Near-Optimal Decremental SSSP in Dense Weighted Digraphs

Aaron Bernstein Rutgers University New Brunswick bernstei@gmail.com Maximilian Probst Gutenberg BARC, University of Copenhagen maximilian.probst@outlook.com

Christian Wulff-Nilsen BARC, University of Copenhagen koolooz@di.ku.dk

Abstract

In the decremental Single-Source Shortest Path problem (SSSP), we are given a weighted directed graph G = (V, E, w) undergoing edge deletions and a source vertex $r \in V$; let n = |V|, m = |E| and W be the aspect ratio of the graph. The goal is to obtain a data structure that maintains shortest paths from r to all vertices in V and can answer distance queries in O(1) time, as well as return the corresponding path P in O(|P|) time.

This problem was first considered by Even and Shiloach [JACM'81], who provided an algorithm with total update time O(mn) for unweighted undirected graphs; this was later extended to directed weighted graphs [FOCS'95, STOC'99]. There are conditional lower bounds showing that O(mn) is in fact near-optimal [ESA'04, FOCS'14, STOC'15, STOC'20]. In a breakthrough result, Forster et al. showed that total update time $\min\{m^{7/6}n^{2/3+o(1)}, m^{3/4}n^{5/4+o(1)}\}$ polylog(W) = $mn^{0.9+o(1)}$ polylogW is possible if the algorithm is allowed to return $(1 + \epsilon)$ -approximate paths, instead of exact ones [STOC'14, ICALP'15]. No further progress was made until Probst Gutenberg and Wulff-Nilsen [SODA'20] provided a new approach for the problem, which yields total time $\tilde{O}(\min\{m^{2/3}n^{4/3}\log W, (mn)^{7/8}\log W\}) = \tilde{O}(\min\{n^{8/3}\log W, mn^{3/4}\log W\})$.

Our result builds on this recent approach, but overcomes its limitations by introducing a significantly more powerful abstraction, as well as a different core subroutine. Our new framework yields a decremental $(1 + \epsilon)$ -approximate SSSP data structure with total update time $\tilde{O}(n^2 \log^4 W/\epsilon)$. Our algorithm is thus *near-optimal* for dense graphs with polynomial edge-weights. Our framework can also be applied to sparse graphs to obtain total update time $\tilde{O}(mn^{2/3} \log^3 W/\epsilon)$. Combined, these data structures dominate all previous results. Like all previous o(mn) algorithms that can return a path (not just a distance estimate), our result is randomized and assumes an oblivious adversary.

Our framework effectively allows us to reduce SSSP in general graphs to the same problem in directed acyclic graphs (DAGs). We believe that our framework has significant potential to influence future work on directed SSSP, both in the dynamic model and in others.

1 Introduction

In the Single-Source Shortest Paths (SSSP) problem, the input is a directed weighted graph G = (V, E, w) and a dedicated source vertex $r \in V$, and the goal is to compute shortest paths from r to every other vertex v in V. Let n = |V|, m = |E|, and W be the aspect ratio of the graph, which is the ratio of maximum to minimum edge weight. The problem can be solved in $O(m + n \log(n))$ time using Dijkstra's algorithm. In this article, we study the dynamic version of the problem, where the graph changes over time. The most general model is the *fully dynamic* one, where the graph is subject to a sequence of edge insertions and deletions. Unfortunately, there are extremely strong conditional lower bound for this model [RZ04; AW14; Hen+15; PVW20].

For this reason, much of the research on this problem has focused on the *decremental* setting. Formally, the algorithm is given a graph G = (V, E, w) subject to a sequence of edge deletions and edge weight increases, and the goal is is to maintain shortest distances/paths from r to every $v \in V$. Although decremental SSSP only applies to a more restricted model, it is an extremely common subroutine in other dynamic algorithms (including fully dynamic ones), and has recently been used to make progress on long-standing *static* problems such as computing max-flow [CK19; Chu+19], multi-commodity flow [Mad10], expanders [Chu+19], and sparse cuts [Chu+19]. For this reason, decremental SSSP is one of the most well-studied problems in dynamic algorithms [ES81; BR11; HKN14a; HKN14b; HKN15; HKN16; BC16; BC17; Ber17; CK19; Chu+19; GW20a; GW20b; BPS20].

The first algorithm for decremental SSSP was the Even and Shiloach tree [ES81], which dates back to 1981 and has total update time O(mn) over the entire sequence of deletions; it was later extended to directed weighted graphs [HK95; Kin99]. O(mn) is conditionally optimal for the exact version [RZ04; AW14; Hen+15; PVW20], but one can do better with a $(1 + \epsilon)$ -approximation. In fact, recent research culminated in a near-optimal algorithm for undirected graphs [HKN14a] with total update time $m^{1+o(1)} \log(W)^1$; there has also been more recent work improving upon O(mn) with adaptive or even deterministic algorithms (see e.g. [BC16; BC17; Ber17; CK19; Chu+19; GW20b; Ber+20]).

In directed graphs, however, the decremental SSSP problem remains poorly understood. The first algorithms to improve upon the classic O(mn) bound (with a $(1 + \epsilon)$ -approximation) were by Henzinger, Forster and Nanongkai [HKN14b] in 2014; the total update time is $\min\{m^{7/6}n^{2/3+o(1)}, m^{3/4}n^{5/4+o(1)}\}$ polylog $(W) = mn^{0.9+o(1)}$ polylog(W) [HKN15]. Since then, the only progress on the problem is a very recent algorithm of Probst Gutenberg and Wulff-Nilsen [GW20a] with total update time $\tilde{O}(\min\{m^{2/3}n^{4/3}, mn^{3/4}\}\log(W))$, or slightly better in unweighted graphs. Very recently, it was further shown in [BPS20] that a total update time of $n^{2+2/3+o(1)}$ can even be obtained deterministically. Besides this recent progress, state-of-the-art algorithms for directed graphs still lag far behind the ones for undirected ones, and only achieve small improvements beyond O(mn).

For more related work on the problems of maintaining Single-Source Reachability, Strongly-Connected Components, Single-Source Shortest Paths and All-Pairs Shortest Paths in incre-

¹In the decremental setting m is taken to be the number of edges in the *initial* graph G

mental, decremental and fully-dynamic graphs, we refer the reader to Appendix A.

Our Contribution. We make significant progress on this problem and present the first *near-optimal* algorithm for decremental SSSP in directed dense graphs.

Theorem 1.1. Given a decremental input graph G = (V, E, w) with n = |V|, m = |E| and aspect ratio W, a dedicated source $r \in V$ and $\epsilon > 0$, there is a randomized algorithm that maintains a distance estimate $\widetilde{\operatorname{dist}}(r, x)$, for every $x \in V$, such that

$$\operatorname{dist}_G(r, x) \leq \operatorname{dist}(r, x) \leq (1 + \epsilon) \operatorname{dist}_G(r, x)$$

at any stage w.h.p. The algorithm has total expected update time $\tilde{O}(n^2 \log^4 W/\epsilon)$. Distance queries are answered in O(1) time, and a corresponding path P can be returned in O(|P|) time.

We also present the currently fastest algorithm for sparse graphs. Combined, our two results significantly improve upon all previous work.

Theorem 1.2. Given a decremental input graph G = (V, E, w) with n = |V|, m = |E| and aspect ratio W, a dedicated source $r \in V$ and $\epsilon > 0$, there is a randomized algorithm that maintains a distance estimate $\widetilde{\operatorname{dist}}(r, x)$, for every $x \in V$, such that

$$\operatorname{dist}_G(r, x) \le \operatorname{dist}(r, x) \le (1 + \epsilon) \operatorname{dist}_G(r, x)$$

at any stage w.h.p. The algorithm has total expected update time $\tilde{O}(mn^{2/3}\log^3 W/\epsilon)$. Distance queries are answered in O(1) time, and a corresponding path P can be returned in O(|P|) time.

Adaptive Versus Oblivious Adversaries Both our results are randomized, and assume an oblivious adversary whose update sequence does not depend on the distance estimates and query-paths returned by the algorithm. This assumption is also necessary for all previous algorithms in directed graphs that go beyond the classic O(mn) bound of Even and Shiloach [HKN14b; HKN15; GW20a], and achieving such a result adaptively remains a major open problem. (The result of [GW20a] allows for adaptive distance queries, but still assumes the adversary is oblivious to the paths returned by the algorithm.) More generally, achieving truly fast adaptive algorithms seems implausible until we understand the problem well enough to at least have fast oblivious algorithms, and our result makes significant progress on this front. In fact, arguably more than any other dynamic graph algorithm, our $\tilde{O}(n^2)$ result highlights the gap between these two models: if one could design an *adaptive* algorithm that matches this efficiency, then plugging it into the existing framework of [CK19] (based in turn on the multiplicative-weight update method as described in [Fle00; GK07; Mad10]) would yield an $\tilde{O}(n^2)$ algorithm for *static* directed maximum flow, which would constitute an enormous breakthrough in the field of graph algorithms. **Techniques** The key contribution of our paper is a general technique for converting algorithms on directed acyclic graphs (DAGs) into algorithms on general graphs. Earlier techniques in [Ber17; GW20a] lead to two simple algorithms for DAGs with total update times $\tilde{O}(n^2)$ and $\tilde{O}(mn^{2/3})$; our conversion then extends these bounds to general graphs. We first introduce the concept *approximate topological order* (\mathcal{ATO}), which loosely speaking imposes a DAG-like structure on *any* graph. We then show that an \mathcal{ATO} always exists and can be maintained efficiently.

At a high-level, the conversion is as follows. Let \mathcal{A}_{DAG} be a decremental SSSP algorithm for DAGs and let $T(\mathcal{A}_{DAG})$ be the total update time. The first (easier) step is to convert \mathcal{A}_{DAG} to an algorithm \mathcal{A}^*_{DAG} that works on any graph with an \mathcal{ATO} and has total update time $T(\mathcal{A}^*_{DAG}) \sim T(\mathcal{A}_{DAG})$. The second (harder) step is to build an algorithm \mathcal{A}_{ATO} that maintains an \mathcal{ATO} in G by recursively applying \mathcal{A}^*_{DAG} as a subroutine: the basic idea is that an \mathcal{ATO} of better "quality" can be built by using \mathcal{A}^*_{DAG} to maintain shortest paths in an \mathcal{ATO} of worse quality. Using this layered approach, we can achieve $T(\mathcal{A}_{ATO}) \sim T(\mathcal{A}^*_{DAG}) \sim T(\mathcal{A}_{DAG})$, and combining \mathcal{A}_{ATO} with \mathcal{A}^*_{DAG} gives an algorithm for general graphs.

Neither the step from \mathcal{A}_{DAG} to \mathcal{A}_{DAG}^* nor the step from \mathcal{A}_{DAG}^* to \mathcal{A}_{ATO} are black-box, but the techniques are quite modular and flexible, as evidenced by the fact that we were able to apply this conversion to both of the state-of-the-art algorithms for DAGs. We believe our conversion has strong potential to influence future work on directed shortest paths, in both the dynamic model and in others, by allowing researchers to focus on the simpler case of DAGs.

2 Preliminaries

We let a graph H refer to a weighted, directed graph with vertex set denoted by V(H), edge set E(H) and weight function $w_H : E(H) \to [1, W] \cup \{\infty\}^2$. We say that H is a *decremental* graph if it is undergoing a sequence of edge deletions and edge weight increases (also referred to as updates), and refer to *version* t of H, or H at *stage* t as the graph H obtained after the first t updates have been applied. In this article, we denote the (decremental) input graph by G = (V, E, w) with n = |V| and m = |E| (where m refers to the number of edges of G at stage 0). In all subsequent definitions, we often use a subscript to indicate which graph we refer to, however, when we refer to G, we often omit the subscript.

Cuts, Neighborhoods and Subgraphs. For graph H, and any two disjoint subsets $X, Y \subseteq V(H)$, we let $E_H(X)$ be the set of edges in E(H) with an endpoint in X, and $E_H(X, Y)$ denote the set of edges in E(H) with tail in X and head in Y. We let $\mathcal{N}_H^{in}(v) = \{u \mid (u, v) \in E(H)\}$ and $\mathcal{N}_H^{out}(v) = \{u \mid v \in \mathcal{N}_H^{in}(u)\}$ denote the in-neighborhood and out-neighborhoods of $v \in V$. We further let H[X] refer to the subgraph of H induced by X, i.e. $H[X] = (X, E_H(X, X), w_H)$. We use $H \subseteq G$ to denote that V(H) = V(G) and $E(H) \subseteq E(G)$.

Contractions. We define the graph H/X to be the graph obtained from H by contracting all vertices in X into a single node (we use the word *node* instead of vertex if it was obtained

²In this article, we restrict our attention to integers and let [a, b] denote the set of integers $a, a+1, a+2, \ldots, b$.

by contractions). Similarly, for a set of pairwise disjoint vertex sets X_1, X_2, \ldots, X_k , we let $H/\{X_1, X_2, \ldots, X_k\}$ denote the graph $((((H/X_1)/X_2) \ldots)/X_k)$. If \mathcal{V} forms a partition of V, we use the convention to denote by X^v the node in G/\mathcal{V} that contains $v \in V$, i.e. $v \in X^v$.

Reachability and Strong-Connectivity. For graph H and any two vertices $u, v \in V(H)$, we let $u \rightsquigarrow_H v$ denote that u can reach v in H, and $u \rightleftharpoons_H v$ that u can reach v and vice versa (in the latter case, we also say u and v are strongly-connected). For any sets $X, Y \subseteq V(H)$, we say that $X \rightsquigarrow_H Y$ if there exists some $x \in X, y \in Y$ such that $x \rightsquigarrow_H y$; we define $X \rightleftharpoons_H Y$ analogously. We say that the partition of V(H) induced by the equivalence relation \rightleftharpoons_H is the set of strongly-connected components (SCCs).

Generalized Topological Order. We define a generalized topological order

GENERALIZEDTOPORDER(H) to be a tuple (\mathcal{V}, τ) where \mathcal{V} is the set of SCCs of H and $\tau : \mathcal{V} \to [0, n)$ is a function that maps any sets $X, Y \in \mathcal{V}$ such that $\tau(X) < \tau(Y)$, if $X \rightsquigarrow_H Y$ and such that $[\tau(X), \tau(X) + |X|) \cap [\tau(Y), \tau(Y) + |Y|) = \emptyset$. Thus, τ effectively establishes a one-to-one correspondence between |X|-sized intervals and SCCs X in H. We point out that a GENERALIZEDTOPORDER(H) can always be computed in O(|E(H)|) time [Tar72]. In fact, a generalized topological order can also be maintained efficiently in a decremental graph H. Here, we say that (\mathcal{V}, τ) is a dynamic tuple that forms a generalized topological order for all versions of H. Further, we say that (\mathcal{V}, τ) has the *nesting* property, if for any set $X \in \mathcal{V}$ and a set $Y \supseteq X$ that was in \mathcal{V} at an earlier stage, we have $\tau(X) \in [\tau(Y), \tau(Y) + |Y| - |X|]$; in other words, the interval $[\tau(X), \tau(X) + |X|)$ is entirely contained in the interval $[\tau(Y), \tau(Y) + |Y|)$. Thus, the associated interval with X is contained in the interval associated with Y. We refer to the following result that can be obtained straight-forwardly by combining the data structure given in [BPW19] and the static procedure by Tarjan [Tar72] as described in [GW20a].

Theorem 2.1 (see [Tar72; BPW19; GW20a]). Given a decremental digraph H, there exists an algorithm that can maintain the generalized topological order (\mathcal{V}, τ) of H where τ has the nesting property. The algorithm runs in expected total update time $O(m \log^4 n)$, is randomized and works against an adaptive adversary.

Distances and Diameter. We let $\operatorname{dist}_H(u, v)$ denote the distance from vertex u to vertex v in graph H and denote by $\pi_{u,v,H}$ the corresponding shortest path (we assume uniqueness by implicitly referring to the lexicographically shortest path). We define the weak diameter of $X \subseteq V(H)$ in H by $\operatorname{diam}(X, H) = \max_{x, y \in X} \operatorname{dist}_H(x, y)$.

Exponential Distribution. Finally, we make use of the exponential distribution, that is we use random variables X with cumulative distribution function $F_X(x,\lambda) = 1 - e^{-\lambda x}$ for all $x \ge 0, \lambda > 0$, which we denote by the shorthand $X \sim \text{Exp}(\lambda)$. If $X \sim \text{Exp}(\lambda)$ is clear, we also use $F_X(x)$ in place of $F_X(x,\lambda)$. The exponential distribution has the special property of being *memoryless*, that is if $X \sim \text{Exp}(\lambda)$, then

$$\mathbb{P}[X > s + t \mid X > t] = \mathbb{P}[X > s].$$

3 Overview

We now give an overview of our algorithm. In order to illustrate the main concepts, we start by giving a simple algorithm to obtain total update time $O(n^2 \log n/\epsilon)$ in directed acyclic graphs (DAGs). While this algorithm was previously not explicitly mentioned, it follows rather directly from the techniques developed in [Ber17; GW20a]. We present this algorithm to provide intuition for our approach and motivates our novel notion of approximate topological orders. In the light of approximate topological orders, we then shed light on limitations of the previous approach in [GW20a] and present techniques to surpass these limitations to obtain Theorem 1.1. In this overview, we focus on obtaining an SSSP algorithm that runs in $\tilde{O}(n^2 \log^4 W/\epsilon)$ expected update time. The final paragraph then sketches our improvement for sparse graphs (Theorem 1.2); this result combines our directed framework in a non-trivial way with ideas from an earlier result for sparse SSSP in *undirected, unweighted* graphs [BC17].

3.1 A Fast Algorithm for DAGs

The Topological Order Difference. Let G = (V, E, w) be a DAG and let τ be the function returned by GENERALIZEDTOPORDER(G) computed on the initial version of G (since G is a DAG, this is just a standard topological order). Let us now make an almost trivial observation: for any shortest s-to-t path $\pi_{s,t}$, in any version of G, the sum of topological order differences is bounded by n. More formally:

$$\mathcal{T}(\pi_{s,t},\tau) \stackrel{\text{def}}{=} \sum_{(u,v)\in\pi_{s,t}} \tau(v) - \tau(u) = \tau(t) - \tau(s) \le n.$$
(1)

Observe that every path in G can only contain few edges (u, v) with large topological order difference, i.e. with $\tau(v) - \tau(u)$ large, by the pigeonhole principle.

Reviewing the ES-tree. To understand how this fact can be exploited, let us review the classic ES-tree (see [ES81]): We maintain distances from a fixed source $r \in V$ to each vertex v in V up to distance $\delta > 0$ by storing a distance estimate $\operatorname{dist}(r, v)$ that is initialized to the distance between r and v in G along with a shortest-path tree T rooted at r. On update (u, v), we delete the edge from G and possibly from T. Then, if possible, we extract $w_{min} \in V \setminus \{r\}$, the vertex with smallest distance estimate among vertices without incoming edge in T. (For the first extraction, we always have $w_{min} = v$.) For this vertex w_{min} , we then try to find a vertex $x \in \mathcal{N}^{in}(w_{min})$ such that adding (x, w_{min}) to T implies $\operatorname{dist}_T(r, w_{min}) \leq \operatorname{dist}(r, w_{min})$. We therefore search in $\mathcal{N}^{in}(v)$ for an x that satisfies

$$\operatorname{dist}(r, x) + w(x, w_{\min}) \le \operatorname{dist}(r, w_{\min}).$$

$$\tag{2}$$

If no such x exists, then $\widetilde{\operatorname{dist}}(r, w_{min})$ has to be incremented, and we set T to $T \setminus N^{out}(w_{min})$. If $\widetilde{\operatorname{dist}}(r, w_{min}) > \delta$, we set it to ∞ and remove w_{min} from the tree. We iterate the process until T is spanning for vertices with distance estimate $< \infty$. Since for each distance estimate value $\widetilde{\operatorname{dist}}(r, v)$, the in-neighbourhood $\mathcal{N}^{in}(v)$ has to be scanned once (in an efficient implementation), and since estimates increase monotonically, the total update time can be bound by $O(\sum_{v \in V} |\mathcal{N}^{in}(v)|\delta) = O(m\delta)$. Further, observe that if a vertex v was only allowed to scan a certain vertex $x \in \mathcal{N}^{in}(v)$ every i distance estimates (for example whenever $\widetilde{\operatorname{dist}}(r, v)$ is divisible by i) then this corresponds to enforcing that at all times $\widetilde{\operatorname{dist}}(r, x) + w(x, v) + i - 1 \leq \widetilde{\operatorname{dist}}(r, v)$ since we check equation 2 every i steps (in particular, for i = 1, we get an exact algorithm). Consequently, we get at most i - 1 additive error in the distance estimate for any t whose shortest r-to-t path $\pi_{r,t}$ contains (x, v). On the other hand, we only need to scan and check the edge (x, v), by the argument above, δ/i times³.

Improving the Running Time. Let us now exploit Inequality 1. We therefore define

$$B_j(v) = \{ u \in \mathcal{N}^{in}(v) \text{ with } 2^j \le \tau(v) - \tau(u) < 2^{j+1} \}$$

for every $v \in V$ and $0 \leq j \leq \lg n$; the $B_j(v)$ partition the in-neighborhood of v according to topological order difference to v. Observe that $|B_j(v)| \leq 2^j$. Now, consider the algorithm as above where for every vertex v, instead of checking all $\mathcal{N}^{in}(v)$, we only check edge (x, v) for $x \in B_j(v)$ if $\operatorname{dist}(r, v)$ is divisible by $\lfloor 2^j \frac{\epsilon \delta}{n} \rfloor$. By the arguments above the total running time now sums to

$$O\left(\sum_{v \in V} \sum_{0 \le j \le \lg n} |B_j(v)| \frac{\delta}{2^j \frac{\epsilon \delta}{n}}\right) = O\left(\sum_{v \in V} \sum_{0 \le j \le \lg n} 2^{j+2} \frac{\delta}{2^j \frac{\epsilon \delta}{n}}\right) = \tilde{O}(n^2/\epsilon).$$

Bounding the Error. Fix a shortest r-to-t path $\pi_{r,t}$, and consider any edge $(u, v) \in \pi_{r,t}$ with $u \in B_{j+1}(v)$. We observe that the edge (u, v) contributes at most an additive error of $2^{j+1}\frac{\epsilon\delta}{n}$ to $\widetilde{\operatorname{dist}}(r,t)$ since it is scanned every $\lceil 2^j \frac{\epsilon\delta}{n} \rceil$ distance values and if $\lceil 2^j \frac{\epsilon\delta}{n} \rceil$ is equal to 1 it does not induce any error.

