

Single-Source Shortest Paths and Strong Connectivity in Dynamic Planar Graphs¹

Panagiotis Charalampopoulos^{a,b,c}, Adam Karczmarz^a

^a*Institute of Informatics, University of Warsaw, Warsaw, Poland*

^b*Department of Informatics, King's College London, London, UK*

^c*Efi Arazi School of Computer Science, The Interdisciplinary Center Herzliya, Herzliya, Israel*

Abstract

We give a fully dynamic single-source shortest paths data structure for planar weighted digraphs with $\tilde{O}(n^{4/5})$ worst-case update time and $O(\log^2 n)$ query time. Here, a single update can either change the graph by inserting or deleting an edge, or reset the source s of interest. All known non-trivial planarity-exploiting *exact* dynamic single-source shortest paths algorithms to date had polynomial query time. We then extend our approach, obtaining a data structure that can maintain a planar digraph under edge insertions and deletions, and is capable of returning the identifier of the strongly connected component of any query vertex. The worst-case update and query time bounds are the same as for our single-source distance oracle. To the best of our knowledge, this is the first fully dynamic strong-connectivity algorithm achieving both sublinear update time and polylogarithmic query time for an important class of digraphs.

Keywords: dynamic graph algorithms, planar graphs, single-source shortest paths, strong connectivity

1. Introduction

The dynamic shortest paths problem seeks for a data structure that maintains a graph under updates and supports querying for shortest paths and/or distances. Depending on the set of supported updates, we call such a graph data structure *fully dynamic* if both edge insertions and deletions are allowed, *incremental* if only edge insertions (or edge weight decreases) are supported, or *decremental* if only edge deletions (or weight increases) are allowed. In the *all-pairs* variant of dynamic shortest paths problem one has to support shortest path queries between any pair of vertices of the graph. In the *single-source* variant all shortest paths queries have to originate in a fixed distinguished vertex and the only input to a query is the target vertex.

For the most general setting where one requires *exact* answers, Demetrescu and Italiano [2] gave a fully dynamic algorithm (improved slightly by Thorup [3]) recomputing the all-pairs shortest paths matrix in nearly optimal $\tilde{O}(n^2)$ amortized time even if real edge weights are allowed.² Note that the distance matrix can be recomputed from scratch in $\tilde{O}(nm)$ time [4]. Fully dynamic all-pairs shortest paths data structures with subcubic *worst-case* update bounds are also known [5, 6, 7]. There exist faster algorithms if the input graph is unweighted and *partially dynamic* (i.e., incremental or decremental) [8, 9, 10]. However, none of the known results improves upon a trivial, recompute-from-scratch algorithm with $\tilde{O}(mn)$ update time and $O(1)$ query time if the graph is *sparse*, i.e., $m = \tilde{O}(n)$. For the single-source variant, all known non-trivial exact dynamic shortest paths algorithms [11] are partially dynamic and yield no improvement over the respective recompute-from-scratch algorithm in the sparse case either.

¹An earlier version of this work was presented at ESA 2020 [1].

Email addresses: pcharalampo@gmail.com (Panagiotis Charalampopoulos), a.karczmarz@mimuw.edu.pl (Adam Karczmarz)

²Throughout this work, the \tilde{O} -notation suppresses $\log^{O(1)} n$ factors.

The lack of progress on obtaining an exact fully dynamic single-source shortest paths algorithms with $O(n^{3-\epsilon})$ initialization time, $O(m^{1-\epsilon})$ amortized update time and $O(n^{1-\epsilon})$ query time at the same time can be explained by a matching lower bound conditional on the (static) APSP conjecture [12]. In fact, breaking this barrier even in the partially dynamic setting for undirected unweighted graphs would be a large breakthrough [13].

As a result, since finding good exact algorithms for general graphs seems hopeless, one needs to look for approximate solutions, and/or restrict their attention to more structured graph classes and/or undirected graphs. Indeed, a large body of research has been devoted to designing *approximate* dynamic shortest paths algorithms, especially in partially dynamic settings [14, 15, 16, 17, 18, 19, 20, 21, 13, 22, 23, 24, 25, 26, 27, 28, 10, 29, 30], which find many applications, e.g., in various maximum flow related problems [21, 31]. Unfortunately, many of the known fastest approximate dynamic shortest path algorithms (e.g., [14, 23, 25]) suffer from assuming an oblivious adversary, which significantly limits their applicability (cf. e.g., [21]).

Similarly, dynamic shortest paths problems have also been studied for important graph classes like planar graphs [32, 33, 34, 35, 36, 37] or low-treewidth graphs [33, 38]. The primary reason why faster dynamic shortest paths algorithms in these cases are possible is the existence of non-trivial *distance oracles* for these classes. A distance oracle is a compact representation of the graph’s shortest paths such that the distance (or a distance estimate) between any pair of vertices can be retrieved efficiently. For general graphs, such non-trivial distance oracles exist only for undirected graphs and assuming an approximation ratio of at least 3 [39, 40, 41]. On the contrary, for planar graphs many non-trivial exact distance oracles have been proposed [42, 43, 44, 34, 35, 36, 45, 46]. The first exact oracles with *polylogarithmic* query time and subquadratic space have been obtained only recently [47, 48, 49, 50], following Cabello’s breakthrough of employing Voronoi diagrams for the planar diameter problem [51]. In addition, near-optimal (in terms of query time, construction time, and used space) $(1 + \epsilon)$ -approximate distance oracles have been known for nearly two decades [52, 53], and a lot of effort has been put to push the known bounds as close to optimal as possible [54, 55, 56, 57].

Dynamic shortest paths in planar graphs. In this paper our focus is on computing shortest paths in dynamic planar graphs and applications. Klein and Subramanian were the first to give a planarity-exploiting dynamic shortest paths algorithm [37] – their data structure worked for undirected graphs, was fully dynamic, $(1 + \epsilon)$ -approximate and had $\tilde{O}(n^{2/3})$ update and query time bounds. A data structure with the same bounds (up to polylogarithmic factors), but for *exact* distances in *directed* graphs was obtained in the breakthrough work of Fakcharoenphol and Rao [34] (later extended and slightly improved in [58, 59, 60, 35, 36]). Abraham et al. [32] gave a faster $(1 + \epsilon)$ -approximate dynamic algorithm for undirected graphs with $\tilde{O}(n^{1/2})$ update and query times. Karczmarz [38] matched this bound for directed planar graphs, albeit only in the $(1 + \epsilon)$ -approximate decremental setting.³ Abboud and Dahlgaard [61] showed that by the APSP conjecture, one should not expect an exact dynamic all-pairs shortest paths data structure for planar graphs with strongly sublinear⁴ product of update time and query time. However, no exact data structure with strongly sublinear update time and near-linear product of update time and query time is known.

The single-source scenario is much less studied for dynamic planar graphs. Karczmarz [38] showed a decremental $(1 + \epsilon)$ -approximate single-source shortest paths algorithm for minor-free (and thus also planar) digraphs with $\tilde{O}(n^{1/2})$ update time and $O(1)$ query time. Although not explicitly stated in the literature, the all-pairs data structure of [36] can be easily converted to a fully dynamic exact single-source distance oracle with $\tilde{O}(n/x)$ update time and $\tilde{O}(x)$ query time for any parameter $x \in [1, n^{1/3}]$ (we formally show this in Section 3). However, no *fully dynamic* single-source shortest paths algorithm for planar graphs to date has been able to achieve sublinear update time and at the same time either support queries in polylogarithmic time or at least break through the $\tilde{O}(n)$ update-query time product barrier, even in the approximate setting.

Recall that for general graphs, Roditty and Zwick [12] proved a lower bound for the dynamic single-source shortest paths problem, conditional on the conjectured hardness of the (static) APSP problem.

³In the statements of the results of [32, 38], we have assumed that the edge weights are in $[1, n^{O(1)}]$ for ease of presentation.

⁴In this paper, by *strongly sublinear* we mean $O(n^{1-\epsilon})$ for some $\epsilon > 0$.

Conjecture 1 (APSP Conjecture). *There exists no algorithm for solving the all-pairs shortest paths problem in general weighted graphs in time $O(n^{3-\epsilon})$ for any constant $\epsilon > 0$.*

Abboud and Dahlgaard [61] proved a lower bound, also conditional on the APSP conjecture, for single-source, single-sink dynamic distance oracles for planar graphs.

Theorem 2 ([61, Theorem 5]). *The APSP conjecture implies that there is no algorithm that maintains a planar graph of size n under edge insertions and deletions in $u(n)$ amortized time per update and answers distance queries from a fixed source s to a fixed sink t in $q(n)$ amortized time, such that $u(n) + q(n) = O(n^{1/2-\epsilon})$ for any constant $\epsilon > 0$. Furthermore, if $q(n) \geq u(n)$, the algorithm cannot have $q(n) \cdot u(n) = O(n^{1-\epsilon})$. This holds even if the planar embedding of the graph never changes.*

Our results. In this paper we show the first exact dynamic single-source shortest paths algorithm for planar graphs with *strongly sublinear* update time and *polylogarithmic* query time. Our algorithm, summarized by the following theorem and described in Section 3, is deterministic and can be easily extended to report the corresponding shortest path.

Theorem 3. *Let G be a real-weighted planar digraph with a source $s \in V(G)$. There exists an $O(n \log n)$ -space data structure maintaining G under edge insertions, edge deletions, and source changes with worst-case update time $O(n^{4/5} \log^2 n)$ that can compute $\text{dist}_G(s, v)$ for any $v \in V(G)$ in $O(\log^2 n)$ time. The initialization time is $O(n \log^2 n)$.*

To the best of our knowledge, this result constitutes the first known application of additively weighted Voronoi diagrams machinery (first introduced by Cabello [51]) in dynamic graph algorithms. More specifically, it is obtained by combining fully dynamic maintenance of r -divisions [37, 62], the multiple-source shortest paths data structure [36, 63], the shortest paths algorithm for dense distance graphs [34], the recent efficient construction of dual Voronoi diagrams [48] via FR-Dijkstra [34], and the efficient point location data structure for Voronoi diagrams [50].

We now provide a brief overview of the data structure underlying Theorem 3. We maintain distances in G from the source vertex s to each boundary vertex of each piece of an r -division of G using FR-Dijkstra. For each piece of the r -division, we maintain an additively weighted Voronoi diagram augmented with a point location data structure, with weights equal to the distances from s . Upon a query for $\text{dist}_G(s, v)$, we perform a point location query on the Voronoi diagram of a piece of the r -division that contains v .

It is worth noting that our data structure (and in fact all data structures obtained in this paper) works in the most general model of dynamic planar graphs where we only require that G remains planar after each update. Note that one can test whether each issued update violates the planarity of the graph in $O(\log^3 n)$ amortized time per update [65]. Some fully dynamic planar graph algorithms assume a weaker *plane* model (e.g., [66, 67, 60]) where some plane embedding of G is fixed and we only allow inserting edges connecting vertices that lie on a common face of (that embedding of) G .

We also generalize our single-source data structure to the case when, instead of a single source s , a set of *facilities* $F \subseteq V$ is given, and our goal is to locate the closest (i.e., minimizing $\text{dist}_G(f, v)$) facility $f \in F$ for a given query vertex.⁵ We show that maintaining such a data structure under edge updates issued to G , or updates to the facilities set F , is possible using $\tilde{O}(n^{3/4} \cdot |F|^{1/4} + n^{4/5})$ worst-case update time. The query time remains $O(\log^2 n)$. Note that even though multiple-source shortest paths or maximum flow problems can be typically easily reduced to the single-source case by adding a super-source, such a reduction does not preserve planarity and indeed handling multiple sources tends to be challenging in planar graphs (cf. e.g., [68, 69]). Our generalized data structure handles up to $O(n^{1/5})$ sources as efficiently as the single-source case. Moreover, the update time remains strongly sublinear unless the number of facilities is not strongly sublinear.

Surprisingly, we show that the same framework that we use to prove Theorem 3 can be applied to obtain interesting results not directly related to the shortest paths problem. Namely, in Section 4 we show a fully dynamic strong-connectivity algorithm for planar graphs, encapsulated in the following theorem.

⁵One can also view F as a set of sites of a graphic Voronoi diagram – then the query locates the cell of the Voronoi diagram wrt. F that a given vertex v belongs to.

