

Decremental APSP in Unweighted Digraphs Versus an Adaptive Adversary

Jacob Evald ✉

BARC, University of Copenhagen, Denmark

Viktor Fredslund-Hansen ✉ 

BARC, University of Copenhagen, Denmark

Maximilian Probst Gutenberg ✉ 

ETH Zurich, Switzerland

Christian Wulff-Nilsen ✉ 

BARC, University of Copenhagen, Denmark

Abstract

Given an unweighted digraph $G = (V, E)$, undergoing a sequence of edge deletions, with $m = |E|, n = |V|$, we consider the problem of maintaining all-pairs shortest paths (APSP).

Whilst this problem has been studied in a long line of research [ACM'81, FOCS'99, FOCS'01, STOC'02, STOC'03, SWAT'04, STOC'13] and the problem of $(1 + \epsilon)$ -approximate, weighted APSP was solved to near-optimal update time $\tilde{O}(mn)$ by Bernstein [STOC'13], the problem has mainly been studied in the context of an *oblivious* adversary which fixes the update sequence before the algorithm is started. In this paper, we make significant progress on the problem for an adaptive adversary which can perform updates based on answers to previous queries:

- We first present a *deterministic* data structure that maintains the *exact* distances with total update time $\tilde{O}(n^3)^1$.
- We also present a *deterministic* data structure that maintains $(1 + \epsilon)$ -approximate distance estimates with total update time $\tilde{O}(\sqrt{mn}^2/\epsilon)$ which for sparse graphs is $\tilde{O}(n^{2+1/2}/\epsilon)$.
- Finally, we present a randomized $(1 + \epsilon)$ -approximate data structure which works against an adaptive adversary; its total update time is $\tilde{O}(m^{2/3}n^{5/3} + n^{8/3}/(m^{1/3}\epsilon^2))$ which for sparse graphs is $\tilde{O}(n^{2+1/3}/\epsilon^2)$.

Our exact data structure matches the total update time of the best *randomized* data structure by Baswana et al. [STOC'02] and maintains the distance matrix in near-optimal time. Our approximate data structures improve upon the best data structures against an adaptive adversary which have $\tilde{O}(mn^2)$ total update time [JACM'81, STOC'03].

2012 ACM Subject Classification Theory of computation → Data structures design and analysis; Theory of computation → Shortest paths; Theory of computation → Dynamic graph algorithms

Keywords and phrases Dynamic Graph Algorithm, Data Structure, Shortest Paths

Digital Object Identifier 10.4230/LIPIcs.ICALP.2021.64

Category Track A: Algorithms, Complexity and Games

Related Version *Extended Version:* [arXiv:2010.00937](https://arxiv.org/abs/2010.00937) [5]

Funding *Jacob Evald:* The author is supported by the Starting Grant 7027-00050B from the Independent Research Fund Denmark under the Sapere Aude research career programme.

Viktor Fredslund-Hansen: The author is supported by the Starting Grant 7027-00050B from the Independent Research Fund Denmark under the Sapere Aude research career programme.

Maximilian Probst Gutenberg: Supported by Rasmus Kyng's Start-up Grant at ETH. Work done while at BARC, supported by Thorup's Investigator Grant No. 16582.

¹ We use \tilde{O} -notation to hide logarithmic factors.



Christian Wulff-Nilsen: The author is supported by the Starting Grant 7027-00050B from the Independent Research Fund Denmark under the Sapere Aude research career programme.

Acknowledgements We thank anonymous reviewers for their comments and remarks that helped improve the presentation of the paper.

1 Introduction

Shortest paths is a classical algorithmic problem dating back to the 1950s. The two main variants are the all-pairs shortest paths (APSP) problem and the single-source shortest paths (SSSP) problem, both of which have been extensively studied in various models, including the partially and fully-dynamic setting.

A dynamic graph algorithm is an algorithm that maintains information about a graph that is subject to updates such as insertions and deletions of edges or vertices. Such a graph can model real-world networks that change over time, such as road networks where traffic changes and roads are blocked from time to time. We say that a dynamic graph problem is *decremental* if it only allows deletions, *incremental* if it only allows insertions and *fully-dynamic* if it allows both. Incremental and decremental graphs are referred to as being *partially-dynamic*. A dynamic graph algorithm aims to efficiently process a sequence of online updates interspersed with queries about some property of the underlying dynamic graph.

1.1 Problem Definition

In this paper, we consider the *decremental all-pairs shortest-paths* problem where the goal is to efficiently maintain shortest path distances between all pairs of vertices in a decremental directed graph $G = (V, E)$. We shall restrict our attention to the case where G is unweighted. Letting m denote the initial number of edges and $n = |V|$, we want a data-structure which for any $u, v \in V$ supports the following operations:

- $\text{DIST}(u, v)$: reports the distance $d_G(u, v)$ from u to v in the current version of G ,
- $\text{DELETE}(u, v)$: deletes an edge (u, v) from E .

We furthermore consider the problem also in its relaxed version where we only aim to maintain *approximate* distance estimates which can then be queried. We denote by $\tilde{d}_G(u, v)$ a distance estimate for the distance from u to v and we say that an APSP algorithm has an *approximation ratio* (or *stretch*) of $t > 1$ if for any $u, v \in V$, we have that $d_G(u, v) \leq \tilde{d}_G(u, v) \leq t \cdot d_G(u, v)$. This paper will be concerned with both the exact and the $(1 + \epsilon)$ -approximate version of the problem.

Another focus of this article is the *adversarial model*; the adversarial model defines the model under which the sequence of updates and queries are assumed to be made by an *adversary*. We say that a performance guarantee of an algorithm works against an *oblivious* adversary if the adversary must define the sequence of updates before the algorithm starts for the guarantee to hold. Thus the sequence of updates is independent of any random bits used by the algorithm. This is opposed to algorithms that work against an *adaptive* adversary, where the adversary is allowed to create the update sequence “on the go”, e.g. based on answers to previous queries made to the data structure. Depending on the data structure, these choices may not be independent of the random choices made, which may result in the data structure performing poorly. One key advantage of a data structure that works against an adaptive adversary is that it can be used inside an algorithm as a black box, regardless of whether that algorithm adapts its updates to answers to queries. We point out that *deterministic* data structures always work against an adaptive adversary.

The performance of a *partially-dynamic* algorithm is usually measured in terms of the *total update time*. That is, the accumulated time it takes to process all updates (edge deletions). The *query time*, on the other hand, is the time to answer a single distance query. A natural goal is to minimize the total update time while keeping the stretch and query time small. Since all the structures presented in this paper explicitly maintain a distance matrix, the query time is constant.

1.2 Prior Work

The naive approach to dynamic APSP is to recompute the shortest path distances after each update using the best static algorithm. The query time is then constant and the time for a single update is $\tilde{O}(mn)$ for APSP and $\tilde{O}(m)$ for SSSP. At the other end of the spectrum one could achieve optimal update time by simply updating the input graph and only running an SSSP algorithm whenever a query is processed. Running a static algorithm each time, however, fails to reuse any information between updates whatsoever and gives a high query time, motivating more efficient dynamic approaches that do this.

In 1981, Even and Shiloach [6] gave a deterministic data-structure for maintaining a shortest path tree to given depth d in an undirected, unweighted decremental graph in total time $O(md)$. Henzinger and King [7] and King [10] later adapted this to directed graphs with integer weights. Running their structure for each vertex solves the decremental all-pairs shortest paths problem in $O(mn^2W)$ time, where edge weights are integers in $[1, W]$.

Henzinger and King were the first to improve upon this bound, giving an algorithm with total update time $\tilde{O}(mn^{2.5}\sqrt{W})$ [10] which is an improvement for $W = \omega(n)$. Demetrescu and Italiano [4] improved this data structure slightly and showed that the restriction to integral edge weights can be removed. Finally, the same authors [3] presented a data structure with total update time $\tilde{O}(mn^2)$ which is the state of the art for any data structure against an adaptive adversary up to today. In fact, their algorithm can be extended to a fully-dynamic algorithm with $\tilde{O}(n^2)$ amortized update time and which can handle vertex updates². We also point out that this data structure was later simplified and generalized by Thorup [11].

Around the same time Baswana, Hariharan, and Sen [1] gave an *oblivious* Monte-Carlo construction with total update time $\tilde{O}(n^3)$ for *unweighted* graphs. Further, they showed that their data structure could be adapted to give an $(1 + \epsilon)$ -approximate APSP algorithm for *weighted* graphs with total update time of $\tilde{O}(\sqrt{mn^2}/\epsilon)$. In the exact setting, the oblivious adversary assumption is only required when paths are to be reported rather than just shortest path distances which are unique. Finally, Bernstein presented a $(1 + \epsilon)$ -approximate algorithm with total running time $\tilde{O}(mn \log(W)/\epsilon)$ by using a clever approach of shortcutting paths [2]. Whilst his algorithm achieves near-optimal running time, again, the algorithm has to assume an oblivious adversary.

More recently, Karczmarz and Łącki [9] gave a deterministic $(1 + \epsilon)$ -approximate APSP algorithm for decremental graphs that runs in total time $\tilde{O}(n^3 \log(W)/\epsilon)$. They also presented the first non-trivial algorithm for incremental graphs [8] achieving total update time $\tilde{O}(mn^{4/3} \log(W)/\epsilon)$.

We refer the reader to the full version of the paper [5] for a more comprehensive treatment of related work which also includes algorithms for *undirected* graphs and algorithms with larger stretch.

² In this case, vertex updates refers to insertions or deletions of vertices with up to $n - 1$ incident edges.

1.3 Our Contributions

In this paper, we present three new data structures for the all-pairs shortest paths problem. Our first theorem gives a *deterministic* data structure for the exact variant of the problem with near-optimal $\tilde{O}(n^3)$ total update time. It also matches the best *randomized* algorithm by Baswana et al. [1] and constitutes a significant improvement over the previous best deterministic bound of $\tilde{O}(mn^2)$ which is obtained by running an ES-tree [6] from every source or by the data structure Italiano et al. [3] (that also works in weighted graphs) and improves over all but the sparsest graph densities.

Our exact data structure is near-optimal as there is an $\Omega(n^3)$ lower bound on the total update time of any decremental data structure that explicitly maintains the distance matrix. The lower bound follows by considering an initial undirected, unweighted graph consisting of a simple path $v_0, v_1, v_2, \dots, v_{n-1}$ plus additional edges (v_i, v_{i+2}) for each even $i \leq n-3$. Deleting these additional edges in any order creates $\Omega(n^3)$ distance matrix changes in total.