On the other hand, since $u \in B_{j+1}(v)$ we also have $\tau(v) - \tau(u) \ge 2^j$. We can thus charge $n/(2\epsilon\delta)$ units from $\mathcal{T}(\pi_{r,t},\tau)$ for each additive error unit; we know from Equation 1 that $\mathcal{T}(\pi_{r,t},\tau) \le n$, so the total additive error is at most $\frac{n}{n/(2\epsilon\delta)} = 2\epsilon\delta$. Thus, for all distances $\approx \delta$ (say in $[\delta/2, \delta)$), we obtain a $(1 + 4\epsilon)$ -multiplicative distance estimate⁴.

Working with Multiple Distance Scales. Observe that the data structure above has no running time dependency on δ . Thus, to obtain a data structure that maintains a $(1 + 2\epsilon)$ -approximate distance estimate from r to any vertex x, we can simply use $\lg(nW)$ data structures in parallel where the i^{th} data structure has $\delta = 2^i$. A query can then be answered by returning the smallest distance estimate from any data structure, using a min-heap data structure to obtain this smallest estimate in constant time⁵. The running time for all data structures is then bounded by $\tilde{O}(n^2 \log W/\epsilon)$.

³This trade-off was first observed in [Ber17].

⁴Technically, we run to depth $(1+4\epsilon)\delta$ to ensure that vertices' distance estimates are not set to ∞ too early.

⁵Here, we exploit that all distance estimates are overestimates, and at least one of them is $(1 + 2\epsilon)$ -approximate.

3.2 Extending the Result to General Graphs

We now encourage the reader to verify that in the data structure for DAGs, we used at no point that the graph was acyclic, but rather only used that $\mathcal{T}(\pi_{s,t},\tau)$ is bounded by n for any path $\pi_{s,t}$. In this light, it might be quite natural to ask whether such a function τ might exist for general decremental graphs. Surprisingly, it turns out that after carrying out some contractions in G that only distort distances slightly, we can find such a function τ that comes close in terms of guarantees. We call such a function τ an *approximate topological order* (this function will no longer encode guarantees about reachability, but it helps for intuition to think of τ as being similar to a topological order).

The Approximate Topological Order We start with the formal definition:

Definition 3.1. Given a decremental weighted digraph G = (V, E, w) and parameter $\eta_{diam} \geq 0$. We call a dynamic tuple (\mathcal{V}, τ) an approximate topological order of G of quality q > 1(abbreviated $\mathcal{ATO}(G, \eta_{diam})$ of quality q), if at any stage

- 1. $\mathcal{V} = \{X_1, X_2, ..., X_k\}$ forms a partition of V and a refinement of all earlier versions of \mathcal{V} , and
- 2. $\tau: \mathcal{V} \to [0, n)$ is a function that maps each $X \in \mathcal{V}$ to a value $\tau(X)$. If some set $X \in \mathcal{V}$ is split at some stage into disjoint subsets $X_1, X_2, ..., X_k$, then we let $\tau(X_{\pi(1)}) = \tau(X)$ and $\tau(X_{\pi(j+1)}) = \tau(X_{\pi(j)}) + |X_{\pi(j)}|$ for each j < k and some permutation π of [1, k], and
- 3. each $X \in \mathcal{V}$ has weak diameter $\operatorname{diam}(X,G) \leq \frac{|X|\eta_{diam}}{n}$, and
- 4. for any two vertices $s, t \in V$, the shortest path $\pi_{s,t}$ in G satisfies $\mathcal{T}(\pi_{s,t},\tau) \leq q \cdot w(\pi_{s,t}) + n$ where we define $\mathcal{T}(\pi_{s,t},\tau) \stackrel{\text{def}}{=} \sum_{(u,v) \in \pi_{s,t}} |\tau(X^u) - \tau(X^v)|.$

We say that (\mathcal{V}, τ) is an $\mathcal{ATO}(G, \eta_{diam})$ of expected quality q, if (\mathcal{V}, τ) satisfies properties 1-3, and at any stage, for every $s, t \in V$, $\mathbb{E}[\mathcal{T}(\pi_{s,t}, \tau)] \leq q \cdot w(\pi_{s,t}) + n$.

Let us expound the ideas captured by this definition. We remind the reader that such a function τ is required by a data structure that only considers distances in $[\delta/2, \delta)$. Let us consider a tuple (\mathcal{V}, τ) that forms an $\mathcal{ATO}(G, \epsilon \delta)$ of quality q. Then, for any *s*-to-*t* shortest path $\pi_{s,t} = \langle s = v_1, v_2, \ldots v_\ell = t \rangle$ in G, let s_i and t_i be the first and last vertex on the path in $X_i \in \mathcal{V}$ (see property 1) if there are any. Observe that by property 3, the vertices s_i and t_i are at distance at most $\frac{|X_i|\epsilon \delta}{n}$ in G. It follows that if we contract the SCC X_i , the distance $\operatorname{dist}_{G/X_i}(s,t)$ is at least the distance from s to t in G minus an additive error of at most $\frac{|X_i|\epsilon \delta}{n}$. It follows straight-forwardly, that after contracting all sets in \mathcal{V} , we have that distances in G/\mathcal{V} correspond to distances in G up to a negative additive error of at most $\sum_{X_i \in \mathcal{V}} \frac{|X_i|\epsilon \delta}{n} = \frac{n \cdot \delta \epsilon}{n} = \epsilon \delta$. Thus, maintaining the distance sin G/\mathcal{V} $(1 + 2\epsilon)$ -approximately is still sufficient for getting a $(1 \pm 2\epsilon)$ -approximate distance estimate.

Property 1 simply ensures that the vertex sets forming the elements of \mathcal{V} decompose over time. Property 2 states that τ assigns each node in G/\mathcal{V} a distinct number in [0, n). It also ensures that if a set $X \in \mathcal{V}$ receives $\tau(X)$ that every later subset of X will obtain a number in the interval $[\tau(X), \tau(X) + |X|)$. Moreover, τ effectively establishes a one-to-one correspondence between nodes in a version of G/\mathcal{V} and intervals in [0, n) of size equal to their underlying vertex set. Once a set $X \in \mathcal{V}$ decomposes into sets X_1, X_2, \ldots , property 2 stipulates that the intervals that τ maps X_1, X_2, \ldots to are disjoint subintervals of $[\tau(X), \tau(X) + |X|)$. We point out that once \mathcal{V} consists of singletons, each vertex is essentially assigned a single number.

Finally, property 4 gives an upper bound on the topological order difference. Observe that we redefine $\mathcal{T}(\pi_{s,t},\tau)$ in a way that is consistent with Definition 1. In our algorithm, for $\eta_{diam} \approx \epsilon \delta$, we obtain a quality of $\tilde{O}(n/\epsilon \delta)$. Thus, any path π of weight $\approx \delta$ has $\mathcal{T}(\pi/\mathcal{V},\tau) \leq \tilde{O}(n/\epsilon)$ which is very close to the upper bound obtained by the topological order function in DAGs. We summarize this result in the theorem below which is one of our main technical contributions.

Theorem 3.2. [see Section 4 and Section 5.2.] For any $0 \le i \le \lg(Wn)$, given a decremental digraph G = (V, E, w), we can maintain an $\mathcal{ATO}(G, 2^i)$ of expected quality $\tilde{O}(n/2^i)$. The algorithm runs in total expected update time $\tilde{O}(n^2)$ against a non-adaptive adversary with high probability.

Combining the theorem above and the theorem below which is obtained by generalizing the above decremental SSSP algorithm for DAGs, we obtain our main result Theorem 1.1.

Theorem 3.3. [see Section 5.1.] Given G = (V, E, w) and (\mathcal{V}, τ) an $\mathcal{ATO}(G, \eta_{diam} \approx \epsilon \delta)$, for some depth parameter $\delta > 0$, of quality q, a dedicated source r in V, and an approximation parameter $\epsilon > 0$. Then, there exists a deterministic data structure \mathcal{E}_r that maintains a distance estimate $\widetilde{\operatorname{dist}}(r, v)$ for each $v \in V$, which is guaranteed to be $(1+\epsilon)$ -approximate if $\operatorname{dist}(r, v) \in$ $[\delta, 2\delta)$. Distance queries are answered in O(1) time and a corresponding path P can be returned in O(|P|) time. The total update time is $\tilde{O}(n\delta q/\epsilon + n^2)$.

3.3 The Framework by [GW20a]

Before we describe our new result, we review the framework in [GW20a] to construct and maintain an $\mathcal{ATO}(G, \eta_{diam})$. We point out that while the abstraction of an approximate topological order is new to our paper, analyzing the technique in [GW20a] through the \mathcal{ATO} lens is straight-forward and gives a first non-trivial result. Throughout this review section, we assume that the graph G is unweighted to simplify presentation. This allows us to make use of the following result which states that for vertex sets that are far apart, one can find deterministically a vertex separator that is small compared to the smaller side of the induced partition (to obtain an algorithm for weighted graphs a simple edge rounding trick is sufficient to generalize the ideas presented below).

Definition 3.4. Given graph G = (V, E, w), then we say a partition of V into sets A, S_{Sep}, B is a one-way vertex separator if $A \nleftrightarrow_{G \setminus S_{Sep}} B$ and A and B are non-empty.

Lemma 3.5 (see Definition 5 and Lemma 6 in [Che+16]). Given an unweighted graph G of diameter diam(G). Then we can find sets A, S_{Sep}, B that form a one-way vertex separator such that $|S_{Sep}| \leq \tilde{O}(\frac{\min\{|A|,|B|\}}{\operatorname{diam}(G)})$ in time O(m).

High-level Framework. The main idea of [GW20a] is to maintain a tuple (\mathcal{V}, τ) which is an $\mathcal{ATO}(G, \eta_{diam})$ by setting (\mathcal{V}, τ) to be GENERALIZEDTOPORDER(G') of some decremental graph $G' \subseteq G$ (over the same vertex set, i.e. V(G') = V(G)). It is straight-forward to see that (\mathcal{V}, τ) satisfies property 1 in Definition 3.1, since SCCs in the decremental graph G'decompose. Further, it is not hard to extend the existing algorithm for maintaining SCCs in a decremental graph G' given in [BPW19] to also maintain function τ that obeys property 2 in Definition 3.1. The algorithm to maintain (\mathcal{V}, τ) given G' runs in total update time $\tilde{O}(m)$ (same as in [BPW19]).

So far, we have not given any reason why G' needs to be a subgraph of G. To see why we cannot use the above strategy on G directly, recall property 3 in the \mathcal{ATO} -definition 3.1, which demands that each SCC X has weak diameter at most $\frac{|X|\eta_{diam}}{n}$. This property might not hold in the main graph G. In order to resolve this issue, G' is initialized to G and then the diameter of SCCs in G' is monitored. Whenever an SCC X violates property 3, a vertex separator S_{Sep} is found in the graph G'[X] as described in Lemma 3.5 and all edges incident to S_{Sep} are removed from G'. Letting S denote the union of all such separators S_{Sep} , we can now write $G' = G \setminus E(S)$.

Establishing the Quality Guarantee. To establish a quality q of the \mathcal{ATO} as described in property 4 in Definition 3.1, let us first partition the set S into sets $S_0, S_1, \ldots, S_{\lg n}$ where each S_i contains all separator vertices found on a graph of size $[n/2^i, n/2^{i+1})$, thus it was found when the procedure from Lemma 3.5 was invoked on a graph with diameter at least $\frac{(n/2^{i+1})\eta_{diam}}{n} = \frac{\eta_{diam}}{2^{i+1}}$. Since separators are further balanced, i.e. there size is controlled by the smaller side of the induced cut, we can further use induction and Lemma 3.5 to establish that there are at most $\tilde{O}(\frac{2^{i}n}{\eta_{diam}})$ vertices in S_i . Next, observe that since any separator that was added to S_i was found in a graph G'[X] with $|X| \in [n/2^i, n/2^{i+1})$, we have by property 2 that nodes $X' \subseteq X$ that are in the current version of \mathcal{V} are assigned a $\tau(X')$ from the interval $[\tau(X), \tau(X) + |X|)$. Thus, any edge (x, s) or (s, x) with $x \in X, s \in S_i \cap X$ has topological order difference $|\tau(X^x) - \tau(X^s)| \leq |X| \leq n/2^i$.

Finally, let us define

$$\mathcal{T}'(\pi_{s,t},\tau) \stackrel{\text{def}}{=} \sum_{(u,v)\in\pi_{s,t}} \min\{0,\tau(X^v) - \tau(X^u)\}$$
(3)

the function similar to $\mathcal{T}(\pi_{s,t},\tau)$ that only captures negative terms, i.e. sums only over edges that go "backwards" in τ . Observe that $\mathcal{T}(\pi_{s,t},\tau) \leq 2\mathcal{T}'(\pi_{s,t},\tau) + n$. Now, since (\mathcal{V},τ) is a GENERALIZEDTOPORDER(G'), we have that (u,v) occurs in the sum of $\mathcal{T}'(\pi_{s,t},\tau)$ if and only if $(u,v) \in G \setminus G'$, so one endpoint is in a set S_i and therefore $\tau(X^v) - \tau(X^u) \leq n/2^i$. Since we only have two edges on any shortest path incident to the same vertex, we can establish that

$$\mathcal{T}'(\pi_{s,t},\tau) \le \sum_{i} 2|S_i|n/2^i = \tilde{O}(n^2/\eta_{diam}).$$

We obtain that (\mathcal{V}, τ) is a $\mathcal{ATO}(G, \eta_{diam})$ of quality $\tilde{O}(\frac{n^2}{2^i \eta_{diam}})$ for all paths of weight at least 2^i . Thus, when the distance scale $\delta \geq \sqrt{n}/\epsilon$, Theorem 3.3 requires total update time $\approx n^{2.5}$

to maintain $(1 + \epsilon)$ -approximate SSSP. For distance scales where $\delta < \sqrt{n}/\epsilon$, a classic ES-tree has total update time $\approx m\sqrt{n} \leq n^{2.5}$.

Limitations of the Framework. Say that the goal is to maintain shortest paths of length around \sqrt{n} . The first step in the framework of [GW20a] is to find separator S such that all SCCs of $G' = G \setminus E(S)$ have diameter at most $\epsilon \sqrt{n}$ and then maintain $(\mathcal{V}, \tau) = \text{GENERALIZED-TOPORDER}(G')$. Every edge $(u, v) \notin E(S)$ will only go forward in τ , but each edge (u, s), for $s \in S$, can go "backwards" in τ . By the nesting property of generalized topological orders, the amount that (u, s) goes backwards – i.e. the quantity $|\tau(X^u) - \tau(X^s)|$ – is upper bounded by the size of the SCC in G' from which s was chosen: the original SCC has size n, but as we add vertices to S, the SCCs of $G' = G \setminus E(S)$ decompose and new vertices added to S may belong to smaller SCCs. Define $S^* \subseteq S$ to contain all vertices $s \in S$ that were chosen in an SCC of size $\Omega(n)$. Intuitively, S^* is the top-level separator chosen in G, before SCCs decompose into significantly smaller pieces. Every edge in $E(S^*)$ may go backwards by as much as n in τ , so for any path $\pi_{x,y}$ in G, the best we can guarantee is that $\mathcal{T}(\pi_{x,y}) \sim n \cdot |\pi_{x,y} \cap S^*|$.

The framework of [GW20a] tries to find a *small* separator S^* and then uses the trivial upper bound $|\pi_{x,y} \cap S^*| \leq |S^*|$. In fact, one can show that given any deterministic separator procedure, the adversary can pick a sequence of updates where $|\pi_{x,y} \cap S^*| \sim |S^*|$. But now, say that G is a $\sqrt{n} \times \sqrt{n}$ -grid graph with bidirectional edges. It is not hard to check that $|S^*| = \Omega(\sqrt{n})$, because every balanced separator of a grid has $\Omega(\sqrt{n})$ vertices. The framework of [GW20a] can thus at best guarantee $\mathcal{T}(\pi_{x,y}) \sim n \cdot |\pi_{x,y} \cap S^*| \sim n|S^*| = \Omega(n^{1.5})$, which is a \sqrt{n} factor higher than it would be in a DAG, and thus leads to running time $\tilde{O}(n^{2.5})$ instead of $\tilde{O}(n^2)$.

Our algorithm uses an entirely different random separator procedure. We allow S^* to be arbitrarily large, but use randomness to ensure that $|\pi_{x,y} \cap S^*|$ is nonetheless small.

3.4 Our Improved Framework

We now introduce our new separator procedure and then show how it can be used in a recursive algorithm that uses ATOs of worse quality (large q) to compute ATOs of better quality (small q). (By contrast, the framework of [GW20a] could not benefit from a multi-layered algorithm because it would still hit upon the fundamental limitation outlined above.)

A New Separator Procedure. Before we describe the separator procedure, let us formally define the guarantees that we obtain. In the lemma below, think of $\zeta = \Theta(\log(n))$.

Lemma 3.6. There exists a procedure OUTSEPARATOR (r, G, d, ζ) where G is a weighted graph, $r \in V$ a root vertex, and integers $d, \zeta > 0$. Then, with probability at least $1 - e^{-\zeta}$, the procedure computes a tuple (E_{Sep}, V_{Sep}) where edges $E_{Sep} \subseteq E$, and vertices $V_{Sep} = \{v \in V | r \rightsquigarrow_{G \setminus E_{Sep}} v\}$ such that

- 1. for every vertex $v \in V_{Sep}$, $\operatorname{dist}_{G \setminus E_{Sep}}(r, v) \leq d$, and
- 2. for every $e \in E$, we have $\mathbb{P}[e \in E_{Sep} | r \rightsquigarrow_{G \setminus E_{Sep}} \mathbf{tail}(e)] \leq \frac{\zeta}{d} w(e)$.

Otherwise, it reports Fail. The running time of OUTSEPARATOR(\cdot) can be bounded by $O(|E(V_{Sep})|\log n)$.

In fact, Algorithm 1 gives a simple implementation of procedure OUTSEPARATOR(·). Here, we pick a ball $B = B^{out}(r, X)$ in the graph G from r to random depth X, and then simply return the tuple $(E_{Sep}, V_{Sep}) = (E(B, \overline{B}), B)$ where $E(B, \overline{B})$ are the edges (u, v) with $u \in B$ but $v \notin B$. The procedure thus only differs from a standard edge separator procedure in that we choose X according to the exponential distribution $\text{Exp}(\frac{\zeta}{d})$.

Algorithm 1: OUTSEPARATOR (r, G, d, ζ)

1 Choose $X \sim \text{Exp}(\frac{\zeta}{d})$.

```
2 if X \ge d then return Fail
```

- **3** Compute the Ball $B = B^{out}(r, X) = \{v \in V | \mathbf{dist}(r, v) \le X\}$
- 4 return $(E(B,\overline{B}),B)$

A proof of Lemma 3.6 is now straightforward. We return **Fail** in Line 2 with probability $\mathbb{P}[X \ge d] = 1 - F_X(d) = 1 - (1 - e^{-\frac{\zeta}{d} \cdot d}) = e^{-\zeta}$ (recall from section 2 that $F_X(d)$ is shorthand for $F(x, \frac{\zeta}{d})$, the cumulative distribution function of an exponential distribution with parameter $\frac{\zeta}{d}$). Assuming no failure, we have $E_{Sep} = E(B, \overline{B})$, and it is easy to see that $V_{Sep} = \{v \in V | r \rightsquigarrow_{G \setminus E_{Sep}} v\} = B$, so Property 1 of Lemma 3.6 holds by definition of B. Moreover, we can compute B in the desired $O(|E(B)|\log n)$ time by using Dijkstra's algorithm by only extracting a vertex from the heap if it is at distance at most X. Finally, for property 2, note that $e \in E_{sep}$ iff $\operatorname{dist}(r, \operatorname{tail}(e)) \le X < \operatorname{dist}(r, \operatorname{tail}(e)) + w(e)$. Thus,

$$\mathbb{P}[e \in E_{Sep} | r \rightsquigarrow_{G \setminus E_{Sep}} \mathbf{tail}(e)] = \mathbb{P}[X < \mathbf{dist}(r, \mathbf{tail}(e)) + w(e) | X \ge \mathbf{dist}(r, \mathbf{tail}(e))]$$
$$= \mathbb{P}[X < w(e)] = F_X(w(e)) = 1 - e^{-\frac{\zeta}{d}w(e)}$$
$$\le 1 - \left(1 - \frac{\zeta}{d}w(e)\right) = \frac{\zeta}{d}w(e)$$

where the second equality follows from the memory-less property of the exponential distribution, and the inequality holds because $1 + x \leq e^x$ for all $x \in \mathbb{R}$. We point out that the technique of random ball growing using the exponential distribution is not a novel contribution in itself and has be previously used in the context of low-diameter decompositions [LS93; Bar96; MPX13; Pac+18] which have recently also been adapted to dynamic algorithms [FG19; CZ20].

A New Framework. Let us now outline how to use Lemma 3.6 to derive Theorem 3.2 which is stated below again. The full details and a rigorous proof are provided in Section 4.

Theorem 3.2. [see Section 4 and Section 5.2.] For any $0 \le i \le \lg(Wn)$, given a decremental digraph G = (V, E, w), we can maintain an $\mathcal{ATO}(G, 2^i)$ of expected quality $\tilde{O}(n/2^i)$. The algorithm runs in total expected update time $\tilde{O}(n^2)$ against a non-adaptive adversary with high probability.

As in [GW20a], we maintain a graph $G' \subseteq G$ and its generalized topological order (\mathcal{V}, τ) . Whenever the diameter of an SCC X in G' is larger than $\frac{|X|\eta_{diam}}{n}$, we now use the separator procedure described in Lemma 3.6 with $d = \frac{|X|\eta_{diam}}{2n}$ from some vertex r in X with $|B^{out}(r, d = \frac{|X|\eta_{diam}}{2n})| \leq |X|/2$. Such a vertex exists since by definition of diameter, as we can find two vertices with disjoint balls. We obtain an edge separator E_{Sep} and update G' by removing the edges in E_{Sep} . Let the union of all edge separators be denoted by F and again observe that $G' = G \setminus F$. It is not hard to see that our scheme still ensures properties 1-3 in Definition 3.1. We now argue that the quality improved to $\tilde{O}(n/2^i)$.

Partition F into sets $F_0, F_1, \ldots, F_{\lg n}$ where each F_j contains all separator edges found on a graph of size $[n/2^{j+1}, n/2^j)$. We again have that every edge $(u, v) \in F_j$ has $|\tau(X^u) - \tau(X^v)| \leq n/2^j$. Let us establish an upper bound on the number of edges in F_j on any shortest path $\pi_{s,t}$. Consider therefore any edge $e \in \pi_{s,t}$, at a stage where both endpoints of e are in a SCC X of size $[n/2^{j+1}, n/2^j)$ and where we compute a tuple (V_{Sep}, E_{Sep}) from some $r \in X$. Now, observe that if $\operatorname{tail}(e) \notin V_{Sep}$, then e can not be in E_{Sep} . So assume that $\operatorname{tail}(e) \in V_{Sep}$. Then, we have e joining F_j with probability

$$\mathbb{P}[e \in E_{Sep} | r \rightsquigarrow_{H \setminus E_{Sep}} \mathbf{tail}(e)] = \frac{\zeta}{2^i |X|/n} w(e) = \tilde{O}\left(\frac{2^j w(e)}{2^i}\right)$$

according to Lemma 3.6, where we set $\zeta = \tilde{O}(1)$ to obtain high success probability. But if e did not join F_j at that stage, then it is now in a SCC of size at most $n/2^{j+1}$ (recall we chose $|B^{out}(r,d)| \leq |X|/2$, and we have $V_{Sep} \subseteq B^{out}(r,d)$). Thus, e cannot join F_j at any later stage.