Theorem 4. *Let G be a planar digraph. There exists an $O(n \log n)$ -space data structure maintaining G under edge insertions and deletions with $O(n^{4/5} \log^2 n)$ worst-case update time that can compute the identifier of the strongly connected component of any $v \in V(G)$ in $O(\log^2 n)$ time. The initialization time is $O(n \log^2 n)$.*

We now sketch the main ideas behind our fully dynamic strong-connectivity algorithm. As in Subramanian’s dynamic all-pairs reachability algorithm [62], the base of our data structure is a graph X , called a *reachability certificate*, that sparsifies the reachability information between boundary vertices $\partial\mathcal{R}$ of a fully dynamic r -division \mathcal{R} with few holes. Naively recomputing the strongly connected components of X gives us the restriction of the strongly connected components of G to the boundary vertices $\partial\mathcal{R}$. The main challenge, of course, is to compute the identifier of a strongly connected component (SCC) of an arbitrary non-boundary vertex v of G , internal to some piece P of the r -division \mathcal{R} . To this end, we use the following observation: suppose b_1, \dots, b_k are some vertices of G lying in distinct strongly connected components of G . Then, v is strongly connected to some b_j if and only if b_j is in the topologically earliest SCC of G reachable from v and b_j is in the topologically latest SCC of G that can reach v . Roughly speaking, this observation applied to the boundary vertices of P labeled using the topological order of their respective SCCs in the certificate X , allows us to identify the SCC of v using two point location queries on the Voronoi diagram of piece P . Each such point location query, computes, instead of the nearest site of v , the highest (or lowest) priority site that can reach v (that v can reach, resp.), and can be simulated using a standard point location query on a Voronoi diagram [50].

Whereas maintaining strongly connected components is a well-studied problem in partially dynamic settings [70, 71, 72, 73, 29], we are not aware of any non-trivial fully dynamic strongly connected components data structures designed specifically for this problem for *any* digraph class – note that one could use a fully dynamic transitive closure data structure for this task: for example, the dynamic plane transitive closure data structure of [66] which has $\tilde{O}(n^{1/2})$ update and query time. Such a strongly connected components data structure (i.e., with both update and query bounds $O(n^{1-\epsilon})$) for *general graphs* is in fact ruled out by a conditional (on SETH) lower bound [74]. As a result, to the best of our knowledge, we obtain the first fully dynamic strongly connected components algorithm to achieve sublinear update-query time product for *any* important class of digraphs.

The undirected counterpart of the dynamic strongly connected components problem, the *dynamic connectivity* problem, is very well-studied. Near-optimal deterministic amortized update bounds [75, 76, 77] and randomized worst-case update bounds [78, 79] (see also [80]) are known for fully dynamic general graphs. An almost optimal deterministic worst-case update bound was very recently achieved [81, 82]. For fully dynamic planar graphs polylogarithmic worst-case update bounds are known to be achievable even deterministically [64].

2. Preliminaries

Throughout the paper we consider as input a simple, directed and weighted planar graph G with n vertices, and no negative weight cycles. We call a planar graph G *plane* if some embedding of G is assumed. We use $|G|$ to denote the number of vertices of G . Since simple planar graphs are sparse, $|E(G)| = O(|G|)$ as well.

We use the terms weight and length for edges and paths interchangeably throughout the paper. For any two vertices $u, v \in V(G)$, we denote by $\text{dist}_G(u, v)$ the length of some shortest $u \rightarrow v$ path in the graph G .

Multiple-source shortest paths. The multiple-source shortest paths (MSSP) data structure [63, 36] represents all shortest path trees rooted at the vertices of a single face f in a weighted plane digraph using a persistent dynamic tree. It can be constructed in $O(n \log n)$ time, requires $O(n \log n)$ space, and can report any distance between a vertex of f and any other vertex in the graph in $O(\log n)$ time. MSSP can be augmented to also return the first edge of this path (and each of its subsequent edges) in $O(\log \log n)$ time (cf. [35]).

An alternative implementation of MSSP due to Long and Pettie [47], that uses Euler Tour trees instead of Link-Cut trees, can be constructed in time $n^{1+o(1)}$, takes space $n^{1+o(1)}$, and answers all queries we are interested in within $\log^{o(1)} n$ time.

Throughout the paper, we use the MSSP data structure of [63, 36]. Substituting the one of [47] would improve the query times in Theorems 3 and 4 to $\log^{1+o(1)} n$ at the cost of a multiplicative $n^{o(1)}$ factor in each of the other measures of efficiency: update time, space, and initialization time.

Separators and recursive decompositions. Miller [83] showed how to compute, in a triangulated plane graph with n vertices, a simple cycle of size $2\sqrt{2}\sqrt{n}$ that separates the graph into two subgraphs, each with at most $2n/3$ vertices. Simple cycle separators can be used to recursively separate a planar graph until pieces have constant size. The authors of [84] show how to obtain a complete recursive decomposition tree $\mathcal{T}(G)$ of a triangulated graph G using cycle separators in $O(n)$ time. $\mathcal{T}(G)$ is a binary tree whose nodes correspond to subgraphs of G (pieces), with the root being all of G and the leaves being pieces of constant size. We identify each piece P with the node representing it in $\mathcal{T}(G)$. We can thus abuse notation and write $P \in \mathcal{T}(G)$. The *boundary vertices* ∂P of a *non-leaf piece* P are vertices that P shares with some other piece $Q \in \mathcal{T}(G)$ that is not P 's ancestor. For convenience we define the boundary set ∂L of a leaf piece L to be its entire vertex set $V(L)$. We assume P to inherit the embedding of G . The faces of P that are faces of G are called *natural*, whereas the faces of P that are not natural are the *holes* of P . The construction of [84] additionally guarantees that for each piece $H \in \mathcal{T}(G)$, (a) H is connected, (b) if H is a non-leaf piece, then each natural face f of H is a face of a unique child of H , (c) H has $O(1)$ holes containing precisely the vertices ∂H . Throughout, to avoid confusion, we use *nodes* when referring to $\mathcal{T}(G)$ and *vertices* when referring to G or its subgraphs. It is well-known [85, 50, 73, 84] that by suitably choosing cycle separators one can also guarantee that (1) $\sum_{H \in \mathcal{T}(G)} |H| = O(n \log n)$, (2) $\sum_{H \in \mathcal{T}(G)} |\partial H|^2 = O(n \log n)$, and (3) $|\partial H| = O(\sqrt{n}/c^d)$, where node H of $\mathcal{T}(G)$ has depth d and $c > 1$ is some constant.

The recursive decomposition algorithm of [84] works with no changes and maintains all the properties of $\mathcal{T}(G)$ even if the initial graph G has a predefined set of boundary vertices ∂G of size $O(\sqrt{|G|})$ located on $O(1)$ of G 's faces. These faces are predefined as holes of G and are the only faces of G that are allowed to be non-triangular.

An *r-division* [86] \mathcal{R} of a planar graph, for $r \in [1, n]$, is a decomposition of the graph into $O(n/r)$ pieces, each of size $O(r)$, such that each piece P has $O(\sqrt{r})$ boundary vertices (denoted ∂P), i.e., vertices shared with some other piece of \mathcal{R} . We denote by $\partial \mathcal{R}$ the set $\bigcup_{P \in \mathcal{R}} \partial P$. If additionally all pieces are connected, and the boundary vertices of each piece P of the *r-division* \mathcal{R} are distributed among $O(1)$ faces of P (also called holes⁶ of P), we call \mathcal{R} an *r-division with few holes*.

In [84] it was shown that for every r larger than some constant, $\mathcal{T}(G)$ *admits* an *r-division with few holes*, i.e., there exists a subset of nodes of $\mathcal{T}(G)$ forming an *r-division with few holes* of G . Using this property, it is shown in [84] that an *r-division with few holes* of a triangulated graph can be computed in linear time. More generally, given a geometrically decreasing sequence of numbers $(r_m, r_{m-1}, \dots, r_1)$, where r_1 is a sufficiently large constant, $r_{i+1}/r_i \geq b$ for all i for some $b > 1$, and $r_m = n$, we can obtain *r_i-divisions with few holes* for all i in time $O(n)$ in total. For convenience, we define the only piece in the r_m -division to be G itself. These *r-division*s satisfy the property that a piece in the r_i -division is a – not necessarily strict – descendant (in $\mathcal{T}(G)$) of a piece in the r_j -division for each $j > i$. We also call such sequence of *r_i-divisions* obtained from $\mathcal{T}(G)$ a *recursive (r_m, \dots, r_1) -division* of G .

We assume for simplicity that all holes we ever encounter are simple cycles. Unfortunately, this is not true in general. However, non-simple holes do not pose a significant obstacle, and can be avoided by suitably extending the graphs, as discussed numerous times in the past, see e.g., [48, 35, 73, 87]. Roughly speaking, for each non-simple part of each hole, we perform an incision along this part, we duplicate the vertices and edges we cut along, and add zero-length edges between the two copies of each duplicated vertex.

Dense distance graphs and FR-Dijkstra. For a plane digraph H with weights from $\mathbb{R}_{\geq 0} \cup \{\infty\}$ and a distinguished set $\partial H \subseteq V(H)$ of boundary vertices lying on $O(1)$ faces of H , we denote by DDG_H the complete weighted graph on ∂H whose edge weights represent distances between all pairs of vertices of ∂H in H .

⁶This definition is slightly more general than the definition of a hole of a piece $P \in \mathcal{T}(G)$. Namely, the definition of an *r-division* does not assume a fixed embedding of the entire G ; it only assumes some fixed embeddings of individual pieces.

DDG_H can be computed in $O((|H| + |\partial H|^2) \log n)$ time using MSSP [36]. In particular, dense distance graphs for all pieces $H \in \mathcal{T}(G)$ can be computed in $O(n \log^2 n)$ time.

When $\mathcal{H} = \{H_1, \dots, H_q\}$ is a collection of plane graphs, we set $\text{DDG}(\mathcal{H}) := \bigcup_{H \in \mathcal{H}} \text{DDG}_H$.

Lemma 5 (FR-Dijkstra [34, 59]). *Given all DDG_{H_i} , one can compute a single-source shortest paths tree from any source s in $\text{DDG}(\mathcal{H})$ in $O(\sum_{i=1}^q |\partial H_i| \log^2 n)$ time, where $n = |V(\text{DDG}(\mathcal{H}))|$.⁷*

We next state some fairly standard definitions and lemmas about representing distances between some vertices of G of interest using unions of dense distance graphs of a recursive decomposition's (or r -division's) pieces (for instance cf. [58]), adapted to our notation. We include their proofs for completeness, but the reader may safely skip them. For example, Lemma 7 captures the well-known observation of [34] that in order to compute a shortest path between any pair of vertices of G , it is enough to compute a shortest path in a union of dense distance graphs from $\mathcal{T}(G)$ with only $O(\sqrt{n})$ vertices in total.

Lemma 6. *Let \mathcal{Q} be some collection of pieces from $\mathcal{T}(G)$ such that:*

1. *for each leaf piece $L \in \mathcal{T}(G)$, either L or some ancestor of L is in \mathcal{Q} ,*
2. *for each $H \in \mathcal{Q}$, if some ancestor of H is in \mathcal{Q} , then the parent of H is also in \mathcal{Q} .*

Then for any $u, v \in V(\text{DDG}(\mathcal{Q}))$, $\text{dist}_{\text{DDG}(\mathcal{Q})}(u, v) = \text{dist}_G(u, v)$.

Proof. Let us note that if for some $H \in \mathcal{T}(G)$ no ancestor of H belongs to \mathcal{Q} , then $\partial H \subseteq V(\text{DDG}(\mathcal{Q}))$. We prove this claim by induction on the level⁸ ℓ of piece H . For $\ell = 0$, we get $H \in \mathcal{Q}$, so clearly $\partial H \subseteq V(H) \subseteq V(\text{DDG}(\mathcal{Q}))$. Suppose $\ell \geq 1$. The statement is trivial if $H \in \mathcal{Q}$. Otherwise, consider the children H_1, \dots, H_k of H . By induction we get $\partial H_i \subseteq V(\text{DDG}(\mathcal{Q}))$. So in fact we have $\partial H \subseteq \bigcup_{i=1}^k \partial H_i \subseteq V(\text{DDG}(\mathcal{Q}))$.

Since the edges of each DDG_H encode lengths of some paths in G , we have that $\text{dist}_{\text{DDG}(\mathcal{Q})}(u, v) \geq \text{dist}_G(u, v)$. We now prove that there is a path of length at most $\text{dist}_G(u, v)$ in $\text{DDG}(\mathcal{Q})$. Let P be some shortest $u \rightarrow v$ path in G . Let H be a piece of $\mathcal{T}(G)$ of minimum level that contains P . We prove our claim by induction on the level ℓ of H .

