

# A New Approach to Dynamic All Pairs Shortest Paths <sup>\*</sup>

Camil Demetrescu <sup>†</sup>      Giuseppe F. Italiano <sup>‡</sup>

## Abstract

We study novel combinatorial properties of graphs that allow us to devise a completely new approach to dynamic all pairs shortest paths problems. Our approach yields a fully dynamic algorithm for general directed graphs with non-negative real-valued edge weights that supports any sequence of operations in  $O(n^2 \log^3 n)$  amortized time per update and unit worst-case time per distance query, where  $n$  is the number of vertices. We can also report shortest paths in optimal worst-case time. These bounds improve substantially over previous results and solve a long-standing open problem. Our algorithm is deterministic, uses simple data structures, and appears to be very fast in practice.

## 1 Introduction

In this paper we present fully dynamic algorithms for maintaining all pairs shortest paths (APSP) in directed graphs with *real-valued* edge weights. A dynamic graph algorithm maintains a given property  $\mathcal{P}$  on a weighted graph subject to dynamic changes, such as edge insertions, edge deletions and edge weight updates. Note that edge deletions and edge insertions can be easily formulated as edge weight updates, by setting to  $+\infty$  the weight of edges not in the graph. A dynamic graph algorithm should process queries on property  $\mathcal{P}$  quickly, and must perform update operations faster than recomputing the solution from scratch, as

---

<sup>\*</sup>This work has been partially supported by the IST Programme of the EU under contract n. IST-1999-14186 (ALCOM-FT), by the HPRN Programme of the EU under contract n. HPRN-CT-1999-00104 (AMORE), and by the Italian Ministry of University and Research (Project “ALINWEB: Algorithmics for Internet and the Web”). A preliminary version of this paper was presented at the *35th Annual ACM Symposium on Theory of Computing (STOC'03)* [6]. The final remarks of the conference version claimed that an  $O(n^2 \log^2 n)$  update bound could be achieved for fully dynamic all pairs shortest paths; however, this claim was supported with inaccurate arguments.

<sup>†</sup>Dipartimento di Informatica e Sistemistica, Università di Roma “La Sapienza”, Roma, Italy. Email: demetres@dis.uniroma1.it. URL: <http://www.dis.uniroma1.it/~demetres>.

<sup>‡</sup>Dipartimento di Informatica, Sistemi e Produzione, Università di Roma “Tor Vergata”, Roma, Italy and Centro “Vito Volterra”, Università di Roma “Tor Vergata”, Roma, Italy. Email: italiano@disp.uniroma2.it. URL: <http://www.info.uniroma2.it/~italiano>.

carried out by the fastest static algorithm. We say that an algorithm is *fully dynamic* if it can handle both edge weight increases and edge weight decreases. A *partially dynamic* algorithm can handle either edge weight increases or decreases, but not both.

**The Problem.** In the fully dynamic APSP problem that we consider we wish to maintain a directed graph with real-valued edge weights under an intermixed sequence of the following operations:

- `update`( $v, w'$ ): update the weights of all edges incident to vertex  $v$  according to weight function  $w'$ .
- `distance`( $x, y$ ): return the distance from vertex  $x$  to vertex  $y$ .
- `path`( $x, y$ ): report a shortest path from vertex  $x$  to vertex  $y$ , if any.

Notice that in this paper we consider a generalized version of the dynamic all pairs shortest path problem where the weights of all edges incident to a given vertex can be changed with just one update operation. In the following we will call such operation *vertex update*. We recall that in this setting, edge deletions can be modeled by raising edge weights to  $+\infty$ , while edge insertions can be realized by decreasing edge weights from  $+\infty$  to a finite value. Throughout the paper, we denote by  $n$  the number of vertices in the graph and by  $m$  the number of edges with weights  $< +\infty$  in the graph.

**Previous Work.** The dynamic maintenance of shortest paths has a remarkably long history, as the first papers date back to over 35 years ago [13, 14, 17]. After that, many dynamic shortest paths algorithms have been proposed (see e.g., [7, 9, 10, 15, 16, 18]), but their running times in the worst case were comparable to recomputing APSP from scratch.

The first dynamic shortest path algorithms which are provably faster than recomputing APSP from scratch only worked on graphs with small integer weights. In particular, Ausiello *et al.* [1] proposed a decrease-only shortest path algorithm for directed graphs having positive integer weights less than  $C$ : the amortized running time of their algorithm is  $O(Cn \log n)$  per edge insertion. Henzinger *et al.* [11] designed a fully dynamic algorithm for APSP on planar graphs with integer weights, with a running time of  $O(n^{4/3} \log(nC))$  per operation. Recently, Fakcharoemphol and Rao in [8] designed a fully dynamic algorithm for maintaining single-source shortest paths in planar directed graphs that supports both queries and edge weight updates in  $O(n^{4/5} \log^{13/5} n)$  amortized time per edge operation.

The first big step on general graphs and integer weights was made by King [12], who presented a fully dynamic algorithm for maintaining all pairs shortest paths in directed graphs with positive integer weights less than  $C$ : the running time of her algorithm is  $O(n^{2.5} \sqrt{C \log n})$  per update. In previous work [4, 5], we have solved fully dynamic APSP on general directed graphs with real weights. In particular, given a directed graph  $G$ , subject to dynamic operations, and such that each edge weight can take at most  $S$  different *real* values, we proposed a deterministic algorithm that supports each update in  $O(n^{2.5} \sqrt{S \log^3 n})$

amortized time and each query in  $O(1)$  worst-case time. Other deletions-only algorithms for APSP, in the simpler case of unweighted graphs, are presented in [2].

**Our Results.** We study novel combinatorial properties of graphs that allow us to devise a completely new approach to dynamic all pairs shortest paths. This approach yields a fully dynamic algorithm for APSP on general directed graphs with non-negative real-valued edge weights that achieves the following time bounds: it supports any sequence of operations in  $O(n^2 \log^3 n)$  amortized time per update and one look-up in the worst case per distance query; it can also report shortest paths in optimal worst-case time. We remark that our algorithm improves substantially over previous bounds [4, 5, 12]. Furthermore, and differently from all the previous approaches, it solves fully dynamic APSP in its generality. Indeed, it runs on (non-negative) real weights, and each weight has no limit on the number of different values it can take. Finally, we note that when the distance matrix has to be maintained explicitly, i.e., distance queries have to be answered with exactly one look-up, as many as  $\Omega(n^2)$  entries of the distance matrix can change during each update. Thus, in this model our algorithm is only a polylogarithmic factor away from the best possible bound. In the special case of increase-only update sequences, our techniques yield a faster update algorithm that runs in  $O(n^2 \log n)$  amortized time per operation. Similarly to the fully dynamic case, no previous general solution was known for this problem.

Another interesting feature of our techniques is that both weight increases and weight decreases can be supported with exactly the same code. Surprisingly, our algorithms are rather simple and thus amenable to efficient implementations: indeed, according to a recent experimental study [3], the techniques described in this paper are not only asymptotically efficient, but can yield very fast implementations in many practical scenarios.

**Notation.** Let  $G = (V, E)$  be a directed graph with real edge weights and no negative-length cycles. A path  $\pi_{xy} = \langle x_0, x_1, \dots, x_k \rangle$  from vertex  $x$  to vertex  $y$  in  $G$  is a sequence of vertices such that  $x_0 = x$ ,  $x_k = y$ , and  $(x_i, x_{i+1}) \in E$ , for each  $i$ ,  $0 \leq i < k$ . Let  $(u, v)$  be an edge in  $E$ : we denote by  $w_{uv}$  the weight of edge  $(u, v)$ . Let  $\pi_{xy} = \langle x_0, x_1, \dots, x_k \rangle$  be a path in  $G$ : we denote by

$$w(\pi_{xy}) = \sum_{i=0}^{k-1} w_{x_i x_{i+1}}$$

the weight of  $\pi_{xy}$ . We assume that as a special case,  $\pi_{xx} = \langle x \rangle$  is a path of weight zero.

Given  $\pi_{xv} = \langle x, \dots, x', v \rangle$  and  $\pi_{vy} = \langle v, y', \dots, y \rangle$ , we denote by  $\pi_{xv} \cdot \pi_{vy}$  the path  $\langle x, \dots, x', v, y', \dots, y \rangle$  obtained by concatenating  $\pi_{xv}$  and  $\pi_{vy}$  at  $v$ . Moreover, we denote by  $\ell(\pi_{xy})$  the path  $\pi_{xb}$  such that  $\pi_{xy} = \pi_{xb} \cdot \langle b, y \rangle$ . Similarly, we denote by  $r(\pi_{xy})$  the path  $\pi_{ay}$  such that  $\pi_{xy} = \langle x, a \rangle \cdot \pi_{ay}$ .

Let the graph  $G$  be subject to a sequence of (vertex) updates  $\Sigma = \langle \sigma_1, \sigma_2, \dots, \sigma_k \rangle$ . We denote by  $t_{\sigma_i}$  the time immediately after update  $\sigma_i$ , with  $t_{\sigma_i} = i$  for any  $i$ ,  $1 \leq i \leq k$ . We also denote by  $v_\sigma$  the vertex affected by the update  $\sigma$ : i.e., if  $\sigma = \text{update}(x, w')$ , then  $v_\sigma = x$ . The notation used in this paper is summarized in Table 1.

$G = (V, E)$	weighted directed graph with vertex set $V$ and edge set $E$
$w_{uv}$	weight of edge $(u, v)$
$\pi_{xy} = \langle x, \dots, y \rangle$	path from vertex $x$ to vertex $y$
$w(\pi_{xy})$	weight of path $\pi_{xy}$ (sum of the weights of the edges in $\pi_{xy}$ )
$\ell(\pi_{xy})$	subpath $\pi_{xb}$ of $\pi_{xy}$ such that $\pi_{xy} = \pi_{xb} \cdot \langle b, y \rangle$
$r(\pi_{xy})$	subpath $\pi_{ay}$ of $\pi_{xy}$ such that $\pi_{xy} = \langle x, a \rangle \cdot \pi_{ay}$
$\Sigma = \langle \sigma_1, \sigma_2, \dots, \sigma_k \rangle$	sequence of update operations
$t_\sigma$	time at which update $\sigma$ occurs
$v_\sigma$	vertex affected by update $\sigma$

Table 1: Notation used in the paper.

Throughout the paper we assume that there is only one shortest path between each pair of vertices in  $G$ . This is without loss of generality, since ties can be broken consistently as we will discuss in Section 3.4.

**Organization of the Paper.** The remainder of this paper is organized as follows. Section 2 studies some properties of special classes of paths in a graph, while Section 3 shows how to exploit them to devise the first general increase-only update algorithm for all pairs shortest paths. To deal with fully dynamic sequences, Section 4 addresses more path properties, which are next used in Section 5 to devise the first general algorithm for fully dynamic all pairs shortest paths. We remark that, while Section 2 and Section 4 contain combinatorial results on graphs, algorithmic aspects are treated in Section 3 and Section 5. We conclude the paper in Section 6 with some remarks and directions for further research.