► **Theorem 1.** *Let G be an unweighted directed graph with n vertices and initially m edges. Then there exists a deterministic data structure which maintains all-pairs shortest path distances in G undergoing an online sequence of edge deletions using a total time of $O(n^3 \log^3 n)$. The $n \times n$ distance matrix is explicitly maintained so that at any point, a shortest path distance query can be answered in constant time. The data structure can report a shortest path between any query pair in time proportional to the length of the path.*

Our second result is concerned with maintaining $(1 + \epsilon)$ -approximate all-pairs shortest path distances. This constitutes the first deterministic data structure that solves the problem in subcubic time with small approximation error (except for graphs that are not extremely dense).

► **Theorem 2.** *Let G be an unweighted directed graph with n vertices and initially m edges. Then given $\epsilon > 0$, there exists a deterministic data structure that maintains all-pairs $(1 + \epsilon)$ -approximate shortest path distances in G undergoing an online sequence of edge deletions using a total time of $O(\sqrt{mn} \log^2(n)/\epsilon)$. At any point, a $(1 + \epsilon)$ -approximate shortest path distance query can be answered in constant time and a $(1 + \epsilon)$ -approximate shortest path between the query pair can be reported in time proportional to the length of the path.*

Our third result gives a data structure achieving a better time bound. While we use randomization to achieve the improved time bound, our algorithm again works against an adaptive adversary.

► **Theorem 3.** *Let G be an unweighted directed graph with n vertices and initially m edges. Then given any $\epsilon > 0$, there exists a Las Vegas data structure that maintains all-pairs $(1 + \epsilon)$ -approximate shortest path distances in G under an online sequence of edge deletions using a total expected time of $\tilde{O}(m^{2/3}n^{5/3}/\epsilon + n^{8/3}/(m^{1/3}\epsilon^2))$. This bound holds w.h.p. and the data structure works against an adaptive adversary. At any point, a $(1 + \epsilon)$ -approximate shortest path distance query can be answered in constant time.*

We summarize our results as well as previous state-of-the-art results in Table 1.

1.4 Overview

Our overall approach for the deterministic data structures is similar to that of Baswana et al. [1] but with a key difference that allows us to avoid using a randomized hitting set and instead rely on deterministic separators. The idea of the construction by Baswana et

■ **Table 1** Our results and previous state-of-the-art results for decremental APSP.

Time	Approximation	Adversary/ Deterministic	Reference
$O(mn^2)$	exact	deterministic	[6, 3]
$\tilde{O}(n^3)$	exact	deterministic	New Result
$\tilde{O}(n^3)$	exact	adaptive	[1]
$\tilde{O}(\sqrt{mn^2}/\epsilon)$	$(1 + \epsilon)$	deterministic	New Result
$\tilde{O}(m^{2/3}n^{5/3}/\epsilon + n^{8/3}/(m^{1/3}\epsilon^2))$	$(1 + \epsilon)$	adaptive	New Result
$\tilde{O}(\sqrt{mn^2}/\epsilon)$	$(1 + \epsilon)$	oblivious	[1]
$\tilde{O}(nm)$	$(1 + \epsilon)$	oblivious	[2]

al. relies on a well-known result which says that if we sample a subset H_i^ρ of the vertices of size $\tilde{O}(n/\rho^i)$ (where ρ is some constant strictly larger than 1), each with uniform probability, then, w.h.p. we “hit” each shortest-path of length $[\rho^i, \rho^{i+1})$ between any pair of vertices in any version of the graph G .

Phrased differently, given vertices $u, v \in V$, we have that if a shortest path from u to v is of length $\ell \in [\rho^i, \rho^{i+1})$, then there is some vertex $w \in H_i^\rho$, such that the concatenation of a shortest path from u to w and a shortest path from w to v is of length ℓ . For each such w , we say w is a *witness* for the tuple (u, v) for distance ℓ .

Now for each $u, v \in V$, if the initial distance from u, v was $\ell \in [\rho^i, \rho^{i+1})$, we can check H_i^ρ to find a witness w . If the length of the path from u to w to v is increased, we can continue our scanning of H_i^ρ to see whether another witness exists. If there is no witness $w \in H_i^\rho$ left at some stage, we know that there is no path of length ℓ left in G w.h.p. and increase our guess by setting $\ell \mapsto \ell + 1$.

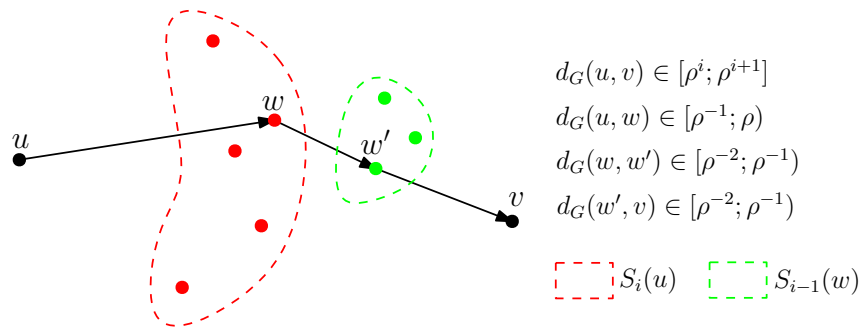
Sampling initially a hitting set H_i^ρ for every $i \in [0, \log_\rho n]$, we can find the “right” hitting set for each distance ℓ . Observe now that for each tuple $(u, v) \in V^2$, we have to scan a hitting set of size $\tilde{O}(n/\rho^i)$ for $\rho^{i+1} - \rho^i \sim \rho^{i+1}$ levels before the hitting set index i is increased which only occurs $O(\log n)$ times, thus we only spend time $\tilde{O}(n)$ for each vertex tuple (u, v) . Thus, the total running time of the searches for witnesses can be bound by $\tilde{O}(n^3)$.

The Deterministic Exact Data Structure. Our construction is similar in the sense that we maintain witnesses for each distance scale $[\rho^i, \rho^{i+1})$ for every $i \in [0, \log_\rho n]$ such that each distance ℓ is in one such distance scale. The key difference is that instead of using a randomized *global* hitting set H_i^ρ for a distance scale $[\rho^i, \rho^{i+1})$, our construction relies on deterministically maintaining a small local vertex separator $S_i(u)$ for every vertex $u \in V$ of size $\tilde{O}(n/\rho^i)$ separating all shortest paths starting in u with a distance in $[\rho^i, \rho^{i+1})$.

More precisely, for each distance scale $[\rho^i, \rho^{i+1})$ and vertex $u \in V$, we maintain a separator $S_i(u)$ that satisfies the invariant that every shortest path from u to a vertex v at distance at least ρ^i is intersected by a vertex in $S_i(u)$. If this invariant is violated after an adversarial update, then we find such a vertex v and need to add additional vertices to $S_i(u)$ during the time step. The challenge is to take these additional separator vertices such that the total size of $S_i(u)$ is not increased beyond $\tilde{O}(n/\rho^i)$. The separator procedure makes use of sparse layers of BFS trees and here is where we rely on the assumption that the graph is unweighted. We defer the details of the separator procedure to a later section and continue our discussion of the APSP data structure.

Since we need to detect whether vertices have distance less than ρ^i from u or not in G , we further have to use a bottom-up approach to compute distances after an edge deletion, i.e. we start with the smallest possible distance range and update all distances in this range and then update larger distances using the information already computed. This issue did not arise in Baswana et al. [1] but can be handled by a careful approach. The distances computed for one distance scale include all distances to and from witnesses for the next larger distance scale.

It is now easy to see that the scanning for witnesses can be implemented in the same time as in the analysis sketched above by scanning the list of local separator vertices which serve as witnesses instead of the hitting set. Further, we can maintain local vertex separators using careful arguments in total time $\tilde{O}(mn)$ giving our result in Theorem 1.



■ **Figure 1** Illustration of separators and path “hierarchy”. Here $u \rightsquigarrow v$ goes through a witness w , and $w \rightsquigarrow v$ goes through w' . If the length of the path $w' \rightsquigarrow v$ is increased by Δ , the distance estimates of all 2-hop-paths that use $w' \rightsquigarrow v$ as a sub-path are increased by that amount. In this case, the estimate for $w \rightsquigarrow w' \rightsquigarrow v$ is increased and is propagated to the next level where subsequently the estimate for $u \rightsquigarrow w \rightsquigarrow v$ is possibly increased.

The Deterministic Approximate Data Structure. In order to improve the running time for sparse graphs, we can further focus on only considering distances that are roughly at a $(1 + \epsilon)$ -multiplicative factor from each other. More concretely, instead of increasing the expected distance from ℓ to $\ell + 1$ when we cannot find a witness for some path from u to v for distance ℓ , we can increase the next expected distance level ℓ' to $\sim (1 + \epsilon)\ell$ and consider every vertex w a witness if there is a path $u \rightsquigarrow w \rightsquigarrow v$ of length at most ℓ' . Thus, we handle fewer distances and can thereby reduce the time to maintain distances that are at least d in total time $\tilde{O}(n^3/d + mn)$. Again, a careful approach is necessary to ensure that approximations do not add up over distance scales.

This is no faster than the data structure for exact distances when d is small so in order to get Theorem 2, we use the $O(mnd)$ data structure of Even and Shiloach [6] to maintain distances up to d . Picking d such that $mnd = n^3/d$ gives the result of Theorem 2 (the term $\tilde{O}(mn)$ vanishes since it is subsumed by the two other terms, also we assumed $\epsilon > 0$ to be a constant to simplify the presentation).

Maintaining Separators. We now sketch how to deterministically maintain a “small” local separator for a vertex $s \in V$ with some useful invariants.

Let S be the local separator for s . The first invariant that will be useful is that any vertex $t \in V$ that is reachable from s in $G \setminus S$, is “close” to s or roughly within distance d . As edges are deleted from G , the distances from s to such vertices t may increase. If a vertex

t moves too far away from s , the invariant is re-established by growing BFS trees in parallel, one layer at a time, from s in $G \setminus S$ and from t in the graph obtained from $G \setminus S$ by reversing the orientations of all edges. The search halts when a layer (corresponding to the leaves of the BFS tree at the current iteration) that is “thin” is found, and its vertices are added to S ; vertices that are on the opposite side of the separator than s are cut off as they must all be too far away from s . Here, “thin” refers to a BFS layer such that the number of vertices added to the separator is only a factor $\tilde{O}(1/d)$ times the number of vertices cut off. It is well known that such a layer exists (cfr. Lemma 5 for the details). Summing up, it follows that $|S| = \tilde{O}(n/d)$ at all times. By marking vertices as they are searched (according to the side of the BFS layer on which they are found), the vertices that are “cut off” from s by the augmented separator will never be searched again, and the cost of searching the edges of either side of the search can be charged to sum of the degree of these vertices, for a total update time of $O(m)$.

For our randomized data structure, we need an additional property that essentially allows us to take a snapshot of the current separator and use it in later updates rather than having to repeatedly update the separator. This will be key to getting an improved randomized time bound. Details can be found in Lemma 6 which states our separator result.

