

**Original citation:**

Łącki, Jakub, Oćwieja, Jakub, Pilipczuk, Marcin, Sankowski, Piotr and Zych, Anna (2015) The power of dynamic distance oracles. In: STOC '15 Forty-Seventh Annual ACM on Symposium on Theory of Computing, Portland, OR, 14-17 Jun 2015. Published in: STOC '15 Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing pp. 11-20.

**Permanent WRAP url:**

<http://wrap.warwick.ac.uk/71320>

**Copyright and reuse:**

The Warwick Research Archive Portal (WRAP) makes this work of researchers of the University of Warwick available open access under the following conditions. Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

**A note on versions:**

The version presented here is a working paper or pre-print that may be later published elsewhere. If a published version is known of, the above WRAP url will contain details on finding it.

For more information, please contact the WRAP Team at: [publicatons@warwick.ac.uk](mailto:publicatons@warwick.ac.uk)



<http://wrap.warwick.ac.uk/>

# The Power of Dynamic Distance Oracles: Efficient Dynamic Algorithms for the Steiner Tree

Jakub Łącki\*    Jakub Oćwieja<sup>†</sup>    Marcin Pilipczuk<sup>‡</sup>    Piotr Sankowski<sup>§</sup>  
Anna Zych<sup>¶</sup>

## Abstract

In this paper we study the Steiner tree problem over a dynamic set of terminals. We consider the model where we are given an  $n$ -vertex graph  $G = (V, E, w)$  with positive real edge weights, and our goal is to maintain a tree which is a good approximation of the minimum Steiner tree spanning a terminal set  $S \subseteq V$ , which changes over time. The changes applied to the terminal set are either terminal additions (*incremental* scenario), terminal removals (*decremental* scenario), or both (*fully dynamic* scenario). Our task here is twofold. We want to support updates in sublinear  $o(n)$  time, and keep the approximation factor of the algorithm as small as possible.

We show that we can maintain a  $(6 + \varepsilon)$ -approximate Steiner tree of a general graph in  $\tilde{O}(\sqrt{n} \log D)$  time per terminal addition or removal. Here,  $D$  denotes the stretch of the metric induced by  $G$ . For planar graphs we achieve the same running time and the approximation ratio of  $(2 + \varepsilon)$ . Moreover, we show faster algorithms for incremental and decremental scenarios. Finally, we show that if we allow higher approximation ratio, even more efficient algorithms are possible. In particular we show a polylogarithmic time  $(4 + \varepsilon)$ -approximate algorithm for planar graphs.

One of the main building blocks of our algorithms are dynamic distance oracles for vertex-labeled graphs, which are of independent interest. We also improve and use the online algorithms for the Steiner tree problem.

---

\*University of Warsaw, [j.lacki@mimuw.edu.pl](mailto:j.lacki@mimuw.edu.pl). Jakub Łącki is a recipient of the Google Europe Fellowship in Graph Algorithms, and this research is supported in part by this Google Fellowship.

<sup>†</sup>University of Warsaw, [j.ocwieja@mimuw.edu.pl](mailto:j.ocwieja@mimuw.edu.pl). Partially supported by ERC grant PAAI no. 259515.

<sup>‡</sup>University of Bergen, [malcin@mimuw.edu.pl](mailto:malcin@mimuw.edu.pl). The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 267959. Partially supported by ERC grant PAAI no. 259515.

<sup>§</sup>University of Warsaw, [sank@mimuw.edu.pl](mailto:sank@mimuw.edu.pl). Partially supported by ERC grant PAAI no. 259515, NCN grant N206 567940 and the Foundation for Polish Science.

<sup>¶</sup>University of Warsaw, [anka@mimuw.edu.pl](mailto:anka@mimuw.edu.pl). Partially supported by ERC grant PAAI no. 259515.