## 2 Locally Shortest Paths

In this section we study the properties of a class of paths in a graph that we call *locally shortest paths*. Using this notion, in Section 3 we will show how to maintain efficiently all pairs shortest paths in a graph subject to partially dynamic edge weight updates. Locally shortest paths are defined as follows.

**Definition 1** A path  $\pi_{xy}$  is **LOCALLY SHORTEST**<sup>1</sup> in  $G$  if either:

- (i)  $\pi_{xy}$  consists of a single vertex, or
- (ii) every proper subpath of  $\pi_{xy}$  is a shortest path in  $G$ .

This definition is inspired by the optimal-substructure property of shortest paths: all subpaths of a shortest path are shortest paths. Here, we relax this property by considering only

---

<sup>1</sup>Since every internal vertex of a locally shortest path has the same sum of distances to the endpoints, in an earlier version of this paper we called such paths *uniform paths*.

*proper* subpaths. Indeed, in a locally shortest path, all proper subpaths are shortest paths: however, the path itself may not necessarily be shortest. We assume that every trivial path formed by a single vertex or a single edge is locally shortest. Notice that, by the optimal substructure of shortest paths, it is possible to check whether a non-trivial path  $\pi_{xy}$  is locally shortest by just verifying that  $\ell(\pi_{xy})$  and  $r(\pi_{xy})$  are shortest paths in  $G$ . In the remainder of this section we discuss some properties of locally shortest paths.

**Lemma 1** *If we denote by  $SP$  and  $LSP$  respectively the sets of shortest paths and locally shortest paths in  $G$ , then  $SP \subseteq LSP$ .*

**Proof.** Every subpath of a shortest path is a shortest path itself. Thus every shortest path is trivially a locally shortest path.  $\square$

**Lemma 2** *If shortest paths are unique in  $G$ , then for each pair of vertices  $x$  and  $y$ , the locally shortest paths connecting  $x$  and  $y$  in  $G$  are internally vertex-disjoint, i.e., except for the endpoints, they use different vertices.*

**Proof.** Suppose by contradiction that there exist two distinct locally shortest paths  $\pi_{xy}^1$  and  $\pi_{xy}^2$  that are not internally vertex-disjoint. This means that there is some vertex  $v$ , with  $v \neq x$  and  $v \neq y$ , that belongs to both  $\pi_{xy}^1$  and  $\pi_{xy}^2$ . Since shortest paths are unique, then there is only one shortest path  $\pi_{xv}$  from  $x$  to  $v$ , and only one shortest path  $\pi_{vy}$  from  $v$  to  $y$ . Since every proper subpath of  $\pi_{xy}^1$  and  $\pi_{xy}^2$  is shortest, then  $\pi_{xv}$  and  $\pi_{vy}$  are necessarily subpaths of both  $\pi_{xy}^1$  and  $\pi_{xy}^2$ . Thus,  $\pi_{xy}^1 = \pi_{xy}^2$ , which contradicts our initial assumption.  $\square$

**Lemma 3** *If shortest paths are unique in  $G$ , then there can be at most  $mn$  locally shortest paths in  $G$ .*

**Proof.** Fix an edge  $(x, v)$  and a vertex  $y$  in  $G$ . We first note that there can be at most one locally shortest path  $\pi_{xy} = \langle x, v, \dots, y \rangle$  starting from edge  $(x, v)$ . This derives from the fact that every proper subpath of  $\pi_{xy}$  must be shortest (Definition 1), and from the uniqueness of shortest paths. Since the first edge  $(x, v)$  can be chosen in  $m$  different ways and the destination vertex  $y$  can be chosen among  $n$  different vertices, at any time there can be at most  $mn$  locally shortest paths in  $G$ .  $\square$

We now study how the set of locally shortest paths changes in a graph subject to partially dynamic updates such as vertex increases, i.e., operations that increase the weights of all edges incident to it, or vertex decreases, i.e., operations that decrease the weights of all edges incident to it.

**Lemma 4** *Let  $G$  be a graph subject to a sequence  $\Sigma$  of vertex updates. If shortest paths are unique in  $G$ , then in the worst case at most  $O(n^2)$  paths can stop being locally shortest due to a vertex increase.*

**Proof.** We observe that a path can stop being locally shortest only if any of its proper subpaths stops being shortest. In case of increases, this can happen only if that subpath contains the updated vertex, say vertex  $v$ . By Lemma 2, there can be at most  $O(n^2)$  locally shortest paths that contain  $v$  as an internal vertex. Furthermore, there can be at most  $O(n^2)$  locally shortest paths starting or ending in  $v$ . This yields a total of at most  $O(n^2)$  paths that can stop being locally shortest because of a weight increase in  $v$ .  $\square$

**Theorem 1** *Let  $G$  be a graph subject to a sequence  $\Sigma$  of increase-only vertex update operations and let  $m$  be the maximum number of edges in  $G$  throughout sequence  $\Sigma$ . If shortest paths are unique in  $G$ , then the number of paths that start being locally shortest after each update is:*

- (1)  $O(mn)$  in the worst case.
- (2)  $O(n^2)$  amortized over  $\Omega(m/n)$  update operations.

**Proof.** Claim (1) follows immediately from Lemma 3. To prove Claim (2), we assign a debit to each locally shortest path in the graph. The debit is paid for by the operation that makes it stop being locally shortest. By Lemma 4, at most  $O(n^2)$  paths can stop being locally shortest at each update. Moreover, by Lemma 3 there can be at most  $mn$  locally shortest paths in a graph at any time, so the total unpaid debit at the end of the sequence never exceeds  $mn$ . Thus, the amortized number of paths that start being locally shortest after each update in any sequence of  $\Omega(m/n)$  operations can be at most  $O(n^2)$ .  $\square$

Notice that both the statements and the proofs of Lemma 4 and Theorem 1 hold symmetrically if we replace “increase” by “decrease” and swap “start” with “stop”. This can be intuitively explained by observing that, if we replay a decrease-only sequence backwards starting from the final graph, each vertex decrease in the forward sequence corresponds to a symmetric vertex increase in the backward sequence that undoes edge weights back to their previous values. In this scenario, a path starts being locally shortest during a decrease in the forward sequence if and only if it stops being locally shortest during the corresponding increase in the backward sequence, and thus the counting argument holds in both directions.

### 3 Partially Dynamic Shortest Paths

We now show how to exploit the properties of locally shortest paths discussed in Section 2 to devise an increase-only update algorithm for all pairs shortest paths that runs in  $O(n^2 \log n)$  amortized time per operation. To the best of our knowledge, this is the first general result for increase-only all pairs shortest that is faster than recomputing the solution from scratch after each update. This is rather surprising compared to the decrease-only case, where an  $O(n^2)$  bound can be immediately obtained by just running a single-source computation from the updated vertex  $v$  to every other vertex, and a single-sink computation from every vertex to  $v$ ; by doing so, for each pair of vertices  $x$  and  $y$ , we find a shortest path  $\pi_{xv}^*$  from  $x$  to  $v$ ,

and a shortest path  $\pi_{vy}^*$  from  $v$  to  $y$ : if  $\pi'_{xy} = \pi_{xv}^* \cdot \pi_{vy}^*$  is shorter than the previous shortest path  $\pi_{xy}$  from  $x$  to  $y$ , we simply replace  $\pi_{xy}$  with  $\pi'_{xy}$ .

Although the update algorithm that we describe below works also for decreases, for the sake of simplicity in this section we analyze the algorithm in the case of increases only. The approach is very simple: we maintain all the locally shortest paths of the underlying graph. By Theorem 1, changes in the data structure will be  $O(n^2)$  per update in an increase-only sequence of  $\Omega(m/n)$  operations. We will maintain locally shortest paths in priority queues, and thus we will pay  $O(\log n)$  for each path that starts/stops being locally shortest. This will yield an  $O(n^2 \log n)$  amortized time per update. Even if the combinatorial results discussed in Section 2 impose no restrictions on the edge weights (provided that shortest paths are unique) our algorithm requires that all the edge weights in the graph are non-negative.

### 3.1 Data Structure

For each pair of vertices  $x$  and  $y$  in  $G$  we maintain the weight  $w_{xy} \geq 0$  of edge  $(x, y)$  (or  $+\infty$  if no such edge exists) and the following two data structures:

$$\begin{aligned} P_{xy} &= \{ \pi_{xy} : \pi_{xy} \text{ is a locally shortest path in } G \} \\ P_{xy}^* &= \{ \pi_{xy} : \pi_{xy} \text{ is a shortest path in } G \} \end{aligned}$$

We maintain each  $P_{xy}$  as a priority queue where item  $\pi_{xy} \in P_{xy}$  has priority  $w(\pi_{xy})$ . We note that, if shortest paths are unique,  $|P_{xy}^*| \leq 1$ . Furthermore, since by Lemma 1 any shortest path is locally shortest, then for each pair of vertices  $x$  and  $y$ ,  $P_{xy}^* \subseteq P_{xy}$ . Therefore, a minimum weight path in  $P_{xy}$  is a shortest path. We also observe that each path  $\pi_{xy}$  in  $P_{xy}$  and  $P_{xy}^*$  can be represented implicitly with constant space by just storing two pointers to the subpaths  $\ell(\pi_{xy})$  and  $r(\pi_{xy})$ . Finally, for each path  $\pi_{xy}$  stored in  $P_{xy}$  we maintain  $w(\pi_{xy})$  and the following four lists:

$$\begin{aligned} L(\pi_{xy}) &= \{ \pi_{x'y} = \langle x', x \rangle \cdot \pi_{xy} : (x', x) \in E \text{ and } \pi_{x'y} \text{ is a locally shortest path in } G \} \\ L^*(\pi_{xy}) &= \{ \pi_{x'y} = \langle x', x \rangle \cdot \pi_{xy} : (x', x) \in E \text{ and } \pi_{x'y} \text{ is a shortest path in } G \} \\ R(\pi_{xy}) &= \{ \pi_{xy'} = \pi_{xy} \cdot \langle y, y' \rangle : (y, y') \in E \text{ and } \pi_{xy'} \text{ is a locally shortest path in } G \} \\ R^*(\pi_{xy}) &= \{ \pi_{xy'} = \pi_{xy} \cdot \langle y, y' \rangle : (y, y') \in E \text{ and } \pi_{xy'} \text{ is a shortest path in } G \} \end{aligned}$$

In other words,  $L(\pi_{xy})$  and  $L^*(\pi_{xy})$  represent pre-extensions of  $\pi_{xy}$ , while  $R(\pi_{xy})$  and  $R^*(\pi_{xy})$  represent post-extensions of  $\pi_{xy}$ . Once again, by Lemma 1 any shortest path is locally shortest, and thus for each path  $\pi_{xy}$  stored in  $P_{xy}$ ,  $L^*(\pi_{xy}) \subseteq L(\pi_{xy})$  and  $R^*(\pi_{xy}) \subseteq R(\pi_{xy})$ . For the sake of simplicity, in the following we will sometimes write  $P$  or  $P^*$  instead of  $P_{xy}$  or  $P_{xy}^*$  whenever the meaning is clear from the context. The notation introduced in this section is summarized in Table 2.

