

Deterministic Partially Dynamic Single Source Shortest Paths in Weighted Graphs*

Aaron Bernstein

Technical University of Berlin, Berlin, Germany
bernstei@gmail.com

Abstract

In this paper we consider the decremental single-source shortest paths (SSSP) problem, where given a graph G and a source node s the goal is to maintain shortest distances between s and all other nodes in G under a sequence of online adversarial edge deletions. In their seminal work, Even and Shiloach [JACM 1981] presented an exact solution to the problem in unweighted graphs with only $O(mn)$ total update time over all edge deletions. Their classic algorithm was the state of the art for the decremental SSSP problem for three decades, even when approximate shortest paths are allowed.

The first improvement over the Even-Shiloach algorithm was given by Bernstein and Roditty [SODA 2011], who for the case of an unweighted and undirected graph presented a $(1 + \epsilon)$ -approximate algorithm with constant query time and a total update time of $O(n^{2+o(1)})$. This work triggered a series of new results, culminating in a recent breakthrough of Henzinger, Krinninger and Nanongkai [FOCS 14], who presented a $(1 + \epsilon)$ -approximate algorithm for undirected weighted graphs whose total update time is near linear: $O(m^{1+o(1)} \log(W))$, where W is the ratio of the heaviest to the lightest edge weight in the graph. In this paper they posed as a major open problem the question of derandomizing their result.

Until very recently, all known improvements over the Even-Shiloach algorithm were randomized and required the assumption of a non-adaptive adversary. In STOC 2016, Bernstein and Chechik showed the first *deterministic* algorithm to go beyond $O(mn)$ total update time: the algorithm is also $(1 + \epsilon)$ -approximate, and has total update time $\tilde{O}(n^2)$. In SODA 2017, the same authors presented an algorithm with total update time $\tilde{O}(mn^{3/4})$. However, both algorithms are restricted to undirected, unweighted graphs. We present the *first* deterministic algorithm for *weighted* undirected graphs to go beyond the $O(mn)$ bound. The total update time is $\tilde{O}(n^2 \log(W))$.

1998 ACM Subject Classification G.2.2 Graph Algorithms

Keywords and phrases Shortest Paths, Dynamic Algorithms, Deterministic, Weighted Graph

Digital Object Identifier 10.4230/LIPIcs.ICALP.2017.44

1 Introduction

The objective of dynamic graph algorithms is to handle an online sequence of update operations while maintaining a desirable functionality on the graph, e.g., the ability to answer shortest path queries. An update operation may involve a deletion or insertion of an edge or a node, or a change in an edge's weight. In case the algorithm can handle only deletions and weight increases it is called decremental, if it can handle only insertions and weight decreases it is called incremental, and if it can handle both it is called fully dynamic.

* This work was supported by the Einstein Grant at Technical University Berlin.



© Aaron Bernstein;

licensed under Creative Commons License CC-BY

44th International Colloquium on Automata, Languages, and Programming (ICALP 2017).

Editors: Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl;

Article No. 44; pp. 44:1–44:14



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



In this paper we consider the problem of (approximate) single source shortest paths (SSSP) in unweighted undirected graphs, in the decremental setting. Specifically, given an undirected graph G with positive weights and a source node s , our algorithm needs to perform the following operations:

1. $\text{Delete}(e)$ – delete the edge e from the graph
2. $\text{Increase-Weight}(e)$ – increase the weight of e
3. $\text{Distance}(v)$ – return the distance between s and v , i.e., $\mathbf{dist}(s, v)$, in the current graph G .

Fully dynamic shortest paths has a very clear motivation, as computing shortest paths in a graph is one of the fundamental problems of graph algorithms, and many shortest path applications must deal with a graph that is changing over time. The incremental setting is somewhat more restricted, but is applicable to any setting in which the network is only expanding. The decremental setting is often very important from a theoretical perspective, as decremental shortest paths (and decremental *single source* shortest paths especially) are used as a building block in a very large variety of fully dynamic shortest paths algorithms; see e.g. [16, 17, 3, 2, 20, 1] to name just a few. Decremental shortest paths can also have applications to non-dynamic graph problems; see e.g. Madry’s paper on efficiently computing multicommodity flows [18].

We say that an algorithm has an approximation guarantee of α if its output to the query $\text{Distance}(v)$ is never smaller than the actual shortest distance and is not more than α times the shortest distance. Dynamic algorithms are typically judged by two parameters: the time it takes the algorithm to adapt to an update (the Delete or Increase-Weight operation), and the time to process a query (the Distance operation). Typically one tries to keep the query time small (polylog or constant), while getting the update time as low as possible. All the algorithms discussed in this paper have constant query time, unless noted otherwise. In the decremental setting, which is the focus of this paper, one usually considers the aggregate sum of update times over the *entire* sequence of deletions, which is referred to as the *total update time*.

Related Work

The most naive solution to dynamic SSSP is to simply invoke a static SSSP algorithm after every deletion, which requires $\tilde{O}(m)$ time, using e.g. Dijkstra’s algorithm. (The \tilde{O} notation suppresses polylogarithmic factors.) For unweighted graphs, since there can be a total of m deletions, the total update time for the naive implementation is thus $\Omega(m^2)$.

For fully dynamic SSSP, nothing better than the trivial $O(m)$ time per update is known. That is, we do not know how to do better than reconstructing from scratch after every edge update. For this reason, researchers have turned to the decremental (or incremental) case in search of a better solution. The first improvement stems all the way back to 1981, when Even and Shiloach [8] showed how to achieve total update time $O(mn)$ in unweighted undirected graphs. A similar result was independently found by Dinitz [7]. This was later generalized to directed graphs by Henzinger and King [9]. This $O(mn)$ total update time bound is still the state of art, and there are conditional lower bounds [19, 15] showing that it is in fact optimal up to log factors. (The reductions are to boolean matrix multiplication and the online matrix-vector conjecture respectively).

These lower bounds motivated the study of the approximate version of this problem. In 2011, Bernstein and Roditty [6] presented the first algorithm to go beyond the $O(mn)$ bound of Even and Shiloach [8]: they presented a $(1 + \epsilon)$ -approximate decremental SSSP algorithm for undirected unweighted graphs with $O(n^{2+O(1/\sqrt{\log n})}) = O(n^{2+o(1)})$ total

update time. Henzinger, Krinninger and Nanongkai [13] later improved the total update time to $O(n^{1.8+o(1)} + m^{1+o(1)})$, and soon after the same authors [11] achieved a close to optimal total update time of $O(m^{1+o(1)} \log W)$ in undirected weighted graphs, where W is the largest weight in the graph (assuming the minimum edge weight is 1). The same authors also showed that one can go beyond $O(mn)$ total update time bound in *directed* graphs (with a $(1 + \epsilon)$ approximation) [12, 14], although the state of art is still only a small improvement: total update time $O(mn^{0.9+o(1)} \log W)$.

