

Decremental Single-Source Shortest Paths on Undirected Graphs in Near-Linear Total Update Time*

Monika Henzinger[†] Sebastian Krinninger[†] Danupon Nanongkai[‡]

Abstract

In the decremental single-source shortest paths (SSSP) problem we want to maintain the distances between a given source node s and every other node in an n -node m -edge graph G undergoing edge deletions. While its static counterpart can be easily solved in near-linear time, this decremental problem is much more challenging even in the *undirected unweighted* case. In this case, the classic $O(mn)$ total update time of Even and Shiloach [JACM 1981] has been the fastest known algorithm for three decades. At the cost of a $(1 + \epsilon)$ -approximation factor, the running time was recently improved to $O(n^{2+o(1)})$ by Bernstein and Roditty [SODA 2011]. In this paper, we bring the running time down to near-linear: We give a $(1 + \epsilon)$ -approximation algorithm with $O(m^{1+o(1)})$ total update time, thus obtaining *near-linear time*. Moreover, we obtain $O(m^{1+o(1)} \log W)$ time for the weighted case, where the edge weights are integers from 1 to W . The only prior work on weighted graphs in $o(mn \log W)$ time is the $O(mn^{0.9+o(1)})$ -time algorithm by Henzinger, Krinninger, and Nanongkai [STOC 2014, ICALP 2015] which works for the general weighted directed case.

In contrast to the previous results which rely on maintaining a sparse emulator, our algorithm relies on maintaining a so-called *sparse (h, ϵ) -hop set* introduced by Cohen [JACM 2000] in the PRAM literature. An (h, ϵ) -hop set of a graph $G = (V, E)$ is a set F of weighted edges such that the distance between any pair of nodes in G can be $(1 + \epsilon)$ -approximated by their h -hop distance (given by a path containing at most h edges) on $G' = (V, E \cup F)$. Our algorithm can maintain an $(n^{o(1)}, \epsilon)$ -hop set of near-linear size in near-linear time under edge deletions. It is the first of its kind to the best of our knowledge. To maintain approximate distances using this hop set, we extend the monotone Even-Shiloach tree of Henzinger, Krinninger, and Nanongkai [FOCS 2013] and combine it with the bounded-hop SSSP technique of Bernstein [FOCS 2009, STOC 2013] and Mądry [STOC 2010]. These two new tools might be of independent interest.

*A preliminary version of this paper was presented at the 55th IEEE Symposium on Foundations of Computer Science (FOCS 2014).

[†]University of Vienna, Faculty of Computer Science, Austria. Supported by the Austrian Science Fund (FWF): P23499-N23 and the University of Vienna (IK I049-N). The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement no. 340506 and from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 317532.

[‡]KTH Royal Institute of Technology. Work partially done while at ICERM, Brown University, USA, and Nanyang Technological University, Singapore 637371, and while supported in part by the following research grants: Nanyang Technological University grant M58110000, Singapore Ministry of Education (MOE) Academic Research Fund (AcRF) Tier 2 grant MOE2010-T2-2-082, and Singapore MOE AcRF Tier 1 grant MOE2012-T1-001-094.

Contents

1	Introduction	3
2	Preliminaries	5
3	Technical Overview	6
4	From Approximate SSSP to Approximate Balls	15
4.1	Relation to Exact Balls	16
4.2	Properties of Approximate Balls	18
5	From Approximate Balls to Approximate SSSP	21
5.1	Algorithm Description	22
5.2	Running Time Analysis	23
5.3	Definitions of Values for Approximation Guarantee	23
5.4	Analysis of Approximation Guarantee	27
6	Putting Everything Together	34
6.1	Approximate SSSP	35
6.2	Approximate APSP	40
7	Conclusion	42
	References	43

1 Introduction

Dynamic graph algorithms refer to data structures on graphs that support update and query operations. They are classified according to the type of update operations they allow: *decremental* algorithms allow only edge deletions, *incremental* algorithms allow only edge insertions, and *fully dynamic* algorithms allow both insertions and deletions. In this paper, we consider decremental algorithms for the *single-source shortest paths (SSSP)* problem on *undirected* graphs. The *unweighted* case of this problem allows the following operations.

- DELETE(u, v): delete the edge (u, v) from the graph, and
- DISTANCE(x): return the distance $\text{dist}_G(s, x)$ between node s and node x in the current graph G .

The *weighted* case allows an additional operation INCREASE(u, v, Δ) which increases the weight of the edge (u, v) by Δ . We allow positive integer edge weights in the range from 1 to W , for some parameter W . For any $\alpha \geq 1$, we say that an algorithm is an α -*approximation* algorithm if, for any distance query DISTANCE(x), it returns a distance estimate $\delta(s, x)$ such that $\text{dist}_G(s, x) \leq \delta(s, x) \leq \alpha \text{dist}_G(s, x)$. There are two time complexity measures associated with this problem: *query time* denoting the time needed to answer *each* distance query, and *total update time* denoting the time needed to process *all* edge deletions. The running time will be in terms of n , the number of nodes in the graph, and m , the number of edges *before* the first deletion. For the weighted case, we additionally have W , the maximum edge weight. We use \tilde{O} -notation to hide $O(\text{polylog } n)$ terms. In this paper, we focus on algorithms with small ($O(1)$ or $O(\text{polylog } n)$) query time, and the main goal is to minimize the total update time, which will simply be referred to as *time* when the context is clear.

Related Work. The static version of SSSP can be easily solved in $\tilde{O}(m)$ time using, e.g., Dijkstra’s algorithm. Moreover, due to the deep result of Thorup [Tho99], it can even be solved in linear ($O(m)$) time in undirected graphs with positive integer edge weights. This implies that in our setting we can naively solve decremental SSSP in $O(m^2)$ time by running the static algorithm after every deletion. The first non-trivial decremental algorithm is due to Even and Shiloach [ES81] from 1981 and takes $O(mn)$ time in unweighted undirected graphs. This algorithm will be referred to as *ES-tree* throughout this paper. It has many applications such as for decremental strongly-connected components [Rod13] and multicommodity flow problems [Mad10]; yet, the ES-tree has resisted many attempts of improving it for decades. Roditty and Zwick [RZ11] explained this phenomenon by providing evidence that the ES-tree is optimal for maintaining exact distances even on *unweighted undirected* graphs, unless there is a major breakthrough for Boolean matrix multiplication and many other long-standing problems [VWW10]. After the preliminary version of this work appeared, Henzinger et al. [HKN⁺15] showed that $O(mn)$ is essentially the best possible total update time for maintaining exact distances under the assumption that there is no “truly subcubic” algorithm for a problem called online Boolean matrix-vector multiplication. It is thus natural to shift the focus to *approximation algorithms*.

The first improvement for unweighted undirected graphs was due to Bernstein and Roditty [BR11] who presented a $(1 + \epsilon)$ -approximation algorithm with $O(n^{2+O(1/\sqrt{\log n})})$

expected total update time.¹ This time bound is only slightly larger than quadratic and beats the $O(mn)$ time of the ES-tree unless the input graph is very sparse. For the more general cases, Henzinger and King [HK95] observed that the ES-tree can be easily adapted to directed graphs. King [Kin99] later extended the ES-tree to an $O(mnW)$ -time algorithm for weighted directed graphs. A rounding technique used in recent algorithms of Bernstein [Ber09, Ber13] and Mądry [Mad10], as well as earlier papers on approximate shortest paths [Coh98, Zwi02], gives a $(1 + \epsilon)$ -approximate $\tilde{O}(mn \log W)$ -time algorithm for weighted directed graphs. Very recently, we obtained a $(1 + \epsilon)$ -approximation algorithm with total update time $O(mn^{0.9+o(1)})$ for decremental approximate SSSP in weighted directed graphs [HKN14b, HKN15] if $W \leq 2^{\log^c n}$ for some constant c . This gives the first $o(mn)$ time algorithm for the directed case, as well as other important problems such as single-source reachability and strongly-connected components [RZ08, Lac13, Rod13]. Also very recently, Abboud and Vassilevska Williams [AVW14] showed that “deamortizing” our algorithms in [HKN14b] might not be possible: a combinatorial algorithm with *worst case* update time and query time of $O(n^{2-\delta})$ (for some $\delta > 0$) per deletion implies a faster combinatorial algorithm for Boolean matrix multiplication and, for the more general problem of maintaining the number of reachable nodes from a source under deletions (which our algorithms in [HKN14b] can do) a worst case update and query time of $O(m^{1-\delta})$ (for some $\delta > 0$) will falsify the strong exponential time hypothesis.

Our Results. Given the significance of the decremental SSSP problem, it is important to understand its time complexity.

In this paper, we obtain a near-linear time algorithm for decremental $(1 + \epsilon)$ -approximate SSSP in weighted undirected graphs. Its total update time is $O(m^{1+O(\sqrt{\log \log n / \log n})} \log W)$ and maintains an estimate of the distance between the source node and every other node, guaranteeing constant worst-case query time. The algorithm is randomized and assumes an oblivious adversary who fixes the sequence of updates in advance; it is correct with high probability and the bound on its total update time holds in expectation. In the unweighted case, our algorithm significantly improves our previous algorithm in [HKN14a] as discussed above. There was no previous algorithm designed specifically for weighted undirected graphs, and the previous best running time for this case comes from our $O(mn^{0.9+o(1)})$ time for weighted directed graphs [HKN14b, HKN15].

As a consequence of our techniques we also obtain an algorithm for the all-pairs shortest paths (APSP) problem. For every integer $k \geq 2$ and every $0 < \epsilon \leq 1$, we obtain a decremental $((2 + \epsilon)^k - 1)$ -approximate APSP algorithm with query time $O(k^k)$ and total update time $O(m^{1+1/k+O(\log^{5/4}((\log n)/\epsilon)/\log^{1/4} n)} \log^2 W)$. We remark that for $k = 2$ and $1/\epsilon = O(\text{polylog } n)$ our result gives a $(3 + \epsilon)$ -approximation with constant query time and total update time $O(m^{1+1/2+o(1)} \log W)$. For very sparse graphs with $m = \Theta(n)$, this is almost optimal in the sense that it almost matches the static running time [TZ05] of $O(m\sqrt{n})$, providing stretch of $3 + \epsilon$ instead of 3 as in the static setting. Our result on approximate APSP has to be compared with the following prior work. For weighted directed graphs Bernstein [Ber13] gave a decremental $(1 + \epsilon)$ -approximate APSP algorithm with constant query time and total update time $\tilde{O}(mn \log W)$. For unweighted undirected graphs there are

¹To enhance readability we assume that ϵ is a constant when citing related work, thus omitting the dependence on ϵ in the running times.

two previous results that improve upon this update time at the cost of larger approximation error. First, for any fixed integer $k \geq 2$, Bernstein and Roditty gave a decremental $(2k - 1 + \epsilon)$ -approximate APSP algorithm with constant query time and total update time $\tilde{O}(mn^{1/k})$. Second, for any integer $k \geq 2$, Abraham, Chechik, and Talwar [ACT14] gave a decremental $2^{O(\rho k)}$ -approximate APSP algorithm for unweighted undirected graphs with query time $O(k\rho)$ and total update time $\tilde{O}(mn^{1/k})$, where $\rho = (1 + \lceil (\log n^{1-1/k}) / \log(m/n^{1-1/k}) \rceil)$.

2 Preliminaries

In this paper we want to maintain approximate shortest paths in an undirected graph $G = (V, E)$ with positive integer edge weights in the range from 1 to W , for some parameter W . The graph undergoes a sequence of *updates*, which might be edge deletions or edge weight increases. This is called the *decremental setting*. We denote by V the set of nodes of G and by E the set of edges of G before the first edge deletion. We set $n = |V|$ and $m = |E|$.

For every weighted undirected graph $G = (V, E)$, we denote the weight of an edge (u, v) in G by $w_G(u, v)$. The *distance* $\text{dist}_G(u, v)$ between a node u and a node v in G is the weight of the shortest path, i.e., the minimum-weight path, between u and v in G . If there is no path between u and v in G , we set $\text{dist}_G(u, v) = \infty$. For every set of nodes $U \subseteq V$ we denote by $E[U]$ the set of edges incident to the nodes of U , i.e., $E[U] = \{(u, v) \in E \mid u \in U\}$.² Furthermore, for every set of nodes $U \subseteq V$, we denote by $G|U$ the subgraph of G induced by the nodes in U , i.e., $G|U$ contains all edges (u, v) such that (u, v) is contained in E and u and v are both contained in U , or short: $G|U = (U, E \cap U^2)$. Similarly, for every set of edges $F \subseteq V^2$ and every set of nodes $U \subseteq V$ we denote by $F|U$ the subset of F induced by U .

We say that a distance estimate $\delta(u, v)$ is an (α, β) -approximation of the true distance $\text{dist}_G(u, v)$ if $\text{dist}_G(u, v) \leq \delta(u, v) \leq \alpha \text{dist}_G(u, v) + \beta$, i.e., $\delta(u, v)$ never underestimates the true distance and overestimates it with a multiplicative error of at most α and an additive error of at most β . If there is no additive error, we simply say α -approximation instead of $(\alpha, 0)$ -approximation.

In our algorithms we will use graphs that do not only undergo edge deletions and edge weight increases, but also edge insertions. For such a graph H , we denote by $\mathcal{E}(H)$ the number of edges ever contained in H , i.e., the number of edges contained in H before any deletion or insertion plus the number of inserted edges. We denote by $\mathcal{W}(H)$ the number of edge weight increases in H . Similarly, for a set of edges F , we denote by $\mathcal{E}(F)$ the number of edges ever contained in F and by $\mathcal{W}(F)$ the number of edge weight increases in F .

The central data structure in decremental algorithms for exact and approximate shortest paths is the Even-Shiloach tree (short: ES-tree). This data structure maintains a shortest paths tree from a root node up to a given depth D .

Lemma 2.1 ([ES81, HK95, Kin99]). *There is a data structure called ES-tree that, given a weighted directed graph G undergoing deletions and edge weight increases, a root node s , and a depth parameter D , maintains, for every node v a value $\delta(v, s)$ such that $\delta(v, s) = \text{dist}_G(v, s)$ if $\text{dist}_G(v, s) \leq D$ and $\delta(v, s) = \infty$ if $\text{dist}_G(v, s) > D$. It has constant query time and a total update time of $O(mD)$.*

²Since G is an undirected graph, this definition is equivalent to $E[U] = \{(u, v) \in E \mid u \in U \text{ or } v \in U\}$.

Recent approaches for solving approximate decremental SSSP and APSP use special graphs called *emulators*. An (α, β) -emulator H of a graph G is a graph containing the nodes of G such that $\text{dist}_G(u, v) \leq \text{dist}_H(u, v) \leq \alpha \text{dist}_G(u, v) + \beta$ for all nodes u and v .³ Maintaining exact distances on H provides an (α, β) -approximation of distances in G . As good emulators are sparser than the original graph this is usually more efficient than maintaining exact distances on G . However, the edges of H also have to be maintained while G undergoes updates. For unweighted, undirected graphs undergoing edge deletions, the emulator of Thorup and Zwick (based on the second spanner construction in [TZ06]), which provides a relatively good approximation, can be maintained quite efficiently [BR11]. However the definition of this emulator requires the occasional insertion of edges into the emulator. Thus, it is not possible to run a purely decremental algorithm on top of it.

There have been approaches to design algorithms that mimic the behavior of the classic ES-tree when run on an emulator that undergoes insertions. The first approach by Bernstein and Roditty [BR11] extends the ES-tree to a fully dynamic algorithm and analyzes the additional work incurred by the insertions. The second approach was introduced by us in [HKN13] and is called *monotone ES-tree*. It basically ignores insertions of edges into H and never decreases the distance estimate it maintains. However, this algorithm does not provide an (α, β) -approximation on *any* (α, β) -approximate emulator as it needs to exploit the structure of the emulator. In [HKN13] we gave an analysis of the monotone ES-tree when run on a specific $(1 + \epsilon, 2)$ -emulator and in the current paper we use a different analysis for our new algorithms. If we want to use the monotone ES-tree to maintain (α, β) -approximate distances up to depth D we will set the maximum level in the monotone ES-tree to $L = \alpha D + \beta$. The running time of the monotone ES-tree as analyzed in [HKN13] is as follows.

Lemma 2.2. *For every $L \geq 1$, the total update time of a monotone ES-tree up to maximum level L on a graph H undergoing edge deletions, edge insertions, and edge weight increases is $O(\mathcal{E}(H) \cdot L + \mathcal{W}(H))$.*

Our algorithms will heavily use randomization. It is well-known, and exploited by many other algorithms for dynamic (approximate) shortest paths and reachability, that by sampling a set of nodes with a sufficiently large probability we can guarantee that certain sets of nodes contain at least one of the sampled nodes whp. To the best of our knowledge, the first use of this technique in graph algorithms goes back to Ullman and Yannakakis [UY91].

Lemma 2.3. *Let T be a set of size t and let S_1, S_2, \dots, S_l be subsets of T of size at least s . Let U be a subset of T that was obtained by choosing each element of T independently with probability $p = (a \ln lt)/s$, where a is a constant. Then, for every $1 \leq i \leq l$, the set S_i contains a node of U with high probability (whp), i.e. probability at least $1 - 1/t^a$, and the size of U is $O((t \log(lt))/s)$ in expectation.*

3 Technical Overview

In the following we explain the main ideas of this paper, which lead to an algorithm for maintaining a hop set of a graph undergoing edge deletions.

³For the related notion of a spanner we additionally have to require that H is a subgraph of G .

General Idea. With the well-known algorithm of Even and Shiloach we can maintain a shortest paths tree from a source node up to a given depth D under edge deletions in time $O(mD)$. In unweighted graphs, all simple paths have length at most n and therefore we can set $D = n$ to maintain a full shortest paths tree. In weighted graphs with positive integer edge weights from 1 to W , all simple paths have length at most nW and therefore we can set $D = nW$ to maintain a full shortest paths tree. Using an established rounding technique [Coh98, Zwi02, Ber09, Mad10, Ber13, Nan14], one can use this algorithm to maintain $(1 + \epsilon)$ -approximate single-source shortest paths up to h hops in time $O(mh \log(nW)/\epsilon)$. With this algorithm we can set $h = n$ to maintain a full-length approximate shortest paths tree, even in weighted graphs. This algorithm would be very efficient if the graph had a small hop diameter, i.e., if for any pair of nodes there is a shortest path with a small number of edges. Our idea is to artificially construct such a graph.

To this end we will use a so-called *hop set*. An (h, ϵ) -hop set F of a graph $G = (V, E)$ is a set of weighted edges $F \subseteq V^2$ such that in the graph $H = (V, E \cup F)$ there exists, for every pair of nodes u and v , a path from u to v of weight at most $(1 + \epsilon) \text{dist}_G(u, v)$ and with at most h edges. If we run the approximate SSSP algorithm on H , we obtain a running time of $O((m + |F|)h \log(nW)/\epsilon)$. In our algorithm we will obtain an $(O(n^{o(1)}), \epsilon)$ -hop set of size $O(m^{1+o(1)})$ and thus the running time will be $O(m^{1+o(1)} \log(nW)/\epsilon)$. It is however not enough to simply construct the hop set at the beginning. We also need a dynamic algorithm for maintaining the hop set under edge deletions in G . We will present an algorithm that performs this task also in almost linear time over all deletions.

Roughly speaking, we achieve the following. Given a graph $G = (V, E)$ undergoing edge deletions, we can maintain a restricted hop set F such that, for all pairs of nodes u and v if the shortest path from u to v in G has $h \geq n^{1/q}$ hops, in the shortcut graph $H = (V, E \cup F)$ there is a path from u to v of weight at most $(1 + \epsilon) \text{dist}_G(u, v)$ and with at most $\lceil h/n^{1/q} \rceil \log n$ hops. Our high-level idea for maintaining an (unrestricted) $(n^{o(1)}, \epsilon)$ hop set is the following hierarchical approach. We start with $H_0 = G$ to maintain a hop set F_1 of G , which reduces the number of hops by a factor of $\log n/n^{1/q}$ at the cost of a multiplicative error of $1 + \epsilon$. Given F_1 , we use the shortcut graph $H_1 = (V, E \cup F_1)$ to maintain a hop set F_2 of G that reduces the number of hops by another factor of $\log n/n^{1/q}$ introducing another error of $1 + \epsilon$. By repeating this process q times we arrive at a hop set that guarantees, for all pairs of nodes u and v , a path of weight at most $(1 + \epsilon)^q \text{dist}_G(u, v)$ and with at most $(\log n)^q n^{1/q}$ hops. Figure 1 visualizes this hierarchical approach.