The Randomized Approximate Data Structure. The randomized approximate data structure of Theorem 3 follows the same overall approach but is technically more involved. Instead of keeping track of all 2-hop paths $u \rightsquigarrow s \rightsquigarrow v^3$ for every $s \in S_i(u)$, the randomized data structure samples a subset of these by picking each vertex of $S_i(u)$ independently with some probability p . It only keeps track of approximate shortest path distances going through this subset rather than the full set $S_i(u)$. This will speed up the above since the subset of the separator we need to scan is smaller by a factor p . However, this approach fails once no short 2-hop path intersects the sampled subset. At this point, w.h.p. there should only be short 2-hop paths through $O(\log n/p)$ vertices of $S_i(u)$ so also in this case, the subset can be kept small. However, scanning linearly through $S_i(u)$ to find this small subset will take $\tilde{O}(n/d)$ time and happen over all pairs (u, v) .

Our solution is roughly the following. Suppose no sampled vertex certifies an approximate short path from u to v . Then v scans linearly through $S_i(u)$ to find the small size $O(\log n/p)$ subset $S'_i(u)$. Consider the set W of vertices w such that $d_G(w, v)$ is small compared to d , i.e., $d_G(w, v) \leq \epsilon d$ for some small constant $\epsilon > 0$. Then we show that the small subset $S'_i(u)$ found for v can also be used for each vertex $w \in W$. The intuition is that for any vertex $s \in S_i(u) \setminus S'_i(u)$, the approximate shortest path distance from u to w through s must be large since otherwise we get a short path $u \rightsquigarrow s \rightsquigarrow w \rightsquigarrow v$ from u to v through s , contradicting that $s \notin S'_i(u)$.

It follows that if $|W|$ is large, the $\tilde{O}(n/d)$ cost of scanning $S_i(u)$ can be distributed among a large number of vertices of W . Dealing with the case where $|W|$ is small is more technical so we omit it here.

The way we deal with an adaptive adversary is roughly as follows. Consider a deterministic data structure that behaves like the randomized data structure above, except that it maintains 2-hop paths $u \rightsquigarrow s \rightsquigarrow v$ for all $S_i(u)$ rather than only through a sampled subset. The slack from the approximation allows us to round up all “short” approximate distances to the same value. Hence, as long as the randomized data structure has short 2-hop paths, it maintains exactly the same approximate distances as the deterministic structure and hence the approximate distances output to the adversary is independent of the random bits used.

³ Note that such a path may have more than one intermediate vertex, but it is useful to think of it as a path of two weighted edges/hops (u, s) and (s, v) since this is what is maintained by the data structure.

2 Definitions and Notation

In the following, let $G = (V, E)$ be a directed unweighted graph. The graph G_{rev} is obtained from G by reversing the orientation of each edge. For any two vertices $u, v \in V$, we denote by $u \rightsquigarrow v$ a shortest path from u to v in G and let $d_G(u, v)$ denote the length of such a path. We extend this notation to sets so that, e.g., $d_G(u, V') = \min\{d_G(u, v) \mid v \in V'\}$ for $V' \subseteq V$.

We define a BFS-layer to mean the set of nodes at some fixed distance from some v in G . An *in-tree* in G is a BFS tree in G_{rev} .

We will need notation to refer to dynamically changing data at specific points in time. Consider a sequence of updates to some object X where each update takes place at a time step $t \in \mathbb{N}$. We denote by $X^{(t)}$ the object just after update t . Here, X could be a graph, a shortest path distance, etc.

For handling small distances, we rely on the data-structure of Even and Shiloach [6], the properties of which we will state in the following lemma:

► **Lemma 4** ([6]). *Given a directed unweighted graph G undergoing a sequence of edge deletions, a source vertex $s \in V$, and $d > 0$, a shortest path tree in G rooted at s can be maintained up to distance d in total time $O(md)$. The structure requires $O(m)$ space and can be constructed in time $O(m + n)$.*

3 Maintaining Separators

Lemma 6 below provides a key tool used in all of our data structures. It gives an efficient data structure that maintains a growing separator set S of small size in a decremental graph G . To prove it, we need the following well-known result.

► **Lemma 5.** *Given a directed unweighted n -vertex graph $G = (V, E)$, given $d_1, d_2 \in \mathbb{N}_0$ with $d_2 - d_1 + 1 \geq \lg n$, and given vertices $u, v \in V$ with $d_G(u, v) \geq d_2$, a BFS tree in G with root u contains a layer $L \subseteq V$ with $d_1 \leq d_G(u, L) \leq d_2$ and $|L| \leq |L_-| \lg n / (d_2 - d_1 + 1)$ where $L_- = \{w \in V \mid d_G(u, w) < d_G(u, L)\}$ is the union of layers closer to u than L .*

Proof. Denote by L_i the i th layer of the BFS tree from u . For each i , let $L_{<i} = \cup_{j < i} L_j$. Let $q = (d_2 - d_1 + 1) / \lg n$. Assume for contradiction that L does not exist. Then for $i = d_1, \dots, d_2$, $|L_i| > |L_{<i}| / q$ so $|L_{<i+1}| = |L_i| + |L_{<i}| > (1 + 1/q)|L_{<i}|$. Since $q \geq 1$, we have $(1 + 1/q)^q \geq 2$ so

$$|L_{<d_2+1}| > (1 + 1/q)^{d_2 - d_1 + 1} |L_{<d_1}| \geq 2^{(d_2 - d_1 + 1)/q} = n,$$

contradicting that there are only n vertices in G . ◀

We now state and prove Lemma 6. It gives an efficient data structure that maintains a growing separator set S of small size in a decremental graph G with the following guarantees. Let s be a fixed vertex and let d be some given threshold distance. Then at every time step, vertices reachable from s in $G \setminus S$ are of distance slightly less than d from s in G . Conversely, for vertices v not reachable from s in $G \setminus S$, we have $d_G(s, v) = \Omega(d)$; furthermore, if $d_G(s, v)$ is larger than d by some small constant factor then any shortest path $s \rightsquigarrow v$ in G can be decomposed into $s \rightsquigarrow w \rightsquigarrow v$ such that $w \in S$, $d_G(s, w) \leq d$, and $d_G(w, v) \leq d$. In fact, the lemma states that w can be chosen in S^{t_0} where t_0 is the first time step in which $d_G(s, v)$ became (slightly) larger than d ; note that this is a stronger statement since S is growing over time.

► **Lemma 6.** *Let $G = (V, E)$ be an n -vertex unweighted digraph undergoing a sequence of edge deletions, let $s \in V$ be a source, and let $d \in \mathbb{N}$ with $d > 33 \lg n$. Let \mathcal{O} be a data structure that maintains for each $v \in V$ a distance estimate $\tilde{d}(s, v) \geq d_G(s, v)$ such that if $d_G(s, v) \leq d$ then $\tilde{d}(s, v) \leq \frac{4}{3}d_G(s, v)$. Whenever an estimate $\tilde{d}(s, v)$ grows to a value of at least $\frac{32}{33}d$, \mathcal{O} outputs v . Then there is a data structure \mathcal{S} with access to \mathcal{O} which maintains a growing set $S \subseteq V$ such that for each $v \in V$,*

1. *if v is reachable from s in $G \setminus S$ then $d_G(s, v) < \frac{32}{33}d$ and otherwise $d_G(s, v) > \frac{2}{3}d$,*
2. *if t_0 is a time step in which $d < d_G^{(t_0)}(s, v) \leq \frac{34}{33}d$ then for every time step $t_1 \geq t_0$ in which $d_G^{(t_1)}(s, v) \leq \frac{34}{33}d$, any shortest s -to- v path P in $G^{(t_1)}$ intersects $S^{(t_0)}$ and for the first such intersection vertex w along P , $d_G^{(t_1)}(s, w) \leq d$, and $d_G^{(t_1)}(w, v) \leq d$.*

At any time, $|S| = O(n \log n/d)$ and \mathcal{S} has total update time $O(m)$, excluding the time spent by \mathcal{O} .

Proof. Let $\epsilon = \frac{1}{33}$. For each $v \in V$, let $\hat{d}(v)$ be obtained from the degree of v in the initial graph G by rounding up to the nearest multiple of $\Delta = \lceil m/n \rceil$. In the description of \mathcal{S} below, processing one edge takes at most one unit of time.

Data structure \mathcal{S} initializes $S = \emptyset$ and unmarks all vertices of V . Whenever \mathcal{O} outputs an unmarked vertex v (marked output vertices are ignored), \mathcal{S} runs a modified BFS from s in $G_S = G \setminus S$ which for each vertex w spends $\hat{d}(w)$ time to process its outgoing edges; this can always be achieved by busy-waiting at w if needed. In parallel, \mathcal{S} runs a similar modified BFS from v in $G'_S = (G \setminus S)_{\text{rev}}$ ⁴. The search from s halts if a layer L_s is found such that $\frac{2}{3}d < d_{G_S}(s, L_s) \leq (\frac{2}{3} + \epsilon)d$ and $|L_s| = O((x \log n)/d)$ (for a suitable hidden constant to be specified) where x is the number of vertices visited by the search, excluding L_s . Similarly, the search from v halts if a layer L_v is found such that $d_{G'_S}(v, L_v) < \epsilon d$ and $|L_v| = O((y \log n)/d)$ (again, for a suitable hidden constant) where y is the number of vertices visited by the search excluding L_v . Let L be the first of the two layers found. \mathcal{S} halts both searches when L is found. Then L is added to S and all vertices visited by the search from v in G'_S are marked. The hidden constants are chosen such that the existence of L follows from Lemma 5; this lemma applies since by assumption, $\epsilon d > \lg n$.

Observe that when \mathcal{O} outputs v , we have $d_G(s, v) \geq (1 - \epsilon)d/(4/3) = (\frac{2}{3} + 2\epsilon)d$ as otherwise we have $d_G(s, v) \leq d$ and hence $\tilde{d}(s, v) \leq \frac{4}{3}d_G(s, v) < (1 - \epsilon)d = \frac{32}{33}d$. This shows the existence of L_s and L_v and that no vertex or edge is visited by both searches. We have $d_{G_S}(s, v) \geq d_G(s, v) \geq (\frac{2}{3} + 2\epsilon)d$ and $d_{G_S}(s, L_s) > \frac{2}{3}d$ and for every $w \in L_v$,

$$\begin{aligned} d_{G_S}(s, w) &\geq d_{G_S}(s, v) - d_{G_S}(w, v) \geq \left(\frac{2}{3} + 2\epsilon\right)d - d_{G'_S}(v, w) = \left(\frac{2}{3} + 2\epsilon\right)d - d_{G'_S}(v, L_v) \\ &> \left(\frac{2}{3} + \epsilon\right)d, \end{aligned}$$

implying that $d_{G_S}(s, L_v) > (\frac{2}{3} + \epsilon)d$. Hence $d_{G_S}(s, L) \geq \min\{d_{G_S}(s, L_s), d_{G_S}(s, L_v)\} > \frac{2}{3}d$.