---

<b>distance</b> ( $x, y$ ):
1. <b>if</b> $P_{xy} = \emptyset$ <b>then return</b> $+\infty$
2. <b>else return</b> the weight of a minimum weight path in $P_{xy}$

---

<b>path</b> ( $x, y$ ):
1. <b>if</b> $P_{xy} = \emptyset$ <b>then return</b> $\emptyset$
2. <b>else return</b> a minimum weight path in $P_{xy}$

---

Figure 1: Implementation of `distance` and `path` operations.

## 3.2 Implementation of Operations

The `distance`( $x, y$ ) and `path`( $x, y$ ) operations can be implemented as shown in Figure 1, by simply accessing the minimum weight path in  $P_{xy}$ . Since each  $P_{xy}$  is a subset of the paths in  $G$ , the correctness of query operations follows directly from Lemma 1.

The implementation of the `update`( $v, w'$ ) operation is shown in Figure 2. The update works in two steps: `cleanup` and `fixup`. To simplify the description, we say that a path that is shortest (resp., locally shortest) after the update is *new* either if it was not shortest (resp., locally shortest) before the update, or if it contains the updated vertex  $v$ . We now describe procedures `cleanup` and `fixup` in more detail; pseudo-code is given in Figure 2.

---

### `cleanup`( $v$ )

The procedure removes every path  $\pi_{xy}$  containing vertex  $v$  from  $P_{xy}, P_{xy}^*, L(r(\pi_{xy})), L^*(r(\pi_{xy})), R(\ell(\pi_{xy})),$  and  $R^*(\ell(\pi_{xy}))$ . Namely, we remove from the data structure all the paths that would stop being locally shortest if we deleted  $v$  from the graph. This task can be accomplished iteratively by first removing paths of the form  $\langle u, v \rangle$  and  $\langle v, u \rangle$ , and then by removing all paths listed in  $L(\pi_{xy})$  and  $R(\pi_{xy})$  for each path  $\pi_{xy}$  removed in the previous iterations.

---

### `fixup`( $v, w'$ )

The procedure adds to the data structure all the new shortest and locally shortest paths. It works in three phases, as follows.

*Phase 1:* Sets the weight of every edge  $(u, v)$  entering  $v$  to the new value  $w'_{uv}$  (line 2). If

$P_{xy}$	set of locally shortest paths from $x$ to $y$ in $G$
$P_{xy}^*$	set of shortest paths from $x$ to $y$ in $G$
$L(\pi_{xy})$	set of pre-extensions $\langle x', x \rangle \cdot \pi_{xy}$ of $\pi_{xy}$ that are locally shortest paths in $G$
$L^*(\pi_{xy})$	set of pre-extensions $\langle x', x \rangle \cdot \pi_{xy}$ of $\pi_{xy}$ that are shortest paths in $G$
$R(\pi_{xy})$	set of post-extensions $\pi_{xy} \cdot \langle y, y' \rangle$ of $\pi_{xy}$ that are locally shortest paths in $G$
$R^*(\pi_{xy})$	set of post-extensions $\pi_{xy} \cdot \langle y, y' \rangle$ of $\pi_{xy}$ that are shortest paths in $G$

Table 2: Notation introduced in Section 3.1.



---

```

update( $v, w'$ ):
1.   cleanup( $v$ )
2.   fixup( $v, w'$ )

```

---

```

cleanup( $v$ ):
1.    $Q \leftarrow \{\langle v \rangle\}$ 
2.   while  $Q \neq \emptyset$  do
3.     extract any  $\pi$  from  $Q$ 
4.     for each  $\pi_{xy} \in L(\pi) \cup R(\pi)$  do
5.       add  $\pi_{xy}$  to  $Q$ 
6.       remove  $\pi_{xy}$  from  $P_{xy}$ ,  $L(r(\pi_{xy}))$ , and  $R(\ell(\pi_{xy}))$ 
7.       if  $\pi_{xy} \in P_{xy}^*$  then remove  $\pi_{xy}$  from  $P_{xy}^*$ ,  $L^*(r(\pi_{xy}))$ , and  $R^*(\ell(\pi_{xy}))$ 

```

---

```

fixup( $v, w'$ ):
1.   for each  $u \neq v$  do {Phase 1}
2.      $w_{uv} \leftarrow w'_{uv}$ ;  $w_{vu} \leftarrow w'_{vu}$ 
3.     if  $w_{uv} < +\infty$  then
4.        $w(\langle u, v \rangle) \leftarrow w_{uv}$ ;  $\ell(\langle u, v \rangle) \leftarrow \langle u \rangle$ ;  $r(\langle u, v \rangle) \leftarrow \langle v \rangle$ 
5.       add  $\langle u, v \rangle$  to  $P_{uv}$ ,  $L(\langle v \rangle)$ , and  $R(\langle u \rangle)$ 
6.     if  $w_{vu} < +\infty$  then
7.        $w(\langle v, u \rangle) \leftarrow w_{vu}$ ;  $\ell(\langle v, u \rangle) \leftarrow \langle v \rangle$ ;  $r(\langle v, u \rangle) \leftarrow \langle u \rangle$ 
8.       add  $\langle v, u \rangle$  to  $P_{vu}$ ,  $L(\langle u \rangle)$ , and  $R(\langle v \rangle)$ 
9.    $H \leftarrow \emptyset$  {Phase 2}
10.  for each  $(x, y)$  do
11.    add  $\pi_{xy} \in P_{xy}$  with minimum  $w(\pi_{xy})$  to  $H$ 
12.  while  $H \neq \emptyset$  do {Phase 3}
13.    extract  $\pi_{xy}$  from  $H$  with minimum  $w(\pi_{xy})$ 
14.    if  $\pi_{xy}$  is the first extracted path for pair  $(x, y)$  then
15.      if  $\pi_{xy} \notin P_{xy}^*$  then
16.        add  $\pi_{xy}$  to  $P_{xy}^*$ ,  $L^*(r(\pi_{xy}))$ , and  $R^*(\ell(\pi_{xy}))$ 
17.        for each  $\pi_{x'b} \in L^*(\ell(\pi_{xy}))$  do
18.           $\pi_{x'y} \leftarrow \langle x', x \rangle \cdot \pi_{xy}$ 
19.           $w(\pi_{x'y}) \leftarrow w_{x'x} + w(\pi_{xy})$ ;  $\ell(\pi_{x'y}) \leftarrow \pi_{x'b}$ ;  $r(\pi_{x'y}) \leftarrow \pi_{xy}$ 
20.          add  $\pi_{x'y}$  to  $P_{x'y}$ ,  $L(\pi_{xy})$ ,  $R(\pi_{x'b})$ , and  $H$ 
21.        for each  $\pi_{ay'} \in R^*(r(\pi_{xy}))$  do
22.           $\pi_{xy'} \leftarrow \pi_{xy} \cdot \langle y, y' \rangle$ 
23.           $w(\pi_{xy'}) \leftarrow w(\pi_{xy}) + w_{yy'}$ ;  $\ell(\pi_{xy'}) \leftarrow \pi_{xy}$ ;  $r(\pi_{xy'}) \leftarrow \pi_{ay'}$ 
24.          add  $\pi_{xy'}$  to  $P_{xy'}$ ,  $L(\pi_{ay'})$ ,  $R(\pi_{xy})$ , and  $H$ 

```

---

Figure 2: Implementation of the update operation.

$w'_{uv} < +\infty$ , then the trivial path  $\pi_{uv} = \langle u, v \rangle$  is added to  $P_{uv}$ ,  $L(r(\pi_{uv}))$ , and  $R(\ell(\pi_{uv}))$  (lines 3–5). Similar steps are performed for every edge of the form  $(v, u)$  (line 2 and lines 6–8). Thus, all the new locally shortest paths formed by one edge are added to the data structure. Longer new paths will be added in *Phase 3*.

*Phase 2*: Initializes a priority queue  $H$  with the minimum weight path  $\pi_{xy} \in P_{xy}$  for each pair of vertices  $(x, y)$  (lines 9–11);

*Phase 3*: Repeatedly extracts paths  $\pi_{xy}$  from  $H$  in increasing weight order (line 13). The first extracted path for each pair  $(x, y)$  is a shortest path between  $x$  and  $y$  (see Invariant 1 below), while paths extracted later on for the same pair are ignored (line 14). Whenever a shortest path  $\pi_{xy}$  is extracted, we check whether  $\pi_{xy}$  is already in  $P_{xy}^*$  (line 15). If not, we add  $\pi_{xy}$  to  $P_{xy}^*$ ,  $L^*(r(\pi_{xy}))$ , and  $R^*(\ell(\pi_{xy}))$  (line 16), and we combine it with existing shortest paths to form new locally shortest paths (lines 17–24). This is done by scanning all paths  $\pi_{x'b}$  listed in  $L^*(\ell(\pi_{xy}))$  (line 17) and all paths  $\pi_{ay'}$  listed in  $R^*(r(\pi_{xy}))$  (line 21), forming the new locally shortest paths  $\pi_{x'y} = \langle x', x \rangle \cdot \pi_{xy}$  (lines 18–19) and  $\pi_{xy'} = \pi_{xy} \cdot \langle y, y' \rangle$  (lines 22–23). Each new locally shortest path  $\pi_{ij}$  is added to  $P_{ij}$ ,  $L(r(\pi_{ij}))$ ,  $R(\ell(\pi_{ij}))$ , and  $H$  as soon as it is discovered (line 20 and line 24).

### 3.3 Analysis

To prove the correctness of **update**, we assume that  $P$  and  $P^*$  are correct before the operation, and we show that they are also correct afterwards. Operations on the other data structures are simple bookkeeping operations and their correctness can be easily checked. We first discuss an invariant maintained by procedure **fixup**.