The notion of hop set was first introduced by Cohen [Coh00] in the PRAM literature and is conceptually related to the notion of emulator. It is also related to the notion of *shortest-paths diameter* used in distributed computing (e.g., [KKM⁺12, Nan14]). To the best of our knowledge, the only place that this hop set concept was used before in the dynamic algorithm literature (without the name being mentioned) is Bernstein’s fully dynamic $(2 + \epsilon)$ -approximation algorithm for all-pairs shortest paths [Ber09]. There, an $(n^{o(1)}, \epsilon)$ -hop set is essentially recomputed *from scratch* after every edge update, and a shortest-paths data structure is maintained on top of this hop set.

Static Hop Set. We first assume that $G = (V, E)$ is an unweighted undirected graph and for simplicity we also assume that ϵ is a constant. We explain how to obtain a hop set of G using a randomized construction of Thorup and Zwick [TZ06] based on the notion of balls of

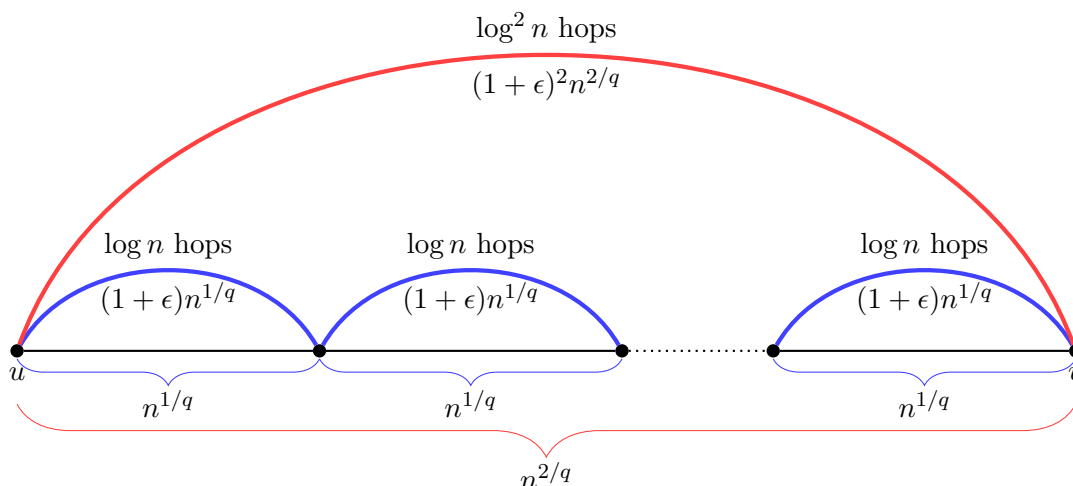


Figure 1: Illustration of the hierarchical approach for maintaining the hop set reduction. Here $q = \Theta(\sqrt{\log n})$ and u and v are nodes that are at distance $n^{2/q}$ from each other. In the first layer we find a hop set that shortcuts all subpaths of weight $n^{1/q}$ by paths of weight at most $(1 + \epsilon)n^{1/q}$ and with at most $\log n$ hops. In the second layer, we use the shortcuts of the first layer to find a hop set that shortcuts the path from u to v of weight $n^{2/q}$ by a path of weight at most $(1 + \epsilon)^2 n^{2/q}$ and with at most $\log^2 n$ hops.

nodes. We describe this construction and the hop-set analysis in the following.

Let $2 \leq p \leq \log n$ be a parameter and consider a sequence of sets of nodes A_0, A_1, \dots, A_p obtained as follows. We set $A_0 = V$ and $A_p = \emptyset$ and for $1 \leq i \leq p - 1$ we obtain the set A_i by picking each node of V independently with probability $1/n^{i/p}$. The expected size of A_i is $n^{1-i/p}$. For every node u we define the priority of u as the maximum i such that $u \in A_i$. For a node u of priority i we define

$$Ball(u) = \{v \in V \mid \text{dist}_G(u, v) < \text{dist}_G(u, A_{i+1})\} \quad (1)$$

where $\text{dist}_G(u, A_{i+1}) = \min_{v \in A_{i+1}} \text{dist}_G(u, v)$. Note that $\text{dist}_G(u, A_p) = \infty$ and thus if $u \in A_{p-1}$, then $Ball(u) = V$. For each node u of priority i the size of $Ball(u)$ is $n^{(i+1)/p}$ in expectation by the following argument: Order the nodes in non-decreasing distance from u . Each of these nodes belongs to A_{i+1} with probability $1/n^{(i+1)/p}$ and therefore, in expectation, we need to see $n^{(i+1)/p}$ nodes until one of them is contained in A_{i+1} . It follows that the expected size of all balls of priority i is at most $n^{1+1/p}$ (the expected size of A_i times the expected size of $Ball(u)$ for each node u of priority i) and the expected size of all balls, i.e., $\sum_{u \in V} |Ball(u)|$, is at most $pn^{1+1/p}$.

Let F be the set of edges $F = \{(u, v) \in V^2 \mid v \in Ball(u)\}$ and give each edge $(u, v) \in F$ the weight $w_F(u, v) = \text{dist}_G(u, v)$. By the argument above, the expected size of F is at most $pn^{1+1/p}$. An argument of Thorup and Zwick [TZ06] shows that the weighted graph $H = (V, F)$ has the following property for every pair of nodes u and v and any $0 < \epsilon \leq 1$

such that $1/\epsilon$ is integer:⁴

$$\text{dist}_G(u, v) \leq \text{dist}_H(u, v) \leq (1 + \epsilon) \text{dist}_G(u, v) + 2 \left(2 + \frac{2}{\epsilon}\right)^{p-2}.$$

Note that the choice of ϵ gives a trade-off in the error between the multiplicative part $(1 + \epsilon)$ and the additive part $2(2 + 2/\epsilon)^{p-2}$. In the literature, such a graph H is known as an *emulator* of G with multiplicative error $(1 + \epsilon)$ and additive error $2(2 + 2/\epsilon)^{p-2}$.⁵ Roughly speaking, the strategy in their proof is as follows. Let u' be the node following u on the shortest path from u to v in G . If the edge (u, u') is also contained in H , then we can shorten the distance to v by 1 without introducing any approximation error (recall that we assume that G is unweighted). Otherwise, one can show that there is a path π' with at most p edges in H from u to a node v' closer to v than u such that the ratio between the weight of π' and the distance from u to v' is at most $(1 + \epsilon)$, and, if $v' = v$, then the weight of π' is at most $2(2 + 2/\epsilon)^{p-2}$. The proof needs the following property of the balls: for every node u of priority i and every node v either $v \in \text{Ball}(u)$ or there is some node v' of priority $j > i$ such that $u \in \text{Ball}(v')$. We illustrate the proof strategy in Figure 2.

Observe that the same strategy can be used to show the following: Given any integer $\Delta \leq n$, let u' be the node that is at distance Δ from u on the shortest path from u to v in G . If the edge (u, u') is contained in H , then we can shorten the distance to v by Δ without introducing any approximation error. Otherwise, one can show that there is a path π' with at most p edges in H from u to a node v' closer to v than u such that the ratio between the weight of π' and the distance from u to v' is at most $(1 + \epsilon)$, and, if $v' = v$, then the weight of π' is at most $2(2 + 2/\epsilon)^{p-2} \cdot \Delta$. Every time we repeat this argument the distance to v is shortened by at least Δ . Therefore there is a path from u to v in H with at most $p \lceil \text{dist}_G(u, v) / \Delta \rceil$ edges that has weight at most $(1 + \epsilon) \text{dist}_G(u, v) + 2(2 + 2/\epsilon)^{p-2} \cdot \Delta$. One can show that this statement would also be true if we had removed all edges from F of weight more than $(1 + 2/\epsilon)(2 + 2/\epsilon)^{p-2}$, which is the maximum weight of the edge to v' in the proof strategy above outlined in Figure 2. We will need this fact in the dynamic algorithm as it allows us to limit the depth of the balls.

By a suitable choice of $p = \Theta(\sqrt{\log n})$ (as a function of n and ϵ) we can guarantee that $2(2 + 2/\epsilon)^{p-2} \leq \epsilon n^{1/p}$ and $n^{1/p} = n^{o(1)}$. Now define $q = p$ and $\Delta_k = n^{k/q}$ for each $0 \leq k \leq q - 2$. Then we have, for every $0 \leq k \leq q - 2$ and all pairs of nodes u and v

$$\begin{aligned} \text{dist}_G(u, v) \leq \text{dist}_H(u, v) &\leq (1 + \epsilon) \text{dist}_G(u, v) + 2 \left(2 + \frac{2}{\epsilon}\right)^{p-2} \cdot \Delta_k \\ &\leq (1 + \epsilon) \text{dist}_G(u, v) + \epsilon n^{1/p} \cdot \Delta_k \\ &= (1 + \epsilon) \text{dist}_G(u, v) + \epsilon \Delta_{k+1}. \end{aligned}$$

Thus, if $\Delta_{k+1} \leq \text{dist}_G(u, v) \leq \Delta_{k+2}$, then there is a path from u to v in H of weight at most

$$(1 + \epsilon) \text{dist}_G(u, v) + \epsilon \Delta_{k+1} \leq (1 + \epsilon) \text{dist}_G(u, v) + \epsilon \text{dist}_G(u, v) = (1 + 2\epsilon) \text{dist}_G(u, v)$$

⁴The requirement that $1/\epsilon$ must be integer is not needed in the paper of Thorup and Zwick; we have added it here to simplify the exposition.

⁵In their paper, Thorup and Zwick [TZ06] actually define a graph H' whose set of edges is the union of the shortest paths trees from every node u to all nodes in its ball. This graph has the same approximation error and the same size as H ; since H' is a subgraph of G it is called a *spanner* of G .

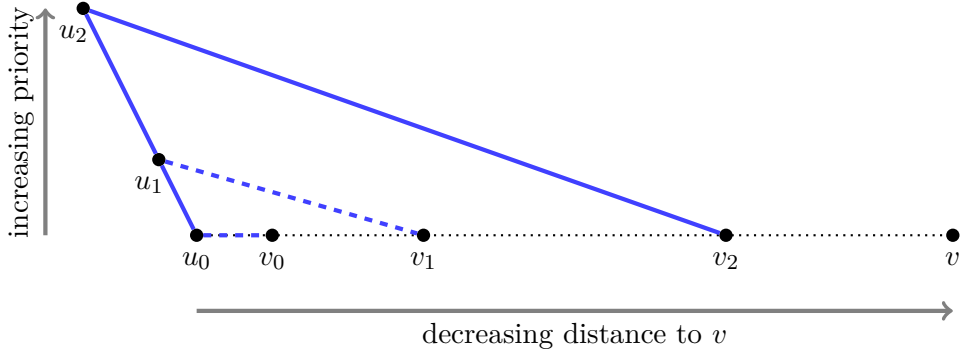


Figure 2: Illustration of the hop reduction for $p = 3$ priorities. The dotted line is the shortest path π from u_0 to v in G . The thick, blue edges are the edges of F used to shorten the distance to v . The dashed, blue edges are not contained in F and imply the existence of edges to nodes of increasing priority. Starting from u_0 , a node of priority 0, we let v_0 be the node on π such that $\text{dist}_G(u_0, v_0) = r_0 := 1$, i.e., the neighbor of u_0 on π . If the edge (u_0, v_0) is not contained in F , then F contains an edge (u_0, u_1) to a node u_1 of priority at least 1 such that $\text{dist}_G(u_0, u_1) \leq r_0$. Let v_1 be the node on π such that $\text{dist}_G(u_1, v_1) = r_1 := 1 + 2/\epsilon$. If the edge (u_1, v_1) is not contained in F , then F contains an edge (u_1, u_2) to a node u_2 of priority at least 2 such that $\text{dist}_G(u_1, u_2) \leq r_1$. Let v_2 be the node on π such that $\text{dist}_G(u_2, v_2) = r_2 := (1 + 2/\epsilon)(2 + 2/\epsilon)$. Since 2 is the highest priority, u_2 contains the edge (u_2, v_2) . Note that the weight of these three edges from F is at most $r_0 + r_1 + r_2$ and $\text{dist}_G(u_0, v_2) \geq r_2 - (r_0 + r_1)$. Since $r_2 = (1 + 2/\epsilon)(r_0 + r_1)$, the ratio between these two quantities is $(1 + \epsilon)$.

and with at most $p \lceil \text{dist}_G(u, v) / \Delta_k \rceil \leq (p+1) \Delta_{k+2} / \Delta_k = (p+1) n^{2/q} = n^{o(1)}$ edges. It follows that F is a $(2\epsilon, n^{o(1)})$ hop set of size $O(pn^{1+1/p}) = O(n^{1+o(1)})$. By running the algorithm with $\epsilon' = \epsilon/2$ we obtain an $(\epsilon, n^{o(1)})$ hop set of size $O(n^{1+o(1)})$.

Efficient Construction. So far we have ignored the running time for computing the balls and thus constructing F , even in the static setting. Thorup and Zwick [TZ06] have remarked that a naive algorithm for computing the balls takes time $O(mn)$. We can reduce this running time by sampling edges instead of nodes.

We modify the process for obtaining the sequence of sets A_0, A_1, \dots, A_p as follows. We set $A_0 = V$ and $A_p = \emptyset$ and for $1 \leq i \leq p-1$ we obtain the set A_i by picking each edge of E independently with probability $1/m^{i/p}$ and adding both endpoints of each sampled edge to A_i . The priority of a node u is the maximum i such that $u \in A_i$. We define, for every node u of priority i , $Ball(u)$ just like in Equation (1), but using the new definition of A_i . Note that the expected size of A_i is $O(m^{1-i/p})$ for every $1 \leq i \leq p-1$.

The balls can now be computed as follows. First, we use the following process to compute $\text{dist}_G(u, A_i) = \min_{v \in A_i} \text{dist}_G(u, v)$ for every node u and every $1 \leq i \leq p-1$. Using Dijkstra's algorithm on a graph where we add an artificial source node s_i that is connected to every node in A_i by an edge of weight 0, this takes time $O(p(m+n \log n))$. Second, we compute for every node u of priority i a shortest paths tree *up to depth* $\text{dist}_G(u, A_{i+1}) - 1$ to obtain all nodes contained in $Ball(u)$. Using an implementation of Dijkstra's algorithm that only puts nodes into its queue upon their first visit this takes time $O(|E[Ball(u)]| \log n)$ where $E[Ball(u)] = \{(u, v) \in E \mid u \in Ball(u) \text{ or } v \in Ball(u)\}$ is the set of edges incident to $Ball(u)$. By the sampling of edges we have, using the same ordering argument as before, that the expected size of $Ball(u)$ is $m^{(i+1)/p}$. For $0 \leq i \leq p-1$ the expected size of A_i is $O(m^{1-i/p})$ and thus these Dijkstra computations take time $O(m^{1+1/p} \log n)$ for all nodes of priority i . By choosing $p = \Theta(\sqrt{\log n})$ as described above we have $m^{1/p} = O(m^{o(1)})$ and thus the balls can be computed in time $O(m^{1+o(1)})$.

We define F as the set of edges $F = \{(u, v) \in V^2 \mid v \in Ball(u)\}$ and give each edge $(u, v) \in F$ the weight $w_F(u, v) = \text{dist}_G(u, v)$. The distance-preserving and hop-reducing properties of F still hold as stated above and its expected size is $O(pm^{1+1/p})$. Note that F is not a sparsification of G anymore (as the bound on its size is even more than m). For our purposes the sparsification aspect is not relevant, we only need the hop reduction. Thus in the static setting, we can compute an $(\epsilon, m^{o(1)})$ -hop set (which is also an $(\epsilon, n^{o(1)})$ -hop set) of expected size $O(m^{1+o(1)})$ in expected time $O(m^{1+o(1)})$.

Maintaining Balls Under Edge Deletions. As the graph G undergoes deletions the hop set has to be updated as well. We can achieve this by maintaining the balls w.r.t. a fixed sequence of randomly chosen sets A_0, A_1, \dots, A_p , where $A_0 = V$ and $A_p = \emptyset$ and for $1 \leq i \leq p-1$ we obtain A_i by picking each edge of E independently with probability $c \ln n / m^{i/p}$, for some large enough constant c , and adding both endpoints of each sampled edge to A_i . Note that now, for every $1 \leq i \leq p-1$, the expected size of A_i is $O(m^{1-i/p} \log n)$ and the size of $E[Ball(u)]$ for every node u of priority i is at most $m^{(i+1)/p}$ *with high probability* (whp) at any time. This holds because by deleting edges there can only be a polynomial number of different graphs in the whole sequence of updates. Unfortunately, we do not know how to maintain the balls efficiently. However we can maintain for all nodes u

the approximate ball

$$Ball(u, D) = \{v \in V \mid \log \text{dist}_G(u, v) < \lfloor \log \text{dist}_G(u, A_{i+1}) \rfloor \text{ and } \text{dist}_G(u, v) \leq D\}$$

(where u has priority i) in time $O(pm^{1+1/p}D \log D)$. Note that $Ball(u, D)$ differs from the definition of $Ball(u)$ in the following ways. First, we use the inequality $\log \text{dist}_G(u, v) < \lfloor \log \text{dist}_G(u, A_{i+1}) \rfloor$ instead of the inequality $\text{dist}_G(u, v) < \text{dist}_G(u, A_{i+1})$. This alone increases the additive error of the hops set from $2(2 + 2/\epsilon)^{p-2}$ to $4(3 + 4/\epsilon)^{p-2}$, which can easily be compensated. Second, we limit the balls to a certain depth D . By using a small value of D we will only obtain a restricted hop set that provides sufficient hop reduction for nodes that are relatively close to each other. We will show later that this is enough for our purposes. Despite of these modifications we clearly have $Ball(u, D) \subseteq Ball(u)$ and therefore all size bounds still apply.

In the first part of the algorithm for maintaining the balls we maintain $\text{dist}_G(u, A_i)$ every $1 \leq i \leq p - 1$ and every node u . We do this by adding an artificial source node s_i that has an edge of weight 0 to every node in A_i and maintain an ES-tree up to depth D from s_i . This step takes time $O(pmD)$.

Now, for every node u of priority i we maintain $Ball(u, D)$ as follows. We maintain an ES-tree up to depth

$$\min \left(2^{\lfloor \log \text{dist}_G(u, A_{i+1}) \rfloor} - 1, D \right)$$

and every time $2^{\lfloor \log \text{dist}_G(u, A_{i+1}) \rfloor}$ increases, we restart the ES-tree. Naively, we incur a cost of $O(mD)$ for each instance of the ES-tree. However we can easily implement the ES-tree in a way that it never looks at edges that are not contained in $E[Ball(u)]$.⁶ Thus, the cost of each instance of the ES-tree is $O(|Ball(u)|D)$. Remember that $Ball(u, D) \subseteq Ball(u)$ and that by the random sampling of edges the size of $E[Ball(u)]$ is at most $m^{(i+1)/p}$ whp. As $2^{\lfloor \log \text{dist}_G(u, A_{i+1}) \rfloor}$ can increase at most $\log D$ times until it exceeds D , we initialize at most $\log D$ ES-trees for the node u . Therefore the total time needed for maintaining $Ball(u, D)$ is $O(m^{(i+1)/p}D \log D)$. As there are at most $\tilde{O}(m^{1-i/p})$ nodes of priority i in expectation, the total time needed for maintaining all approximate balls is $\tilde{O}(pm^{1+1/p}D \log D)$ in expectation.