First note that by property 1 of \mathcal{Q} , if $H \in \mathcal{Q}$ or no descendant of H belongs to \mathcal{Q} , then H has a nearest ancestor H^* (possibly $H^* = H$) such that $H^* \in \mathcal{Q}$. Observe that $V(\text{DDG}(\mathcal{Q})) \cap V(H) \subseteq \partial H^*$. Hence, $u, v \in \partial H^*$ and thus $\text{dist}_{\text{DDG}(\mathcal{Q})}(u, v) \leq \text{dist}_{\text{DDG}_{H^*}}(u, v) = \text{dist}_{H^*}(u, v) \leq \text{dist}_H(u, v) = \text{dist}_G(u, v)$.

So we can assume that $H \notin \mathcal{Q}$ and some descendant of H belongs to \mathcal{Q} . Note that then we have $\ell \geq 1$ and thus H is not a leaf. By property 2 of \mathcal{Q} , we have that no ancestor of H belongs to \mathcal{Q} . So for any child H_i of H , no ancestor of H_i belongs to \mathcal{Q} . We have proved that this implies $\partial H_i \subseteq V(\text{DDG}(\mathcal{Q}))$. Let us split P into maximal subpaths P_1, \dots, P_k such that each $P_i = u_i \rightarrow v_i$ is entirely contained in a single child H_i of H . Then for each i we have $\{u_i, v_i\} \subseteq \{u, v\} \cup \partial H_i \subseteq V(\text{DDG}(\mathcal{Q}))$ so by induction we get that $\text{dist}_{\text{DDG}(\mathcal{Q})}(u_i, v_i) \leq \text{dist}_G(u_i, v_i)$ which implies $\text{dist}_{\text{DDG}(\mathcal{Q})}(u, v) \leq \text{dist}_G(u, v)$. \square

Let L be some leaf of $\mathcal{T}(G)$. We define the *cone* of L , denoted $\text{cone}_G(L)$, to be the collection of pieces of $\mathcal{T}(G)$ containing L , all ancestors of L , and all siblings of (weak) ancestors of L . For some collection \mathcal{L} of leaf pieces of $\mathcal{T}(G)$, we define $\text{cone}_G(\mathcal{L}) = \bigcup_{L \in \mathcal{L}} \text{cone}_G(L)$.

Lemma 7 ([58, 34]). *Let \mathcal{L} be some collection of leaf pieces of $\mathcal{T}(G)$. Then:*

1. *For any $u, v \in V(\text{DDG}(\text{cone}_G(\mathcal{L})))$, $\text{dist}_G(u, v) = \text{dist}_{\text{DDG}(\text{cone}_G(\mathcal{L}))}(u, v)$.*
2. $\sum_{H \in \text{cone}_G(\mathcal{L})} |\partial H| = O(\sqrt{n|\mathcal{L}|})$.

⁷In particular H_i may be a single edge. This way, this lemma captures also the case when we compute shortest paths in a collection of DDGs with some auxiliary vertices and edges.

⁸Here, by level of H we mean the depth of the subtree of $\mathcal{T}(G)$ rooted at H .

Proof. To obtain item 1 it is enough to note that $\text{cone}_G(\mathcal{L})$ satisfies the requirements posed on the collection \mathcal{Q} in Lemma 6.

Let \mathcal{A} be the set containing all ancestors of all the leaf pieces $L \in \mathcal{L}$. We show item 2 by bounding the sum $X = \sum_{H \in \mathcal{A}} |\partial H|$. Since the number of boundary vertices of a piece is bounded by the sum of numbers of boundary vertices of its parent and its sibling, the sum $\sum_{H \in \text{cone}_G(\mathcal{L})} |\partial H|$ of our interest can be larger from X only by a constant factor.

Recall that $\mathcal{T}(G)$ admits an r -division for any $r \in [1, n]$, i.e., there exists such r -division \mathcal{R} that $\mathcal{R} \subseteq \mathcal{T}(G)$ and the boundary of each piece of \mathcal{R} equals the boundary of that piece in $\mathcal{T}(G)$. Let us split \mathcal{A} into two parts: let \mathcal{A}_1 contain those $H \in \mathcal{A}$ that are descendants of some piece $P \in \mathcal{R}$, and let $\mathcal{A}_2 = \mathcal{A} \setminus \mathcal{A}_1$. Let $X_i = \sum_{H \in \mathcal{A}_i} |\partial H|$.

Since each $L \in \mathcal{L}$ is a descendant of a unique piece $P \in \mathcal{R}$, we now bound the sum of $|\partial H|$ over all ancestors of L that are descendants of P . Recall (Section 2) that if H is a piece in a recursive decomposition of an n -vertex graph, then $|\partial H| = O(\sqrt{n}/c^d)$ where d is the depth of H in that decomposition. Consider the subtree of $\mathcal{T}(G)$ rooted at P – it forms a recursive decomposition $\mathcal{T}(P)$ of P with some initial boundary vertex set $|\partial P|$ of size $O(\sqrt{r})$. So the sum of $|\partial H|$ over all ancestors of L that are descendants of P is actually equal to the sum of $|\partial H|$ over all ancestors H of L in $\mathcal{T}(P)$. Since each ancestor has distinct integral depth, this sum is $O(\sum_{i=0}^{\infty} \sqrt{|P|}/c^i) = O(\sqrt{|P|}) = O(\sqrt{r})$. Hence $X_1 = O(|\mathcal{L}| \cdot \sqrt{r})$. On the other hand, it is known that the sum of boundary nodes of all (weak) ancestors of the pieces that constitute \mathcal{R} is $O(\sum_{P \in \mathcal{R}} |\partial P|) = O(n/\sqrt{r})$ (cf. the discussion after Corollary 5.1 in [50]) and thus this is also an upper bound on X_2 . So we obtain $X = X_1 + X_2 = O(n/\sqrt{r} + |\mathcal{L}| \cdot \sqrt{r})$. By choosing $r = n/|\mathcal{L}|$, we obtain $X = O(\sqrt{n \cdot |\mathcal{L}|})$ as desired. \square

Let \mathcal{R} be an r -division with few holes of G and $\mathcal{T}(P)$ be a recursive decomposition of $P \in \mathcal{R}$ with the root boundary set to ∂P . For any $v \in V(G) \setminus \partial \mathcal{R}$, let L_v be some leaf containing v in the unique piece $P_v \in \mathcal{R}$ containing v . For any $X \subseteq V(G)$ let us define $\text{cone}_{\mathcal{R}}(X) = \mathcal{R} \cup \bigcup_{v \in X \setminus \partial \mathcal{R}} \text{cone}_{P_v}(L_v)$.

Lemma 8 ([58, 34]). *Let $X \subseteq V(G)$ be non-empty. Then:*

1. For any $u, v \in V(\text{DDG}(\text{cone}_{\mathcal{R}}(X)))$, $\text{dist}_G(u, v) = \text{dist}_{\text{DDG}(\text{cone}_{\mathcal{R}}(X))}(u, v)$.
2. $\sum_{H \in \text{cone}_{\mathcal{R}}(X)} |\partial H| = O\left(n/\sqrt{r} + \min\left(\sqrt{n \cdot |X|}, |X| \cdot \sqrt{r}\right)\right)$.

Proof sketch. The proof is completely analogous to that of items 1 and 2 of Lemma 7. It is enough to glue the individual decompositions $\mathcal{T}(P)$ into a single decomposition $\mathcal{T}'(G)$ such that the root has $O(n/r)$ children instead of just 2: the individual pieces of \mathcal{R} . Then item 1 follows by Lemma 6. Let $Y = \sum_{H \in \text{cone}_{\mathcal{R}}(X)} |\partial H|$. Note that $Y = O(n/\sqrt{r} + |X| \cdot \sqrt{r})$. If we have $r = n/|X| \leq r$, then we can obtain the bound $Y = O(\sqrt{n \cdot |X|})$ in the proof of Lemma 7. Otherwise, $|X| \leq n/r$, so $Y = O(n/\sqrt{r} + |X| \cdot \sqrt{r}) = O(n/\sqrt{r})$. So $Y = O(n/\sqrt{r} + \sqrt{n \cdot |X|})$ in all cases as well. \square

Fully dynamic r -divisions. Many dynamic algorithms for planar graphs maintain r -divisions and useful auxiliary data structures under dynamic updates. The exact set of supported updates to G varies; e.g., [34, 36, 35] support only edge weight changes, [60] assumes embedding-preserving insertions, whereas [37, 62] only assume that the graph G remains planar at all times. We stick to the last, most general setting. The core of the construction behind the following theorem is due to Klein and Subramanian [37, 62].

Let us slightly generalize the definition of an r -division to non-planar graphs. An r -division \mathcal{R} of a (not necessarily planar) graph, for $r \in [1, n]$, is a decomposition of the graph into $O(n/r)$ pieces, each of size $O(r)$, such that each piece P is plane and has $O(\sqrt{r})$ boundary vertices (denoted ∂P), i.e., vertices shared with some other piece of \mathcal{R} . The definitions of $\partial \mathcal{R}$ and r -divisions with few holes remain unchanged. The reason for this generalization will become clear in the following theorem, where we show how to maintain an r -division with few holes for a dynamic planar graph, augmented with some infinite-weight edges that might violate planarity, but not affect any other properties of the graph that we care about.

Theorem 9. Let $G = (V, E)$ be a weighted planar graph that undergoes edge deletions and edge insertions (assumed to preserve the planarity of G). Let $r \in [1, n]$.

There is a data structure maintaining an r -division with few holes \mathcal{R} of some G^+ , where G^+ can be obtained from G by adding infinite-weight edges⁹, such that each piece $P \in \mathcal{R}$:

1. has all its faces except its holes triangular, and
2. is accompanied with some auxiliary data structures that can be constructed in $T(r)$ time given P and use $S(r)$ space.

The data structure uses $O\left(n + \frac{n}{r} \cdot S(r)\right)$ space and can be initialized in $O\left(n + \frac{n}{r} \cdot T(r)\right)$ time. After each edge insertion/deletion, it can be updated in $O(r + T(r))$ worst-case time.

Proof. On initialization, we first connect and triangulate G using infinite-weight edges, thus obtaining G^+ . Then, we compute an r -division \mathcal{R} with few holes of G^+ in linear time [84] and subsequently initialize the auxiliary data structures. Let $h \geq 2$ and $c \geq 8$ be constants such that a single piece of the computed r -division has at most $c\sqrt{r}$ boundary vertices distributed over no more than h holes.

We will guarantee at all times that for any single piece $P \in \mathcal{R}$, $|\partial P| \leq 3c\sqrt{r}$, and there exist at most $3h$ faces of P such that any $v \in \partial P$ lies on one of these faces, called holes of P . Moreover, each edge of G^+ is contained in at most two pieces of \mathcal{R} : this is satisfied initially since the r -division of [84] forms a partition of faces of G^+ .

Suppose that the removal of an edge $e = uv$ is issued to G . We then replace e in each of the at most two pieces P containing it with an infinite-weight edge. If, on the other hand, a new edge $e = uv$ is inserted, we add a new piece P_e consisting of a single edge e to \mathcal{R} . Adding P_e may cause an endpoint of e , say u , to become a boundary vertex of \mathcal{R} . If u was not a boundary vertex before the insertion, it had to be a vertex of a single piece P_u . At this point, since a new boundary vertex emerges in P_u , we might have $\partial P_u > 3c\sqrt{r}$ or ∂P might no longer lie on at most $3h$ faces of P_u . However, there surely exist some $3h + 1$ faces whose vertices include the whole set ∂P_u , and $|\partial P_u| \leq 3c\sqrt{r} + 1$. To fix our invariants, we first compute a cycle separator C of P_u wrt. the boundary vertices of P_u in $O(r)$ time, and replace P_u with two pieces $P_{u,1}, P_{u,2}$ – the two subgraphs of P_u induced by vertices weakly on one side of C . Clearly, the vertices of C become new boundary vertices afterwards. Subsequently, we similarly break each of $P_{u,1}, P_{u,2}$ further into two parts using a cycle separator wrt. the holes of this piece (see [84] for details). Each of the at most four resulting pieces has at most $\frac{2}{3}(3c\sqrt{r} + 1) + 2\sqrt{2}\sqrt{r} \leq (2c + 2\sqrt{2} + 2)\sqrt{r} \leq 3c\sqrt{r}$ boundary vertices, and at most $\frac{2}{3} \cdot (3h + 1) + 1 \leq 2h + 2 \leq 3h$ holes.