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	Our results . . . . .	4
1.2	Related results . . . . .	6
1.3	Organization of the paper . . . . .	7
<b>2</b>	<b>Preliminaries</b>	<b>7</b>
<b>3</b>	<b>Overview of our algorithms</b>	<b>8</b>
3.1	Online algorithms . . . . .	8
3.1.1	Decremental online algorithm . . . . .	10
3.1.2	Incremental online algorithm . . . . .	10
3.1.3	Fully dynamic online algorithm . . . . .	11
3.2	Dynamic vertex-color distance oracles . . . . .	12
3.2.1	Oracle interface . . . . .	13
3.2.2	Generic distance oracles . . . . .	14
3.2.3	From a generic oracle to a dynamic vertex-color distance oracle . . . . .	15
3.3	Implementing the online algorithms efficiently . . . . .	16
3.3.1	Efficient decremental algorithm . . . . .	16
3.3.2	Efficient incremental algorithm . . . . .	17
3.3.3	Efficient fully dynamic algorithm . . . . .	19
<b>4</b>	<b>Online algorithms</b>	<b>20</b>
4.1	$\mu$ -near metric spaces . . . . .	21
4.1.1	Analysis of the greedy approach in near metric spaces . . . . .	22
4.2	Properties guaranteeing good approximation factors . . . . .	23
4.3	Decremental online algorithm . . . . .	26
4.4	Incremental online algorithm . . . . .	28
4.5	Fully dynamic online algorithm . . . . .	30
4.5.1	Decreasing the dependence on the stretch of the metric . . . . .	33
<b>5</b>	<b>Dynamic vertex-color distance oracles</b>	<b>35</b>
5.1	Generic construction . . . . .	35
5.2	3-approximate oracles for general graphs . . . . .	41
5.3	$(1 + \epsilon)$ -approximate oracles for planar graphs . . . . .	43
5.4	Additional operations and properties . . . . .	48
<b>6</b>	<b>Sublinear time dynamic algorithms</b>	<b>51</b>
6.1	Decremental algorithm . . . . .	51
6.2	Incremental algorithm . . . . .	55
6.3	Fully dynamic algorithm . . . . .	61
<b>7</b>	<b>Steiner tree via bipartite emulators</b>	<b>65</b>
7.1	Constructing bipartite emulators . . . . .	68
7.1.1	General graphs . . . . .	68
7.1.2	Planar graphs . . . . .	68
<b>8</b>	<b>Conclusions</b>	<b>71</b>

# 1 Introduction

Imagine a network and a set of users that want to maintain a cheap multicast tree in this network during a conference call [17]. The users can join and leave, but in the considered time scale the network remains static. In other words we are considering the following problem. We are given a graph  $G = (V, E, w)$  with positive edge weights  $w : E \rightarrow \mathbb{R}^+$ . The goal is to maintain information about approximate Steiner tree in  $G$  for a dynamically changing set  $S \subseteq V$  of terminals.

This problem was first introduced in the pioneering paper by Imase and Waxman [28] and its study was later continued in [30, 20, 21]. However, all these papers focus on minimizing the number of changes to the tree that are necessary to maintain a good approximation, and ignore the problem of efficiently finding these changes. The problem of maintaining a Steiner tree is also one of the important open problems in the network community [14], and while it has been studied for many years, the research resulted only in several heuristic approaches [6, 2, 27, 33] none of which has been formally proved to be efficient. In this paper we show the first sublinear time algorithm which maintains an approximate Steiner tree under terminal additions and deletions.

Our paper deals with two variants of the problem of maintaining the Steiner tree. Throughout this paper, we assume that in the *online Steiner tree* problem (we also say *online setting*) the goal is to maintain a Steiner tree making few changes to the tree after each terminal addition or removal. On the other hand, in the *dynamic Steiner tree* problem (*dynamic setting*), the requirement is that each update is processed faster than by recomputing the tree from scratch. This aligns with the usual definition of a dynamic algorithm used in the algorithmic community (see e.g. [18]). In this paper we study both settings and use the techniques for the online setting to show new results in the dynamic setting.

As our point of reference we observe that it is possible to construct  $\tilde{O}(n)$  time dynamic algorithm for Steiner tree using the dynamic polylogarithmic time minimum spanning forest (dynamic MSF) algorithm [26]. This solution is obtained by first computing the metric closure  $\overline{G}$  of the graph  $G$ , and then maintaining the minimum spanning tree (MST) over the set of terminals  $S$  in  $\overline{G}[S]$ . It is a well-known fact that this yields a 2-approximate Steiner tree. In order to update  $\overline{G}[S]$  we need to insert and remove terminals together with their incident edges, which requires  $\Theta(n)$  calls to the dynamic MSF structure. However, such a linear bound is far from being satisfactory as does not lead to any improvement in the running time for sparse networks where  $m = O(n)$ .<sup>1</sup> In such networks after each update we can actually compute the 2-approximate Steiner tree in  $O(n \log n)$  time from scratch [31]. Hence, the main challenge is to break the linear time barrier for maintaining constant approximate tree. Only algorithms with such sublinear complexity could be of some practical importance and could potentially be used to reduce the communication cost of dynamic multicast trees. Our paper aims to be a theoretical proof of concept that from algorithmic complexity perspective this is indeed possible.

As observed by [11] and [12], the dynamic problems with vertex updates are much more challenging, but are actually closer to real-world network models than problems with edge updates. In computer networks, vertex updates happen more often, as they correspond to software events (server reboot, misconfiguration, or simply activation of a user), whereas edge updates are much less likely, as they correspond to physical events (cable cut/repair).

Finally, we note that the Steiner tree problem is one of the most fundamental problems in combinatorial optimization. It has been studied in many different settings, starting from classical

---

<sup>1</sup>It is widely observed that most real-world networks are sparse [15].