Decremental Approximate SSSP. Let us first sketch an algorithm for maintaining shortest paths from a source node s with a running time of $O(m^{1+1/2+o(1)})$ for which we use $q = 2$ layers and $p = \Theta(\sqrt{\log n})$ priorities. We set $\Delta = \sqrt{n}$, $D = 2(3 + 4/\epsilon)^{p-2}\Delta$, and p such that $2(3 + 4/\epsilon)^{p-2} \leq \epsilon n^{1/p}$ and $n^{1/p} = O(n^{o(1)})$. We maintain single-source shortest paths up to depth D from s using the ES-tree, which takes time $O(mD) = O(mn^{1/2+o(1)})$. To maintain approximate shortest paths to nodes that are at distance more than D from s we use the following approach. We maintain $Ball(u, D)$ for every node u , as sketched above, which takes time $\tilde{O}(pm^{1+1/p}D \log D) = O(m^{1+1/2+o(1)})$ in expectation. At any time, we set the hop set F to be the set of edges $F = \{(u, v) \in V^2 \mid \exists v \in Ball(u, D)\}$ and give each edge $(u, v) \in F$ the weight $w_F(u, v) = \text{dist}_G(u, v)$. By our arguments above, the weighted graph $H = (V, F)$ has the following property: for every pair of nodes u and v such that

⁶If we prefer to use the ES-tree as a “black box” we can, in a preprocessing step, find the initial set $Ball(u, D)$ and only build an ES-tree for this ball. All other nodes will never be contained in $Ball(u, D)$ anymore as long as the value of $2^{\lfloor \log \text{dist}_G(u, A_{i+1}) \rfloor}$ remains unchanged and therefore we can remove them. This can be done in time $O(|E[Ball(u, D)]| \log n)$ by using an implementation of Dijkstra’s algorithm that only puts nodes into its queue upon their first visit.

$\text{dist}_G(u, v) \geq D \geq \Delta$ there is a path π in H of weight at most $(1 + 2\epsilon) \text{dist}_G(u, v)$ and with at most $p \text{dist}_G(u, v) / \Delta$ edges.

To maintain approximate shortest paths for nodes at distance more than D from s we will now use the hop reduction in combination with the following rounding technique. We set $\varphi = \epsilon \Delta / p$ and let H' be the graph resulting from rounding up every edge weight in H to the next multiple of φ . By using H' instead of H we incur an error of φ for every edge on the approximate shortest path π . Thus in H' , π has weight at most

$$\begin{aligned} (1 + 2\epsilon) \text{dist}_G(u, v) + (p \text{dist}_G(u, v) / \Delta) \cdot \varphi &= (1 + 2\epsilon) \text{dist}_G(u, v) + \epsilon \text{dist}_G(u, v) \\ &\leq (1 + 3\epsilon) \text{dist}_G(u, v). \end{aligned}$$

The efficiency now comes from the observation that we can run the algorithm on the graph H'' where every edge weight in H' is scaled down by a factor of $1/\varphi$. The graph H'' has integer weights and the weights of all paths in H' and H'' differ exactly by the factor $1/\varphi$. Thus, instead of maintaining a shortest paths tree up to depth n in H we only need to maintain a shortest paths tree in H'' up to depth $n/\varphi = p\sqrt{n}/\epsilon$. In this way we obtain a $(1 + 3\epsilon)$ -approximation for all nodes such that $\text{dist}_G(u, v) \geq D$.

However, we cannot simply use the ES-tree on H'' because as edges are deleted from G , nodes might join the approximate balls and therefore edges might be inserted into F and thus into H'' . This means that a dynamic shortest paths algorithm running on H'' would not find itself in a purely decremental setting. However the insertions have a “nice” structure. We can deal with them by using a previously developed technique, called *monotone ES-tree*. The main idea of the monotone ES-tree is to ignore the level decreases suggested by inserting edges. The hop-set proof still goes through, even though we are not arguing about the current distance in H'' anymore, but the level of a node u in the monotone ES-tree. Maintaining the monotone ES-tree for distances up to D in H'' takes time $O(|\mathcal{E}(H)|D)$ where $\mathcal{E}(H'')$ is the set of edges ever contained in H'' (including those that are inserted over time) and $D = O(n^{1/2+1/p})$ as explained above. Each insertion of an edge into F corresponds to a node joining $\text{Ball}(u, D)$ for some node u . For a fixed node u of priority i there are at most $\log D$ possibilities for nodes to join $\text{Ball}(u, D)$ (namely each time $\lfloor \log \text{dist}_G(u, A_{i+1}) \rfloor$ increases) and every time at most $m^{(i+1)/p}$ nodes will join whp. It follows that $|\mathcal{E}(H)|$ is $O(m^{1+o(1)})$ whp and the running time of this step is $O(m^{1+1/2+o(1)})$ in expectation.

The almost linear-time algorithm is just slightly more complicated. Here we use $p = \Theta(\sqrt{\log n})$ priorities and $q = \sqrt{p}$ layers and set $\Delta_k = n^{k/q}$ for each $0 \leq k \leq q - 2$. In the algorithm we will maintain, for each $0 \leq k \leq q - 2$ a hop set F_k such that for every pair of nodes u and v with $\Delta_{k+1} \leq \text{dist}_G(u, v) \leq \Delta_{k+2}$ there is a path from u to v in $H_k = (V, F_k)$ of weight at most $(1 + 2\epsilon) \text{dist}_G(u, v)$ and with at most $p \text{dist}_G(u, v) / \Delta_k \leq pn^{2/q}$ hops. To achieve this we use the following hierarchical approach. Given the hop set F_k we can maintain approximate shortest paths up to depth Δ_{k+2} in time $O(m^{1+o(1)})$ and given a data structure for maintaining approximate shortest paths up to depth Δ_k we can maintain approximate balls and thus the hop set F_{k+1} in time $O(m^{1+o(1)})$. The hierarchy “starts” with using the ES-tree as an algorithm for maintaining an (exact) shortest paths tree up to depth $n^{2/q}$. Thus, running efficient monotone ES-trees on top of the hop sets and maintaining the hop sets (using efficient monotone ES-trees) go hand in hand.

There are two obstacles in implementing this hierarchical approach when we want to maintain the approximate balls in each layer. First, in our algorithm for maintaining the

approximate balls we have used the ES-tree as an exact decremental SSSP algorithm. In the multilayer approach we have to replace the ES-tree with the monotone ES-tree which only provides approximate distance estimates. This will lead to approximation errors that increase with the number of layers. Second, by the arguments above the number of edges in F_k is $O(m^{1+1/p})$ for each $0 \leq k \leq q - 2$. In the algorithm for maintaining the approximate balls for the next layer, this bound however is not good enough because we run a separate instance of the monotone ES-tree for each node u . We deal with this issue by running the monotone ES-tree in the subgraph of G induced by the nodes initially contained in $Ball(u)$. For a node u of priority i this subgraph contains $m_i = m^{(i+1)/p}$ edges whp and we can recursively run our algorithm on this smaller graph. By this process we incur a factor of $m^{1/p}$ in the running time each time each time we increase the depth of the recursion. This results in a total update time of $\tilde{O}(m^{1+q/p})$ which is $\tilde{O}(m^{1+1/q}) = O(m^{1+o(1)})$ since $q = \sqrt{p}$.

Extension to Weighted Graphs. The hop set construction described above was only for unweighted graphs. However, the main property that we needed was $\text{dist}_G(u, v) \leq n$ for any pair of nodes u and v . Using the rounding technique mentioned above, we can construct for each $0 \leq i \leq \lfloor \log nW \rfloor$ a graph G_i such that for all pairs of nodes u and v with $2^i \leq \text{dist}_G(u, v) \leq 2^{i+1}$ we have $\text{dist}_{G_i}(u, v) \leq 4n/\epsilon$ for some small γ and the shortest path in G_i can be turned into a $(1 + \epsilon)$ -approximate shortest path in G by scaling up the edge weights. We now run $O(\log(nW))$ instances of our algorithm, one for each graph G_i , and maintain the hop set and approximate SSSP for each of them.

We only need to refine the analysis of the hop-set property in the following way. Remember that in the analysis we considered the shortest path π from u to v and defined the node u' that is at distance Δ from u on π . If the hop set contained the edge (u, u') we could reduce the distance to v by Δ . In weighted graphs (even after the scaling), we cannot guarantee there is a node at distance Δ from u on π . Therefore we define u' as the furthest node that is at distance *at most* Δ from u on π . Furthermore we define u'' as the neighbor of u' on π , i.e., u'' is at distance *at least* Δ from u . Now if the hop set contains the edge (u, u') we first use the edge (u, u') from the hop set, and then the edge (u', u'') from the original graph to reduce the distance to v by at least Δ with only 2 hops. Note that for unweighted graphs it was sufficient to only use the edges of the hop set. For weighted graphs we really have to add the edges of the hop set to the original graph in our algorithm.

Outline. As sketched above, our algorithm uses the following hierarchical approach: Given a decremental approximate SSSP algorithm for distances up to D_i with total update time $O(m^{1+o(1)})$, we can maintain approximate balls for distances up to D_i with total time $O(m^{1+o(1)})$ as well. And given a decremental algorithm for maintaining approximate balls for distances up to D_i with total update time $O(m^{1+o(1)})$ we can use the approximate balls to define a hop set which allows us to maintain approximate shortest paths for distances up to $D_{i+1} = n^{o(1)}D_i$ with total update time $O(m^{1+o(1)})$. This scheme is repeated until D_i is large enough to cover the full distance range.

We have formulated the two parts of this scheme as reductions. In Section 4 we give a decremental algorithm for maintaining approximate balls that internally uses a decremental approximate SSSP algorithm. In Section 5 we give a decremental approximate SSSP algorithm that internally uses a decremental algorithm for maintaining approximate balls. In Section 6

we explain the hierarchical approach for putting these two parts together and obtain the decremental $(1 + \epsilon)$ -approximate SSSP algorithm with a total update time of $O(m^{1+o(1)})$ for the full distance range. In addition to this result, the algorithm for maintaining approximate balls, together with a suitable query algorithm, gives us a decremental approximate APSP algorithm. This algorithm is also given in Section 6.

4 From Approximate SSSP to Approximate Balls

In the following we show how to maintain the approximate balls of every node if we already have an algorithm for maintaining approximate shortest paths. In our reduction we will use the algorithm for maintaining approximate shortest paths as a “black box”, requiring only very few properties. Formally, we prove the following statement in this section.

Proposition 4.1. *Assume there is a decremental approximate SSSP algorithm APPROXSSSP with the following properties, using fixed values $\alpha \geq 1$, $\beta \geq 0$, and $D \geq 1$: Given a weighted graph $G = (V, E)$ undergoing edge deletions and edge weight increases and a fixed source node $s \in V$, APPROXSSSP maintains for every node $v \in V$ a distance estimate $\delta(s, v)$ such that:*

A1 $\delta(s, v) \geq \text{dist}_G(s, v)$

A2 If $\text{dist}_G(s, v) \leq D$, then $\delta(s, v) \leq \alpha \text{dist}_G(s, v) + \beta$.

A3 After every update in G , APPROXSSSP returns, for every node v such that $\delta(s, v)$ has changed, v together with the new value of $\delta(s, v)$.

We denote the total update time of APPROXSSSP by $T(m, n)$.

Then there is a decremental algorithm APPROXBALLS for maintaining approximate balls with the following properties: Given a weighted graph $G = (V, E)$ undergoing edge deletions and edge weight increases and parameters $2 \leq k \leq \log n$ and $0 < \epsilon \leq 1$, it assigns to every node $u \in V$ a number from 0 to $k - 1$, called the priority of u , and maintains for every node $u \in V$ a set of nodes $B(u)$ and a distance estimate $\delta(u, v)$ for every node $v \in B(u)$ such that:

B1 For every node u and every node $v \in B(u)$ we have $\text{dist}_G(u, v) \leq \delta(u, v) \leq \alpha \text{dist}_G(u, v) + \beta$.

B2 Let $s(\cdot, \cdot)$ be a non-decreasing function such that, for all $x \geq 1$, and $l \geq 1$,

$$s(x, l) \geq a(a + 1)^{l-1} \text{dist}_G(u, v) + ((a + 1)^l - 1)b/a,$$

where $a = (1 + \epsilon)\alpha$ and $b = (1 + \epsilon)\beta + 1$. Then for every node u of priority i and every node v such that $s(\text{dist}_G(u, v), p - 1 - i) \leq D$, either (1) $v \in B(u)$ or (2) there is some node v' of priority $j > i$ such that $u \in B(v')$ and $\text{dist}_G(u, v') \leq s(\text{dist}_G(u, v), j - i)$.

B3 If, for every node u , $\mathcal{B}(u)$ denotes the set of nodes ever contained in $B(u)$, then $\sum_{u \in V} |\mathcal{B}_u| = \tilde{O}(m^{1+1/k} \log D / \epsilon)$ in expectation.

B4 The update time of APPROXBALLS is

$$t(m, n, k, \epsilon) = \tilde{O} \left(m^{1+1/k} \log D/\epsilon + \sum_{0 \leq i \leq k-1} m^{1-i/k} \cdot T(m_i, n_i) \log D/\epsilon + T(m, n) \right),$$

where, for each $0 \leq i \leq k-1$, $m_i = m^{(i+1)/k}$ and $n_i = \min(m_i, n)$.

B5 After every update in G , APPROXBALLS returns all pairs of nodes u and v such that v joins $B(u)$, v leaves $B(u)$, or $\hat{\delta}(u, v)$ changes.

Our algorithm for maintaining the approximate balls $B(u)$ for every node $u \in V$ is as follows:

1. At the initialization we set $F_0 = E$ and $F_k = \emptyset$ and for $1 \leq i \leq k-1$, a set of edges F_i is obtained from sampling each edge of E independently with probability $(c \ln n)/m^{i/k}$ (for a large enough constant c). For every $0 \leq i \leq k-1$ we set $A_i = \{v \in V \mid \exists (v, w) \in F_i\}$ and for every node $v \in V$, we set the *priority* of u to be the maximum i such that $v \in A_i$.
2. For each $1 \leq i \leq k-1$ we run an instance of APPROXSSSP with depth D from an artificial source node s_i that has an edge of weight 0 to every node in A_i . We denote the distance estimate provided by APPROXSSSP $\delta(u, A_i)$ and set $\delta(u, A_k) = \infty$ for every node $u \in V$.
3. For every $0 \leq i \leq k-1$ and every node $u \in V$ of priority i , we maintain the value

$$r(u) = \min \left(\frac{(1 + \epsilon)^{\lfloor \log_{1+\epsilon}(\delta(u, A_{i+1}) - 1) \rfloor} - \beta}{\alpha}, D \right).$$

and at the initialization and each time $r(u)$ increases we do the following:

- (a) Compute the set of nodes $R(u) = \{v \in V \mid \text{dist}_G(u, v) \leq r(u)\}$.
- (b) Run an instance of APPROXSSSP with depth D from u in $G|R(u)$, the subgraph of G induced by $R(u)$. Let $\delta(u, v)$ denote the estimate of the distance between u and v in $G|R(u)$ maintained by APPROXSSSP.
- (c) Maintain $B(u) = \{v \in V \mid \delta(u, v) \leq \alpha D + \beta\}$: every time $\delta(u, v)$ changes for some node v we check whether v is still contained in $B(u)$.

Note that APPROXBALLS has Property **B5**, i.e., it returns changes in the approximate balls and the distance estimates, which is possible because APPROXSSSP has Property **A3**.

4.1 Relation to Exact Balls

In the following we compare the approximate balls maintained by our algorithm to the exact balls, as used by Thorup and Zwick [TZ06]. We show how the main properties of exact balls translate to approximate balls. We use the following definition of the (exact) ball of a node u of priority i :

$$\text{Ball}(u) = \{v \in V \mid \text{dist}_G(u, v) < \text{dist}_G(u, A_{i+1})\}.$$

The balls have the following simple property: If $v \notin \text{Ball}(u)$, then there is a node v' of priority $j > i$ such that $\text{dist}_G(u, v') \leq \text{dist}_G(u, v)$. We show that a relaxed version of this statement also holds for the approximate balls.

Lemma 4.2. *Let u be a node of priority i and let v be a node such that $\text{dist}_G(u, v) \leq D$. If $v \notin B(u)$, then there is a node v' of priority $j > i$ such that $\text{dist}_G(u, v') \leq a \text{dist}_G(u, v) + b$, where $a = (1 + \epsilon)\alpha$ and $b = (1 + \epsilon)\beta + 1$.*

Proof. We show the following: If $\text{dist}_G(u, A_{i+1}) \geq a \text{dist}_G(u, v)$, then $v \in B(u)$. The claim then follows from contraposition: If $v \notin B(u)$, then $\text{dist}_G(u, A_{i+1}) < a \text{dist}_G(u, v) + b$ and thus there exists some node $v' \in A_{i+1}$ that has priority $j \geq i + 1$ such that $\text{dist}_G(u, v') < a \text{dist}_G(u, v) + b$.

Assume that $\text{dist}_G(u, A_{i+1}) \geq a \text{dist}_G(u, v) + b$. Remember that we have set

$$r(u) = \min \left(\frac{(1 + \epsilon)^{\lfloor \log_{1+\epsilon}(\delta(u, A_{i+1}) - 1) \rfloor} - \beta}{\alpha}, D \right).$$

Since $\delta(u, A_{i+1}) \geq \text{dist}_G(u, A_{i+1})$ by Property **A1** we have

$$\delta(u, A_{i+1}) \geq \text{dist}_G(u, A_{i+1}) \geq a \text{dist}_G(u, v) + b = (1 + \epsilon)(\alpha \text{dist}_G(u, v) + \beta) + 1$$

which is equivalent to

$$\text{dist}_G(u, v) \leq \frac{\frac{\delta(u, A_{i+1}) - 1}{1 + \epsilon} - \beta}{\alpha}.$$

Since $(1 + \epsilon)^{\lfloor \log_{1+\epsilon}(\delta(u, A_{i+1}) - 1) \rfloor} \geq (1 + \epsilon)^{\log_{1+\epsilon}(\delta(u, A_{i+1}) - 1) - 1} = (\delta(u, A_{i+1}) - 1)/(1 + \epsilon)$ it follows that

$$\text{dist}_G(u, v) \leq \frac{(1 + \epsilon)^{\lfloor \log_{1+\epsilon}(\delta(u, A_{i+1}) - 1) \rfloor} - \beta}{\alpha}$$

Since we have assumed that $\text{dist}_G(u, v) \leq D$ we get $\text{dist}_G(u, v) \leq r(u)$ which implies that $\text{dist}_{G|R(u)}(u, v) \leq r(u) \leq D$ as well. Thus, by Property **A2**, it follows that $\delta(u, v) \leq \alpha \text{dist}_{G|R(u)} + \beta \leq \alpha D + \beta$ and $v \in B(u)$ as desired. \square

We now show that the approximate balls are contained in the exact balls. The exact balls are useful in our analysis because we can easily bound their size.

Lemma 4.3. *At any time $B(u) \subseteq \text{Ball}(u)$ for every node u .*

Proof. Let $R(u) = \{v \in V \mid \text{dist}_G(u, v) \leq r(u)\}$ denote the set of nodes in distance at most $r(u)$ from u at the last time $r(u)$ has increased. Note that $B(u)$ is a set of nodes of the graph $G|R(u)$ and therefore $B(u) \subseteq R(u)$. It remains to show that $R(u) \subseteq \text{Ball}(u)$.

If $i = k - 1$, then the claim is trivially true because $\text{Ball}(u)$ contains all nodes that are connected to u in G . In the case $0 \leq i < k - 1$ remember that

$$r(u) = \min \left(\frac{(1 + \epsilon)^{\lfloor \log_{1+\epsilon}(\delta(u, A_{i+1}) - 1) \rfloor} - \beta}{\alpha}, D \right).$$

If $\text{dist}_G(u, A_{i+1}) > r(u)$, we trivially have $\text{dist}_G(u, v) \leq r(u) < \text{dist}_G(u, A_{i+1})$. If on the other hand $\text{dist}_G(u, A_{i+1}) \leq r(u)$, then in particular $\text{dist}_G(u, A_{i+1}) \leq D$ and by Property **A2** we have $\delta(u, A_{i+1}) \leq \alpha \text{dist}_G(u, A_{i+1}) + \beta$. It follows that

$$\begin{aligned} \text{dist}_G(u, v) \leq r(u) &\leq \frac{(1 + \epsilon)^{\lfloor \log_{1+\epsilon}(\delta(u, A_{i+1}) - 1) \rfloor} - \beta}{\alpha} \\ &\leq \frac{(1 + \epsilon)^{\log_{1+\epsilon}(\delta(u, A_{i+1}) - 1)} - \beta}{\alpha} \leq \frac{\delta(u, A_{i+1}) - 1 - \beta}{\alpha} \leq \text{dist}_G(u, A_{i+1}) - \frac{1}{\alpha}. \end{aligned}$$

This implies that $\text{dist}_G(u, v) < \text{dist}_G(u, A_{i+1})$ because the distances are integer and $1/\alpha > 0$. In both cases we get $\text{dist}_G(u, v) < \text{dist}_G(u, A_{i+1})$ and as this is the defining property of $\text{Ball}(u)$ we have $v \in \text{Ball}(u)$. \square

Lemma 4.4. *At any time the size of $\text{Ball}(u)$ is $O(m^{(i+1)/k})$ whp for every node u of priority i .*

Proof. Note that the claim is trivially true for $i = k - 1$. Therefore assume $i < k - 1$ in the following. For every edge $e = (v, w) \in E$ we define $\text{dist}_G(u, e) = \min(\text{dist}_G(u, v), \text{dist}_G(u, w))$. Let $F \subseteq E$ be the set consisting of the $m^{(i+1)/k}$ edges that are closest to u according to this definition of $\text{dist}_G(u, e)$ for each edge e , where ties are broken in an arbitrary but fixed order. (Note that if less than $m^{(i+1)/k}$ edges are connected to u , then the claim is true anyway).