**Invariant 1** *If shortest paths are unique and edge weights are non-negative, then for each pair of vertices  $x$  and  $y$  in  $G$ , the first path connecting  $x$  and  $y$  extracted from  $H$  in Phase 3 of **fixup** is a shortest path.*

**Proof.** Suppose by contradiction that at some extraction the invariant is violated, and the first path  $\hat{\pi}_{xy}$  extracted for some pair  $(x, y)$  is not a shortest path. Consider the earliest of these events, and let  $\pi_{xy}$  be the unique shortest path between  $x$  and  $y$ , with  $w(\pi_{xy}) < w(\hat{\pi}_{xy})$ . Clearly,  $\pi_{xy} \notin H$ , otherwise it would have been extracted in place of  $\hat{\pi}_{xy}$ . Moreover,  $\pi_{xy} \notin P_{xy}$ , otherwise it would have been inserted in  $H$  in *Phase 2*, being a shortest path. Therefore,  $\pi_{xy}$  has to be necessarily a new locally shortest path, but since in *Phase 1* we add to  $P$  all the edges incident to  $v$ , it cannot be one of them. This implies that  $\pi_{xy}$  contains at least two edges and either one of  $\ell(\pi_{xy})$  or  $r(\pi_{xy})$  is a new shortest path and was not in  $P^*$  at the beginning of **fixup**. As edge weights are non-negative, then  $w(\ell(\pi_{xy})) \leq w(\pi_{xy}) < w(\hat{\pi}_{xy})$  and  $w(r(\pi_{xy})) \leq w(\pi_{xy}) < w(\hat{\pi}_{xy})$ . Since paths are extracted from  $H$  in increasing weight order, shortest paths are unique, and all the extractions before the wrong extraction are correct, then either  $\ell(\pi_{xy})$  or  $r(\pi_{xy})$  should have been extracted from  $H$  and added to  $P^*$  before the wrong extraction. This implies that  $\pi_{xy}$  should have been formed and inserted in  $H$  at the extraction time of either one of  $\ell(\pi_{xy})$  or  $r(\pi_{xy})$ . So  $\pi_{xy}$  should have been extracted before  $\hat{\pi}_{xy}$ , contradicting the assumption that the invariant is violated.  $\square$

**Theorem 2** *If the operation is an increase, shortest paths are unique, and edge weights are non-negative, then algorithm `update` correctly updates  $P_{xy}$  and  $P_{xy}^*$  for each pair of vertices  $x$  and  $y$ .*

**Proof.** We first observe that `cleanup`( $v$ ) removes from the data structures every locally shortest path that contains the updated vertex  $v$ . Thus, since only paths containing  $v$  can stop being shortest after an increase, then every path  $\pi_{xy}$  that is no longer locally shortest after the update is removed from  $P_{xy}$ , and possibly from  $P_{xy}^*$ .

We now show that every new locally shortest path  $\pi_{xy}$  is added to  $P_{xy}$  by `fixup`( $v, w'$ ). Paths  $\pi_{uv}$  and  $\pi_{vu}$  made of one edge are correctly added to  $P_{uv}$  and  $P_{vu}$  in *Phase 1*: therefore, we only focus on paths  $\pi_{xy}$  with at least two edges. We recall that each such path  $\pi_{xy}$  has to be added to  $P_{xy}$  if both  $\ell(\pi_{xy})$  and  $r(\pi_{xy})$  are shortest paths after the update, but at least one of them (say  $\ell(\pi_{xy})$ ) was not in  $P^*$  before `fixup`. By Invariant 1, the first extracted path for each pair is shortest, and then  $\ell(\pi_{xy})$  is certainly extracted and added to  $P^*$  at some iteration in *Phase 3*. At that time,  $\ell(\pi_{xy})$  is correctly combined with  $r(\pi_{xy}) \in R^*(r(\ell(\pi_{xy})))$  to form  $\pi_{xy}$ , which is added to  $P_{xy}$ .

To conclude the proof, we observe that by Invariant 1 all shortest paths are extracted at some iteration and are added to  $P^*$ , if they are not already there.  $\square$

To conclude the analysis, we address the running time and the space usage of our implementation.

**Theorem 3** *In an increase-only sequence of  $\Omega(m/n)$  operations, if shortest paths are unique and edge weights are non-negative, our data structure supports each `update` operation in  $O(n^2 \log n)$  amortized time, and each `distance` and `path` query in optimal time. The space used is  $O(mn)$ .*

**Proof.** The bounds for queries are immediate, since the minimum weight path in the priority queue  $P_{xy}$  can be retrieved in optimal time.

Since each iteration of `cleanup` removes a path  $\pi_{xy}$  from a constant number of lists and priority queues, it requires  $O(\log n)$  time in the worst case. By Lemma 2, at most  $O(n^2)$  locally shortest paths can go through the updated vertex in the worst case, and thus `cleanup` will perform at most  $O(n^2)$  iterations, thus spending at most  $O(n^2 \log n)$  time.

Next, it is easy to see that *Phase 1* of `fixup` requires  $O(n \log n)$  time in the worst case. To bound the time required by *Phase 2* of `fixup`, we consider that finding the minimum in  $P_{xy}$  takes constant time and each insertion in  $H$  takes  $O(\log n)$  time. This gives a total bound of  $O(n^2 \log n)$  for this phase. Finally, we note that a path  $\pi_{xy}$  is a new locally shortest path if and only if both  $\ell(\pi_{xy}) = \langle x, \dots, b \rangle \in P_{xb}^*$  and  $r(\pi_{xy}) = \langle a, \dots, y \rangle \in P_{ay}^*$  after the update, and either  $\ell(\pi_{xy}) \notin P_{xb}^*$  or  $r(\pi_{xy}) \notin P_{ay}^*$  before calling `fixup`. Thus, since in *Phase 3* we process a shortest path  $\pi_{xy}$  extracted from  $H$  only if it was not already in  $P_{xy}^*$  before calling `fixup`, then we spend time only for the new locally shortest paths. This implies that *Phase 3* requires  $O(n^2 \log n)$  time for extracting the  $O(n^2)$  pairs which are initially in  $H$ , plus  $O(\log n)$  time to insert/delete each new locally shortest path in  $H$  and  $P$ : let  $p$  the number of such paths. By Theorem 1, there can be  $O(n^2)$  paths that start being locally shortest after




Figure 3: If shortest paths are not unique, using an arbitrary tie-breaking strategy may lead to pairs of vertices connected in  $G$ , but disconnected in  $P^*$ . In this example, if  $P^*$  contains  $\langle x, a, c \rangle$ ,  $\langle x, b, d \rangle$ ,  $\langle a, d, y \rangle$ , and  $\langle b, c, y \rangle$  as the unique shortest paths of length 2, then no locally shortest path between  $x$  and  $y$  can be obtained as the union of any two of them.

each update, amortized over an increase-only sequence of  $\Omega(m/n)$  operations. Moreover, by Lemma 2, there can be at most  $O(n^2)$  locally shortest paths that were removed by `cleanup` since they contain the updated vertex, and are added back by `fixup`: thus,  $p = O(n^2)$  and `update` requires  $O(n^2 \log n)$  overall time.

At any time, each locally shortest path  $\pi_{xy}$  can only appear in a constant number of data structures: namely, in  $H$ ,  $P_{xy}$ ,  $L(r(\pi_{xy}))$ ,  $R(\ell(\pi_{xy}))$ , and possibly in  $P^*$ ,  $L^*(r(\pi_{xy}))$ ,  $R^*(\ell(\pi_{xy}))$ . Thus, the space usage of our data structure is bounded by the number of locally shortest paths in a graph, which is  $O(mn)$  in the worst case by Lemma 3.  $\square$

By Theorem 2, algorithm `update` maintains correctly the locally shortest paths of the graph, and thus the shortest paths, under increase-only update sequences. As we will see in Section 5, the above algorithm is indeed correct also in case of decreases, even if it does not maintain  $P$  and  $P^*$  as they were defined in Section 3.1: specifically, in case of decreases, paths that stop being shortest may not be removed from the data structure by `cleanup`, and thus  $P^*$  might contain paths that were shortest in the past, even if they are no longer shortest. In Section 4 we will call such “old” paths *historical* and show that they can play a crucial role in dealing with fully dynamic sequences.

### 3.4 Breaking Ties

If shortest paths in the graph are not unique, the properties of locally shortest paths studied in Section 2 do not hold. Furthermore, using an arbitrary tie-breaking strategy may lead to incorrect results. Consider the example of Figure 3, where all edges have weight 1, and suppose that at some point the unique shortest paths of length 2 in  $P^*$  are:  $\langle x, a, c \rangle$ ,  $\langle x, b, d \rangle$ ,  $\langle a, d, y \rangle$ , and  $\langle b, c, y \rangle$ . Notice that we can find no path of length 3 such that every subpath of length 2 is in  $P^*$ : therefore, no path of length 3 can be formed by `update`, leaving the pair  $(x, y)$  incorrectly disconnected in  $P^*$ . In order `update` to be correct, we thus need a tie-breaking strategy that both ensures uniqueness and leads to a complete set of shortest paths closed under subpaths.

How can we break ties consistently? A first simple approach is to add some tiny random fraction to the weight of each edge: in this way, we can make the probability of finding two paths with the same weight arbitrarily small.

A more robust deterministic tie-breaking approach is the following. Without loss of generality, assume that  $V = \{1, 2, \dots, n\}$ . We assign to each edge  $(u, v)$  a unique number  $ID(u, v)$  (e.g.,  $ID(u, v) = u + nv$ ), and for each path  $\pi$  we define  $ID(\pi)$  as the maximum  $ID$  of its edges. This allows us to define the *extended weight*  $ew(\pi)$  of a path  $\pi$  as follows.

**Definition 2** For any path  $\pi$  with at least one edge we define the extended weight of  $\pi$  as  $ew(\pi) = \langle w(\pi), ID(\pi) \rangle$ , where  $w(\pi)$  is the real weight of  $\pi$  in the graph and  $ID(\pi)$  is defined as follows:

$$ID(\pi) = \begin{cases} u + nv & \text{if } \pi = \langle u, v \rangle \\ \max \{ ID(\ell(\pi)), ID(r(\pi)) \} & \text{otherwise} \end{cases}$$

The extended weights of two paths  $\pi_1$  and  $\pi_2$  can be compared “lexicographically” in constant time by assuming that:

$$ew(\pi_1) \leq ew(\pi_2) \Leftrightarrow w(\pi_1) < w(\pi_2) \text{ or} \\ w(\pi_1) = w(\pi_2) \text{ and } ID(\pi_1) \leq ID(\pi_2)$$

We now show that using extended weights we can break ties consistently so that exactly one shortest path (if any) between each pair of vertices can be identified in a deterministic way. To do so, we define the following set of paths in a graph:

**Definition 3** Let  $G$  be a graph with real-valued edge weights and let  $ew$  be the extended weight function of Definition 2. We define  $S_{ew}$  as follows:

$$S_{ew} = \{ \pi \in G : \forall \pi_{xy} \subseteq \pi, \forall \pi'_{xy} \in G, ew(\pi_{xy}) \leq ew(\pi'_{xy}) \}.$$

According to this definition,  $S_{ew}$  contains every path that has minimum extended weight and whose subpaths have minimum extended weight as well (optimal-substructure property). As a consequence,  $S_{ew}$  is closed under subpaths, i.e., every subpath of a path in  $S_{ew}$  belongs to  $S_{ew}$ . It is easy to prove that there can be only one path in  $S_{ew}$  between each pair of vertices.

**Lemma 5** For each pair of vertices  $x$  and  $y$  in  $G$ , there can be at most one path  $\pi_{xy} \in S_{ew}$  connecting them.