Showing part 1. Let $v \in V$ and consider any point during the sequence of updates. Assume first that v is reachable from s in G_S . Then \mathcal{O} has not yet output v . For suppose otherwise. If v was unmarked when it was output by \mathcal{O} , the above procedure would separate v from s with S , making v unreachable from s in G_S . Conversely, if v was marked, v would already be unreachable from s in G_S since a vertex is only marked when it is separated from s in G_S . In both cases, we have a contradiction. It follows that \mathcal{O} did not output v so $d_G(s, v) \leq \tilde{d}(s, v) < \frac{32}{33}d$, as desired.

⁴ By “parallel”, we mean that \mathcal{S} alternates between spending one unit of time in one search, then one unit of time in the other search, and so on.

64:10 Decremental APSP in Unweighted Digraphs Versus an Adaptive Adversary

Now, assume that v is not reachable from s in G_S . We may assume that there is a shortest path P from s to v in G since otherwise $d_G(s, v) = \infty > \frac{2}{3}d$. Let w be the first vertex of S along P . It suffices to show that for the prefix P' of P from s to w , $|P'| > \frac{2}{3}d$. At some earlier point in time, the procedure added w to S ; just prior to this, P' was contained in G_S so from the above, $|P'| > \frac{2}{3}d$, as desired.

Showing part 2. Let $t_0 \leq t_1$ satisfy the second part of the lemma. Since $d_G^{(t_0)}(s, v) > d$ by assumption, the first part of the lemma implies that v is not reachable from s in $G_S^{(t_0)}$ and hence v is also not reachable from s in $G_S^{(t_1)}$.

Let P be a shortest path from s to v in $G^{(t_1)}$. From what we have just shown, P must intersect $S^{(t_0)}$. Let w be the first vertex of $S^{(t_0)}$ along P . Then clearly, $d_G^{(t_1)}(s, v) = d_G^{(t_1)}(s, w) + d_G^{(t_1)}(w, v)$. Since the vertex w' preceding w on P is reachable from s in $G_S^{(t_0)}$, the first part of the lemma implies that $d_G^{(t_0)}(s, w) \leq d_G^{(t_0)}(s, w') + 1 < \frac{32}{33}d + 1$ and $d_G^{(t_0)}(s, w) > \frac{2}{3}d$. The latter implies that $d_G^{(t_1)}(w, v) = d_G^{(t_1)}(s, v) - d_G^{(t_1)}(s, w) \leq \frac{34}{33}d - d_G^{(t_0)}(s, w) < \frac{34}{33}d - \frac{2}{3}d < d$, showing one of the two inequalities in the second part of the lemma.

We show the other inequality by contradiction so assume that $d_G^{(t_1)}(s, w) > d$. Then $d_G^{(t_1)}(s, w) \geq d + 1$ so by the above $d_G(s, w)$ would have increased by more than $d + 1 - (\frac{32}{33}d + 1) = \frac{1}{33}d$ from time step t_0 to t_1 . Combining this with $d_G^{(t_1)}(s, v) = d_G^{(t_1)}(s, w) + d_G^{(t_1)}(w, v)$, $d_G^{(t_0)}(w, v) \leq d_G^{(t_1)}(w, v)$, and the triangle inequality, we get

$$d_G^{(t_1)}(s, v) - d_G^{(t_0)}(s, v) \geq d_G^{(t_1)}(s, w) + d_G^{(t_1)}(w, v) - (d_G^{(t_0)}(s, w) + d_G^{(t_0)}(w, v)) > \frac{1}{33}d$$

This contradicts the assumption $d < d_G^{(t_0)}(s, v) \leq d_G^{(t_1)}(s, v) \leq \frac{34}{33}d$. We conclude that $d_G^{(t_1)}(s, w) \leq d$ and $d_G^{(t_1)}(w, v) \leq d$ which shows the second part of the lemma.

Bounding $|S|$ and running time. To bound, $|S|$, consider the two parallel searches from s and from v , respectively, in some update. As argued earlier, there cannot be an edge or vertex visited by both searches. Let X resp. Y be the set of vertices visited by the BFS from s resp. v , excluding L_s resp. L_v and let $x = |X|$ and $y = |Y|$.

Assume first that $L = L_s$. Then all vertices in $Y \cup L_v$ become unreachable in G_S once L has been added to S . For each $w \in V$, $\hat{d}(w)/\Delta \geq 1$. Since each BFS spends $\hat{d}(w)$ time to process edges incident to w and since the two searches run in parallel, we have

$$|L| = O((x \log n)/d) = O\left(\frac{\log n}{d} \sum_{w \in X} \frac{\hat{d}(w)}{\Delta}\right) = O\left(\frac{\log n}{d} \sum_{w \in Y \cup L_v} \frac{\hat{d}(w)}{\Delta}\right)$$

Now, assume that $L = L_v$. Then all vertices of $Y \cup L_v$ become unreachable in G_S once L has been added to S so again,

$$|L| = O((y \log n)/d) = O\left(\frac{\log n}{d} \sum_{w \in Y \cup L_v} \frac{\hat{d}(w)}{\Delta}\right)$$

In both cases, the cost $|L|$ of adding L to S can be paid for by charging each vertex w no longer reachable from s in G_S a cost of $O(\frac{\log n}{d} \hat{d}(w)/\Delta)$. Since a vertex is only charged once during the course of the algorithm, we get that for the final separator S (and hence for each intermediate separator),

$$\begin{aligned}
 |S| &= O\left(\frac{\log n}{d} \sum_{w \in V} \frac{\hat{d}(w)}{\Delta}\right) = O\left(\frac{\log n}{d} \sum_{w \in V} \frac{d(w) + \Delta}{\Delta}\right) = O\left(\frac{\log n(m + n \lceil m/n \rceil)}{d \lceil m/n \rceil}\right) \\
 &= O\left(\frac{n \log n}{d}\right)
 \end{aligned}$$

where the last bound follows since we may assume that all vertices are initially reachable from s in G , implying $m \geq n - 1$ and hence $\lceil m/n \rceil = \Theta(m/n)$. This shows the desired bound on $|S|$.

The running time cost of any two parallel searches can be charged to the total degree of the vertices that get marked since they all become unreachable in G_S (this is the vertex set Y above). Since a marked vertex is never visited again by a BFS search, the total running time of parallel searches over all updates is $O(m)$, as desired. ◀

The lemma is somewhat technical and its full strength is only needed for the randomized data structure. For the deterministic data structures, the second part of the lemma will only be applied to the current time step $t_1 = t_0$ so it can be simplified to:

2. if $d < d_G(s, v) \leq \frac{34}{33}d$ then any shortest s -to- v path P in G intersects S and for the first such intersection vertex w along P , $d_G(s, w) \leq d$, and $d_G(w, v) \leq d$.

4 Deterministic Decremental APSP

In this section, we present our deterministic data structures for the exact resp. $(1 + \epsilon)$ -approximate decremental APSP problem and show Theorems 1 and 2. In the following, let $G = (V, E)$ denote the decremental graph.

4.1 Exact distances

Let $\rho = \frac{34}{33}$ and $D_i = \rho^i$ for $i = 0, \dots, \lfloor \log_\rho n \rfloor$. For each i and each $u \in V$, we give a data structure $\mathcal{D}_i(u)$ which for any query vertex v maintains a value $\tilde{d}_i(u, v) \geq d_G(u, v)$ with equality if $d_G(u, v) \in (D_i, D_{i+1}]$. In each update, these data structures will be updated in order of increasing i .

Handling all-pairs shortest path distances up to at most $33 \lg n$ can be done in $O(mn \lg n)$ time using the data structure of Even and Shiloach so we only consider i such that $D_i > 33 \lg n$. This allows us to apply Lemma 6. Consider such an i and assume that we already have data structures for all values smaller than i .

Data structure $\mathcal{D}_i(u)$ maintains a separator set $S_i(u)$ using an instance $\mathcal{S}_i(u)$ of the data structure of Lemma 6 with $s = u$ and $d = D_i$; inductively, we have an exact data structure for distances smaller than d and this data structure plays the role of \mathcal{O} with $\tilde{d} = d_G$. At the beginning of each update, $\mathcal{S}_i(u)$ updates $S_i(u)$. Then for each v , if \mathcal{O} reports that $\tilde{d}(u, v)$ has increased from a value of at most D_i to a value strictly greater than D_i , $\mathcal{D}_i(u)$ initializes $S_i(u, v)$ to be the current separator set $S_i(u)$; $\mathcal{D}_i(u)$ then initializes a priority queue $Q_i(u, v)$ where elements are all $s \in S_i(u, v)$ with corresponding keys $\tilde{d}_{i-1}(u, s) + \tilde{d}_{i-1}(s, v)$. During updates, whenever $\mathcal{D}_{i-1}(u)$ resp. $\mathcal{D}_{i-1}(s)$ reports that $\tilde{d}_{i-1}(u, s)$ resp. $\tilde{d}_{i-1}(s, v)$ increases, the key value of s in $Q_i(u, v)$ increases by the same amount. Note that after initialization, $S_i(u, v)$ remains fixed and so does $Q_i(u, v)$ (except for key value changes).

For each vertex v , $\mathcal{D}_i(u)$ maintains $\tilde{d}_i(u, v)$ as the min key value in $Q_i(u, v)$ after $Q_i(u, v)$ has been initialized; prior to this, $\tilde{d}_i(u, v) = \infty$. This completes the description of each $\mathcal{D}_i(u)$.

The overall data structure \mathcal{D} maintains a priority queue $Q(u, v)$ for each vertex pair (u, v) with an element for each i of key value $\tilde{d}_i(u, v)$. For i in increasing order, \mathcal{D} updates $\mathcal{D}_i(u)$ for each u . Whenever a data structure $\mathcal{D}_i(u)$ increases a value $\tilde{d}_i(u, v)$, the corresponding key in $Q(u, v)$ is increased accordingly. On a query (u, v) , \mathcal{D} reports the min key value in $Q(u, v)$.

Correctness. We prove that for each i , each time step t_1 , and each vertex pair (u, v) , $\tilde{d}_i(u, v) \geq d_G(u, v)$ with equality if $d_G(u, v) \in (D_i, D_{i+1}]$. The inequality is clear since every estimate corresponds to the length of some path in the current graph. The equality part is shown by induction on i .

The base cases where $D_i < 33 \lg n$ are clear so pick i with $D_i \geq 33 \lg n$ and $d_G^{(t_1)}(u, v) \in (D_i, D_{i+1}]$ and assume that correctness holds for all vertex pairs and time steps for $i - 1$. Let $t_0 \leq t_1$ be the first time step such that $d_G^{(t_0)}(u, v) \in (D_i, D_{i+1}]$. Note that $S_i(u, v) = S_i(u)^{(t_0)}$. The induction hypothesis and the second part of Lemma 6 combined with the observation that no key value in $Q_i(u, v)$ is below $d_G(u, v)$, it follows that the min key value in $Q_i(u, v)$ equals $d_G^{(t_1)}(u, v)$. This shows correctness.