Observe that a single edge update can introduce $O(1)$ new pieces of size $O(r)$ in the maintained r -division, and the sizes of the existing pieces do not increase. For each of the affected pieces we recompute the auxiliary data structures in $T(r)$ time. As a result, after $O(n/r)$ updates, there are still $O(n/r)$ pieces, each of size $O(r)$ and with $O(\sqrt{r})$ boundary vertices distributed over $O(1)$ holes of that piece. Consequently, after every $\Omega(n/r)$ updates, we reinitialize \mathcal{R} for the current graph G in $O\left(n + \frac{n}{r} \cdot T(r)\right)$ time. Hence, the amortized time to update \mathcal{R} is $O(r + T(r))$.

Finally, observe that our data structure can be modified in a standard way (see e.g., [5]) to have $O(r + T(r))$ worst-case update time bound instead of just an amortized one. This is possible since our update procedure actually takes $O(r + T(r))$ worst-case time apart from once every $k = \Omega(n/r)$ updates when the whole data structure is rebuilt in $O\left(n + \frac{n}{r} \cdot T(r)\right)$ worst-case time. To this end we apply the time-slicing technique. We use two copies of our data structure switching their roles every $k/2$ updates. One copy is for handling at most $k/2$ updates and answering queries, and another is being gradually reinitialized in chunks of $\Omega(r + T(r))$ time (of either initialization or updates replayed) in the background. \square

Additively weighted Voronoi diagrams. Let G be a directed planar graph of size n with real edge-lengths, and no negative-length cycles. Assume that all faces of G are triangles except, perhaps, a single face f . Let S be the set of vertices that lie on f , called *sites*, i.e., $S = V(f)$. Let us assign to each site $s \in S$ a

⁹Note that G^+ need not be planar.

weight $\omega(s) \in \mathbb{R}_{\geq 0} \cup \{\infty\}$. The additively weighted distance $\text{dist}_G^\omega(s, v)$ between a site $s \in S$ and a vertex $v \in V(G)$ is defined as $\omega(s) + \text{dist}_G(s, v)$.

The additively weighted Voronoi diagram of (S, ω) within G , denoted by $\text{VD}(S, \omega)$, is a partition of $V(G)$ into pairwise disjoint sets, one set $\text{Vor}(s)$ for each site $s \in S$. The set $\text{Vor}(s)$, which is called the Voronoi cell of s , contains all vertices in $V(G)$ that are closer (wrt. dist_G^ω) to s than to any other site in S .

In the following and throughout, whenever we work with Voronoi diagrams we assume that (1) G is strongly connected, (2) shortest paths in G are unique, and (3) additively weighted shortest paths in G are unique, i.e., for each $v \in V(G)$ there is a unique site s minimizing $\text{dist}_G^\omega(s, v)$. Note that these assumptions make the Voronoi cells well-defined and simply connected, and guarantee that they indeed form a partition of $V(G)$. We will explicitly ensure that these requirements are met for G, S and the weight function whenever we define a Voronoi diagram on G .

There is a dual representation $\text{VD}^*(S, \omega)$ of Voronoi diagram $\text{VD}(S, \omega)$ as a tree of constant degree with $O(|S|)$ vertices and edges [50]. An efficient FR-Dijkstra based algorithm for computing $\text{VD}^*(S, \omega)$ was presented by Charalampopoulos et al. [48].

Theorem 10 ([48]). *Suppose that we have at hand a recursive decomposition $\mathcal{T}(G)$ of G , with the only hole of G being f and $S = V(f)$. Further suppose that we have DDG_H computed for each piece $H \in \mathcal{T}(G)$. Then, we can compute $\text{VD}^*(S, \omega)$ in $O(\sqrt{n} \cdot |S| \log^2 n)$ time.*

Remark 11. *The algorithm underlying Theorem 10 implicitly assumes that $\text{Vor}(s)$ is non-empty for all $s \in S$. We will explicitly ensure that this is the case whenever employing the above lemma.*

In a *point location* query for some Voronoi diagram $\text{VD}(S, \omega)$, we are given a vertex $v \in V(G)$ and are requested to find the site $s \in S$ such that $v \in \text{Vor}(s)$ and also the value of $\text{dist}_G^\omega(s, v)$. Gawrychowski et al. [50] showed the following result.

Theorem 12 ([50]). *Suppose that we have at hand an MSSP data structure for G with sources the vertices that lie on face f . Given some dual representation $\text{VD}^*(S, \omega)$, we can preprocess it in $O(|S|)$ time, so that point location queries for $\text{VD}(S, \omega)$ can be answered in $O(\log^2 n)$ time.*

3. Fully Dynamic Single-Source Shortest Paths

In this section we show our single-source exact distance oracle for planar graphs with $O(n^{4/5} \log^2 n)$ update time and $O(\log^2 n)$ query time and thus prove Theorem 3. For simplicity, let us assume for now that G is non-negatively weighted. We will describe how to handle negative edges in Section 3.1; we treat them the same way as Kaplan et al. [35].

Theorem 3. *Let G be a real-weighted planar digraph with a source $s \in V(G)$. There exists an $O(n \log n)$ -space data structure maintaining G under edge insertions, edge deletions, and source changes with worst-case update time $O(n^{4/5} \log^2 n)$ that can compute $\text{dist}_G(s, v)$ for any $v \in V(G)$ in $O(\log^2 n)$ time. The initialization time is $O(n \log^2 n)$.*

The base of our data structure is a dynamic r -division \mathcal{R} with few holes, as given in Theorem 9. Note that in our shortest paths problem, indeed adding infinite-weight edges to G does no harm. Hence, in the following we work with the graph G^+ from Theorem 9 when computing distances, but identify it, without loss of generality, with our original graph G .

For technical reasons, related with the uniqueness of shortest paths, we would like to avoid dealing with infinite weights in some of our data structures handling individual pieces. In the real-weighted fully dynamic setting, however, we cannot fix a sufficiently large finite number, larger than all edge weights that will ever appear in the future graph G , beforehand. Instead, we do the following. For each $P \in \mathcal{R}$, let M_P be a sufficiently large finite number, e.g., larger than the sum of finite edge weights in P . Consider M_P to be an auxiliary data structure of P as in Theorem 9. We will use M_P to simulate infinite edge weights in P , and also for detecting paths non-existent in the original graph G (but having infinite weight in $G^+ \cap P$). As a result, below we assume that each infinite weight in P (or any auxiliary data structure related to

P) is replaced by M_P in all the computations performed locally on the piece P , whereas globally (when performing some computation for many pieces at once, like the shortest paths algorithm of Lemma 5) we treat all edge or path weights in P that are at least M_P as infinite.

For each piece $P \in \mathcal{R}$ we store the following additional data structures.

- We store the recursive decomposition $\mathcal{T}(P)$ with the initial boundary set to ∂P , and also DDGs for all the pieces $H \in \mathcal{T}(P)$.
- For each hole h of P , let P_h be the piece P after applying the following standard augmentations. First, P is extended into a graph P'_h using $O(r)$ vertices and edges of weight M_P embedded inside either the piece or other (than h) holes of P that would make P strongly connected and triangulated (except for the hole h) without changing the distance between any pair of reachable vertices in P . The graph P_h is in turn obtained from P'_h by changing P'_h 's edge weights into $O(1)$ -size vectors as described in [88] so that there is a *unique* shortest path (wrt. the lexicographical order on path weights, defined as the coordinate-wise sum of the path's individual edge weights) between any $u, v \in V(P_h)$ with cost of the form $(\text{dist}_{P'_h}(u, v), \cdot)$. As proven in [88], one can compute P_h from P deterministically in linear time. The $O(1)$ -size vector weights, in turn, can be easily packed into usual positive single-number weights.

For each P_h , we store an MSSP data structure initialized for the hole h . Recall that an MSSP data structure can be computed in $O(r \log r)$ time. Moreover, we store a recursive decomposition $\mathcal{T}(P_h)$ of P_h with the boundary ∂P_h of the root piece set to $\partial P \cap V(h)$ of size $O(\sqrt{r})$. For each node (piece) $H \in \mathcal{T}(P_h)$, we also store DDG_H . Since the sum of sizes of all the pieces of $\mathcal{T}(P_h)$ is $O(r \log r)$, computing all these dense distance graphs takes $O(r \log^2 r)$ time (see Section 2).

Note that computing piecewise auxiliary data structures defined so far takes $O(r \log^2 r)$ time. So, by Theorem 9, they can be updated in $O(r \log^2 r)$ worst-case time after G undergoes an update.

We also perform the following computations. We compute the values $\text{dist}_G(s, b)$ for all $b \in \partial R$ in $O(n/\sqrt{r} \cdot \log^2 n)$ time by running the single-source shortest paths algorithm of Lemma 5 (FR-Dijkstra) on the graph $\text{DDG}(\text{cone}_{\mathcal{R}}(s))$ – the correctness follows by Lemma 8. We also compute $\text{dist}_{P_s}(s, u)$ for all $u \in P_s$, where P_s is an arbitrary piece containing s using Dijkstra's algorithm in $O(r \log r)$ time.

Consider a query $\text{dist}_G(s, v)$. If the shortest $s \rightarrow v$ path in G does not go through $\partial \mathcal{R}$, then it is fully contained in P_s and therefore $v \in P_s$ and $\text{dist}_G(s, v) = \text{dist}_{P_s}(s, v)$, i.e., we have $\text{dist}_G(s, v)$ already computed. Otherwise, let P_v be an arbitrary piece containing v . Observe that we have $\text{dist}_G(s, v) = \min_{b \in \partial P_v} \{\text{dist}_G(s, b) + \text{dist}_{P_v}(b, v)\}$ where the minimizer b corresponds to a boundary vertex on some shortest $s \rightarrow v$ path in G .

Let us first consider a preliminary solution. Recall that we have already computed $\text{dist}_G(s, b)$ for all $b \in \partial P_v \subseteq \partial R$. Then, we can find the minimum by iterating over all $O(\sqrt{r})$ vertices $b \in \partial P_v$, computing $\text{dist}_{P_v}(b, v)$ using one of the MSSP data structures stored for P_v in $O(\log n)$ time. This yields the following result, as updates are processed in $O(r \log^2 r + n/\sqrt{r} \cdot \log^2 n)$ time, queries are answered in $O(\sqrt{r} \log n)$ time, while the space usage per piece is $O(r \log r)$.

Proposition 13. *Let G be a real-weighted planar digraph with a source $s \in V(G)$. For any $x \in [1, n^{1/3}]$, there exists an $O(n \log n)$ -space data structure maintaining G under edge insertions, edge deletions, and source changes with worst-case update time $O(n \log^2 n/x)$ that can compute $\text{dist}_G(s, v)$ for any $v \in V(G)$ in $O(x \log n)$ time. The initialization time is $O(n \log^2 n)$.*

Remark 14. *The update time in Theorem 3 dominates the update time of the data structure underlying the above proposition for $x = \Omega(n^{1/5})$. Further, note that the data structure underlying Proposition 13 is essentially the same as Klein's all-pairs dynamic distance oracle for planar graphs [36], with part of the query procedure offloaded to the update procedure.*

Next, we will employ the efficient construction algorithm of dual Voronoi diagrams of [48] and the efficient point location mechanism of [50], in order to be able to compute the minimizer $b \in \partial P_v$ in $O(\log^2 n)$ time, while keeping the update time strongly sublinear.

After the initialization and each update, once \mathcal{R} and all auxiliary data structures are updated, and we have run FR-Dijkstra from s , we compute for each $P \in \mathcal{R}$ a point location data structure.

Lemma 15. *Given a weight function $\omega : \partial P \rightarrow \mathbb{R} \cup \{\infty\}$, one can compute in $O(r^{3/4} \log^2 r)$ time a data structure $L(P)$ answering the following queries in $O(\log^2 r)$ time: given any $v \in V(P)$, compute the value $\min_{b \in \partial P} \{\omega(b) + \text{dist}_P(b, v)\}$ along with the minimizer b .*

Proof. Since ∂P is a union of $O(1)$ sets ∂P_h , we can compute the desired minimum over each ∂P_h separately and then take the minimum over all h .

Let us first note that negative values of ω are not a problem. We can turn negative weights into non-negative by adding some common large value to the weights of all sites. We can thus suppose wlog. that all values of ω are non-negative.