Now, it suffices to sum over edges on the path $\pi_{s,t}$ and indices j to obtain that

$$\mathcal{T}'(\pi_{s,t},\tau) \le \sum_{e \in \pi_{s,t}} \sum_{j} \mathbb{P}[e \in F_j] \cdot n/2^j = \sum_{e \in \pi_{s,t}} \sum_{j} \tilde{O}\left(\frac{2^j w(e)}{2^i}\right) \cdot n/2^j = \tilde{O}\left(\frac{w(\pi_{s,t})n}{2^i}\right)$$

giving quality $\tilde{O}\left(\frac{n}{2^{i}}\right)$.

Efficiently Maintaining G'. As shown above, maintaining an $\mathcal{ATO}(G, 2^i)$ requires detecting when any SCC in $G' = G \setminus E(S)$ has diameter above 2^i . We start by showing how to do this efficiently if we are given a black-box algorithm \mathcal{A}_{SSSP} that maintains distance estimates up to depth threshold 2^i (i.e. if a vertex is at distance less than 2^i from the source vertex, there is a distance estimate with good approximation ratio).

We use the random source scheme introduced in [RZ08] along with some techniques developed in [Che+16; BPW19; GW20a]: we choose for each SCC X in G' a center vertex CENTER(X) \in X uniformly at random, and use \mathcal{A}_{SSSP} to maintain distances from CENTER(X) to depth $2^i |X|/n \leq 2^i$. Since the largest distances between the vertex CENTER(X) and any other vertex in X is a 2-approximation on the diameter of G[X], this is sufficient to monitor the diameter and to trigger the separator procedure in good time.

Cast in terms of our new ATO-framework, the previous algorithm of [GW20a] used a regular Even and Shiloach tree for the algorithm \mathcal{A}_{SSSP} . We instead use a recursive structure,

where \mathcal{ATOs} of bad quality (large q) are used to build \mathcal{ATOs} of better quality (small q). Recall that our goal is to build an $\mathcal{ATO}(G, 2^i)$ of quality $\tilde{O}(n/2^i)$ and say that X is some SCC of $G' = G \setminus E(S)$ whose diameter we are monitoring. Now, using the lower level of the recursion, we inductively assume that we can maintain an $\mathcal{ATO}(G'[X], 2^{i-1})$ of quality $\tilde{O}(n/2^{i-1})$ in time $\tilde{O}(|X|^2)$. Plugging this \mathcal{ATO} into Theorem 3.3 gives us an algorithm for maintaining distances up to depth 2^i in X with total update time $\tilde{O}(|X|^2)$, Summing over all components X in G', we get an $\tilde{O}(n^2)$ total update time to maintain $\mathcal{ATO}(G, 2^i)$, as desired.

We actually cheated a bit in the last calculation, because the scheme above could incur an additional logarithmic factor for computing $\mathcal{ATO}(G, 2^i)$ from all the $\mathcal{ATO}(G'[X], 2^{i-1})$, so we can only afford a sublogarithmic number of levels, which leads to an extra $n^{o(1)}$ factor in the running time. However, a careful bootstrapping argument allows us to avoid this extra term.

3.5 A Framework for Sparse SSSP

Similarly to Theorem 3.2, we can prove a similar theorem with better running time for sparse graphs.

Theorem 3.7. [see Section 4 and Section 6.2.] For any $0 \le i \le \log(Wn)$, given a decremental digraph G = (V, E, w), we can maintain an $\mathcal{ATO}(G, 2^i)$ of expected quality $\tilde{O}(n/2^i)$. The algorithm runs in total expected update time $\tilde{O}(mn^{2/3})$ against a non-adaptive adversary with high probability.

In the remaining section, let us sketch how we obtain an efficient algorithm to compute SSSP given the result from Theorem 3.7. To simplify the presentation let us assume for the rest of the section that the graph G = (V, E) is unweighted.

Reducing SSSP to Hopset Maintenance. Let us introduce the notion of a $(1 + \epsilon, h)$ -hopset H which is a weighted decremental graph on the same vertex set as G such that at any stage, for any two vertices $s, t \in V$, we have

$$\mathbf{dist}_G(s,t) \le \mathbf{dist}_{G \cup H}^h(s,t) \le (1+\epsilon)\mathbf{dist}_G(s,t) \tag{4}$$

where we use the notation $\operatorname{dist}_{F}^{\ell}(s,t)$ to denote the shortest s to t path in the graph F consisting of at most ℓ edges. Using a well-known rounding technique, we can then run a slightly modified ES-tree data structure from a root vertex r on the graph $G \cup H$ to depth h to obtain $(1 + \epsilon)$ -approximate distances in G from r. The data structure only requires time $\tilde{O}((m + |E(H)|)h)$.

In our paper, we show how to construct for every $\lg(n^{2/3}\log n) \leq i \leq \lg n$, a $(1 + \epsilon, n^{2/3})$ hopset F_i with $\tilde{O}(n)$ edges that satisfies equation Equation (4) for all pairs at distance $[2^i, 2^{i+1})$, so the total update time of the ES-tree becomes $\tilde{O}(mn^{2/3})$, as desired. For $i < \lg(n^{2/3}\log n)$, we can maintain distances up to 2^i in total update time $\tilde{O}(mn^{2/3})$ by simply running a classic EStree, without any hop-set. We obtain final distance estimates by running the above algorithm for each *i*; then, for any pair *s*, *t*, we find the smallest distance estimate among these $O(\log(n))$ data structures and output that as the final distance estimate. **Maintaining a Hopset** F_i . For each $\lg(n^{2/3} \log n) \le i \le \lg n$, we want to maintain F_i as a weighted graph with the guarantee that for every vertices s and t at distance at $[2^i, 2^{i+1})$, we have a $(1 + \epsilon)$ -approximate shortest path in $F_i \cup G$ of hop at most $\tilde{O}(n^{2/3})$.

To maintain such a graph F_i , we use (\mathcal{V}_i, τ_i) an $\mathcal{ATO}(G, \eta_{diam} = \epsilon 2^i)$ as given in Theorem 3.7. We first sample each vertex $v \in V$ with probability $\tilde{\Theta}(n^{1/3}/2^i)$. We let S be the set of sampled vertices and have $|S| = \tilde{O}(n^{1+1/3}/2^i)$ with high probability. We then run from each vertex $s \in S$, with node $X^s \in \mathcal{V}_i$, an ES-tree to depth $2^i/n^{2/3}$ on the graph G/\mathcal{V}_i induced by the set of nodes $Y \in \mathcal{V}_i$ such that $|\tau_i(X^s) - \tau_i(Y)| = \tilde{O}(n^{2/3})$. We then add an edge (s, t)for every vertex t that is in a node in the ES-tree of the vertex s where we use corresponding distance estimate as an edge weight.

Hopset Sparsity and Running Time. It can be shown that every edge only participates in $\tilde{O}(|S|)$ ES-trees and therefore the total running time of all ES-trees can be bound by $\tilde{O}(m|S| \cdot 2^i) = \tilde{O}(mn^{2/3})$. Further, each ES-tree from a vertex $s \in S$ has with high probability at most $\tilde{O}(n^{1/3}/2^i)$ vertices in S in its tree and therefore the set F_i has at most $\tilde{O}(|S|n^{1/3}/2^i) = \tilde{O}(n)$ edges.

Correctness. Finally, to show that $G \cup F_i$ contains a path of hop at most $\tilde{O}(n^{2/3})$ between any two vertices s, t at distance 2^i let us focus on their corresponding shortest path $\pi_{s,t}$. We can partition the path into segments of length at most $2^{i-1}/n^{2/3}$, and by a standard hitting set argument, we have with high probability a vertex in S in every segment. Let s_1, s_2, \ldots, s_k be such that each s_j is a vertex in S in the j^{th} segment. We observe that $k \leq \frac{2^i}{2^{i-1}/n^{2/3}} = 2n^{2/3}$. Now, for every s_j we either have s_{j+1} in the ES-tree, in which case we have a direct edge between the vertices in F_i . Otherwise, some vertex v on the path segment from s_j to s_{j+1} is in a node Y such that $|\tau_i(X^s) - \tau_i(Y)| \gg n^{2/3}$. But since the quality of the \mathcal{ATO} is $\tilde{O}(n/2^i)$, the total sum of topological difference of the path $\pi_{s,t}$ is only $\tilde{O}(n)$. Thus, this can occur at most $n^{1/3}$ times. Every time, we can use the shortest path in G between s_j and s_{j+1} consisting of at most $2^i/n^{2/3} \leq n^{1/3}$ edges. Adding the paths from s to s_1 and from s_k to t, the total number of edges is at most $\tilde{O}(n^{2/3})$ as desired. Finally, we observe that when running the ES-data structure on $G \cup F_i$, we have to add an additive term of η_{diam} since the contractions in the graph G/\mathcal{V}_i might decrease distances by this additive term. But for distances in the range $[2^i, 2^{i+1})$, this term can be subsumed in the $(1 + \epsilon)$ multiplicative error (after rescaling ϵ by a constant factor).

3.6 Organization

We recommend the reader to carefully study Section 3 to gain necessary intuition for our approach. In Section 4, we give an efficient reduction from maintaining an approximate topological order to depth-restricted SSSP. This section is the centerpiece of the article and its main result, Theorem 4.4, is one of our main technical contributions.

We then show how to use Theorem 4.4 to obtain a SSSP data structure for dense graphs in Section 5 and for sparse graphs in Section 6. Both sections follow the same structure: we first show how to obtain a depth-restricted SSSP data structure using an \mathcal{ATO} and then we bootstrap the reductions to obtain the final result.

In Section 7, we draw a conclusion, put our results in perspective and discuss open problems.

4 Reducing Maintenance of an Approximate Topological Order to α -approximate δ -restricted SSSP

In this section, we show how to obtain an \mathcal{ATO} given an α -approximate δ -restricted \mathcal{SSSP} data structure. We start by defining such a data structure and then give a reduction.

Definition 4.1. Let \mathcal{A} be a data structure that given any decremental directed weighted graph G, a dedicated source $r \in V$, an approximation parameter $\alpha > 1$, maintains for each vertex $v \in V$, distance estimates $\widetilde{\operatorname{dist}}(r, v)$ and $\widetilde{\operatorname{dist}}(v, r)$ such that at any stage of G, for every pair $(s,t) \in (\{r\} \times V) \cup (V \times \{r\})$

- we have $\operatorname{dist}(s,t) \leq \widetilde{\operatorname{dist}}(s,t)$, and
- if $\operatorname{dist}(s,t) \leq \delta$, then $\operatorname{dist}(s,t) \leq \alpha \operatorname{dist}(s,t)$.

Then, we say \mathcal{A} is an α -approximate δ -restricted SSSP data structure with running time $T_{SSSP}(m, n, \delta, \alpha)$. We require only that \mathcal{A} runs against a non-adaptive adversary.

Remark 4.2. In the rest of the article, we implicitly assume that all SSSP data structures have $T_{SSSP}(m, n, \delta, \alpha)$ monotonically increasing in the first two parameters.

We also need another definition that makes it more convenient to work with \mathcal{ATO} s that only have *expected* quality (see Definition 3.1). However, we require high probability bounds in our constructions and it will further be easier to work with deterministic objects. This inspires the definition of an \mathcal{ATO} -bundle which is a collection of \mathcal{ATO} s such that for each path of interest, there is at least one \mathcal{ATO} in the bundle that has good quality for the path at-hand.

Definition 4.3 ($\mathcal{ATO}(G, \eta_{diam}, \ell)$ -bundle). Given a decremental weighted directed graph G = (V, E, w) and parameter $\eta_{diam} \geq 0$. We call $\mathcal{S} = \{(\mathcal{V}_i, \tau_i)\}_{i \in [1, \ell]}$ an $\mathcal{ATO}(G, \eta_{diam}, \ell)$ -bundle of quality q if every (\mathcal{V}_i, τ_i) is an $\mathcal{ATO}(G, \eta_{diam})$ and for any two vertices $s, t \in V$, there exists an $i \in [1, \ell]$, such that the shortest path $\pi_{s,t}$ in G satisfies $\mathcal{T}(\pi_{s,t}, \tau_i) \leq q \cdot w(\pi_{s,t}) + n$.

Without further due, let us state and prove the main result of this section.

Theorem 4.4 (\mathcal{ATO} -bundle from SSSP). Given an algorithm \mathcal{A} to solve 2-approximate δ -restricted SSSP on any graph H in time $T_{SSSP}(m, n, \delta)$, and for any c > 0, we can maintain an $\mathcal{ATO}(G, 2\alpha\delta, 40c \log n)$ -bundle of quality $\frac{(c+2)40000n \log^5 n}{\delta}$ in total expected update time

$$O\left(\sum_{j=0}^{\lceil \lg \delta \rceil} \sum_{k=0}^{2^{j+3}c \log^2 n} T_{SSSP}(m_{j,k}, n/2^j, \delta, 2) + m \log^3 n\right)$$
(5)

where $\sum_{j} m_{j,k} \leq 16c \cdot m \log^2 n$ for all k. The algorithm runs correctly with probability $1 - n^{-c}$ for any c > 0.

Remark 4.5. The graphs that the data structure \mathcal{A} runs upon during the algorithm are vertexinduced subgraphs of G. The data structure \mathcal{A} is further allowed to maintain distances on a larger subgraph of G, i.e. when \mathcal{A} is applied to a graph G[X], it can run instead on G[Y] for any set $X \subseteq Y \subseteq V$.

We point out that Remark 4.5 is only of importance at a later point at which the reader will be reminded and can safely be ignored for the rest of this sections (the reader is however invited to verify its correctness which is easy to establish).

We now describe how to obtain an efficient algorithm that obtains an $\mathcal{ATO}(G, 2\alpha\delta)$ henceforth denoted by (\mathcal{V}, τ) . The next sections describe how to initialize (\mathcal{V}, τ) , how to maintain useful data structures to maintain the diameter, give the main algorithm and then a rigorous analysis. Finally, we obtain a $\mathcal{ATO}(G, 2\alpha\delta, 40c \log n)$ -bundle by running $40c \log n$ independent copies of the algorithm below.

4.1 Initializing the Algorithm

As described in Section 3, our goal is to maintain a graph G' that is a subgraph of G and satisfies that no SCC X in G' has weak diameter $\operatorname{diam}(X, G)$ larger than $\frac{\delta|X|}{n}$. Throughout, we maintain the generalized topological order (\mathcal{V}, τ) on G' where τ has the nesting property as described in Theorem 2.1.

To ensure the diameter constraint initially, we use the following partitioning procedure whose proof is deferred to Appendix B.

Lemma 4.6 (Partitioning Procedure). Given an algorithm \mathcal{A} to solve 2-approximate δ -restricted SSSP. There exists a procedure PARTITION (G, d, ζ) that takes weighted digraph G, a depth threshold $d \leq \delta$ and a success parameter $\zeta > 0$, and returns a set $E_{Sep} \subseteq E$ such that

- 1. for each SCC X in $G \setminus E_{Sep}$, we have for any vertices $u, v \in X$ that $\operatorname{dist}_{G \setminus E_{Sep}}(u, v) \leq d$, and
- 2. for $e \in E$, we have $\mathbb{P}[e \in E_{Sep}] \leq \frac{240\zeta \log^2 n}{d} w(e)$.

The algorithm runs in total expected time

$$O\left(\sum_{j=0}^{\lceil \lg \delta \rceil} \sum_{k=0}^{2^{j+1}} T_{SSSP}(m_{j,k}, n/2^j, \delta, 2) + m \log^2 n\right)$$

where we have that $\sum_{k=0} m_{j,k} \leq 2m$ for every *i*. The algorithm terminates correctly with probability $1 - e^{-\zeta}$ for any c > 0.

Remark 4.7. During the execution, the graphs on which we use the SSSP structure upon have the properties as described in Remark 4.5.

Algorithm 2: INIT()

1 Let G' be initialized to G. 2 for i = 0 to $\lceil \lg \delta \rceil$ do 3 Compute the SCCs \mathcal{V} of G' 4 foreach SCC X in \mathcal{V} , $|X| \le n/2^i$ do 5 $E_{Sep} \leftarrow \text{PARTITION}(G[X], \delta/2^i, (c+2) \log n)$ 6 $G' \leftarrow G' \setminus E_{Sep}$ 7 return G'

Using this procedure, it is straight-forward to initialize our algorithm. The pseudo-code of the initialization procedure is given in Algorithm 2. Here, we iteratively apply the partitioning procedure to SCCs of small size to decompose them further if their diameter is too large. It is not hard to establish that the graph G' returned by the procedure, satisfies that every SCC X in G' has $\operatorname{diam}(X,G) \leq \frac{\delta|X|}{n}$.

4.2 Maintaining Information about SCC Diameters

Before we describe how to maintain G' to satisfy the guarantees given above, we address the issue of maintaining information about the diameter of the current SCCs in G'.

Therefore, we maintain a set S of random sources throughout the algorithm, and from each $s \in S$, we run an α -approximate δ -restricted SSSP data structure \mathcal{A}_s . Initially $S = \emptyset$, and whenever there is an SCC X in \mathcal{V} (which is maintained by the data structure on G'), and we find $S \cap X = \emptyset$, we pick a vertex s uniformly at random from X and add it to S. Once added, we initialize and maintain an α -approximate δ -restricted SSSP data structure \mathcal{A}_s on the current version of G[X]. That is, even if X does not form an SCC at later stages, the data structure is run until the rest of the algorithm on the graph G[X]. This ensures that once the algorithm is invoked, all edge updates are determined by the adversary formulating updates to G. Since we assume that the adversary is non-adaptive, we have that the SSSPdata structure only has do deal with updates from a non-adaptive adversary⁶.

We point out that since we maintain G' to be a decremental graph, we have that \mathcal{V} forms a refinement of previous versions at any stage i.e. the SCC sets only decompose over time in G'. Therefore, we can never have multiple center vertices in the same SCC $X \in \mathcal{V}$. For convenience, we let for each $X \in \mathcal{V}$, the vertex $\{s\} = X \cap S$ be denoted by CENTER(X). By the above argument, this function is well-defined.

⁶If we would instead remove vertices from the data structure, we would do so based on the information gathered from the data structure. Thus, the data structure would be required to work against an adaptive adversary. A similar problem arises when running on G'.

4.3 Maintaining G'

Let us now describe the main procedure of our algorithm: the part that efficiently handles violations of the diameter constraint by finding new separators. The implementation of this procedure is given by Algorithm 3. Let us now provide some intuition and detail as to how the algorithm works.

Algorithm 3: RESOLVEDIAMETERVIOLATIONS()

1 while there exists an $X \in \mathcal{V}$, where $\mathcal{A}_{\text{CENTER}(X)}$ has a distance estimate $\widetilde{\operatorname{dist}}(\operatorname{CENTER}(X), t)$ or $\widetilde{\operatorname{dist}}(t, \operatorname{CENTER}(X))$ exceeding $\frac{\delta|X|}{n}$ for some vertex $t \in X$ do /* Find separator sets that decompose X. */ $\mathbf{2}$ 3 else 4 $(E_{Sep}, C) \leftarrow \text{OutSeparator}(t, \overleftarrow{G'[X]}, \frac{|X|\delta}{2n}, (c+2)\log n)$ $\mathbf{5}$
$$\begin{split} E_{Sep}' &\leftarrow \text{PARTITION}(G'[C], \tfrac{|X|\delta}{4n}, (c+2)\log n)) \\ /* \text{ Update } G', \ \mathcal{V} \text{ and } \tau \text{ to reflect the changes.} \end{split}$$
6 */ $G' \leftarrow G' \setminus (E_{Sep} \cup E'_{Sep})$ 7 Wait Until the generalized topological order (\mathcal{V}, τ) of G' was updated, each SCC Z 8 in G' has a center CENTER(Z), and all data structures \mathcal{A}_s are updated.

The algorithm runs a while-loop starting in Line 1 that checks whether there exists a SCC $X \in \mathcal{V}$, such that the α -approximate δ -restricted SSSP data structure $\mathcal{A}_{\text{CENTER}(X)}$ has one of its distance estimates $\widetilde{\text{dist}}(\text{CENTER}(X), t)$ (or $\widetilde{\text{dist}}(t, \text{CENTER}(X))$) exceeding $\frac{\delta|X|}{n}$ for some vertex t in the same SCC X in G'. The goal of the while-loop iteration, is then to find a separator E_{Sep} between CENTER(X) and t and to delete the edges from G'.

Let us describe a loop-iteration where some distance estimate $\operatorname{dist}(\operatorname{CENTER}(X), t)$ was found that exceeded $\frac{\delta|X|}{n}$ and where $t \in X$ (the case where we have a distance estimate $\operatorname{dist}(\operatorname{CENTER}(X), t)$ exceed the threshold value is analogous and therefore omitted). In this case, we find a separator E_{Sep} that separates vertices in C (where $t \in C$) from vertices in $X \setminus C$ (where $\operatorname{CENTER}(X) \in X \setminus C$) in G'. Further, we invoke the procedure $\operatorname{PARTITION}(G'[C], \frac{\delta|X|}{4n}, \zeta)$ on C and obtain a separator E'_{Sep} in G' such that each SCC in $G'[C] \setminus E'_{Sep}$ has small diameter. We point out that while the first separator procedure is necessary to separate the vertices $\operatorname{CENTER}(X)$ and t in G', the partitioning procedure is run for technical reasons only since we cannot ensure an efficient implementation without this step.

Finally, we wait until the data structures that maintain the generalized topological order and the distance estimates from random sources are updated before we continue with the next iteration. On termination of the while-loop, we have that all distance estimates between centers and vertices in their SCC (with regard to G') are small (with regard to G).

4.4 Analysis

We establish Theorem 4.4 by establishing four lemmas establishing for (\mathcal{V}, τ) correctness (Lemma 4.8), running time (Lemma 4.12) and success probability (Lemma 4.13) and finally establishing that $c \log n$ independent copies of (\mathcal{V}, τ) form an $\mathcal{ATO}(G, 2\alpha\delta)$ -bundle with the guarantees given in Theorem 4.4, as required.

Lemma 4.8 (Correctness). Given that no procedure returns **Fail**, we have that the algorithm maintains (\mathcal{V}, τ) to be an $\mathcal{ATO}(G, 2\alpha\delta)$ of expected quality $\frac{(c+2)20000n \log^5 n}{\delta}$.