Running time. Consider an $i \in \{0, \dots, \lfloor \log_\rho n \rfloor\}$ with $D_i \geq 33 \lg n$ and a vertex $u \in V$. We will show that maintaining $\mathcal{D}_i(u)$ takes $O(n^2 \log^2 n)$ time using a standard binary heap. Total time over all i and u will thus be $O(n^3 \log^3 n)$. This dominates the $O(n^3 \log^2 n)$ time to maintain priority queues $Q(u, v)$ and the $O(mn \log n)$ time for the data structure of Even and Shiloach for small values of i .

Maintaining $S_i(u)$ takes a total of $O(m)$ time by Lemma 6. The total number of elements in priority queues $Q_i(u, v)$ over all $v \in S_i(u, v)$ is $O(n^2 \log n / D_i)$, again by Lemma 6. The number of increase-key operations for a single priority queue element s of $Q_i(u, v)$ is $O(D_i)$ which takes a total of $O(D_i \log n)$ time. Over all elements of priority queues $Q_i(u, v)$, this is $O(n^2 \log^2 n)$.

Reporting paths. It is easy to extend our data structure to efficiently answer queries for shortest paths (rather than only shortest path distances) between any vertex pair (u, v) . Associated with the min element of $Q(u, v)$ is a vertex s such that for the associated index i , $\tilde{d}_i(u, v) = d_G(u, v) = d_G(u, s) + d_G(s, v)$, $\tilde{d}_{i-1}(u, s) = d_G(u, s)$, and $\tilde{d}_{i-1}(s, v) = d_G(s, v)$. Hence, by recursively querying for pairs (u, s) and (s, v) , a shortest u -to- v path in G is reported in time proportional to its length.

4.2 Approximate distances

Let $\epsilon > 0$ be given. We now present our deterministic data structure for the $(1+\epsilon)$ -approximate variant of the problem.

The data structure is quite similar to the one for the exact variant so we only describe the changes needed. For $i > 0$ and $u \in V$, we describe data structure $\mathcal{D}_i(u)$ and assume that we have data structures for values less than i . As before, we only consider i with $D_i \geq 33 \lg n$.

Let $\epsilon' > 0$ be a value depending on ϵ such that $(1 + \epsilon')^c = \rho$ for some $c \in \mathbb{N}$; we will specify ϵ' later. For $j = 0, \dots, c = \log_{1+\epsilon'} \rho$, let $d_{i,j} = D_i(1 + \epsilon')^j$. This partitions each interval $(D_i, D_{i+1}]$ into c sub-intervals $(D_i(1 + \epsilon')^j, D_i(1 + \epsilon')^{j+1}]$ for $j = 0, \dots, c - 1$.

$\mathcal{D}_i(u)$ maintains $S_i(u)$ as in the exact version. For each $v \in V$, $\mathcal{D}_i(u)$ maintains an initially empty set $S_i(u, v)$. Once $\mathcal{D}_{i-1}(u)$ reports that $\tilde{d}_{i-1}(u, v)$ increased from a value of at most $D_i(1 + \epsilon')^i$ to a value strictly greater than $D_i(1 + \epsilon')^i$, $\mathcal{D}_i(u)$ sets $S_i(u, v)$ equal to the current set $S_i(u)$.

For each $j = 0, \dots, c - 1$, a data structure $\mathcal{D}_{i,j}(u)$ maintains the following set for each vertex v :

$$Q_{i,j}(u, v) = \{s \in S_i(u, v) \mid \tilde{d}_{i-1}(u, s) + \tilde{d}_{i-1}(s, v) \leq (1 + \epsilon')^i d_{i,j}\}.$$

$Q_{i,j}(u, v)$ is maintained by $\mathcal{D}_{i,j}(u)$ as a queue in which every $s \in Q_{i,j}(u, v)$ has key $\tilde{d}_{i-1}(u, s) + \tilde{d}_{i-1}(s, v)$ and where s is removed from $Q_{i,j}(u, v)$ (or increased to ∞) when this value exceeds $(1 + \epsilon')^i d_{i,j}$.

For each vertex v , define $\tilde{d}_{i,j}(u, v) = (1 + \epsilon')^i d_{i,j}$ if $Q_{i,j}(u, v)$ contains at least one element and otherwise $\tilde{d}_{i,j}(u, v) = \infty$.

Data structure $\mathcal{D}_i(u)$ maintains a min-priority queue $Q_i(u, v)$ for each vertex v with an element of key value $\tilde{d}_{i,j}(u, v)$ for each j . On query v , it outputs $\tilde{d}_i(u, v) = \min\{k, \tilde{d}_{i-1}(u, v)\}$ where k is the min-key value of this queue, i.e., $\tilde{d}_i(u, v) = \min\{\tilde{d}_{i-1}(u, v), \min_j \tilde{d}_{i,j}(u, v)\}$.

The overall data structure \mathcal{D} works in the same manner as for the exact data structure.

Correctness. Consider any point during the sequence of edge deletions. We will show that for suitable choice of ϵ' , the estimate $\tilde{d}(u, v)$ that \mathcal{D} outputs satisfies $d_G(u, v) \leq \tilde{d}(u, v) \leq (1 + \epsilon)d_G(u, v)$ for every vertex pair (u, v) .

We first show that $d_G(u, v) \leq \tilde{d}(u, v)$. It suffices to prove by induction on $i \geq 0$ that $d_G(u, v) \leq \tilde{d}_i(u, v)$. The proof holds for small i such that $D_i < 33 \lg n$ since then we use the data structure of Even and Shiloach, implying $\tilde{d}_i(u, v) = d_G(u, v)$. Now, consider an i such that $D_i \geq 33 \lg n$ and assume that the claim holds for smaller values than i . We have $\tilde{d}_i(u, v) = \min\{\tilde{d}_{i-1}(u, v), \min_j \tilde{d}_{i,j}(u, v)\}$ and $\tilde{d}_{i,j}(u, v) \geq (1 + \epsilon')^i d_{i,j} \geq \tilde{d}_{i-1}(u, s) + \tilde{d}_{i-1}(s, v)$ for each $s \in Q_{i,j}(u, v)$. Additionally, if $Q_{i,j}(v) = \emptyset$ then $\tilde{d}_{i,j}(u, v) = \infty$. The induction hypothesis now implies $\tilde{d}_i(u, v) \geq d_G(u, v)$, showing the induction step. Thus, $d_G(u, v) \leq \tilde{d}(u, v)$.

To show that $\tilde{d}(u, v) \leq (1 + \epsilon)d_G(u, v)$, we prove by induction on $i \geq 0$ that during all updates and for all vertex pairs (u, v) , if $d_G(u, v) \in (0, D_{i+1}]$ then $\tilde{d}_i(u, v) \leq (1 + \epsilon')^i d_G(u, v)$. If we can show this then picking $\epsilon' \leq \ln(1 + \epsilon)/(\lceil \log_\rho n \rceil)$ gives $\tilde{d}(u, v) \leq (1 + \epsilon')^{\lceil \log_\rho n \rceil} d_G(u, v) \leq e^{\epsilon' \lceil \log_\rho n \rceil} d_G(u, v) \leq (1 + \epsilon)d_G(u, v)$ for every vertex pair (u, v) .

We only need to consider i with $D_i \geq 33 \lg n$ since otherwise, we use the data structure of Even and Shiloach. Assume inductively that the claim holds for values less than i .

Let t_1 be the current time step and consider a vertex pair (u, v) with $d_G^{(t_1)}(u, v) \in (0, D_{i+1}]$. By the induction hypothesis, we may assume that $d_G^{(t_1)}(u, v) \in (D_i, D_{i+1}]$. We may further assume that $\tilde{d}_{i-1}^{(t_1)}(u, v) > D_i(1 + \epsilon')^i$ since otherwise,

$$\tilde{d}_i^{(t_1)}(u, v) \leq \tilde{d}_{i-1}^{(t_1)}(u, v) \leq D_i(1 + \epsilon')^i < (1 + \epsilon')^i d_G^{(t_1)}(u, v).$$

Let $t_0 \leq t_1$ be the first time step where $\tilde{d}_{i-1}^{(t_0)}(u, v) > D_i(1 + \epsilon')^i$. We must have $d_G^{(t_0)}(u, v) > D_i$ since otherwise, the induction hypothesis would imply $\tilde{d}_{i-1}^{(t_0)}(u, v) \leq d_G^{(t_0)}(u, v)(1 + \epsilon')^{i-1} \leq D_i(1 + \epsilon')^{i-1}$, contradicting the choice of t_0 . Since also $d_G^{(t_0)}(u, v) \leq d_G^{(t_1)}(u, v) \leq D_{i+1}$, Lemma 6 implies that there is a vertex $s \in S_i^{(t_0)}(u) = S_i^{(t_0)}(u, v) = S_i^{(t_1)}(u, v)$ such that $d_G^{(t_1)}(u, v) = d_G^{(t_1)}(u, s) + d_G^{(t_1)}(s, v)$, $d_G^{(t_1)}(u, s) \leq D_i$, and $d_G^{(t_1)}(s, v) \leq D_i$.

Pick j such that $d_G^{(t_1)}(u, v) \in (d_{i,j}, d_{i,j+1}]$. By the induction hypothesis,

$$\tilde{d}_{i-1}^{(t_1)}(u, s) + \tilde{d}_{i-1}^{(t_1)}(s, v) \leq (1 + \epsilon')^{i-1} d_G^{(t_1)}(u, v) \leq (1 + \epsilon')^{i-1} d_{i,j+1} = (1 + \epsilon')^i d_{i,j}.$$

Hence, $Q_{i,j}(u, v)$ is non-empty at time step t_1 so $\tilde{d}_i^{(t_1)}(u, v) \leq \tilde{d}_{i,j}^{(t_1)}(u, v) = (1 + \epsilon')^i d_{i,j} \leq (1 + \epsilon')^i d_G^{(t_1)}(u, v)$. This shows the induction step.

Running time. The analysis is similar to the one for exact distances. Pick an $i \in \{0, \dots, \lfloor \log_\rho n \rfloor\}$ with $D_i \geq 33 \lg n$. The total time to maintain $S_i(u)$ over all u is $O(mn)$.