If all the weights are infinite, the queries can be answered trivially in $O(1)$ time. So in the following assume that there is at least one site whose weight is finite.

Let $S = \partial P_h = \{s_1, \dots, s_k\}$ be the set of sites. Recall that P_h has finite positive real weights, is strongly connected, has unique shortest paths, and has a single face h that is possibly non-triangular that contains all the sites S .

First, let us take care of the possibility that there are empty Voronoi cells and ensure that additively weighted shortest paths are unique. Consider a site s that belongs to the Voronoi cell of a different site s' . We will substitute $\omega(s)$ by $\text{dist}_{P_h}^\omega(s', s)$ and consider s as a proxy for s' , i.e., whenever we find that some vertex v belongs to the Voronoi cell of s , we will know that in effect v belongs to the Voronoi cell of s' and return s' instead. We do this as follows. Let a be an artificial vertex with k outgoing edges: for $i \in [1, k]$, edge $e_i = as_i$ with weight $\omega(s_i)$. We run FR-Dijkstra on $\{\text{DDG}_{P_h}, e_1, \dots, e_k\}$ with source a in $O(\sqrt{r} \log^2 r)$ time. Since P_h is strongly connected and contains only finite-weight edges, and at least one additive weight is finite, all computed distances are finite. For each i , denote the computed distance from a to s_i by $d(s_i)$. We then consider a new weight function $\hat{\omega}(s_i) = \min\{\omega(s_i), d(s_i)\}$. A site s with $\omega(s) > d(s)$ is then a proxy for its ancestor at depth one in the computed shortest paths tree. Finally, we will break ties wrt. $\text{dist}_{P_h}^{\hat{\omega}}$ in favor of sites with larger additive weights. Let $r(i) = |\{j : j \neq i, \omega(j) > \omega(i)\}|$. We will break ties by considering $(\hat{\omega}(s_i), r(i))$ instead of $\hat{\omega}(s_i)$ as the weight of site s_i , adding a second coordinate to each edge weight in P_h , set to 0, and comparing additively weighted distances lexicographically. Clearly, this extension does not break any of the properties of P_h . Since all edge-weights in P_h are positive, this guarantees that each site belongs to its own Voronoi cell and that for each $v \in V(P_h)$ there is a unique site s minimizing $\text{dist}_{P_h}^{\hat{\omega}}(s, v)$.

Therefore, we can invoke Theorem 10 to construct the dual representation $\text{VD}^*(S, \omega)$ of the Voronoi diagram $\text{VD}(S, \omega)$. This requires $O(\sqrt{r|\partial P_h|} \log^2 n) = O(r^{3/4} \log^2 n)$ time.

Then, we construct the point location data structure of Theorem 12 for $\text{VD}^*(S, \omega)$ in $O(|\partial P_h|) = O(\sqrt{r})$ time. Note that we have an MSSP data structure for P_h for hole h and hence point location queries can then be answered in $O(\log^2 n)$ time. \square

We invoke Lemma 15 with weight function $\omega(b) := \text{dist}_G(s, b)$ in order to construct $L(P)$ for each $P \in \mathcal{R}$. This requires $O(n/r^{1/4} \cdot \log^2 n)$ time in total.

Now, we can compute $\text{dist}_G(s, v)$ for a query vertex v as follows. If the shortest $s \rightarrow v$ path in G does not go through $\partial \mathcal{R}$, then as explained before it has already been computed. Otherwise, recall that we have $\text{dist}_G(s, v) = \min_{b \in \partial P_v} \{\text{dist}_G(s, b) + \text{dist}_{P_v}(b, v)\}$. This case can be reduced to a single query to the data structure $L(P_v)$. Each query can be thus answered in $O(\log^2 n)$ time.

The worst-case update time is $O(r \log^2 r + \frac{n}{r^{1/4}} \log^2 r)$. By setting $r = n^{4/5}$ we get $O(n^{4/5} \log^2 n)$ worst-case update time. Since the space usage per piece is $O(r \log r)$, we need $O(n \log n)$ space.

Remark 16. *Our data structure can be extended to report, following the computation of $\text{dist}_G(s, v)$, a shortest $s \rightarrow v$ path Q in time nearly linear in the number of edges of Q . This follows easily by the fact that the MSSP data structure [36] can report shortest paths efficiently (see e.g., [35] for details). Therefore, we can efficiently expand the used edges of dense distance graphs and the shortest $b \rightarrow v$ path into actual edges.*

Remark 17. *The source-change operation in Theorem 3 can be easily simulated using a single edge insertion and a single edge deletion. For this, it suffices to add a fixed super-source s' to G , with a single outgoing edge to the actual source s of G .*

3.1. Handling Negative Edges

Recall that $p : V(H) \rightarrow \mathbb{R}$ is called a feasible price function of H if $\text{dist}_H(u, v) + p(u) - p(v) \geq 0$ for all $u, v \in V(H)$. A feasible price function is guaranteed to exist if H contains no negative-cost cycle. It is well-known that, provided that a graph H is strongly connected, a vector of distances from any vertex of H constitutes a feasible price function of H .

As in the fully dynamic all-pairs algorithm of [35], we maintain functions $\phi : \partial\mathcal{R} \rightarrow \mathbb{R}$ and $\phi_P : \partial P \rightarrow \mathbb{R}$, where $P \in \mathcal{R}$, such that ϕ is some feasible price function of G restricted to $\partial\mathcal{R}$, and each ϕ_P is a feasible price function of P .

Since single-source shortest paths in planar graphs with negative weights can be computed in $O(n \log^2 n)$ time [89, 90], each ϕ_P can be seen as an accompanying data structure of piece P computable in $O(r \log^2 r)$ time and maintained by the fully dynamic r -division algorithm. The functions ϕ_P allow to treat individual graphs P and P_h as non-negatively weighted when computing all the needed DDGs and MSSP data structures, and also point location data structures $L(P)$.

It is known [35, 59] that the FR-Dijkstra algorithm (as in Lemma 5) can handle negative weights in $\text{DDG}(\mathcal{H})$ with no asymptotic overhead if a feasible price function on $\text{DDG}(\mathcal{H})$ is provided. Therefore, we would like to have a feasible price function on $\text{DDG}(\text{cone}_{\mathcal{R}}(s))$ to compute distances from s in $\text{DDG}(\text{cone}_{\mathcal{R}}(s))$ needed by the update algorithm. We can extend ϕ from $\partial\mathcal{R}$ to all vertices in $\text{DDG}(\text{cone}_{\mathcal{R}}(s))$ as follows. Note that all pieces in $\text{cone}_{\mathcal{R}}(s)$ except of those in \mathcal{R} have their parents also in $\text{cone}_{\mathcal{R}}(s)$. We call those pieces $H \in \text{cone}_{\mathcal{R}}(s)$ for which we know the value of ϕ on all of ∂H *processed*. Initially, only the pieces $P \in \mathcal{R}$ are processed by the definition of ϕ . While there are still unprocessed pieces, we take any unprocessed piece H whose parent $A \in \mathcal{T}(P)$ has already been processed. Let H' be the sibling of H in $\mathcal{T}(P)$. Observe that ϕ_P is a feasible price function of A as well. We extend ϕ to boundary vertices of H, H' by computing shortest paths on $\text{DDG}(\{H, H'\})$ from vertices ∂A , with the initial distance to each $v \in \partial A$ set to $\phi(v)$, and using FR-Dijkstra (Lemma 5) with price function ϕ_P . This way, only the initial distances of ∂A are possibly negative from the point of view of FR-Dijkstra. This does not constitute a problem for either Dijkstra's algorithm or FR-Dijkstra though (see [35]; one can treat the initial distances as weights of edges going out of a super-source; these weights can be all increased by the same large value to be made positive). One can show from the definition of ϕ that the values of ϕ on ∂A will not be altered and the computed distances form a feasible price function on $\partial H \cup \partial H'$ in G . Hence, given that $\partial A \subseteq \partial H \cup \partial H'$, we can process the children H, H' of a processed piece A in $O((|\partial H| + |\partial H'|) \log^2 n)$ time. Summing over all pieces, we obtain by Lemma 8 that extending ϕ to all $V(\text{DDG}(\text{cone}_{\mathcal{R}}(s)))$ takes $O(n/\sqrt{r} \log^2 n)$ time. This cost is therefore negligible.

The “global” price function $\phi : \partial\mathcal{R} \rightarrow \mathbb{R}$ is recomputed from scratch in $O(n \log^2 n)$ time (by running the algorithm of [89] on the entire G) every time \mathcal{R} is reinitialized (i.e., once per each $O(n/r)$ updates). Note that an edge deletion or weight increase cannot break the feasibility of ϕ . We might need to recompute ϕ only upon insertion or weight decrease of some edge uv . As shown by Kaplan et al. [35], the new “global” price function ϕ' (or a negative cycle) can be found by computing distances from v to $\partial\mathcal{R} \cup \{u\}$ in $\text{DDG}(\text{cone}_{\mathcal{R}}(u, v))$ (before applying the insertion) using FR-Dijkstra and the old price function ϕ . This can be done in $O(n/\sqrt{r} \log^2 n)$ time by first extending the old ϕ to $V(\text{DDG}(\text{cone}_{\mathcal{R}}(u, v)))$ as described above and then running the single-source shortest paths algorithm of Lemma 5.

3.2. A Dynamic Closest Facility Data Structure

We can generalize the dynamic single-source shortest paths data structure as follows. Suppose we replace a single source vertex s with a set of *facilities* $F \subseteq V$. Given F , for a query vertex v we would like to compute $\min_{f \in F} \{\text{dist}_G(f, v)\}$, and also possibly $f \in F$ minimizing this expression. A dynamic update would consist of either an edge update or changing the set F . In other words, such a problem can be seen as dynamic point location in a Voronoi diagram wrt. F , where each update either changes the graph or resets the Voronoi diagram of interest.

In this setting, we consider the following simple generalization of the single-source data structure. Let the update procedure first compute the distances $d(b) = \min_{f \in F} \{\text{dist}_G(f, b)\}$ for all $b \in \partial\mathcal{R}$. Note that by Lemma 8 this can be achieved by computing single-source shortest paths in the graph $\text{DDG}(\mathcal{D}_F)$, where

$\mathcal{D}_F = \text{cone}_{\mathcal{R}}(F)$, extended with 0-weight edges sf , where s is an auxiliary super-source. By Lemma 5 this can be done in $O((\sqrt{n|F|} + n/\sqrt{r} + |F|)\log^2 n)$ time. By using weights $\omega := d$ in the individual point-location data structures $L(P)$, $P \in \mathcal{R}$, a single point location query on $L(P_v)$ (recall that P_v is some piece containing v) would compute the desired closest facility f_v minimizing $\text{dist}_G(f_v, v)$ unless the sought (weighted) shortest $f_v \rightarrow v$ does not go through a boundary vertex of \mathcal{R} . We could in principle handle such paths by proceeding as in the single-source case and computing shortest paths naively in each piece containing a facility. However, this could take time $\Omega(r \cdot \min(n/r, |F|))$. Recall that the other components of our update procedure take time $\tilde{O}(r + \sqrt{n|F|} + n/r^{1/4})$ which is $\tilde{O}(\sqrt{n|F|} + n^{4/5})$ if we again set $r = n^{4/5}$. The update time would then be linear in n even for moderately large facility set sizes, e.g., $|F| = \Omega(n^{1/5})$.

To improve upon this simple approach, we proceed as follows. Let (ρ_m, \dots, ρ_1) be such a sequence of integers that $\rho_m = r$, $\rho_1 = O(1)$ and $\rho_{i+1}/\rho_i = 2$ for all $i < m$. For each $P \in \mathcal{R}$ we store a recursive (ρ_m, \dots, ρ_1) -division consisting of pieces of $\mathcal{T}(P)$ (cf. Section 2). Let $\mathcal{R}_{P,i}$ be the ρ_i -division of P . All $\mathcal{R}_{P,i}$ can be computed in linear time given $\mathcal{T}(P)$ [84]. Note that \mathcal{R}_i , defined as the union of $\mathcal{R}_{P,i}$ over all pieces $P \in \mathcal{R}$, actually forms an ρ_i -division with few holes of the entire graph G . In particular, we have $\mathcal{R}_m = \mathcal{R}$.