Setting	apx.	update time	preprocessing time
general, fully dynamic	$6 + \varepsilon$	$\tilde{O}(\sqrt{n} \log D)$	$\tilde{O}(\sqrt{n}(m + n \log D))$
general, incremental	$6 + \varepsilon$	$\tilde{O}(\sqrt{n})$	$\tilde{O}(m\sqrt{n})$
general, decremental	$6 + \varepsilon$	$\tilde{O}(\sqrt{n})$	$\tilde{O}(m\sqrt{n})$
planar, fully dynamic	$2 + \varepsilon$	$\tilde{O}(\sqrt{n} \log D)$	$\tilde{O}(n \log D)$
planar, incremental	$2 + \varepsilon$	$\tilde{O}(\log^3 n \log D)$	$\tilde{O}(n \log D)$

Table 1: Summary of our algorithms for maintaining Steiner tree in general and planar graphs. For an input graph  $G$ ,  $n = |V(G)|$ ,  $m = |E(G)|$ , and  $D$  is the stretch of the metric induced by  $G$ . The dependence of the running times on  $\varepsilon^{-1}$  is polynomial. The update times are amortized, some of them are also in expectation (see the corresponding theorem statements).

approximation algorithms [7, 32, 5, 34, 9, 10], through online [28, 30, 20, 21] and stochastic models [22, 19], ending with game theoretic approaches [4, 8]. Taking into account the significance and the wide interest in this problem, it is somewhat surprising that no efficient dynamic algorithms for this problem have been developed so far. It might be related to the fact that this would require combining ideas from the area of approximation algorithms with the tools specific to dynamic algorithms. This is the first paper that manages to do so.

In this extended abstract we only present the main ideas, skipping the proofs and the formal analysis, which will be presented in the full version of this paper.

## 1.1 Our results

The main result of this paper are sublinear time algorithms for maintaining an approximate Steiner tree. We provide different algorithms for incremental, decremental and fully dynamic scenarios. An *incremental* algorithm allows only to add terminals, a *decremental* algorithm supports removing terminals, whereas a final *fully dynamic* algorithm supports both these operations. Our results are summarized in Table 1. The overall approximation ratio of the algorithms we obtain is  $(6 + \varepsilon)$  for general graphs and  $(2 + \varepsilon)$  for planar graphs. In particular, we can maintain a fully dynamic  $(6 + \varepsilon)$ -approximate tree in  $\tilde{O}(\sqrt{n} \log D)$  amortized time per update in an arbitrary weighted graph, where  $D$  is the stretch of the metric induced by the graph. This extended abstract aims to present a brief overview of this result. The result is a composition and a consequence of many ideas that are of independent interest. We believe that the strength of this paper lies not only in the algorithms we propose, but also in the byproducts of our construction. We outline these additional results below.

**Dynamic vertex-color distance oracles** The algorithms for online Steiner tree assume that the entire metric (i.e., the distances between any pair of vertices) is given explicitly. This assumption is not feasible in case of a metric induced by a graph. Hence, to obtain the necessary distances efficiently, we develop a data structure called *dynamic vertex-color distance oracle*. This oracle is given a weighted undirected graph and maintains an assignment of colors to vertices. While the graph remains fixed, the colors may change. The oracle can answer, among other queries, what is the nearest vertex of color  $c$  to a vertex  $v$ . We develop two variants of approximate vertex-color distance oracles: incremental and fully dynamic. In the first variant, each vertex is initially given

a distinct color and the color sets (i.e., two sets representing vertices of the same color) can be merged. The fully dynamic oracles additionally support (restricted) operations of splitting color sets.<sup>2</sup> Note that these update operations are much more general than the operation of changing the color of a single vertex, which was considered in earlier works [25, 13].

For planar graphs we propose two  $(1 + \varepsilon)$ -approximate oracles. The incremental oracle supports all operations in  $O(\varepsilon^{-1} \log^2 n \log D)$  amortized time (in expectation), whereas the fully dynamic oracle supports operations in worst case time  $O(\varepsilon^{-1} \sqrt{n} \log^2 n \log D)$ . For general graphs we introduce a 3-approximate fully dynamic oracle that works in  $O(\sqrt{n} \log n)$  expected time. Our construction of oracles is generic, that is we show how to extend oracles that may answer vertex-to-vertex queries and satisfy certain conditions into dynamic vertex-color oracles. For that we introduce the concept of a *generic distance oracle*, which captures the common properties of many distance oracles and allows us to use a uniform approach for different oracles for planar and general graphs.