**Proof.** The proof is by induction on the number of edges in the paths. The basis is trivially satisfied, since there can be only one path formed by one edge between each pair of vertices. Now, assume by induction that the claim holds for every path that uses at most  $k$  edges, and suppose by contradiction that there are two different paths  $\pi_{xy} \in S_{ew}$  and  $\pi'_{xy} \in S_{ew}$  that use at most  $k + 1$  edges. Since their extended weights are equal, then both  $\pi_{xy}$  and  $\pi'_{xy}$  must contain the edge  $(u, v)$  such that  $ID(u, v) = ID(\pi_{xy}) = ID(\pi'_{xy})$ . Since every subpath of  $\pi_{xy}$  and  $\pi'_{xy}$  belongs to  $S_{ew}$  and uses at most  $k$  edges, then  $\pi_{xy}$  and  $\pi'_{xy}$  must have the

same subpaths from  $x$  to  $u$  and from  $v$  to  $y$ . Thus,  $\pi_{xy}$  and  $\pi'_{xy}$  must be equal, contradicting the assumption that they are different.  $\square$

Clearly, if  $\pi_{xy} \in S_{\text{ew}}$ , then  $\pi_{xy}$  is a shortest path in  $G$ . We now show that, if there is a shortest path from  $x$  to  $y$  in  $G$ , then there must be a path  $\pi_{xy} \in S_{\text{ew}}$ .

**Lemma 6** *For each pair of connected vertices  $x$  and  $y$  in  $G$ , there is a path  $\pi_{xy} \in S_{\text{ew}}$ .*

**Proof.** Assume that  $y$  is reachable from  $x$  in  $G$ ,  $x \neq y$ . Then, we show how to construct a path  $\pi_{xy}^*$  that satisfies Definition 3. We start with an empty  $\pi_{xy}^*$  and construct it one edge at the time. Let  $\pi_{xy}$  be a path from  $x$  to  $y$  with minimum extended weight in  $G$  and let  $(u, v)$  be the edge with maximum  $ID$  in  $\pi_{xy}$ , i.e., such that  $ID(u, v) = ID(\pi_{xy})$ . We put  $(u, v)$  in  $\pi_{xy}^*$  and build recursively the subpath  $\pi_{xu}^*$  of  $\pi_{xy}^*$  if  $x \neq u$  and the subpath  $\pi_{vy}^*$  of  $\pi_{xy}^*$  if  $v \neq y$ .

We now show that the above procedure constructs paths that belong to  $S_{\text{ew}}$ . We prove this claim proceeding by induction on the path length. Let  $k$  be the number of edges in  $\pi_{xy}^*$ . The base for  $k = 1$  ( $\pi_{xy}^* = \langle x, y \rangle$ ) is trivial. Assume by induction that the claim holds for every path containing less than  $k$  edges and assume by contradiction that there is a subpath  $\pi_{ab} \in \pi_{xy}^*$  that is not of minimum extended weight. Three cases are possible:

- (1) both vertices  $a$  and  $b$  are in  $\pi_{xu}^*$ ;
- (2) both vertices  $a$  and  $b$  are in  $\pi_{vy}^*$ ;
- (3) vertex  $a$  is in  $\pi_{xu}^*$  and vertex  $b$  is in  $\pi_{vy}^*$ .

Case (1) and (2) are not possible, since  $\pi_{xu}^*$  and  $\pi_{vy}^*$  have strictly less than  $k$  edges, and thus by induction are in  $S_{\text{ew}}$ . Assume that we are in case (3). Since  $\pi_{ab}$  is not of minimum extended weight, there must be a path  $\pi'_{ab}$  such that  $\text{ew}(\pi'_{ab}) < \text{ew}(\pi_{ab})$ . Note that, by construction  $\pi_{xy}^*$  contains only edges that lie on some shortest path from  $x$  to  $y$  in  $G$ , and thus  $w(\pi_{ab}) \leq w(\pi'_{ab})$ . Consequently, it must be  $w(\pi_{ab}) = w(\pi'_{ab})$  and  $ID(\pi'_{ab}) < ID(\pi_{ab})$ . This implies that there is a shortest path from  $x$  to  $y$  avoiding edge  $(u, v)$ , with an ID smaller than  $ID(u, v)$ . This contradicts the fact that  $ID(u, v)$  is the ID of a path with minimum extended weight from  $x$  to  $y$ .  $\square$

By Lemma 5 and Lemma 6,  $S_{\text{ew}}$  contains exactly one representative shortest path between each pair of vertices connected in  $G$ . Thus, we can conclude that  $|S_{\text{ew}}| \leq n^2$  and  $S_{\text{ew}}$  is a complete set of unique shortest paths in  $G$ . For this reason, in the remainder of this paper we will assume that we can always ensure uniqueness of shortest paths by dealing with paths in  $S_{\text{ew}}$ .

We now show that using the deterministic approach described above we can break ties consistently and ensure both completeness and uniqueness of shortest paths, avoiding inconsistencies like the one shown in Figure 3. As a matter of fact, we can easily modify the code of the `update` operation given in Figure 2 so that at any time during a sequence of updates,  $P^* = S_{\text{ew}}$ . Specifically, we keep an additional field  $ID(\pi)$  for each path  $\pi$  stored in the data structure, and perform the following additional steps in `fixup` to maintain it as in Definition 2:

$$\left| \begin{array}{l} \text{Line 4} \\ \text{Line 7} \end{array} \left| \begin{array}{l} ID(\langle u, v \rangle) \leftarrow u + nv \\ ID(\langle v, u \rangle) \leftarrow v + nu \end{array} \right| \right| \quad \left| \begin{array}{l} \text{Line 19} \\ \text{Line 23} \end{array} \left| \begin{array}{l} ID(\pi_{x'y}) \leftarrow \max\{ ID(\pi_{x'b}), ID(\pi_{xy}) \} \\ ID(\pi_{xy'}) \leftarrow \max\{ ID(\pi_{xy}), ID(\pi_{ay'}) \} \end{array} \right| \right|$$

Furthermore, we let the priority of any path  $\pi$  in  $P$  and  $H$  be  $\text{ew}(\pi)$ . Thus, in line 11 and line 13 of `fixup` we pick paths with minimum  $\text{ew}$ , instead of minimum  $w$ . To check the correctness of this approach, we prove the following theorem.

**Theorem 4** *If we compare paths according to the extended weight function  $\text{ew}$  given in Definition 2, then  $P^* = S_{\text{ew}}$  after each execution of `update`.*

**Proof.** We first observe that for any path  $\pi$ , if edge weights in the graph are non-negative reals, the extended weight function  $\text{ew}$  given in Definition 2 satisfies the monotonic property  $\text{ew}(\ell(\pi)) \leq \text{ew}(\pi)$  and  $\text{ew}(r(\pi)) \leq \text{ew}(\pi)$ . Moreover, by Lemma 5 and Lemma 6 there is exactly one path in  $S_{\text{ew}}$  between each pair of connected vertices. Therefore, it is easy to check that Invariant 1 continues to hold if we compare paths according to  $\text{ew}$  instead of  $w$ : this implies that each path added to  $P^*$  has minimum extended weight.

Furthermore, by construction, paths in  $P^*$  are closed under subpaths, i.e., every subpath of a path in  $P^*$  is also in  $P^*$ : thus,  $P^* = S_{\text{ew}}$  as in Definition 3.  $\square$

We conclude this section by observing that, since IDs of edges never change, then each increase-only sequence on the weights of  $G$  is also increase-only working with the extended weights. Furthermore, we notice that comparing extended weights still takes constant time. Thus, both Theorem 2 and Theorem 3 continue to hold using extended weights instead of the original weights in the graph.

## 4 Historical and Locally Historical Paths

Figure 4 shows a graph and a pathological fully dynamic update sequence such that each operation causes  $\Theta(n^3)$  changes in the set of locally shortest paths of the graph. This proves that maintaining locally shortest paths in a fully dynamic setting can require as much as  $\Omega(n^3)$  time per operation. To deal with fully dynamic update sequences efficiently, in this section we introduce other two classes of paths that we call *historical paths* and *locally historical paths*: historical paths will play the role of shortest paths, while locally historical paths will play the role of locally shortest paths. Differently from shortest paths and locally shortest paths, which only depend on the graph at a given time, these classes of paths capture the history of the changes in the graph, encompassing the time dimension of the problem. We define them as follows:

**Definition 4** *Let  $\pi_{xy}$  be a path in  $G$  at time  $t$ , and let  $t' \leq t$  be the time of the latest vertex update on  $\pi_{xy}$ . We say that  $\pi_{xy}$  is **HISTORICAL** at time  $t$  if it has been a shortest path at least once in the time interval  $[t', t]$ .*

According to this definition, a shortest path is also a historical path. Although it may stop being a shortest path at some point, it keeps on being a historical path until a vertex update occurs on it.




Figure 4: Pathological update sequence consisting of repeated deletion and re-insertion of edge  $(u, v)$ . Each insertion/deletion causes  $\Theta(n^3)$  changes in the set of locally shortest paths. Locally shortest paths connecting vertices in layer  $L_1$  to vertices in layer  $L_4$  are highlighted as thick lines. Edges are directed from left to right and, unless explicitly shown, their weights are small values close to zero.

**Definition 5** We say that a path  $\pi_{xy}$  is **LOCALLY HISTORICAL**<sup>2</sup> in  $G$  at time  $t$  if either:

- (i)  $\pi_{xy}$  consists of a single vertex, or
- (ii) every proper subpath of  $\pi_{xy}$  is a historical path in  $G$  at time  $t$ .

Differently from Definition 1, proper subpaths do not have to be shortest paths necessarily at the same time. Indeed, in a locally historical path, all proper subpaths have been shortest paths in the past; however, the path itself may have never been locally shortest.

The next lemma addresses the relationship between shortest paths, historical paths, and locally historical paths:

**Lemma 7** If we denote by  $SP$ ,  $HP$ , and  $LHP$  respectively the sets of shortest paths, historical paths, and locally historical paths in  $G$  at any time in a sequence of updates, then the following inclusions hold at that time:  $SP \subseteq HP \subseteq LHP$ .

**Proof.** The inclusion  $SP \subseteq HP$  is straightforward from Definition 4. Let  $\pi$  be a historical path at time  $t$  and let  $t' \leq t$  be the time of the latest vertex update on  $\pi$ . For every proper subpath  $\hat{\pi} \subset \pi$ , let  $\hat{t} \leq t' \leq t$  be the time of the latest vertex update on  $\hat{\pi}$ . By Definition 4,  $\pi$  has been a shortest path in the time interval  $[t', t]$ . By the optimal-substructure property of shortest paths,  $\hat{\pi}$  has been a shortest path in the time interval  $[\hat{t}, t] \supseteq [t', t]$ , and thus it is historical at time  $t$ . Consequently,  $HP \subseteq LHP$ .  $\square$

<sup>2</sup>In an earlier version of this paper, we used terminology *potentially uniform paths* instead of *locally historical paths*.






Figure 5: Inclusions between shortest paths (SP), locally shortest paths (LSP), historical paths (HP), and locally historical paths (LHP).

Figure 5 summarizes the relationship between shortest paths, locally shortest paths, historical paths, and locally historical paths. As we will see later on, the reason why we are considering locally historical paths is that they can be maintained very efficiently in a sequence of fully dynamic updates. Since by Lemma 7 locally historical paths include shortest paths, this implies that we can maintain information about dynamic shortest paths very efficiently. Historical paths, on the other hand, are only used as a tool for defining locally historical paths.

