figure 7) $\mathrm{d}_{G \backslash F}(s, t) \leq$ $\mathrm{d}_{H}(s, t) \leq(1+\epsilon) \mathrm{d}_{G \backslash F}(s, t)$.

Proof. An edge $(u, v)$ in $H$ corresponds to a path $P_{\alpha}$ from $u$ to $v$ in $G \backslash F$ and $w_{H}(u, v)=\left|P_{\alpha}\right|$ (the path must be in $G \backslash F$ because the only possibility for an edge ( $u, v$ ) to be added to $H$ is due to case 2 in Section 2.2 of the query procedure, when building the graph $H$ ). Therefore, if there is a path from $s$ to $t$ in $H$ of length $\ell^{\prime}$, then there is also a path from $s$ to $t$ in $G \backslash F$ of length $\ell^{\prime}$. Thus, $\mathrm{d}_{H}(s, t) \geq \mathrm{d}_{G \backslash F}(s, t)$ and in particular if $s$ and $t$ are disconnected in $G \backslash F$ then they are also disconnected in $H$. So in the rest of the proof we assume that there is (at least one) path from $s$ to $t$ in $G \backslash F$.

We prove by induction on $\mathrm{d}_{G \backslash F}(u, v)$ that for every pair of vertices $u, v \in V(H)$ we have that $\mathrm{d}_{G \backslash F}(u, v) \geq$ $\frac{\mathrm{d}_{H}(u, v)}{1+\epsilon}$ (or, equivalently, $\left.\mathrm{d}_{H}(u, v) \leq(1+\epsilon) \mathrm{d}_{G \backslash F}(u, v)\right)$. We assume that for every $u^{\prime}, v^{\prime} \in V(H)$ such that $\mathrm{d}_{G \backslash F}\left(u^{\prime}, v^{\prime}\right)<\mathrm{d}_{G \backslash F}(u, v)$ we have already established that $\mathrm{d}_{G \backslash F}\left(u^{\prime}, v^{\prime}\right) \geq \frac{\mathrm{d}_{H}\left(u^{\prime}, v^{\prime}\right)}{1+\epsilon}$, and prove the claim for $u, v$.

Let $P=P_{G \backslash F}(u, v)$ and consider the following two cases.
Case $1 \operatorname{tr}_{G \backslash F}(P) \cap V(H)=\emptyset$ : In this case it follows from Lemma 2.5 that $(u, v) \in E(H)$ and $w_{H}(u, v) \leq$ $|P|=\mathrm{d}_{G \backslash F}(u, v)$. So, $w_{H}(u, v)=\mathrm{d}_{G \backslash F}(u, v)$ and the lemma follows.
Case $2 \operatorname{tr}_{G \backslash F}(P) \cap V(H) \neq \emptyset$ : (See Figure 7) In this case by Lemma 2.6 there exist vertices $x, y, z$ such that $x \in\{u, v\}, y \in V(P)$, and $z \in V(H) \cap$ $\operatorname{tr}_{G \backslash F}(P), \mathrm{d}_{G \backslash F}(y, z) \leq \epsilon^{\prime}|P[x, y]|=\epsilon^{\prime} \mathrm{d}_{G \backslash F}(x, y)$, and $\operatorname{tr}\left(P_{G \backslash F}(x, y) \cdot P_{G \backslash F}(y, z)\right) \cap V(H)=\emptyset$. We assume for the rest of the proof that $x=u$. If $x=v$ then the proof is symmetric.

Since $\operatorname{tr}_{G \backslash F}\left(P(u, y) \cdot P_{G \backslash F}(y, z)\right) \cap V(H)=\emptyset$ it follows from Lemma 2.5 that $(u, z) \in E(H)$,

$$
\begin{align*}
& \left|w_{H}(u, z)\right| \leq\left|P_{G \backslash F}(u, y) \cdot P_{G \backslash F}(y, z)\right|=  \tag{2.3}\\
& \operatorname{dist}_{G \backslash F}(u, y)+\mathrm{d}_{G \backslash F}(y, z) .
\end{align*}
$$

Using the triangle inequality in $H$ and Equation


Figure 6: An illustration of the proof of Lemma 2.6.
(2.3) we have (see the edge $(u, z)$ in Figure 7)

$$
\begin{align*}
\mathrm{d}_{H}(u, v) \leq & w_{H}(u, z)+\mathrm{d}_{H}(z, v) \leq  \tag{2.4}\\
& \mathrm{d}_{G \backslash F}(u, y)+\mathrm{d}_{G \backslash F}(y, z)+\mathrm{d}_{H}(z, v)
\end{align*}
$$

Now we apply the induction hypothesis on $u^{\prime}=z$ and $v^{\prime}=v$. We can do so since every point in $\operatorname{tr}_{G \backslash F}(P) \cap V(H)$ is closer to $v$ than to $u$ and in particular, $\mathrm{d}_{G \backslash F}(z, v)<\mathrm{d}_{G \backslash F}(u, v)$ (see dashed edge in Figure 7). This application of the induction hypothesis together with Equation (2.4) gives

$$
\begin{align*}
\mathrm{d}_{H}(u, v) \leq \mathrm{d}_{G \backslash F}(u, y)+\mathrm{d}_{G \backslash F}( & y, z)+  \tag{2.5}\\
& (1+\epsilon) \mathrm{d}_{G \backslash F}(z, v)
\end{align*}
$$

Finally we apply the triangle inequality in $G \backslash F$ and the fact that $\mathrm{d}_{G \backslash F}(y, z) \leq \epsilon^{\prime} \mathrm{d}_{G \backslash F}(u, y)$ to obtain that
$\mathrm{d}_{H}(u, v) \leq$
$\mathrm{d}_{G \backslash F}(u, y)+\epsilon^{\prime} \mathrm{d}_{G \backslash F}(u, y)+(1+\epsilon)\left(\mathrm{d}_{G \backslash F}(z, y)+\mathrm{d}_{G \backslash F}(y, v)\right) \leq$
$\mathrm{d}_{G \backslash F}(u, y)+\epsilon^{\prime} \mathrm{d}_{G \backslash F}(u, y)+(1+\epsilon)\left(\epsilon^{\prime} \mathrm{d}_{G \backslash F}(u, y)+\mathrm{d}_{G \backslash F}(y, v)\right) \leq$
$\mathrm{d}_{G \backslash F}(u, y)+\epsilon^{\prime}(2+\epsilon) \mathrm{d}_{G \backslash F}(u, y)+(1+\epsilon) \mathrm{d}_{G \backslash F}(y, v) \leq$
$(1+\epsilon)\left(\mathrm{d}_{G \backslash F}(u, y)+\mathrm{d}_{G \backslash F}(y, v)\right) \leq$
$(1+\epsilon) \mathrm{d}_{G \backslash F}(u, v)$
where the next to last inequality holds since $\epsilon^{\prime}<$ $\epsilon / 3 \leq \epsilon /(\epsilon+2)$, and the last inequality holds since $P$ is a shortest path between $u$ and $v$ in $G \backslash F$. Therefore, $\mathrm{d}_{G \backslash F}(u, v) \geq \frac{\mathrm{d}_{H}(u, v)}{1+\epsilon}$, and we proved the induction theorem.