However, every single one of these improvements over the $O(mn)$ bound relies on randomization, and has to make the additional assumption of a *non-adaptive* adversary. In particular, they all assume that the updates of the adversary are completely independent from the shortest paths or distances returned to the user, i.e. that the updates are fixed in advance. This makes these algorithms unsuitable for many settings, and also prevents us from using them as a black box data structure. For this reason, it is highly desirable to have deterministic algorithms for the problem.

Very recently, in STOC 2016, Bernstein and Chechik presented the first *deterministic* algorithm to go beyond $O(mn)$ total update time, again with a $(1 + \epsilon)$ approximation. The total update time is $\tilde{O}(n^2)$ [4]. They followed this up with a second deterministic algorithm [5] that has total update time $\tilde{O}(n^{1.5} \sqrt{m}) = \tilde{O}(mn^{3/4})$. Both algorithms rely on the same basic technique, so this is currently the only known technique for deterministically breaking through the $O(mn)$ barrier. However, both results above were limited to undirected, unweighted graphs. In this paper, we show that this core technique can also be applied to weighted graphs.

Our Results

► **Theorem 1.** *Let G be an undirected graph with positive edge weights subject to a sequence of edge deletions and weight increases, let s be a fixed source, and let W be the ratio between the largest and smallest edge weights in the graph. There exists a deterministic algorithm that maintains $(1 + \epsilon)$ -approximate distances from s to every vertex in total update time $O(m \log^2(n) \log(nW) + n^2 \log(n) \log(nW) \epsilon^{-2})$. The query time is $O(1)$. (Like most decremental shortest paths algorithms, our result can very easily be extended with the same update time to the incremental case, where the update sequence contains edge insertions and weight increases.)*

(Technical Note: In weighted graphs the number of updates can be very large – i.e. if each update just increases some edge weight by some small ϵ – and so in addition to the total update time above, the algorithm necessarily requires an additional $O(1)$ per update. This $O(1)$ factor is present in all dynamic algorithms, and is typically omitted.)

Our algorithm does not match the randomized state of the art of $\tilde{O}(m^{1+o(1)} \log W)$ [11], but it is optimal up to log factors for dense graphs, and is the first deterministic algorithm to go beyond the $O(mn)$ barrier in weighted graphs. The algorithm is also much simpler than the randomized algorithm for weighted graphs, and it does not incur the extra $n^{o(1)}$ factor present in the randomized algorithms. Thus, our algorithm is in fact faster than *all* existing randomized algorithms for the problem in dense weighted graphs where $m = \Omega(n^{2-1/\sqrt{\log(n)}})$.

As a final remark, we note that all previous deterministic algorithms to go beyond the $O(mn)$ bound [4, 5] have the strange drawback that there is no obvious way to return an approximate path, only a distance. Our algorithm unfortunately shares this drawback. The reason is that whereas in other dynamic shortest path algorithms the distance returned corresponds to some path in the graph, our distance involves an additive error that is bounded

through a structural claim about the non-existence of certain paths, but does not itself correspond to an actual path (see Lemma 4.4 in [4], Lemma 5.3 in [5], and Lemma 5 in our paper.)

Preliminaries

In our model, an undirected weighted graph is subject to deletions and weight increases. All weights in the original graph are positive. Let $G = (V, E)$ always refer to the *current* versions of the graph. Let m refer to the number of edges in the original graph, and n to the number of vertices. Let $w(u, v)$ refer to the weight of edge (u, v) . *Let us assume for simplicity that the minimum weighted in the original graph is at least 1*, and note that since edge weights only increase, this will be true of all version of the graph. Let W be the largest edge weight to ever appear in the graph. For any pair of vertices u, v , let $\pi(u, v)$ be the shortest $u - v$ path in G (ties can be broken arbitrarily), and let $\mathbf{dist}(u, v)$ be the length of $\pi(u, v)$. Let s be the fixed source from which our algorithm must maintain approximate distances, and let ϵ refer to our approximation parameter; when the adversary queries the distance to a vertex v , the algorithm's guarantee is to return a distance $\mathbf{dist}'(v)$ such that $\mathbf{dist}(s, v) \leq \mathbf{dist}'(s, v) \leq (1 + O(\epsilon))\mathbf{dist}(s, v)$. Note that the $(1 + O(\epsilon))$ approximation factor can always be reduced to $(1 + \epsilon)$ without affecting the asymptotic running time by simply starting with a suitably smaller ϵ' .

Given any two sets S, T we define the set difference $S \setminus T$ to contain all elements s such that $s \in S$ but $s \notin T$. We will measure the update time of the dynamic subroutines used by our algorithm in terms of their *total* update time over the entire sequence of edge changes. Note that although edges in the main graph G are only being deleted, there may be edge insertions into the auxiliary graphs used by the algorithm.

2 High Level Overview

Our algorithm extends the techniques in the STOC 2016 paper of Bernstein and Chechik [4] from unweighted to weighted graphs. We now briefly summarize their approach for unweighted graphs, aiming in this overview for total update time $\tilde{O}(n^{2.5})$ instead of $\tilde{O}(n^2)$. We start with the well known fact that the classic Even and Shiloach algorithm has total update time $O(n^{2.5})$ [8] if we only care about distances less than \sqrt{n} . So the hard case is long distances in a dense graph. The key observation of [4] is that these two problematic poles cannot coexist. In particular, Bernstein and Chechik showed that any shortest path contains at most $3\sqrt{n}$ vertices of degree more than \sqrt{n} . The basic reason is that if we look at every third high-degree vertex on the shortest path, then these vertices must have mutually disjoint neighborhoods, since otherwise there would be a shorter path between them of length 2; the claim then follows because each high-degree vertex has at least \sqrt{n} neighbors, and there are only n vertices in total. Thus, Bernstein and Chechik show that the total contribution of high-degree vertices to the shortest path is small. They then argue that this allows us to effectively “ignore” all vertices of degree $\geq \sqrt{n}$, at the cost of an additive error of only $O(\sqrt{n})$. (we do not go into details of this ignoring here.) This completed their result because for distances less than $O(\sqrt{n}\epsilon^{-1})$ the classic Even and Shiloach has total update time $O(n^{2.5}\epsilon^{-1})$, whereas for longer distances we can afford the $O(\sqrt{n})$ additive error from ignoring high degree vertices (it is subsumed into the $(1 + \epsilon)$ multiplicative error), so the resulting graph only $O(n^{1.5})$ edges.