We store the extended pieces Q_h (recall how we obtained extended pieces P_h with unique shortest paths from P in the single-source case) plus their recursive decompositions $\mathcal{T}(Q_h)$, DDGs, and an MSSP data structure for all pieces Q of *all* $\mathcal{T}(P)$ instead of just the pieces of $\mathcal{R}_m = \mathcal{R}$ as we did in the single-source case. However, we stress that these auxiliary components for a piece $Q \subseteq P$ where $P \in \mathcal{R}$, are counted as accompanying data structures of the piece P . So, we compute $O(1)$ fresh recursive decompositions $\mathcal{T}(Q_h)$ for each piece $Q \in \mathcal{T}(P)$ – computing each of them takes $O(|Q|\log^2 n)$ time. As a result, by the bound $\sum_{Q \in \mathcal{T}(P)} |Q| = O(|P| \log |P|)$, the time to compute accompanying data structures of piece P increases to $O(r \log^3 n)$.

Given the set of facilities F , let j be such that $\rho_j = \Theta\left(\min\left(\frac{n}{|F|}, r\right)\right)$. Redefine $\mathcal{D}_F = \text{cone}_{\mathcal{R}_j}(F)$. Again, let us compute distances $d(b) = \min_{f \in F} \{\text{dist}_G(f, b)\}$ (and the closest facilities) for all $b \in \bigcup_{H \in \mathcal{D}_F} \partial H$ using FR-Dijkstra on $\text{DDG}(\mathcal{D}_F)$ extended with a super-source s and auxiliary edges sf , $f \in F$. This takes $O\left(\left(\sqrt{n|F|} + n/\sqrt{\rho_j}\right)\log^2 n\right) = O\left(\sqrt{n|F|}\log^2 n\right)$ time by Lemmas 5 and 8.

It only remains to show how to handle computation of closest facilities for $v \in V \setminus \bigcup_{H \in \mathcal{D}_F} \partial H$. Recall that the pieces \mathcal{D}_F cover the entire G . Let H_v be the lowest piece \mathcal{D}_F that contains v . By the definition of \mathcal{D}_F , no facility f is an internal (non-boundary) vertex of H_v . Consequently, by Lemma 15, the closest facility to v can be found in $O(\log^2 n)$ time using a single query to the data structure $L(H_v)$. For this to be possible, upon update we need to build the data structures $L(H)$ for all $H \in \mathcal{D}_F$. By Lemma 15, this takes $O\left(\sum_{H \in \mathcal{D}_F} \sqrt{|H| \cdot |\partial H|} \log^2 n\right)$ time. Let us now bound this sum. First, let us consider the sum restricted to the pieces $H \in \mathcal{D}_F \cap \mathcal{R}_j$, i.e., the pieces of the ρ_j -division \mathcal{R}_j of G . Since $\rho_j = \Omega(n/|F|)$ or $\rho_j = \Omega(r)$ we get:

$$O\left(\sum_{H \in \mathcal{D}_F \cap \mathcal{R}_j} \sqrt{|H| \cdot |\partial H|} \log^2 n\right) = O\left(\frac{n}{\rho_j} \cdot \rho_j^{3/4} \log^2 n\right) = O\left(\left(n^{3/4} \cdot |F|^{1/4} + \frac{n}{r^{1/4}}\right) \log^2 n\right).$$

On the other hand, if $H \notin \mathcal{D}_F \cap \mathcal{R}_j$, then $H \in \text{cone}_{P_f}(L_f)$, where $f \in F \setminus \partial \mathcal{R}_j$, P_f is the unique piece of \mathcal{R}_j containing f , and L_f is some leaf of $\mathcal{T}(P_f)$ containing f . For a fixed f , by the definition of $\text{cone}_{P_f}(L_f)$, there are $O(\log n)$ pieces H satisfying this, at most two per each level i of $\mathcal{T}(P_f)$. Hence, the sum of $\sqrt{|H| \cdot |\partial H|}$ over such pieces can be bounded by $\sum_{i=0}^{\infty} \sqrt{|P_f|} \cdot \sqrt{|P_f|/c^i} = O(\rho_j^{3/4})$. Summing over all f , and using $\rho_j = O(n/|F|)$, we get

$$O\left(\sum_{H \in \mathcal{D}_F \setminus \mathcal{R}_j} \sqrt{|H| \cdot |\partial H|} \log^2 n\right) = O\left(|F| \cdot \rho_j^{3/4} \log^2 n\right) = O\left(n^{3/4} \cdot |F|^{1/4} \log^2 n\right).$$

To conclude, the update time is $O\left(\left(n^{3/4} \cdot |F|^{1/4} + \frac{n}{r^{1/4}}\right) \log^2 n + r \log^3 n\right)$. By setting $r = (n/\log n)^{4/5}$ we obtain the following theorem.

Theorem 18. *Let G be a real-weighted planar digraph with a set $F \subseteq V$ of facilities. There exists an $O(n \log^2 n)$ -space data structure maintaining G under edge insertions, edge deletions, and changes of the facilities set F with $O\left((n^{3/4} \cdot |F|^{1/4} + n^{4/5}) \log^{11/5} n\right)$ worst-case update time that can compute $\min_{f \in F} \{\text{dist}_G(f, v)\}$ along with the respective closest facility of v for any $v \in V(G)$ in $O(\log^2 n)$ time. The initialization time is $O(n \log^3 n)$.*

4. Fully Dynamic Strongly Connected Components

In this section we show that a strategy similar to that of Section 3 can be used to obtain a fully dynamic strong-connectivity algorithm. Again, we maintain an r -division \mathcal{R} and some auxiliary data structures for all the individual pieces. Formally, using a dynamic r -division as in Theorem 9 may require introducing new infinite-weight edges to G which in turn may change the reachability relation in G . We circumvent this problem by setting the weights of the original edges of G to 0, and all auxiliary edges to plus infinity (simulated in the implementation by sufficiently large values M_P inside individual pieces, as in Section 3). This way, u can reach v in G if and only if $\text{dist}_G(u, v) = 0$, and otherwise $\text{dist}_G(u, v) = \infty$. All known properties of reachability in plane graphs also extend to reachability using 0-weight paths (assuming non-negative weights). In the following, whenever we say that v is reachable from u , or there exists a $u \rightarrow v$ path, we formally mean $\text{dist}_G(u, v) = 0$.

As in Section 3, for each piece of \mathcal{R} we store a recursive decomposition, dense distance graphs and MSSP data structures. All these data structures are also maintained for pieces of \mathcal{R} with all edges reversed – for a piece P we call this graph the *reverse* of P and denote it by P^{rev} .

Another ingredient is a collection of *reachability certificates* X_P for all the pieces, as defined in the following lemma due to Subramanian [62], slightly adjusted to certify 0-weight paths.

Lemma 19 ([62]). *Let $P \in \mathcal{R}$ be a (non-negatively weighted) piece. There exists a directed graph X_P , where $\partial P \subseteq V(X_P)$, of size $O(\sqrt{r} \log r)$ satisfying the following property: for any $u, v \in \partial P$, $\text{dist}_P(u, v) = 0$ if and only if there exists a $u \rightarrow v$ path in X_P . The graph X_P can be computed in $O(r \log r)$ time.*

We include the reachability certificate in the set of auxiliary piecewise data structures. Since reachability certificates can be computed in $O(r \log r)$ time, maintaining them does not incur any additional asymptotic cost. The following lemma is a direct consequence of Lemma 19.

Lemma 20. *For any $u, v \in \partial \mathcal{R}$, u can reach v in G if and only if u can reach v in $X = \bigcup_{P \in \mathcal{R}} X_P$.*

Proof. Let $u, v \in \partial \mathcal{R}$. Since each X_P certifies the reachability between ∂P in P , clearly a $u \rightarrow v$ path in X implies existence of a $u \rightarrow v$ path in G . Now suppose there is a $u \rightarrow v$ path Q in G . Split Q into maximal subpaths Q_1, \dots, Q_k , such that each Q_i is fully contained in a single piece $P_i \in \mathcal{R}$. For each i , the endpoints a, b of Q_i are contained in ∂P_i and hence there exists a $a \rightarrow b$ path in $X_P \subseteq X$. Consequently, there exist a $u \rightarrow v$ path in X . \square

To handle an edge update, after \mathcal{R} and auxiliary data structures are updated, we compute the strongly connected components of X (defined as in Lemma 20) in $O(|X|) = O(n/\sqrt{r})$ time using any classical linear-time algorithm. For any $b \in \partial \mathcal{R}$, let $s_X(b)$ denote an integer identifier of b 's strongly connected component in X . By additionally sorting the SCCs of X topologically we can further assume that s_X satisfies the following property: if $a, b \in \partial \mathcal{R}$ are not strongly connected, but a can reach b in X then $s_X(a) < s_X(b)$. By Lemma 20, for $a, b \in \partial \mathcal{R}$, we have $s_X(a) = s_X(b)$ if and only if a and b are strongly connected in G ; moreover, if a can reach b in G , then $s_X(a) \leq s_X(b)$.

We also define and maintain similar SCC-identifiers s_P for individual pieces P , i.e., for $u, v \in V(P)$, $s_P(u) = s_P(v)$ implies u, v are strongly connected in P , whereas $s_P(u) < s_P(v)$ implies there is no $v \rightarrow u$ path in P . Clearly, the identifiers s_P can be recomputed in $O(r)$ time given P , so we also include them into the set of auxiliary per-piece data structures.

For any $Q \in \{X\} \cup \mathcal{R}$, let S_Q be the set of used identifiers of the form $s_Q(\cdot)$. Observe that we can easily guarantee that the sets S_Q are pairwise disjoint, e.g., by using disjoint integer ranges for different sets S_Q .

The final component of our data structure, is, again a collection of per-piece point location data structures. For each $P \in \mathcal{R}$, we have two point location data structures $L(P^{\text{rev}})$ and $L(P)$ of Lemma 15. After each edge update, $L(P^{\text{rev}})$ is computed for P^{rev} with weight function $\omega = s_X$. On the other hand, $L(P)$ is initialized with weight function $\omega = -s_X$. As in Section 3, all these point location data structures are recomputed in $O(n/r^{1/4} \cdot \log^2 n)$ time (over all pieces).

We now describe how our data structure handles a query for an SCC identifier of a vertex v . The returned identifier always comes from the set $\bigcup_{Q \in \{X\} \cup \mathcal{R}} S_Q$. Let P_v be some piece containing v . Let s_{\min} be the value computed by $L(P_v^{\text{rev}})$ for vertex v . Let s_{\max} be *minus* the value computed by $L(P_v)$ for v . If either of s_{\min}, s_{\max} equals $\pm\infty$ or $s_{\min} \neq s_{\max}$ holds, we return $s_{P_v}(v)$. Otherwise, we return $s_{\min} \in S_X$. The following lemma establishes the correctness of this query procedure.

Lemma 21. *Let $u, v \in V(G)$ and let s_u, s_v be the respective identifiers returned by the query procedure. Then, $s_u = s_v$ if and only if u and v are strongly connected in G .*

Proof. Suppose $s_u = s_v$. If $s_u \in S_P$ for some piece P , then $s_u = s_v$ implies that u and v are strongly connected in P and thus also in G . So suppose $s_v \in S_X$. Take any $b \in \partial\mathcal{R}$ such that $s_v = s_X(b)$. We now prove that v and b are strongly connected in G . Similarly we prove that u and b are strongly connected in G . By transitivity it will follow that u and v are indeed strongly connected.

Let P_v, s_{\min}, s_{\max} be defined as in the query procedure's description. Recall that $s_v \in S_X$ implies that s_{\min}, s_{\max} are finite and $s_v = s_{\min} = s_{\max}$. Since all edges of P_v have weight 0, and s_{\max} is finite, s_{\max} in fact represents the maximum value $s_X(a)$ among those $a \in \partial P_v$ such that a path $a \rightarrow v$ exists in P_v . Similarly, observe that s_{\min} represents the minimum value $s_X(c)$ among those $c \in \partial P_v$ such that a path $c \rightarrow v$ exists in P_v^{rev} , i.e., such that a path $v \rightarrow c$ exists in P_v . Let us denote by a and c the respective vertices of ∂P_v attaining the maximum and minimum values of s_X . Since $s_X(a) = s_X(c)$, there exists a path $c \rightarrow a$ in G . However, by the definition of a and c , paths $a \rightarrow v$ and $v \rightarrow c$ also exist in G , and hence a, v and c are strongly connected in G . Since a is clearly strongly connected to b by $s_X(a) = s_X(b)$, indeed v and b are strongly connected in G .