## 3 Reducing Space by Using Decomposable Paths.

Our data-structure in Section 2 requires $O\left(n^{3}(\log D / \epsilon)^{f}\right)$ space. In this section we develop a variant of this data structure which requires only $O\left(n^{2}(\log D / \epsilon)^{f} \cdot f \log D\right)$ space and keeps the query and preprocessing algorithms efficient. The main idea is a method to store the paths $P_{\alpha}$ implicitly rather than explicitly (in a hash table). To specify our new paths $P_{\alpha}$ and their representation we need the following definition.

Definition 3.1. ( $k$-DECOMPosable Path) Let $G$ be a graph and $G \backslash A$ a subgraph of $G$ obtained by discarding the edges in $A$. A $k$-decomposable path in $G \backslash A$ is a path which is the concatenation of at most $k+1$ shortest paths of $G$ interleaved with at most $k$ edges. Denote the shortest $k$-decomposable path from $u$ to $v$ in $G \backslash A$ by $P_{G \backslash A}^{k}(u, v)$.


Figure 7: An illustration of the proof of Theorem 2.1: $(u, z)$ is an edge in $H$; we apply the induction hypothesis on $z$ and $v$.

We also use the following Theorem of Afek et al. [2].
Theorem 3.1. Let $G$ be an undirected graph and let $F$ be a set of $f$ edges of $G$. Each shortest path in $G \backslash F$ is an $f$-decomposable path.

If we keep a representation of all shortest paths in $G$ (which takes $O\left(n^{2}\right)$ space) then we can represent an $f$-decomposable path $P$ using only $2 f+2$ vertices which are the endpoints of the $f+1$ shortest paths and the $f$ edges, whose concatenation is $P$. In particular, a shortest path in $G \backslash F$, for $F$ of size at most $f$, can be represented using $2 f+2$ vertices.

However when we try to apply this approach in our setting we tackle a technical hurdle: Specifically, the graphs $G_{\alpha}$ associated with the nodes $\alpha$ of the trees $F T(u, v)$ are obtained from $G$ by discarding $\leq f$ segments rather than edges, each of which may be long, and together these segments may contain much more than $f$ edges. So we cannot apply Theorem 3.1 to the paths $P_{\alpha}$ which are indeed shortest paths but in graphs that may have obtained from $G$ by discarding more than $f$ edges.

It follows that in order to use Theorem 3.1 we change the definition of the paths $P_{\alpha}$. A natural first attempt is to define $P_{\alpha}$ to be the shortest $f$ decomposable path between $u$ and $v$ in $G_{\alpha}$. The path $P=P_{G \backslash F}(u, v)$ is $f$-decomposable, so as long as $P$ survives in $G_{\alpha}$ it is a suitable candidate for being elected as $P_{\alpha}$.

But, for $P$ to survive as we go from a node $\alpha \in$ $F T(u, v)$ to its child, we need that no edge $e_{P} \in P$ will be contained together with an edge of $F$ in a segment of $P_{\alpha}$. In section 2 we guaranteed that this is indeed the case as long as all edges of $F$ are far from $P$ in $G \backslash F$, specifically, outside of $\operatorname{tr}_{G \backslash F}(P)$. The argument was based on the fact that if some edge $e_{P}=(x, y)$ of $P$
appears also on $P_{\alpha}$ (and assume it is closer to $u$ than to $v$ ) then the length of $\operatorname{seg}\left(e, P_{\alpha}\right)$ is about $\epsilon\left|P_{\alpha}[u, x]\right|$ which is not larger than $\epsilon|P[u, x]|$ since $P_{\alpha}$ was a shortest path in $G_{\alpha}$. Therefore this segment cannot reach an edge of $F$ which must be outside $\operatorname{tr}_{G \backslash F}(P)$. When $P_{\alpha}$ is a shortest $f$-decomposable path then it could be that $P_{\alpha}[u, x]$ is much larger than $P[u, x]$. This could happen, for example, if $P_{\alpha}[x, v]$ is a concatenation of, say $f-2$ shortest path in $G$, but also contains an edge in $F$ so it does not exist in $G \backslash F$. (We ignore for simplicity the interleaving edges in this example.) On the other hand $P[x, v]$ is a concatenation of $\delta \ll f$ shortest paths in $G$. This forces $P_{\alpha}[u, x]$ to be a concatenation of at most 2 shortest paths in $G$ whereas $P[u, x]$ can be assembled from many more shortest paths and therefore be much shorter than $P_{\alpha}[u, x]$.

To overcome this difficulty we introduce the following additional type of paths.

Definition 3.2. ( $k$-EXPATH) Let $G$ be a graph and $G \backslash A$ a subgraph of $G$ obtained by discarding the edges in $A$. $A$-expath in $G \backslash A$ is path in $G \backslash A$ which is a concatenation of $(1+2 \log D)$, $k$-decomposable paths in $G \backslash A$, such that the length of the $i$-th $k$-decomposable path is at most $\ell_{i}=\min \left\{2^{i}, 2^{2 \log D-i}\right\}$.

In the improved structure that we describe in this section, we choose $P_{\alpha}$ to be a shortest $(2 f+1)$-expath from $u$ to $v$ in $G_{\alpha}$. We prove that such a path has the following prefix optimality property. This property will allow us to prove that $P=P_{G \backslash F}(u, v)$ survives in the graphs $G_{\alpha}$ along the path that we traverse in $F T(u, v)$.

Lemma 3.1. Let $A \subset E$ and let $P$ be a shortest $k$-expath from $u$ to $v$ in $G \backslash A$. Then for every $x \in V(P)$ : $|P[u, x]| \leq 4\left|P_{G \backslash A}^{k}(u, x)\right|$ and $|P[x, v]| \leq 4\left|P_{G \backslash A}^{k}(x, v)\right|$.

Proof. Let $P$ be a shortest $k$-expath from $u$ to $v$ in $G \backslash A$. Assume, by contradiction, that there is a vertex $x \in V(P)$ such that $|P[u, x]|>4\left|P_{G \backslash A}^{k}(u, x)\right|$ or $|P[x, v]|>4\left|P_{G \backslash A}^{k}(x, v)\right|$. Assume without loss of generality that $|P[u, x]|>4\left|P_{G \backslash A}^{k}(u, x)\right|$ (the case where $|P[x, v]|>4\left|P_{G \backslash A}^{k}(x, v)\right|$ is symmetric). We show that $P^{\prime}=P_{G \backslash A}^{k}(u, x) \cdot P[x, v]$ is also a $k$-expath in $G \backslash A$, and since $P^{\prime}$ is shorter than $P$ this contradicts the assumption that $P$ is the shortest $k$-expath from $u$ to $v$.