Let us first prove that the diameter of SCCs in G' remains small.

Claim 4.9. After invoking Algorithm 3, we have that each set $X \in \mathcal{V}$ satisfies

$$\mathbf{diam}(X,G) \le \frac{2\alpha\delta|X|}{n}$$

Proof. First, recall that when the while-loop in Line 1 terminates, we have that every $X \in \mathcal{V}$ has that no distance estimate $\widetilde{\mathbf{dist}}(\operatorname{CENTER}(X), t)$ or $\widetilde{\mathbf{dist}}(t, \operatorname{CENTER}(X))$ exceeds $\frac{\delta|X|}{n}$ for any $t \in X$.

Next, observe that the algorithm maintains the following invariant on the while-loop in Line 1: every $X \in \mathcal{V}$ contains exactly one center is only marked in the data structure $\mathcal{E}_{\text{CENTER}(X)}$. This follows by resampling centers in SCCs X that do not have a center yet and the by Line 8 which ensures that at the end of each while-loop iteration, there is time to resample.

Combined, this implies that on termination of the while-loop, for every $x, y \in X$, in any $X \in \mathcal{V}$, we have

$$\operatorname{dist}_{G}(x, y) \leq \operatorname{dist}_{G}(x, \operatorname{CENTER}(X)) + \operatorname{dist}_{G}(\operatorname{CENTER}(X), y)$$
$$\leq \widetilde{\operatorname{dist}}(x, \operatorname{CENTER}(X)) + \widetilde{\operatorname{dist}}(\operatorname{CENTER}(X), y)$$
$$\leq \frac{2\alpha\delta|X|}{n}$$
(6)

where we used the triangle inequality, Definition 4.1 and the fact that $\mathcal{A}_{\text{CENTER}}$ maintains distances with regard to a vertex-induced subgraph of G (adding edges can only decrease distances, thus distances in G are smaller than in $G[Y] \subseteq G$ for any Y).

Let us now bound the quality of the approximate topological order (\mathcal{V}, τ) , i.e. upper bound for any s-to-t path $\pi_{s,t}$ the amount $\mathcal{T}(\pi_{s,t}, \tau)$. As in the overview section, we focus on the "negative" terms in $\mathcal{T}(\pi_{s,t}, \tau)$, which are captured by

$$\mathcal{T}'(\pi_{s,t},\tau) \stackrel{\text{def}}{=} \sum_{(u,v)\in\pi_{s,t}} \min\{0,\tau(X^v) - \tau(X^u)\}$$
(7)

which is the definition of \mathcal{T}' already given in equation 3. It is not hard to see that $\mathcal{T}(\pi_{s,t},\tau) = 2\mathcal{T}'(\pi_{s,t},\tau) + |\tau(X^s) - \tau(X^t)| \leq 2\mathcal{T}'(\pi_{s,t},\tau) + n$. It, thus, only remains to establish the following lemma.

Claim 4.10. At any stage of G, for any path $\pi_{s,t}$ in G, we have

$$\mathbb{E}[\mathcal{T}'(\pi_{s,t},\tau)] \le \frac{(c+2)10000n \log^5 n}{\delta} w_G(\pi)$$

throughout the course of the algorithm.

Before, we provide a proof, let us state the following lemma which has been shown in slightly different forms in various papers before and whose proof is therefore delayed to Appendix C.

Lemma 4.11 (c.f. also [Che+16], Lemma 13; [BPW19], Lemma 7.1). Each vertex $v \in V$ participates in C in at most $2\lceil \lg \delta \rceil$ while-loop iterations of Algorithm 3 during the entire course of the algorithm in expectation. Further, in expectation, the SCC in G' that v is contained in, halves every second time that v participates in C.

Proof of Claim 4.10. We proof this lemma for edges $(u, v) \in E$. Then, the result follows straight-forwardly by summing over the path edges. Let us start by observing that we have $\mathcal{T}'((u, v), \tau) \neq 0$ if and only if X^v strictly precedes X^u in τ (where X^z denotes the set in \mathcal{V} that contains vertex $z \in V$). But since (\mathcal{V}, τ) forms a generalized topological order of G', we have that (u, v) cannot be contained in E(G').

However, we only remove edges from E(G') in Line 7 of our algorithm, after being added to E_{Sep} in Line 3 or 5, or to E'_{Sep} in Line 6. Having $(u, v) \in E_{Sep}$ occurs by Lemma 3.6 only if (u, v) is contained in G'[X] and if at least one of the endpoints is in C (depending on whether the separator is computed on G'[X] or G'[X] it is u or v). In this case, the probability that (u, v) is added to E_{Sep} is at most $\frac{(c+2)\log n2n}{|X|\delta}w_G(u, v)$, again by Lemma 3.6.

However, if (u, v) is not added to E_{Sep} (and not already removed from G') then it is completely contained in G'[C]. Thus, by Lemma 4.6 it is sampled into E'_{Sep} with probability at most $\frac{(c+2)240 \log^4(n) \cdot 4n}{|X|\delta} w_G(u, v) \leq \frac{(c+2)960n \log^4 n}{|X|\delta} w_G(u, v)$. Observe that if (u, v) is sampled into either E_{Sep} or E'_{Sep} , then since it was contained in

Observe that if (u, v) is sampled into either E_{Sep} or E'_{Sep} , then since it was contained in X and by the nesting property of τ which is guaranteed by Theorem 2.1, we have that during the rest of the algorithm, we have $|\tau(X^u) - (X^v)| < |X|$ where X^u (resp. X^v) denotes the set in \mathcal{V} that contains u (resp. v).

Thus, a while-loop iteration where u or v participate in C adds to $\mathbb{E}[\mathcal{T}'(\pi_{s,t},\tau)]$ at most

$$|X| \cdot \frac{(c+2)1000n \log^4 n}{|X|\delta} w_G(u,v) = \frac{(c+2)1000n \log^4 n}{\delta} w_G(u,v)$$

Since by Lemma 4.11 each vertex only occurs during $2 \lg n$ while-loop iterations in C, we can establish the final bound.

Combining the fact that (\mathcal{V}, τ) is a GENERALIZEDTOPOLOGICALORDER(G') at all stages and $G' \subseteq G$ where τ has the nesting property, combined with Claim 4.9 and Claim 4.10, we derive Lemma 4.8. **Lemma 4.12** (Running Time). The algorithm to maintain (\mathcal{V}, τ) requires at most expected time

$$O\left(\sum_{j=0}^{\lceil \lg \delta \rceil} \sum_{k=0}^{2^{j+3} \lceil \lg \delta \rceil} T_{SSSP}(m_{j,k}, n/2^j, \delta, 2) + m \log^2 n\right)$$

where $\sum_{j} m_{j,k} \leq 16m \lceil \lg \delta \rceil$.

Proof. Again, our proof crucially relies on the following lemma.

Lemma 4.11 (c.f. also [Che+16], Lemma 13; [BPW19], Lemma 7.1). Each vertex $v \in V$ participates in C in at most $2\lceil \lg \delta \rceil$ while-loop iterations of Algorithm 3 during the entire course of the algorithm in expectation. Further, in expectation, the SCC in G' that v is contained in, halves every second time that v participates in C.

We first observe that the initialization procedure described in Section 4.1 initializes G' in O(m) time and then runs $O(\log n)$ iterations where in each iteration it invokes the procedure PARTITION(·) on a set of disjoint subgraphs of G to update G'. By Lemma 4.6, we can implement all of these calls in time

$$O\left(\log n\left(\sum_{j=0}^{\lceil \lg \delta \rceil} \sum_{k=0}^{2^{j+1}} T_{SSSP}(m_{j,k}, n/2^j, \delta, 2) + m \log n\right)\right).$$

The latter term in this expression subsumes the time spend on updating G' once a separator is returned.

Next, let us bound the time spend on maintaining the SSSP data structures as described in Section 4.2. It is here that we use Lemma 4.11: we have that initially each vertex (and edges) is in exactly one data structure. Further, every second time a vertex v participates in C as computed in Line 3 or Line 5, the SCC it is contained in in G' is halved in size (i.e. in the number of vertices). Since new SSSP data structures are initialized on the new SCCs that are contained in the C set, we have that each vertex v, in expectation, only participates 2 times in an SSSP structure with running time $T_{SSSP}(m_{j,k}, n/2^j, \delta, 2) + m \log n$ for any j. Since each edge is incident to only two vertices, we have a similar argument on edges and can therefore bound the total amount of time spend on SSSP data structures by

$$O\left(\sum_{j=0}^{\lceil \lg \delta \rceil} \sum_{k=0}^{2^{j+1}} T_{SSSP}(m_{j,k}, n/2^j, \delta, 2)\right).$$

where $\sum_{j} m_{j,k} \leq 4m$.

Finally, let us bound the time spend in calls to Algorithm 3. We observe that each whileloop iteration takes time

$$O\left(\sum_{j=0}^{\lceil \lg \delta \rceil} \sum_{k=0}^{2^{j+1}} T_{SSSP}(m'_{j,k}, n'/2^j, \delta, 2) + m' \log n\right)$$

where $\sum_{i} m'_{i,k} \leq 4m'$ for $m' = |E_G(C)|$ and $n' = |V_G(C)|$. This follows since the OUTSEPARATOR(·) procedure runs in time almost-linear in the number of edges incident to C and afterwards the call of the procedure $PARTITION(\cdot)$ which dominates the costs of the procedure is only on the graph G' induced by the vertices in C. Thus, this insight follows straight-forwardly from Lemma 3.6 and Lemma 4.6 and the insight that the cost of the remaining operations is subsumed in the bounds.

Finally, we again use Lemma 4.11 which gives that summing over all while-loop iterations is at cost at most

$$O\left(\sum_{j=0}^{\lceil \lg \delta \rceil} \sum_{k=0}^{2^{j+2} \lceil \lg \delta \rceil} T_{SSSP}(m_{j,k}, n/2^j, \delta, 2) + m \log^2 n\right)$$

where $\sum_{j} m_{j,k} \leq 8m \lceil \lg \delta \rceil$. Combining the parts of the algorithm, we thus get the total bound.

Lemma 4.13 (Success Probability). The algorithm reports **Fail** with probability at most $2n^{-c-1}$.

Proof. We point out that we can only get a **Fail** due to procedures $OutSeparator(\cdot)$ and PARTITION(\cdot).

Since each separator found in the while-loop in Line 1 refines \mathcal{V} , we can bound the number of while-loop iterations in the course of the algorithm by n-1. Thus, we make at most n-1calls to procedures $OUTSEPARATOR(\cdot)$ and $PARTITION(\cdot)$. Each of the former calls returns Fail with probability at most $n^{-(c+2)}$ and each of the latter with probability at most $n^{-(c+2)}$.

Taking a union bound over all events, the lemma follows.

Finally, let us put everything together and prove our main theorem.

Theorem 4.4 (\mathcal{ATO} -bundle from \mathcal{SSSP}). Given an algorithm \mathcal{A} to solve 2-approximate δ restricted SSSP on any graph H in time $T_{SSSP}(m, n, \delta)$, and for any c > 0, we can maintain an $\mathcal{ATO}(G, 2\alpha\delta, 40c \log n)$ -bundle of quality $\frac{(c+2)40000n \log^5 n}{\delta}$ in total expected update time

$$O\left(\sum_{j=0}^{\lceil \lg \delta \rceil} \sum_{k=0}^{2^{j+3}c \log^2 n} T_{SSSP}(m_{j,k}, n/2^j, \delta, 2) + m \log^3 n\right)$$
(5)

where $\sum_{j} m_{j,k} \leq 16c \cdot m \log^2 n$ for all k. The algorithm runs correctly with probability $1 - n^{-c}$ for any c > 0.

Proof. We maintain a collection of $40c \log n$ independent $\mathcal{ATO}(G, 2\alpha\delta)$ instances

 $(\mathcal{V}_1, \tau_1), (\mathcal{V}_2, \tau_2), \dots, (\mathcal{V}_{40c \log n}, \tau_{40c \log n})$

as described earlier in this section and let \mathcal{S} denote the collection of these instances.

The total running time to maintain these $\mathcal{ATO}(G, 2\alpha\delta)$'s is clearly bounded by the term given in equation 5 by Lemma 4.12.

Now, since by Lemma 4.8, each $\mathcal{ATO}(G, 2\alpha\delta)$ has expected quality $q = \frac{(c+2)20000n \log^5 n}{\delta}$, we have by Markov's inequality and a simple Chernoff bound, that for each shortest path $\pi_{s,t}$ in *G* at some stage *t*, we have that there exists an *i*, such that $\mathcal{T}(\pi_{s,t},\tau_i) \leq 2q$ with probability at least $1 - e^{-40c \log n/8} = 1 - n^{-5c}$. Since c > 1, we have that the probability that any shortestpath at any stage fails, is at most $1 - n^{-c}/2$ by union bounding over at most n^2 stages and at most n^2 shortest-paths, for *n* large enough. Moreover, the total probability that any instance returns **Fail** is at most $n^{-c}/2$ by Lemma 4.13 and a union bound over the instances. Thus, we have established that with probability at least $1 - n^{-c}$, S forms an $\mathcal{ATO}(G, 2\alpha\delta)$ -bundle of quality 2q as defined in Definition 4.3.

5 A SSSP Algorithm for Dense Graphs

We now give a proof of Theorem 3.2 which implies our main result, Theorem 1.1, as a corollary. Our proof is in two steps: we first show how to implement an α -approximate δ -restricted SSSP as described in Definition 4.1 given access to approximate topological orders. We then show how to bootstrap the reductions to maintains different SSSP data structures to cover all depths.

5.1 α -approximate δ -restricted SSSP via Maintaining an Approximate Topological Order

The main objective of this section is to prove the following theorem which gives a reduction from $(1+\epsilon)$ -approximate δ -restricted SSSP to approximate topological orders. In this theorem, we only assume access to an $\mathcal{ATO}(G, \eta_{diam})$ denoted by (\mathcal{V}, τ) where we assess the quality individually for each path. If the quality for a certain tuple is below a threshold q, we show how to exploit the approximate topological order to maintain the distance estimate for the tuple efficiently, otherwise we provide no guarantees.

Theorem 5.1. Given G = (V, E, w), a decremental weighted digraph, a source $r \in V$, a depth threshold $\delta > 0$, a quality parameter q, an approximation parameter $\epsilon > 0$, and access to (\mathcal{V}, τ) an $\mathcal{ATO}(G, \eta_{diam})$.

Then, there exists a deterministic data structure that maintains a distance estimate $\operatorname{dist}(r, v)$ for every vertex $v \in V$ such that at each stage of G, $\operatorname{dist}_G(r, v) \leq \widetilde{\operatorname{dist}}(r, v)$ and if $\operatorname{dist}_G(r, v) \leq \delta$ and $\mathcal{T}(\pi_{r,v}, \tau) \leq q \cdot \delta + n$, then

$$\operatorname{dist}(r, v) \leq \operatorname{dist}_G(r, v) + \eta_{diam} + \epsilon \delta.$$

The total time required by this structure is

$$O(n\delta q \log n/\epsilon + n^2 \log n)$$

Remark 5.2. Technically, we require the approximate topological order (\mathcal{V}, τ) to encode changes efficiently and pass them the SSSP data structure. Since the SSSP data structure is updated

only through delete operations, we require, that with each edge update, the data structure receives changes to (\mathcal{V}, τ) since the last stage. More precisely, we require that the user passes a set of pointers to each set Y that occurred in \mathcal{V} at the previous stage (denoted \mathcal{V}^{OLD}), but did not occur in \mathcal{V} at the current stage (denoted \mathcal{V}^{NEW}), i.e. each $Y \in \mathcal{V}^{OLD} \setminus \mathcal{V}^{NEW}$. Additionally, we require with each such Y that was split into subsets $Y_1, Y_2, \ldots, Y_k \in \mathcal{V}^{NEW}$ that form a partition of Y, pointers to each new element Y_i . We further require worst-case constant query time of τ , and each element $Y \in \mathcal{V}$ (for any version) can be queried for its size in constant time and returns its vertex set in time O(|Y|). For the rest of the paper, this detail will be concealed in order to improve readability.

Before we show how to implement such a data structure, let us emphasize that the above theorem directly implies Theorem 3.3 that we introduced in the overview. It can further also be used to derive the following corollary which is at the heart of our proof in the next section. Its proof is rather straight-forward and can be found in Appendix E.

Corollary 5.3. Given G = (V, E, w), a decremental weighted digraph, a source $r \in V$, a depth threshold $\delta > 0$, an approximation parameter $\epsilon > 0$, and access to a collection $S = \{S_i\}_{1 \leq i \leq \mu}$ for $\mu = \lfloor \lg \delta \rfloor - 1$ where each S_i forms an $\mathcal{ATO}(G, 2^i, 40 \operatorname{clog} n)$ -bundle of quality q_i . Then, there exists an implementation for $(1 + \epsilon)$ -approximate δ -restricted SSSP where $T_{SSSP}(n, m, \delta, \epsilon) = O(n(\max_{1 \leq i \leq \mu} \{\frac{\delta q_i}{2^i}\} + n) \log^3 n/\epsilon^2).$

Let us now describe the implementation of a data structure \mathcal{E}_r that stipulates the guarantees given in Theorem 5.1. Since the proof that this is indeed a valid implementation of Theorem 5.1 is quite similar to the proof sketch we give in Section 3, we refer the reader to Appendix D.

Initialization. Throughout the algorithm, we define $\delta_{max} = \lceil (1 + \epsilon)\delta + \epsilon n/q \rceil$ and define the complete graph $\mathcal{H} = (\mathcal{V}, \mathcal{V}^2, w)^7$ with weight function

$$w(X,Y) = \inf\{w(x,y) | (x,y) \in E(X,Y)\}$$

for $X, Y \in \mathcal{V}$. We use the convention that the infimum of the empty set is ∞ . We use \mathcal{H} to avoid dealing explicitly with G/\mathcal{V} which is a multi-graph, instead in \mathcal{H} we use the same node set with simple edges as the infimum over weights of the multi-edges (even if there is no such edge).

We use a standard min-heap data structure⁸ $Q_{X,Y}$ over the set E(X,Y) for each ordered pair (X,Y) to maintain the weight w(X,Y). We henceforth denote by $Q_{X,Y}$.MINVALUE the value w(X,Y) and by $Q_{X,Y}$.MINELEM a corresponding edge (x,y) with $x \in X, y \in Y$, and use the convention of denoting the node in \mathcal{V} that contains vertex $x \in V$ by X^x . We initialize the data structure \mathcal{E}_r by constructing \mathcal{H} and by running Dijkstra's algorithm⁹ from X^r on \mathcal{H} . We then initialize a distance estimate $\widetilde{\mathbf{dist}}(X^r,Y)$ for each $Y \in \mathcal{V}$ to $\mathbf{dist}_{\mathcal{H}}(X^r,Y)$. If we have

⁷Here, we are again slightly abusing notation by referring to \mathcal{V} as partition and node set, however, context and the fact that this implicitly refers to a one-to-one correspondence between partition sets and nodes ensures that no ambiguity arises.

⁸See for example [Cor+09].

 $^{^{9}}$ See [Cor+09] for an efficient implementation.

 $\operatorname{dist}(X^r, Y) > \delta_{max}$ at any point in the algorithm, we set it to ∞ . Further, we also maintain the distance estimates $\operatorname{dist}(r, u)$ for each $u \in V$ equal to $\operatorname{dist}(X^r, X^u) + \eta_{diam}$, i.e. every time we increase $\operatorname{dist}(X^r, X^u)$, we also increase u's distance estimate¹⁰. This allows us to henceforth focus on the distance estimates of nodes which is easier to describe. We also store the corresponding shortest-path tree T truncated at distance δ_{max} that serves as a certificate of the distance estimates.

Finally, we partition for each node $X \in \mathcal{V}$, the in-neighbors set in \mathcal{H} of X into different buckets based on their τ -distance: for each X, we initialize bucket $B_{-1}(X) = \{X\}$ and for $0 \leq j \leq \lg n$ we initialize the bucket $B_j(X)$ to

$$\{2^{j} \le \chi(X, Y, \tau) < 2^{j+1} | Y \in N_{\mathcal{H}}^{in}(X), X \neq Y\}$$

where we define

$$\chi(X,Y,\tau) \stackrel{\text{def}}{=} \begin{cases} \tau(Y) - (\tau(X) + |X| - 1) & \text{if } \tau(X) < \tau(Y) \\ \chi(Y,X,\tau) & \text{otherwise} \end{cases}$$
(8)

that is $\chi(\cdot)$ is similar to $\mathcal{T}(\cdot)$ (in fact $\chi(X, Y, \tau) \leq \mathcal{T}(X, Y, \tau)$), however, as τ maps nodes X and Y to disjoint intervals, $\mathcal{T}(\cdot)$ measures the distance between the starting points of the intervals, while $\chi(\cdot)$ measures the distance between the intervals (i.e. the closest endpoints of the intervals).

At any stage, we let $B_{\leq j}(X) = \bigcup_{j' \leq j} B_{j'}(X)$. We store each set $B_j(X)$ explicitly as a linked list, store for each $Y \in B_j(X)$ a pointer to the bucket, and maintain the buckets to partition the in-neighbors of each X.

Handling Edge Deletions. The edge deletion procedure takes two parameters: the edge to be deleted (u, v) and a collection of tuples U that encode refinements of \mathcal{V} during this stage. To handle the update, we initialize a min-heap $Q = \emptyset$ that keeps track of the nodes in \mathcal{H} , that cannot be reached from X^r in the truncated shortest-path tree T (i.e. whose certificate for the current distance estimate was compromised).

We start our update procedure by processing updates to (\mathcal{V}, τ) (check the remark of Theorem 5.1 for a description of these updates encoded by U). For any node $X \in \mathcal{V}^{OLD} \setminus \mathcal{V}^{NEW}$, that was split into subsets $X_1, X_2, \ldots, X_k \in V^{NEW}$ (i.e. for every tuple $(X, X_1, X_2, \ldots, X_k) \in U$), we query for each X_i , its size. Then, we let the largest node X_i inherit the original node X(that is the nodes are equal in our data structure at this stage although the partition sets are not), and create a new node $X_{i'}$ in \mathcal{H} for other $i' \neq i$, and new heap structures $Q_{X_{i'},Y}$ for every $Y \in \mathcal{V}$. Then for each $X_{i'}, i' \neq i$, we scan each edge (x, y) in $E(X_{i'}, V \setminus X_{i'})$, remove it from the heap Q_{X_{i},X^y} and add it to the new heap $Q_{X_{i'},X^y}$. We also initialize the distance estimates for each $X_{i'}, \operatorname{dist}(X^r, X_{i'})$ to take the value $\operatorname{dist}(X^r, X)$, also for X_i . We then find the edge (w, x) in T where $X = X^x$. Clearly, we now have that $X_{i'} = X^x$ for some i' and we connect $X_{i'}$ in the tree T since this edge is now a certificate for $X_{i'}$'s distance estimate. The

¹⁰We will show that $\widetilde{\mathbf{dist}}(X^r, X^u)$ is a monotonically increasing value over time.

rest of the nodes, i.e. the nodes $X_1, X_2, \ldots, X_{i'-1}, X_{i'+1}, \ldots, X_k$, we add to Q since we do not have a certificate for them yet.