This technique has no easy extension to weighted graphs via scaling. To see this, say there were two edge weights, 1 and \sqrt{n} , and that every vertex had very few incident edges

of weight 1, but \sqrt{n} edges of weight \sqrt{n} . Then it is easy to see that distances in the graph could still be as large as n (as opposed to \sqrt{n} in the unweighted case). But we cannot scale weights down, or use some sort of hop-distance technique found elsewhere in the literature, because although the heavy edges are the ones that cause the graph to be dense, the shortest path could still contain many edges of weight 1. There is thus no way in a weighted graph to ignore high-degree vertices.

To overcome this, we switch focus to a notion of degree that considers edge weights. In particular, note that in the example above, although the graph was dense, the total number of edges of weight 1 was small. The first step in our result is to formalize this observation, and show that although we cannot simply ignore vertices of high degree, for any vertex v we *can* ignore its low-weight incident edges if there are many of them. The proof of this is similar to the proof that we can afford to ignore dense vertices in unweighted graphs, but because we have to deal with many edge different weights at once, the analysis is more sophisticated.

The problem is that on its own ignoring large neighborhoods of low-weight edges doesn't help: the result is a graph that is guaranteed to only have a small number of low-weight edges (and a medium number of medium-weight edges), but the graph can still be very dense with high-weight edges. To overcome this, we develop a variation on the standard Even and Shiloach algorithm [8] which is more efficient at dealing with heavy edges but incurs a $(1 + \epsilon)$ approximation. The basic intuition is as follows, though the details are somewhat complex. The classic algorithm guarantees that we only touch an edge (u, v) when the distance $\mathbf{dist}(s, u)$ or $\mathbf{dist}(s, v)$ increases. But if (u, v) has high weight, and $\mathbf{dist}(s, u)$ only increases by 1, then the distance increase is insignificant with respect to the total weight on (u, v) , and so we can afford to ignore it. The number of times we touch edge (u, v) thus ends up being proportional to $\epsilon^{-1}/w(u, v)$.

All in all, our algorithm is based on the framework of Bernstein and Chechik for unweighted graphs [4], but requires significant modifications both to the high-level orientation (we need a new notion of sparsity that depends on edge weights), as well as to the technical details. The basic approach introduced in [4] of ignoring certain classes of dense vertices has proved quite versatile and powerful (see the two very different applications of it in [4] and [5]), and is currently the *only* known approach for going beyond $O(mn)$ deterministically, so our extension to weighted graphs might be useful in further applications of this approach. Also, our extension of the classic Even and Shiloach algorithm is presented in extreme generality, and might prove useful in other dynamic shortest path problems where one has a graph that has more edges of high weight than of low weight.

3 The Threshold Graph

In our algorithm, each vertex will have a weight cutoff, denoted $\mathbf{cut-off}(v)$, and will end up “ignoring” the edges below that cutoff. As described in the high-level overview above, the required degree to ignore edges goes up as the edge weights go up.

► **Definition 2.** Let the level of edge (u, v) be $\mathbf{level}(u, v) = \lfloor \log(w(u, v)) \rfloor$. For any vertex v , let $I(v)$ contain all edges incident to v (I for incident), and let $I_i(v)$ contain all edges incident to v of level i or less. Now, for any positive threshold τ (not necessarily integral), and any vertex v , we define $\mathbf{cut-off}_\tau(v)$ to be the *maximum* index i such that $I_i(v)$ contains at least $\tau 2^i$ edges; if no such index exists, $\mathbf{cut-off}_\tau(v)$ is set to 0. We say that an edge (u, v) is τ -heavy if $\mathbf{level}(u, v) \leq \mathbf{cut-off}_\tau(u)$ OR $\mathbf{level}(u, v) \leq \mathbf{cut-off}_\tau(v)$; we say that it is τ -light otherwise. Let $\mathbf{HEAVY}(\tau)$ be the set of all τ -heavy edges in G , and let $\mathbf{LIGHT}(\tau)$ be the set of all τ -light edges in G . Let $G[\mathbf{HEAVY}(\tau)]$ be the graph $(V, \mathbf{HEAVY}(\tau))$. Note that the levels,

cut-offs, heaviness, and lightness are *always* defined with respect to the main graph G , never with respect to auxiliary graphs. Since τ is usually clear from context, we often just write **cut-off**(v), heavy, and light.

► **Definition 3.** Given a graph $G = (V, E)$ with $|V| = n$ and $|E| = m$, and a positive threshold τ , define the threshold graph $G_\tau = (V_\tau, E_\tau)$ as follows

- V_τ contains every vertex $v \in V$.
- V_τ also contains an additional vertex c for each connected component C in the graph $G[\text{HEAVY}(\tau)] = (V, \text{HEAVY}(\tau))$
- E_τ contains all edges in $\text{LIGHT}(\tau)$, with their original weights.
- For every vertex $v \in V$, E_τ contains an edge from v to c of weight $1/2$, where c is the component vertex in $V_\tau \setminus V$ that corresponds to the component C in $G[\text{HEAVY}(\tau)]$ that contains v .

For any pair of vertices $s, t \in V$ define $\pi_\tau(s, t)$ to be the shortest path from s to t in G_τ , and define $\mathbf{dist}_\tau(s, t)$ to be the weight of this path.

► **Lemma 4.** For any i , G_τ contains at most n edges of weight $1/2$ that are not in E (one per vertex in V), as well as at most $O(n \cdot \tau \cdot 2^i)$ edges at level i .

Proof. An edge $(u, v) \in E$ of level i is only in E_τ if (u, v) is τ -light, i.e. if **cut-off**(u) $< i$ AND **cut-off**(v) $< i$. But a vertex with **cut-off** less than i has by definition less than $2^i \tau$ incident edges of level i . ◀

The following Lemma is the weighted analogue of Lemma 4.4 in [4], but the most direct extension of that proof to the weighted case would analyze each of the $\log(W)$ edge levels separately, and thus multiply the error by $\log(W)$, eventually leading to a $\log^3(W)$ factor in the total running time of the algorithm. To avoid this extra error, we need a somewhat more sophisticated analysis than in Lemma 4.4 of [4]

► **Lemma 5.** For any graph $G = (V, E)$, any positive integer threshold τ , and any pair of vertices $s, t \in V$:

$$\mathbf{dist}_\tau(s, t) \leq \mathbf{dist}(s, t) < \mathbf{dist}_\tau(s, t) + \frac{14n}{\tau}.$$

Proof. It is not hard to see that we have $\mathbf{dist}_\tau(s, t) \leq \mathbf{dist}(s, t)$. Simply consider the shortest $s - t$ path $\pi(s, t) \in G$. All the edges in $\pi(s, t) \cap \text{LIGHT}(\tau)$ are also in G_τ . Otherwise, for any τ -heavy edge (u, v) on $\pi(s, t)$, recall that the weight $w(u, v)$ in G is always at least 1, and note that because edge (u, v) is τ -heavy, the vertices u and v must be in the same connected component in $G[\text{HEAVY}(\tau)]$, and so there is a path of length 1 from u to v in G_τ : the edge of weight $1/2$ from u to the component vertex c , and then a second edge from c to v .