**Online Steiner tree** We show an online algorithm that decrementally maintains  $(2 + \varepsilon)$ -approximate Steiner tree, applying after each terminal deletion  $O(\varepsilon^{-1})$  changes to the tree (in amortized sense). This improves over the previous 4-approximate algorithm. In addition to that, we show a fully dynamic  $(2 + \varepsilon)$ -approximate online algorithm, which makes  $O(\varepsilon^{-1} \log D)$  changes to the tree (in amortized sense) after each operation. One of the new techniques used to obtain these results is the lazy handling of high degree Steiner nodes for arbitrary degree threshold. This improves over an algorithm by Imase and Waxman [28], which makes  $O(t^{3/2})$  changes to process  $t$  addition or removal operations, and maintains a 4-approximate tree. An algorithm performing a smaller number of changes was given in [21], but, as we discuss in the next section, it departs slightly from the classical Imase-Waxman model.

**Query Steiner tree** In the query model, as defined in [16], for a fixed graph  $G$ , we are asked queries to compute an approximate Steiner tree for a set  $S \subseteq V$  as fast as possible. This models a situation when many sets of users want to setup a new multicast tree. We obtain an algorithm, which after preprocessing in  $O(\sqrt{n}(m + n \log n))$  expected time uses  $O(n\sqrt{n} \log n)$  space and computes a 6-approximate Steiner tree in  $O(|S|\sqrt{n} \log n)$  expected time. In the planar case, we can compute  $(2 + \varepsilon)$ -approximate tree in  $O(|S|\varepsilon^{-1} \log^2 n \log D)$  expected time, using  $O(\varepsilon^{-1} n \log^2 n \log D)$  preprocessing time and space. In other words, we show a more efficient solution for computing many multicast trees in one fixed graph than computing each tree separately. This preprocessing problem is related to the study initiated in [1] where compact multicast routing schemes were shown. In comparison with [1], our schemes are not only compact but efficient as well.

**Nonrearrangeable incremental Steiner tree** In the nonrearrangeable incremental Steiner tree problem one has to connect arriving terminals to the previously constructed tree without modifying it. We show how to implement the  $O(\log n)$ -approximate online algorithm for this problem given by Imase and Waxman [28] so that it runs in  $O(r\sqrt{n} \log n)$  expected time for non-planar graphs and  $O(r \log^2 n \log D)$  expected time for planar graphs. Here  $r$  denotes the final number of terminals. This gives an improvement over the naive execution of this algorithm that requires  $O(r^2)$  time and resolves one of the open problems in [14].

---

<sup>2</sup>Because of these two operations we believe that it is more natural to change slightly the previously used vocabulary and assign *colors* instead of *labels* to vertices.

**Bipartite emulators** As an interesting side result, we also show a different, simple approach to dynamic Steiner tree, which exposes a trade-off between the approximation ratio and the running time. It is based on *bipartite emulators*: low-degree bipartite graphs that approximate distances in the original graph. We run the dynamic MSF algorithm on top of a bipartite emulator to obtain sublinear time dynamic algorithms. In particular, we obtain a 12-approximate algorithm for general graphs that processes each update in  $\tilde{O}(\sqrt{n})$  expected amortized time and a  $(4 + \varepsilon)$ -approximate algorithm for planar graphs processing updates in  $\tilde{O}(\varepsilon^{-1} \log^6 n)$  amortized time. While our emulators are constructed using previously known distance oracles [37, 36], our contribution lies in introducing the concept of bipartite emulators, whose properties make it possible to solve the Steiner tree problem with a modification of the dynamic MSF algorithm [26].

We want to stress that the construction of our algorithms for the Steiner tree, in particular the approach that combines online Steiner tree algorithm with a distance oracle, is highly modular. Not only any improvement in the construction of the vertex-color distance oracles will result in better Steiner tree algorithms, but the vertex-color distance oracles themselves are constructed in a generic way out of distance oracles in [36, 37]. The approximation factor of  $(6 + \varepsilon)$  for Steiner tree in general graphs comes from using 3-approximate oracles combined with a 2-approximation of the Steiner tree given by the MST in the metric closure and  $(1 + \varepsilon)$ -approximate online MST algorithm. In other words we hit two challenging bounds: in order to improve our approximation factors, one would need either to improve the approximation ratio of the oracles which are believed to be optimal, or devise a framework not based on computing the MST. The second challenge would require to construct simple and fast (e.g., near-linear time) approximation algorithms for Steiner tree that would beat the MST approximation ratio of 2. Constructing such algorithms is a challenging open problem.

## 1.2 Related results

The problems we deal with in this paper and related have received a lot of attention in the literature. We present a brief summary in this section.