Since $P$ is a $k$-expath in $G \backslash A$ then it is a concatenation of $k$-decomposable paths $P_{0}, P_{1}, \ldots, P_{2} \log D$ in $G \backslash A$, such that $\left|P_{i}\right| \leq \ell_{i}=\min \left\{2^{i}, 2^{2 \log D-i}\right\}$. To show that $P^{\prime}$ is a $k$-expath in $G \backslash A$ we define $k$-decomposable paths $P_{0}^{\prime}, P_{1}^{\prime}, \ldots, P_{2 \log D}^{\prime}$ in $G \backslash A$ such that $\left|P_{i}^{\prime}\right| \leq \ell_{i}$ and $P^{\prime}$ is the concatenation of $P_{0}^{\prime}, P_{1}^{\prime}, \ldots, P_{2}^{\prime} \log D$.

Since $\left|P_{i}\right| \leq \ell_{i} \leq 2^{i}$ it follows that $\sum_{i=0}^{t}\left|P_{i}\right|<2^{t+1}$. So if we let $j^{\prime}=\lfloor\log |P[u, x]|\rfloor-1$ then $\sum_{i=0}^{j^{\prime}}\left|P_{i}\right|<$ $|P[u, x]|$ and therefore $x$ must be contained in $P_{j}$ for some $j>j^{\prime}$. Since clearly $j^{\prime}<\log D$ we have that $\ell_{j^{\prime}}=2^{\lfloor\log \mid P[u, x]\rfloor\rfloor-1} \geq|P[u, x]| / 4>\left|P_{G \backslash A}^{k}(x, u)\right|$.

We are now ready to define $P_{0}^{\prime}, P_{1}^{\prime}, \ldots, P_{2 \log D}^{\prime}$. We set $P_{i}^{\prime}=\emptyset$ for every $0 \leq i<j^{\prime}, P_{j^{\prime}}^{\prime}=P_{G \backslash A}^{k}(u, x)$, $P_{i}^{\prime}=\emptyset$ for every $j^{\prime}<i<j, P_{j}^{\prime}$ to be the suffix of $P_{j}$ starting from $x$, and $P_{i}^{\prime}=P_{i}$ for every $i>j$.

It is easy to verify that for every $i, P_{i}^{\prime}$ is a $k$ decomposable path in $G \backslash A$ and that $\left|P_{i}^{\prime}\right| \leq \ell_{i}$. This proves that $P^{\prime}$ is a shorter $k$-expath from $u$ to $v$ than $P$ in $G \backslash A$, which contradicts the assumption that $P$ is the shortest $k$-expath from $u$ to $v$.

As already mentioned the critical change that we make to the structure of section 2 , in order to make it more space efficient, is to take $P_{\alpha}$ in each node $\alpha \in F T(u, v)$ to be a shortest $(2 f+1)$-expath from $u$ to $v$ in $G_{\alpha}$. We represent each such $(2 f+1)$-expath, $P_{\alpha}$, in a straightforward way, by a list of the endpoints of the paths, and the edges in the representation of each of its $\leq 2 \log D(2 f+1)$-decomposable paths. The size of this representation is $O(f \log D)$.

In addition, in order to be able to search the expaths efficiently, and later to report the shortest path from $s$ to $t$ in $G \backslash F$ if required, we store, as part of the data structure, all the shortest paths in $G$. Concretely, we store the values $\mathrm{d}_{G}(u, v)$, and $\pi(u, v)$, which is the predecessor of $v$ on the shortest path from $u$ to $v$, for every pair of vertices $u$ and $v$. We assume that shortest paths in $G$ are unique (which could be easily achieved with high probability, by slightly perturbing the edge weights).

The technical details of the construction and its analysis are given in the following sections, proving Theorem 1.1.
3.1 Reducing Space by Using Decomposable Paths - Technical Details. We first make the following technical change. In Section 2 we used $\epsilon^{\prime}=\epsilon / 3$ to define both the trapezoid of a path $P$, and the netpoints of a path $P$. Here, we change the definition of the netpoints of a path $P$ and use $\epsilon^{\prime \prime}=\epsilon^{\prime} / 4$ there instead of $\epsilon^{\prime}$, but keep $\epsilon^{\prime}$ in the definition of a trapezoid. The definition of netpoints is now:

Definition 3.3. (Path Netpoints) Let $P=\left(v_{1}=\right.$ $u, v_{2}, \ldots, v_{k}=v$ ) be a path from $u$ to $v$ in $G$. Let $L$ be the set of all vertices $v_{j}, v_{j+1} \in P$ such that $\left|P\left[u, v_{j}\right]\right|<$ $\left(1+\epsilon^{\prime \prime}\right)^{i} \leq\left|P\left[u, v_{j+1}\right]\right|$ for some integer $i \geq 0$. Define $R$ analogously to be the set of all vertices $v_{j}, v_{j+1} \in P$ such that $\left|P\left[v_{j+1}, v\right]\right|<\left(1+\epsilon^{\prime \prime}\right)^{i} \leq\left|P\left[v_{j}, v\right]\right|$ for some integer $i$. Let netpoints $(P)=L \cup R \cup\{u, v\}$.

This makes Lemma 2.1 hold with $\epsilon^{\prime \prime}$ rather than $\epsilon^{\prime}$. Specifically, we have

Lemma 3.2. Let $P$ be a path from $u$ to $v, e=$ $(x, y)$ be an edge of $P$ such that $|P[u, x]|<|P[u, y]|$. Then either $\operatorname{seg}(e, P)=\{e\}$ or $|\operatorname{seg}(e, P)| \leq$ $\epsilon^{\prime \prime} \min \{|P[u, x]|,|P[x, v]|\}$.

Clearly, the size of the data structure after these modifications is $O\left(n^{2}(\log D / \epsilon)^{f} \cdot f \log D\right)$. We now describe how to adapt the query algorithm to these modifications of the data structure.
Query. We implement a query as in Section 2 but we have to adapt the way we navigate through each tree $F T(u, v)$, to the new representation of the paths $P_{\alpha}$. This navigation requires an algorithm, that given an edge $e$ and an $(2 f+1)$-expath $P_{\alpha}$ determines if $e \in P_{\alpha}$, and if so, returns the index of $\operatorname{seg}\left(e, P_{\alpha}\right)$ (this is the index of the child of $\alpha$ to which the search continues).

An edge $e=(x, y)$ is on the shortest path in $G$ from $u$ to $v$ if and only if $\mathrm{d}_{G}(u, v)=\min \left\{\mathrm{d}_{G}(u, x)+\right.$ $\left.w(x, y)+\mathrm{d}_{G}(y, v), \mathrm{d}_{G}(u, y)+w(x, y)+\mathrm{d}_{G}(x, v)\right\}$. To test if $e$ is contained in an $f$-expath, $P_{\alpha}$, we simply run the test above for every shortest path (in $G$ ) in the representation of $P_{\alpha}$, and compare $e$ to all the edges in the representation of $P_{\alpha}$. If $e$ is contained in one of the shortest paths in the representation of $P_{\alpha}$, or if $e$ itself belongs to this representation, then $e$ is contained in $P_{\alpha}$ and otherwise it is not.