We now prove that $\mathbf{dist}(s, t) \leq \mathbf{dist}_\tau(s, t) + \frac{14n}{\tau}$. Let $\pi_\tau(s, t)$ be the shortest $s - t$ path in G_τ . Consider the edge set $E^* \subseteq E$ which contains all the τ -light edges of $\pi_\tau(s, t)$ (these edges are also in E), and all τ -heavy edges in E . That is, $E^* = \text{HEAVY}(\tau) \cup (\text{LIGHT}(\tau) \cap \pi_\tau(s, t))$. Let $G^* = (V, E^*)$. We first observe that there exists an $s - t$ path in G^* . We construct this path by following $\pi_\tau(s, t)$. $\pi_\tau(s, t)$ contains τ -light edges, which are by definition in G^* , as well as subpaths of length 2 of the form $(v, c) \circ (c, w)$, where v and w are in the same connected component C in $G[\text{HEAVY}(\tau)]$. But since v and w are in the same connected component, there is a path of only heavy edges connecting them, and all heavy edges are in G^* .

Now, let $\pi^*(s, t)$ be the *shortest* $s - t$ path in G^* . Let $\mathbf{dist}^*(s, t)$ be the length of $\pi^*(s, t)$. Since G^* is a subgraph of G , we know that $\mathbf{dist}(s, t) \leq \mathbf{dist}^*(s, t)$. We now show that

$$\mathbf{dist}^*(s, t) < \mathbf{dist}_\tau(s, t) + \frac{14n}{\tau} \quad (1)$$

which will complete the proof of the lemma.

We prove Equation 1 by showing that all the heavy edges on $\pi^*(s, t)$ have total weight at most $\frac{14n}{\tau}$: since all the light edges on $\pi^*(s, t)$ are by definition also on $\pi_\tau(s, t)$, this proves the equation.

Let V_H^* be the set of vertices on $\pi^*(s, t)$ that are incident to at least one heavy edge on $\pi^*(s, t)$. Now, for any vertex $v \in V_H^*$, let $\text{BALL}(v)$ contain v and all vertices $u \in V$, such that there is an edge (u, v) with $\mathbf{level}(u, v) \leq \mathbf{cut-off}(v)$. Note that by definition of $\mathbf{cut-off}(v)$, we have $|\text{BALL}(v)| \geq 2^{\mathbf{cut-off}(v)} \tau$.

Now, let P be an ordering of the vertices in $\pi^*(s, t)$ in non-increasing order of $\mathbf{cut-off}$: so if $i < j$ then $\mathbf{cut-off}(P[i]) \geq \mathbf{cut-off}(P[j])$. We will now construct a set \mathcal{IV} of *independent* vertices in V_H^* as follows. We start with \mathcal{IV} empty. Now, go through the vertices in V_H^* according to their order in P (so in non-increasing order of $\mathbf{cut-off}$), and for each v do the following: if $\text{BALL}(v)$ is disjoint from $\text{BALL}(u)$ for every $u \in \mathcal{IV}$, add v to \mathcal{IV} . Clearly if v and w are independent then $\text{BALL}(v)$ and $\text{BALL}(w)$ are disjoint. But then recalling that for any v , $|\text{BALL}(v)| \geq 2^{\mathbf{cut-off}(v)} \tau$, and that the total number of vertices in the graph is n , we have

$$\sum_{v \in \mathcal{IV}} 2^{\mathbf{cut-off}(v)} \leq n/\tau \quad (2)$$

We will end up showing that each independent vertex v is “responsible” for at most $14 \cdot 2^{\mathbf{cut-off}(v)}$ weight from heavy edges on $\pi^*(s, t)$, which will complete our proof. We say that a non-independent vertex u belongs to independent vertex v if $\mathbf{cut-off}(v) \geq \mathbf{cut-off}(u)$ and $\text{BALL}(u)$ and $\text{BALL}(v)$ are non-disjoint. By our independence construction, every non-independent vertex u belongs to one or more independent vertices.

► **Claim 6.** *If v is independent, and u belongs to v , then the distance from v to u along $\pi^*(s, t)$ is at most $4 \cdot 2^{\mathbf{cut-off}(v)}$.*

Proof. Say, for contradiction, that this was not the case. Then, since u belongs to v , there must be some vertex $z \in \text{BALL}(v) \cap \text{BALL}(u)$. But by definition of BALL we have $\mathbf{level}(v, z) \leq \mathbf{cut-off}(v)$ and $\mathbf{level}(u, z) \leq \mathbf{cut-off}(u) \leq \mathbf{cut-off}(v)$, so by definition of \mathbf{level} there is a path from u to v through z of weight at most $4 \cdot 2^{\mathbf{cut-off}(v)}$, which contradicts $\pi^*(s, t)$ being the shortest path in G^* . ◀

Now, given any τ -heavy edge (x, y) , let the dominant endpoint of (x, y) be the endpoint that comes earlier in P : so in particular, if x is dominant then $\mathbf{cut-off}(x) \geq \mathbf{cut-off}(y)$. We say that a τ -heavy edge (x, y) belongs to independent vertex v if the dominant endpoint of (x, y) belongs to v . Note that every τ -heavy edge in $\pi^*(s, t)$ belongs to at least one independent vertex. We now show that if v is independent, then the total weight of τ -heavy edges that belong to v is at most $14 \cdot 2^{\mathbf{cut-off}(v)}$. Otherwise (for contradiction), at least half that weight would come before or after v , so w.l.o.g. let us say that there is at least $7 \cdot 2^{\mathbf{cut-off}(v)}$ weight of edges that belong to v and come after v on $\pi^*(s, t)$ – i.e., are closer to t than v is. Let us consider the τ -heavy edge (x, y) that belongs to v and is closest to t , and say that y is closer to t than x . Note that the distance from v to y along $\pi^*(s, t)$ is at least $7 \cdot 2^{\mathbf{cut-off}(v)}$, so by Claim 6, y cannot belong to v . Thus, since edge (x, y) belongs to

v , x must be the dominant endpoint and belong to v . But if x belongs to v and dominates y then $\mathbf{cut-off}(y) \leq \mathbf{cut-off}(x) \leq \mathbf{cut-off}(v)$, and the only way edge (x, y) could be τ -heavy is if $\mathbf{level}(x, y) \leq \mathbf{cut-off}(v)$, so $w(x, y) \leq 2 \cdot 2^{\mathbf{cut-off}(v)}$, so the distance from v to x along $\pi^*(s, t)$ is still at least $7 \cdot 2^{\mathbf{cut-off}(v)} - 2 \cdot 2^{\mathbf{cut-off}(v)} = 5 \cdot 2^{\mathbf{cut-off}(v)}$, which again contradicts Claim 6. Thus, the total weight of heavy edges on $\pi^*(s, t)$ is at most $14 \sum_{v \in \mathcal{TV}} 2^{\mathbf{cut-off}(v)}$, which combined with Equation 2 proves Equation 1, which proves the lemma. \blacktriangleleft

Maintaining the threshold graph G_τ as G changes is easy, because the only hard part is maintaining the connected components C in $G_{[\mathbf{HEAVY}(\tau)]}$, and dynamic connectivity in undirected graphs is a well-studied problem that admits deterministic $O(\log^2(n))$ update time. The details are very similar to those for unweighted graphs (see Lemmas 4.5 and 4.6 in [4]), so we omit the proofs of the following two lemmas.