Observe that each approximate distance $\tilde{d}_{i-1}(u_1, u_2)$ is of the form $(1 + \epsilon')^{i'} d_{i',j}$ for $i' \leq i - 1$. Since each element s in a queue $Q_{i,j}(u, v)$ has key value $\tilde{d}_{i-1}(u, s) + \tilde{d}_{i-1}(s, v)$, it follows that the number of increase-key operations applied to s in $Q_{i,j}(u, v)$ is $O(\log_{1+\epsilon'} D_i) = O(\log D_i / \epsilon') = O(\log n / \epsilon')$. For our purpose, a simplified queue $Q_{i,j}(u, v)$ suffices which keeps a counter of the number of elements of key value at most $(1 + \epsilon')^i d_{i,j}$; this follows since the min key value is at most $(1 + \epsilon')^i d_{i,j}$ if and only if the counter is strictly greater than 0. Every queue operation for $Q_{i,j}(u, v)$ can then be supported in $O(1)$ time. The number of elements in $Q_{i,j}(u, v)$ over all u, v , and j is $O(cn^3 \log n / D_i) = O(n^3 \log n / (D_i \epsilon'))$ by Lemma 6. This gives a total time bound of $O(mn + n^3 \log^2 n / (D_i (\epsilon')^2))$. This dominates the time spent on maintaining priority queues $Q_i(u, v)$.

Recall from above that $\epsilon' \leq \ln(1 + \epsilon) / (\lfloor \log_\rho n \rfloor)$. The only additional constraint on ϵ' is that $(1 + \epsilon')^c = \rho$ for some $c \in \mathbb{N}$. This can be achieved with $\epsilon' = \Theta(\ln(1 + \epsilon) / (\lfloor \log_\rho n \rfloor))$. Hence, we get a time bound of $O(mn + n^3 \log^4 n / (D_i \epsilon^2))$.

Note that this bound is no better than the exact data structure for small D_i . We thus consider a hybrid data structure that only applies our data structure when D_i is above some distance threshold d and otherwise applies the data structure of Even and Shiloach which takes a total of $O(mnd)$ time. Summing over all $D_i > d$ and applying a geometric sums argument, the total time for our hybrid data structure is

$$O(mnd + \sum_{i: D_i > d} n^3 \log^4 n / (D_i \epsilon^2)) = O(mnd + n^3 \log^4 n / (d \epsilon^2))$$

Setting $d = n \log^2 n / (\epsilon \sqrt{m})$ gives Theorem 2. Showing the bound for reporting approximate shortest paths in the theorem is done in the same way as in Section 4.1.

5 Randomized Decremental APSP

In this section, we provide a high-level overview of the randomized $(1 + \epsilon)$ -approximate data structure and analysis to achieve the result presented in Theorem 3. Building on this overview we will then prove the theorem.

Let us start by focusing on maintaining approximate distances close to the value D_i from a single vertex u and for now we assume an oblivious adversary.

Maintaining a sampled separator subset. Instead of maintaining each separator $S_{i,j}(u, v)$ (with associated with priority queue $Q_{i,j}(u, v)$) as the full vertex separator $S_i(u)$, we obtain a speed-up by only maintaining a sampled subset of $S_i(u)$. As long as this sampled subset certifies that there is a short two-hop path from u to v , the data structure proceeds as in the previous section. When this is no longer the case, there might still be a short two-hop path from u to v through a non-sampled vertex s in the full separator set $S_i(u)$. However, since there are no more sampled candidates, the expected number of vertices of $S_i(u)$ that provide a short two-hop path is small and we can update $S_{i,j}(u, v)$ to be this small subset. It follows that in expectation, $S_{i,j}(u, v)$ can be kept small at all times, which is needed to give a speed-up.

A speed-up using shallow in-trees. The problem with the data structure sketched above is that the entire set $S_i(u)$ had to be scanned in order to update $S_{i,j}(u, v)$ which means that the data structure will not be faster than our deterministic structure from the previous section.

To deal with this, consider the following modification. The set $S_{i,j}(u, v)$ is updated as before by scanning over the entire set $S_i(u)$. Now, an in-tree $T(v)$ is grown from v of radius at most $\epsilon' D_i$. Each vertex v' in $T(v)$ then inherits the set of v , i.e., $S_{i,j}(u, v')$ is updated to the set $S_{i,j}(u, v)$ and this update is fast since $S_{i,j}(u, v)$ is small in expectation. This works since v is a proxy for v' in the sense that a short two-hop path from u to v' via $S_{i,j}(u, v)$ can be extended with a short suffix from $T(v)$, giving a short two-hop path from u to v via $S_{i,j}(u, v)$ (as $T(v)$ is an in-tree of small radius). Now, the time spent on the single scan of $S_i(u)$ can be distributed among all vertices of $T(v)$ and the number of such vertices must be at least $\epsilon' D_i + 1$ (if not, v would be within distance $\epsilon' D_i$ from u).

Unfortunately, the time analysis for the above procedure breaks down if the in-trees grown during the sequence of updates overlap too much. We now sketch how to deal with this. Mark vertices of each in-tree grown so far. When the BFS procedure grows a new in-tree $T(v)$, this procedure is modified by having it backtrack at previously marked vertices which thus become leaves of $T(v)$; this set of marked leaves will be referred to as L in the detailed description below.

Case 1, dealing with a large in-tree. If the number of unmarked vertices visited in $T(v)$ is greater than $\epsilon' D_i$, the above procedure and analysis can be applied; this is referred to as Case 1 in the detailed description below.

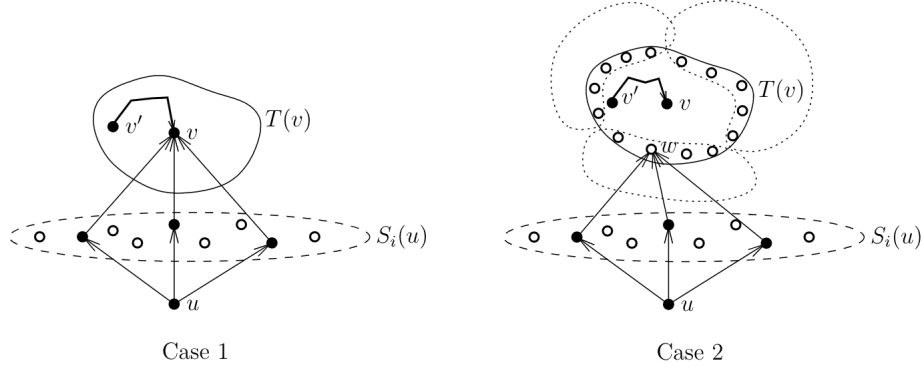
Case 2, dealing with a small in-tree. Otherwise, we are in Case 2; here we recall that $T(v)$ has small radius and observe that the only way to enter $T(v)$ from $G \setminus T(v)$ is through L . Hence, for every vertex s in the union $\cup_{v' \in L} S_{i,j}(u, v')$, there is a good two-hop path from u to v through s . But since we know that there is only a small number of such vertices left (in expectation), this union must be small. Furthermore, the union must contain a good separator for every vertex in $T(v)$ (again because $T(v)$ has small radius and because $T(v)$ must be entered through L) and we thus have an efficient way to update $S_{i,j}(u, w)$ for all $w \in T(v)$.

Handling an adaptive adversary. Above we assumed an oblivious adversary. When the adversary is adaptive, we need to be more careful since the approximate distances reported might reveal information about which vertices have been sampled. To deal with this, we round up every two-hop distance on a given distance scale to the same upper bound value (this will only increase the weight of each two-hop path by a small factor so that the output to a query will still be $(1 + \epsilon)$ -approximate). Hence, the rounded up approximate weight of a two-hop path $u \rightsquigarrow s \rightsquigarrow v$ is the same for every choice of “good” separator vertex s regardless of whether it was sampled or not. It follows that our randomized structure outputs the same distance estimates as a slower deterministic algorithm that maintains the full separator sets. Hence, our randomized algorithm works against an adaptive adversary, as desired.

5.1 The data structure

We now make the the overview formal. First, redefine $\rho = \frac{34 - \frac{1}{2}}{33} = \frac{67}{66}$ and pick ϵ' such that $(1 + \epsilon')^c = \rho$ for some $c \in \mathbb{N}$ and such that $\rho(1 + \epsilon') \leq \frac{34}{33}$. For each u and i such that $D_i \geq 33 \lg n$, a separator $S_i(u)$ is maintained with a data structure $\mathcal{S}_i(u)$ as in Section 4.

We extend the range of index j by 1 so that $j \in \{0, \dots, c + 1\}$. Each structure $\mathcal{D}_{i,j}(u)$ maintains a growing set $M_{i,j}(u)$ of marked vertices; this set is initially empty. In the following, let $U_{i,j}(u) = V \setminus M_{i,j}(u)$ denote the set of unmarked vertices and let $G_{U_{i,j}(u)}$ denote the graph with vertex set V and containing the edges of G having at least one unmarked endpoint.



■ **Figure 2** The two cases in the description of the randomized algorithm. Case 1: black vertices of $S_i(u)$ form the subset of vertices s with $\tilde{d}_{i-1}(u, s) + \tilde{d}_{i-1}(s, v) \leq d_{i,j}(1 + \epsilon')^{2i}$. For each $v' \in V(T(v))$, $Q_{i,j}(u, v')$ is set to be this subset (with key values adjusted). 2-hop paths from u to v through the subset are shown. Case 2: Vertices of L are shown in white inside $T(v)$. Dotted regions are trees touching $T(v)$. For a $w \in L$, black vertices of $S_i(u)$ form the set $Q_{i,j}(u, w)$ and 2-hop paths from u to w through this set are shown. Q is the union of these sets over all $w \in L$. For each $v' \in V(T(v)) \setminus L$, $Q_{i,j}(u, v')$ is a subset of Q .

In each update, $\mathcal{D}_{i,j}(u)$ maintains $S_{i,j}(u, v)$ and $Q_{i,j}(u, v)$ for $v \in V$ in the following way.

For each $v \in V$ and every vertex s added to $S_i(u)$ in the current update, s is added to $S_{i,j}(u, v)$ with some probability p to be fixed later. Note that only vertices v for which s is actually added to $S_{i,j}(u, v)$ need to be processed. In the full version of the paper [5], we employ a different sampling scheme that avoids having to flip a coin for every vertex $v \in V$ in every update.

For vertices v such that $v \in M_{i,j}(u)$ or such that both $v \in U_{i,j}(u)$ and $\tilde{d}_{i-1}(u, v) \leq D_i(1 + \epsilon')^{2i}$, no further processing is done.

Now, assume that $v \in U_{i,j}(u)$ and that $\tilde{d}_{i-1}(u, v) > D_i(1 + \epsilon')^{2i}$. If this inequality did not hold in the previous update, each (sampled) vertex of $S_{i,j}(u, v)$ is added to a new min-queue $Q_{i,j}(u, v)$ with key values as in the previous section. Conversely, if the inequality did hold in the previous update, each new (sampled) vertex added to $S_{i,j}(u, v)$ in the current update is added to $Q_{i,j}(u, v)$.