#### 4.1 Properties of Locally Historical Paths

In this section we study combinatorial properties of locally historical paths, addressing their dependence on the number of historical paths in the graph. These properties will be crucial in Section 5 to designing our efficient algorithm for fully dynamic shortest paths. We start by analyzing the maximum number of paths that can be locally historical at each time during a sequence of updates.

**Lemma 8** *Let  $G$  be a graph subject to a sequence  $\Sigma$  of update operations. If at any time throughout the sequence  $\Sigma$  there are at most  $z$  historical paths between each pair of vertices and  $m$  edges in  $G$ , then at that time there can be at most  $zmn$  locally historical paths in  $G$ .*

**Proof.** We follow the same argument given in the proof of Lemma 3. To build a locally historical path  $\pi_{xy} = \langle x, u, \dots, y \rangle$ , the first edge  $(x, u)$  can be chosen in  $m$  different ways, while the destination vertex  $y$  can be chosen among  $n$  different vertices. Since every proper subpath of  $\pi_{xy}$  has to be historical and there are at most  $z$  historical paths between each pair of vertices, there can be at most  $z$  historical subpaths between  $u$  and  $y$ . Thus there can be at most  $zmn$  locally historical paths in  $G$ .  $\square$

Another interesting question is how many locally historical paths can go through any given vertex in the graph. Since a path stops being locally historical only when it is touched by a vertex update, this yields a worst-case bound on the number of paths that can stop being locally historical at each update. To address this, we first need to prove that, given

any two distinct historical paths between any pair of vertices  $x$  and  $y$ , their latest updated vertices must be different.

**Lemma 9** *Let  $G$  be a graph subject to a sequence  $\Sigma$  of update operations and let  $x, y$  be any two vertices in  $G$ . If shortest paths are unique and we denote by  $\sigma(\pi)$  the latest vertex update on path  $\pi$ , then for any two distinct historical paths  $\pi_{xy}$  and  $\pi'_{xy}$  in  $G$  at any time  $t$ ,  $v_{\sigma(\pi_{xy})} \neq v_{\sigma(\pi'_{xy})}$ .*

**Proof.** Assume by contradiction that  $v_{\sigma(\pi_{xy})} = v_{\sigma(\pi'_{xy})}$ , and thus  $t_{\sigma(\pi_{xy})} = t_{\sigma(\pi'_{xy})}$ . By Definition 4, both  $\pi_{xy}$  and  $\pi'_{xy}$  have to be shortest paths at least once in the time interval  $[t_{\sigma(\pi_{xy})}, t]$ , but this is impossible since neither of them is touched in the time interval  $(t_{\sigma(\pi_{xy})}, t]$ , and at each time there can be only one shortest path between  $x$  and  $y$  in  $G$ .  $\square$

**Lemma 10** *Let  $G$  be a graph subject to a sequence  $\Sigma$  of update operations and let  $v$  be any vertex in  $G$ . If shortest paths are unique and at any time  $t$  throughout the sequence  $\Sigma$  there are at most  $z$  historical paths between each pair of vertices, then at time  $t$  there can be at most  $O(zn^2)$  locally historical paths that contain  $v$ .*

**Proof.** We first show that there can be at most  $O(zn^2)$  locally historical paths having  $v$  as endpoint. Indeed, to form a locally historical path  $\pi_{vy} = \langle v, x, \dots, y \rangle$ , we can choose  $x$  and  $y$  in at most  $O(n^2)$  different ways; since every proper subpath of  $\pi_{vy}$  has to be historical and there are at most  $z$  historical paths between each pair of vertices, there can be at most  $z$  historical subpaths between  $x$  and  $y$ . The argument for paths ending in  $v$  is completely analogous.

To complete the proof, we now show that for each pair of vertices  $x, y$ , with  $x \neq v$  and  $y \neq v$ , there can be at most  $O(z)$  locally historical paths of the form  $\langle x, \dots, v, \dots, y \rangle$ , i.e., having  $v$  as an internal vertex. Denote by  $P$  the set of such paths, and let  $Q = \{\pi_1, \dots, \pi_q\}$  be the subset of  $P$  such that for each  $\pi_i \in Q$ ,  $v$  follows  $v_i$  in  $\pi_i$ , where  $v_i$  is the latest updated vertex in  $\pi_i$ . Now, suppose by contradiction that  $|P| > 2z$ , and assume without loss of generality that  $q = |Q| \geq |P|/2$ , and thus  $q > z$ . Consider the set of subpaths  $\widehat{Q} = \{\widehat{\pi}_i = \langle x, \dots, v_i, \dots, v \rangle \subset \pi_i, \pi_i \in Q\}$ . By Definition 5, paths in  $\widehat{Q}$  are historical at time  $t$ . Furthermore, since for each  $i \in [1, q]$ ,  $v_i$  is the latest updated vertex in  $\pi_i$ , then  $v_i$  is also the latest updated vertex in  $\widehat{\pi}_i$ . Using Lemma 9 on paths in  $\widehat{Q}$ , we obtain that  $v_i \neq v_j$  for each  $i \neq j$  in  $[1, q]$ , and thus  $|\widehat{Q}| = q > z$ . This implies that at time  $t$  there are more than  $z$  historical paths between  $x$  and  $v$  in  $G$ , clearly a contradiction.  $\square$

The next corollary, whose proof follows immediately from Definition 4, Definition 5, and Lemma 10, bounds the number of paths that stop being locally historical at each time during a sequence of updates:

**Corollary 1** *Let  $G$  be a graph subject to a sequence  $\Sigma$  of update operations. If shortest paths are unique and at any time  $t$  throughout the sequence  $\Sigma$  there are at most  $z$  historical paths between each pair of vertices, then the number of paths that stop being locally historical at each update is  $O(zn^2)$  in the worst case.*

We finally bound the number of paths that become locally historical at each time during a sequence of updates:

**Theorem 5** *Let  $G$  be a graph subject to a sequence  $\Sigma$  of update operations and let  $m$  be the maximum number of edges in  $G$  throughout sequence  $\Sigma$ . If shortest paths are unique and at any time throughout the sequence  $\Sigma$  there are at most  $z$  historical paths between each pair of vertices, then the number of paths that become locally historical after each update is:*

- (1)  $O(zmn)$  in the worst case.
- (2)  $O(zn^2)$  amortized over  $\Omega(m/n)$  update operations.

**Proof.** Analogous to the proof of Lemma 1. Indeed, Claim (1) follows immediately from Lemma 8. To prove Claim (2), we assign a debit to each locally historical path in the graph. The debit is payed for by the operation that makes it stop being locally historical. By Corollary 1, at most  $O(zn^2)$  paths can stop being locally historical at each update. Moreover, by Lemma 8 there can be at most  $zmn$  locally historical paths in a graph at any time, so the total unpaid debit at the end of the sequence never exceeds  $zmn$ . Thus, the amortized number of paths that start being locally shortest after each update in any sequence of  $\Omega(m/n)$  operations can be at most  $O(zn^2)$ .  $\square$

## 4.2 A Worst-Case Scenario

We now show that there exists a graph and an update sequence that produces  $\Omega(n^3)$  historical paths. This implies by Theorem 5 an amortized number of  $\Omega(n^3)$  new locally historical paths per update. Consider the graph in Figure 6, where  $L_1 \dots L_6$  are layers containing  $n/6$  vertices each. The graphs induced by  $L_1 \cup L_2$ ,  $L_3 \cup L_4$  and  $L_5 \cup L_6$  are complete bipartite graphs with edge weights equal to 1, while the graphs induced by  $L_2 \cup L_3$  and  $L_4 \cup L_5$  have weights shown below the corresponding edges. The update sequence is as follows:

**Phase I.** Perform the following sequence of  $k = (n/6) - 1$  decreases on edges in  $L_2 \cup L_3$ :

- 1)  $n + 1 \rightarrow n - 1$
- 2)  $n + 2 \rightarrow n - 2$
- 3)  $n + 3 \rightarrow n - 3$
- $\vdots$
- k)  $n + k \rightarrow n - k$

**Phase II.** Perform the following sequence of  $k = (n/6) - 1$  increases on edges in  $L_4 \cup L_5$ :

- 1)  $n - k \rightarrow n + k$
- 2)  $n - (k - 1) \rightarrow n + (k - 1)$
- 3)  $n - (k - 2) \rightarrow n + (k - 2)$
- $\vdots$
- k)  $n - 1 \rightarrow n + 1$

During Phase I,  $(k + 1)^2$  paths connecting  $L_1$  to  $L_5$  become historical at each decrease, and thus at the end of Phase I there are  $\Theta(k^3)$  historical paths in total. During Phase II, at each increase,  $k$  paths connecting each pair of vertices in  $L_1$  and  $L_6$  become locally historical in




Figure 6: Worst-case instance with  $\Omega(n^3)$  historical paths and an amortized number of  $\Omega(n^3)$  new locally historical paths per update.

the graph. Since  $k = (n/6) - 1$ , each increase takes  $\Omega(n^3)$  time. Notice that this update sequence can be made arbitrarily long by repeating Phase I and Phase II back and forth many times. In this scenario, maintaining locally historical paths would require at least  $\Omega(n^3)$  time per update. We will see in the next section how to circumvent this problem.

### 4.3 Keeping Historical Paths Under Control: Smoothing

In this section, we show that it is possible to transform on-line every update sequence into another sequence so that the total number of historical paths in the graph is always very close to  $n^2$ . To guarantee that this transformation does not alter the input update sequence by affecting the result of shortest path queries, we make sure that it produces exactly the same changes in the graph as the input sequence. In the following, we will call such a transformation SMOOTHING. Furthermore, we will call *optimal* a historical path that is also a shortest path.

We first observe that at any time  $t$ , if shortest paths are unique, then there can be at most  $n^2$  optimal historical paths. Any other historical path must have been optimal strictly before time  $t$ , and is no longer optimal at time  $t$ . By Definition 4, if we update a vertex, every non-optimal historical path containing it will immediately stop being historical. To become again historical, it has to be a shortest path once again. Therefore, to reduce the number of historical paths, it is sufficient to add to the input update sequence suitable operations whose only effect is to clean up non-optimal historical paths. For this reason, we will call these additional operations CLEANUP UPDATES. To guarantee that cleanup updates do not alter the input update sequence by affecting the result of shortest path queries, we simply let them overwrite edge weights without changing their values. We now define how our smoothing strategy works:




Figure 7: Graphs of historical paths  $H_{xy}$  for an update sequence  $\Sigma$  and for the smoothed updated sequence  $F(\Sigma)$ . If  $t_{\sigma^{i+1}} < t_{\sigma^i} + T$ , then  $\pi^i$  cannot be historical at time  $t$  in the graph produced by  $F(\Sigma)$ .