► **Lemma 7.** *Given a graph G subject to a decremental update sequence, and a positive integer threshold τ , we can maintain the graph G_τ in total time $O(m \log^2(n))$. Moreover, the total number of edges of weight $1/2$ ever inserted into the graph is $O(n \log(n))$, and the total number of edges (u, v) that have $\mathbf{level}(u, v) = i$ when inserted is at most $n\tau 2^i$.*

► **Lemma 8.** *Given a graph G subject to a decremental update sequence, a positive integer threshold τ , and a pair of vertices u, v in G , the distance $\mathbf{dist}_\tau(u, v)$ in G_τ never decreases as the graph G changes.*

4 An Extension of the Even and Shiloach Algorithm

In the paper of Bernstein and Chechik for unweighted graphs [4], the threshold graph ended up being sparse, and so it was easy to efficiently compute distances within it using the standard Even and Shiloach algorithm with total update time $O(mn)$ [8]. In our paper, however, the threshold graph is only sparse with respect to low-weight edges; that is, because heaviness and lightness is defined in terms of edge weights, the threshold graph can have many edge of high weight. We now present a modification of the Even and Shiloach algorithm that can deal more efficiently with edges of high weight, which makes it perfectly suited to our threshold graph, which has mostly high-weight edges. Like the standard Even and Shiloach algorithm, our extension runs up to a certain depth bound d .

► **Definition 9.** Given any number d , the function $\mathbf{BOUND}_d(x)$ is equal to x if $x \leq d$, and to ∞ otherwise.

Note that to apply to the threshold-graph G_τ , the algorithm must be able to handle insertions, as long as they do not change distances. (See Lemmas 7, 8). To this end we need the following definition:

► **Definition 10.** Let $G = (V, E)$ be a graph with positive weights ≥ 1 , subject a dynamic update sequence of insertions, deletions, and weight-increases. In such a context, let \mathcal{A} contain the set of ALL edges (u, v) to appear in G at any point during the update sequence (\mathcal{A} for all): if an edge (u, v) is deleted and then inserted again, we consider this as two separate edges in \mathcal{A} . For every edge $(u, v) \in \mathcal{A}$, let $w_o(u, v)$ (o for original) be the weight of (u, v) when it first appears in \mathcal{A} (this might be in the original graph itself), and let $\mathbf{level}_o(u, v) = \lfloor \log(w_o(u, v)) \rfloor$.

► **Lemma 11.** *Let $G = (V, E)$ be an undirected graph with positive integer edge weights, s a fixed source, and $d \geq 1$ a depth bound. Say that G is subject to a dynamic sequence of*

edge insertions, deletions, and weight increases, with the property that distances in G never decrease as a result of an update. Then, there exists an algorithm **WSES**(G, s, d) (stands for weight-sensitive Even and Shiloach) that maintains approximate distances $\mathbf{dist}'(s, v)$ with $\mathbf{dist}(s, v) \leq \mathbf{dist}'(s, v) \leq \text{BOUND}_d(\mathbf{dist}(s, v))(1 + \epsilon)$, and has total update time $O(nd + \sum_{(u,v) \in \mathcal{A}} \frac{d}{\epsilon w_o(u,v)})$.

Proof. Throughout the proof of this lemma, we will often rely on the following function:

► **Definition 12.** Given any positive real numbers β and x , let $\text{ROUND}_\beta(x)$ be the smallest number $y > x$ that is an integer multiple of β .

Our algorithm **WSES** follows the basic procedure of the classic Even and Shiloach algorithm [8] (denoted **ES**), with a few modifications. For this reason we can describe many of the guarantees of **ES** without looking under the hood for exactly how they are implemented, since we keep exactly the same implementation. We only modify the classic algorithm in a few key points. Recall the definition of \mathcal{A} and $w_o(u, v)$ from Definition 10. Since weights only increase, we always have $w(u, v) \geq w_o(u, v)$.

The **ES** algorithm maintains for each vertex v a distance label $\delta(v)$ with the property that $\mathbf{dist}(s, v) = \delta(v)$. In particular, it maintains the invariant that $\delta(s) = 0$, and that if $(u, v) \in E$, then $\delta(v) \leq \delta(u) + w(u, v)$. Our algorithm will also maintain a label $\delta(v)$ for every v , but since we only need an approximation, we guarantee a weaker invariant:

► **Invariant 13 (Approximation).** We always have $\delta(s) = 0$. If p is the parent of v in the tree then $\delta(v) \geq \delta(p) + w(p, v)$. Also, for all $(u, v) \in E$, $\delta(v) \leq \delta(u) + w(u, v) + \epsilon w_o(u, v)$.

Approximation Analysis: We now show that as long as the Approximation Invariant is preserved, we always have $\mathbf{dist}(s, v) \leq \delta(v) \leq (1 + \epsilon)\mathbf{dist}(s, v)$. The fact that $\delta(v) \geq \mathbf{dist}(s, v)$ follows trivially for the same reason as in classic **ES**. To prove the other side, consider the shortest path $\pi(s, v)$ in G . Let the vertices in $\pi(s, v)$ be $s_0, s_1, \dots, s_q = v$. We show by induction that $\delta(s_i) \leq (1 + \epsilon)\mathbf{dist}(s, s_i)$. The claim is trivially true for $i = 0$ because $\delta(s) = 0$ at all times. Now, say the claim is true for s_i , and note that $\mathbf{dist}(s, s_{i+1}) = \mathbf{dist}(s, s_i) + w(s_i, s_{i+1})$. Now, by the Approximation Invariant, we have $\delta(s_{i+1}) \leq \delta(s_i) + w(s_i, s_{i+1}) + \epsilon w_o(s_i, s_{i+1}) \leq \delta(s_i) + (1 + \epsilon)w(s_i, s_{i+1})$ (the last step follows from $w(u, v) \geq w_o(u, v)$). But then by the induction hypothesis we have $\delta(s_{i+1}) \leq (1 + \epsilon)(\mathbf{dist}(s, s_i) + w(s_i, s_{i+1})) = (1 + \epsilon)\mathbf{dist}(s, s_{i+1})$.