Finally, when all the node splits for the current stage where processed, we update the buckets $B_j(Y)$ for all j and Y to almost stipulate the initialization rules. We point out that all edges that have to be assigned to a different bucket have to be incident to X_1, X_2, \ldots, X_k by property 2. Then, for each $X_{i'}$ (also X_i), we compute j to be the largest integer such that some number in $|X_{i'}|, |X_{i'}| + 1, \ldots, |X| - 1$ is divisible by 2^j . Then, we update all nodes in $B_{\leq j+1}(X_{i'})$ by scanning and reassigning them, and similarly reassign $X_{i'}$ to a new bucket for each Y where $X_{i'} \in B_{\leq j}(Y)$.

Finally, when all node splits are processed, the node set of \mathcal{H} reflects the current \mathcal{V} , and we can delete the edge (u, v) from \mathcal{H} by deleting it from the heap it is contained in. If (u, v)was equal to Q_{X^u, X^v} . MINELEM, and $(X^u, X^v) \in T$, we delete it from T and insert X^v into Q.

We then rebuild our certificate T: we take the node Y from Q with the smallest distance estimate $\widetilde{\mathbf{dist}}(X^r, Y)$ until Q is empty or the smallest distance estimate ∞ . Now, let j be the largest integer such that $\widetilde{\mathbf{dist}}(X^r, Y)$ is divisible by $\lceil 2^j \cdot \frac{\epsilon}{q} \rceil$. We then check if there exists a node $X \in B_{\leq j}(Y)$ such that

$$\operatorname{dist}(X^r, X) + Q_{X,Y}$$
. MINVALUE $\leq \operatorname{dist}(X^r, Y)$.

In this case, the edge $(x, y) = Q_{X,Y}$. MINELEM serves as a certificate that the distance from X^r to Y is at most $\widetilde{\operatorname{dist}}(X^r, Y)$ and therefore we add (x, y) to T. If there exists no such vertex X, then we increase the value $\widetilde{\operatorname{dist}}(X^r, Y)$ by one or to ∞ if it is currently at least δ_{max} and reinsert Y and children Z_1, Z_2, \ldots, Z_k of Y in T into Q (after deleting the edges (Y, Z_i) from T). This completes the description of the algorithm. Again, we refer the reader interested in the proof of Theorem 5.1 to Appendix D.

5.2 Bootstrapping an Algorithm for Unrestricted Depth

Next, let us prove the following theorem which show is a detailed version of Theorem 3.2. Combined with Corollary 5.3 (where we set the depth threshold parameter δ to Wn), this immediately implies our main result, Theorem 1.1.

Theorem 5.4. For any $0 \le i \le \lg(Wn)$, given a decremental digraph G = (V, E, w), we can maintain a hierarchy $S = \{S_i\}_i$ where each S_i is a $\mathcal{ATO}(G, 2^i, 40c \log n)$ -bundle of expected quality $\tilde{O}(n/2^i)$. The algorithm runs in total expected update time $O(c^5n^2 \log^{17} n \lg^5(Wn))$ against a non-adaptive adversary and is correct with probability at least $1 - n^{-c+2}$ for any failure probability parameter $c \ge 2$.

Proof. In order to prove our theorem formally, we need to fix the constants hidden by the big-O notation in some of our statements. We therefore henceforth denote the constant hidden by Corollary 5.3 to maintain the SSSP data structure by c_{SSSP} , the constant hidden in Theorem 4.4 to obtain an \mathcal{ATO} -bundle from an SSSP data structure by $c_{SSSP\to ATO}$ and finally, we denote the constant hidden in the theorem that we want to prove by c_{Total} where we require that $c_{Total} \geq (c_{SSSP\to ATO} \cdot c_{SSSP})^2 \cdot 2^{49}$.

Without further due, let us prove the theorem by induction on n, the number of vertices in graph G. The base case with $n \leq 1$ is easily established since there are no paths in a graph of only one vertex thus we obtain arbitrarily good quality and the running time is a small constant (at least smaller than c_{Total}).

Let us now give the inductive step $n \mapsto n+1$: for each i, we iteratively construct an $\mathcal{ATO}(G, 2^i, 40c \log n)$ -bundle \mathcal{S}_i as described in Theorem 4.4. Thus, we have to show how to implement a 2-approximate 2^{i-2} -restricted SSSP data structure required in the reduction (note that for i < 2 this task is trivial, so we omit handling it as special levels).

Note that each data structure SSSP that we are asked to implement for an ATO-bundle \mathcal{S}_i at level i is run on a different graph $H \subseteq G$. To obtain an efficient algorithm, we will implement the data structure differently depending on the size of such H. If H has at least $n/2^{\gamma}$ vertices for some $\gamma = \Theta(\lg \log Wn)$ that we fix later, we call H a large graph. Otherwise, we say H is *small*. Now, we implement SSSP as follows:

- if H is small, then we use the induction hypothesis, find a \mathcal{ATO} -bundle S' for H and and invoke Corollary 5.3 on \mathcal{S}' . We note that we need to set the parameter that controls the failure probability for \mathcal{S}' to $c \cdot 4 \log n$ to ensure that it succeeds with high probability (this is since \mathcal{S}' only succeeds with probability polynomial in |V(H)| which might be very small).
- if H is large, we exploit Remark 4.5 which states that when the reduction asks to maintain approximate distances on some graph $H \subseteq G$, it is sufficient to maintain distance estimates on any graph F such that $H \subseteq F \subseteq G$ and in particular, it is ok to simply run on the entire graph G. Therefore, we simply use S_{i-1} in combination with Corollary 5.3 and maintain distances in G.

Let us now analyze the total running time. We start by calculating the running time required by each level i separately. For some fixed i, we have that by Theorem 4.4 we have running time

$$c_{SSSP \to ATO} \left(\sum_{j=0}^{\lceil \lg 2^{i-2} \rceil} \sum_{k=0}^{2^{j+3} c \log^2 n} T_{SSSP}(m_{j,k}, n/2^j, 2^{i-2}, 2) + n^2 \log^3 n \right)$$
(9)

where $\sum_{j} m_{j,k} \leq 16c \cdot m \log^2 n$ for all k, to maintain S_i . Let us analyze the terms $T_{SSSP}(m_{j,k}, n/2^j, 2^{i-2}, 2)$ based on whether $j < \gamma$ or not (i.e. depending on how the SSSP data structure was implemented):

• if $j < \gamma$: then, by the induction hypothesis, we require time at most

$$c_{Total} (c4 \log n)^5 \cdot \left(\frac{n}{2^j}\right)^2 \log^{17}(n) \lg^5(Wn) 2^{-2j}$$

$$< c_{Total} \cdot 2^8 \cdot c^5 n^2 \log^{22}(n) \lg^5(Wn) 2^{-j} 2^{-\gamma}$$

to maintain the new \mathcal{ATO} -bundle \mathcal{S}' on the graph H and again by the induction hypothesis we have that the bundle has quality $\frac{(4 \log nc+2)40000n \log^5 n2^{-j}}{2^{i-3}} \log^3 n$.

Thus, maintaining SSSP on H using S' as described in Corollary 5.3 can be done in time

$$c_{SSSP}((c \cdot 4\log n + 2)40000 \cdot 2^3)(n^2 \log^8 n 2^{-2j}) < c_{SSSP} \cdot c \cdot 2^{22}(n^2 \log^9 n 2^{-j} 2^{-\gamma}).$$

Combined, we obtain that we can implement the entire SSSP data structure with total running time at most

$$c_{Total} \cdot 2^8 \cdot c^5 n^2 \log^{22}(n) \lg^5(Wn) 2^{-j} 2^{-\gamma} + c_{SSSP} \cdot c \cdot 2^{22} (n^2 \log^9 n 2^{-j} 2^{-\gamma})$$

$$\leq c_{Total} \cdot c_{SSSP} \cdot c^5 \cdot 2^{-j} 2^{-\gamma} \cdot 2^{22} \cdot n^2 \log^{22}(n) \lg^5(Wn).$$

Combining these bounds and summing over all small graph terms in equation 9, we obtain that the total contribution is at most

$$c_{SSSP\to ATO} \cdot \lg(nW) \cdot (2^{3}c\log^{2}n) \left(c_{Total} \cdot c_{SSSP} \cdot c^{5} \cdot 2^{-\gamma} \cdot 2^{22} \cdot n^{2}\log^{22}(n) \lg^{5}(Wn) \right)$$

$$\leq c^{6} \cdot c_{SSSP\to ATO} \cdot c_{Total} \cdot c_{SSSP} \left(2^{22} \cdot n^{2}\log^{24}(n) \lg^{6}(Wn) 2^{-\gamma} \right).$$

This completes the analysis of the small graph data structures.

• otherwise $(j \ge \gamma)$: then, we run the SSSP structure from Corollary 5.3 on S_{i-1} which gives running time at most

$$c_{SSSP}\left(n2^{i} \cdot \frac{(c+2)40000n\log^{5} n}{2^{i-3}}\log^{3} n\right) \le c_{SSSP} \cdot c(2^{20} \cdot n^{2}\log^{8} n)$$

where we used $c \ge 2$. Since there are at most $c_{SSSP\to ATO} \cdot c \cdot (2^3 \lg(nW) \log^2 n 2^{\gamma})$ terms for large graphs, where $j \ge \gamma$, we have that the total cost of all SSSP data structures on *large* graphs is at most

$$c_{SSSP \to ATO} \cdot c \cdot (2^{3} \lg(nW) \log^{2}(n) 2^{\gamma}) \cdot \left(c_{SSSP} \cdot c(2^{20} \cdot n^{2} \log^{8} n)\right)$$
$$= c_{SSSP \to ATO} \cdot c^{2} \cdot c_{SSSP} \cdot \left(2^{23} \cdot n^{2} \log^{10}(n) \lg(nW) 2^{\gamma}\right)$$

It now only remains to choose γ and combine the two bounds. We set $\gamma = 24\lceil \lg(c^2 \cdot c_{SSSP \to ATO} \cdot c_{SSSP} \cdot \lg^3(Wn) \log^7(n)) \rceil$, and obtain that the total running time summed over large and small graphs is at most

$$\begin{aligned} c_{SSSP \to ATO}^2 & \cdot c_{SSSP}^2 \cdot c^4 \cdot 2^{48} (n^2 \log^{17} n \lg^4(Wn))) \\ &+ \frac{c^4 c_{Total} (n^2 \log^{17} n \lg^3(Wn))}{2} \\ &\leq c^4 c_{Total} (n^2 \log^{17} n \lg^4(Wn)) \end{aligned}$$

where we used our initial assumption on the size of c_{Total} .

Finally, we point out that there are at most $\lg(nW)$ levels *i* and therefore, the total update time is at most

$$c_{Total}(c^4n^2\log^{17}n\lg^5(Wn))$$

as required.

Further, we point out that every \mathcal{ATO} -bundle \mathcal{S}_i that was constructed runs correctly with high probability at least $1 - n^{-c}$, while every \mathcal{ATO} -bundle \mathcal{S}' is maintained correctly with probability at least $1 - n^{-4c}$ (recall that we set the failure parameter of these data structures to $c \cdot 4 \log n$). Noting that we only have $\lg(nW)$ instances of the former bundles, and at most n^3 of the latter, taking a simple union bound over the events that any bundle instance fails gives a total failure probability of at most n^{-c+2} .

6 A SSSP Algorithm for Sparse Graphs

Finally, we give a construction that is efficient in sparse graphs. Again, we prove Theorem 1.2 in two steps: we first show how to reduce SSSP to maintaining an approximate topological order and then we bootstrap the approach.

6.1 α -approximate δ -restricted SSSP via Maintaining an Approximate Topological Order

In this section, we give the following theorem that takes an $\mathcal{ATO}(G, \eta_{diam})$ and shows how to implement $(1 + \epsilon)$ -approximate δ -restricted \mathcal{SSSP} . As in the previous section, we first prove a simpler theorem and then derive the \mathcal{SSSP} data structure as a corollary.

Theorem 6.1. Given G = (V, E, w), a decremental weighted digraph, a source $r \in V$, a depth threshold $\delta > 0$, a quality parameter q such that $\delta q \ge n$, an approximation parameter $\epsilon > 0$, and access to (\mathcal{V}, τ) an $\mathcal{ATO}(G, \eta_{diam})$.

Then, there exists a data structure that maintains a distance estimate $\operatorname{dist}(r, v)$ for every vertex $v \in V$ such that at each stage of G, $\operatorname{dist}_G(r, v) \leq \widetilde{\operatorname{dist}}(r, v)$ and if $\operatorname{dist}_G(r, v) \leq \delta$ and $\mathcal{T}(\pi_{r,v}, \tau) \leq q \cdot \delta + n$, then

$$\operatorname{dist}(r, v) \le (1 + \epsilon) \operatorname{dist}_G(r, v) + \eta_{diam}.$$

The expected total time required by this structure is

$$O((m\delta q/n^{-1/3} + \delta q n^{2/3}) \log^3 n \log W n/\epsilon)$$

and the data structure runs correctly with probability at least $1 - n^{-c}$ for any constant c > 0.

We point out that Remark 5.2 also applies to this theorem. As before, we use the theorem above to establish a simple corollary where we refer the reader to the proof of Corollary 5.3 which with minor changes proves the below result.

Corollary 6.2. Given G = (V, E, w), a decremental weighted digraph, a source $r \in V$, a depth threshold $\delta > 0$, an approximation parameter $\epsilon > 0$, and access to a collection $S = \{S_i\}_{1 \leq i \leq \mu}$ for $\mu = \lfloor \lg \delta \rfloor - 1$ where each S_i forms an $\mathcal{ATO}(G, 2^i, 40 \operatorname{clog} n)$ -bundle of quality $q \cdot 2^i$. Then, there exists an implementation for $(1 + \epsilon)$ -approximate δ -restricted SSSP where $T_{SSSP}(n, m, \delta, \epsilon) = O((m \delta q / n^{-1/3} + \delta q n^{2/3}) \log^3 n \log^2 Wn / \epsilon)$ with high probability.

As in the previous section, we defer the proof of Theorem 6.1 to Appendix F but present the implementation of a data structure \mathcal{E}_r that satisfies the guarantees of Theorem 6.1. We therefore first show how to maintain a hopset H of G using (\mathcal{V}, τ) and then show that we can run an ES-tree on the hopset H from source s. Before we discuss the internal parts of the data structure, let us define some useful concepts.

Preliminary Concepts. In our algorithm, it is key to distinguish between the hop h and the depth δ that was given. We therefore introduce the superscript h to distances to denote the weight of the h-hop restricted shortest path in G between two vertices $s, t \in V$ by $\operatorname{dist}_{G}^{h}(s, t)$.

Further, we henceforth use $h_i = 2^i$ and we give data structures for every $i \leq \lg n$ that recovers paths from r to vertices at distance at most δ and hop h_i and assume for convenience that n is a power of 2 wlog. In order to describe the data structure for every hop level i, let us discuss some preliminary concepts.

As in equation 8, we also use the function χ_{close} defined by

$$\chi_{close}(X,Y,\tau) \stackrel{\text{def}}{=} \begin{cases} \tau(Y) - (\tau(X) + |X| - 1) & \text{if } \tau(X) < \tau(Y) \\ \chi_{close}(Y,X,\tau) & \text{otherwise} \end{cases}$$

Intuitively, χ_{close} takes two nodes X, Y and maps to distance of the points in their intervals that are closest. We also set up a function χ_{far} that maps to the farthest distance between any such two numbers. We define

$$\chi_{far}(X,Y,\tau) \stackrel{\text{def}}{=} \begin{cases} \tau(Y) - \tau(X) + |Y| - 1 & \text{if } \tau(X) < \tau(Y) \\ \chi_{far}(Y,X,\tau) & \text{otherwise} \end{cases}$$

Using these functions, we define the balls of vertices at small topological order distance.

Definition 6.3 (Topological Order Difference Ball.). For each $v \in V$ and integer $K \ge 0$, we define the closed topological order difference ball by

$$C_{closed}(v, K, (\mathcal{V}, \tau)) \stackrel{aej}{=} \{ Y \in \mathcal{V} \mid \chi_{close}(X^v, Y, \tau) \le K \}$$

1.0

and the corresponding open ball by

$$C_{open}(v, K, (\mathcal{V}, \tau)) \stackrel{def}{=} \{ Y \in \mathcal{V} \mid \chi_{far}(X^v, Y, \tau) \leq K \}$$

We observe that $C_{closed}(v, K, (\mathcal{V}, \tau))$ is a decreasing, refining set over time, i.e. new all sets are subsets of earlier sets¹¹. Finally, let us give a definition for a *decremental refining graph* which eases the description of the hopset.

Definition 6.4 (Decremental Refining Graph). Given a refining partition \mathcal{V} of the universe V, and a graph H on the node set \mathcal{V} , we say that H is a decremental refining graph if at every stage, for any vertices $s, t \in V$, the distance between the nodes X^s and X^t , the nodes containing the vertices s and t respectively, is monotonically increasing in H. In particular, any decremental graph H' along with an arbitrary refining partition \mathcal{V}' of its vertex set forms the decremental refining multi-graph H'/\mathcal{V}' .

¹¹This is however not true for $C_{closed}(v, K, (\mathcal{V}, \tau))$.

It is not hard to adapt the Generalized ES-trees presented in [BPW19] combining it with a standard edge rounding technique for ES-trees (see for example [Ber16]) in order to obtain the following theorem. We refer the reader to [BPW19] for a full proof.

Lemma 6.5 (c.f. Lemma 4.2 and Lemma 4.3 in [BPW19]). Given a decremental refining graph H with refining node set \mathcal{V} part of an $\mathcal{ATO}(\mathcal{V}, \tau)$, an approximation parameter $\epsilon > 0$, a depth threshold $\delta > 0$ and hop threshold h_i for some $0 \le h_i \le \lg n$, a vertex $v \in V$, and a positive integer K.

Let us define the graph H_v of v by $H_v \stackrel{def}{=} H[C_{closed}(v, K, (\mathcal{V}, \tau))].$

Then, there exists a deterministic data structure \mathcal{GES}_v that can maintain the distance estimates $\widetilde{\mathbf{dist}}(X^v, Y)$ from X^v , the node in \mathcal{V} containing v, to every node $Y \in C_{closed}(v, K, (\mathcal{V}, \tau))$ such that $\mathbf{dist}_{G_v}(X^v, Y) \leq \widetilde{\mathbf{dist}}(X^v, Y)$ and

$$\operatorname{dist}(X^v, Y) \le (1+\epsilon)\operatorname{dist}_{H_v}^{h_i}(X^v, Y).$$

The algorithm runs in total update time

$$O(|E_H(C_{open}(v, K, (\mathcal{V}, \tau)))|h_i \log K \log W n/\epsilon + K \log K + \Delta)$$

where W is the largest integer weight in H and the former term is the number of edges in H that is incident to a node in $C_{open}(v, K, (V, \tau))$ throughout the entire course of the algorithm, and Δ is the number of edge weight increases for edges incident to $C_{open}(v, K, (V, \tau))$.

Hopset. For different hop thresholds h_i , we describe how to maintain the hopset H_i which forms a decremental refining graph on the node set \mathcal{V} .

If $h_i < n^{2/3} \log n$, we let H_i be the empty graph throughout the algorithm. Otherwise, we initialize the data structure to maintain the hopset H_i by calculating $l_i = \frac{h_i}{n^{1/3}}$, and sampling each vertex v in V into the set S of sampled vertices with probability $\frac{3(c+6) \cdot \log n}{l_i}$. We run a data structure \mathcal{GES}_s for every vertex $s \in S$, with hop l_i on the graph G/\mathcal{V} where we choose $K_i = \frac{q \cdot \delta}{n^{1/3}}$ as described in Lemma 6.5.

We maintain the hopset H_i so that for any nodes $X, Y \in \mathcal{V}$, there is an edge (X, Y) if

- 1. there exists vertices $s, t \in S$, with $s \in X$ and $t \in Y$, and
- 2. $\operatorname{dist}(X, Y)$ is at most l_i .

If such an edge is in H_i , we assign it weight $\operatorname{dist}(X, Y)$. It is straight-forward to see that we can maintain the required distance estimates $\operatorname{dist}(X, Y)$ for the required depth using the data structure \mathcal{GES}_s (technically there can be multiple s' in X but in this case their $\mathcal{GES}_{s'}$ data structures maintain the same distance estimates so there is no ambiguity in our notation). We observe that H_i is a decremental refining graph which follows since distances in $H_s = (G/\mathcal{V}) [C_{closed}(s, K, (\mathcal{V}, \tau))]$ can only increase over time and only when the distance between the nodes containing s and t exceeds l_i in this graph, we delete (s, t) from H_i .

SSSP on the Hopset. Finally, we can run from the root vertex r, the data structure \mathcal{GES}_r as described in Lemma 6.5 for K = n on the graph $(G/\mathcal{V}) \bigcup (\cup_i H_i)$ with the given weights and (\mathcal{V}, τ) to depth δ and hop $n^{2/3} \log n$. This completes the description of the algorithm.

6.2 Bootstrapping an Algorithm for Unrestricted Depth

We can now prove the theorem below which is a more detailed version of Theorem 3.7. Using Theorem 4.4, this immediately implies Theorem 1.2. Since the proof is quite similar to the proof of Theorem 5.4, we defer the proof to Appendix G.

Theorem 6.6. For any $0 \le i \le \lg(Wn)$, given a decremental digraph G = (V, E, w), we can maintain a hierarchy $S = \{S_i\}_i$ where each S_i is a $\mathcal{ATO}(G, 2^i, 40c \log n)$ -bundle of expected quality $n/2^i$. The algorithm runs in total expected update time $O(mn^{2/3} \log^{16} n \log^3(Wn))$ against a non-adaptive adversary and is correct with high probability.

7 Conclusion

In this article, we gave the first near-optimal algorithm to solve the decremental SSSP problem on dense digraphs. Combined with the recent result by Gutenberg et al [PVW20], this establishes an $\tilde{O}(n^2 \log^4 W/\epsilon)$ complexity for the partially-dynamic SSSP problem in directed weighted graphs. Moreover, we gave a simple new technique to derive a data structure for sparse graphs using our framework, which runs in total time $\tilde{O}(mn^{2/3} \log^3 W/\epsilon)$ and thereby vastly improves over the best existing result by Probst Gutenberg and Wulff-Nilsen as recently shown in [GW20a].