If the min key value of $Q_{i,j}(u, v)$ is greater than $d_{i,j}(1 + \epsilon')^{2i}$, $\mathcal{D}_{i,j}(u)$ grows an in-tree $T(v)$ from v in $G_{U_{i,j}(u)}$ up to radius $\epsilon' D_i$.

There are now two cases (see Figure 2): $|V(T(v)) \setminus M_{i,j}(u)| > \epsilon' D_i$ and $|V(T(v)) \setminus M_{i,j}(u)| \leq \epsilon' D_i$.

Case 1: If $|V(T(v)) \setminus M_{i,j}(u)| > \epsilon' D_i$ then $\mathcal{D}_{i,j}(u)$ scans once over $S_i(u)$ to find the subset of vertices $s \in S_i(u)$ for which $\tilde{d}_{i-1}(u, s) + \tilde{d}_{i-1}(s, v) \leq d_{i,j}(1 + \epsilon')^{2i}$. For each $v' \in V(T(v))$, $Q_{i,j}(u, v')$ is set to contain exactly this subset of vertices s but with key value $\tilde{d}_{i-1}(u, s) + \tilde{d}_{i-1}(s, v')$.

Case 2: If $|V(T(v)) \setminus M_{i,j}(u)| \leq \epsilon' D_i$ then let $L = V(T(v)) \cap M_{i,j}(u)$ and let $Q = \cup_{v' \in L} Q_{i,j}(u, v')$. For each $v' \in V(T(v)) \setminus L$, $\mathcal{D}_{i,j}(u, v)$ sets $Q_{i,j}(u, v')$ to contain the elements $s \in Q$ with $\tilde{d}_{i-1}(u, s) + \tilde{d}_{i-1}(s, v) \leq d_{i,j}(1 + \epsilon')^{2i}$; their key values are $\tilde{d}_{i-1}(u, s) + \tilde{d}_{i-1}(s, v')$.

In both cases, $\mathcal{D}_{i,j}(u)$ then marks all vertices of $T(v)$, i.e., $M_{i,j}(u) \leftarrow M_{i,j}(u) \cup V(T(v))$.

Approximate distances $\tilde{d}_{i,j}(u, v)$ are maintained by $\mathcal{D}_{i,j}(u)$ in a way similar to that in Section 4.2: $\tilde{d}_{i,j}(u, v) = (1 + \epsilon')^{2i} d_{i,j}$ if the min key value of $Q_{i,j}(u, v)$ is at most $(1 + \epsilon')^{2i} d_{i,j}$ and otherwise $\tilde{d}_{i,j}(u, v) = \infty$.

Data structures $\mathcal{D}_i(u)$ as well as the overall data structure \mathcal{D} work exactly as in Section 4.2.

5.2 Correctness

Consider any point during the sequence of edge deletions. We will show that for suitable choice of ϵ' , we have $d_G(u, v) \leq \tilde{d}(u, v) \leq (1 + \epsilon)d_G(u, v)$.

We do this by proving that during all updates and for all vertex pairs (u, v) , if $d_G(u, v) \in (0, D_{i+1}(1 + \epsilon')]$ then $\tilde{d}_i(u, v) \leq (1 + \epsilon')^{2i} d_G(u, v)$. By picking $\epsilon' = \ln(1 + \epsilon)/(2\lceil \log_\rho n \rceil)$, this will give $d_G(u, v) \leq \tilde{d}_G(u, v) \leq (1 + \epsilon')^{2\lceil \log_\rho n \rceil} d_G(u, v) \leq e^{2\lceil \log_\rho n \rceil \epsilon'} d_G(u, v) \leq (1 + \epsilon)d_G(u, v)$, as desired.

The proof is by induction on i . The claim is clear for i with $D_i < 33 \lg n$ since then we use the data structure of Even and Shiloach. Now, consider an i with $D_i \geq 33 \lg n$ and assume that the claim holds for values less than i . By the induction hypothesis, we only need to consider pairs (u, v) with $d_G(u, v) \in (D_i(1 + \epsilon'), D_{i+1}(1 + \epsilon')]$, i.e., $d_G(u, v) \in (d_{i,j}, d_{i,j+1}]$ with $j > 0$.

We first show the following invariant for marked vertices that holds prior to each update over the entire sequence of updates:

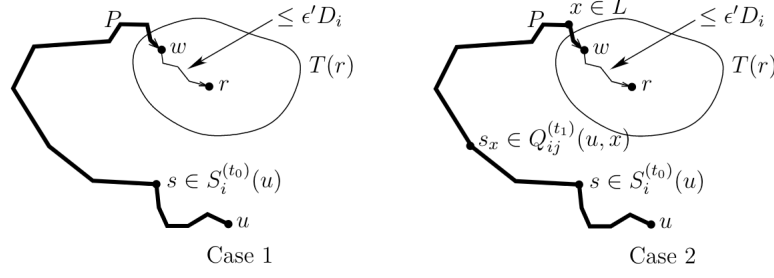
► **Invariant 7.** *At the end of each update, for every $w \in M_{i,j}(u)$ with $d_G(u, w) \in (d_{i,j}, d_{i,j+1}]$, each shortest u -to- w path in G intersects a vertex $s \in Q_{i,j}(u, w)$ such that $d_G(u, s) \leq D_i$ and $d_G(s, w) \leq D_i$.*

Proof. The invariant is shown by induction on the rank of w in the order in which vertices are marked. Note that this is a proof by induction inside a step of the main proof by induction on i ; in addition to the induction hypothesis stated above, we may thus assume that the invariant holds for values less than i . Additionally, for the current value of i , we may assume by induction that the invariant holds for vertices of lower rank than w .

Let t_1 be a time step with $w \in M_{i,j}(u)^{(t_1)}$ and $d_G^{(t_1)}(u, w) \in (d_{i,j}, d_{i,j+1}]$, let $t_0 \leq t_1$ be the time step in which w was marked, and let r be the vertex from which an in-tree $T(r) \ni w$ was grown in time step t_0 . Let P be a shortest u -to- w path in $G^{(t_1)}$.

We must have $\tilde{d}_{i-1}^{(t_0)}(u, r) > D_i(1 + \epsilon')^{2i}$ since otherwise, no processing would be done for r in time step t_0 , contradicting that $T(r)$ is grown in that time step. We also have $d_G^{(t_0)}(u, r) > D_i(1 + \epsilon')$ since otherwise the induction hypothesis would give the contradiction $D_i(1 + \epsilon') \geq d_G^{(t_0)}(u, r) \geq \tilde{d}_{i-1}^{(t_0)}(u, r)/(1 + \epsilon')^{2(i-1)} > D_i(1 + \epsilon')^{2i-2(i-1)} = D_i(1 + \epsilon')^2$.

By the triangle inequality and the fact that $w \in T(r)$ and $T(r)$ has radius at most $\epsilon' D_i$, we get $d_G^{(t_0)}(u, w) \geq d_G^{(t_0)}(u, r) - d_G^{(t_0)}(w, r) > D_i(1 + \epsilon') - \epsilon' D_i = D_i$. Hence, $D_i < d_G^{(t_0)}(u, w) \leq d_G^{(t_1)}(u, w) \leq d_{i,j+1} \leq \frac{34}{33} D_i$ so by Lemma 6, P intersects $S_i^{(t_0)}(u)$ and for the first such intersection vertex s along P , $d_G^{(t_0)}(u, s) \leq d_G^{(t_1)}(u, s) \leq D_i$ and $d_G^{(t_0)}(s, w) \leq d_G^{(t_1)}(s, w) \leq D_i$. We consider the two cases in the description of $\mathcal{D}_{i,j}(u)$ (see Figure 3):



■ **Figure 3** The two cases in the proof of Invariant 7. The path P is marked in bold.

Case 1. It suffices to show that $s \in Q_{i,j}^{(t_1)}(u, w)$. We have $d_G^{(t_0)}(s, r) \leq d_G^{(t_0)}(s, w) + d_G^{(t_0)}(w, r) \leq (1 + \epsilon')D_i$. By the induction hypothesis,

$$\begin{aligned}
 \tilde{d}_{i-1}^{(t_0)}(u, s) + \tilde{d}_{i-1}^{(t_0)}(s, r) &\leq (1 + \epsilon')^{2(i-1)} d_G^{(t_0)}(u, r) \\
 &\leq (1 + \epsilon')^{2i-2} (d_G^{(t_0)}(u, w) + \epsilon' D_i) \\
 &\leq (1 + \epsilon')^{2i-2} (d_G^{(t_1)}(u, w) + \epsilon' D_i) \\
 &\leq (1 + \epsilon')^{2i-2} (d_{i,j+1} + \epsilon' d_{i,j+1}) \\
 &= (1 + \epsilon')^{2i} d_{i,j},
 \end{aligned}$$

so $s \in Q_{i,j}^{(t_0)}(u, w) = Q_{i,j}^{(t_1)}(u, w)$, showing maintenance of the invariant.

Case 2. We first show that P must intersect the set L formed when growing $T(r)$ in time step t_0 . Since we are in Case 2, every leaf of $T(r)$ either belongs to L or has no ingoing edges from vertices not in $T(r)$; otherwise, $T(r)$ would contain more than $\epsilon' D_i$ vertices since it is grown up to radius $\epsilon' D_i$. Hence, the only way that P could not intersect L would be if P were fully contained in $T(r)$. But this is not possible since then $T(r)$ would contain at least $|P| + 1 \geq d_{i,j} + 1 > D_i \geq \epsilon' D_i$ unmarked vertices at the beginning of time step t_0 , contradicting that we are in Case 2.