We must now show how to efficiently maintain the Approximation Invariant. The classic **ES** algorithm does $O(1)$ time per update for basic setup, and otherwise all the work comes from charging constant time operations to various edges (u, v) and vertices v . The update time invariant in classic **ES** is as follows: we only charge vertex v when $\delta(v)$ increases, and we only charge edge (u, v) when $\delta(u)$ or $\delta(v)$ increases. Since weights are positive integers, every $\delta(u)$ must increase by at least 1, and since we only go up to distance d , each $\delta(v)$ increases at most d times, which yields $O(md)$. To improve upon this, our algorithm maintains a slightly different update time invariant

► **Invariant 14 (Update-Time).** A vertex v is only charged when $\delta(v)$ increases. An edge (u, v) is only charged when either $\text{ROUND}_{\epsilon w_o(u,v)}(\delta(u))$ or $\text{ROUND}_{\epsilon w_o(u,v)}(\delta(v))$ increases.

Since for any β , $\text{ROUND}_\beta(\delta(v))$ can increase at most d/β before it exceeds d , it is clear that the Update-Time Invariant guarantees the total update bound of the lemma.

Maintaining the Two Invariants: Now, let us recall how the standard **ES** algorithm works. This algorithm only handles deletions, so that is what we focus on first: we later show how to extend our algorithm to handle weight-increases and insertions that do not change distances. We will emphasize the parts of classic **ES** that we will later change in our modified algorithm. First off, each vertex v always maintains label information about its neighbors: *so in particular, for each edge (u, v) , vertex v stores a copy $\delta_v(u) = \delta(u)$* (FIRST-DIFFERENCE). The algorithm maintains a shortest path tree T from the root s . We say an edge (u, v) is a tree edge if $(u, v) \in T$, and a non-tree edge if $(u, v) \in E \setminus T$. Whenever a non-tree edge is deleted, shortest distances do not change, so the algorithm does not have to do much. Now, consider the deletion of a tree edge (u, v) where u is the parent of v in T . The algorithm checks in constant time, by cleverly storing local label information $\delta_v(z)$, if v has another edge to a vertex z with $\delta(v) = \delta_v(z) + w(z, v)$. If yes, v can be reconnected without a label increase, the edge (z, v) is added to the tree and all distance labels remain the same so we are done. If not, label $\delta(v)$ must increase, which in turn affects of the children of v , so *The algorithm then examines all the children of v and checks if they can be attached to the tree without increasing their distance in a similar manner.* (SECOND-DIFFERENCE) Some vertices will fail to be reattached with the same label, so the algorithm keeps a list of all vertices that need to be re-attached at a higher label. Each time the algorithm discovers that a label $\delta(v)$ must increase, it increases it by 1 (which is the most optimistic new label it can get) and returns it to the needs-reattachment list in an attempt to reattach v with this higher label. The algorithm examines the vertices in the list in an increasing order of their label. Note that the label of a vertex may be increased many times as a result of an update. *Whenever the label $\delta(v)$ of a vertex v changes, the algorithm must adjust $\delta_u(v)$ for every edge (u, v) in the graph.* (THIRD-DIFFERENCE.)

Note that the algorithm above fails to satisfy the Update-Time Invariant, because although a vertex v is only charged when $\delta(v)$ increases (otherwise the search stops at v and can be charged to the edge that led to v), edge (u, v) is charged whenever $\delta(u)$ or $\delta(v)$ increase even just by 1. In particular, the invariant is not preserved because of the work in SECOND-DIFFERENCE and THIRD-DIFFERENCE. The changes we implement are simple. We start with FIRST-DIFFERENCE: each vertex v will still store local label information $\delta_v(u)$, but now it might be slightly inaccurate. In particular, we always have

► **Invariant 15 (Local-Information).** $\delta(u) \leq \delta_v(u) \leq \text{ROUND}_{\epsilon w_o(u,v)}(\delta(u))$.

A vertex v will now choose its parent based on the slightly inaccurate local labels $\delta_u(v)$ instead of the true label $\delta(u)$. Let $p(v)$ be the parent of v in the tree.

► **Invariant 16 (Parent-Choice).** $p(v) = \text{argmax}_{u \in \text{NEIGHBORS}(v)} \{\delta_v(u) + w(u, v)\}$ and $\delta(v) = \delta_v(p(v)) + w(p(v), v)$.

The Parent-Choice and Local-Information invariants combined clearly guarantee the Approximation Invariant. We must now show how to maintain these efficiently. This leads us to THIRD-DIFFERENCE: when $\delta(v)$ changes, we only update $\delta_u(v)$ for every vertex u when $\text{ROUND}_{\epsilon w_o(u,v)}(\delta(v))$ increases. This ensures that we satisfy the Update-Time Invariant for this step of the algorithm, while still ensuring that all local information adheres to the Local-Information Invariant. We omit the full details, but it very easy to maintain a data structure for each vertex v , that returns in $O(1)$ time per edge all edges (u, v) for which $\text{ROUND}_{\epsilon w_o(u,v)}(\delta(v))$ increased. This is because $w_o(u, v)$ is completely fixed for each edge, while $\delta(v)$ only rises from 1 to d , so we only need d slots per vertex (slot i corresponding to $\delta(v)$ going up from i to $i + 1$), and $O(1)$ time per edge affected in each slot.

The last violation of the Update-Time Invariant was in SECOND-DIFFERENCE. When $\delta(v)$ changes, instead of processing all children u , we only need to process those for which $\delta_u(v)$ changes. But in our modified algorithm this only occurs when $\text{ROUND}_{\epsilon w_o(u,v)}(\delta(v))$ increases, in keeping with the Update-Time Invariant.

We have thus shown how to modify classical **ES** in a way that maintains the Approximation Invariant and the Update-Time Invariant while processing deletions. Increase-Weight(u, v) is processed in the same way: if (u, v) was not a tree edge we do nothing, and otherwise if u was the parent of v , then some invariants will be violated for v , and we fix these in exactly the same way as for a deletion.