This substantial progress on the problem motivates the following two open questions:

- Can we obtain a near-optimal algorithm for partially-dynamic SSSP in directed graphs for any sparsity? Recent work by Bernstein et al. [BPW19] has shown that this is at least possible for the simpler problem of maintaining single-source reachability.
- Can we derandomize our results or make them work against an adaptive adversary? Partial progress on this question has been made in [PVW20] where the incremental SSSP data structure presented is already deterministic which gave the first deterministic improvement in directed graphs over the ES-tree. In decremental graphs, a recent result by Bernstein et al. [BPS20] gives deterministic total update time $O(n^{2+2/3})$ following a result by Probst Gutenberg and Wulff-Nilsen [GW20a] that broke the O(mn) bound randomized but against an adaptive adversary. However, this is far from the near-optimal bound achieved in this paper. We point out that even in the undirected, unweighted setting, the best data structure requires total update time $\tilde{O}(\min\{mn^{0.5+o(1)}, n^2\})$ [GW20b; BC16] as opposed to $m^{1+o(1)}$ total update time in the non-oblivious setting [HKN14b].

Acknowledgements. Aaron Bernstein is supported by NSF CAREER Grant 1942010 and the Simons Group for Algorithms & Geometry. Maximilian Probst Gutenberg is supported by Basic Algorithms Research Copenhagen (BARC), supported by Thorup's Investigator Grant

from the Villum Foundation under Grant No. 16582. Christian Wulff-Nilsen is supported by the Starting Grant 7027-00050B from the Independent Research Fund Denmark under the Sapere Aude research career programme. The authors thank Thatchaphol Saranurak for insightful comments and corrections, and anonymous FOCS reviewers for their helpful feedback.

References

- [AC13] Ittai Abraham and Shiri Chechik. "Dynamic Decremental Approximate Distance Oracles with $(1 + \epsilon, 2)$ stretch". In: *arXiv preprint arXiv:1307.1516* (2013) (cit. on p. 39).
- [ACK17] Ittai Abraham, Shiri Chechik, and Sebastian Krinninger. "Fully dynamic all-pairs shortest paths with worst-case update-time revisited". In: Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms. SIAM. 2017, pp. 440– 452 (cit. on p. 39).
- [AW14] Amir Abboud and Virginia Vassilevska Williams. "Popular conjectures imply strong lower bounds for dynamic problems". In: *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on.* IEEE. 2014, pp. 434–443 (cit. on p. 1).
- [Bar96] Yair Bartal. "Probabilistic approximation of metric spaces and its algorithmic applications". In: Proceedings of 37th Conference on Foundations of Computer Science. IEEE. 1996, pp. 184–193 (cit. on p. 11).
- [BC16] Aaron Bernstein and Shiri Chechik. "Deterministic decremental single source shortest paths: beyond the o (mn) bound". In: *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. ACM. 2016, pp. 389–397 (cit. on pp. 1, 32).
- [BC17] Aaron Bernstein and Shiri Chechik. "Deterministic partially dynamic single source shortest paths for sparse graphs". In: *Proceedings of the Twenty-Eighth Annual* ACM-SIAM Symposium on Discrete Algorithms. SIAM. 2017, pp. 453–469 (cit. on pp. 1, 5).
- [Ber+20] Aaron Bernstein, Jan van den Brand, Maximilian Probst Gutenberg, Danupon Nanongkai, Thatchaphol Saranurak, Aaron Sidford, and He Sun. "Fully-Dynamic Graph Sparsifiers Against an Adaptive Adversary". In: arXiv preprint arXiv:2004.08432 (2020) (cit. on p. 1).
- [Ber16] Aaron Bernstein. "Maintaining shortest paths under deletions in weighted directed graphs". In: SIAM Journal on Computing 45.2 (2016), pp. 548–574 (cit. on pp. 31, 39).
- [Ber17] Aaron Bernstein. "Deterministic partially dynamic single source shortest paths in weighted graphs". In: 44th International Colloquium on Automata, Languages, and Programming, ICALP 2017. Schloss Dagstuhl-Leibniz-Zentrum fur Informatik GmbH, Dagstuhl Publishing. 2017, p. 44 (cit. on pp. 1, 3, 5, 6).
- [BHS07] Surender Baswana, Ramesh Hariharan, and Sandeep Sen. "Improved decremental algorithms for maintaining transitive closure and all-pairs shortest paths". In: *Journal of Algorithms* 62.2 (2007), pp. 74–92 (cit. on p. 39).

- [BN19] Jan van den Brand and Danupon Nanongkai. "Dynamic Approximate Shortest Paths and Beyond: Subquadratic and Worst-Case Update Time". In: *arXiv* preprint arXiv:1909.10850 (2019) (cit. on p. 39).
- [BPS20] Aaron Bernstein, Maximilian Probst Gutenberg, and Thatchaphol Saranurak. "Deterministic Decremental Reachability, SCC, and Shortest Paths via Directed Expanders and Congestion Balancing". In: 2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS). IEEE. 2020 (cit. on pp. 1, 32, 39).
- [BPW19] Aaron Bernstein, Maximilian Probst, and Christian Wulff-Nilsen. "Decremental strongly-connected components and single-source reachability in near-linear time". In: Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing. ACM. 2019, pp. 365–376 (cit. on pp. 4, 9, 12, 20, 21, 31, 32, 39, 43).
- [BR11] Aaron Bernstein and Liam Roditty. "Improved dynamic algorithms for maintaining approximate shortest paths under deletions". In: *Proceedings of the twenty*second annual ACM-SIAM symposium on Discrete Algorithms. Society for Industrial and Applied Mathematics. 2011, pp. 1355–1365 (cit. on pp. 1, 39).
- [Che+16] Shiri Chechik, Thomas Dueholm Hansen, Giuseppe F Italiano, Jakub Łącki, and Nikos Parotsidis. "Decremental single-source reachability and strongly connected components in O (m√n) total update time". In: Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on. IEEE. 2016, pp. 315–324 (cit. on pp. 8, 12, 20, 21, 39, 43).
- [Che18] Shiri Chechik. "Near-Optimal Approximate Decremental All Pairs Shortest Paths".
 In: 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS). IEEE. 2018, pp. 170–181 (cit. on p. 39).
- [Chu+19] Julia Chuzhoy, Yu Gao, Jason Li, Danupon Nanongkai, Richard Peng, and Thatchaphol Saranurak. "A Deterministic Algorithm for Balanced Cut with Applications to Dynamic Connectivity, Flows, and Beyond". In: arXiv preprint arXiv:1910.08025 (2019) (cit. on p. 1).
- [CK19] Julia Chuzhoy and Sanjeev Khanna. "A new algorithm for decremental singlesource shortest paths with applications to vertex-capacitated flow and cut problems". In: Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing. ACM. 2019, pp. 389–400 (cit. on pp. 1, 2).
- [Cor+09] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. Introduction to algorithms. MIT press, 2009 (cit. on p. 24).
- [CZ20] Shiri Chechik and Tianyi Zhang. "Dynamic low-stretch spanning trees in subpolynomial time". In: Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms. SIAM. 2020, pp. 463–475 (cit. on p. 11).
- [DI01] Camil Demetrescu and Giuseppe F Italiano. "Fully dynamic all pairs shortest paths with real edge weights". In: Proceedings 42nd IEEE Symposium on Foundations of Computer Science. IEEE. 2001, pp. 260–267 (cit. on p. 39).

- [DI04] Camil Demetrescu and Giuseppe F Italiano. "A new approach to dynamic all pairs shortest paths". In: *Journal of the ACM (JACM)* 51.6 (2004), pp. 968–992 (cit. on p. 39).
- [ES81] Shimon Even and Yossi Shiloach. "An on-line edge-deletion problem". In: *Journal* of the ACM (JACM) 28.1 (1981), pp. 1–4 (cit. on pp. 1, 5, 49).
- [FG19] Sebastian Forster and Gramoz Goranci. "Dynamic low-stretch trees via dynamic low-diameter decompositions". In: Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing. 2019, pp. 377–388 (cit. on p. 11).
- [Fle00] Lisa K Fleischer. "Approximating fractional multicommodity flow independent of the number of commodities". In: SIAM Journal on Discrete Mathematics 13.4 (2000), pp. 505–520 (cit. on p. 2).
- [GK07] Naveen Garg and Jochen Koenemann. "Faster and simpler algorithms for multicommodity flow and other fractional packing problems". In: *SIAM Journal on Computing* 37.2 (2007), pp. 630–652 (cit. on p. 2).
- [GW20a] Maximilian Probst Gutenberg and Christian Wulff-Nilsen. "Decremental sssp in weighted digraphs: Faster and against an adaptive adversary". In: Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms. SIAM. 2020, pp. 2542–2561 (cit. on pp. 1–5, 8–10, 12, 32).
- [GW20b] Maximilian Probst Gutenberg and Christian Wulff-Nilsen. "Deterministic algorithms for decremental approximate shortest paths: Faster and simpler". In: Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms. SIAM. 2020, pp. 2522–2541 (cit. on pp. 1, 32).
- [Hen+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: Proceedings of the forty-seventh annual ACM symposium on Theory of computing. ACM. 2015, pp. 21–30 (cit. on p. 1).
- [HK95] Monika Rauch Henzinger and Valerie King. "Fully dynamic biconnectivity and transitive closure". In: Foundations of Computer Science, 1995. Proceedings., 36th Annual Symposium on. IEEE. 1995, pp. 664–672 (cit. on pp. 1, 39).
- [HKN14a] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. "Decremental single-source shortest paths on undirected graphs in near-linear total update time". In: Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on. IEEE. 2014, pp. 146–155 (cit. on pp. 1, 39).
- [HKN14b] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. "Sublineartime decremental algorithms for single-source reachability and shortest paths on directed graphs". In: Proceedings of the forty-sixth annual ACM symposium on Theory of computing. ACM. 2014, pp. 674–683 (cit. on pp. 1, 2, 32, 39).

- [HKN15] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. "Improved algorithms for decremental single-source reachability on directed graphs". In: *International Colloquium on Automata, Languages, and Programming.* Springer. 2015, pp. 725–736 (cit. on pp. 1, 2, 39).
- [HKN16] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. "Dynamic Approximate All-Pairs Shortest Paths: Breaking the O(mn) Barrier and Derandomization". In: SIAM Journal on Computing 45.3 (2016), pp. 947–1006 (cit. on pp. 1, 39).
- [Ita+17] Giuseppe F Italiano, Adam Karczmarz, Jakub Łącki, and Piotr Sankowski. "Decremental single-source reachability in planar digraphs". In: Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing. ACM. 2017, pp. 1108– 1121 (cit. on p. 39).
- [Kar18] Adam Karczmarz. "Decrementai transitive closure and shortest paths for planar digraphs and beyond". In: Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms. SIAM. 2018, pp. 73–92 (cit. on p. 39).
- [Kin99] Valerie King. "Fully Dynamic Algorithms for Maintaining All-Pairs Shortest Paths and Transitive Closure in Digraphs". In: 40th Annual Symposium on Foundations of Computer Science, FOCS '99, 17-18 October, 1999, New York, NY, USA. IEEE Computer Society, 1999, pp. 81–91 (cit. on pp. 1, 39).
- [KŁ20] Adam Karczmarz and Jakub Łącki. "Simple Label-Correcting Algorithms for Partially Dynamic Approximate Shortest Paths in Directed Graphs". In: Symposium on Simplicity in Algorithms. SIAM. 2020, pp. 106–120 (cit. on p. 39).
- [Łąc13] Jakub Łącki. "Improved deterministic algorithms for decremental reachability and strongly connected components". In: ACM Transactions on Algorithms (TALG) 9.3 (2013), p. 27 (cit. on p. 39).
- [LS93] Nathan Linial and Michael Saks. "Low diameter graph decompositions". In: Combinatorica 13.4 (1993), pp. 441–454 (cit. on p. 11).
- [Mad10] Aleksander Madry. "Faster approximation schemes for fractional multicommodity flow problems via dynamic graph algorithms". In: *Proceedings of the forty-second ACM symposium on Theory of computing*. ACM. 2010, pp. 121–130 (cit. on pp. 1, 2).
- [MPX13] Gary L Miller, Richard Peng, and Shen Chen Xu. "Parallel graph decompositions using random shifts". In: Proceedings of the twenty-fifth annual ACM symposium on Parallelism in algorithms and architectures. ACM. 2013, pp. 196–203 (cit. on p. 11).
- [Pac+18] Jakub Pachocki, Liam Roditty, Aaron Sidford, Roei Tov, and Virginia Vassilevska Williams. "Approximating cycles in directed graphs: Fast algorithms for girth and roundtrip spanners". In: Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms. SIAM. 2018, pp. 1374–1392 (cit. on p. 11).

- [PVW20] Maximilian Probst Gutenberg, Virginia Vassilevska Williams, and Nicole Wein. "New Algorithms and Hardness for Incremental Single-Source Shortest Paths in Directed Graphs". In: Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing. 2020 (cit. on pp. 1, 32, 39).
- [PW20] Maximilian Probst Gutenberg and Christian Wulff-Nilsen. "Fully-Dynamic All-Pairs Shortest Paths: Improved Worst-Case Time and Space Bounds". In: Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms. SIAM. 2020 (cit. on p. 39).
- [RZ04] Liam Roditty and Uri Zwick. "On dynamic shortest paths problems". In: *European Symposium on Algorithms*. Springer. 2004, pp. 580–591 (cit. on pp. 1, 39).
- [RZ08] Liam Roditty and Uri Zwick. "Improved dynamic reachability algorithms for directed graphs". In: SIAM Journal on Computing 37.5 (2008), pp. 1455–1471 (cit. on pp. 12, 39).
- [RZ12] Liam Roditty and Uri Zwick. "Dynamic approximate all-pairs shortest paths in undirected graphs". In: SIAM Journal on Computing 41.3 (2012), pp. 670–683 (cit. on p. 39).
- [RZ16] Liam Roditty and Uri Zwick. "A fully dynamic reachability algorithm for directed graphs with an almost linear update time". In: SIAM Journal on Computing 45.3 (2016), pp. 712–733 (cit. on p. 39).
- [Tar72] Robert Tarjan. "Depth-first search and linear graph algorithms". In: *SIAM journal* on computing 1.2 (1972), pp. 146–160 (cit. on p. 4).
- [Tho05] Mikkel Thorup. "Worst-case update times for fully-dynamic all-pairs shortest paths". In: *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*. ACM. 2005, pp. 112–119 (cit. on p. 39).

A Related Work

We also point out that the simpler problem of Single-Source Reachability where the data structure only has to report whether there exists a path from the fixed vertex s to some vertex $v \in V$ has been solved to near-optimality in [BPW19] which improved on the breakthrough results in [RZ08; Łąc13; HKN14b; HKN15; Che+16]. The algorithm in [BPW19] further even allows to maintain the Strongly-Connected Components of G. A recent result by Bernstein et al. [BPS20] further attempts to derandomize the algorithm in [BPW19] and deterministically achieves total update time $mn^{2/3+o(1)}$ to maintain Strongly-Connected Components, which constitutes the first deterministic improvement over the long-standing O(mn) bound.

We further point out that the problem of maintain Single-Source Shortest Paths has also been considered quite recently in the incremental setting, i.e. the setting where a graph only undergoes edge insertions. In [PVW20], the authors give a deterministic algorithm that achieves total update time $\tilde{O}(n^2 \text{polylog}(W))$. Unfortunately, the techniques proposed in this algorithm cannot be extended to the more interesting decremental setting.

There is also a wide literature on the related problems of All-Pairs Shortest Paths and All-Pairs Reachability both in the decremental [BHS07; BR11; HKN14a; HKN16; Ber16; Che18; KŁ20] and in the fully-dynamic setting [HK95; Kin99; DI01; DI04; RZ04; Tho05; RZ12; AC13; RZ16; HKN16; ACK17; BN19; PW20].

Finally, for planar graphs, decremental algorithms are known to solve Single-Source Reachability deterministically in near-linear update time [Ita+17] and SSSP in directed graphs in total update time $\tilde{O}(n^{4/3})$ as shown by [Kar18].

B Efficient Partitioning

Let us now prove Lemma 4.6 which is restated below for convenience.

Lemma 4.6 (Partitioning Procedure). Given an algorithm \mathcal{A} to solve 2-approximate δ -restricted SSSP. There exists a procedure PARTITION (G, d, ζ) that takes weighted digraph G, a depth threshold $d \leq \delta$ and a success parameter $\zeta > 0$, and returns a set $E_{Sep} \subseteq E$ such that

- 1. for each SCC X in $G \setminus E_{Sep}$, we have for any vertices $u, v \in X$ that $\operatorname{dist}_{G \setminus E_{Sep}}(u, v) \leq d$, and
- 2. for $e \in E$, we have $\mathbb{P}[e \in E_{Sep}] \leq \frac{240\zeta \log^2 n}{d} w(e)$.

The algorithm runs in total expected time

$$O\left(\sum_{j=0}^{\lceil \lg \delta \rceil} \sum_{k=0}^{2^{j+1}} T_{SSSP}(m_{j,k}, n/2^j, \delta, 2) + m \log^2 n\right)$$

where we have that $\sum_{k=0} m_{j,k} \leq 2m$ for every *i*. The algorithm terminates correctly with probability $1 - e^{-\zeta}$ for any c > 0.

Let us start by claiming that Algorithm 4 is an efficient implementation of the procedure $PARTITION(G, d, \zeta)$ that gives the guarantees stated in Lemma 4.6. In the algorithm, we first initialize the separator set to the empty set and let H throughout the algorithm be the graph G without the vertices that are already contained in an SCC that satisfies the constraints. Thus, we run multiple iterations reducing the size of H until it is empty and therefore all vertices satisfy our constraints.

Algorithm 4: PARTITION (G, d, ζ)
Input: A weighted digraph G , and integers d and ζ .
Output: Returns a set of edges E_{Sep} such that for every SCC X in $G \setminus E_{Sep}$, any two
vertices $u, v \in X$, satisfy $\operatorname{dist}_{G \setminus E_{Sep}}(u, v) \leq d$.
1 $E_{Sep} \leftarrow \emptyset; H \leftarrow G;$
2 while $H \neq \emptyset$ do
3 Pick an arbitrary vertex r in $V(H)$.
4 Run OUTSEPARATOR $(r, H, d/8, 3\zeta \log n)$ and OUTSEPARATOR $(r, H, d/8, 3\zeta \log n)$,
and let (E'_{Sep}, V'_{Sep}) be the tuple of the procedure such that $ E(V_{Sep}) $ is minimized.
5 if $ V'_{Sep} \leq \frac{2}{3}n$ then
$6 \qquad E_{small} \leftarrow \text{PARTITION}(H[V'_{Sep}], d, \zeta)$
7 $E_{Sep} \leftarrow E_{Sep} \cup E_{Small} \cup E'_{Sep}$
8 $H \leftarrow H \setminus V'_{Sep}$
9 else
10 Initialize \mathcal{A}_r to be a 2-approximate δ -restricted $SSSP$ data structure on H
/* Find a good separator for every vertex that is far from $r.$ */
11 while there exists a vertex $v \in V(H)$ such that \mathcal{A}_r has distance estimate
$\operatorname{dist}(r,v)$ or $\operatorname{dist}(v,r)$ exceeding $\frac{d}{4}$ do
12 if $dist(v,r) > d/2$ then
13 $(E''_{Sep}, V''_{Sep}) \leftarrow \text{OUTSEPARATOR}(v, H, d/8, 3\zeta \log n)$
14 else // If $\widetilde{\operatorname{dist}}(r,v) > d/2$
15 $(E''_{Sep}, V''_{Sep}) \leftarrow \text{OUTSEPARATOR}(v, \overleftarrow{H}, d/8, 3\zeta \log n)$
16 $H \leftarrow H \setminus V_{Sep}''$
17 $E_{Sep}^{\prime\prime\prime} \leftarrow \text{PARTITION}(H[V_{Sep}^{\prime\prime}], d, \zeta)$
$18 \qquad \qquad E_{Sep} \leftarrow E_{Sep} \cup E_{Sep}^{\prime\prime} \cup E_{Sep}^{\prime\prime\prime}$
$19 \qquad \qquad$
20 return E_{Sep}

In each iteration, we first pick an arbitrary vertex $r \in V(H)$ and then run the separator procedure described in lemma 3.6 on H and on H. We let (E_{Sep}, V_{Sep}) denote the tuple returned by the procedure where the size of $E(V_{Sep})$, that is the number of edges in H incident to V_{Sep} , is minimized and break ties arbitrary if they are of equal size. We then check whether $E(V_{Sep})$ contains less than a 2/3-fraction of the edges in G in line 5 (in fact, we could also compare to the number of edges in H, without affecting the asymptotic running time). In this case, we recurse on the graph $H[V_{Sep}]$, add the resulting edges to our edge set E_{Sep} and remove the vertices V_{Sep} from H.

Otherwise, we know that most edges have their tail at small distance from and to r. We therefore initialize a SSSP data structure \mathcal{A}_r from r on H in Line 10. Then, we successively extract a vertex v that is far from r, compute a separator from/to them, and prune the set of vertices separated (after recursing on them). Since on termination of the while-loop in Line 11 each vertex is close to and from r, we have that all vertices are in an SCC of small diameter. Thus, we set $H = \emptyset$.

Let us now analyze the algorithm more formally. We start by showing that no induced subgraph that we recurse on in Line 17 can contain more than $\frac{2}{3}n$ vertices. This will be crucial to bound our running time, since it implies that every time that we recurse, we recurse on subgraphs that are significantly smaller, and therefore we make significant progress.

Claim B.1. If the algorithm enters Line 9, then on termination of the while-loop starting in Line 11, we have that the set vertices V(H) in H before it is set to \emptyset is of size at least $\frac{1}{3}n$.

Proof. Observe first that since we did not enter the if-case in Line 6, that $|V'_{Sep}| > \frac{2}{3}n$. Since tuple (E'_{Sep}, V'_{Sep}) was chosen among the tuples $(E^{out}_{Sep}, V^{out}_{Sep})$ and $(E^{in}_{Sep}, V^{in}_{Sep})$ (the first corresponding to the first separator procedure, the second is the output of the second procedure in Line 4 which are computed on H), we have

$$|V_{Sep}^{out}| > \frac{2}{3}n$$
 and $|V_{Sep}^{in}| > \frac{2}{3}n$

Further, we have by Lemma 3.6 that $V_{Sep}^{out} \subseteq B_H^{out}(r, d/8)$ and $V_{Sep}^{in} \subseteq B_H^{in}(r, d/8)$.

Thus, it is straight-forward to see that

$$|B_H^{out}(r, d/8) \cap B_H^{in}(r, d/8)| \ge \frac{1}{3}n \tag{10}$$

Finally, observe that during the while-loop starting Line 11, we only prune vertices at distance at least d/8 in at least one direction. This is since distance estimates are overestimates so each vertex v from which we start the procedure OUTSEPARATOR(·) is at distance at least d/4 from r and by Lemma 3.6 we only remove vertices that are at most at distance d/8 from v, so a straight-forward application of the triangle inequality implies our claim (a slightly subtle issue is that H evolves, however distance estimates are with regard to the graph that A_r is initialized upon and H has monotonically increasing distances over time so this might only help us).