Let U be the set of nodes $U = \{v \in V \mid \exists (v, w) \in F\}$. By the random sampling (Lemma 2.3), F contains an edge (v, w) of F_{i+1} with high probability. Assume without loss of generality that $\text{dist}_G(u, v) \leq \text{dist}_G(u, w)$, i.e., $\text{dist}_G(u, e) = \text{dist}_G(u, v)$. Let $v' \in \text{Ball}(u)$ and let $e' = (v', w')$ be some edge incident to v' . Since the node v is contained in A_{i+1} we have

$$\text{dist}_G(u, e') \leq \text{dist}_G(u, v') < \text{dist}_G(u, A_{i+1}) \leq \text{dist}_G(u, v) = \text{dist}_G(u, e).$$

Therefore we have $e' \in F$ and thus $v' \in U$. It follows that $\text{Ball}(u) \subseteq U$. Observe that $|U| \leq 2|F|$ and thus $|\text{Ball}(u)| \leq |U| \leq 2|F| = 2m^{(i+1)/k}$ whp. \square

4.2 Properties of Approximate Balls

We now show that the approximate balls and the corresponding distance estimate have the Properties **B1–B4**. We first show that the distance estimates for nodes in the approximate balls have the desired approximation guarantee, although they have been computed in subgraphs of G .

Lemma 4.5 (Property **B1**). *For every pair of nodes u and v such that $v \in \text{Ball}(u)$ we have $\text{dist}_G(u, v) \leq \delta(u, v) \leq \alpha \text{dist}_G(u, v) + \beta$.*

Proof. By Property **A1** we have $\delta(u, v) \geq \text{dist}_{G|R(u)}(u, v)$ and since $G|R(u)$ is a subgraph of G we have $\text{dist}_{G|R(u)}(u, v) \geq \text{dist}_G(u, v)$. Therefore the inequality $\delta(u, v) \geq \text{dist}_G(u, v)$ follows.

Since $v \in \text{Ball}(u)$ we have $\delta(u, v) \leq \alpha D + \beta$. If $\text{dist}_G(u, v) \geq D$, then trivially $\delta(u, v) \leq \alpha D + \beta \leq \alpha \text{dist}_G(u, v) + \beta$. If $\text{dist}_G(u, v) < D$, then there is a path k from u to v in G of weight at most D . This path was also contained in previous versions of G , possibly with smaller weight, and in particular at the time the algorithm has computed $R(u)$, the set of

nodes that are in distance at most D from u in G , for the last time. It follows that k is also contained in $G|R(u)$ and thus $\text{dist}_{G|R(u)}(u, v) = \text{dist}_G(u, v) \leq D$. By Property **A2** we then have $\delta(u, v) \leq \alpha \text{dist}_{G|R(u)} + \beta = \alpha \text{dist}_G + \beta$. \square

We show now that the approximate balls have a certain structural property that either allows us shortcut the path between two nodes or helps us in finding a nearby node of higher priority.

Lemma 4.6 (Property **B2**). *Let $s(\cdot, \cdot)$ be a non-decreasing function such that, for all $x \geq 1$, and $l \geq 1$,*

$$s(x, l) \geq a(a+1)^{l-1} \text{dist}_G(u, v) + ((a+1)^l - 1)b/a,$$

where $a = (1 + \epsilon)\alpha$ and $b = (1 + \epsilon)\beta + 1$. Then for every node u of priority i and every node v such that $s(\text{dist}_G(u, v), p - 1 - i) \leq D$, either (1) $v \in B(u)$ or (2) there is some node v' of priority $j > i$ such that $u \in B(v')$ and $\text{dist}_G(u, v') \leq s(\text{dist}_G(u, v), j - i)$.

Proof. We first define the following series: let $f(1) = a \text{dist}_G(u, v) + b$ and for all $l \geq 1$ let $f(l+1) = f(l) + af(l) + b$. It is easy to verify that for all $l \geq 1$ we have

$$f(l) = \frac{a^2(a+1)^{l-1} \text{dist}_G(u, v) + ((a+1)^l - 1)b}{a} \leq s(\text{dist}_G(u, v), l).$$

Note that since $s(\cdot, \cdot)$ is non-decreasing we have, for all $1 \leq l \leq k - 1 - i$, $f(l) \leq s(\text{dist}_G(u, v), l) \leq s(\text{dist}_G(u, v), p - 1 - i) \leq D$.

If $v \in B(u)$, then we are done. Otherwise, by Lemma 4.2, there is some node v_1 of priority $p_1 \geq i + 1$ such that

$$\text{dist}_G(v_1, u) \leq a \text{dist}_G(u, v) + b = f(1) \leq D.$$

Thus, if $u \in B(v_1)$, then we are done. Otherwise, by Lemma 4.2, there is some node v_2 of priority $p_2 \geq p_1 + 1 \geq i + 2$ such that

$$\text{dist}_G(v_2, v_1) \leq a \text{dist}_G(v_1, u) + b \leq af(1) + b.$$

By the triangle inequality we have

$$\text{dist}_G(v_2, u) \leq \text{dist}_G(v_2, v_1) + \text{dist}_G(v_1, u) \leq af(1) + b + f(1) = f(2) \leq D.$$

We now repeat this argument to obtain nodes v_1, v_2, \dots, v_l of priorities p_1, p_2, \dots, p_l such that $\text{dist}_G(v_l, u) \leq f(l) \leq D$ and $p_l \geq i + l$ until the inequality $\text{dist}_G(v_l, A_{p_l+1}) \geq a \text{dist}_G(v_l, u) + b$ is fulfilled. This happens eventually since $A_k = \emptyset$ and thus $\text{dist}_G(u, A_k) = \infty$. \square

Next, we bound the size of the system of approximate balls we maintain. Here we use the fact that we can easily bound the size of the exact ball $Ball(u)$ for every node u and that by our definitions we ensure that the approximate balls are subsets of the exact balls.

Lemma 4.7 (Size of Approximate Balls (Property **B3**)).

If, for every node u , $\mathcal{B}(u)$ denotes the set of nodes ever contained in $B(u)$, then $\sum_{u \in V} |\mathcal{B}_u| = \tilde{O}(m^{1+1/k} \log_D / \epsilon)$ in expectation.

Proof. We first bound \mathcal{B}_u , the number of nodes ever contained in the approximate ball $B(u)$, of some node u of priority i . Remember that nodes are joining $B(u)$ only when $r(u)$ increases and that

$$r(u) = \min \left(\frac{(1 + \epsilon)^{\lfloor \log_{1+\epsilon}(\delta(u, A_{i+1}) - 1) \rfloor} - \beta}{\alpha}, D \right).$$

Thus, $r(u)$ can only increase if $\lfloor \log_{1+\epsilon}(\delta(u, A_{i+1}) - 1) \rfloor$ increases and the left term in the minimum is at most D . Since $1 + \epsilon \geq 1$ and $\beta \geq 0$ it follows that $r(u)$ increases only $O(\log_{1+\epsilon} D) = O(\log D/\epsilon)$ times. As $B(u) \subseteq \text{Ball}(u)$ by Lemma 4.3, after every increase of $r(u)$ only nodes contained in $\text{Ball}(u)$ can join $B(u)$. By Lemma 4.4 the size of $\text{Ball}(u)$ is $O(m^{(i+1)/k})$ whp. Thus, the number of nodes ever contained in $B(u)$ is $|\mathcal{B}(u)| = O(m^{(i+1)/k} \log D/\epsilon)$ whp.

As the number of nodes of priority i is $\tilde{O}(m/m^{i/k})$ in expectation, the number of nodes ever contained in the approximate balls is $\sum_{u \in V} |\mathcal{B}(u)| \tilde{O}(m^{1+1/k} \log D/\epsilon)$ in expectation. \square

Finally, we analyze the running time of our algorithm for maintaining the approximate balls. Since we use the data structure APPROXSSSP as a black box, the running time of our algorithm depends on the running time of APPROXSSSP.

Lemma 4.8 (Running Time (Property B4)). *The total time needed for maintaining the sets $B(u)$ for all nodes $u \in V$ is*

$$\tilde{O} \left(m^{1+1/k} \log D/\epsilon + \sum_{0 \leq i \leq k-1} T(m_i, n_i) \log D/\epsilon + T(m, n) \right),$$

where, for each $0 \leq i \leq k-1$, $m_i = m^{(i+1)/k}$ and $n_i = \min(m_i, n)$.

Proof. The initialization in Step 1 of the algorithm, where we determine the sets A_0, \dots, A_p takes time $O(pm)$. In Step 2, we run for each $1 \leq i \leq k-1$ an instance of APPROXSSSP with depth D . This takes time $kT(m, n)$. Step 3, where we maintain for every node u of priority i the approximate ball and corresponding distance estimates can be analyzed as follows. Remember that every time $r(u)$ increases we first compute $R(u)$, the set of nodes that are in distance at most $r(u)$ from u . Using a suitable implementation of Dijkstra's algorithm, this takes time $O(|E[R(u)]| \log n)$, where $E[R(u)]$ is the set of edges incident to $R(u)$. By Lemma 4.3 we have $R(u) \subseteq \text{Ball}(u)$ and by Lemma 4.4 we have $|\text{Ball}(u)| = O(m^{(i+1)/k})$ whp. Thus, computing $R(u)$ takes time $\tilde{O}(m^{(i+1)/k})$ whp. We then maintain an instance of APPROXSSSP up to depth D on $G|R(u)$, the subgraph of G induced by $R(u)$. Note that $G|R(u)$ has $m_i = m^{(i+1)/k}$ edges and $n_i = \min(m_i, n)$ nodes and therefore this takes time $T(m_i, n_i)$. As $r(u)$ increases $O(\log D/\epsilon)$ times and the number of nodes of priority i is $\tilde{O}(m/m^{i/k})$ in expectation, maintaining $B(u)$ for all nodes u of priority i (and the corresponding distance estimates) takes time

$$\sum_{0 \leq i \leq k-1} \tilde{O} \left(\frac{m}{m^{i/k}} \log D/\epsilon \cdot (m^{(i+1)/k} + T(m_i, n_i)) \right) = \tilde{O} \left(m^{1+1/k} \log D/\epsilon + \sum_{0 \leq i \leq k-1} m^{1-i/k} \cdot T(m_i, n_i) \log D/\epsilon \right).$$

Since $k \leq \log n$, the claimed running time follows. \square

5 From Approximate Balls to Approximate SSSP

In the following we show how to maintain an approximate shortest paths tree if we already have an algorithm for maintaining approximate balls. Our main tool in this reduction is a hop set that we define from the approximate balls. We will add the “shortcut” edges of the hop set to the graph and scale down the edge weights, maintaining the approximate shortest paths with a monotone ES-tree. Formally, we prove the following statement in this section.

Proposition 5.1. *Assume there is a decremental algorithm APPROXBALLS for maintaining approximate balls with the following properties, using fixed values $a \geq \alpha \geq 1$, $b \geq \beta \geq 0$, and $\hat{D} \geq 1$. Given a weighted graph $G = (V, E)$ undergoing edge deletions and edge weight increases and a parameters $2 \leq k \leq \log n$, it assigns to every node $u \in V$ a number from 0 to $k - 1$, called the priority of u , and maintains for every node $u \in V$ a set of nodes $B(u)$ and, for every node $v \in B(u)$, a distance estimate $\hat{\delta}(u, v)$ such that:*

B1 *For every node u and every node $v \in B(u)$ we have $\text{dist}_G(u, v) \leq \hat{\delta}(u, v) \leq \alpha \text{dist}_G(u, v) + \beta$.*

B2 *There is a function $s(\cdot, \cdot)$ such that, for all $x \geq 1$, $s(x, 1) = ax + b$ for some $a \geq \alpha$ and $b \geq \beta$ and, for all $l \geq 1$,*

$$s(x, l + 1) \leq \frac{a(\alpha + 1 + \epsilon)(\alpha s(x, l) + \beta) + \beta}{\epsilon} + b.$$

guaranteeing the following: For every node u of priority i and every node v such that $s(\text{dist}_G(u, v), p - 1 - i) \leq \hat{D}$, either (1) $v \in B(u)$ or (2) there exist some node $v' \in V$ of priority $j > i$ such that $u \in B(v')$ and $\text{dist}_G(u, v') \leq s(\text{dist}_G(u, v'), j - i)$.

B3 *After every update in G , APPROXBALLS returns all pairs of nodes u and v such that v joins $B(u)$, v leaves $B(u)$, or $\hat{\delta}(u, v)$ changes.*

For every node u , let $\mathcal{B}(u)$ denote the set of nodes ever contained in $B(u)$ and let $t(m, n, k)$ denote the total update time of APPROXBALLS.

Then there is an approximate SSSP data structure APPROXSSSP with the following properties: Given a weighted graph G undergoing edge deletions and edge weight increases, a fixed source node s , and parameters p , Δ , D , and ϵ such that

$$2 \leq p \leq \frac{\sqrt{\log n}}{\sqrt{\log\left(\frac{4a^3}{\epsilon}\right)}},$$

$\Delta \geq b$, $n^{1/p}\Delta \leq \hat{D}$, $D \geq \Delta$ and $0 < \epsilon \leq 1$, it maintains a distance estimate $\delta(s, v)$ for every node $v \in V$ such that:

A1 $\delta(s, v) \geq \text{dist}_G(s, v)$

A2 If $\text{dist}_G(s, v) \leq D$, then $\delta(s, v) \leq (\alpha + \epsilon) \text{dist}_G(s, v) + \epsilon n^{1/p} \Delta$

A3 The total update time of APPROXSSSP is

$$T(m, n, \Delta, D, \epsilon) = \tilde{O}\left((\alpha D / \Delta + n^{1/p}) \sum_{u \in V} (m + |\mathcal{B}(u)|) / \epsilon + t(m, n, p)\right).$$

A4 After every update in G , APPROXSSSP returns, for every node v such that $\delta(s, v)$ has changed, v together with the new value of $\delta(s, v)$.

We assume without loss of generality that the distance estimate maintained by APPROXBALLS is non-decreasing. If APPROXBALLS ever reports a decrease we can ignore it because Property **B1** will still hold as distances in G are non-decreasing under edge deletions and edge weight increases.

5.1 Algorithm Description

The algorithm APPROXSSSP maintains the set of edges $F = \{(u, v) \in V^2 \mid v \in B(u)\}$ such that each edge $(u, v) \in F$ has weight $w_F(u, v) = \delta(u, v)$. We update F every time in APPROXBALLS a node joins or leaves an approximate ball or if the distance estimate $\delta(u, v)$ increases for some pair of nodes u and v . By Property **B3** this information is returned by APPROXBALLS after every update in G . Thus the set of edges F undergoes insertions, deletions, and weight increases.

In the following we will define a shortcut graph H'' with scaled-down edge weights and our algorithm APPROXSSSP will simply run a monotone ES-tree [HKN13] from s in H'' . Note that the monotone ES-tree trivially has properties **A1** and **A4**. We denote the weight of an edge (u, v) in G by $w_G(u, v)$ and define H as a graph that has the same nodes as G and contains all edges of G and F that have weight at most $D + n^{1/p}\Delta$. We set the weight of every edge (u, v) in G to $w_H = \min(w_G(u, v), w_F(u, v))$. We set

$$\varphi = \frac{\epsilon\Delta}{p+1}$$

and define H' as the graph that has the same nodes and edges as H and in which every edge (u, v) has weight

$$w_{H'}(u, v) = \left\lceil \frac{w_H(u, v)}{\varphi} \right\rceil \varphi,$$

i.e., we round every edge weight to the next multiple of φ . Furthermore, we define H'' as the graph that has the same nodes and edges as H' and in which every edge (u, v) has weight

$$w_{H''}(u, v) = \frac{w_{H'}(u, v)}{\varphi} = \left\lceil \frac{w_H(u, v)}{\varphi} \right\rceil,$$

i.e., we scale down every edge weight by a factor of $1/\varphi$. We maintain a monotone ES-tree with maximum level

$$L = (\alpha + 2\epsilon)D/\varphi + (p+1)n^{1/p}$$

from s and denote the level of a node v in this tree by $\ell(v)$. For every node v our algorithm returns the distance estimate $\delta(s, v) = \ell(v) \cdot \varphi$. Note that the graph H'' has integer edge weights and, as F might undergo insertions, deletions, and edge weight increases, the same type of updates might occur in H'' . Furthermore, observe that the rounding guarantees that

$$w_H(u, v) \leq w_{H'}(u, v) \leq w_H(u, v) + \varphi$$

for every edge (u, v) of H' .

5.2 Running Time Analysis

We first provide the running time analysis. We run the algorithm in a graph in which we scale down the edge weights by a factor of φ . This makes the algorithm efficient.

Lemma 5.2 (Running Time (Property **A3**)). *The expected total update time of a monotone ES-tree with maximum level $L = (\alpha + 2\epsilon)D/\varphi + (p + 1)n^{1/p}$ on H'' is*

$$\tilde{O}\left(\left(\alpha D/\Delta + n^{1/p}\right)\left(m + \sum_{u \in V} |\mathcal{B}(u)|\right)/\epsilon\right).$$

Proof. By Lemma 2.2 the total time needed for maintaining the monotone ES-tree with maximum level L on H'' is

$$O(\mathcal{E}(H'') \cdot L + \mathcal{W}(H''))$$

where $\mathcal{E}(H'')$ is the number of edges ever contained in H'' and $\mathcal{W}(H'')$ is the number of updates (i.e., edge deletions, edge weight increases, and edge insertions) on H'' .