**Vertex-color distance oracles** Our vertex-color distance oracles fall into the model studied in the literature under the name of vertex-label distance oracles. Dynamic vertex-color oracles for general graphs have been introduced by Hermelin et. al [25] and improved by Chechik [13]. These oracles allow only to change a color (label) of a single vertex, as opposed to our split and merge operations. The oracle by Chechik [13] has expected size  $\tilde{O}(n^{1+1/k})$ , and reports  $(4k - 5)$ -approximate distances in  $O(k)$  time. This oracle can support changes of vertices' colors in  $O(n^{1/k} \log n)$  time. Our results have much better approximation guarantee and more powerful update operations, at the cost of higher query time. A vertex-color oracle for planar graphs has been shown by Li, Ma and Ning [29], but it does not support updating colors. The incremental variant of our oracle allows merging colors, and has only slightly higher running time.

**Online Steiner tree** There has been an increasing interest in the online Steiner tree and the related online MST problem in recent years, which started with a paper by Megow et al. [30]. They showed that in the incremental case one can maintain an approximate online MST in  $\bar{G}[S]$  (and consequently an approximate Steiner tree) with only a constant number of changes to the tree per terminal insertion (in amortized sense), which resolved a long standing open problem posed in [28].

The result of [30] was improved to worst-case constant by Gu, Gupta and Kumar [20]. Then, Gupta and Kumar [21] have shown that constant worst-case number of changes is sufficient in the decremental case. Their paper also shows a fully dynamic algorithm that performs constant number of changes in amortized sense, but in a slightly different model. In [21], a newly added terminal is treated as a new vertex in the graph, with new distances given *only* to all *currently active* (not yet deleted) terminals; the remaining distances are assumed implicitly by the triangle inequality. However, in the classical Imase-Waxman model, the entire algorithm runs on a fixed host graph that is given at the beginning, and terminals are only activated and deactivated. After a terminal is deleted, it may still be used in the maintained tree as a Steiner node. In particular, in our algorithms it is crucial that the newly added terminal may be connected directly to a Steiner node. This is not allowed in [21], so it seems that their analysis cannot be used in the Imase-Waxman model that is studied here.

While the online algorithms for Steiner tree have been studied extensively, our paper is the first one to show that they can be turned into algorithms with low running time. Moreover, in the decremental case and fully dynamic case we are the first to show  $(2 + \varepsilon)$ -approximate algorithms.

### 1.3 Organization of the paper

The remainder of the paper is organized as follows. Section 2 introduces the notation and recalls some results used further in the paper. In Section 3 we describe a shortened version of our results, which is supposed to sketch the ideas without diving into technicalities and proofs of correctness. Section 4 presents online algorithms for Steiner tree. It covers our improvements in this area and analyses the algorithms in the setting where we only have access to approximate distances in the graph. Section 5 introduces dynamic approximate vertex-color distance oracles, starting with a generic oracle, through various constructions of particular oracles, ending with the summary of our results in this area. In Section 6 we show how to combine the online algorithms with distance oracles to obtain efficient algorithms for dynamic Steiner tree. In Section 7 we show alternative algorithms for dynamic Steiner tree problem based on bipartite emulators. These algorithms are slightly faster compared to the previously presented, but the approximation ratio doubles. Finally, Section 8 concludes the paper and suggests directions for future research.

## 2 Preliminaries

Let  $G = (V, E, d_G)$  be a graph. Throughout the paper, we assume that each graph is undirected and has positive edge weights. By  $V(G)$  and  $E(G)$  we denote the vertex and edge sets of  $G$ , respectively. For  $u, v \in V$ , we write  $\delta_G(u, v)$  to denote the distance between  $u$  and  $v$  in  $G$ . The graph we refer to may be omitted if it is clear from the context. For  $v \in V$ ,  $\Gamma_G(v) = \{u \in V : uv \in E\}$  is the set of neighbors of  $v$  in  $G$ .

Whenever we give some statement with accuracy parameter  $\varepsilon > 0$ , we implicitly assume that  $\varepsilon$  is a small fixed constant. In particular, we assume  $\varepsilon < 1$ .

We define  $\overline{G} = (V, \binom{V}{2}, d_{\overline{G}})$  to be the metric closure of  $G$ , i.e., for  $u, v \in V$ ,  $d_{\overline{G}}(u, v) = \delta_G(u, v)$ . The *stretch* of the metric induced by  $G$  is the ratio between the longest and the shortest edge of  $\overline{G}$ . Moreover, let  $G[X]$  denote the subgraph of  $G$  induced by  $X \subseteq V$  and let  $MST(G)$  stand for the minimum spanning tree in  $G$ . For a given graph  $G$  and a terminal set  $S$  we define  $ST(G, S)$  to be an optimum Steiner tree in  $G$  that spans  $S$ . Our algorithms are based on the following well-known



fact.