We now turn to processing insertions. In classical **ES**, it is trivial to process an insertion of (u, v) that does decrease distances because the distance labels do not change, so we just spend $O(1)$ time fixing update $\delta_v(u)$ and $\delta_u(v)$. In our algorithm, however, a strange difficulty arises: because our labels are approximate, even if the insertion of (u, v) does not decrease distances, it might nonetheless violate the Approximation Invariant: say that $\mathbf{dist}(u) = \mathbf{dist}(v) = 1000$ but $\delta(v) = 1000(1 + \epsilon)$ while $\delta(u) = 1000$ and we insert an edge of weight 1. In this case we cannot afford to decrease $\delta(v)$ because our update time analysis relied on non-decreasing labels. But note that although v now violates the Approximation Invariant, $\delta(v)$ is still a $(1 + \epsilon)$ approximation to $\mathbf{dist}(s, v)$, because $\mathbf{dist}(s, v)$ never decreases. We thus rely on the idea used in the “Monotone” **ES** tree of Henzinger et al. [10]: we simply never decrease distance labels. More formally, our distance labels will satisfy the following *relaxed* Approximation Invariant: after every update, for every vertex v , either $\delta(v)$ satisfies the standard Approximation Invariant, or $\delta(v) = \delta^{\text{OLD}}(v)$, where $\delta^{\text{OLD}}(v)$ was the distance label before the update. By an induction on time, it is easy to see that with this relaxed Approximation Invariant, we still always have that $\delta(v) \leq (1 + \epsilon)\mathbf{dist}(s, v)$. If $\delta(v)$ satisfies the approximation invariant, the proof is the same as in the Approximation Analysis above. Otherwise, we have $\delta(v) = \delta^{\text{OLD}}(v)$, in which case by our time induction we have $\delta^{\text{OLD}}(v) \leq (1 + \epsilon)\mathbf{dist}^{\text{OLD}}(s, v)$, and since our updates are guaranteed not to decrease distances, we have $\mathbf{dist}^{\text{OLD}}(s, v) \leq \mathbf{dist}(s, v)$ and we are done. ◀

5 The decremental SSSP algorithm

We now put together all our ingredients to proving the main Theorem 1.

► **Lemma 17.** *For any threshold τ , and depth bound $d = n\epsilon^{-1}\tau^{-1}$, we can maintain approximate distances $\mathbf{dist}'_{\tau}(s, v)$ in total time $O(m \log^2(n) + n^2 \log(n)\epsilon^{-2})$, with the property that $\mathbf{dist}'_{\tau}(s, v) \leq (1 + \epsilon)\text{BOUND}_d(\mathbf{dist}_{\tau}(s, v)) + d\epsilon$*

Proof. First we use Lemma 7 to maintain the threshold graph G_{τ} in $O(m \log^2(n))$ total update time. Then, setting $\beta = \frac{d\epsilon}{2n}$, for each edge $(u, v) \in G_{\tau}$, we round the edge weight $w(u, v)$ up to the nearest multiple of β . It is easy to see that since every shortest path in the graph contains at most n edges, this weight change incurs an additive error of at most $n\beta = d\epsilon/2$. Since all edge weights are now multiple of β , we can divide all edge weights by factor of β without changing shortest paths, which results in a scaled graph G_{τ}^s with integral weights for which we have to maintain shortest distances up to depth $d^s = d/\beta = 2n\epsilon^{-1}$. We now run **WSES**(G_{τ}^s, s, d^s) in the scaled graph, and then scale the resulting distances back up by β . Since **WSES** is a $(1 + \epsilon)$ approximation, we get: $\mathbf{dist}'_{\tau}(s, v) \leq (1 + \epsilon)\text{BOUND}_d(\mathbf{dist}_{\tau}(s, v) + d\epsilon/2) \leq (1 + \epsilon)\text{BOUND}_d\mathbf{dist}_{\tau}(s, v) + d\epsilon$.

Recall that the total update time of **WSES**(G_{τ}^s, s, d^s) is $O(nd^s + \sum_{(u,v) \in \mathcal{A}} \frac{d^s}{\epsilon w_o(u,v)})$. We know that $O(nd^s) = O(n^2\epsilon^{-1})$. Now, G_{τ}^s contains two types of edges: edges between a vertex

$v \in V$ and a component vertex $c \in V_\tau \setminus C$, and τ -light edges from the original graph. By Lemma 7, the total number of component edges ever inserted is $O(n \log(n))$, and because of our scaling they all have weight at least 1, so their total contribution to the sum is at most $n \log(n) \epsilon^{-1} \cdot d^s = O(n^2 \log(n) \epsilon^{-2})$. For edges of the original graph, let us look at the contribution of all edges (u, v) with $\mathbf{level}_o(u, v) = i$. By Lemmas 4 and 7, the total number of such edges is at most $O(n\tau 2^i)$. Each such edge has $w_o(u, v) \geq 2^i$, which has been scaled down by β in G_τ^s , and so contributes $\frac{d^s \beta}{\epsilon 2^i} = \frac{d}{\epsilon 2^i}$ to the sum. Thus in total we have that each level i contributes a total of $O(n\tau d \epsilon^{-1}) = O(n^2 \epsilon^{-2})$.

We complete the proof by arguing that there are only $\log(n)$ contributing levels. For simplicity, let us look at the graph G_τ before scaling. First off, edges of weight more than $d = n\epsilon^{-1}/\tau$ can be ignored, since we do not care about distances greater than d (we are working with BOUND_d). Secondly, from the definition of $\mathbf{cut-off}(v)$, an edge (u, v) of weight less than $1/(2\tau)$ will always be heavy and so never appear in G_τ : if $2\tau w(u, v) \leq 2^{\mathbf{level}(u, v)} \tau < 1$, then $\mathbf{cut-off}(u)$ and $\mathbf{cut-off}(v)$ are always at least as large as $\mathbf{level}(u, v)$, so (u, v) will remain heavy. Thus, the only edges that appear in our graph have weight between $1/(2\tau)$ and $n/(\epsilon\tau)$, which corresponds to $\log(n/\epsilon) = O(\log(n))$ different levels of edge weights. \blacktriangleleft

► **Lemma 18.** *For any distance bound d , with total update time $O(m \log^2(n) + n^2 \log(n) \epsilon^{-2})$ we can maintain approximate distances $\mathbf{dist}_d(s, v)$ to all vertices v such that $\mathbf{dist}(s, v) \leq \mathbf{dist}_d(s, v) \leq \text{BOUND}_d(\mathbf{dist}(s, v)) + 15\epsilon d$*