Thus, P intersects L and we have $w \notin L$ since w was an unmarked vertex of $T(r)$ when growing this tree. Let x be the last vertex of P belonging to L . Since x was marked earlier than w , the induction hypothesis implies that the subpath of P from u to x intersects $Q_{i,j}^{(t_1)}(u, x) = Q_{i,j}^{(t_0)}(u, x)$ in a vertex s_x such that $d_G^{(t_0)}(u, s_x) \leq d_G^{(t_1)}(u, s_x) \leq D_i$ and $d_G^{(t_0)}(s_x, x) \leq d_G^{(t_1)}(s_x, x) \leq D_i$. The latter implies $d_G^{(t_0)}(s_x, r) \leq (1 + \epsilon')D_i$. By the induction hypothesis, $\tilde{d}_{i-1}^{(t_0)}(u, s_x) + \tilde{d}_{i-1}^{(t_0)}(s_x, r) \leq (1 + \epsilon')^{2(i-1)} (d_G^{(t_0)}(u, s_x) + d_G^{(t_0)}(s_x, r)) = (1 + \epsilon')^{2(i-1)} d_G^{(t_0)}(u, r)$ which by the same calculations as in Case 1 is at most $(1 + \epsilon')^{2i} d_{i,j}$. Inspecting the execution of $\mathcal{D}_{i,j}(u)$ in Case 2, it follows that $s_x \in Q_{i,j}^{(t_0)}(u, w) = Q_{i,j}^{(t_1)}(u, w)$. We have $s_x \in Q_{i,j}^{(t_0)}(u, x) \subseteq S_i^{(t_0)}(u)$. Since s is the first vertex of $S_i^{(t_0)}(u)$ along P , P can thus be decomposed into $u \rightsquigarrow s \rightsquigarrow s_x \rightsquigarrow x \rightsquigarrow w$ and we get $d_G^{(t_1)}(u, s_x) \leq D_i$ (as shown above) and $d_G^{(t_1)}(s_x, w) \leq d_G^{(t_1)}(s, w) \leq D_i$. This shows maintenance of the invariant with s_x in place of s . ◀

Now, we continue with our proof by induction on i . Consider any vertex pair (u, v) at the end of an update with $d_G(u, v) \in (d_{i,j}, d_{i,j+1}]$ and $j > 0$. If $v \notin M_{i,j}(u)$ and $\tilde{d}_{i-1}(u, v) \leq (1 + \epsilon')^{2i} D_i$ then $d_G(u, v) \leq \tilde{d}_{i,j}(u, v) \leq (1 + \epsilon')^{2i} D_i < (1 + \epsilon')^{2i} d_G(u, v)$, as desired.

Now assume that $v \notin M_{i,j}(u)$ and $\tilde{d}_{i-1}(u, v) > (1 + \epsilon')^{2i} D_i$. Since v was not marked in the current update, the min key value of $Q_{i,j}(u, v)$ at the end of the update is at most $d_{i,j}(1 + \epsilon')^{2i}$ so $d_G(u, v) \leq \tilde{d}_{i,j}(u, v) \leq (1 + \epsilon')^{2i} d_{i,j} < (1 + \epsilon')^{2i} d_G(u, v)$, as desired.

Finally assume that $v \in M_{i,j}(u)$. By Invariant 7, there is an $s \in Q_{i,j}(u, v)$ such that $d_G(u, v) = d_G(u, s) + d_G(s, v)$, $d_G(u, s) \leq D_i$, and $d_G(s, v) \leq D_i$. By the induction hypothesis, $d_G(u, v) \leq \tilde{d}_i(u, v) \leq \tilde{d}_{i-1}(u, s) + \tilde{d}_{i-1}(s, v) \leq (1 + \epsilon')^{2(i-1)} d_G(u, v)$, as desired. This completes the inductive proof and correctness follows.

5.3 Running time

Maintaining separators $S_i(u)$ over all u and i takes $O(mn \log_p n) = O(mn \log n)$ time by Lemma 6. For the remaining time analysis, we focus on a single data structure $\mathcal{D}_{i,j}(u)$. It is useful in the following to regard this structure as handling an adversarial sequence of updates consisting of changes to approximate distances maintained by structures $\mathcal{D}_{i'}(v)$ for $i' < i$ and $v \in V$. We will give an expected time bound for $\mathcal{D}_{i,j}(u)$ and we shall rely on the following key lemma; the proof can be found in the full version of the paper [5].

► **Lemma 8.** *Let $r \in V$. If at some point in the sequence of updates, $\mathcal{D}_{i,j}(u)$ grows an in-tree from r then at the end of that update, the expected number of vertices $s \in S_i(u)$ satisfying $\tilde{d}_{i-1}(u, s) + \tilde{d}_{i-1}(s, r) \leq D_i(1 + \epsilon')^{2i}$ is $O(\ln n/p)$. This bound holds against an adaptive adversary.*

► **Corollary 9.** *When a vertex v is marked, $E[|Q_{i,j}(u, v)|] = O(\ln n/p)$ and this bound holds against an adaptive adversary.*

Proof. Consider the update in which v is marked and let r be the root of the in-tree $T(r)$ containing v . If $|T(r)| \geq \epsilon' D_i$ then $Q_{i,j}(u, v) = Q_{i,j}(u, r) \subseteq S_i(u)$ and all $s \in Q_{i,j}(u, r)$ satisfy the inequality of Lemma 8. In the case where $|T(r)| < \epsilon' D_i$, let Q be as defined in the description of the data structure. Then vertices $s \in Q \subseteq S_i(u)$ are only added to $Q_{i,j}(u, v)$ if they satisfy the inequality of Lemma 8. The corollary now follows. ◀

Now, we can bound the time spent by $\mathcal{D}_{i,j}(u)$. The total time spent on growing in-trees is $O(m)$ since every edge (w_1, w_2) visited must have $w_1 \notin M_{i,j}(u)$ at the beginning of the BFS search and $w_1 \in M_{i,j}(u)$ immediately afterwards and a vertex can never be unmarked. This also bounds the time spent on marking vertices.

The total expected number of sampled vertices added to $Q_{i,j}(u, v)$ prior to v being marked is at most xp where x is the size of the set $S_i(u)$ after the final update. By Lemma 6, $x = O(n \log n / D_i)$. By Corollary 9, the expected size of $Q_{i,j}(u, v)$ after v is marked is $O(\ln n/p)$. Using the same argument as in the running time analysis of Section 4.2, the number of increase-key operations applied to a single element of $Q_{i,j}(u, v)$ is $O(\log n / \epsilon')$. Hence, the total expected time spent on operations on $Q_{i,j}(u, v)$ is $O((n \log n \cdot p / D_i + \log n / p) \log^3 n / \epsilon)$.

Whenever $\mathcal{D}_{i,j}(u)$ grows an in-tree $T(r)$ with $|V(T(r)) \setminus M_{i,j}(u, v)| > \epsilon' D_i$, scanning $S_i(u)$ takes $O(n \log n / D_i)$ time by Lemma 6. Since all vertices of $V(T(r)) \setminus M_{i,j}(u, v)$ are marked just after $T(r)$ is grown and since vertices are never unmarked, the number of such trees over the course of the updates is at most $n / (\epsilon' D_i)$ so the total time for all these scans is $O(n^2 \log^2 n / (\epsilon D_i^2))$.

Whenever $\mathcal{D}_{i,j}(u)$ grows an in-tree $T(r)$ with $|V(T(r)) \setminus M_{i,j}(u, v)| \leq \epsilon' D_i$, the set $\cup_{x \in L} Q_{i,j}(u, x)$ needs to be computed. Note that for each $x \in L$, $E[|Q_{i,j}(u, x)|] = O(\log n / p)$ by Corollary 9. At least one edge (y, x) ingoing to x belongs to $T(r)$ and this edge is not part of any later grown in-tree since x is marked immediately after $T(r)$ is grown. We charge a cost of $O(\log n / p)$ to (y, x) for computing $Q_{i,j}(u, x)$. Over all $x \in L$, this pays for computing $\cup_{x \in L} Q_{i,j}(u, x)$ and we get a total expected time bound for this part of $O(m \log n / p)$.

Summing the above over all $u, v, i,$ and $j,$ we get a total expected time bound for our data structure of

$$\tilde{O}(mn/\epsilon + \sum_i \sum_j (n^3 \cdot p/(D_i \epsilon) + n^2/(p\epsilon) + n^3/(\epsilon D_i^2) + mn/p)).$$

Since this bound is only fast for sufficiently large $i,$ we pick a distance threshold d and apply the algorithm of Even and Shiloach for distances of at most d and our data structure for distances above $d.$ By a geometric sums argument, our hybrid algorithm has a expected total time bound of

$$\begin{aligned} &\tilde{O}(mnd + mn/\epsilon + n^3 \cdot p/(d\epsilon^2) + n^2/(p\epsilon^2) + n^3/(\epsilon^2 d^2) + mn/(p\epsilon)) \\ &= \tilde{O}(mnd + n^3 \cdot p/(d\epsilon^2) + n^2/(p\epsilon^2) + n^3/(d^2 \epsilon^2) + mn/(p\epsilon)) \end{aligned}$$

Setting the second and fifth terms equal to each other, we get $p = \tilde{\Theta}(\sqrt{m\epsilon d}/n)$ and the time bound simplifies to

$$\tilde{O}(mnd + \sqrt{mn^2}/(\sqrt{d}\epsilon^{3/2}) + n^3/(\sqrt{md}\epsilon^{5/2}) + n^3/(d^2 \epsilon^2)).$$

We balance the first two terms by setting $d = \tilde{\Theta}(n^{2/3}/(m^{1/3}\epsilon))$ and we get a time bound of

$$\tilde{O}(m^{2/3}n^{5/3}/\epsilon + n^{8/3}/(m^{1/3}\epsilon^2)),$$

which shows the time bound of Theorem 3.

References

- 1 Surender Baswana, Ramesh Hariharan, and Sandeep Sen. Improved decremental algorithms for maintaining transitive closure and all-pairs shortest paths. In *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, pages 117–123. ACM, 2002.
- 2 Aaron Bernstein. Maintaining shortest paths under deletions in weighted directed graphs. *SIAM Journal on Computing*, 45(2):548–574, 2016.
- 3 Camil Demetrescu and Giuseppe F Italiano. A new approach to dynamic all pairs shortest paths. *Journal of the ACM (JACM)*, 51(6):968–992, 2004.
- 4 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.
- 5 Jacob Evald, Viktor Fredslund-Hansen, Maximilian Probst Gutenberg, and Christian Wulff-Nilsen. Decremental APSP in directed graphs versus an adaptive adversary. *CoRR*, abs/2010.00937, 2020. [arXiv:2010.00937](https://arxiv.org/abs/2010.00937).
- 6 Shimon Even and Yossi Shiloach. An on-line edge-deletion problem. *Journal of the ACM (JACM)*, 28(1):1–4, 1981.
- 7 M.R. Henzinger and Valerie King. Fully dynamic biconnectivity and transitive closure. In *Proceedings of IEEE 36th Annual Foundations of Computer Science*, pages 664–672. IEEE Comput. Soc. Press, 1995. doi:10.1109/SFCS.1995.492668.
- 8 Adam Karczmarz and Jakub Lacki. Reliable hubs for partially-dynamic all-pairs shortest paths in directed graphs. In *27th Annual European Symposium on Algorithms (ESA 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.
- 9 Adam Karczmarz and Jakub Łącki. Simple label-correcting algorithms for partially dynamic approximate shortest paths in directed graphs. In *Symposium on Simplicity in Algorithms*, pages 106–120. SIAM, 2020.
- 10 Valerie King. Fully dynamic algorithms for maintaining all-pairs shortest paths and transitive closure in digraphs. In *Foundations of Computer Science, 1999. 40th Annual Symposium on*, pages 81–89. IEEE, 1999.
- 11 Mikkel Thorup. Fully-dynamic all-pairs shortest paths: Faster and allowing negative cycles. In *Scandinavian Workshop on Algorithm Theory*, pages 384–396. Springer, 2004.