**Definition 6** Let  $\Sigma = \langle \sigma_1, \sigma_2, \dots, \sigma_k \rangle$  be an update sequence and let

$$F(\Sigma) = \langle \sigma_1, \delta_1^1, \dots, \delta_1^{q_1}, \sigma_2, \delta_2^1, \dots, \delta_2^{q_2}, \dots, \sigma_k, \delta_k^1, \dots, \delta_k^{q_k} \rangle$$

be another sequence derived from  $\Sigma$  such that each  $\delta_i^j$  is a cleanup update, and such that, if a vertex is updated by operation  $\sigma_t$  in  $\Sigma$ , it is touched by cleanup update  $\delta_i^j$  in  $F(\Sigma)$  when  $i = t + 1, t + 2, t + 4, t + 8, \dots, t + 2^p, \dots$ , for some  $j$ . We call  $F(\Sigma)$  a SMOOTHED SEQUENCE derived from  $\Sigma$ , we refer to  $\langle \delta_i^1, \dots, \delta_i^{q_i} \rangle$  as the BURST OF CLEANUP UPDATES AFTER  $\sigma_i$ , and we say that cleanup update  $\delta_{t+2^p}^j$  is TRIGGERED BY THE ORIGINAL UPDATE  $\sigma_t$ .

Note that a smoothed sequence is obtained by squeezing a burst of cleanup updates between each pair of consecutive updates in the original sequence. Throughout the paper, we assume that the burst of cleanup updates  $\langle \delta_i^1, \dots, \delta_i^{q_i} \rangle$  occurs exactly at the same time as update  $\sigma_i$ , i.e.,  $t_{\delta_i^j} = t_{\sigma_i}$  for any  $j$ . We now show that a smoothed sequence is operationally equivalent to the sequence it is derived from, and produces a smaller number of historical paths.

**Theorem 6** Let  $\Sigma = \langle \sigma_1, \sigma_2, \dots, \sigma_k \rangle$  be an update sequence of length  $k$ . Then the smoothed sequence

$$F(\Sigma) = \langle \sigma_1, \delta_1^1, \dots, \delta_1^{q_1}, \sigma_2, \delta_2^1, \dots, \delta_2^{q_2}, \dots, \sigma_k, \delta_k^1, \dots, \delta_k^{q_k} \rangle$$

of Definition 6 has the following properties:

1. For  $1 \leq i \leq k$ ,  $q_i \leq O(\log k)$ ;
2. For  $1 \leq i \leq k$ , after performing operation  $\sigma_i$ , the graph produced by  $\Sigma$  and the graph produced by  $F(\Sigma)$  are equal;

3. At any time there are at most  $O(\log k)$  historical paths between each pair of vertices in the graph produced by  $F(\Sigma)$ .

**Proof.** Properties 1 and 2 follow directly from Definition 6. To prove Property 3, we first note that a cleanup update on a vertex  $v$  does not create new historical paths: its only effect is to remove non-optimal historical paths containing vertex  $v$ . For each pair of vertices  $(x, y)$  in  $G$ , consider the *graph of historical paths*  $H_{xy}$ . There is a vertex in  $H_{xy}$  for each update operation  $\tau$  in  $F(\Sigma)$  and there is an edge  $(\tau_a, \tau_b)$  in  $H_{xy}$ , with  $t_{\tau_a} < t_{\tau_b}$ , for each path  $\pi_{xy}$  in  $G$  having the following three properties:

- (i)  $\pi_{xy}$  contains both  $v_{\tau_a}$  and  $v_{\tau_b}$  (possibly  $v_{\tau_a} = v_{\tau_b}$ );
- (ii) no vertex in  $\pi_{xy}$  is updated in the time interval  $(t_{\tau_a}, t_{\tau_b})$ ;
- (iii)  $\pi_{xy}$  becomes a shortest path at some time  $t \in [t_{\tau_a}, t_{\tau_b})$ .

It is easy to see that at any time at most one path  $\pi_{xy}$  in  $G$  can satisfy Properties (i), (ii) and (iii) above.

Let  $\pi^1, \dots, \pi^z$  be the historical paths between vertices  $x$  and  $y$  in the graph  $G$  obtained right after performing update  $\sigma_t$  in  $F(\Sigma)$ , i.e., at time  $t$ . Let  $(\tau^i, \hat{\tau}^i)$  be the edge in  $H_{xy}$  corresponding to  $\pi^i$ ,  $1 \leq i \leq z$ . Since  $\pi^i$  is historical at time  $t$ , it must be  $t_{\tau^i} \leq t < t_{\hat{\tau}^i}$ . If  $\tau^i$  is a cleanup update, let  $\sigma^i$  be the original update that triggered  $\tau^i$ ,  $t_{\sigma^i} < t_{\tau^i}$ . Thus, to each historical path  $\pi^i$  we can associate an original update  $\sigma^i$ , with  $t_{\sigma^1} < t_{\sigma^2} < \dots < t_{\sigma^z} \leq t$ . Let  $\sigma^i$  and  $\sigma^{i+1}$  be two such updates, and let  $T$  be the largest power of 2 smaller than  $t - t_{\sigma^i}$ , i.e.,  $T = 2^{\lfloor \log_2(t - t_{\sigma^i}) \rfloor}$ . We claim that it must be  $t_{\sigma^i} + T \leq t_{\sigma^{i+1}}$ .

Indeed, assume by contradiction that  $t_{\sigma^{i+1}} < t_{\sigma^i} + T$  (see Figure 7). Since  $T = 2^{\lfloor \log_2(t - t_{\sigma^i}) \rfloor}$  is a power of 2, by Definition 6 there is a cleanup update triggered by  $\sigma^i$  that touches a vertex of  $\pi^i$  at time  $t_{\sigma^i} + T$ : thus  $\pi^i$  cannot be historical at time  $t_{\sigma^i} + T$  by Definition 4. Notice that, since  $\pi^{i+1}$  must be a shortest path between  $x$  and  $y$  after  $\pi^i$ , and no vertex of  $\pi^i$  and  $\pi^{i+1}$  is updated in the time interval  $[t_{\sigma^{i+1}}, t]$ ,  $\pi^i$  cannot become again a shortest path between  $x$  and  $y$  in the time interval  $[t_{\sigma^{i+1}}, t]$ . Thus  $\pi^i$  cannot become historical again in the time interval  $[t_{\sigma^i} + T, t]$ , contradicting the assumption that  $\pi^i$  was historical at time  $t$ .

We have thus shown that if  $\pi^1, \dots, \pi^z$  are historical paths at time  $t$ , there must be  $z$  updates  $\sigma^1, \dots, \sigma^z$  in the original sequence  $\Sigma$  such that  $t_{\sigma^i} + 2^{\lfloor \log_2(t - t_{\sigma^i}) \rfloor} \leq t_{\sigma^{i+1}} \leq t$ , for  $1 \leq i \leq z - 1$ . This implies  $z = O(\log t) = O(\log k)$ .  $\square$

We note that  $\Theta(\log k)$  historical paths between a pair of vertices in  $G$  can be actually realized, as illustrated in Figure 8.

## 5 Fully Dynamic Shortest Paths

In this section we show that the properties of locally historical paths given in Section 4 allow us to use the update algorithm of Section 3 in a fully dynamic setting, and to achieve  $\tilde{O}(n^2)$  amortized time per update. This is the first general algorithm for fully dynamic all pairs




Figure 8: Worst case instance of  $\Sigma$  such that at time  $k$  there are  $\Theta(\log k)$  historical paths between pair  $x, y$  in the graph produced by  $F(\Sigma)$ .

shortest path that is faster than recomputing the solution from scratch after each update. The approach we follow here differs in just two points from Section 3:

1. We maintain locally *historical* paths instead of locally *shortest* paths. Since by Lemma 7 shortest paths are locally historical, this guarantees that we are still maintaining information about shortest paths.
2. We provide a new front-end **fully-update** for the update operation, which calls the **update** procedure of Figure 2, implementing the smoothing strategy of Section 4.3.

We now describe in more detail how to achieve these modifications, and then discuss correctness and time/space requirements of the new procedure **fully-update**.

## 5.1 Data Structure

We maintain exactly the same data structure described in Section 3.1, with just one difference: we replace “shortest” with “historical”.

$$\begin{aligned}
 P_{xy} &= \{ \pi_{xy} & : \pi_{xy} \text{ is a locally historical path in } G \} \\
 P_{xy}^* &= \{ \pi_{xy} & : \pi_{xy} \text{ is a historical path in } G \} \\
 L(\pi_{xy}) &= \{ \pi_{x'y} = \langle x', x \rangle \cdot \pi_{xy} & : (x', x) \in E \text{ and } \pi_{x'y} \text{ is a locally historical path in } G \} \\
 L^*(\pi_{xy}) &= \{ \pi_{x'y} = \langle x', x \rangle \cdot \pi_{xy} & : (x', x) \in E \text{ and } \pi_{x'y} \text{ is a historical path in } G \} \\
 R(\pi_{xy}) &= \{ \pi_{xy'} = \pi_{xy} \cdot \langle y, y' \rangle & : (y, y') \in E \text{ and } \pi_{xy'} \text{ is a locally historical path in } G \} \\
 R^*(\pi_{xy}) &= \{ \pi_{xy'} = \pi_{xy} \cdot \langle y, y' \rangle & : (y, y') \in E \text{ and } \pi_{xy'} \text{ is a historical path in } G \}
 \end{aligned}$$

In order to support the smoothing operation, we also keep for each vertex  $v$  the time  $t_v$  of its latest update and a counter *time* of the number of **fully-update** operations performed. The notation introduced in this section is summarized in Table 3.

$P_{xy}$	set of locally historical paths from $x$ to $y$ in $G$
$P_{xy}^*$	set of historical paths from $x$ to $y$ in $G$
$L(\pi_{xy})$	set of pre-extensions $\langle x', x \rangle \cdot \pi_{xy}$ of $\pi_{xy}$ that are locally historical paths in $G$
$L^*(\pi_{xy})$	set of pre-extensions $\langle x', x \rangle \cdot \pi_{xy}$ of $\pi_{xy}$ that are historical paths in $G$
$R(\pi_{xy})$	set of post-extensions $\pi_{xy} \cdot \langle y, y' \rangle$ of $\pi_{xy}$ that are locally historical paths in $G$
$R^*(\pi_{xy})$	set of post-extensions $\pi_{xy} \cdot \langle y, y' \rangle$ of $\pi_{xy}$ that are historical paths in $G$

Table 3: Notation introduced in Section 5.1.

## 5.2 Implementation of Operations

The operations are exactly the same as described in Section 3.2. The only difference is that we do not call directly the `update` procedure shown in Figure 2, but rather provide a new front-end procedure named `fully-update` and shown in Figure 9, which implements the smoothing strategy of Section 4.3 by calling the `update` procedure of Figure 2 for both original and cleanup updates.