Remember that $\varphi = \epsilon\Delta/(p + 1)$. Since $\epsilon \leq 1$ and $p \leq \log n$ we have $L = \tilde{O}(\alpha D/(\epsilon\Delta) + n^{1/p})$. We now bound $\mathcal{E}(H'')$ and $\mathcal{W}(H'')$. Note that at any time H'' has the same edges as H and each edge of H either is also an edge in G , which contains m edges, or is an edge from F . As F is defined via the approximate balls (i.e., $(u, v) \in F$ if and only if $v \in \mathcal{B}(u)$), the number of edges ever contained in F is at most $\sum_{u \in V} |\mathcal{B}(u)|$, the total number of nodes ever contained in the approximate balls. It follows that $\mathcal{E}(H'') = \tilde{O}(m + \sum_{u \in V} |\mathcal{B}(u)|)$ in expectation. Note that every edge contained in H'' can be inserted and deleted at most once and its weight can increase at most $(D + n^{1/p}\Delta)/\varphi$ times as we have limited the maximum edge weight in H to $D + n^{1/p}\Delta$. Note that $(D + n^{1/p}\Delta)/\varphi = (D + n^{1/p}\Delta)(p + 1)/(\epsilon\Delta) = \tilde{O}((D/\Delta + n^{1/p})/\epsilon)$. Therefore we have

$$\mathcal{W}(H'') \leq 2\mathcal{E}(H'') + \mathcal{E}(H'') \cdot (D + n^{1/p}\Delta)/\varphi = \tilde{O}\left((m + \sum_{u \in V} |\mathcal{B}(u)|) \cdot (D/\Delta + n^{1/p})/\epsilon\right).$$

We conclude that

$$\mathcal{E}(H'') \cdot L + \mathcal{W}(H'') = \tilde{O}\left((\alpha D/\Delta + n^{1/p})(m + \sum_{u \in V} |\mathcal{B}(u)|)/\epsilon\right)$$

and thus the claimed running time follows. \square

5.3 Definitions of Values for Approximation Guarantee

Before we analyze the approximation guarantee we define the following values. We set

$$r_0 = \Delta$$

and for every $0 \leq i \leq p - 1$ we set

$$\begin{aligned} s_i &= ar_i + b, \\ w_i &= \alpha s_i + \beta, \text{ and} \\ r_i &= \frac{(\alpha + 1 + \epsilon) \sum_{0 \leq j \leq i-1} w_j + \beta}{\epsilon} \quad (\text{if } i \geq 1). \end{aligned}$$

Finally, we set

$$\gamma_{p-1} = \beta$$

and, for every $0 \leq i \leq p-2$,

$$\gamma_i = \gamma_{i+1} + (\alpha + 1 + \epsilon)w_i = (\alpha + 1 + \epsilon) \sum_{i \leq j \leq p-2} w_j + \beta.$$

We also set

$$\gamma = \gamma_0 + 2\epsilon\Delta.$$

Lemma 5.3. For all $0 \leq i \leq p-1$, $\epsilon r_i = \gamma_0 - \gamma_i + \beta$

Proof. First, observe that for all $0 \leq i \leq p-1$ we have

$$\gamma_i = (\alpha + 1 + \epsilon) \sum_{i \leq j \leq p-2} w_j + \beta.$$

Thus, for all $0 \leq i \leq p-1$, we get

$$\gamma_0 - \gamma_i + \beta = (\alpha + 1 + \epsilon) \sum_{0 \leq j \leq p-2} w_j - (\alpha + 1 + \epsilon) \sum_{i \leq j \leq p-2} w_j + \beta = (\alpha + 1 + \epsilon) \sum_{0 \leq j \leq i-1} w_j + \beta = \epsilon r_i. \quad \square$$

Lemma 5.4. $(4a^3/\epsilon)^p = n^{1/p}$

Proof. Remember that we have

$$p \leq \frac{\sqrt{\log n}}{\sqrt{\log\left(\frac{4a^3}{\epsilon}\right)}}.$$

We only need to simplify both expressions as follows:

$$\begin{aligned} n^{1/p} &= 2^{1/p \cdot \log n} \geq 2^{\frac{\sqrt{\log\left(\frac{4a^3}{\epsilon}\right)} \cdot \log n}{\sqrt{\log n}}} = 2^{\sqrt{\log\left(\frac{4a^3}{\epsilon}\right)} \cdot \sqrt{\log n}} \\ \left(\frac{4a^3}{\epsilon}\right)^p &= 2^{p \cdot \log\left(\frac{4a^3}{\epsilon}\right)} \leq 2^{\frac{\sqrt{\log n} \cdot \log\left(\frac{4a^3}{\epsilon}\right)}{\sqrt{\log\left(\frac{4a^3}{\epsilon}\right)}}} = 2^{\sqrt{\log n} \cdot \sqrt{\log\left(\frac{4a^3}{\epsilon}\right)}}. \end{aligned} \quad \square$$

Lemma 5.5. For all $0 \leq i \leq p-1$ we have

$$r_i \leq \frac{3 \cdot 4^{i-1} a^{3i} \Delta + (9 \cdot 4^{i-1} - 2) a^{3i-1} b}{\epsilon^i}$$

and

$$\sum_{0 \leq j \leq i} w_j \leq \frac{4^i a^{3i+2} \Delta + (3 \cdot 4^i - 1) a^{3i+1} b}{\epsilon^i}.$$

Proof. Remember that $\epsilon \leq 1 \leq \alpha \leq a$. Now observe that for all $1 \leq i \leq p-1$ we have

$$r_i = \frac{(\alpha + 1 + \epsilon) \sum_{0 \leq j \leq i-1} w_j + \beta}{\epsilon} \leq \frac{(a + 1 + \epsilon) \sum_{0 \leq j \leq i-1} w_j + b}{\epsilon} \leq \frac{3a \sum_{0 \leq j \leq i-1} w_j + b}{\epsilon}$$

and for all $0 \leq i \leq p-1$ we have

$$w_i = \alpha s_i + \beta \leq a s_i + b = a(ar_i + b) + b = a^2 r_i + ab + b \leq a^2 r_i + 2ab.$$

We now prove the inequalities induction on i . We begin with the base case $i = 0$ where $r_0 = \Delta$ and

$$\sum_{0 \leq j \leq 0} w_j = w_0 \leq a^2 r_0 + 2ab = a^2 \Delta + 2ab = \frac{4^0 a^{3 \cdot 0 + 2} \Delta + (3 \cdot 4^0 - 1) a^{3 \cdot 0 + 1} b}{\epsilon^0}.$$

In the induction step we assume that $i \geq 1$:

$$\begin{aligned} r_i &\leq \frac{3a \sum_{0 \leq j \leq i-1} w_j + b}{\epsilon} \\ &\leq \frac{3a(4^{i-1} a^{3(i-1)+2} \Delta + (3 \cdot 4^{i-1} - 1) a^{3(i-1)+1} b) + b}{\epsilon^i} \\ &= \frac{3 \cdot 4^{i-1} a^{3i} \Delta + (9 \cdot 4^{i-1} - 2) a^{3i-1} b}{\epsilon^i} \end{aligned}$$

$$\begin{aligned} \sum_{0 \leq j \leq i} w_j &= \sum_{0 \leq j \leq i-1} w_j + w_i \\ &\leq \sum_{0 \leq j \leq i-1} w_j + a^2 r_i + 2ab \\ &\leq \frac{4^{i-1} a^{3(i-1)+2} \Delta + (3 \cdot 4^{i-1} - 1) a^{3(i-1)+1} b}{\epsilon^{i-1}} + a^2 r_i + 2ab \\ &\leq \frac{4^{i-1} a^{3(i-1)+2} \Delta + (3 \cdot 4^{i-1} - 1) a^{3(i-1)+1} b}{\epsilon^{i-1}} + \frac{3 \cdot 4^{i-1} a^{3i+2} \Delta + (9 \cdot 4^{i-1} - 2) a^{3i+1} b}{\epsilon^i} + 2ab \\ &\leq \frac{4^{i-1} a^{3(i-1)+2} \Delta + (3 \cdot 4^{i-1} - 1) a^{3(i-1)+1} b + 3 \cdot 4^{i-1} a^{3i+2} \Delta + (9 \cdot 4^{i-1} - 2) a^{3i+1} b + 2ab}{\epsilon^i} \\ &\leq \frac{4^{i-1} a^{3i+2} \Delta + (3 \cdot 4^{i-1} - 1) a^{3i+1} b + 3 \cdot 4^{i-1} a^{3i+2} \Delta + (9 \cdot 4^{i-1} - 2) a^{3i+1} b + 2a^{3i+1} b}{\epsilon^i} \\ &= \frac{(1+3) \cdot 4^{i-1} a^{3i+2} \Delta + (3 \cdot 4^{i-1} - 1 + 9 \cdot 4^{i-1} - 2 + 2) a^{3i+1} b}{\epsilon^i} \\ &= \frac{4^i a^{3i+2} \Delta + (3 \cdot 4^i - 1) a^{3i+1} b}{\epsilon^i}. \quad \square \end{aligned}$$

Lemma 5.6. $a\gamma + b \leq \epsilon n^{1/p} \Delta$.

Proof. Remember that we have $\epsilon \leq 1 \leq \alpha \leq a$ and $\beta \leq b \leq \Delta$. By Lemma 5.5 we have

$$\sum_{0 \leq j \leq p-2} w_j \leq \frac{4^{p-2} a^{3(p-2)+2} \Delta + (3 \cdot 4^{p-2} - 1) a^{3(p-2)+1} b}{\epsilon^{p-2}} \leq \frac{4^{p-2} a^{3p-2} \Delta + 3 \cdot 4^{p-2} a^{3p-2} \Delta}{\epsilon^{p-1}}$$

We now get:

$$\begin{aligned}
\frac{a\gamma + b}{\epsilon} &= \frac{a\gamma_0 + 2\epsilon a\Delta + \beta}{\epsilon} \\
&= \frac{a(\alpha + 1 + \epsilon) \sum_{0 \leq j \leq p-2} w_j + \alpha\beta + 2\epsilon a\Delta + \beta}{\epsilon} \\
&\leq \frac{a(a + 1 + \epsilon) \sum_{0 \leq j \leq p-2} w_j + a\Delta + 2a\Delta + \Delta}{\epsilon} \\
&\leq \frac{3a^2 \sum_{0 \leq j \leq p-2} w_j + 4a\Delta}{\epsilon} \\
&\leq \frac{3 \cdot 4^{p-2} a^{3p} \Delta + 9 \cdot 4^{p-2} a^{3p} \Delta + 4a\Delta}{\epsilon^p} \\
&\leq \frac{3 \cdot 4^{p-2} a^{3p} \Delta + 9 \cdot 4^{p-2} a^{3p} \Delta + 4^{p-2} a^{3p} \Delta}{\epsilon^p} \\
&= \frac{(3 + 9 + 1) \cdot 4^{p-2} a^{3p} \Delta}{\epsilon^p} \\
&\leq \frac{4^p a^{3p} \Delta}{\epsilon^p} \\
&= (4a^3/\epsilon)^p \Delta \leq n^{1/p} \Delta.
\end{aligned}$$

The last inequality follows from Lemma 5.4. \square

Lemma 5.7. $ar_{p-1} + b \leq n^{1/p} \Delta$.

Proof. By the definitions of r_{p-1} and γ_0 we have $r_{p-1} = \gamma_0/\epsilon$. Since $\gamma_0 \leq \gamma$ and $a\gamma + b \leq \epsilon n^{1/p} \Delta$ by Lemma 5.6, we have

$$ar_{p-1} + b = a \frac{\gamma_0}{\epsilon} + b \leq \frac{a\gamma_0 + b}{\epsilon} \leq \frac{a\gamma + b}{\epsilon} \leq n^{1/p} \Delta. \quad \square$$

Lemma 5.8. For all $0 \leq i < j \leq p-1$, $s(r_i, j-i) \leq s(r_{j-1}, 1)$

Proof. Fix some $0 \leq i \leq p-2$. The proof is by induction on j . In the base case $j = i+1$ the claim holds trivially. Consider now the induction step where we assume that the inequality holds for $j \geq i+1$ and have to show that it also holds for $j+1$. First, observe that

$$r_j = \frac{(\alpha + 1 + \epsilon) \sum_{0 \leq j' \leq j-1} w_{j'} + \beta}{\epsilon} \geq \frac{(\alpha + 1 + \epsilon) w_{j-1} + \beta}{\epsilon} = \frac{(\alpha + 1 + \epsilon)(\alpha s_{j-1} + \beta) + \beta}{\epsilon}$$

and thus

$$s(r_j, 1) = ar_j + b \geq \frac{a(\alpha + 1 + \epsilon)(\alpha s_{j-1} + \beta) + \beta}{\epsilon} + b = \frac{a(\alpha + 1 + \epsilon)(\alpha s(r_{j-1}, 1) + \beta) + \beta}{\epsilon} + b$$

By the Property **B2** we have

$$s(r_i, j-1) \leq \frac{a(\alpha + 1 + \epsilon)(\alpha s(r_i, j-i-1) + \beta) + \beta}{\epsilon} + b$$

and by the induction hypothesis we have $s(r_i, j-i-1) \leq s(r_j, j-1)$. Therefore it follows that

$$s(r_i, j-1) \leq \frac{a(\alpha + 1 + \epsilon)(\alpha s(r_j, j-1) + \beta) + \beta}{\epsilon} + b \leq s(r_j, 1). \quad \square$$

5.4 Analysis of Approximation Guarantee

We now analyze the approximation error of a monotone ES-tree maintained on H'' . This approximation error consists of two parts. The first part is an approximation error that comes from the fact that the monotone ES-tree only considers paths from s with a relatively small number of edges and therefore has to use edges from the hop set F . The second part is the approximation error we get from rounding the edge weights. We first give a formula for the approximation error that depends on the priority of the nodes and their distance to the root of the monotone ES-tree.

Before we give the proof we review a few properties of the monotone ES-tree (see [HKN13] for the full algorithm). Similar to the classic ES-tree, the monotone ES-tree with root s maintains a level $\ell(v)$ for every node v . The monotone ES-tree is initialized by computing a shortest paths tree up to depth L from s in H'' and thus, initially, $\ell(v) = \text{dist}_{H''}(s, v)$. A single deletion or edge weight increase in G might result in a sequence of deletions, weight increases and insertions in F , and thus H'' . The monotone ES-tree first processes the insertions and then the deletions and edge weight increases. It handles deletions and edge weights increases in the same way as the classic ES-tree. The procedure for handling the insertion of an edge (u, v) is trivial: it only stores the new edge and in particular does *not* change $\ell(u)$ or $\ell(v)$. Once the level $\ell(u)$ of a node u exceeds the maximum level L , we set $\ell(u) = \infty$. For completeness we list the pseudocode of the monotone ES-tree in Algorithm 1.

For the analysis of the monotone ES-tree we will use the following notions. We say that an edge (u, v) is *stretched* if $\ell(u) > \ell(v) + w_{H''}(u, v)$. We say that a node u is *stretched* if it is incident to an edge (u, v) that is stretched. Note that for a node u that is not stretched we have $\ell(u) \leq \ell(v) + w_{H''}(u, v)$ for every edge (u, v) contained in H'' . In our proof we will use the following properties of the monotone ES-tree.

Observation 5.9 ([HKN13]). *The following holds for the monotone ES-tree:*

- (1) *The level of a node never decreases.*
- (2) *An edge can only become stretched when it is inserted.*
- (3) *As long as a node x is stretched, its level does not change.*
- (4) *For every tree edge (u, v) (where v is the parent of u), $\ell(u) \geq \ell(v) + w(u, v)$.*

A second prerequisite from [HKN13] tells us when we may apply a variant of the triangle inequality to argue about the levels of nodes.

Lemma 5.10 ([HKN13]). *Let (u, v) be an edge of H'' such that $\ell(v) + w_{H''}(u, v) \leq L$. If (u, v) is not stretched and after the previous update in G the level of u was less than ∞ , then for the current level of u we have $\ell(u) \leq \ell(v) + w_{H''}(u, v)$.*

Note that condition (2) simply captures the property of the monotone ES-tree that once the level of a node exceeds L it is set to ∞ and will never be decreased anymore. At the initialization (i.e., before the first update in H''), conditions (1) and (2) are fulfilled automatically.

Algorithm 1: Monotone ES-tree

```
// Internal data structures:
//  $N(u)$ : for every node  $u$  a heap  $N(u)$  whose intended use is to store
// for every neighbor  $v$  of  $u$  in the current graph the value of
//  $\ell(v) + w_{H''}(u, v)$ 
//  $Q$ : global heap whose intended use is to store nodes whose levels
// might need to be updated

1 Procedure INITIALIZE()
2   Compute shortest paths tree from  $s$  in  $H''$  up to depth  $L$ 
3   foreach  $u \in V$  do
4     Set  $\ell(u) = \text{dist}_{H''}(s, u)$ 
5     for every edge  $(u, v)$  in  $H''$  do insert  $v$  into heap  $N(u)$  of  $u$  with key
6        $\ell(v) + w_{H''}(u, v)$ 

6 Procedure DELETE( $u, v$ )
7   INCREASE( $u, v, \infty$ )

8 Procedure INCREASE( $u, v, w(u, v)$ )
9   // Increase weight of edge  $(u, v)$  to  $w(u, v)$ 
10  Insert  $u$  and  $v$  into heap  $Q$  with keys  $\ell(u)$  and  $\ell(v)$  respectively
11  Update key of  $v$  in heap  $N(u)$  to  $\ell(v) + w(u, v)$  and key of  $u$  in heap  $N(v)$  to
12     $\ell(u) + w(u, v)$ 
13  UPDATELEVELS()

12 Procedure INSERT( $u, v, w(u, v)$ )
13  // Increase edge  $(u, v)$  of weight  $w(u, v)$ 
14  Insert  $v$  into heap  $N(u)$  with key  $\ell(v) + w(u, v)$  and  $u$  into heap  $N(v)$  with key
15     $\ell(u) + w_{H''}(u, v)$ 

14 Procedure UPDATELEVELS()
15  while heap  $Q$  is not empty do
16    Take node  $u$  with minimum key  $\ell(u)$  from heap  $Q$  and remove it from  $Q$ 
17     $\ell'(u) \leftarrow \min_v (\ell(v) + w_{H''}(u, v))$ 
18    //  $\min_v (\ell(v) + w_{H''}(u, v))$  can be retrieved from the heap  $N(u)$ .
19     $\arg \min_v (\ell(v) + w_{H''}(u, v))$  is  $u$ 's parent in the ES-tree.
20    if  $\ell'(u) > \ell(u)$  then
21       $\ell(u) \leftarrow \ell'(u)$ 
22      if  $\ell'(u) > L$  then  $\ell(u) \leftarrow \infty$ 
23      foreach neighbor  $v$  of  $u$  do
24        update key of  $u$  in heap  $N(v)$  to  $\ell(u) + w_{H''}(u, v)$ 
25        insert  $v$  into heap  $Q$  with key  $\ell(v)$  if  $Q$  does not already contain  $v$ 
```

To count the additive error from rounding the edge weights, we define, for every node u and every $0 \leq i \leq p-1$, the function $h(u, i)$ as follows:

$$h(u, i) = \begin{cases} 0 & \text{if } u = s \\ (p+1) \left\lceil \frac{\max(\text{dist}_G(u, s) - r_i, 0)}{\Delta} \right\rceil + p + 1 - i & \text{otherwise} \end{cases}.$$

The intuition is that $h(u, i)$ bounds the number of hops from u to s , i.e., the number of edges required to go from u to s while at the same time providing the desired approximation guarantee. The approximation guarantee can now formally be stated as follows

Lemma 5.11 (Approximation Guarantee). *For every node u of priority i with $\text{dist}_G(u, s) \leq D + \sum_{0 \leq i' \leq i-1} s_{i'}$ we have*

$$\delta(s, u) \leq (\alpha + \epsilon) \text{dist}_G(u, s) + \gamma_i + h(u, i) \cdot \varphi.$$

Once we have proved this lemma, the desired bound on the approximation error (Property **A2**) follows easily because $h(u, i) \cdot \varphi \leq \epsilon \text{dist}_G(u, v) + 2\epsilon\Delta$ (as we show below) and thus

$$\begin{aligned} \delta(s, u) &\leq (\alpha + \epsilon) \text{dist}_G(u, s) + \gamma_i + h(u, i) \cdot \varphi \\ &\leq (\alpha + \epsilon) \text{dist}_G(u, s) + \gamma_0 + h(u, i) \cdot \varphi \\ &\leq (\alpha + \epsilon) \text{dist}_G(u, s) + \gamma_0 + \epsilon \text{dist}_G(u, v) + 2\epsilon\Delta \\ &= (\alpha + 2\epsilon) \text{dist}_G(u, s) + \gamma. \end{aligned}$$

Lemma 5.12. *For every node u and every $0 \leq i \leq p-1$,*

$$h(u, i) \cdot \varphi \leq \epsilon \text{dist}_G(u, s) + 2\epsilon\Delta$$

Proof.

$$\begin{aligned} \left((p+1) \left\lceil \frac{\max(\text{dist}_G(u, s) - r_i, 0)}{\Delta} \right\rceil + p + 1 - i \right) \varphi &\leq \left((p+1) \left\lceil \frac{\text{dist}_G(u, s)}{\Delta} \right\rceil + p + 1 \right) \varphi \\ &\leq \left((p+1) \left(\frac{\text{dist}_G(u, s)}{\Delta} + 1 \right) + p + 1 \right) \varphi \\ &= \left(\frac{(p+1) \text{dist}_G(u, s)}{\Delta} + 2(p+1) \right) \varphi \\ &= \left(\frac{(p+1) \text{dist}_G(u, s)}{\Delta} + 2(p+1) \right) \cdot \frac{\epsilon\Delta}{p+1} \\ &= \epsilon \text{dist}_G(u, s) + 2\epsilon\Delta. \quad \square \end{aligned}$$

Proof of Lemma 5.11. The proof is by double induction first on the number of updates in G and second on $h(u, i)$. Let u be a node of priority i such that $\text{dist}_G(u, s) \leq D + \sum_{0 \leq i' \leq i-1} s_{i'}$. Remember that $\delta(u, s) = \ell(u) \cdot \varphi$, where $\ell(u)$ is the level of u in the monotone ES-tree of s . We know that after the last previous in G the distance estimate gave an approximation of the true distance in G . Since distances in G are non-decreasing it must have been the case that the level of u was less than ∞ after the previous.