Now let us move to proving the " \Leftarrow " direction. Suppose u and v are strongly connected in G . First consider the case when there exists some vertex $b \in \partial\mathcal{R}$ located in the same strongly connected component of G as u and v . In this case we prove that $s_v = s_X(b)$. An analogous proof that $s_u = s_X(b)$ will establish $s_u = s_v$. Since v and b are strongly connected, there exist some paths $Q_1 = v \rightarrow b$ and $Q_2 = b \rightarrow v$ in G . Let v_1 be the first vertex on Q_1 such that $v_1 \in \partial P_v$ – note that v_1 necessarily exists since $b \in \partial\mathcal{R}$. Similarly set v_2 to be the last vertex on Q_2 such that $v_2 \in \partial P_v$. Observe that the subpaths $v \rightarrow v_1$ and $v_2 \rightarrow v$ of Q_1 and Q_2 respectively lie entirely inside P_v . Hence, s_{\min} and s_{\max} are finite and we have $s_{\max} \geq s_X(v_2)$ and $s_{\min} \leq s_X(v_1)$. Recall that there exists a walk $v \rightarrow v_1 \rightarrow b \rightarrow v_2 \rightarrow v$, so in fact v_1, v_2, b are strongly connected, i.e., $s_X(v_1) = s_X(v_2) = s_X(b)$. Thus, we obtain $s_{\min} \leq s_X(b) \leq s_{\max}$.

On the other hand, let $a \in \partial P_v$ be such that a path $a \rightarrow v$ exists in P_v and $s_{\max} = s_X(a)$ (a exists by $s_{\max} \neq \pm\infty$). Similarly, let $c \in \partial P_v$ be such that a path $v \rightarrow c$ exists in P_v and $s_{\min} = s_X(c)$. Since a path $a \rightarrow c$ through v exists in P_v (so also in G), we have that $s_{\max} \leq s_{\min}$ by the fact that the identifiers S_X respect the topological order of the SCCs of X . Recall that we have already proved $s_{\min} \leq s_X(b) \leq s_{\max}$ so in fact we have $s_{\min} = s_{\max} = s_X(b)$, and consequently $s_v = s_X(b)$.

Finally, suppose there is no vertex of $\partial\mathcal{R}$ in the SCC of G containing u and v . First, this implies that $u, v \notin \partial\mathcal{R}$ and all $u \rightarrow v$ and $v \rightarrow u$ paths are contained in a single, unique piece P . This implies that $s_P(u) = s_P(v)$. Hence it is sufficient to prove $s_u = s_P(u)$ and $s_v = s_P(v)$. We prove the latter equality; proving the former is analogous. Recall that s_v is *not* set to $s_P(v)$ only if both s_{\min}, s_{\max} are finite and $s_{\min} = s_{\max}$. This can only happen if there exists vertices $a, c \in \partial P$ such that a can reach v in P , v can reach c in P and $s_X(c) = s_{\min} = s_{\max} = s_X(a)$, i.e., a and c are strongly connected in G . But this implies that a, c and v are strongly connected in G , which contradicts the fact that the SCC of v in G does not contain vertices of $\partial\mathcal{R}$. \square

The running time analyses of both the update and query procedures are identical to the analyses of Section 3. Hence, we have proved the following theorem.

Theorem 4. *Let G be a planar digraph. There exists an $O(n \log n)$ -space data structure maintaining G under edge insertions and deletions with $O(n^{4/5} \log^2 n)$ worst-case update time that can compute the identifier of the strongly connected component of any $v \in V(G)$ in $O(\log^2 n)$ time. The initialization time is $O(n \log^2 n)$.*

Acknowledgments

Panagiotis Charalampopoulos was supported by ERC Starting Grant 677651 TOTAL, Israel Science Foundation grant ISF no. 592/17, a Studentship from the Faculty of Natural and Mathematical Sciences at King’s College London, and an A. G. Leventis Foundation Educational Grant.

Adam Karczmarz was supported by ERC Consolidator Grant 772346 TUGBOAT, the Polish National Science Centre 2018/29/N/ST6/00757 grant, and by the Foundation for Polish Science (FNP) via the START programme.

References

- [1] P. Charalampopoulos, A. Karczmarz, Single-source shortest paths and strong connectivity in dynamic planar graphs, in: 28th Annual European Symposium on Algorithms, ESA 2020, 2020, pp. 31:1–31:23. doi:10.4230/LIPIcs.ESA.2020.31.
- [2] C. Demetrescu, G. F. Italiano, A new approach to dynamic all pairs shortest paths, *J. ACM* 51 (6) (2004) 968–992. doi:10.1145/1039488.1039492.
- [3] M. Thorup, Fully-dynamic all-pairs shortest paths: Faster and allowing negative cycles, in: Algorithm Theory - SWAT 2004, 9th Scandinavian Workshop on Algorithm Theory, Proceedings, 2004, pp. 384–396. doi:10.1007/978-3-540-27810-8_33.
- [4] D. B. Johnson, Efficient algorithms for shortest paths in sparse networks, *J. ACM* 24 (1) (1977) 1–13. doi:10.1145/321992.321993.
- [5] I. Abraham, S. Chechik, S. Krininger, Fully dynamic all-pairs shortest paths with worst-case update-time revisited, in: Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, 2017, pp. 440–452. doi:10.1137/1.9781611974782.28.
- [6] M. P. Gutenberg, C. Wulff-Nilsen, Fully-dynamic all-pairs shortest paths: Improved worst-case time and space bounds, in: Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, SIAM, 2020, pp. 2562–2574. doi:10.1137/1.9781611975994.156.
- [7] M. Thorup, Worst-case update times for fully-dynamic all-pairs shortest paths, in: Proceedings of the 37th Annual ACM Symposium on Theory of Computing, STOC 2005, 2005, pp. 112–119. doi:10.1145/1060590.1060607.
- [8] G. Ausiello, G. F. Italiano, A. Marchetti-Spaccamela, U. Nanni, Incremental algorithms for minimal length paths, in: Proceedings of the First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 1990, 1990, pp. 12–21. URL <http://dl.acm.org/citation.cfm?id=320176.320178>
- [9] S. Baswana, R. Hariharan, S. Sen, Improved decremental algorithms for maintaining transitive closure and all-pairs shortest paths, *J. Algorithms* 62 (2) (2007) 74–92. doi:10.1016/j.jalgor.2004.08.004.
- [10] J. Ewald, V. Fredslund-Hansen, M. P. Gutenberg, C. Wulff-Nilsen, Decremental APSP in unweighted digraphs versus an adaptive adversary, in: 48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12–16, 2021, Glasgow, Scotland (Virtual Conference), 2021, pp. 64:1–64:20. doi:10.4230/LIPIcs.ICALP.2021.64.
- [11] S. Even, Y. Shiloach, An on-line edge-deletion problem, *J. ACM* 28 (1) (1981) 1–4. doi:10.1145/322234.322235.
- [12] L. Roditty, U. Zwick, On dynamic shortest paths problems, *Algorithmica* 61 (2) (2011) 389–401. doi:10.1007/s00453-010-9401-5.
- [13] M. P. Gutenberg, V. V. Williams, N. Wein, New algorithms and hardness for incremental single-source shortest paths in directed graphs, in: Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, ACM, 2020, pp. 153–166. doi:10.1145/3357713.3384236.
- [14] A. Bernstein, Maintaining shortest paths under deletions in weighted directed graphs, *SIAM J. Comput.* 45 (2) (2016) 548–574. doi:10.1137/130938670.
- [15] A. Bernstein, J. van den Brand, M. P. Gutenberg, D. Nanongkai, T. Saranurak, A. Sidford, H. Sun, Fully-dynamic graph sparsifiers against an adaptive adversary, *CoRR* abs/2004.08432. arXiv:2004.08432.
- [16] A. Bernstein, S. Chechik, Deterministic decremental single source shortest paths: beyond the $o(mn)$ bound, in: Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, 2016, pp. 389–397. doi:10.1145/2897518.2897521.
- [17] A. Bernstein, Deterministic partially dynamic single source shortest paths in weighted graphs, in: 44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, 2017, pp. 44:1–44:14. doi:10.4230/LIPIcs.ICALP.2017.44.
- [18] A. Bernstein, S. 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, 2017, pp. 453–469. doi:10.1137/1.9781611974782.29.

- [19] A. Bernstein, M. P. Gutenberg, C. Wulff-Nilsen, Near-optimal decremental SSSP in dense weighted digraphs, in: 61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, 2020, pp. 1112–1122. doi:10.1109/FOCS46700.2020.00107.
- [20] S. Chechik, Near-optimal approximate decremental all pairs shortest paths, in: 59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, IEEE Computer Society, 2018, pp. 170–181. doi:10.1109/FOCS.2018.00025.
- [21] J. Chuzhoy, S. Khanna, A new algorithm for decremental single-source shortest paths with applications to vertex-capacitated flow and cut problems, in: Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, ACM, 2019, pp. 389–400. doi:10.1145/3313276.3316320.
- [22] M. P. Gutenberg, C. Wulff-Nilsen, Deterministic algorithms for decremental approximate shortest paths: Faster and simpler, in: Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, SIAM, 2020, pp. 2522–2541. doi:10.1137/1.9781611975994.154.
- [23] M. P. Gutenberg, C. Wulff-Nilsen, Decremental SSSP in weighted digraphs: Faster and against an adaptive adversary, in: Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, SIAM, 2020, pp. 2542–2561. doi:10.1137/1.9781611975994.155.
- [24] M. Henzinger, S. Krinninger, D. Nanongkai, Sublinear-time decremental algorithms for single-source reachability and shortest paths on directed graphs, in: Symposium on Theory of Computing, STOC 2014, 2014, pp. 674–683. doi:10.1145/2591796.2591869.
- [25] M. Henzinger, S. Krinninger, D. Nanongkai, Decremental single-source shortest paths on undirected graphs in near-linear total update time, in: 55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, 2014, pp. 146–155. doi:10.1109/FOCS.2014.24.
- [26] M. Henzinger, S. Krinninger, D. Nanongkai, Dynamic approximate all-pairs shortest paths: Breaking the $O(mn)$ barrier and derandomization, *SIAM J. Comput.* 45 (3) (2016) 947–1006. doi:10.1137/140957299.
- [27] A. Karczmarz, J. Łącki, Reliable hubs for partially-dynamic all-pairs shortest paths in directed graphs, in: 27th Annual European Symposium on Algorithms, ESA 2019, 2019, pp. 65:1–65:15. doi:10.4230/LIPIcs.ESA.2019.65.
- [28] A. Karczmarz, J. Łącki, Simple label-correcting algorithms for partially dynamic approximate shortest paths in directed graphs, in: 3rd Symposium on Simplicity in Algorithms, SOSA@SODA 2020, SIAM, 2020, pp. 106–120. doi:10.1137/1.9781611976014.15.
- [29] A. Bernstein, M. P. Gutenberg, T. Saranurak, Deterministic decremental reachability, scc, and shortest paths via directed expanders and congestion balancing, in: 61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, 2020, pp. 1123–1134. doi:10.1109/FOCS46700.2020.00108.
- [30] A. Bernstein, M. P. Gutenberg, T. Saranurak, Deterministic decremental SSSP and approximate min-cost flow in almost-linear time, *CoRR* abs/2101.07149. arXiv:2101.07149.
- [31] A. 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, ACM, 2010, pp. 121–130. doi:10.1145/1806689.1806708.
- [32] I. Abraham, S. Chechik, C. Gavoille, Fully dynamic approximate distance oracles for planar graphs via forbidden-set distance labels, in: Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, 2012, pp. 1199–1218. doi:10.1145/2213977.2214084.
- [33] I. Abraham, S. Chechik, D. Delling, A. V. Goldberg, R. F. Werneck, On dynamic approximate shortest paths for planar graphs with worst-case costs, in: Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, 2016, pp. 740–753. doi:10.1137/1.9781611974331.ch53.
- [34] J. Fakcharoenphol, S. Rao, Planar graphs, negative weight edges, shortest paths, and near linear time, *J. Comput. Syst. Sci.* 72 (5) (2006) 868–889. doi:10.1016/j.jcss.2005.05.007.
- [35] H. Kaplan, S. Mozes, Y. Nussbaum, M. Sharir, Submatrix maximum queries in monge matrices and partial monge matrices, and their applications, *ACM Trans. Algorithms* 13 (2) (2017) 26:1–26:42. doi:10.1145/3039873.
- [36] P. N. Klein, Multiple-source shortest paths in planar graphs, in: Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2005, 2005, pp. 146–155. URL <http://dl.acm.org/citation.cfm?id=1070432.1070454>
- [37] P. N. Klein, S. Subramanian, A fully dynamic approximation scheme for shortest paths in planar graphs, *Algorithmica* 22 (3) (1998) 235–249. doi:10.1007/PL00009223.
- [38] A. Karczmarz, Decremental transitive closure and shortest paths for planar digraphs and beyond, in: Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, SIAM, 2018, pp. 73–92. doi:10.1137/1.9781611975031.5.
- [39] S. Chechik, Approximate distance oracles with improved bounds, in: Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, 2015, pp. 1–10. doi:10.1145/2746539.2746562.
- [40] C. Wulff-Nilsen, Approximate distance oracles with improved query time, in: Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, SIAM, 2013, pp. 539–549. doi:10.1137/1.9781611973105.39.
- [41] M. Thorup, U. Zwick, Approximate distance oracles, *J. ACM* 52 (1) (2005) 1–24. doi:10.1145/1044731.1044732.
- [42] S. Cabello, Many distances in planar graphs, *Algorithmica* 62 (1-2) (2012) 361–381. doi:10.1007/s00453-010-9459-0.
- [43] D. Z. Chen, J. Xu, Shortest path queries in planar graphs, in: Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, STOC 2000, 2000, pp. 469–478. doi:10.1145/335305.335359.
- [44] H. Djidjev, On-line algorithms for shortest path problems on planar digraphs, in: Graph-Theoretic Concepts in Computer Science, 22nd International Workshop, WG '96, 1996, pp. 151–165. doi:10.1007/3-540-62559-3_14.
- [45] S. Mozes, C. Sommer, Exact distance oracles for planar graphs, in: Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, 2012, pp. 209–222. doi:10.1137/1.9781611973099.19.