To find $\operatorname{seg}\left(e, P_{\alpha}\right)$ we compute $\left|P_{\alpha}[u, x]\right|$ and $\left|P_{\alpha}[x, v]\right|$ and $\left|P_{\alpha}\right|$. To compute $\left|P_{\alpha}[u, x]\right|$ we sum the lengths of the shortest paths and the edges, in the representation of $P_{\alpha}$, between $u$ and $e$, or the shortest path $Q$, in the representation of $P_{\alpha}$, containing $e$. If $e$ is contained in a shortest path $Q$ in the representation of $P_{\alpha}$, we also add $\mathrm{d}_{G}(w, x)$ where $w$ is the endpoint closer to $x$ of $Q$. The computation of $\left|P_{\alpha}[x, v]\right|$ is analogous. We
compute $\left|P_{\alpha}\right|$ by summing up the lengths of all the paths and edges in the representation of $P_{\alpha}$. Clearly, computing $\left|P_{\alpha}[u, x]\right|$ and $\left|P_{\alpha}[x, v]\right|$ and $\left|P_{\alpha}\right|$ takes $O(f \log D)$ time.

Once we have $\left|P_{\alpha}[u, x]\right|$ and $\left|P_{\alpha}[x, v]\right|$ we compute the largest index $i_{1}$ such that $\left(1+\epsilon^{\prime \prime}\right)^{i_{1}} \leq\left|P_{\alpha}[u, x]\right|$. The number of netpoints in $L$ that are between $u$ and $x$ on $P_{\alpha}$ (possibly including $x$ ) is $2 i_{1}$. Similarly, we compute the largest index $i_{2}$ such that $\left(1+\epsilon^{\prime \prime}\right)^{i_{2}} \leq\left|P_{\alpha}[x, v]\right|$ and then deduce that $2 i_{2}$ netpoints of $R$ are between $x$ and $v$ on $P_{\alpha}$. We also compute the total number of endpoints on $P_{\alpha}$ by computing the largest $i$ such that $\left(1+\epsilon^{\prime \prime}\right)^{i} \leq\left|P_{\alpha}\right|$. It is easy to see how to find $\operatorname{seg}\left(e, P_{\alpha}\right)$ from the indices $i_{1}, i_{2}$ and $i .{ }^{2}$ We compute $i_{1}, i_{2}$, and $i$ by a binary search of $\leq-n g t h S P_{\alpha}[u, x],\left|P_{\alpha}[x, v]\right|$, and $\left|P_{\alpha}\right|$, among the $\log _{1+\epsilon} D=(\log D) / \epsilon$ possible powers of $1+\epsilon^{\prime \prime} .{ }^{3}$

In total it takes us $O(f \log D+\log (\log D / \epsilon))=$ $O\left(f \log D+\log \frac{1}{\epsilon}\right)=O(f \log D)$ (note that the last equation holds since we assume $\epsilon>1 / D)$ time to decide if $e \in P_{\alpha}$ and to find $\operatorname{seg}\left(e, P_{\alpha}\right)$ if indeed $e \in P_{\alpha}$. Now recall that to perform a query we traverse $O\left(f^{2}\right)$ trees $F T(u, v)$, in each tree we traverse a path of length $f$, and in each node $\alpha$ we may have to scan through each edge of $F$ to find one that is contained in $P_{\alpha}$. For each such edge we apply the procedure just described. Therefore, the total query time is $O\left(f^{5} \log D\right)$.
Preprocessing. For every pair of vertices $u$ and $v$ and for every node $\alpha \in F T(u, v)$ we have to compute the shortest $(2 f+1)$-expath from $u$ to $v$ in $G_{\alpha}$. We do this in three steps as follows.

1) We compute all pairs shortest paths in $G_{\alpha}$ and identify the pairs of vertices $x$ and $y$ for which $\mathrm{d}_{G}(x, y)=$ $\mathrm{d}_{G_{\alpha}}(x, y)$. These are the pairs for which the shortest path in $G_{\alpha}$ is the same as the shortest path in $G$.
2) In this step we compute for every pair of vertices $x$ and $y$ the shortest $(2 f+1)$-decomposable path in $G_{\alpha}$. We form a layered directed graph $G_{\alpha}^{2}$ consisting of $2(2 f+1)-1$ layers $V_{1}, V_{2}, \ldots, V_{2(2 f+1)-1}$. Each layer contains a copy of each vertex of $G$. We denote the copy of a vertex $x$ at layer $i$ by $x_{i}$. The arcs of this layered graph are defined as follows. For $x \neq y$, we add an arc $\left(x_{1}, y_{2}\right)$ if and only if $\mathrm{d}_{G}(x, y)=\mathrm{d}_{G_{\alpha}}(x, y)$. We set the weight of the $\operatorname{arc}\left(x_{1}, y_{2}\right)$ to be equal to $\mathrm{d}_{G}(x, y)$. We put an $\operatorname{arc}\left(x_{2}, y_{3}\right)$ if there is an an edge $(x, y)$ in $G_{\alpha}$. We set the length of the $\operatorname{arc}\left(x_{2}, y_{3}\right)$ to be equal to the length of $(x, y)$. We add arcs between subsequent levels so that for odd $i$, the arcs between $V_{i}$ and $V_{i+1}$ are the same as the arcs between $V_{1}$ and $V_{2}$, and for even $i$,

[^1]the arcs between $V_{i}$ and $V_{i+1}$ are the same as the arcs between $V_{2}$ and $V_{3}$. I.e. if we have an $\operatorname{arc}\left(x_{1}, y_{2}\right)$ we also have an arc $\left(x_{i}, y_{i+1}\right)$ of the same length for odd $i$, and if we have an $\operatorname{arc}\left(x_{2}, y_{3}\right)$ we also have an $\operatorname{arc}$ $\left(x_{i}, y_{i+1}\right)$ of the same length for even $i$. We also put an $\operatorname{arc}\left(x_{i}, x_{i+1}\right)$ of length 0 for every $x$ and $i$.

We compute the shortest path in $G_{\alpha}^{2}$ from every vertex in the first layer to every vertex in the last layer. It is easy to verify that the length of the shortest path from $x_{1}$ to $y_{2(2 f+1)-1}$ in $G_{\alpha}^{2}$ is equal to the length of the shortest $(2 f+1)$-decomposable path from $x$ to $y$ in $G_{\alpha}$. So following this step we know $P_{G_{\alpha}}^{2 f+1}(x, y)$ for every pair $x, y$ in $G_{\alpha}$.
3) In this step we also build a directed layered network $G_{\alpha}^{3}$ with $2 \log D+1$ layers $V_{1}, V_{2}, \ldots, V_{2 \log D+1}$. Each layer contains a copy of every vertex of $G$. We denote the copy of a vertex $x$ at level $i$ by $x_{i}$. For $x \neq$ $y$ we add an $\operatorname{arc}\left(x_{i}, y_{i+1}\right)$ of length $\left|P_{G_{\alpha}}^{2 f+1}(x, y)\right|$ if $\left|P_{G_{\alpha}}^{2 f+1}(x, y)\right| \leq \ell_{i}$. For every vertex $u$ and index $i$ we add an $\operatorname{arc}\left(x_{i}, x_{i+1}\right)$ of length 0.