## 5.3 Analysis

The correctness of query operations follows directly from Lemma 7. To prove the correctness of `fully-update`( $v, w'$ ), we observe that it calls `update`( $v, w'$ ) to update the data structure after changing the weights of the edges incident to  $v$  to the new values  $w'$  (line 3). Then, it repeatedly calls `update`( $v, w$ ) on vertices that have to be smoothed (lines 4–5): notice that the call overwrites the edge weights without changing their values, and the only effect of this operation is to remove non-optimal historical paths from the data structure as discussed in Section 4.3. Therefore, it is sufficient to prove the correctness of `update`. To this aim, it is easy to check that both Invariant 1 and its proof continue to hold using the new data structure of Section 5.1 instead of the data structure of Section 3.1. We can thus prove the following theorem:

**Theorem 7** *If shortest paths are unique and edge weights are non-negative, then algorithm `update` correctly updates  $P_{xy}$  and  $P_{xy}^*$  for each pair of vertices  $x$  and  $y$ .*

---

<b>fully-update</b> ( $v, w'$ ):	
1. $time \leftarrow time + 1$	
2. $t_v \leftarrow time$	
3. <b>update</b> ( $v, w'$ )	{original update $\sigma$ }
4. <b>for each</b> $u \in V$ : $time - t_u = 2^{\lceil \log_2(time - t_u) \rceil}$ <b>do</b>	{smoothing}
5. $\text{update}(u, w)$	{cleanup update $\delta$ }

---

Figure 9: Implementation of the `fully-update` operation.



**Proof.** The proof is exactly the same as the proof of Theorem 2, provided that we suitably replace “shortest” with “historical” and “increase” with “update”.  $\square$

We conclude our analysis of the fully dynamic algorithm by discussing time and space requirements of our implementation. To this aim, we first analyze procedure `update` in a fully dynamic setting.

**Lemma 11** *In a fully dynamic sequence of  $\Omega(m/n)$  operations, if shortest paths are unique, edge weights are non-negative, and at any time there are at most  $z$  historical paths between each pair of vertices, our data structure supports each `update` operation in  $O(zn^2 \log n)$  amortized time. The space used is  $O(zmn)$ .*

**Proof.** We rephrase the proof of Theorem 3, by suitably replacing “shortest” with “historical”, and by using Lemma 10, Theorem 5, and Lemma 8 instead of Lemma 2, Theorem 1, and Lemma 3, respectively.

The bounds for queries are immediate, since the minimum weight path in the priority queue  $P_{xy}$  can be retrieved in optimal time.

Since each iteration of `cleanup` removes a path  $\pi_{xy}$  from a constant number of lists and priority queues, it requires  $O(\log n)$  time in the worst case. By Lemma 10, at most  $O(zn^2)$  locally historical paths can go through the updated vertex in the worst case, and thus `cleanup` will perform at most  $O(zn^2)$  iterations, thus speeding at most  $O(zn^2 \log n)$  time.

Next, it is easy to see that *Phase 1* of `fixup` requires  $O(n \log n)$  time in the worst case. To bound the time required by *Phase 2* of `fixup`, we consider that finding the minimum in  $P_{xy}$  takes constant time and each insertion in  $H$  takes  $O(\log n)$  time. This gives a total bound of  $O(n^2 \log n)$  for this phase. Finally, we note that a path  $\pi_{xy}$  is a new locally historical path if and only if both  $\ell(\pi_{xy}) = \langle x, \dots, b \rangle \in P_{xb}^*$  and  $r(\pi_{xy}) = \langle a, \dots, y \rangle \in P_{ay}^*$  after the update, and either  $\ell(\pi_{xy}) \notin P_{xb}^*$  or  $r(\pi_{xy}) \notin P_{ay}^*$  before calling `fixup`. Thus, since in *Phase 3* we process a shortest path  $\pi_{xy}$  extracted from  $H$  only if it was not already in  $P_{xy}^*$  before calling `fixup`, then we spend time only for the new locally historical paths. This implies that *Phase 3* requires  $O(n^2 \log n)$  time for extracting the  $O(n^2)$  pairs which are initially in  $H$ , plus  $O(\log n)$  time to insert/delete each new locally historical path in  $H$  and  $P$ : let  $p$  the number of such paths. By Theorem 5, there can be  $O(zn^2)$  paths that start being locally historical after each update, amortized over a sequence of  $\Omega(m/n)$  operations. Moreover, by Lemma 10, there can be at most  $O(zn^2)$  locally historical paths that were removed by `cleanup` since they contain the updated vertex, and are added back by `fixup`: thus,  $p = O(zn^2)$  and `update` requires  $O(zn^2 \log n)$  overall time.

At any time, each locally historical path  $\pi_{xy}$  can only appear in a constant number of data structures: namely, in  $H$ ,  $P_{xy}$ ,  $L(r(\pi_{xy}))$ ,  $R(\ell(\pi_{xy}))$ , and possibly in  $P_{xy}^*$ ,  $L^*(r(\pi_{xy}))$ ,  $R^*(\ell(\pi_{xy}))$ . Thus, the space usage of our data structure is bounded by the number of locally historical paths in a graph, which is  $O(zmn)$  in the worst case by Lemma 8.  $\square$

**Theorem 8** *In a fully dynamic sequence of  $\Omega(m/n)$  operations, if shortest paths are unique and edge weights are non-negative, our data structure supports each `fully-update` operation*

in  $O(n^2 \log^3 n)$  amortized time, and each distance and path query in optimal time. The space used is  $O(mn \log n)$ .

**Proof.** As in the proof of Theorem 3, the bounds for queries are immediate. To prove our update bound, we observe that procedure `fully-update` generates the smoothed sequence of Definition 6: thus, by Theorem 6 at any time in a sequence of  $k$  updates there are at most  $z = O(\log k)$  historical paths between each pair of vertices in  $G$ .

By Lemma 11, each call to `update` requires  $O(n^2 \log n \log k)$  amortized time. Since each `fully-update` causes as many as  $O(\log k)$  `update` operations, then the amortized running time of `fully-update` is  $O(n^2 \log n \log^2 k) = O(n^2 \log^3 n)$  on any sequence of length  $k$ , with  $k$  polynomial in  $n$ . To get rid of this assumption, we can simply restart our data structure from scratch every  $\Theta(n^2)$  operations. The space bound follows immediately from Lemma 11.  $\square$

## 6 Conclusions

In this paper we have presented a new algorithm that achieves nearly-quadratic update bounds for fully dynamic all-pairs shortest paths on graphs with non-negative real edge weights. The algorithm is deterministic and uses simple data structures. According to a recent computational study [3], the techniques described in this paper are not only asymptotically efficient, but can yield very fast implementations in many practical scenarios. The interested reader can download the C package used in the experiments at the URL: <http://www.dis.uniroma1.it/~demetres/experim/dsp/>.

The algorithms presented in this paper hinge upon the novel notion of locally-defined path properties. Using this approach, we have considered *locally shortest paths*, i.e., paths whose proper subpaths are shortest paths. Since at most  $O(n^2)$  such paths (amortized) can appear or disappear in the graph at each update in a partially dynamic setting, maintaining them in a data structure can be done in  $O(n^2 \log n)$  amortized time per update. In a fully dynamic setting though, changes in locally shortest paths may be much higher. To overcome this problem, we have used again the locality approach, considering *locally historical paths*, i.e., paths whose proper subpaths are historical, thus achieving an  $O(n^2 \log^3 n)$  amortized time per update with a space usage of  $O(mn \log n)$ . Using our approach, but with a different smoothing strategy, Thorup [19] has shown how to achieve  $O(n^2(\log n + \log^2(m/n)))$  amortized time per update and  $O(mn)$  space.

There are several issues that seem worth of further investigation. First, can we reduce the space usage to  $O(n^2)$ ? Second, and perhaps more importantly, can we solve efficiently fully dynamic *single-source* shortest paths on general graphs?

## Acknowledgements

We are indebted to Valerie King, Mikkel Thorup, and Uri Zwick for many useful comments and suggestions on previous versions of this work.

## References

- [1] G. Ausiello, G.F. Italiano, A. Marchetti-Spaccamela, and U. Nanni. Incremental algorithms for minimal length paths. *Journal of Algorithms*, 12(4):615–38, 1991.
- [2] S. Baswana, R. Hariharan, and S. Sen. Improved decremental algorithms for transitive closure and all-pairs shortest paths. In *Proc. 34th ACM Symposium on Theory of Computing (STOC'02)*, pages 117–123, 2002.
- [3] C. Demetrescu, S. Emiliozzi, and G.F. Italiano. Experimental analysis of dynamic all pairs shortest path algorithms. In *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'04) (to appear)*, 2004.
- [4] C. Demetrescu and G.F. Italiano. Fully dynamic all pairs shortest paths with real edge weights. In *Proc. of the 42nd IEEE Annual Symposium on Foundations of Computer Science (FOCS'01), Las Vegas, Nevada*, pages 260–267, 2001.
- [5] C. Demetrescu and G.F. Italiano. Improved bounds and new trade-offs for dynamic all pairs shortest paths. In *Proc. of the 29th International Colloquium on Automata, Languages, and Programming (ICALP'02), Málaga, Spain*, pages 633–643, 2002.
- [6] C. Demetrescu and G.F. Italiano. A new approach to dynamic all pairs shortest paths. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC'03), San Diego, CA*, pages 159–166, 2003.
- [7] S. Even and H. Gazit. Updating distances in dynamic graphs. *Methods of Operations Research*, 49:371–387, 1985.
- [8] J. Fakcharoemphol and S. Rao. Planar graphs, negative weight edges, shortest paths, and near linear time. In *Proc. of the 42nd IEEE Annual Symposium on Foundations of Computer Science (FOCS'01), Las Vegas, Nevada*, pages 232–241, 2001.
- [9] D. Frigioni, A. Marchetti-Spaccamela, and U. Nanni. Semi-dynamic algorithms for maintaining single source shortest paths trees. *Algorithmica*, 22(3):250–274, 1998.
- [10] D. Frigioni, A. Marchetti-Spaccamela, and U. Nanni. Fully dynamic algorithms for maintaining shortest paths trees. *Journal of Algorithms*, 34:351–381, 2000.
- [11] M.R. Henzinger, P. Klein, S. Rao, and S. Subramanian. Faster shortest-path algorithms for planar graphs. *Journal of Computer and System Sciences*, 55(1):3–23, August 1997.
- [12] V. King. Fully dynamic algorithms for maintaining all-pairs shortest paths and transitive closure in digraphs. In *Proc. 40th IEEE Symposium on Foundations of Computer Science (FOCS'99)*, pages 81–99, 1999.
- [13] P. Loubal. A network evaluation procedure. *Highway Research Record 205*, pages 96–109, 1967.

- [14] J. Murchland. The effect of increasing or decreasing the length of a single arc on all shortest distances in a graph. Technical report, LBS-TNT-26, London Business School, Transport Network Theory Unit, London, UK, 1967.
- [15] G. Ramalingam and T. Reps. An incremental algorithm for a generalization of the shortest path problem. *Journal of Algorithms*, 21:267–305, 1996.
- [16] G. Ramalingam and T. Reps. On the computational complexity of dynamic graph problems. *Theoretical Computer Science*, 158:233–277, 1996.
- [17] V. Rodionov. The parametric problem of shortest distances. *U.S.S.R. Computational Math. and Math. Phys.*, 8(5):336–343, 1968.
- [18] H. Rohnert. A dynamization of the all-pairs least cost problem. In *Proc. 2nd Annual Symposium on Theoretical Aspects of Computer Science, (STACS'85), LNCS 182*, pages 279–286, 1985.
- [19] M. Thorup. Tighter fully-dynamic all pairs shortest paths, 2003. Personal communication.