If $u = s$, the claim is trivially true because $\ell(s) = 0$. Assume that $u \neq s$. If u is stretched in the monotone ES-tree, then the level of u has not changed since the previous deletion in G and thus the claim is true by induction. If u is not stretched, then $\ell(u) \leq \ell(v) + w_{H''}(u, v)$ for every edge (u, v) in H'' . Define the nodes v and x as follows. If $\text{dist}_G(u, s) \leq r_i$, then $v = s$. If $\text{dist}_G(u, s) > r_i$, then consider a shortest path π from u to s in G and let v be the furthest node from u on π such that $\text{dist}_G(u, v) \leq r_i$ (which implies $\text{dist}_G(v, s) \geq \text{dist}_G(u, s) - r_i$). Furthermore let x be the neighbor of v on the shortest path π that is closer to s than v . Note that $\text{dist}_G(u, x) \geq r_i$ (and thus $\text{dist}_G(x, s) \leq \text{dist}_G(u, s) - r_i$) and in particular G contains the edge (v, x) . Note that (v, x) is also contained in H (and thus in H' and H'') because for $\text{dist}_G(u, s) \leq D + \sum_{0 \leq i' \leq i-1} s_{i'}$ to hold it has to be the case that $w_G(v, x) \leq D + \sum_{0 \leq i' \leq i-1} s_{i'}$. Note that $\sum_{0 \leq i' \leq i-1} s_{i'} \leq \sum_{0 \leq i' \leq i-1} w_{i'} \leq r_{p-1} \leq n^{1/p} \Delta$ by Lemma 5.7. Thus, $w_G(v, x) \leq D + n^{1/p} \Delta$, which by the definition of H means that the edge (v, x) is contained in H .

Note that $s(\text{dist}_G(u, v), p-1-i) \leq s(r_i, p-1-i)$ since the function $s(\cdot, \cdot)$ is non-decreasing. By Lemma 5.8 we have $s(r_i, p-1-i) \leq s(r_{p-2}, 1)$ and by the definition of $s(\cdot, 1)$ and Lemma 5.7 we have $s(r_{p-2}, 1) = ar_{p-2} + b \leq ar_{p-1} + b \leq n^{1/p} \Delta \leq \hat{D}$. Thus, by Property B2 we know that either $v \in B(u)$ or there is a node v' of priority $j' > i$ such that $\text{dist}_G(u, v') \leq s(\text{dist}_G(u, v), j-i)$. Note that in the first case the set of edges F contains the edge (u, v) and in the second case it contains the edge (u, v') .

Case 1: $v \in B(u)$

If $v \in B(u)$, then F contains an edge (u, v) such that

$$w_F(u, v) = \hat{\delta}(u, v) \leq \alpha \text{dist}_G(u, v) + \beta \quad (2)$$

Since $\text{dist}_G(u, v) \leq r_i$ we have $w_F(u, v) \leq \alpha r_i + \beta \leq \alpha r_{p-1} + \beta \leq n^{1/p} \Delta$, where the last inequality holds by Lemma 5.7. Thus, (u, v) is contained in H and thus also in H' and H'' .

If $\text{dist}_G(u, s) \leq r_i$, then we have $v = s$. First observe that by the definition of H'' we have $w_{H''}(u, s) = w_{H'}(u, s)/\varphi$. Furthermore the rounding of the edge weights in H' guarantees that $w_{H'}(u, s) \leq w_H(u, s) + \varphi$. We therefore get

$$\begin{aligned} w_{H''}(u, s) &\leq \frac{w_F(u, s) + \varphi}{\varphi} \\ &\leq \frac{\alpha \text{dist}_G(u, s) + \beta + \varphi}{\varphi} \\ &\leq \frac{\alpha \left(D + \sum_{0 \leq i' \leq i-1} s_{i'} \right) + \beta + \varphi}{\varphi} \\ &\leq \frac{\alpha D + (\alpha + 1 + \epsilon) \sum_{0 \leq i' \leq p-2} w_{i'} + \beta + \varphi}{\varphi} \\ &= \frac{\alpha D + \gamma_0 + \varphi}{\varphi} \\ &= \frac{\alpha D + \gamma_0 + \frac{\epsilon \Delta}{p+1}}{\varphi} \\ &\leq \frac{\alpha D + \gamma_0 + 2\epsilon \Delta}{\varphi} \\ &= \frac{\alpha D + \gamma}{\varphi} \leq \frac{\alpha D + \epsilon n^{1/p} \Delta}{\varphi} \leq \frac{(\alpha + 2\epsilon) D}{\varphi} + (p+1)n^{1/p} = L. \end{aligned}$$

Here we have used the inequality $\gamma \leq \epsilon n^{1/p} \Delta$ from Lemma 5.6. Since the maximum level in the monotone ES-tree is L and u is not stretched, it follows from Lemma 5.10 that $\ell(u) \leq \ell(s) + w_{H''}(u, s) = w_{H''}(u, s)$. Together with the observation that $h(u, i) \geq 1$ since $u \neq s$ and $\beta \leq \gamma_0$ we therefore get

$$\begin{aligned} \delta(s, u) &= \ell(u) \cdot \varphi \leq w_{H''}(u, s) \cdot \varphi \leq \alpha \operatorname{dist}_G(u, s) + \beta + \varphi \\ &\leq \alpha \operatorname{dist}_G(u, s) + \beta + h(u, i) \cdot \varphi \leq (\alpha + \epsilon) \operatorname{dist}_G(u, s) + \gamma_0 + h(u, i) \cdot \varphi. \end{aligned}$$

Consider now the case $\operatorname{dist}_G(u, s) > r_i$. Let j denote the priority of x . We first prove the following inequality, which will allow us among other things to use the induction hypothesis on x .

Claim 5.13. *If $\operatorname{dist}_G(u, s) > r_i$, then $h(x, j) + 2 \leq h(u, i)$.*

Proof. Remember that $i \leq p - 1$. The assumption $\operatorname{dist}_G(u, s) > r_i$ implies that $\operatorname{dist}_G(x, s) \leq \operatorname{dist}_G(u, s) - r_i$. If $\operatorname{dist}_G(x, s) < r_j$, we have

$$\begin{aligned} h(x, j) + 2 &\leq p + 1 - j + 2 \leq p + 1 + 2 \leq p + 1 + p + 1 - i \\ &\leq (p + 1) \left\lceil \frac{\operatorname{dist}_G(u, s) - r_i}{\Delta} \right\rceil + p + 1 - i = h(u, i). \end{aligned}$$

Here we use the inequality $\lceil (\operatorname{dist}_G(u, s) - r_j) / \Delta \rceil \geq 1$ which follows from the assumption $\operatorname{dist}_G(u, s) > r_i$.

If $\operatorname{dist}_G(x, s) \geq r_j$, then, using $r_j \geq r_0 \geq \Delta$, we get

$$\begin{aligned} h(x, j) + 2 &= (p + 1) \left\lceil \frac{\operatorname{dist}_G(x, s) - r_j}{\Delta} \right\rceil + p + 1 - j + 2 \\ &\leq (p + 1) \left\lceil \frac{\operatorname{dist}_G(x, s) - \Delta}{\Delta} \right\rceil + p + 1 + 2 \\ &= (p + 1) \left\lceil \frac{\operatorname{dist}_G(x, s)}{\Delta} - 1 \right\rceil + p + 1 + 2 \\ &= (p + 1) \left(\left\lceil \frac{\operatorname{dist}_G(x, s)}{\Delta} \right\rceil - 1 \right) + p + 1 + 2 \\ &= (p + 1) \left\lceil \frac{\operatorname{dist}_G(x, s)}{\Delta} \right\rceil + 2 \\ &\leq (p + 1) \left\lceil \frac{\operatorname{dist}_G(x, s)}{\Delta} \right\rceil + p + 1 - i \\ &\leq (p + 1) \left\lceil \frac{\operatorname{dist}_G(u, s) - r_i}{\Delta} \right\rceil + p + 1 - i \\ &\leq (p + 1) \left\lceil \frac{\max(\operatorname{dist}_G(u, s) - r_i, 0)}{\Delta} \right\rceil + p + 1 - i = h(u, i). \end{aligned}$$

Here the last inequality follows from the trivial observation $\operatorname{dist}_G(u, s) - r_i \leq \max(\operatorname{dist}_G(u, s) - r_i, 0)$. \square

Having proved this claim, we go on with the proof of the lemma. We will now show that

$$\ell(x) + w_{H''}(v, x) + w_{H''}(u, v) \leq \frac{(\alpha + \epsilon) \operatorname{dist}_G(u, s) + \gamma_i + h(u, i) \cdot \varphi}{\varphi} \quad (3)$$

as follows. If $\text{dist}_G(u, s) > r_i$, then we have $\text{dist}_G(u, x) \geq r_i$ by the choice of x . Remember that the edge (v, x) lies on a shortest path from u to s in G . It is therefore contained in G since before the first deletion and thus will never be stretched. We also may apply the induction hypothesis on x since

$$\text{dist}_G(x, s) = \text{dist}_G(u, s) - \text{dist}_G(u, x) \leq \text{dist}_G(u, s) - r_i \leq D + \sum_{0 \leq i' \leq i-1} s_{i'} - r_i \leq D$$

due to $\sum_{0 \leq i' \leq i-1} s_{i'} \leq r_i$ by the definition of r_i . Therefore we get

$$\begin{aligned}
& (\ell(x) + w_{H''}(v, x) + w_{H''}(u, v)) \cdot \varphi \\
& \leq \delta(s, x) + w_{H''}(v, x) \cdot \varphi + w_{H''}(u, v) \cdot \varphi && \text{(definition of } \delta(s, x)\text{)} \\
& = \delta(s, x) + w_{H'}(v, x) + w_{H'}(u, v) && \text{(definition of } H'\text{)} \\
& \leq \delta(s, x) + w_H(v, x) + \varphi + w_H(u, v) + \varphi && \text{(property of } w_{H'}\text{)} \\
& \leq \delta(s, x) + w_G(v, x) + \varphi + w_F(u, v) + \varphi && ((v, x) \in E \text{ and } (u, v) \in F) \\
& \leq (\alpha + \epsilon) \text{dist}_G(x, s) + \gamma_j + h(x, j) \cdot \varphi + w_G(v, x) + \varphi + w_F(u, v) + \varphi && \text{(induction hypothesis)} \\
& = (\alpha + \epsilon) \text{dist}_G(x, s) + \gamma_j + w_F(u, v) + w_G(v, x) + (h(x, j) + 2) \cdot \varphi && \text{(rearranging terms)} \\
& \leq (\alpha + \epsilon) \text{dist}_G(x, s) + \gamma_j + w_F(u, v) + w_G(v, x) + h(u, i) \cdot \varphi && \text{(Claim 5.13)} \\
& \leq (\alpha + \epsilon) \text{dist}_G(x, s) + \gamma_0 + w_F(u, v) + w_G(v, x) + h(u, i) \cdot \varphi && (\gamma_j \leq \gamma_0) \\
& \leq (\alpha + \epsilon) \text{dist}_G(x, s) + \gamma_0 + \alpha \text{dist}_G(u, v) + \beta + w_G(v, x) + h(u, i) \cdot \varphi && \text{(by Inequality (2))} \\
& = (\alpha + \epsilon) \text{dist}_G(x, s) + \gamma_0 + \alpha \text{dist}_G(u, v) + \beta + \text{dist}_G(v, x) + h(u, i) \cdot \varphi && ((v, x) \text{ on shortest path)} \\
& \leq (\alpha + \epsilon) \text{dist}_G(x, s) + \gamma_0 + \alpha \text{dist}_G(u, v) + \beta + \alpha \text{dist}_G(v, x) + h(u, i) \cdot \varphi && (\alpha \geq 1) \\
& = (\alpha + \epsilon) \text{dist}_G(x, s) + \alpha(\text{dist}_G(u, v) + \text{dist}_G(v, x)) + \beta + \gamma_0 + h(u, i) \cdot \varphi && \text{(rearranging terms)} \\
& = (\alpha + \epsilon) \text{dist}_G(x, s) + \alpha \text{dist}_G(u, x) + \beta + \gamma_0 + h(u, i) \cdot \varphi && (v \text{ on shortest path)} \\
& = (\alpha + \epsilon) \text{dist}_G(x, s) + \alpha \text{dist}_G(u, x) + \beta + \gamma_0 - \gamma_i + \gamma_i + h(u, i) \cdot \varphi && \text{(rearranging terms)} \\
& = (\alpha + \epsilon) \text{dist}_G(x, s) + \alpha \text{dist}_G(u, x) + \epsilon r_i + \gamma_i + h(u, i) \cdot \varphi && \text{(by Lemma 5.3)} \\
& \leq (\alpha + \epsilon) \text{dist}_G(x, s) + \alpha \text{dist}_G(u, x) + \epsilon \text{dist}_G(u, x) + \gamma_i + h(u, i) \cdot \varphi && (\text{dist}_G(u, x) \geq r_i) \\
& = (\alpha + \epsilon)(\text{dist}_G(u, x) + \text{dist}_G(x, s)) + \gamma_i + h(u, i) \cdot \varphi && \text{(rearranging terms)} \\
& = (\alpha + \epsilon) \text{dist}_G(u, s) + \gamma_i + h(u, i) \cdot \varphi && (x \text{ on shortest path)}.
\end{aligned}$$

By Lemma 5.12 we have $h(u, i) \cdot \varphi \leq \epsilon \text{dist}_G(u, s) + 2\epsilon\Delta$ and thus Inequality (3) implies that

$$\begin{aligned}
\ell(x) + w_{H''}(v, x) + w_{H''}(u, v) & \leq \frac{(\alpha + 2\epsilon) \text{dist}_G(u, s) + \gamma_i + 2\epsilon\Delta}{\varphi} \\
& \leq \frac{(\alpha + 2\epsilon) \left(D + \sum_{0 \leq i' \leq i-1} s_{i'} \right) + \gamma_i + 2\epsilon\Delta}{\varphi} \\
& \leq \frac{(\alpha + 2\epsilon)D + (\alpha + 1 + \epsilon) \left(\sum_{0 \leq i' \leq i-1} w_{i'} \right) + \gamma_i + 2\epsilon\Delta}{\varphi} \\
& \leq \frac{(\alpha + 2\epsilon)D + \gamma_i + 2\epsilon\Delta}{\varphi} \\
& = \frac{(\alpha + 2\epsilon)D + \gamma}{\varphi}
\end{aligned}$$

$$\leq \frac{(\alpha + 2\epsilon)D}{\varphi} + (p+1)n^{1/p} = L.$$

As the maximum level in the monotone ES-tree is L and the edge (v, x) is not stretched, it follows from Lemma 5.10 that $\ell(v) \leq \ell(x) + w_{H''}(v, x)$ and since u is not stretched, we have

$$\ell(u) \leq \ell(v) + w_{H''}(u, v) \leq \ell(x) + w_{H''}(v, x) + w_{H''}(u, v).$$

and thus

$$\delta(s, u) = \ell(u) \cdot \varphi \leq (\ell(x) + w_{H''}(v, x) + w_{H''}(u, v)) \cdot \varphi \leq (\alpha + \epsilon) \text{dist}_G(u, s) + \gamma_i + h(u, i) \cdot \varphi$$

Case 2: $v \notin B(u)$

By Property **B2** we know that there is some node v' of priority $j' > i$ such that $u \in B(v')$ and $\text{dist}_G(u, v') \leq s(\text{dist}_G(u, v), j' - i)$. By Lemma 5.8 we therefore have

$$\text{dist}_G(u, v') \leq s(r_i, j' - i) \leq s(r_{j'-1}, j' - 1) = s_{j'-1}.$$

From the definition of F and Property **B1** it now follows that F contains the edge (u, v') of weight

$$\text{dist}_G(u, v') \leq w_F(u, v') = \hat{\delta}(u, v') \leq \alpha \text{dist}_G(u, v') + \beta \leq \alpha s_{j'-1} + \beta = w_{j'-1}$$

Since $j' \leq p-1$ we have $w_{j'-1} \leq w_{p-2} \leq r_{p-1}$. As $r_{p-1} \leq n^{1/p}\Delta$, by Lemma 5.7, we conclude that the edge (u, v') is contained H and thus also in H' and H'' .

We first prove the following inequality, which will allow us among other things to use the induction hypothesis on x .

Claim 5.14. $h(v', j') + 1 \leq h(u, i)$

Proof. Remember that $j' \geq i + 1$. If $\text{dist}_G(v', s) < r_{j'}$, we get

$$h(v', j') + 1 \leq p + 1 - j' + 1 \leq p + 1 - i \leq h(u, i).$$

If $\text{dist}_G(v', s) \geq r_{j'}$, then we use the inequality $r_i + s_{j'-1} \leq r_{j'}$ (which easily follows from the definition of $r_{j'}$) and get

$$\begin{aligned} h(v', j') + 1 &= (p+1) \left\lceil \frac{\text{dist}_G(v', s) - r_{j'}}{\Delta} \right\rceil + p + 1 - j' + 1 \\ &\leq (p+1) \left\lceil \frac{\text{dist}_G(v', s) - r_{j'}}{\Delta} \right\rceil + p + 1 - i - 1 + 1 \\ &\leq (p+1) \left\lceil \frac{\text{dist}_G(u, s) + \text{dist}_G(v', u) - r_{j'}}{\Delta} \right\rceil + p + 1 - i \\ &\leq (p+1) \left\lceil \frac{\text{dist}_G(u, s) + s_{j'-1} - r_{j'}}{\Delta} \right\rceil + p + 1 - i \\ &\leq (p+1) \left\lceil \frac{\text{dist}_G(u, s) - r_i}{\Delta} \right\rceil + p - i \\ &\leq (p+1) \left\lceil \frac{\max(\text{dist}_G(u, s) - r_i, 0)}{\Delta} \right\rceil + p + 1 - i = h(u, i). \end{aligned}$$

□

Having proved this claim, we go on with the proof of the lemma. Note that we may apply the induction hypothesis on v' because by the triangle inequality we have

$$\begin{aligned} \text{dist}_G(v', s) &\leq \text{dist}_G(u, s) + \text{dist}_G(v', u) \leq D + \sum_{0 \leq i' \leq i-1} s_{i'} + \text{dist}_G(v', u) \\ &\leq D + \sum_{0 \leq i' \leq i-1} s_{i'} + s_{j'-1} \leq D + \sum_{0 \leq i' \leq j'-1} s_{i'}. \end{aligned}$$

We will now show that

$$\ell(v') + w_{H''}(u, v') \leq \frac{(\alpha + \epsilon) \text{dist}_G(u, s) + \gamma_i + h(u, i) \cdot \varphi}{\varphi} \quad (4)$$

as follows:

$$\begin{aligned} (\ell(v') + w_{H''}(u, v')) \cdot \varphi & && (u \text{ not stretched}) \\ &= \delta(v', s) + w_{H''}(u, v') \cdot \varphi && (\text{definition of } \delta(v', s)) \\ &= \delta(v', s) + w_{H'}(u, v') && (\text{definition of } H'') \\ &\leq \delta(v', s) + w_H(u, v') + \varphi && (\text{property of } w_{H'}(u, v')) \\ &\leq \delta(v', s) + w_F(u, v') + \varphi && (\text{definition of } H) \\ &\leq (\alpha + \epsilon) \text{dist}_G(v', s) + \gamma_{j'} + h(v', j') \cdot \varphi + w_F(u, v') + \varphi && (\text{induction hypothesis}) \\ &= (\alpha + \epsilon) \text{dist}_G(v', s) + \gamma_{j'} + w_F(u, v') + (h(v', j') + 1) \cdot \varphi && (\text{rearranging terms}) \\ &\leq (\alpha + \epsilon) \text{dist}_G(v', s) + \gamma_{j'} + w_F(u, v') + h(u, i) \cdot \varphi && (\text{Claim 5.14}) \\ &\leq (\alpha + \epsilon)(\text{dist}_G(v', u) + \text{dist}_G(u, s)) + \gamma_{j'} + w_F(u, v') + h(u, i) \cdot \varphi && (\text{triangle inequality}) \\ &\leq (\alpha + \epsilon)(w_F(u, v') + \text{dist}_G(u, s)) + \gamma_{j'} + w_F(u, v') + h(u, i) \cdot \varphi && (\text{by Inequality 23}) \\ &= (\alpha + \epsilon) \text{dist}_G(u, s) + \gamma_{j'} + (\alpha + \epsilon + 1)w_F(u, v') + h(u, i) \cdot \varphi && (\text{rearranging terms}) \\ &\leq (\alpha + \epsilon) \text{dist}_G(u, s) + \gamma_{j'} + (\alpha + \epsilon + 1)w_{j'-1} + h(u, i) \cdot \varphi && (\text{by Inequality 23}) \\ &= (\alpha + \epsilon) \text{dist}_G(u, s) + \gamma_{j'-1} + h(u, i) \cdot \varphi && (\text{definition of } \gamma_{j'-1}) \\ &\leq (\alpha + \epsilon) \text{dist}_G(u, s) + \gamma_i + h(u, i) \cdot \varphi && (\gamma_i \geq \gamma_{j'-1} \text{ as } j' \geq i + 1). \end{aligned}$$

By Lemma 5.12 we have $h(u, i) \cdot \varphi \leq \epsilon \text{dist}_G(u, s) + 2\epsilon\Delta$ and thus Inequality (4) implies that

$$\ell(v') + w_{H''}(u, v') \leq \frac{(\alpha + 2\epsilon) \text{dist}_G(u, s) + \gamma_i + 2\epsilon\Delta}{\varphi} \leq \frac{(\alpha + 2\epsilon)D}{\varphi} + (p + 1)n^{1/p} = L.$$

As the maximum level in the monotone ES-tree is L and u is not stretched, it follows from Lemma 5.10 that $\ell(u) \leq \ell(v') + w_{H''}(u, v')$ and thus

$$\delta(s, u) = \ell(u) \cdot \varphi \leq (\ell(v') + w_{H''}(u, v')) \cdot \varphi \leq (\alpha + \epsilon) \text{dist}_G(u, s) + \gamma_i + h(u, i) \cdot \varphi. \quad \square$$

6 Putting Everything Together

In the following we combine our results of Section 4 and Section 5 to obtain decremental algorithms for approximate SSSP and approximate APSP.