Recall that we want to set $P_{\alpha}$ to be the shortest $(2 f+1)$-expath from $u$ to $v$ in $G_{\alpha}$. We compute the shortest path $P^{3}$ from $u_{1}$ to $v_{2} \log D$ in $G_{\alpha}^{3}$. The length of $P^{3}$ is equal to the length of $P_{\alpha}$. Let $P_{1}, P_{2}, \ldots, P_{2} \log D$ be the $(2 f+1)$-decomposable paths composing $P_{\alpha}$. The $i$ th edge of $P^{3}$ corresponds to $P_{i}$. If this edge is of the form $\left(x_{i}, x_{i+1}\right)$ then $P_{i}$ is empty and if it is of the form $\left(x_{i}, y_{i+1}\right)$ for $y \neq x$ then $P_{i}=P_{G_{\alpha}}^{2 f+1}(x, y)$. Following these observations we extract the representation of $P_{\alpha}$ from $G_{\alpha}^{2}$ and $G_{\alpha}^{3}$ and store it at $\alpha$.

The graph $G_{\alpha}^{2}$ has $O\left(f n^{2}\right)$ arcs so we can compute all the shortest paths in $G_{\alpha}^{2}$ from vertices of the first layer to vertices of the last layer in $O\left(f n^{3}\right)$ time (using Dijkstra's algorithm to compute shortest paths from every vertex in the first layer). The graph $G_{\alpha}^{3}$ has $O\left(n^{2} \log D\right) \operatorname{arcs}$ so we can compute the shortest path from $u_{1}$ to $v_{2} \log D$ in $G_{\alpha}^{3}$ in $O\left(n^{2} \log D+\right.$ $n \log D \log n \log D)=O\left(n^{2} \log D+n \log D \log \log D\right)$ time. It follows that the preprocessing time at each node $\alpha$ is $O\left(f n^{3}+n^{2} \log D+n \log D \log \log D\right)$ and by multiplying by the total number of nodes in all the trees $F T(u, v)$ we get that the total preprocessing time is $O\left(\left(f n^{5}+n^{4} \log D+n^{3} \log D \log \log D\right)(\log D / \epsilon)^{f}\right)$.

The following theorem summarizes the properties of the data structure that we have presented and proves its correctness.

ThEOREM 3.2. The $f$-sensitive distance oracle presented in this section requires $O\left(n^{2}(\log D / \epsilon)^{f}\right.$. $f \log D)$ space, can be built in $O\left(\left(f n^{5}+n^{4} \log D+\right.\right.$ $\left.\left.n^{3} \log D \log \log D\right)(\log D / \epsilon)^{f}\right)$ time, and can answer a query in $O\left(f^{5} \log D\right)$ time. Our distance oracle can also report an approximate shortest path from s to $t$ in $G \backslash F$
in time proportional to the length of this path.
Proof. The running times specified in the statement of theorem are clear from the preceding discussion. We are only left to prove the correctness of the data structure.

For correctness we follow the footsteps of the correctness proof in Section 2. A close inspection of the proof shows that the modifications which we did in the structure affects mainly Lemma 2.5.

In Section 2 we proved Lemma 2.5 for an arbitrary path $P$ such that $\operatorname{tr}_{G \backslash F}(P) \cap V(H)=\emptyset$. However, when we established correctness in Theorem 2.1, we used Lemma 2.5 only for paths $P$ which are shortest paths in $G \backslash F$ or concatenations of two shortest paths in $G \backslash F$. A shortest path in $G \backslash F$ is $f$-decomposable by Theorem 3.1. It follows immediately that a path which is a concatenation of two shortest paths in $G \backslash F$ is $(2 f+1)$-decomposable. Therefore it suffices if we reprove Lemma 2.5 for $(2 f+1)$-decomposable paths.

We established Lemma 2.5 in Section 2 by showing that $P$ is contained in all graphs $G_{\alpha}$, for nodes $\alpha$ along the path that the query algorithm follows in $F T(u, v)$. We showed this by induction on the nodes of this path starting from the root. We do the same here and indicate the modifications needed in the induction step using the same notation as in the proof of Lemma 2.5. Here we assume that $P$ is a $(2 f+1)$-decomposable path from $u$ to $v$ in $G \backslash F$ such that $\operatorname{tr}(P) \cap V(H)=\emptyset$.

The critical Equation there in the proof of Lemma 2.5 is Equation (2.1). This equation does not hold here since $P_{\alpha}$ is not a shortest path anymore. But since $P_{\alpha}$ is a shortest $(2 f+1)$-expath then by Lemma 3.1 we have

$$
\begin{array}{r}
\min \left\{\left|P_{\alpha}[u, x]\right|,\left|P_{\alpha}[x, v]\right|\right\} \leq  \tag{3.6}\\
4 \min \left\{\left|P_{G \backslash F}^{2 f+1}(u, x)\right|,\left|P_{G \backslash F}^{2 f+1}(x, v)\right|\right\}
\end{array}
$$

Continuing exactly the same reasoning as in the proof of Lemma 2.5 we get

$$
\begin{array}{r}
\left|\operatorname{seg}\left(e_{F}, P_{\alpha}\right)\right| \geq\left|P_{\alpha}[x, y]\right| \geq \mathrm{d}_{G \backslash F}(x, y) \\
>\epsilon^{\prime} \min (|P(u, x)|,|P(x, v)|) \\
\geq \epsilon^{\prime} \min \left\{\left|P_{G \backslash F}^{2 f+1}(u, x)\right|,\left|P_{G \backslash F}^{2 f+1}(x, v)\right|\right. \\
\geq  \tag{3.9}\\
\frac{\epsilon^{\prime}}{4} \min \left\{\left|P_{\alpha}[u, x]\right|,\left|P_{\alpha}[x, v]\right|\right\} \\
= \\
=\epsilon^{\prime \prime} \min \left\{\left|P_{\alpha}[u, x]\right|,\left|P_{\alpha}[x, v]\right|\right\}
\end{array}
$$

where Equation (3.7) follows since $\operatorname{tr}(P) \cap V(H)=$ $\emptyset$, Equation (3.8) follows since $P$ is a $(2 f+1)$ decomposable path, and Equation (3.9) follows by Equation (3.6).