Since vertices in $B_H^{out}(r, d/8) \cap B_H^{in}(r, d/8)$ are close to r, none of them are pruned away. This is since distances do not change because even though we remove vertices from H, the SSSP data structure is run on the initial graph H, i.e. the one it was initialized upon. Thus, the vertices in the intersection of these balls remain in H until the end. Combined with Equation 10, we derive the claim. This claim is indeed the only crucial ingredient for our proof of Lemma 4.6, which we can now carry out.

Establishing Property 1. Let us prove the property by induction on the size of E. For the base case, assume $E = \emptyset$. Then, the claim is vacuously true.

For |E| = i + 1 > 0, we observe that every time, we enter the if-case in line 5 for some set V'_{Sep} , we add the edges E'_{Sep} to E_{Sep} , thus for any vertex $u \in V \setminus V_{Sep}, v \in V_{Sep}$, we have that they are not strongly-connected in $G \setminus E_{Sep}$. Thus, we do not have to establish any guarantee for these pairs. We then recurse on $H[V_{Sep}] \subseteq G[V_{Sep}]$ which has fewer edges than G by the if-condition. Thus, our claim for SCCs contained in this subgraph is true by the induction hypothesis.

Otherwise, we enter the else-case in Line 9. By Claim B.1, on termination of the while-loop starting in Line 11, the graph H is still incident to a third of the edges in G. Thus, whenever we prune away a subgraph due to taking a separator chosen in Line 12, we can again invoke the induction hypothesis when we recurse in Line 17.

Thus, it only remains to establish on termination of the while-loop in Line 11, the graph H satisfies the property. Let H' denote the graph that the data structure \mathcal{A}_r was initialized on, so $H \subseteq H' \subseteq G$. It is clear that the vertices remaining in H are at distance at most d/2 to and from r (with regard to H'). Thus, for any two such vertices $x, y \in V(H)$, we have

$$\operatorname{dist}_{G}(x,y) \leq \operatorname{dist}_{H'}(x,y) \leq \operatorname{dist}_{H'}(x,r) + \operatorname{dist}_{H'}(r,y) \leq d/2 + d/2 \leq d$$

as desired.

Establishing Property 2. Let us first observe that we only add edges to the separator set E_{Sep} in lines 7 and 18. There are further four different places that emerge where an edge e could have been added to a set that is then added to the separator set. We next observe that the edge e is either added due to a recursive call or due to invoking procedure $(E_{Sep}^{\prime\prime\prime\prime}, V_{Sep}^{\prime\prime\prime\prime}) \leftarrow \text{OutSEPARATOR}(\cdot)$ (here $(E_{Sep}^{\prime\prime\prime\prime}, V_{Sep}^{\prime\prime\prime\prime})$ is a placeholder for the tuple returned in either one of the procedures). However, if $E_{Sep}^{\prime\prime\prime\prime}$ is really added to E_{Sep} and edge e has its tail in $V_{Sep}^{\prime\prime\prime\prime\prime}$, then we have

- 1. that by Lemma 3.6, we added e to E_{Sep} with probability at most $\frac{8\cdot3\zeta\log n}{d}w(e)$, and
- 2. that $V_{Sep}^{\prime\prime\prime\prime\prime}$ is removed from H and since separator edges are computed in H, e can afterwards only be added to E_{Sep} due to a subsequent recursive call to PARTITION(\cdot), and
- 3. the graph $G[V_{Sep}^{\prime\prime\prime\prime}]$ contains at most a $\frac{2}{3}$ -fraction of the vertices in G.

Combining these facts, it is straight-forward to argue that the total probability of e being added can be bound by $\log_{3/2} n \cdot \frac{24\zeta \log n}{d} w(e) \leq \frac{240\zeta \log^2 n}{d} w(e)$.

Bounding the Error Probability. First, observe that the probability of failure each procedure $OutSeparator(\cdot)$ in our algorithm is at most $e^{-3\zeta \log n}$ by Lemma 3.6. Now observe that every separator computed by this procedure that is used, separates two vertex sets that have so far been part of the same graph. Thus, throughout all subcalls, there can be at most n-1 such separators. Further observe, that if a separator computed by the procedure is not used, then we enter the else-case in Line 9. Thus, this can only happen for each recursive call once. But since the number of recursive calls is exactly the number of used separators, this also occurs only n-1 times. Finally, we point out that in order to get a separator, we might run two separator procedures to get a single separator (namely in Line 4). Thus, the total number of separator procedures used throughout all the algorithm (including all subcalls) is at most 4(n-1). Using a simple union bound over all bad events, we can bound the probability of failure by $4n \cdot e^{-3\zeta \log n} \leq e^{-\zeta}$.

Bounding the Running Time. Let us first bound the running time of the procedure excluding recursive calls. We have that for each iteration of the while-loop in line 2, the separator procedures in line 4 can be implemented efficiently, by running in parallel (that is they are executed such that the machine interleaves operations from the separator procedures). Then, if one of them terminates with tuple (E'_{Sep}, V'_{Sep}) , it runs in time $O(|E(V'_{Sep})| \log n)$ by lemma Lemma 3.6. Further observe that if the other procedure runs longer than $O(|E(V'_{Sep})| \log n)$ by some large constant, we can abort it since it will not produce a tuple that is eligible for becoming $(E'_{Sep}, V'_{Sep})| \log n)$. If we enter the if-case in line 5, then we prune the set $E(V'_{Sep})$ from H. Thus, we can amortize all while-iterations entering the if-case over the edges removed from H, leading to total update time $O(|E| \log n)$.

The final while-iteration might enter the else-case in line 9. Then, maintaining the data structure \mathcal{A}_r takes total time $T_{SSSP}(n, m, \delta, 2)$ (here we could use $|V(H)| \leq n$ and $|E(H)| \leq m$, however this will not have significant impact).

Finally, summing over the recursive calls where we use again the argument that the number of vertices in each induced subgraph is at most $\frac{2}{3}n$ by Claim B.1. We thus obtain total running time

$$O\left(\sum_{j=0}^{\lceil \lg \delta \rceil} \sum_{k=0}^{2^{j+1}} T_{SSSP}(m_{j,k}, n/2^j, \delta, 2) + m \log^2 n\right)$$

with $\sum_{k=0}^{2^{j+1}} m_{j,k} \leq 2m$. To derive the last bound observe that each vertex (and therefore incident edge) is in a single instance on each recursive level and after at most two recursive levels, the number of vertices has decreased by factor at least 2.

C Proof of Lemma 4.11

Lemma 4.11 (c.f. also [Che+16], Lemma 13; [BPW19], Lemma 7.1). Each vertex $v \in V$ participates in C in at most $2\lceil \lg \delta \rceil$ while-loop iterations of Algorithm 3 during the entire course

of the algorithm in expectation. Further, in expectation, the SCC in G' that v is contained in, halves every second time that v participates in C.

Proof. For $v \in V$, consider any two iterations of the while-loop in Line 1 where $v \in C$, say iteration t_1 and iteration t_2 , where $t_1 < t_2$. Let $X^{v,t}$ refer to the SCC containing v in G' after the while-loop iteration t ended, where $t_1 \leq t \leq t_2$.

We claim that with probability at least 1/2, we have $|X^{v,t_2}| \leq \frac{1}{2}|X^{v,t_1}|$. This implies the theorem, since every time this event occurs the size of the SCC that v is contained in is halved (observe that sequence of trials until one halving event occurs constitutes a geometric random variable with expectation 2), and the SCC of v can be halved at most $\lg n$ times.

Let us define by G^t the graph G after the t^{th} iteration of the while loop and by H^t the graph G' after the t^{th} iteration of the while-loop. Further let us define the graph F^t , for $t \ge t_1$, to be the defined $F^t = H^{t_1} \setminus (G^t \setminus G^{t_1})$. That is F^t is the graph G^{t_1} to which only adversarial updates where applied after iteration t_1 . So that is the edge deletions to G' due to separator edges are ignored after iteration t_1 and on.

Next, let $X_{max,t} \subseteq X^{v,t_1}$ denote some maximal set of vertices such that $H^t[X_{max,t}]$ is a SCC containing v of diameter at most $\frac{\delta|X^{v,t_1}|}{16n}$. Similarly, let $Y_{max,t} \subseteq X^{v,t_1}$ denote some maximal set of vertices such that $F^t[Y_{max,t}]$ is a SCC containing v of diameter at most $\frac{\delta|X^{v,t_1}|}{16n}$. Observe that $|X_{max,j}| \leq |Y_{max,j}|$ since $F^j \supseteq H^j$.

Next, let $t_1 \leq t'$ be the index such that,

$$|Y_{max,t_1}| < \frac{1}{2} |X^{v,t'}| \le |Y_{max,t_1}|$$

that is t' is the first iteration after which $G'[X^{v,t_1}]$ has no SCC of diameter $\frac{\delta|X^{v,t_1}|}{16n}$ containing v of size more than $\frac{1}{2}|X^{v,t_1}|$.

Now consider the case where some vertex in $Y_{max,t'-1}$ was chosen as a center CENTER (X^{v,t_1}) in iteration t_1 . Then, observe that in each while-loop iteration, we always have a vertex t that is at least at distance $\frac{\delta|X^{v,t_1}|}{8n}$ from CENTER (X^{v,t_1}) . Thus, the separator is taken at distance at least $\frac{\delta|X^{v,t_1}|}{16n}$ in Line 3 or Line 5. Thus, we never remove a vertex at distance smaller-equal-than $\frac{\delta|X^{v,t_1}|}{16n}$ from CENTER (X^{v,t_1}) due to a separator procedure (and since we reuse the distance threshold until the component decreases by factor 2 in size). Thus, if such a center was chosen, none of the vertices in $Y_{max,t'-1}$ would have participated in C before iteration t'. Thus, if $t_2 < t'$, then the set C, with $v \in C$, was of size at most $\frac{1}{2}|X^{v,t_1}|$ and therefore the new SCC containing v would have at most half the size of the previous one. One the other hand, if $t_2 \ge t'$, then G'[C] would not contain a single induced SCC containing v of diameter $\frac{\delta|X^{v,t_1}|}{16n}$ that is of size larger than $\frac{1}{2}|X^{v,t_1}|$. Since we invoke PARTITION(\cdot) in that iteration on G'[C], and by the guarantees given in Lemma 4.6, the SCC X^{v,t_2} certainly satisfies $|X^{v,t_2}| \le \frac{1}{2}|X^{v,t_1}|$.

Finally, we observe that since $\frac{1}{2}|X^{v,t_1}| \leq |Y_{max,t'-1}|$, we would have chosen a center in $Y_{max,t'-1}$, with probability at least 1/2 initially. This completes the proof.

D The SSSP Data Structure

Let us now show how to prove Theorem 5.1 which is restated below for convenience.

Theorem 5.1. Given G = (V, E, w), a decremental weighted digraph, a source $r \in V$, a depth threshold $\delta > 0$, a quality parameter q, an approximation parameter $\epsilon > 0$, and access to (\mathcal{V}, τ) an $\mathcal{ATO}(G, \eta_{diam})$.

Then, there exists a deterministic data structure that maintains a distance estimate $\operatorname{dist}(r, v)$ for every vertex $v \in V$ such that at each stage of G, $\operatorname{dist}_G(r, v) \leq \widetilde{\operatorname{dist}}(r, v)$ and if $\operatorname{dist}_G(r, v) \leq \delta$ and $\mathcal{T}(\pi_{r,v}, \tau) \leq q \cdot \delta + n$, then

 $\operatorname{dist}(r, v) \leq \operatorname{dist}_G(r, v) + \eta_{diam} + \epsilon \delta.$

The total time required by this structure is

$$O(n\delta q \log n/\epsilon + n^2 \log n)$$

Remark 5.2. Technically, we require the approximate topological order (\mathcal{V}, τ) to encode changes efficiently and pass them the SSSP data structure. Since the SSSP data structure is updated only through delete operations, we require, that with each edge update, the data structure receives changes to (\mathcal{V}, τ) since the last stage. More precisely, we require that the user passes a set of pointers to each set Y that occurred in \mathcal{V} at the previous stage (denoted \mathcal{V}^{OLD}), but did not occur in \mathcal{V} at the current stage (denoted \mathcal{V}^{NEW}), i.e. each $Y \in \mathcal{V}^{OLD} \setminus \mathcal{V}^{NEW}$. Additionally, we require with each such Y that was split into subsets $Y_1, Y_2, \ldots, Y_k \in \mathcal{V}^{NEW}$ that form a partition of Y, pointers to each new element Y_i . We further require worst-case constant query time of τ , and each element $Y \in \mathcal{V}$ (for any version) can be queried for its size in constant time and returns its vertex set in time O(|Y|). For the rest of the paper, this detail will be concealed in order to improve readability.

For a description of the algorithm, we refer the reader to Section 5.1.

D.1 Correctness of the Algorithm

We first point out that subsequently in the correctness proof, the term η_{diam} actually appears in the lower bound instead of the upper bound on distance estimates. This however can be rectified by simply adding η_{diam} to every distance estimate. Further, we observe that by property 4, we obtain that $\mathcal{T}(\pi_{r,u},\tau)$ is upper bounded by $q\delta+n$, thus $\frac{\epsilon \mathcal{T}(\pi_{r,u},\tau)}{q} \leq \epsilon(\delta+n/q) \leq 2\epsilon\delta (n/q)$ is always upper bounded by δ). Thus, we get the guarantees described in Theorem 5.1 if we invoke the algorithm with $\epsilon' = \epsilon/2$ (which only affects the running time by a constant factor).

Lemma D.1 (Proof of Correctness). For each $u \in V$, we have that dist(r, u) is a monotonically increasing distance estimate such that

• $\operatorname{dist}_G(r, u) \leq \operatorname{dist}(r, u)$, at each stage, and

• if $\operatorname{dist}_G(r, u) \leq \delta$ and $\mathcal{T}(\pi_{r, u}, \tau) \leq q \cdot \delta + n$, then $\widetilde{\operatorname{dist}}(r, u) \leq \operatorname{dist}_G(r, u) + \frac{\epsilon \mathcal{T}(\pi_{r, u}, \tau)}{q} + \eta_{diam}$.

Proof. In this proof, we use the superscript t to refer to the version of a variable at the end of stage t (it will be second if there already is a superscript indicating the node X^x that some vertex x is contained in, i.e. $X^{x,t}$). We recall that we maintain the distance estimate $\widetilde{\operatorname{dist}}(r, u)$ to be equal to $\widetilde{\operatorname{dist}}(X^r, X^u) + \eta_{diam}$ and observe that the distance estimates of X^u and u are monotonically increasing over time since $X^{u,t+1}$ is a subset of the set $X^{u,t}$, and if it is a strict subset, then its distance estimate was initialized to the distance estimate $\widetilde{\operatorname{dist}}^t(X^{r,t}, X^{u,t})$ and afterwards, the update procedure can only have increase the estimate.

Now, let us establish that the distance estimate $\widetilde{\operatorname{dist}}^t(r, u) = \widetilde{\operatorname{dist}}^t(X^{r,t}, X^{u,t}) + \eta_{diam}$, at any stage t, satisfies the stretch guarantee. We start by establishing the lower bound: we therefore observe that the certificate T^t is formed by a subset of edges E^t , and the node set always corresponds to the partition \mathcal{V}^t . Thus, every path in T must be at least of length equal to the shortest-path in G^t/\mathcal{V}^t . But since property 3 enforces that any path has at most η_{diam} of its length contracted, the lower bound follows.

We now prove the upper bound by induction on the distance $\mathbf{dist}^t(r, u)$ for every $u \in V$ where we assume that $\mathbf{dist}^t(r, u) \leq \delta + \eta_{diam}$ and $\mathcal{T}(\pi_{r.u.G^t}, \tau^t) \leq q \cdot \delta + n$.

- Base case $\operatorname{dist}^t(r, u) = \eta_{diam}$: since edge weights are positive, this is only true if u = r. Since the distance estimate of r is never changed by our update procedure (since r never looses an incoming edge in T and therefore is never added to Q), the distance estimate is 0, as desired.
- Inductive step for $\operatorname{dist}^t(r, u) = d + 1 + \eta_{diam}$: Let x be the vertex in π_{r,u,G^t} that precedes \overline{u} (i.e. the second last vertex on the shortest-path). Since (x, u) is in G at stage t, we have that at every stage $t' \leq t$, the edge (x, u) is in $Q_{X^{x,(t')},X^{u,(t')}}$ and since the function $\chi(X^{x,(t')}, X^{u,(t')})$ is monotonically increasing in t', we therefore also have that $X^{x,(t')} \in B_{\leq j}(X^{u,(t')})$ for $j = \lfloor \lg \chi(X^{x,t}, X^{u,t}) \rfloor$ since our bucket update procedure always puts $X^{x,(t')}$ into the bucket $B_{j'}(X^{u,(t')})$ with the largest j' and then $\chi(X^{x,t}, X^{u,t})$ only increases during the ensuing stages.

We retell that $X^{u,0}$ had its distance estimate initialized to $\operatorname{dist}^0(X^r, X^u) \leq \operatorname{dist}^t(X^r, X^u) = d + 1$. Every time t'', it was increased ever since by at most $\lceil 2^j \cdot \frac{\epsilon}{q} \rceil$, we scanned $X^{x,t''}$ since it is in $B_{\leq j}(X^{u,t''})$. Thus, $X^{u,t}$'s distance estimate is at most $\frac{2^j \epsilon}{q} + w_{G^t}(x,u)$ larger than the distance estimate of X^x (here we use that $w_{G^t}(x,u) \geq w_{G^{t'}}(x,u)$ for all $t' \leq t$).

Therefore,

$$\begin{split} \widetilde{\mathbf{dist}}^{t}(X^{r,t}, X^{u,t}) &\leq \widetilde{\mathbf{dist}}^{t}(X^{r,t}, X^{x,t}) + w_{G^{t}}(x, u) + \frac{2^{j}\epsilon}{q} \\ &\leq \mathbf{dist}^{t}(r, x) + \eta_{diam} + \frac{\epsilon \mathcal{T}(\pi_{r,x,G^{t}}, \tau)}{q} + w_{G^{t}}(x, u) + \frac{2^{j}\epsilon}{q} \\ &\leq \mathbf{dist}^{t}(r, x) + \eta_{diam} + \frac{\epsilon \mathcal{T}(\pi_{r,x,G^{t}}, \tau)}{q} + w_{G^{t}}(x, u) + \frac{2^{j}\epsilon}{q} \\ &\leq (\mathbf{dist}^{t}(r, x) + \eta_{diam} + w_{G^{t}}(x, u)) + \frac{\epsilon \mathcal{T}(\pi_{r,x,G^{t}}, \tau)}{q} + \frac{\chi(X^{x,t}, X^{u,t}, \tau)\epsilon}{q} \\ &\leq (\mathbf{dist}^{t}(r, x) + \eta_{diam} + w_{G^{t}}(x, u)) + \frac{\epsilon \mathcal{T}(\pi_{r,x,G^{t}}, \tau)}{q} + \frac{\mathcal{T}(X^{x,t}, X^{u,t}, \tau)\epsilon}{q} \\ &= \mathbf{dist}^{t}(r, u) + \eta_{diam} + \frac{\epsilon \mathcal{T}(\pi_{r,x,G^{t}}, \tau)}{q} + \frac{\mathcal{T}((x, u), \tau)\epsilon}{q} \\ &= \mathbf{dist}^{t}(r, u) + \eta_{diam} + \frac{\epsilon \mathcal{T}(\pi_{r,u,G^{t}}, \tau)}{q} \end{split}$$

where we use the induction hypothesis on x in the second inequality, then plug in the value of j, observe that χ is dominated by \mathcal{T} and finally observe that the function \mathcal{T} is linear. We finally observe that by our assumptions $\widetilde{\operatorname{dist}}^t(X^{r,t}, X^{u,t}) \leq \operatorname{dist}^t(r, u) - \eta_{diam} + \frac{\epsilon \mathcal{T}(\pi_{r,u,G^t}, \tau)}{q} \leq \delta + \frac{\epsilon(q\delta+n)}{q} = (1+\epsilon)\delta + \epsilon n/q \leq \delta_{max}$ and since it was monotonically increasing, there was no previous stage at which the distance estimate could have been reassigned ∞ due to exceeding δ_{max} .

D.2 Running Time

It remains to establish that the algorithm given above is efficient. Let us start by analyzing the bucket update procedure.

Claim D.2. At any stage t, for any $X, Y \in \mathcal{V}$, with $2^j \leq \chi(X, Y, \tau) < 2^{j+1}$, the update procedure ensures that we have we have X in $B_j(Y)$ or $B_{j-1}(Y)$.

Proof. We observe that X should be moved to a different bucket if $\chi(X, Y, \tau)$ changes by a significant amount. We observe further, that $\chi(X, Y, \tau)$ is monotonically increasing over stages and increases between two stages $t' \leq t''$ by at most $|X^{t'} \setminus X^{t''}| + |Y^{t''} \setminus Y^{t''}|$. This follows since $\chi(X, Y, \tau)$ maps to $\tau(Y) - (\tau(X) + |X| - 1)$ which increases if either Y is mapped to a larger τ -value due to a split or due to a lower size of X, again, due to a split.

Now, let t_{last} be the last stage that X was reassigned to a bucket $B_{j'}(Y)$. Then, at that stage, X was assigned according to the initialization rule, so $2^{j'} \leq \chi(X, Y, \tau) < 2^{j'+1}$. However, this implies that since t_{last} neither X nor Y have decreased in size by $2^{j'}$ or more since otherwise

the pair would be scanned (a simple proof by contradiction establishes this claim). Thus, we can upper bound the increase in $\chi(X, Y, \tau)$ since this stage by

$$|X^t \setminus X^{t_{last}}| + |Y^t \setminus Y^{t_{last}}| < 2^{j'+1}$$

and therefore we have that $\chi(X, Y, \tau) < 2^{j'+2}$.

Claim D.3. The total time required to update the graph \mathcal{H} and the corresponding buckets is $O(n^2 \log n).$

Proof. Every time we split a node X into nodes X_1, X_2, \ldots , the largest new node X_i inherits the node of X, which can be implemented in constant time by reassigning pointers. Then, for each $X_{i'}$, we scan the incident edges to construct the new vertices. However, $X_{i'}$ is of size at most half the size of X. Thus, each vertex (and incident edge) participates only $O(\log n)$ times in such a node split. Therefore, the total time required by such scans is $O(m \log n) = O(n^2 \log n)$. Claim D.2 implies that after the update procedure, each bucket $B_i(Y)$ contains at most 2^{j+2} vertices. Since each vertex can only be in $n/2^j$ nodes that scan $B_i(Y)$ due to the size decrease (or node versions at which a size update is triggered), and since there are only $O(\log n)$ values for j, we scan at most $O(n^2 \log n)$ pairs without moving a vertex to another bucket.