6.1 Approximate SSSP

We first show how to obtain an algorithm for approximate SSSP. First, we obtain an algorithm that provides approximate distance for all nodes that are in distance at most R from the source, where R is some range parameter. We use a hierarchical approach to obtain this algorithm: Given an algorithm for maintaining approximate shortest paths, we obtain an algorithm for maintaining approximate balls, which in turn gives us an algorithm for maintaining approximate shortest paths for a larger range of distances than the initial algorithm. This scheme is repeated several times and can be “started” with the (exact) ES-tree.

Lemma 6.1. *For every $R \geq n$ and every $0 < \epsilon \leq 1$, there is a decremental approximate SSSP algorithm that, given a fixed source node s , maintains, for every node v , a distance estimate $\delta(s, v)$ such that $\delta(s, v) \geq \text{dist}_G(s, v)$ and if $\text{dist}_G(s, v) \leq R$, then $\delta(s, v) \leq (1 + \epsilon) \text{dist}_G(s, v)$. It has a total update time of $\tilde{O}(m^{1+3(\log \log R)/q} R^{2/q})$, where*

$$q = \left\lceil \sqrt{\left\lceil \frac{\sqrt{\log n}}{\sqrt{\log \left(\frac{8 \cdot 4^3 \log n}{\epsilon} \right)}} \right\rceil} \right\rceil$$

and, after every update in G , returns, for every node v such that $\delta(s, v)$ has changed, v together with the new value of $\delta(s, v)$.

Proof. In the proof we will use the following values. We set $a = 4$,

$$p = \left\lceil \frac{\sqrt{\log n}}{\sqrt{\log \left(\frac{8a^3 \log n}{\epsilon} \right)}} \right\rceil$$

and $q = \lfloor \sqrt{p} \rfloor$. Furthermore we set $\epsilon' = \epsilon/2(q - 2)$ and for every $0 \leq k \leq q - 2$ we set $\alpha_k = 1 + 2k\epsilon' \leq 1 + \epsilon$, $\Delta_k = R^{k/p}$, and $D_k = R^{(k+2)/p}$.

The heart of our proof is the following claim which gives us decremental approximate SSSP algorithms for larger and larger depths, until finally the full range R is covered.

Claim 6.2. *For every $0 \leq k \leq q - 2$, there is a decremental approximate SSSP algorithm APPROXSSSP_k with the following properties:*

- A1** $\delta(s, v) \geq \text{dist}_G(s, v)$
- A2** If $\text{dist}_G(s, v) \leq D_k$, then $\delta(s, v) \leq \alpha_k \text{dist}_G(s, v)$.
- A3** The total update time of APPROXSSSP_k is

$$T_k(m) = \tilde{O}(2^k m^{1+k/p} R^{2/q} (\log R)^k / \epsilon').$$

- A4** After every update in G , APPROXSSSP_k returns, for every node v such that $\delta(s, v)$ has changed, v together with the new value of $\delta(s, v)$.

Proof. We prove the claim by induction on k . In the base case $k = 0$ we use the (exact) ES-tree, which for distances up to $D \leq D_0$ has a total update time of $O(mD_0) = O(mR^{2/q})$ and thus has all claimed properties

We now consider the induction step. We apply Proposition 4.1 to obtain a decremental algorithm APPROXBALLS $_k$ (with parameters $\hat{k} = p$ and $\hat{\epsilon} = 1$) that maintains for every node $u \in V$ a set of nodes $B_k(u)$ and a distance estimate $\hat{\delta}_k(u, v)$ for every node $v \in B(u)$ such that:

B1 For every node u and every node $v \in B_k(u)$ we have $\text{dist}_G(u, v) \leq \hat{\delta}_k(u, v) \leq \alpha_{k-1} \text{dist}_G(u, v)$.

B2 Let $s(x, l) = a(a+1)^l x$. For every node u of priority $i \leq k-1$ and every node v such that $s(\text{dist}_G(u, v), p-1-i) \leq D_k$ either $v \in B(u)$ or there is some node v' of priority $j > i$ such that $\text{dist}_G(u, v') \leq s(\text{dist}_G(u, v), j-i)$.

B3 In expectation, $\sum_{u \in V} |\mathcal{B}_k(u)| = \tilde{O}(m^{1+1/k} \log D_k)$, where $\mathcal{B}_k(u)$ denotes the set of nodes ever contained in $B_k(u)$.

B4 The update time of APPROXBALLS $_k$ is

$$t_k(m) = \tilde{O} \left(m^{1+1/p} \log D_k + \sum_{0 \leq i \leq p-1} m^{1-i/p} \cdot T_{k-1}(m^{(i+1)/p}) \log D_k + T(m) \right).$$

Note that $D_k \leq R$ and thus $\log D_k \leq \log R$ and remember that by the induction hypothesis we have

$$T_{k-1}(m) = \tilde{O}(2^{k-1} m^{1+(k-1)/p} R^{2/q} (\log R)^{k-1} / \epsilon').$$

We now analyze $m^{1-i/p} \cdot T_{k-1}(m^{(i+1)/p})$ for each $0 \leq i \leq p-1$. The algorithm APPROXSSSP $_{k-1}$ is run on a graph with $m_i = m^{(i+1)/p}$ edges and $n_i = n$ nodes. Using the parameter $p_i = p$, it has a total update time of $\tilde{O}(m_i^{1+1/p_i}) = \tilde{O}(m_i^{1+1/p})$. Furthermore, we have

$$1-i/p + ((i+1)/p) \cdot (1+(k-1)/p) = 1+1/p + ((i+1)/p)((k-1)/p) \leq 1+1/p + (k-1)/p = 1+k/p$$

Thus, $m^{1-i/p} \cdot T_{k-1}(m^{(i+1)/p}) = \tilde{O}(2^{k-1} m^{1+k/p} R^{2/q} (\log R)^{k-1} / \epsilon^k)$ and since $k \geq 1$ it follows that

$$t_k(m) = \tilde{O}(2^k m^{1+k/p} R^{2/q} (\log R)^k / \epsilon').$$

We now want to argue that we may apply Proposition 5.1 to obtain an approximate decremental SSSP algorithm APPROXSSSP' $_k$ (with parameters p , $\Delta_k D_k$, and ϵ'). We first show that

$$p \leq \frac{\sqrt{\log n}}{\sqrt{\log \left(\frac{4a^3}{\epsilon} \right)}},$$

First note that $q \leq \log n$ and thus $\epsilon' = \epsilon(2(q-2)) \geq \epsilon/(2q) \geq \epsilon/(2 \log n)$. It follows that

$$\frac{\sqrt{\log n}}{\sqrt{\log \left(\frac{4a^3}{\epsilon} \right)}} \geq \frac{\sqrt{\log n}}{\sqrt{\log \left(\frac{8a^3 \log n}{\epsilon} \right)}} \geq \left\lfloor \frac{\sqrt{\log n}}{\sqrt{\log \left(\frac{8 \cdot 4^3 \log n}{\epsilon} \right)}} \right\rfloor = p.$$

Note also that for all $x \geq 1$ and $l \geq 1$ we have

$$s_k(x, l+1) = (a+1)s_k(x, l) \geq 2as_k(x, l) \geq a(\alpha_{k-1} + 1 + \epsilon')\alpha_{k-1}s(x, l)/\epsilon'.$$

We therefore may apply Proposition 5.1 to obtain an approximate decremental SSSP algorithm APPROXSSSP'_k (with parameters p , D_k , and ϵ') that maintains, for every node $v \in V$, a distance estimate $\delta'(s, v)$ such that:

A1' $\delta'(s, v) \geq \text{dist}_G(s, v)$

A2' If $\text{dist}_G(s, v) \leq D_k$, then $\delta'(s, v) \leq (\alpha_k + \epsilon') \text{dist}_G(s, v) + \epsilon' n^{1/p} \Delta_k$

A3' The total update time of APPROXSSSP'_k is

$$T'_k(m) = \tilde{O}((\alpha_k D_k / \Delta_k + n^{1/p}) \sum_{u \in V} (m + |\mathcal{B}(u)|) / \epsilon' + t_k(m, n, p, n^{1/p} \Delta)).$$

A4' After every update in G , APPROXSSSP'_k returns, for every node v such that $\delta(s, v)$ has changed, v together with the new value of $\delta(s, v)$.

Its total update time is

$$T_k(m) = \tilde{O}((\alpha_k D_k / \Delta_k + n^{1/p}) \sum_{u \in V} (m + |\mathcal{B}_k(u)|) / \epsilon' + t_k(m, n, n^{1/p} \Delta_k))$$

Note that $\alpha_k \leq 1 + \epsilon \leq 2$ and $D_k / \Delta_k = R^{2/q}$. Since $q \leq p$ and $R \geq n$ we have $n^{1/p} \leq R^{2/q}$. We also have $\sum_{u \in V} |\mathcal{B}_k(u)| = \tilde{O}(m^{1+1/p} \log R)$. Therefore the total update time of APPROXSSSP'_k is

$$T'_k(m) = \tilde{O}(m^{1+1/p} R^{2/q} \log R / \epsilon' + 2^k m^{1+k/p} R^{2/q} (\log R)^k / \epsilon')$$

Since $k \geq 1$ it follows that

$$T_k(m) = \tilde{O}(2^k m^{1+k/p} R^{2/q} (\log R)^k / \epsilon').$$

Let APPROXSSSP_k denote the algorithm that internally runs both APPROXSSSP'_k and APPROXSSSP_{k-1} and additionally maintains, for every node v , the value $\delta_k(s, v) = \min(\delta'_k(s, v), \delta_{k-1}(s, v))$. Since both APPROXSSSP'_k and APPROXSSSP_{k-1} return, after each update in G , every node v for which $\delta(s, v)$ has changed, and the minimum can be computed in constant time, APPROXSSSP_k has the same asymptotic total update time as APPROXSSSP'_k. It remains to show that $\delta_k(s, v)$ fulfills the desired approximation guarantee for every node v . Since both $\delta'_k(s, v) \geq \text{dist}_G(s, v)$ and $\delta_{k-1}(s, v) \geq \text{dist}_G(s, v)$ also $\delta_k(s, v) \geq \text{dist}_G(s, v)$. Furthermore, we know that if $\text{dist}_G(s, v) \leq D_k$, then $\delta'_k(s, v) \leq \epsilon n^{1/p} \Delta_k$. Let v be a node such that $\text{dist}_G(s, v) \leq D_k$. If $\text{dist}_G(s, v) \leq D_{k-1}$, then $\delta_k(s, v) \leq \delta_{k-1}(s, v) \leq \alpha_{k-1} \text{dist}_G(s, v) \leq \alpha_k \text{dist}_G(s, v)$. If $\text{dist}_G(s, v) \geq D_{k-1}$, then

$$\begin{aligned} \delta_k(s, v) &\leq \delta'_k(s, v) \leq (\alpha_{k-1} + \epsilon') \text{dist}_G(s, v) + \epsilon' n^{1/p} \Delta_k \\ &\leq (\alpha_{k-1} + \epsilon') \text{dist}_G(s, v) + \epsilon' D_{k-1} \leq (\alpha_{k-1} + 2\epsilon') \text{dist}_G(s, v) \\ &= \alpha_k \text{dist}_G(s, v). \end{aligned}$$

This finishes the proof of the claim. □

The lemma now follows from the claim by observing that APPROXSSSP_{q-2} is the desired decremental approximate SSSP algorithm. The correctness simply follows from $D_{q-2} = R$. The total update time is

$$T_{q-2}(m) = \tilde{O}(2^{q-2} m^{1+(q-2)/p} R^{2/q} (\log R)^{q-2} / \epsilon').$$

Remember that $q = \lfloor \sqrt{p} \rfloor$ and thus $(q-2)/p \leq q/p \leq 1/\sqrt{p} \leq 1/q$. By the definition of p we have $(2/\epsilon')^p \leq n^{1/p}$ and thus $(2/\epsilon')^q \leq (2/\epsilon')^p \leq n^{1/p} \leq n^{1/q}$ and furthermore $(\log R)^q \leq (\log R)^p = (2^p)^{\log \log R} \leq (n^{1/p})^{\log \log R} = n^{(\log \log R)/p} \leq n^{(\log \log R)/q}$. It follows that the total update time is

$$T_{q-2}(m) = \tilde{O}(m^{1+3(\log \log R)/q} R^{2/q}). \quad \square$$

We can turn the algorithm above into an algorithm for the full distance range by using the rounding technique once more.

Theorem 6.3. *For every $0 < \epsilon \leq 1$, there is a decremental approximate SSSP algorithm that, given a fixed source node s , maintains, for every node v , a distance estimate $\delta(s, v)$ such that $\text{dist}_G(s, v) \leq \delta(s, v) \leq (1 + \epsilon) \text{dist}_G(s, v)$. It has constant query time and a total update time of*

$$O(m^{1+O(\log^{5/4}((\log n)/\epsilon)/\log^{1/4} n)} \log W).$$

If $1/\epsilon = O(\text{polylog } n)$, then the total update time is $O(m^{1+o(1)} \log W)$.

Proof. For every $0 \leq i \leq \lfloor \log(nW) \rfloor$ we define

$$\varphi_i = \frac{\epsilon 2^i}{n}.$$

Let G'_i be the graph that has the same nodes and edges as G and in which every edge weight is rounded to the next multiple of φ_i , i.e., every edge (u, v) in G'_i has weight

$$w_{G'_i}(u, v) = \left\lceil \frac{w_G(u, v)}{\varphi_i} \right\rceil \cdot \varphi_i$$

where $w_G(u, v)$ is the weight of (u, v) in G . This rounding guarantees that

$$w_G(u, v) \leq w_{G'_i}(u, v) \leq w_G(u, v) + \varphi_i$$

for every edge (u, v) of G . Furthermore we define G''_i to be the graph that has the same nodes and edges as G'_i and in which every edge weight is scaled down by a factor of $1/\varphi_i$, i.e., every edge (u, v) in G''_i has weight

$$w_{G''_i}(u, v) = \frac{w_{G'_i}(u, v)}{\varphi_i} = \left\lceil \frac{w(u, v)}{\varphi_i} \right\rceil.$$

The algorithm is as follows: For every $0 \leq i \leq \lfloor \log(nW) \rfloor$ we use the algorithm of Lemma 6.1 on the graph G''_i with $R = 4n/\epsilon$ to maintain a distance estimate $\delta_i(s, v)$ for every node v that satisfies

- $\delta_i(s, v) \geq \text{dist}_{G''_i}(s, v)$ and

- if $\text{dist}_{G'_i}(s, v) \leq R$, then $\delta_i(s, v) \leq (1 + \epsilon) \text{dist}_{G'_i}(s, v)$.

We let our algorithm return the distance estimate

$$\delta(s, v) = \min_{0 \leq i \leq \lfloor \log nW \rfloor} \varphi_i \delta_i(s, v).$$

We now show that there is some $0 \leq i \leq \lfloor \log(nW) \rfloor$ such that $\varphi_i \delta_i(s, v) \leq (1 + 3\epsilon) \text{dist}_G(s, v)$. As $\delta(s, v)$ is the minimum of all the distance estimates, this implies that $\delta(s, v) \leq (1 + 3\epsilon) \text{dist}_G(s, v)$. In particular, we know that there is some $0 \leq i \leq \lfloor \log(nW) \rfloor$ such that $2^i \leq \text{dist}_G(s, v) \leq 2^{i+1}$ since W is the maximum edge weight and all paths consist of at most n edges. Consider a shortest path π from v to s in G whose weight is equal to $\text{dist}_G(s, v)$. Let $w_G(\pi)$ and $w_{G'_i}(\pi)$ denote the weight of the path π in G and G'_i , respectively. Since π consists of at most n edges we have $w_{G'_i}(\pi) \leq w(\pi) + n\varphi_i$. Therefore we get

$$\begin{aligned} \text{dist}_{G'_i}(s, v) &\leq w_{G'_i}(\pi) \leq w(\pi) + n\varphi_i = \text{dist}_G(s, v) + \epsilon 2^i \leq \text{dist}_G(s, v) + \epsilon \text{dist}_G(s, v) \\ &= (1 + \epsilon) \text{dist}_G(s, v). \end{aligned}$$

Now observe the following:

$$\begin{aligned} \text{dist}_{G'_i}(s, v) &= \frac{\text{dist}_{G'_i}(s, v)}{\varphi_i} \leq \frac{(1 + \epsilon) \text{dist}_G(s, v)}{\varphi_i} \leq \frac{2 \text{dist}_G(s, v)}{\varphi_i} = \frac{2 \text{dist}_G(s, v)n}{\epsilon 2^i} \\ &\leq \frac{2 \cdot 2^{i+1}n}{\epsilon 2^i} = \frac{4n}{\epsilon} = R. \end{aligned}$$

Since $\text{dist}_{G'_i}(s, v) \leq R$ we get $\delta_i(s, v) \leq (1 + \epsilon) \text{dist}_{G'_i}(s, v)$ by Lemma 6.1. Thus, we get

$$\varphi_i \delta_i(s, v) \leq \varphi_i ((1 + \epsilon) \text{dist}_{G'_i}(s, v)) = (1 + \epsilon) \text{dist}_{G'_i}(s, v) \leq (1 + \epsilon)^2 \text{dist}_G(s, v) \leq (1 + 3\epsilon) \text{dist}_G(s, v)$$

as desired.