However, since we modified our definition of netpoints to use an exponential scale with base $1+\epsilon^{\prime \prime}$, it
follows from Lemma 3.2 that

$$
\begin{equation*}
\left|\operatorname{seg}\left(e_{F}, P_{\alpha}\right)\right| \leq \epsilon^{\prime \prime} \min \left\{\left|P_{\alpha}[u, x]\right|,\left|P_{\alpha}[x, v]\right|\right\} \tag{3.10}
\end{equation*}
$$

which is a contradiction.
To report an approximate shortest path from $s$ to $t$ in $G \backslash F$ we follow $P_{H}(s, t)$. Each edge $e=(u, v)$ on $P_{H}(s, t)$ corresponds to some path $P_{\alpha}$ in $G \backslash F$ which is stored at node $\alpha$ of $F T(u, v)$. We report the concatenation of these paths $P_{\alpha}$ for all the edges along $P_{H}(s, t)$. Each such path $P_{\alpha}$ is by itself a concatenation of $O(f \log D)$ shortest paths and edges of $G$. To report each shortest path in the representation of $P_{\alpha}$, we use the table $\pi$. Recall that for every pair of vertices $u$ and $v, \pi(u, v)$ is the predecessor of $v$ on the shortest path from $u$ to $v$.

## 4 A Reduction From Arbitrary Weights To Bounded Weights.

In this section we reduce size of the data-structure by replacing the $\log ^{f} D$ factor in the size complexity to $\log ^{f} n$. We do this by a reduction from the problem of $f$ sensitive distance oracles on a graph with arbitrary edge weights to $\frac{\log W}{\log n}$ independent problems of $f$-sensitive distance oracles on graphs with bounded weights, where the minimum edge weight is 1 and the heaviest edge weight is $O(\operatorname{poly}(n))$.

Lemma 4.1. Assume we have an $f$-sensitive distance oracle with stretch $\alpha$ for bounded weights graphs (where the minimum edge weight is 1 and the heaviest edge weight $W=O(\operatorname{poly}(n)))$ with $f_{s}(n)$ size, $f_{q}(n)$ query time and $f_{p}(n)$ preprocessing time. Then we can build an $f$-sensitive distance oracle for a graph with arbitrary edge weights whose size is $O\left(f_{s}(n) \cdot \frac{\log W}{\log n}\right)$, with $(1+$ $o(1)) \alpha$ stretch, $O\left(f_{q}(n) \cdot \log \log W\right)$ query time and $\left.O\left(f_{p}(n) \cdot \frac{\log W}{\log n}\right)\right)$ preprocessing time.

Proof. We first define graphs $G^{i}$ and $\widetilde{G^{i}}$ for $1 \leq i \leq$ $\frac{\log W}{\log n}$, which we derive from $G$ using different rounding scales. We show how to find the index $i$ such that the shortest $s$ to $t$ path in $G^{i}$ has $(1+o(1)) \alpha$ stretch compared to the shortest $s$ to $t$ path in $G \backslash F$.

Define $G^{i}$ to be the following graph. Replace edge weights $\leq n^{i-2}$ by $n^{i-2}$. Discard all edges of weight $\geq n^{i+1}$. We show the following monotonicity property on the graphs $G^{i}$. There is some index $i$ such that $s$ and $t$ are disconnected in $G^{i^{\prime}}$ for every $i^{\prime} \leq i-2$, while $s$ and $t$ are connected in $G^{i^{\prime}}$ for every $i^{\prime} \geq i$.

Let $W_{(F, s, t)}$ be the weight of the heaviest edge on the shortest path from $s$ to $t$ in the graph $G \backslash F$. Let $i=\left\lfloor\log _{n} W_{(F, s, t)}\right\rfloor$, then $n^{i} \leq W_{(F, s, t)}<n^{i+1}$. Note that this implies $\mathrm{d}_{G \backslash F}(s, t) \geq n^{i}$.

First we show that $s$ and $t$ are disconnected in $G^{i^{\prime}} \backslash F$ for every $i^{\prime} \leq i-2$. To see this, note that the maximal edge weight in $G^{i^{i}}$ is less than $n^{i^{\prime}+1} \leq n^{i-1}$. hence, if $s$ and $t$ are connected in $G^{i^{i}} \backslash F$ then they are connected by a path length less than $n^{i}$, contradiction to the fact that $\mathrm{d}_{G \backslash F}(s, t) \geq n^{i}$.

Next, we show that $s$ and $t$ are connected in $G^{i^{i}} \backslash F$ for every $i^{\prime} \geq i$. Consider $i^{\prime} \geq i$. Recall that $G^{i^{\prime}}$ contains all edges whose weight is $<n^{i^{\prime}+1}-$ in particular $G^{i^{i}}$ contains all edges whose weight is $<n^{i+1}$. It follows that $G^{i^{\prime}} \backslash F$, contains all edges of $G \backslash F$ whose weight is $<n^{i+1}$. Hence, $G^{i^{\prime}} \backslash F$ contains all edges along the shortest path from $s$ to $t$ in $G \backslash F$. Thus, $s$ and $t$ are connected in $G^{i^{\prime}} \backslash F$.

We now show that querying $(F, s, t)$ in $G^{i}$ returns a $(1+o(1)) \alpha$ approximate shortest path in the graph $G \backslash F$. We first show that $\mathrm{d}_{G^{i} \backslash F}(s, t) \leq(1+1 / n) \mathrm{d}_{G \backslash F}(s, t)$. Recall that $\mathrm{d}_{G \backslash F}(s, t) \geq n^{i}$. Rounding up all edges of weight less than $n^{i-2}$ to $n^{i-2}$ increases the shortest path by at most $(n-1) \cdot n^{i-2}<n^{i-1}$. Hence $\mathrm{d}_{G^{i} \backslash F}(s, t) \leq$ $\mathrm{d}_{G \backslash F}(s, t)+n^{i-1} \leq(1+1 / n) \mathrm{d}_{G \backslash F}(s, t)$. Recall that the stretch of $D^{i}$ is $\alpha$ which means it returns an $\alpha$ approximation of $\mathrm{d}_{G^{i} \backslash F}(s, t)$. It follows that the length we return $\leq \alpha \mathrm{d}_{G^{i} \backslash F}(s, t) \leq(1+1 / n) \alpha \mathrm{d}_{G \backslash F}(s, t)=$ $(1+o(1)) \alpha \mathrm{d}_{G \backslash F}(s, t)$.

The graph $G_{i}$ has minimum edge weight $n^{i-2}$ and maximum edge weight $n^{i+1}$, we want to scale it so that the minimum edge weigh will be 1. Divide all edge weights of $G^{i}$ by $n^{i-2}$ and round upwards to the closest integer, call the resulting graph $\widetilde{G^{i}}$. Obviously, in $\widetilde{G^{i}}$, the minimum edge weight is 1 and maximum edge weight is $n^{3}$. We now show that $n^{i-2} \cdot \mathrm{~d}_{\widetilde{G^{i}} \backslash F}(s, t) \leq$ $(1+1 / n) \mathrm{d}_{G^{i} \backslash F}(s, t)$. The length of any simple $s$ to $t$ path in $G^{i}$ is at least $n^{i}$, and if we round the edge weights upwards to the closest multiple of $n^{i-2}$ we add to it a length of at most $(n-1) \cdot n^{i-2}<n^{i-1}$. Thus, the length of the path in $\widetilde{G^{i}}$ multiplied by $n^{i-2}$ is at most $(1+1 / n)$ the length in $G^{i}$. We conclude that $n^{i-2} \cdot \mathrm{~d}_{\widetilde{G^{i}} \backslash F}(s, t) \leq(1+1 / n) \mathrm{d}_{G^{i} \backslash F}(s, t) \leq(1+$ $o(1)) \alpha \mathrm{d}_{G \backslash F}(s, t)$, where in the last inequality we used that $\mathrm{d}_{G^{i} \backslash F}(s, t) \leq(1+o(1)) \alpha \mathrm{d}_{G \backslash F}(s, t)$.