Proof. Set $\tau = n/(\epsilon d)$. By Lemma 17, in total update time $O(m \log^2(n) + n^2 \log(n) \epsilon^{-2})$ we can maintain values $\mathbf{dist}'_\tau(s, v)$ with $\mathbf{dist}_\tau(s, v) \leq \mathbf{dist}'_\tau(s, v) \leq (1 + \epsilon) \text{BOUND}_d(\mathbf{dist}_\tau(s, v)) + \epsilon d$. Now, By Lemma 5 we have $\mathbf{dist}(s, v) - 14n/\tau \leq \mathbf{dist}_\tau(s, v) \leq \mathbf{dist}(s, v)$. Plugging in our value $\tau = n/(\epsilon d)$ we get $\mathbf{dist}(s, v) - 14\epsilon d \leq \mathbf{dist}_\tau(s, v) \leq \mathbf{dist}(s, v)$. Thus, we have $\mathbf{dist}(s, v) - 14\epsilon d \leq \mathbf{dist}'_\tau(s, v) \leq \text{BOUND}_d(\mathbf{dist}(s, v)) + \epsilon d$. If we return $\mathbf{dist}_d(s, v) = \mathbf{dist}'_\tau(s, v) + 14\epsilon d$, we get the bound in the lemma. \blacktriangleleft

It is now easy to prove our main Theorem 1. Observe that the maximum distance we could possibly see is nW . Thus, for each $i = 0, 1, 2, 3, \dots, \lceil \log(nW) \rceil$, we use Lemma 18 to maintain $\delta_i(v) = \mathbf{dist}_{2^i}(s, v)$. By Lemma 18, the total update time is the desired $O(m \log^2(n) \log(nW) + n^2 \log(n) \log(nW) \epsilon^{-2})$. We then output as our final answer: $\mathbf{dist}'(v) = \min_i \{\delta_i(v)\}$. Since for each i we have $\delta_i \geq \mathbf{dist}(v)$, we have $\mathbf{dist}'(v) \geq \mathbf{dist}(s, v)$. We now need to show that $\mathbf{dist}'(v) \leq (1 + O(\epsilon)) \mathbf{dist}(s, v)$. It is not hard to see that for $i = \lceil \log(\mathbf{dist}(s, v)) \rceil$, we end up with $\delta_i(v) \leq \mathbf{dist}(s, v) + 30\epsilon \mathbf{dist}(s, v)$, as desired.

6 Conclusions

In this paper we presented the first *deterministic* decremental SSSP algorithm for *weighted* undirected graphs that goes beyond the Even-Shiloach bound of $O(mn)$ total update time. Previously such a result was only known for unweighted undirected graphs. The two main open questions are further improving this total update time, and going beyond $O(mn)$ deterministically for *directed* graphs. Finally, we recall from the introduction that all existing deterministic $o(mn)$ results including ours are only able to return approximate shortest distances, not the paths themselves.

References

- 1 Ittai Abraham, Shiri Chechik, and Kunal Talwar. Fully dynamic all-pairs shortest paths: Breaking the $o(n)$ barrier. In *Approximation, Randomization, and Combinatorial Optimiz-*

- ation. *Algorithms and Techniques, APPROX/RANDOM 2014, September 4-6, 2014, Barcelona, Spain*, pages 1–16, 2014. doi:10.4230/LIPIcs.APPROX-RANDOM.2014.1.
- 2 Surender Baswana, Ramesh Hariharan, and Sandeep Sen. Maintaining all-pairs approximate shortest paths under deletion of edges. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 394–403. Society for Industrial and Applied Mathematics, 2003. URL: <http://dl.acm.org/citation.cfm?id=644108.644171>.
 - 3 Aaron Bernstein. Fully dynamic $(2 + \epsilon)$ approximate all-pairs shortest paths with fast query and close to linear update time. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science, FOCS*, pages 693–702, 2009. doi:10.1109/FOCS.2009.16.
 - 4 Aaron Bernstein and Shiri Chechik. Deterministic decremental single source shortest paths: beyond the $o(mn)$ bound. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, pages 389–397, 2016. doi:10.1145/2897518.2897521.
 - 5 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, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 453–469, 2017. doi:10.1137/1.9781611974782.29.
 - 6 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, SODA*, pages 1355–1365, 2011. doi:10.1137/1.9781611973082.104.
 - 7 Yefim Dinitz. Dinitz’ algorithm: The original version and even’s version. In *Theoretical Computer Science, Essays in Memory of Shimon Even*, pages 218–240, 2006. doi:10.1007/11685654_10.
 - 8 Shimon Even and Yossi Shiloach. An on-line edge-deletion problem. *Journal of the ACM*, 28(1):1–4, 1981. doi:10.1145/322234.322235.
 - 9 Monika Henzinger and Valerie King. Fully dynamic biconnectivity and transitive closure. In *36th Annual Symposium on Foundations of Computer Science, Milwaukee, Wisconsin, 23-25 October 1995*, pages 664–672, 1995. doi:10.1109/SFCS.1995.492668.
 - 10 Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Dynamic approximate all-pairs shortest paths: Breaking the $o(mn)$ barrier and derandomization. In *Proceedings of the 54th Annual Symposium on Foundations of Computer Science, FOCS*, pages 538–547, 2013. doi:10.1109/FOCS.2013.64.
 - 11 Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Decremental single-source shortest paths on undirected graphs in near-linear total update time. In *Proceedings of the 55th Annual Symposium on Foundations of Computer Science, FOCS*, pages 146–155, 2014. doi:10.1109/FOCS.2014.24.
 - 12 Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Sublinear-time decremental algorithms for single-source reachability and shortest paths on directed graphs. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC)*, pages 674–683, 2014. doi:10.1145/2591796.2591869.
 - 13 Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. A subquadratic-time algorithm for decremental single-source shortest paths. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 1053–1072, 2014. URL: <http://dl.acm.org/citation.cfm?id=2634074.2634153>.
 - 14 Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Improved algorithms for decremental single-source reachability on directed graphs. In *Proceedings of the 42nd International Colloquium, ICALP*, pages 725–736, 2015. doi:10.1007/978-3-662-47672-7_59.

- 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 on Symposium on Theory of Computing (STOC)*, pages 21–30, 2015. doi:10.1145/2746539.2746609.
- 16 Monika Rauch Henzinger and Valerie King. Maintaining minimum spanning forests in dynamic graphs. *SIAM J. Comput.*, 31(2):364–374, 2001. doi:10.1137/S0097539797327209.
- 17 Valerie King. Fully dynamic algorithms for maintaining all-pairs shortest paths and transitive closure in digraphs. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science, FOCS*, pages 81–91, 1999. URL: <http://dl.acm.org/citation.cfm?id=795665.796487>.
- 18 Aleksander Madry. Faster approximation schemes for fractional multicommodity flow problems via dynamic graph algorithms. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 121–130, 2010. doi:10.1145/1806689.1806708.
- 19 Liam Roditty and Uri Zwick. On dynamic shortest paths problems. *Algorithmica*, 61(2):389–401, 2011. doi:10.1007/s00453-010-9401-5.
- 20 Liam Roditty and Uri Zwick. Dynamic approximate all-pairs shortest paths in undirected graphs. *SIAM J. Comput.*, 41(3):670–683, 2012. doi:10.1137/090776573.