Further since each X in $B_i(Y)$ can only be assigned to the same bucket or a bucket with larger index, each vertex X can change $O(\log n)$ times the bucket at Y, thus we have at most $O(n^2 \log n)$ scans overall that where some node is moved to another bucket.

Since each scan only takes constant time, our claim now follows.

Finally, we establish that our procedure to reconstruct
$$T$$
 is efficient

Claim D.4. The total update time to construct and maintain the certificate T is $O(n \log n \frac{\delta_{max}q}{\epsilon})$.

Proof. We first observe that we can initialize T by Dijkstra's algorithm in time $O(|E(\mathcal{H})| \log |E(\mathcal{H})|) =$ $O(n^2 \log n)$. On updates, for each distance value $\operatorname{dist}(X^r, Y)$, a node $Y \in \mathcal{V}$ computes j to be the largest integer such that the distance estimate is divisible by $\left[2^{j}\frac{\epsilon}{q}\right]$. Thus, for a specific j, there are at most $\frac{\delta_{max}}{2j} \cdot \frac{q}{\epsilon}$ distance values where j is computed. We then scan at the distance value, the current set $B_i(Y)$ which is of size $O(2^j)$ by Claim D.2. Note that the node Y might be in Q multiple times until its distance value increases. However, every node X in $B_i(Y)$, once scanned and not used to repair the certificate T can be ignored for this distance value since distance estimates and edge weights are monotonically increasing, thus it can never be used at a later stage to repair T at the current distance value. Thus, the total cost to scan the nodes in $B_j(Y)$ for all j can be bounded by $O(\delta_{max} \cdot \frac{q}{\epsilon})$ and since we have at most n nodes at any point and $O(\log n)$ buckets at each node, we have total cost $O(n \log n \frac{\partial_{max} q}{\epsilon})$ for these scans. We also point out that the set $B_i(Y)$ might be changed between two fixing procedures of T while the distance value of Y has not changed. However, we can amortize the scans of new items over the bucket update procedure and obtain that at most $O(n^2 \log n)$ additional scans are necessary due to changes of $B_i(Y)$ for all nodes Y.

Finally, we point out that $\frac{\delta_{max}q}{\epsilon} = \Omega(n)$ by definition and therefore the latter bound is subsumed.

Using $\epsilon' = \epsilon/2$ in the algorithms, we can combine, Lemma D.1, Claims D.3 and D.4, and the value of δ_{max} establish Theorem 5.1.

E SSSP from ATO-bundles

In this section, we prove Corollary 5.3.

Corollary 5.3. Given G = (V, E, w), a decremental weighted digraph, a source $r \in V$, a depth threshold $\delta > 0$, an approximation parameter $\epsilon > 0$, and access to a collection $S = \{S_i\}_{1 \leq i \leq \mu}$ for $\mu = \lfloor \lg \delta \rfloor - 1$ where each S_i forms an $\mathcal{ATO}(G, 2^i, 40 \operatorname{clog} n)$ -bundle of quality q_i . Then, there exists an implementation for $(1 + \epsilon)$ -approximate δ -restricted SSSP where $T_{SSSP}(n, m, \delta, \epsilon) = O(n(\max_{1 \leq i \leq \mu} \{\frac{\delta q_i}{2^i}\} + n) \log^3 n/\epsilon^2).$

Proof. We run for each $(\mathcal{V}, \tau) \in S_i$ a data structure as described in Theorem 5.1 from r on G and \overleftarrow{G} , respectively, with approximation parameter ϵ and depth threshold $\delta/(\epsilon 2^i)$. By Definition 4.3 and the guarantees given in Theorem 5.1, we have for any pair $(s,t) \in (\{r\} \times V) \cap (V \times \{r\})$, that some data structure has a distance estimate $\widetilde{\mathbf{dist}}(s,t)$ such that if $\delta/(\epsilon 2^{i-1}) \leq \mathbf{dist}(s,t) \leq \delta/(\epsilon 2^i)$, then

$$\mathbf{dist}(s,t) \le (1+4\epsilon)\mathbf{dist}(s,t).$$

Further, we have that each distance estimate is an overestimate of the actual distance. We observe that the only distance range uncovered is distances that are smaller than $1/\epsilon$. We therefore run a single instance of an ES-tree as given in [ES81] which runs in time $O(n^2/\epsilon)$.

Thus, letting $\operatorname{dist}(s,t)$ be the minimal value of any distance estimate, in any of the data structures, we have that our data structure maintains for each relevant tuple (s,t), a $(1 + 4\epsilon)$ distance estimate (this can be done by maintaining a heap over all such distance estimates). Rescaling ϵ slightly, gives the desired result.

The total update time can be bound straight-forwardly since we run $O(\log^2 n)$ instances of the data structure in Theorem 5.1 where each is at cost

$$O(n\frac{\delta}{\epsilon 2^{i}} \cdot q_{i} \log n/\epsilon + n^{2} \log n) = O(n(\max_{1 \le i \le \mu} \{\frac{\delta q_{i}}{2^{i}}\} + n) \log n/\epsilon^{2})$$

and a single ES-tree at cost $O(n^2/\epsilon)$. The costs of maintaining a heap over all distance estimates is easily subsumed by the running time.

F Proof of Theorem 6.1

In this section we are concerned with proving Theorem 6.1.

Theorem 6.1. Given G = (V, E, w), a decremental weighted digraph, a source $r \in V$, a depth threshold $\delta > 0$, a quality parameter q such that $\delta q \ge n$, an approximation parameter $\epsilon > 0$, and access to (\mathcal{V}, τ) an $\mathcal{ATO}(G, \eta_{diam})$.

Then, there exists a data structure that maintains a distance estimate $\operatorname{dist}(r, v)$ for every vertex $v \in V$ such that at each stage of G, $\operatorname{dist}_G(r, v) \leq \widetilde{\operatorname{dist}}(r, v)$ and if $\operatorname{dist}_G(r, v) \leq \delta$ and $\mathcal{T}(\pi_{r,v}, \tau) \leq q \cdot \delta + n$, then

$$\operatorname{dist}(r, v) \leq (1 + \epsilon) \operatorname{dist}_G(r, v) + \eta_{diam}.$$

The expected total time required by this structure is

$$O((m\delta q/n^{-1/3} + \delta q n^{2/3})\log^3 n\log W n/\epsilon)$$

and the data structure runs correctly with probability at least $1 - n^{-c}$ for any constant c > 0.

We start the proof by showing a lemma that bounds construction time and size of H_i . We then proceed to prove for each level *i* that we can find distance estimates for vertices at distance at most δ from the root vertex *r* using the hopset H_i and finally we combine the results to prove Corollary 6.2. Before we establish the first lemma, let us state a version of the Chernoff bound that we use frequently in our analysis.

Theorem F.1 (Multiplicative Chernoff Bound). Suppose X_1, X_2, \ldots, X_k are independent O/1random variables. Let $X = \sum_{i=1}^k$ and let the expected value of X be denoted by $\mu = E[X]$. Then for any $\delta \ge 1$, we have that

$$\mathbb{P}\left[X > (1+\delta)\mu\right] \le e^{-\frac{\delta\mu}{3}}.$$

Equipped with this theorem, let us now start the analysis.

Lemma F.2. For every $0 \le i < \lg n$, the hopset H_i has at most $O(\delta q \log n)$ edges and running time $O(m\delta q \log n \log nW/(n^{1/3}\epsilon))$ to maintain the hopset H_i with probability at least $1 - O(n^{-(c+1)})$.

Proof. The claim is trivially true for i where $h_i < n^{2/3} \log n$ since then H_i is empty by definition. Otherwise, consider any vertex $v \in V$, then we have that v is contained in a node in $C_{open}(w, K_i, (\mathcal{V}, \tau))$ during the entire course of the algorithm of at most $2K_i$ vertices $w \in V$. This follows since by definition of χ_{far} all nodes Y in the set $C_{open}(w, K_i, (\mathcal{V}, \tau))$ have their interval $[\tau(Y), \tau(Y) + |Y|)$ contained in the interval $[\tau(X^w) + |X^w| - K_i, \tau(X^w) + K_i)$. But this implies that there is only $2K_i$ vertices in such nodes.

It is not hard to establish that every set $C_{open}(w, K_i, (\mathcal{V}, \tau) \text{ contains at most } \frac{18(c+6)K_i \log^2 n}{l_i}$ vertices in S throughout the entire course of the algorithm with probability at least $1 - n^{-(c+2)}$ by a straight-forward Chernoff bound application (see Theorem F.1). Using a union bound, we can upper bound the probability that any vertex w has more than $\frac{18(c+6)K_i \log^2 n}{l_i}$ vertices in S in the set $C_{open}(w, K_i, (\mathcal{V}, \tau))$ by $1 - n^{-(c+1)}$. Henceforth in the proof, we condition on the event that no such event occurs.

Since, for every vertex $s \in S$, we have at most $\frac{18(c+6)K_i \log^2 n}{l_i} + 2$ nodes throughout the algorithm for which an edge is inserted into H_i . The additive plus 2 term stems from the fact that the sets $C_{open}(s, K_i, (\mathcal{V}, \tau))$ and $C_{closed}(s, K_i, (\mathcal{V}, \tau))$ differ by at most two nodes as pointed

out before (the two nodes that have overlapping τ -intervals with $[\tau(X^s) + |X^s| - K_i, \tau(X^s) + K_i)$).

By a similar argument as above, there are at most $\frac{3(c+1)n\log n}{l_i}$ vertices in S with probability at least $1-n^{-(c+1)}$. Conditioning on the event that the number of vertices in S does not exceed this bound, we can upper bound the number of edges in H_i by

$$\sum_{s \in S} |C_{closed}(s, K_i, (\mathcal{V}, \tau))| \le \frac{3(c+1)n\log n}{l_i} \cdot \left(\frac{18(c+2)K_i\log^2 n}{l_i} + 2\right) = O\left(\frac{K_i \cdot n\log^3 n}{l_i^2}\right).$$

We recall that $l_i = \frac{h_i}{n^{1/3}}$ and $K_i = \frac{q \cdot \delta}{n^{1/3}}$, and therefore the total number of edges can be bound by $O\left(\frac{\delta q n^{1+1/3} \cdot \log^3 n}{h_i^2}\right)$ and since $h_i \ge n^{2/3} \log n$, we have at most $O(\delta q \log n)$ edges in H_i .

Finally, let us bound the running time to maintain H_i . Using an similar argument as before, we have that every vertex $v \in V$ is only in the sets $C_{open}(s, K_i, (\mathcal{V}, \tau))$ of at most $\frac{3(c+1)K_i \log n}{l_i}$ vertices $s \in S$ with probability $1 - n^{-(c+1)}$. Conditioning on this event, we have that we have every edge in G, is contained in at most $\frac{6(c+1)K_i \log n}{l_i}$ sets $E_{G/\mathcal{V}}(C_{open}(s, K_i, (\mathcal{V}, \tau)))$ for $s \in S$. But then, the total update time to maintain all data structures \mathcal{GES}_s can be bound by

$$O\left(m \cdot \frac{6(c+1)K_i \log n \log Wn/\epsilon}{l_i} \cdot l_i\right) = O(m\delta q \log n \log nW/(n^{1/3}\epsilon))$$

It is not hard to verify that the time to maintain the hopset H_i is subsumed in the running time of the data structures \mathcal{GES}_s since after every time a node is split in such a data structure we need to can check in constant time whether the new nodes contain a vertex in S and include a new edge, and again in case the distance to a node exceeds l_i , we can in constant time remove the corresponding edge if present in H_i . Edge weight changes further coincide with increases of distance estimates. We point out that the events that we did not condition upon occur with total probability at most $O(n^{-(c+1)})$.

Lemma F.3. For any $0 \le i < \lg n$, the graph $(G/\mathcal{V}) \cup H_i$ maintains for any pair of vertices $s, t \in V$ with $\operatorname{dist}_G(s, t) = \operatorname{dist}_G^{h_i}(s, t) \le \delta$ and $\mathcal{T}(\pi_{s,t}), \tau) \le q\delta + n$, that

$$\operatorname{dist}_{(G/\mathcal{V})\cup H_i}^{n^{2/3}\log n}(X^s, X^t) \le (1+\epsilon)\operatorname{dist}_G(s, t).$$

where X^s and X^t are the nodes in \mathcal{V} that contain the vertices s and t respectively.

Proof. Clearly, for $h_i < n^{2/3} \log n$, we have that the claim is vacuously true since this implies that the path corresponding to $\operatorname{dist}_G^{h_i}(s,t)$ has less than $n^{2/3} \log n$ edges in G and therefore the path in the graph G/\mathcal{V} where some vertices are contracted into nodes can have at most the same hop.

For $h_i \ge n^{2/3} \log n$, let $\pi_{s,t}$ be a shortest path in G between s and t of hop at most h_i . Let us partition the path into segments $\pi_1, \pi_2, \ldots, \pi_k$ of exactly $l_i/2$ edges where only the last segment is possible smaller. We observe that $k \le \lceil 2h_i/l_i \rceil \le n^{2/3} + 1$. We claim that with exception of the last segment, each segment contains some vertex v that was sampled into S. We can prove that this is indeed true by using a Chernoff bound as described in theorem F.1 where for each vertex in a fixed segment we have the experiment X_i that is 1 if the vertex is sampled into S and 0 otherwise. It is not hard that for a fixed segment, there is at least one such vertex with probability $1 - n^{-(c+6)}$. Taking a union bound over the segments, we have that with probability at least $1 - n^{-(c+5)}$ every segment (except the last one) contains a vertex in S.

Next, let us select such a vertex $s_j \in S$ for each segment π_j . We observe that any two such vertices s_j and s_{j+1} in consecutive segments have a shortest path consisting of at most l_i edges between them by the optimal substructure property of shortest paths, in particular they have the shortest path $\pi_{s,t}^i[s_j, s_{j+1}]$ with at most l_i edges.

We further observe that since s_j has a \mathcal{GES}_s data structure, by the argument above, we have that there is an edge $(X^{s_j}, X^{s_{j+1}})$ in H_i of weight at most $(1 + \epsilon)\mathbf{dist}_G(s_j, s_{j+1})$ if the shortest path $\pi_{s,t}^i[s_j, s_{j+1}]$ is contained in the graph H_s . We remind the reader that the graph H_s was defined to be the graph $(G/\mathcal{V})[C_{closed}(s, K, (\mathcal{V}, \tau))]$.

It is straight-forward to observe that for every segment between two vertices s_j and s_{j+1} we have

- either the edge $(X^{s_j}, X^{s_{j+1}})$ of weight at most $(1 + \epsilon)w(\pi_{s,t}^i[s_j, s_{j+1}])$, or
- some vertex $v \in \pi_{s,t}^{i}[s_{j}, s_{j+1}]$ that is not in $C_{closed}(s, K_{i}, (\mathcal{V}, \tau))$.

Now, we observe that the latter case we have by definition of $C_{closed}(s, K_i, (\mathcal{V}, \tau))$ that $|\tau(X^v) - \tau(X^s)| \geq K_i = \frac{q \cdot \delta}{n^{1/3}}$. But since we have by the definition of quality, that $\mathcal{T}(\pi_{s,t}), \tau) \leq q\delta + n$, we have that there can be at most $\frac{q\delta+n}{q \cdot \delta/n^{1/3}} = 2n^{1/3}$ such indices j by the pigeonhole principle (here we use $\delta q \geq n$).

Finally, we can put everything together: we have that the shortest path from s to s_1 contains at most $l_i/2$ edges, analogously for the last vertex s_{k-1} to t. In between for each segment, we either have the direct edge in the hopset or we can take at most l_i edges in G/\mathcal{V} whenever we do not. Since this can occur at most $n^{1/3}$ times, we have that the path constructed by this method has hop at most

$$2 \cdot l_i + k + 2n^{1/3} l_i < 4n^{2/3}$$

since $l_i \leq n^{1/3}$ and $k \leq n^{2/3}$ for all choices of *i*. This proves the lemma.

Proof of Theorem 6.1. It is now rather straight-forward to observe that the data structure \mathcal{GES}_r maintains $(1 + \epsilon)^2$ -approximate distance estimates from X^r to every node $Y \in \mathcal{V}$ that contains a vertex t at distance at most δ from s in G and with $\mathcal{T}(\pi_{r,t}), \tau) \leq q\delta + n$ by combining Lemma 6.5 and Lemma F.3. Rescaling ϵ by a constant factor, we obtain again a $(1 + \epsilon)$ -approximation on the distance estimates without affecting the running time asymptotically.

Due to contractions according to \mathcal{V} , we might slightly underestimate the distances. However, since (\mathcal{V}, τ) form an $\mathcal{ATO}(G, \eta_{diam})$, we underestimate by at most η_{diam} and adding η_{diam} to every distance estimate remedies this problem without violation our constraints on the upper bound for the approximation. We further have that the running time of the data structure \mathcal{GES}_r is at most the number of edges in H_i and G. But by combining Lemma 6.5 and Lemma F.2, we obtain that the running time is at most

$$O((m + \sum_{i} \delta q \log n) n^{2/3} \log n \log W n/\epsilon) = O((m + \delta q) n^{2/3} \log^3 n \log W n/\epsilon).$$

Since we can maintain each H_i in time $O(m\delta q \log n \log n W/(n^{1/3}\epsilon))$ by Lemma F.2, we can bound the total update time that the data structure spends by

$$O(m\delta q \log^2 n \log nW/(n^{1/3}\epsilon) + (m+\delta q)n^{2/3} \log^3 n \log Wn/\epsilon)$$

= $O((m\delta q/n^{-1/3} + \delta q n^{2/3}) \log^3 n \log Wn/\epsilon)$

where we use again $\delta q \geq n$.

Finally, we argue that each hopset H_i is correctly maintained with probability $1-O(n^{-(c+1)})$ and therefore by a simple union bound, we can deduce that the overall algorithm runs correctly with probability at least $1 - n^{-c}$. The theorem follows.

G Proof of Theorem 6.6

In this section we prove Theorem 6.6.

Theorem 6.6. For any $0 \le i \le \lg(Wn)$, given a decremental digraph G = (V, E, w), we can maintain a hierarchy $S = \{S_i\}_i$ where each S_i is a $\mathcal{ATO}(G, 2^i, 40c \log n)$ -bundle of expected quality $n/2^i$. The algorithm runs in total expected update time $O(mn^{2/3} \log^{16} n \log^3(Wn))$ against a non-adaptive adversary and is correct with high probability.

Proof. We prove the theorem by induction on n, the number of vertices in graph G. The base case with $n \leq 1$ is easily established since we can subsume any term using a large enough constant.

Let us now give the inductive step $n \mapsto n+1$: we maintain a hierarchy of \mathcal{ATO} 's, for levels $0 \leq i \leq \log(Wn)$. At level $i \leq 2$, we let \mathcal{S}_i be a $\mathcal{ATO}(G, 2^i, 40c \log n)$ -bundle where each $(\mathcal{V}, \tau) \in \mathcal{S}_i$ is an $\mathcal{ATO}(G, 2^i)$ obtained from letting \mathcal{V} be the trivial partition of V into singletons and by letting τ be an arbitrary permutation of the elements in \mathcal{V} . It is straightforward to see that because the maximal topological difference is n, that each (\mathcal{V}, τ) has quality $n \leq \frac{(c+2)40000n \log^5 n}{2^2}$.

For each i > 2, we then iteratively construct an $\mathcal{ATO}(G, 2^i, 40c \log n)$ -bundle S_i as described in Theorem 4.4. Thus, we have to show how to implement a 2-approximate 2^{i-2} -restricted SSSP data structure required in the reduction.

We implement the SSSP data structure depending on the size of the graph that it is run upon, where we call a graph *large* if it has at least $n/2^{\gamma}$ vertices and *small* otherwise. It is here that we exploit Remark 4.5 which states that when the reduction asks to maintain approximate distances on some graph $H \subseteq G$, it is sufficient to maintain distance estimates on any graph F such that $H \subseteq F \subseteq G$ and in particular, it is ok to simply run on the entire graph G.

Now, if the graph is *large*, we run the data structure given in Corollary 6.2 on the full graph G using $S_{i-3}, S_{i-3}, \ldots, S_0$: thus each such data structure exploits the access to the existing \mathcal{ATOS} . If the graph is *small*, then we construct a new hierarchy $S' = \{S'_i\}_i$ on the graph by invoking the induction hypothesis and then run an unrestricted 2-approximate SSSP data structure derived by using Corollary 6.2 on S'. This completes the description of level i.

Let us now bound the costs to maintain a single level *i*. For $i \leq 2$, we can bound the total running time by $O(n \log n)$ since we only have to publish \mathcal{S}_0 initially.

For i > 2, we have running time

$$O\left(\sum_{j=0}^{\lceil \lg \delta \rceil} \sum_{k=0}^{2^{j+3}c \log^2 n} T_{SSSP}(m_{j,k}, n/2^j, \delta, 2) + m \log^3 n\right)$$
(11)

where $\sum_{j} m_{j,k} \leq 16c \cdot m \log^2 n$ for all k, as given by Theorem 4.4.

For each term $T_{SSSP}(m_{j,k}, n/2^j, 2^j, 2)$, where the graph is *large*, i.e. $j \leq \gamma$, then we have

$$T_{SSSP}(m_{j,k}, n/2^j, 2^j, 2) = O(mn^{2/3} \log^8 n \log^2 Wn/\epsilon)$$

by Corollary 6.2. Since we have less than $O(\log^4 n 2^{\gamma})$ terms, where $j \ge \gamma$, we have that the total cost of all SSSP data structures on *large* graphs is at most

$$O(2^{\gamma}mn^{2/3}\log^{12}n\log^2 Wn/\epsilon).$$

Further, using the induction hypothesis, we have that every term $T_{SSSP}(m_{j,k}, n/2^j, \delta, 2)$ where $j < \gamma$, requires time $O(m_{j,k}n^{2/3}\log^{16}n\log^3 Wn/\epsilon)$. Summing over all such small graph terms, we obtain total time

$$O(2^{-2/3 \cdot \gamma} m \log^{18} n \log^3(Wn)).$$

Finally, we balance both terms by setting $\gamma = \Theta(\lg(\log^2(Wn)\log^4 n))$ with a sufficiently large constant to obtain total running time $O(mn^{2/3}\log^{16} n\log^3(Wn))$. Thus, summing over $\log(Wn)$ values for *i*, the total running time is established (slightly better trade-off values can be achieved but in order to keep bounds simple we use this trade-off).

Further, we point out that every \mathcal{ATO} that was constructed runs correctly with high probability $1 - n^{-c'}$ for constant c' > 0, and since there are only polynomially many instances, we can set c' large enough to ensure that the entire algorithm runs correctly with probability $1 - n^{-c}$ for any c > 0 by taking a union bound over the events that an instance fails. \Box