So far, we have described the graphs $G^{i}$ and $\widetilde{G^{i}}$. We now describe the data-structure we build. In the preprocessing stage we compute all $G^{i}, \widetilde{G^{i}}$ for all $1 \leq i \leq$ $\frac{\log W}{\log n}$, and store the an $f$-sensitive distance oracle $\bar{D}_{i}$ for every $\widetilde{G^{i}}$. Thus, space requirements are $O\left(f_{s}(n) \cdot \frac{\log W}{\log n}\right)$ and preprocessing time is $\left.O\left(f_{p}(n) \cdot \frac{\log W}{\log n}\right)\right)$. We now describe the procedure for answering query $(F, s, t)$. We first do a binary search in the range $1 \leq i \leq \frac{\log W}{\log n}$, to find the minimal value $i$ such that $s$ and $t$ are connected in $G^{i}$ (querying the distance oracle $D_{i}$ with
$(F, s, t)$ answers if $s$ and $t$ are connected or not in $\left.G^{i}\right)$. Then, using the distance oracles $D_{i}$ and $D_{i+1}$ we query $(F, s, t)$ and multiply the result with $n^{i-2}$ and $n^{i-1}$, respectively. We return the minimum of both alternatives. The claims above prove that the stretch is $(1+o(1)) \alpha$. Due to the binary search, total query time is $O\left(f_{q}(n) \cdot \log \log W\right)$.

Observe that when we substitute $D=\operatorname{poly}(n)$ in Theorem 1.1, we obtain an $f$-sensitive distance oracle with stretch $\alpha=1+\epsilon$ with the following properties:

- space complexity: $f_{s}(n)=O\left(n^{2}(\log n / \epsilon)^{f} \cdot f \log n\right)$.
- query time: $f_{q}(n)=O\left(f^{5} \log n\right)$.
- preprocessing time: $f_{p}(n)=O\left(f n^{5}(\log n / \epsilon)^{f}\right)$.

By substituting these functions $f_{s}(n), f_{q}(n), f_{p}(n)$ in Lemma 4.1, we obtain the following Theorem.

THEOREM 4.1. Let $G$ be a weighted undirected graph. There exists $f$-sensitive distance oracle for $G$ with $(1+\epsilon)$ stretch, $O\left(n^{2}(\log n / \epsilon)^{f} \cdot f \log W\right)$ space, and $O\left(f^{5} \log n \log \log W\right)$ query time. The oracle can be built in $O\left(f n^{5}(\log n / \epsilon)^{f} \frac{\log W}{\log n}\right)$ time.

## 5 Conclusion.

In this paper, we present the first $f$-sensitive distance oracle with $(1+\epsilon)$-stretch with small space and fast query time that handles multiple edge failures. For any $f=o\left(\frac{\log n}{\log \log n}\right)$ and fixed $\epsilon>0$ our $f$-sensitive distance oracle has $n^{2+o(1)}$ space, $\widetilde{O}(1)$ query time and $(1+\epsilon)$ stretch.

This is a big step forward in our understanding of this problem. We leave however many open questions. Firstly, our data-structure has large space requirements when $f \gg \frac{\log n}{\log \log n}$ which are dominated by the factor of $O\left(\left(\frac{\log n}{\epsilon}\right)^{f}\right)$ in the space complexity. Can the space be reduced? Secondly, our data-structure deals with at most $f$ edge-failures, what can be done in the case of vertex failures? Thirdly, can one develop efficient distance-sensitive oracles in the case of directed graphs for multiple failures? Another interesting direction is: can we develop efficient oracles for the seemingly easier problem of reachability ("is there some path from $u$ to $v$ after the failures?") for multiple edge/vertex failures?

## References

[1] Ittai Abraham, Shiri Chechik, and Kunal Talwar. Fully dynamic all-pairs shortest paths: Breaking the $o(n)$ barrier. In International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX), pages 1-16, 2014.
[2] 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.
[3] Giorgio Ausiello, Giuseppe F. Italiano, Alberto Marchetti-Spaccamela, and Umberto Nanni. Incremental algorithms for minimal length paths. J. Algorithms, 12(4):615-638, 1991.
[4] Surender Baswana, Sumeet Khurana, and Soumojit Sarkar. Fully dynamic randomized algorithms for graph spanners. ACM Transactions on Algorithms, 8(4):35, 2012.
[5] Aaron Bernstein. Fully dynamic ( $2+$ epsilon) approximate all-pairs shortest paths with fast query and close to linear update time. In Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science, FOCS, pages 693-702, 2009.
[6] Aaron Bernstein and David Karger. Improved distance sensitivity oracles via random sampling. In Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, pages 34-43, 2008.
[7] Aaron Bernstein and David Karger. A nearly optimal oracle for avoiding failed vertices and edges. In Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing, STOC, pages 101-110, 2009.
[8] Shiri Chechik, Michael Langberg, David Peleg, and Liam Roditty. f-sensitivity distance oracles and routing schemes. Algorithmica, 63(4):861-882, 2012.
[9] Camil Demetrescu and Giuseppe F. Italiano. Fully dynamic all pairs shortest paths with real edge weights. J. Comput. Syst. Sci., 72(5):813-837, 2006.
[10] Camil Demetrescu and Mikkel Thorup. Oracles for distances avoiding a link-failure. In Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, pages 838-843, 2002.
[11] Camil Demetrescu, Mikkel Thorup, Rezaul Alam Chowdhury, and Vijaya Ramachandran. Oracles for distances avoiding a failed node or link. SIAM J. Comput., 37(5):1299-1318, January 2008.
[12] Ran Duan and Seth Pettie. Dual-failure distance and connectivity oracles. In Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, pages 506-515, 2009.
[13] Shimon Even and Yossi Shiloach. An on-line edgedeletion problem. J. ACM, 28(1):1-4, 1981.
[14] Fabrizio Grandoni and Virginia Vassilevska Williams. Improved distance sensitivity oracles via fast singlesource replacement paths. In Proceedings of the 53rd Annual Symposium on Foundations of Computer Science, FOCS, pages 748-757, 2012.
[15] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Dynamic approximate all-pairs shortest paths: Breaking the $o(m n)$ barrier and derandomization. In Proceedings of the 54th Annual Symposium on Foundations of Computer Science, FOCS, pages 538-547, 2013.
[16] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Sublinear-time maintenance
of breadth-first spanning tree in partially dynamic networks. In Proceedings of the 40th International Colloquium, ICALP, pages 607-619, 2013.
[17] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Sublinear-time decremental algorithms for single-source reachability and shortest paths on directed graphs. In Proceedings of the $46 t h$ Annual ACM Symposium on Theory of Computing (STOC), pages 674-683, 2014.
[18] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. A subquadratic-time algorithm for decremental single-source shortest paths. In Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, pages 1053-1072, 2014.
[19] Neelesh Khanna and Surender Baswana. Approximate shortest paths avoiding a failed vertex: Optimal size data structures for unweighted graphs. In ${ }^{27}$ th International Symposium on Theoretical Aspects of Computer Science, STACS, pages 513-524, 2010.
[20] Valerie King. Fully dynamic algorithms for maintaining all-pairs shortest paths and transitive closure in digraphs. In Proceedings of the 40 th Annual Symposium on Foundations of Computer Science, FOCS, pages 8191, 1999.
[21] Valerie King and Mikkel Thorup. A space saving trick for directed dynamic transitive closure and shortest path algorithms. In COCOON, pages 268-277, 2001.
[22] Liam Roditty and Uri Zwick. On dynamic shortest paths problems. Algorithmica, 61(2):389-401, 2011.
[23] Liam Roditty and Uri Zwick. Dynamic approximate all-pairs shortest paths in undirected graphs. SIAM J. Comput., 41(3):670-683, 2012.
[24] Mikkel Thorup. Fully-dynamic all-pairs shortest paths: Faster and allowing negative cycles. In $S W A T$, pages 384-396, 2004.
[25] Mikkel Thorup. Worst-case update times for fullydynamic all-pairs shortest paths. In Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing (STOC), pages 112-119, 2005.
[26] Oren Weimann and Raphael Yuster. Replacement paths and distance sensitivity oracles via fast matrix multiplication. ACM Trans. Algorithms, 9(2):14:114:13, March 2013. Announced at FOCS'10.
[27] Ran Duan. Algorithms and dynamic data structures for basic graph optimization problems. PhD Thesis, University of Michigan, 2011.