We now analyze the running time of this algorithm. By Lemma 6.1, for every $0 \leq i \leq \lfloor \log(nW) \rfloor$, maintaining $\delta_i(s, v)$ on G'_i for every node s takes time $\tilde{O}(m^{1+3(\log \log R)/q} R^{2/q})$, where

$$q = \left\lceil \sqrt{\left\lceil \frac{\sqrt{\log n}}{\sqrt{\log \left(\frac{8 \cdot 4^3 \log n}{\epsilon} \right)}} \right\rceil} \right\rceil$$

By our choice of $R = 4n/\epsilon$, the total update time for maintaining all these $\lfloor \log(nW) \rfloor$ distance estimates is $\tilde{O}(m^{1+5(\log \log (4n/\epsilon))/q} \log W/\epsilon)$, where

$$q = \left\lceil \sqrt{\left\lceil \frac{\sqrt{\log n}}{\sqrt{\log \left(\frac{8 \cdot 4^3 \log n}{\epsilon} \right)}} \right\rceil} \right\rceil$$

To obtain a $(1 + \epsilon)$ approximation (instead of a $(1 + 3\epsilon)$ -approximation, we simply run the whole algorithm with $\epsilon' = \epsilon/3$. This results in a total update time of $\tilde{O}(m^{1+5(\log \log (12n/\epsilon))/q} \log W/\epsilon)$,

where

$$q = \left\lfloor \sqrt{\left\lfloor \frac{\sqrt{\log n}}{\sqrt{\log\left(\frac{24 \cdot 4^3 \log n}{\epsilon}\right)}} \right\rfloor} \right\rfloor$$

Now observe that $1/\epsilon \leq n^{1/q}$ and that

$$\frac{5 \left(\log \log \left(\frac{12n}{\epsilon} \right) \right)}{q} = O \left(\frac{(\log \log \left(\frac{n}{\epsilon} \right)) \left(\log \left(\frac{\log n}{\epsilon} \right) \right)^{1/4}}{(\log n)^{1/4}} \right) = O \left(\frac{\left(\log \left(\frac{\log n}{\epsilon} \right) \right)^{5/4}}{(\log n)^{1/4}} \right).$$

Since $\tilde{O}(1) = O(\text{polylog } n) = O(n^{O(\log^{5/4}(\log n))})$ the total update time therefore is

$$O(m^{1+O(\log^{5/4}((\log n)/\epsilon)/\log^{1/4} n)} \log W).$$

If $1/\epsilon = O(\text{polylog } n)$, then the total update time is $O(m^{1+\log^{5/4} \log n / \log^{1/4} n} \log W)$, which is $O(m^{1+o(1)} \log W)$ since $\lim_{x \rightarrow \infty} (\log^{5/4} \log n / \log^{1/4} n) = 0$.

The query time of the algorithm described above is $O(\log(nW))$ as it has to compute $\delta(s, v) = \min_{0 \leq i \leq \lfloor \log nW \rfloor} \delta_i(s, v) \cdot \varphi_i$ when asked for the approximate distance from v to s . We can reduce the query time to $O(1)$ by using a min-heap for every node v that stores $\delta_i(s, v)$ for all $0 \leq i \leq \lfloor \log(nW) \rfloor$. This allows us to query for $\delta(s, v)$ in constant time. \square

6.2 Approximate APSP

We now show how to use our techniques to obtain a decremental approximate APSP algorithm. This is conceptually simple now. We simply use the approximate SSSP algorithm from Theorem 6.3 and plug it into the algorithm for maintaining approximate balls from Proposition 4.1. By using an adequate query procedure we can use the distance estimates maintained for the approximate balls to return the approximate distances between any two nodes.

Theorem 6.4. *There is a decremental approximate APSP algorithm that upon a query for the approximate between any pair of nodes u and v returns a distance estimate $\delta(u, v)$ such that $\text{dist}_G(u, v) \leq \delta(u, v) \leq ((2 + \epsilon)^k - 1) \text{dist}_G(u, v)$. It has a query time of $O(k^k)$ and a total update time of*

$$O(m^{1+1/k+O(\log^{5/4}((\log n)/\epsilon)/\log^{1/4} n)} \log^2 W)$$

If $1/\epsilon = O(\text{polylog } n)$, then the total update time is $O(m^{1+o(1)} \log^2 W)$.

Proof. We use the approximate SSSP algorithm of Theorem 6.3 that provides a $(1 + \epsilon)$ -approximation and has a total update time of

$$T(m, n) = O(m^{1+O(\log^{5/4}((\log n)/\epsilon)/\log^{1/4} n)} \log W)$$

and if $1/\epsilon = O(\text{polylog } n)$, then the total update time is $T(m, n) = O(m^{1+o(1)} \log W)$. By Proposition 4.1 we can maintain approximate balls with a total update time of

$$t(m, n, k, \epsilon) = \tilde{O} \left(m^{1+1/k} \log D / \epsilon + \sum_{0 \leq i \leq k-1} m^{1-i/k} \cdot T(m_i, n_i) \log D / \epsilon + T(m, n) \right),$$

where, for each $0 \leq i \leq k-1$, $m_i = m^{(i+1)/k}$ and $n_i = \min(m_i, n)$. Using similar arguments as above we get that $t(m, n, k, \epsilon) = O(m^{1+1/k+O(\log^{5/4}((\log n)/\epsilon)/\log^{1/4} n)} \log^2 W)$ and $O(m^{1+o(1)} \log^2 W)$ if $1/\epsilon = O(\text{polylog } n)$.

Additionally we maintain, for every node $v \in V$, the node $c_i(v)$ which is a node with minimum $\delta(u, v)$ among all nodes u of priority j such that $v \in B(u)$. This can be done as follows. For every node v we maintain a heap containing all nodes u of priority i such that $v \in B(u)$ using the key $\delta(u, v)$. Every time v joins or leaves $B(u)$ we insert or remove u from the heap of v . Every time $\delta(u, v)$ changes, we update the key of u in the heap of v . After each insert, remove, or update in the heap of some node v , we find the minimal element $c_i(v)$ of the heap. As each heap operation takes logarithmic time, the total update time of the algorithm of Proposition 4.1 only increases by a logarithmic factor.

Procedure 2: QUERY(u, v)

```

1 if  $v \in B(u)$  then
2   |  $\delta'(u, v) \leftarrow \delta(u, v)$ 
3 else
4   | Set  $i$  to the priority of  $u$ 
5   | foreach  $j = i + 1$  to  $k - 1$  do
6     | if  $c_j(u)$  exists then
7       |  $v'' \leftarrow c_j(u)$ 
8       |  $\delta'(v'', v) \leftarrow \text{QUERY}(v'', v)$ 
9       |  $\delta'_j(u, v) \leftarrow \delta(v, v'') + \delta'(v'', v)$ 
10      | else
11        |  $\delta'_j(u, v) \leftarrow \infty$ 
12      |  $\delta'(u, v) \leftarrow \min_{i+1 \leq j \leq k-1} \delta'_j(u, v)$ 
13 return  $\delta'(u, v)$ 

```

To answer a query for the approximate distance between a pair of nodes u and v we use Procedure 2. This procedure first tests whether $v \in B(u)$ and if yes returns $\delta(u, v)$. Otherwise it does the following for every $j \geq i+1$, where i is the priority of u : It first computes the node $c_j(u)$, which among the nodes v' of priority j with $u \in B(v')$ is the one with the minimum value of $\delta(v', u)$. Then it recursively queries for the approximate distance $\delta'(c_j(u), v)$ from $c_j(u)$ to v and sets the distance estimate via $c_j(u)$ to $\delta'_j(u, v) = \delta(v, c_j(u)) + \delta'(c_j(u), v)$. Finally, it returns the minimum of all distance estimates $\delta'_j(u, v)$.

Note that in each instance there are $O(k)$ recursive calls and with each recursive call the priority of u increases by at least one. Thus the running time of the query procedure is $O(k^k)$.

Claim 6.5. *For every pair of nodes u and v the distance estimate $\delta'(u, v)$ computed by Procedure 2 satisfies $\delta'(u, v) \leq (((1 + \epsilon)^2 + 1)^{k-i} - 1) \text{dist}_G(u, v)$, where i is the priority of u .*

Proof. The proof is by induction on the priority i of u . Let $\delta'(u, v)$ denote the distance estimate returned by Procedure 2. If $i = k-1$, then we know that $v \in B(u)$ and thus $\delta'(u, v) = \delta(u, v) \leq (1 + \epsilon) \text{dist}_G(u, v)$. If $i < k-1$ we distinguish between the two cases $v \in B(u)$ and $v \notin B(u)$. If $v \in B(u)$, then $\delta'(u, v) = \delta(u, v) \leq (1 + \epsilon) \text{dist}_G(u, v)$. If

$v \notin B(u)$, then by Proposition 4.1 there is a node v' of priority $j > i$ such that $u \in B(v')$ and $\text{dist}_G(u, v') \leq (1 + \epsilon)^2((1 + \epsilon)^2 + 1)^{j-i-1} \text{dist}_G(u, v)$.

We will now argue that $\delta'_j(u, v) \leq 2((1 + \epsilon)^3 + 1)^{k-1-i} \text{dist}_G(u, v)$, which implies the same upper bound for $\delta'(u, v)$. Set $v'' \leftarrow c_j(u)$. Since both v'' and v' have priority j and $u \in B(v')$ as well as $v \in B(v'')$ we have $\delta(u, v'') \leq \delta(u, v')$ by the definition of v'' . Since $\delta(u, v') \leq (1 + \epsilon) \text{dist}_G(u, v')$, we have

$$\begin{aligned} \delta(u, v'') &\leq (1 + \epsilon) \text{dist}_G(u, v') \leq (1 + \epsilon)^3((1 + \epsilon)^2 + 1)^{j-i-1} \text{dist}_G(u, v) \\ &\leq (1 + \epsilon)^3((1 + \epsilon)^3 + 1)^{j-i-1} \text{dist}_G(u, v). \end{aligned}$$

To simplify the presentation in the following we set $a = (1 + \epsilon)^3$ and thus have $\delta(u, v'') \leq a(a + 1)^{j-i-1} \text{dist}_G(u, v)$. By the triangle inequality we have

$$\begin{aligned} \text{dist}_G(v'', v) &\leq \text{dist}_G(v'', u) + \text{dist}_G(u, v) \leq \delta(v'', u) + \text{dist}_G(u, v) \\ &\leq (a(a + 1)^{j-i-1} + 1) \text{dist}_G(u, v) \end{aligned}$$

and by the induction hypothesis we have

$$\begin{aligned} \delta'(v'', v) &\leq (2(a + 1)^{k-1-j} - 1) \text{dist}_G(v'', v) \\ &\leq (2(a + 1)^{k-1-j} - 1)(a(a + 1)^{j-i-1} + 1) \text{dist}_G(u, v). \end{aligned}$$

Since $j \geq i + 1$ we get

$$\begin{aligned} \delta'_j(u, v) &= \delta(u, v'') + \delta'(v'', v) \\ &\leq \left(a(a + 1)^{j-i-1} + (2(a + 1)^{k-1-j} - 1)(a(a + 1)^{j-i-1} + 1) \right) \text{dist}_G(u, v) \\ &= \left(2(a + 1)^{k-1-j}(a(a + 1)^{j-i-1} + 1) - 1 \right) \text{dist}_G(u, v) \\ &= \left(2a(a + 1)^{k-1-(i+1)} + 2(a + 1)^{k-1-j} - 1 \right) \text{dist}_G(u, v) \\ &\leq \left(2a(a + 1)^{k-1-(i+1)} + 2(a + 1)^{k-1-(i+1)} - 1 \right) \text{dist}_G(u, v) \\ &= \left(2(a + 1)^{k-1-(i+1)}(a + 1) - 1 \right) \text{dist}_G(u, v) \\ &= (2(a + 1)^{k-1-i} - 1) \text{dist}_G(u, v). \quad \square \end{aligned}$$

Note that $2 \leq ((1 + \epsilon)^3 + 1)$ and therefore we have $\delta'(u, v) \leq (((1 + \epsilon)^3 + 1)^{k-i} - 1) \text{dist}_G(u, v)$. Furthermore, $(1 + \epsilon)^3 \leq 1 + 7\epsilon$ and in the worst case $i = 0$. Thus, by running the whole algorithm with $\epsilon' = \epsilon/7$, we can guarantee that $\delta'(u, v) \leq ((2 + \epsilon)^k - 1) \text{dist}_G(u, v)$. \square

7 Conclusion

In this paper, we show that single-source shortest paths in undirected graphs can be maintained under edge deletions with near-linear total update time and constant query time. The main approach is to maintain an $(n^{o(1)}, \epsilon)$ -hop set of near-linear size in near-linear time. We leave two major open problems. The first problem is whether the same total update time can be achieved for directed graphs. This problem is very challenging because such a hop set is not known even in the static setting. Moreover, improving the current $\tilde{O}(mn^{0.9+o(1)})$ total

update time by [HKN14b, HKN15] for the decremental reachability problem is already very interesting. The second major open problem is to derandomize our algorithm. The major task here is to deterministically maintain the priorities and corresponding balls of the nodes, which is the key to maintaining the hop set. A related question is whether the algorithm of Roditty and Zwick [RZ12] for decrementally maintaining the original distance oracle of Thorup and Zwick (and the corresponding spanners and emulators) can be derandomized. (Note however that the distance oracle of Thorup and Zwick can be constructed deterministically in the static setting [RTZ05].)

Acknowledgement

The authors would like to thank the anonymous reviewers of FOCS 2014 for their valuable feedback.

References

- [ACT14] 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)*. 2014, pp. 1–16 (cit. on p. 5).
- [AVW14] Amir Abboud and Virginia Vassilevska Williams. “Popular conjectures imply strong lower bounds for dynamic problems”. In: *Symposium on Foundations of Computer Science (FOCS)*. 2014, pp. 434–443 (cit. on p. 4).
- [BR11] Aaron Bernstein and Liam Roditty. “Improved Dynamic Algorithms for Maintaining Approximate Shortest Paths Under Deletions”. In: *Symposium on Discrete Algorithms (SODA)*. 2011, pp. 1355–1365 (cit. on pp. 1, 3, 6).
- [Ber09] Aaron Bernstein. “Fully Dynamic $(2 + \epsilon)$ Approximate All-Pairs Shortest Paths with Fast Query and Close to Linear Update Time”. In: *Symposium on Foundations of Computer Science (FOCS)*. 2009, pp. 693–702 (cit. on pp. 1, 4, 7).
- [Ber13] Aaron Bernstein. “Maintaining Shortest Paths Under Deletions in Weighted Directed Graphs”. In: *Symposium on Theory of Computing (STOC)*. 2013, pp. 725–734 (cit. on pp. 4, 7).
- [Coh00] Edith Cohen. “Polylog-Time and Near-Linear Work Approximation Scheme for Undirected Shortest Paths”. In: *Journal of the ACM* 47.1 (2000). Announced at STOC’94, pp. 132–166 (cit. on pp. 1, 7).
- [Coh98] Edith Cohen. “Fast Algorithms for Constructing t -Spanners and Paths with Stretch t ”. In: *SIAM Journal on Computing* 28.1 (1998). Announced at FOCS’93, pp. 210–236 (cit. on pp. 4, 7).
- [ES81] Shimon Even and Yossi Shiloach. “An On-Line Edge-Deletion Problem”. In: *Journal of the ACM* 28.1 (1981), pp. 1–4 (cit. on pp. 1, 3, 5).

- [HK95] Monika Henzinger and Valerie King. “Fully Dynamic Biconnectivity and Transitive Closure”. In: *Symposium on Foundations of Computer Science (FOCS)*. 1995, pp. 664–672 (cit. on pp. 4, 5).
- [HKN13] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. “Dynamic Approximate All-Pairs Shortest Paths: Breaking the $O(mn)$ Barrier and Derandomization”. In: *Symposium on Foundations of Computer Science (FOCS)*. 2013, pp. 538–547 (cit. on pp. 1, 6, 22, 27).
- [HKN14a] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. “A Subquadratic-Time Algorithm for Dynamic Single-Source Shortest Paths”. In: *Symposium on Discrete Algorithms (SODA)*. 2014, pp. 1053–1072 (cit. on p. 4).
- [HKN14b] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. “Sublinear-Time Incremental Algorithms for Single-Source Reachability and Shortest Paths on Directed Graphs”. In: *Symposium on Theory of Computing (STOC)*. 2014, pp. 674–683 (cit. on pp. 1, 4, 43).
- [HKN15] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. “Improved Algorithms for Incremental Single-Source Reachability on Directed Graphs”. In: *International Colloquium on Automata, Languages and Programming (ICALP)*. 2015, pp. 725–736 (cit. on pp. 4, 43).
- [HKN⁺15] Monika Henzinger, Sebastian Krinninger, Danupon Nanongkai, and Thatchaphol Saranurak. “Unifying and Strengthening Hardness for Dynamic Problems via the Online Matrix-Vector Multiplication Conjecture”. In: *Symposium on Theory of Computing (STOC)*. 2015, pp. 21–30 (cit. on p. 3).
- [KKM⁺12] Maleq Khan, Fabian Kuhn, Dahlia Malkhi, Gopal Pandurangan, and Kunal Talwar. “Efficient distributed approximation algorithms via probabilistic tree embeddings”. In: *Distributed Computing* 25.3 (2012). Announced at PODC’08, pp. 189–205 (cit. on p. 7).
- [Kin99] Valerie King. “Fully Dynamic Algorithms for Maintaining All-Pairs Shortest Paths and Transitive Closure in Digraphs”. In: *Symposium on Foundations of Computer Science (FOCS)*. 1999, pp. 81–91 (cit. on pp. 4, 5).
- [Lac13] Jakub Łacki. “Improved Deterministic Algorithms for Incremental Reachability and Strongly Connected Components”. In: *ACM Transactions on Algorithms* 9.3 (2013). Announced at SODA’11, p. 27 (cit. on p. 4).
- [Mad10] Aleksander Mądry. “Faster Approximation Schemes for Fractional Multicommodity Flow Problems via Dynamic Graph Algorithms”. In: *Symposium on Theory of Computing (STOC)*. 2010, pp. 121–130 (cit. on pp. 1, 3, 4, 7).
- [Nan14] Danupon Nanongkai. “Distributed Approximation Algorithms for Weighted Shortest Paths”. In: *Symposium on Theory of Computing (STOC)*. 2014, pp. 565–573 (cit. on p. 7).
- [RTZ05] Liam Roditty, Mikkel Thorup, and Uri Zwick. “Deterministic Constructions of Approximate Distance Oracles and Spanners”. In: *International Colloquium on Automata, Languages and Programming (ICALP)*. 2005, pp. 261–272 (cit. on p. 43).

- [RZ08] Liam Roditty and Uri Zwick. “Improved Dynamic Reachability Algorithms for Directed Graphs”. In: *SIAM Journal on Computing* 37.5 (2008). Announced at FOCS’02, pp. 1455–1471 (cit. on p. 4).
- [RZ11] Liam Roditty and Uri Zwick. “On Dynamic Shortest Paths Problems”. In: *Algorithmica* 61.2 (2011). Announced at ESA’04, pp. 389–401 (cit. on p. 3).
- [RZ12] Liam Roditty and Uri Zwick. “Dynamic Approximate All-Pairs Shortest Paths in Undirected Graphs”. In: *SIAM Journal on Computing* 41.3 (2012). Announced at FOCS’04, pp. 670–683 (cit. on p. 43).
- [Rod13] Liam Roditty. “Decremental maintenance of strongly connected components”. In: *Symposium on Discrete Algorithms (SODA)*. 2013, pp. 1143–1150 (cit. on pp. 3, 4).
- [TZ05] Mikkel Thorup and Uri Zwick. “Approximate Distance Oracles”. In: *Journal of the ACM* 52.1 (2005). Announced at STOC’01, pp. 74–92 (cit. on p. 4).
- [TZ06] Mikkel Thorup and Uri Zwick. “Spanners and emulators with sublinear distance errors”. In: *Symposium on Discrete Algorithms (SODA)*. 2006, pp. 802–809 (cit. on pp. 6–9, 11, 16).
- [Tho99] Mikkel Thorup. “Undirected Single-Source Shortest Paths with Positive Integer Weights in Linear Time”. In: *Journal of the ACM* 46.3 (1999). Announced at FOCS’97, pp. 362–394 (cit. on p. 3).
- [UY91] Jeffrey D. Ullman and Mihalis Yannakakis. “High-Probability Parallel Transitive-Closure Algorithms”. In: *SIAM Journal on Computing* 20.1 (1991). Announced at SPAA’90, pp. 100–125 (cit. on p. 6).
- [VWW10] Virginia Vassilevska Williams and Ryan Williams. “Subcubic Equivalences between Path, Matrix and Triangle Problems”. In: *Symposium on Foundations of Computer Science (FOCS)*. 2010, pp. 645–654 (cit. on p. 3).
- [Zwi02] Uri Zwick. “All Pairs Shortest Paths using Bridging Sets and Rectangular Matrix Multiplication”. In: *Journal of the ACM* 49.3 (2002). Announced at FOCS’98, pp. 289–317 (cit. on pp. 4, 7).