- [46] V. Fredslund-Hansen, S. Mozes, C. Wulff-Nilsen, Truly subquadratic exact distance oracles with constant query time for planar graphs, CoRR abs/2009.14716. arXiv:2009.14716.
- [47] Y. Long, S. Pettie, Planar distance oracles with better time-space tradeoffs, in: Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 2517–2537. doi:10.1137/1.9781611976465.149.
- [48] P. Charalampopoulos, P. Gawrychowski, S. Mozes, O. Weimann, Almost optimal distance oracles for planar graphs, in: Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, 2019, pp. 138–151. doi:10.1145/3313276.3316316.
- [49] V. Cohen-Addad, S. Dahlgaard, C. Wulff-Nilsen, Fast and compact exact distance oracle for planar graphs, in: 58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, 2017, pp. 962–973. doi:10.1109/FOCS.2017.93.
- [50] P. Gawrychowski, S. Mozes, O. Weimann, C. Wulff-Nilsen, Better tradeoffs for exact distance oracles in planar graphs, in: Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, 2018, pp. 515–529. doi:10.1137/1.9781611975031.34.
- [51] S. Cabello, Subquadratic algorithms for the diameter and the sum of pairwise distances in planar graphs, ACM Trans. Algorithms 15 (2) (2019) 21:1–21:38. doi:10.1145/3218821.
- [52] P. N. Klein, Preprocessing an undirected planar network to enable fast approximate distance queries, in: Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2002, 2002, pp. 820–827. URL <http://dl.acm.org/citation.cfm?id=545381.545488>
- [53] M. Thorup, Compact oracles for reachability and approximate distances in planar digraphs, J. ACM 51 (6) (2004) 993–1024. doi:10.1145/1039488.1039493.
- [54] Q. Gu, G. Xu, Constant query time $(1+\epsilon)$ -approximate distance oracle for planar graphs, in: Algorithms and Computation - 26th International Symposium, ISAAC 2015, Proceedings, Vol. 9472 of Lecture Notes in Computer Science, Springer, 2015, pp. 625–636. doi:10.1007/978-3-662-48971-0_53.
- [55] K. Kawarabayashi, C. Sommer, M. Thorup, More compact oracles for approximate distances in undirected planar graphs, in: Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, 2013, pp. 550–563. doi:10.1137/1.9781611973105.40.
- [56] C. Wulff-Nilsen, Approximate distance oracles for planar graphs with improved query time-space tradeoff, in: Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, SIAM, 2016, pp. 351–362. doi:10.1137/1.9781611974331.ch26.
- [57] T. M. Chan, D. Skrepetos, Faster approximate diameter and distance oracles in planar graphs, Algorithmica 81 (8) (2019) 3075–3098. doi:10.1007/s00453-019-00570-z.
- [58] P. Charalampopoulos, S. Mozes, B. Tebeka, Exact distance oracles for planar graphs with failing vertices, in: Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, 2019, pp. 2110–2123. doi:10.1137/1.9781611975482.127.
- [59] P. Gawrychowski, A. Karczmarz, Improved bounds for shortest paths in dense distance graphs, in: 45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, 2018, pp. 61:1–61:15. doi:10.4230/LIPIcs.ICALP.2018.61.
- [60] G. F. Italiano, Y. Nussbaum, P. Sankowski, C. Wulff-Nilsen, Improved algorithms for min cut and max flow in undirected planar graphs, in: Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, 2011, pp. 313–322. doi:10.1145/1993636.1993679.
- [61] A. Abboud, S. Dahlgaard, Popular conjectures as a barrier for dynamic planar graph algorithms, in: IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 2016, pp. 477–486. doi:10.1109/FOCS.2016.58.
- [62] S. Subramanian, A fully dynamic data structure for reachability in planar digraphs, in: Algorithms - ESA '93, First Annual European Symposium, Bad Honnef, Germany, September 30 - October 2, 1993, Proceedings, 1993, pp. 372–383. doi:10.1007/3-540-57273-2_72.
- [63] S. Cabello, E. W. Chambers, J. Erickson, Multiple-source shortest paths in embedded graphs, SIAM J. Comput. 42 (4) (2013) 1542–1571. doi:10.1137/120864271.
- [64] D. Eppstein, Z. Galil, G. F. Italiano, T. H. Spencer, Separator based sparsification. I. Planarity testing and minimum spanning trees, J. Comput. Syst. Sci. 52 (1) (1996) 3–27. doi:10.1006/jcss.1996.0002.
- [65] J. Holm, E. Rotenberg, Fully-dynamic planarity testing in polylogarithmic time, in: Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, ACM, 2020, pp. 167–180. doi:10.1145/3357713.3384249. URL <https://doi.org/10.1145/3357713.3384249>
- [66] K. Diks, P. Sankowski, Dynamic plane transitive closure, in: Algorithms - ESA 2007, 15th Annual European Symposium, Proceedings, 2007, pp. 594–604. doi:10.1007/978-3-540-75520-3_53.
- [67] D. Eppstein, G. F. Italiano, R. Tamassia, R. E. Tarjan, J. R. Westbrook, M. Yung, Maintenance of a minimum spanning forest in a dynamic plane graph, J. Algorithms 13 (1) (1992) 33–54. doi:10.1016/0196-6774(92)90004-V.
- [68] G. Borradaile, P. N. Klein, S. Mozes, Y. Nussbaum, C. Wulff-Nilsen, Multiple-source multiple-sink maximum flow in directed planar graphs in near-linear time, SIAM J. Comput. 46 (4) (2017) 1280–1303. doi:10.1137/15M1042929.
- [69] G. Borradaile, P. N. Klein, An $O(n \log n)$ algorithm for maximum st -flow in a directed planar graph, J. ACM 56 (2) (2009) 9:1–9:30. doi:10.1145/1502793.1502798.
- [70] M. A. Bender, J. T. Fineman, S. Gilbert, R. E. Tarjan, A new approach to incremental cycle detection and related problems, ACM Trans. Algorithms 12 (2) (2016) 14:1–14:22. doi:10.1145/2756553.
- [71] A. Bernstein, M. Probst, C. Wulff-Nilsen, Decremental strongly-connected components and single-source reachability in near-linear time, in: Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, ACM, 2019, pp. 365–376. doi:10.1145/3313276.3316335.
- [72] B. Haeupler, T. Kavitha, R. Mathew, S. Sen, R. E. Tarjan, Incremental cycle detection, topological ordering, and strong

- component maintenance, *ACM Trans. Algorithms* 8 (1) (2012) 3:1–3:33. doi:10.1145/2071379.2071382.
- [73] G. F. Italiano, A. Karczmarz, J. Lacki, P. Sankowski, Decremental single-source reachability in planar digraphs, in: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017*, ACM, 2017, pp. 1108–1121. doi:10.1145/3055399.3055480.
- [74] A. Abboud, V. V. Williams, Popular conjectures imply strong lower bounds for dynamic problems, in: *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014*, IEEE Computer Society, 2014, pp. 434–443. doi:10.1109/FOCS.2014.53.
- [75] J. Holm, K. de Lichtenberg, M. Thorup, Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity, *J. ACM* 48 (4) (2001) 723–760. doi:10.1145/502090.502095. URL <https://doi.org/10.1145/502090.502095>
- [76] S. Huang, D. Huang, T. Kopelowitz, S. Pettie, Fully dynamic connectivity in $O(\log n(\log \log n)^2)$ amortized expected time, in: *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017*, 2017, pp. 510–520. doi:10.1137/1.9781611974782.32.
- [77] C. Wulff-Nilsen, Faster deterministic fully-dynamic graph connectivity, in: *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013*, SIAM, 2013, pp. 1757–1769. doi:10.1137/1.9781611973105.126.
- [78] B. M. Kapron, V. King, B. Mountjoy, Dynamic graph connectivity in polylogarithmic worst case time, in: *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013*, 2013, pp. 1131–1142. doi:10.1137/1.9781611973105.81.
- [79] Z. Wang, An improved randomized data structure for dynamic graph connectivity, *CoRR* abs/1510.04590. arXiv:1510.04590.
- [80] D. Nanongkai, T. Saranurak, C. Wulff-Nilsen, Dynamic minimum spanning forest with subpolynomial worst-case update time, in: *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017*, IEEE Computer Society, 2017, pp. 950–961. doi:10.1109/FOCS.2017.92.
- [81] J. Chuzhoy, Y. Gao, J. Li, D. Nanongkai, R. Peng, T. Saranurak, A deterministic algorithm for balanced cut with applications to dynamic connectivity, flows, and beyond, in: *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020*, 2020, pp. 1158–1167. doi:10.1109/FOCS46700.2020.00111.
- [82] G. Goranci, H. Räcke, T. Saranurak, Z. Tan, The expander hierarchy and its applications to dynamic graph algorithms, in: *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 2212–2228. doi:10.1137/1.9781611976465.132.
- [83] G. L. Miller, Finding small simple cycle separators for 2-connected planar graphs, in: *Proceedings of the 16th Annual ACM Symposium on Theory of Computing, STOC 1984*, 1984, pp. 376–382. doi:10.1145/800057.808703.
- [84] P. N. Klein, S. Mozes, C. Sommer, Structured recursive separator decompositions for planar graphs in linear time, in: *Symposium on Theory of Computing Conference, STOC’13*, 2013, pp. 505–514. doi:10.1145/2488608.2488672.
- [85] G. Borradaile, P. Sankowski, C. Wulff-Nilsen, Min st -cut oracle for planar graphs with near-linear preprocessing time, *ACM Trans. Algorithms* 11 (3) (2015) 16:1–16:29. doi:10.1145/2684068.
- [86] G. N. Frederickson, Fast algorithms for shortest paths in planar graphs, with applications, *SIAM J. Comput.* 16 (6) (1987) 1004–1022. doi:10.1137/0216064.
- [87] Y. Nussbaum, Network flow problems in planar graphs, Ph.D. thesis, Tel Aviv University (2014).
- [88] J. Erickson, K. Fox, L. Lkhamsuren, Holiest minimum-cost paths and flows in surface graphs, in: *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018*, 2018, pp. 1319–1332. doi:10.1145/3188745.3188904.
- [89] P. N. Klein, S. Mozes, O. Weimann, Shortest paths in directed planar graphs with negative lengths: A linear-space $O(n \log^2 n)$ -time algorithm, *ACM Trans. Algorithms* 6 (2) (2010) 30:1–30:18. doi:10.1145/1721837.1721846.
- [90] S. Mozes, C. Wulff-Nilsen, Shortest paths in planar graphs with real lengths in $O(n \log^2 n / \log \log n)$ time, in: *Algorithms - ESA 2010, 18th Annual European Symposium. Proceedings, Part II*, 2010, pp. 206–217. doi:10.1007/978-3-642-15781-3_18.