## A A Simple But Space Inefficient Construction.

We present the following construction in an attempt to make the construction in the body of the paper more accessible. This data-structure can accommodate up to $f$ edge failures and returns the exact shortest path in the graph without the failed edges. Unfortunately, the space requirements are quite large, $O\left(n^{f+3}\right)$. It may be helpful to consider Figure 8 when reading the description of the following data structure and query algorithm:

- for every pair of vertices $u, v \in V$, construct a tree BruteForce $(u, v)$, of height at most $f$ and every node has at most $n-1$ children.
- The root $r$ stores the shortest path from $u$ to $v$ in $G$.
- Let $\alpha$ be an arbitrary node of the tree, and let $G_{\alpha}$ be a subgraph of $G$ associated with the node $\alpha$ (we define $G_{\alpha}$ recursively below).
- Let $P_{\alpha}$ be the shortest path between $u$ and $v$ in $G_{\alpha}$ (denote by $P_{\alpha}=\emptyset$ if $u$ and $v$ are in different connected components of $G_{\alpha}$ ). Store the path $P_{\alpha}$ in node $\alpha$.
- If $P_{\alpha}=\emptyset$ or $\alpha$ is at depth $f$ then $\alpha$ is a leaf of BruteForce $(u, v)$.
- If the depth of $\alpha<f$ and $P_{\alpha} \neq \emptyset$, then we create a child node $\alpha^{\prime}$ for every edge $e \in P_{\alpha}$. We associate the edge $e$ with the tree arc from $\alpha$ to its child $\alpha^{\prime}$ and denote it by $S_{\alpha^{\prime}}=\{e\}$.
- We now recursively define $\operatorname{avoid}(\alpha)$, a set of edges we discard from $G$ at node $\alpha$. For the root $r, \operatorname{avoid}(r)=\emptyset$. For a node $\alpha^{\prime}$ which is a child of $\alpha$, we define $\operatorname{avoid}\left(\alpha^{\prime}\right)=$ $\operatorname{avoid}(\alpha) \cup S_{\alpha^{\prime}}$. In other words, $\operatorname{avoid}(\alpha)$ is the set of edges associated with the arcs along the tree path from the root to $\alpha$.
- We define $G_{\alpha}=G \backslash \operatorname{avoid}(\alpha)$, that is the graph $G$ after all edges in avoid $(\alpha)$ are discarded.
- BruteForce $(u, v)$ is a tree of height at most $f$, with degree at most $n$, and hence it contains $O\left(n^{f}\right)$ nodes. In each node $\alpha$ we store the path $P_{\alpha}$ using $O(n)$ space per node. Thus, each tree occupies $O\left(n^{3}\right)$ space. There are $O\left(n^{2}\right)$ such trees (one for each pair of vertices $u, v \in V$ ), so total space is $O\left(n^{f+3}\right)$.
- To answer the query $(F, s, t)$ we search the tree BruteForce ( $s, t$ ) using the set of failed edges $F$ as follows: We start with the root of BruteForce( $s, t$ ), if the path stored in the root does not contain
any of the edges $F$ - the answer to the query is the path stored in the root. Otherwise, we branch to a child node using an arbitrary edge in the intersection of $F$ and the path stored at the root of BruteForce( $s, t)$. We continue this search procedure until we either reach a leaf where $s$ and $t$ are disconnected (in which case we reply that $s$ and $t$ are disconnected in the graph $G$ after all edges in $F$ fail) or until we reach a node containing a path which is edge disjoint from $F$. This path is a shortest path from $s$ to $t$ in the graph $G$ after all edges in $F$ fail. The complexity of such a query is $O\left(f^{2}\right)$ since the query follows a downwards path from the root of BruteForce $(s, t)$, (i.e., of length $\leq f$ ), and we can check the intersection of $F$ with the path stored in any node in time $O(f)$ using a hash table.


Figure 8: Illustration of the simple and exact $O\left(n^{f+3}\right)$ space algorithm. Every node has at most $n-1$ children, one for each edge in the path stored within the node. Node $\beta$ is a child of node $r$, the path associated with node $r$ is the path $u \rightarrow x \rightarrow v$, the shortest path from $u$ to $v$ in the graph $G$. The path associated with node $\beta$, $u \rightarrow y \rightarrow z \rightarrow v$, is the shortest path from $u$ to $v$ after the edge $(x, v)$ is deleted from $G$. The path associated with node $\gamma$ is the shortest path from $u$ to $v$ after both edges $(x, v)$ and $(u, y)$ have been deleted from $G$, if such a path exists, or $\emptyset$ if there is no path from $u$ to $v$ in this graph.


[^0]:    ${ }^{1}$ We use the notation $V(F)$ for a subset of edges $F$, to denote the endpoints of the edges of $F$. We use the notation $V(H)$ for a graph $H$ to denote the set of vertices of the graph.

[^1]:    ${ }^{2} \mathrm{Here}$ we assume that we do not remove duplicate netpoints but introduce empty segments between them.
    ${ }^{3}$ These powers are independent of the query and can be stored with the data structure.