**Lemma 2.1.** *Let  $G = (V, E, d_G)$  be a graph and  $S \subseteq V$ . Then  $d_G(ST(G, S)) \leq d_{\overline{G}}(MST(\overline{G}[S])) \leq 2d_G(ST(G, S))$ .*

In each of our algorithms the ultimate goal is to maintain a good approximation of  $MST(\overline{G}[S])$ . Note that, as we focus on  $MST(\overline{G}[S])$ , our algorithms can also maintain a solution to Subset TSP on the same (dynamic) terminal set with the same approximation guarantee. Since we work with a metric closure of  $G$ , the solution that we maintain is actually a collection of (approximately) shortest paths in  $G$ . As we precompute these paths during initialization, our algorithms can be extended to report each such path in time that is linear in its length.

We distinguish three dynamic scenarios: *incremental*, *decremental* and *fully dynamic*. All three scenarios assume that the input graph remains the same, while the set  $S$  of terminals changes over time. In the incremental scenario we consider how to efficiently update the tree after a vertex is added to  $S$  (it is announced to be a terminal), in the decremental scenario the update operation is removing a vertex from  $S$ , and the fully dynamic scenario allows both these operations intermixed with each other. As discussed already in the introduction, we are interested in performing one update operation in sublinear (amortized) time, as a near-linear time algorithm follows easily from the works on dynamic minimum spanning forest [26].

**Dynamic MSF algorithm.** In a number of our algorithms we use an algorithm that may dynamically maintain a minimum spanning forest of a graph, subject to edge insertions and removals. In the following, we call it a dynamic MSF algorithm.

**Theorem 2.2** ([26]). *There exists a fully dynamic MSF algorithm, that for a graph on  $n$  vertices supports  $m$  edge additions and removals in  $O(m \log^4 n)$  total time.*

Whenever we use the algorithm of Theorem 2.2 on some graph, we implicitly make sure that all the edges have pairwise distinct weights. We ensure this by breaking ties between weights using the timestamp of the addition, and some arbitrary total ordering of edges present at initialization. In particular, if we compare two edges of the same weight, we prefer to use the older one (i.e., we treat it as slightly cheaper). If, after an edge  $e$  is added to the graph, the maintained forest is altered, then the new forest includes  $e$ , and either the cost of the forest strictly decreased in this operation or the insertion of  $e$  connected two connected components of the graph.

### 3 Overview of our algorithms

In this section we present a high-level overview of our algorithm for dynamic Steiner tree. This is done in three steps. First, we describe the online algorithms that we use (section 3.1). Then, we introduce dynamic vertex-color distance oracles (section 3.2), and finally (section 3.3) we show how to combine both these tools to obtain efficient algorithms.

#### 3.1 Online algorithms

We start with presenting online algorithms for Steiner tree. In the online setting the goal is to maintain an approximate Steiner tree and make small number of changes to the tree after a

terminal is added or deleted. We later use these algorithms to obtain efficient algorithms with low approximation factors.

By Lemma 2.1, in order to maintain an approximate Steiner tree, we may maintain a tree  $T$  that spans the current set of terminals  $S$  in the metric closure of the graph. Although in this section we assume that we work with a complete graph, it should not be thought of as a metric closure of  $G$ , but instead as its approximation. In later sections, this approximation will be given by a distance oracle. A  $\mu$ -approximate oracle for a graph  $G = (V, E, d_G)$  yields a complete graph  $GD = (V, \binom{V}{2}, d_{GD})$ , such that  $\delta_G(u, v) \leq d_{GD}(u, v) \leq \mu\delta_G(u, v)$ . We view  $GD$  as a complete graph, but it should be noted that it is not metric, since it does not satisfy triangle inequality. Therefore, we call  $GD$  a  $\mu$ -near metric space.

We essentially follow the core ideas of Imase and Waxman [28], with more modern improvements of [30, 20, 21]: to maintain good approximation ratio, it suffices to (a) as long as we can replace an edge  $e'$  in the tree with a another one  $e$  of significantly lower cost, proceed with the replacements; (b) defer deletions of large-degree nonterminal vertices from the tree. However, we need some technical work to formally state these properties, especially in the context of *near* metric spaces. Furthermore, in [28] it is only argued that deferral of deletion of non-terminal vertices of degree at least three does not influence the approximation factor much; we generalize this result to arbitrary threshold, in particular showing that larger thresholds give better approximation guarantee.

Following [28], when a terminal  $t$  is added to  $S$ , we can update the current tree by connecting  $t$  with  $T$  using the cheapest connecting edge. This alone does not lead to a good approximation of  $MST(GD[S \cup \{t\}])$  because some other edges incident to  $t$  could possibly be used to replace expensive edges in  $T$ . In our algorithms, we repeatedly replace tree edges with non-tree edges (assuring that after the replacement, we obtain a tree spanning  $S$ ) until we reach a point at which we can be sure that the current tree is a good approximation of  $MST(GD[S])$ . We now formalize the notion of replacement.

Let  $T$  be a tree in  $GD$ . For any edge  $e$  with both endpoints in  $V(T)$ , we say that an edge  $e_T \in E(T)$  is a *friend* of  $e$  (with respect to  $T$ ) if  $e_T$  lies on the unique path between the endpoints of  $e$ . For a friend  $e_T$  of  $e$  with regards to  $T$ ,  $T' := (T \setminus \{e_T\}) \cup \{e\}$  is a tree that spans  $V(T)$  as well, with cost  $d_{GD}(T) - (d_{GD}(e_T) - d_{GD}(e))$ . We say that  $T'$  is created from  $T$  by *replacing* the edge  $e_T$  with  $e$ , and  $(e, e_T)$  is a *replacement pair* in  $T$ .

**Definition 3.1** (heavy, efficient and good replacement). *Let  $c > \varepsilon \geq 0$  be constants and let  $T$ ,  $e$  and  $e_T$  be defined as above. We say that the  $(e, e_T)$  pair is a*

1.  $\varepsilon$ -heavy replacement if  $d_{GD}(e_T) > \varepsilon d_{GD}(T)/|V|$ ;
2.  $\varepsilon$ -efficient replacement if  $(1 + \varepsilon)d_{GD}(e) < d_{GD}(e_T)$ ;
3.  $(\varepsilon, c)$ -good replacement if it is both  $\varepsilon$ -efficient and  $(\varepsilon/c)$ -heavy.

Standard arguments show that if no  $(\varepsilon, c)$ -good replacement is present in the tree, then it is close to a minimum spanning tree of its vertex set.

**Lemma 3.2.** *Let  $c > \varepsilon \geq 0$  be constants and let  $T$  be a tree in a complete graph  $GD$ . If  $T$  does not admit any  $(\varepsilon, c)$ -good replacement, then it is a  $c(1 + \varepsilon)/(c - \varepsilon)$ -approximation of minimum spanning tree of  $V(T)$ . In particular, if  $T$  does not admit any  $\varepsilon$ -efficient replacement, then it is a  $(1 + \varepsilon)$ -approximation of minimum spanning tree on  $V(T)$ .*



Figure 1: In panel (a), all dashed edges of the tree, in particular  $e_T$ , are friends of the blue thick edge  $e$ . Panel (b) illustrates a tree with some nonterminal vertices that are of large degree.

In the next lemma we observe that having nonterminals of high degree does not influence the approximation ratio significantly.

**Lemma 3.3.** *Let  $\varepsilon \geq 0$  be a constant,  $\eta \geq 2$ ,  $\overline{G} = (V, \binom{V}{2}, d_{\overline{G}})$  be a complete weighted graph and  $S \subseteq V$ . Let  $T_{MST} = MST(\overline{G}[S])$  be a MST of  $S$  and let  $T$  be a tree in  $\overline{G}$  that spans  $S \cup N$  and does not admit  $(1 + \varepsilon)$ -efficient replacements. Furthermore, assume that each vertex of  $V(T) \cap N$  is of degree larger than  $\eta$  in  $T$ . Then  $d_{\overline{G}}(T) \leq (1 + \varepsilon) \frac{\eta}{\eta - 1} d_{\overline{G}}(T_{MST})$ .*

The special case  $\eta = 2$  of Lemma 3.3 has been proven by Imase and Waxman [28]; we provide a different, simpler proof of this result that additionally extends to arbitrary  $\eta \geq 2$ .

### 3.1.1 Decremental online algorithm

In this section we describe the decremental scheme, i.e., how to handle the case when the terminals are removed from the terminal set. In the decremental scheme, the main idea is to maintain the minimum spanning tree on the terminal set, but to postpone the deletion of terminals that are of degree above some fixed threshold  $\eta \geq 2$  in this spanning tree.

When a vertex  $v$  of degree  $s \leq \eta$  is deleted from the tree  $T$ , the tree breaks into  $s$  components  $\mathcal{T}^v = \{T_1, T_2, \dots, T_s\}$  that need to be reconnected. The natural idea is to use for this task a set of edges of  $GD$  of minimum possible total weight. That is, we first for each  $1 \leq i < j \leq s$  identify an edge  $e_{ij}$  of minimum possible cost in  $GD$  among edges between  $T_i$  and  $T_j$ . Then, we construct an auxiliary complete graph with vertex set  $\mathcal{T}^v$  and edge cost  $d(T_i T_j) := d_{GD}(e_{ij})$ , find a minimum spanning tree  $T^\circ$  of this graph, and use edges of this tree to reconnect  $T$ .

Take  $\eta = 1 + \lceil \varepsilon^{-1} \rceil$ . Lemma 3.3 implies that the decremental scheme maintains a  $2(1 + \varepsilon)\mu$ -approximation of minimum Steiner tree of  $G$ . In the course of the first  $r$  deletions, it removes a vertex from the tree at most  $r$  times and each removed vertex has degree at most  $\eta = O(\varepsilon^{-1})$  in the currently maintained tree.

### 3.1.2 Incremental online algorithm

We now consider the case when new terminals are added to the terminal set and we are to update the currently maintained tree. Following Imase and Waxman [28] and Megow et al [30], we first connect the new terminal to the tree we already have, and then try to apply all occurring  $(\varepsilon/2, 1 + \varepsilon)$ -good replacement pairs, for some  $\varepsilon > 0$ .




Figure 2: Panel (a) illustrates the addition step: when a new terminal  $t_i$  is added, it is first connected to the closest terminal with an edge  $e_0$ , and then a replacement  $(e, e')$  is applied. Panel (b) illustrates the deletion step: when a vertex  $v$  is deleted, the tree splits into subtrees  $T_i$ , and between every pair of subtrees the shortest reconnecting edge  $e_{i,j}$  is found (depicted as blue thick edge). Then, we find the set of reconnecting edges by computing a minimum spanning tree of the auxiliary graph  $G_c$  with vertex set  $\{T_1, T_2, \dots, T_s\}$  and edges  $e_{i,j}$  connecting  $T_i$  and  $T_j$ .

At one step, given a tree  $T$  and a new terminal vertex  $v \in V \setminus V(T)$ , we:

1. add an edge to  $T$  connecting  $v$  with  $V(T)$  of cost  $\min_{u \in V(T)} d_{GD}(uv)$ ;
2. apply a sequence of  $\varepsilon$ -efficient replacement pairs  $(e, e')$  that satisfy the following additional property:  $e'$  is a friend of  $e$  in the currently maintained tree of maximum possible cost;
3. after all the replacements, we require that there does not exist a  $(\varepsilon/2, 1 + \varepsilon)$ -good replacement pair  $(e, e')$  with  $e$  incident to  $v$ .

We emphasize here the the incremental algorithm requires only to care about replacement edges incident to the newly added vertex; we show that no other efficient replacement pair can be introduced in the above procedure.

Lemma 3.2 implies that the incremental algorithm for some  $\varepsilon > 0$  maintains  $2(1 + \varepsilon)\mu$  approximation of the Steiner tree. With the help of the recent analysis of [20], we show that, in the course of the first  $r$  additions, it performs  $O(r\varepsilon^{-1}(1 + \log \mu))$  replacements.

### 3.1.3 Fully dynamic online algorithm

In order to obtain a scheme for a fully dynamic algorithm, we merge the ideas of two previous sections. We fix an accuracy parameter  $\varepsilon > 0$ . We aim at maintaining a  $(1 + \varepsilon)^2$ -approximation

of  $MST(GD[S])$  where  $S$  is the current set of terminals. As in the decremental step, we define  $\eta = 1 + \lceil \varepsilon^{-1} \rceil = O(\varepsilon^{-1})$ , that is,  $\eta$  is the minimum positive integer with  $\frac{\eta}{\eta-1} \leq 1 + \varepsilon$ .

In a deletion step, we behave in exactly the same manner as in the decremental scheme. If we want to delete  $v$  from the terminal set, we mark it as non-terminal and, if its degree is at most  $\eta$ , it is removed from the tree.

In an addition step, we perform similarly as in the incremental scheme, but there are two significant differences. First, we do not have the guarantee that the cost of the tree will not decrease much in the future, so we cannot stop replacing edges at some cost threshold: although the low cost edges may contribute only a little to the weight of the tree, this may change after many terminals are removed and the weight of the tree drops.

Second, we need to watch out for non-terminal vertices whose degree may drop to the threshold  $\eta$  as a consequence of a replacement. Formally, to add a vertex  $v$  to the terminal set  $S$  we perform the following operations on the currently maintained tree  $T$ .

1. If  $v \in V(T)$ , mark  $v$  as a terminal and finish.
2. Otherwise, connect  $v$  to any vertex of  $T$  using the cheapest connection in  $GD$ .
3. Apply a sequence of  $\varepsilon$ -efficient replacement pairs, where for each such pair  $(e, e')$  we require that  $e'$  has the maximum possible cost among the friends of  $e$ .
4. Once all replacement pairs are applied, we require that there exists no  $\varepsilon$ -efficient replacement pair  $(e, e')$  with  $e$  incident to  $v$ ;
5. At the end remove all non-terminal vertices whose degree dropped to at most  $\eta$  due to replacements.

Observe that, again as in the incremental scheme, we show that it suffices to care only about replacement pairs  $(e, e')$  with  $e$  incident to the newly added vertex  $v$ .

Lemmas 3.2 and 3.3 imply that the scheme maintains a  $2(1 + \varepsilon)^2 \mu$  approximation of the Steiner tree. In the course of the first  $r$  operations, where  $r_+$  of this operations are additions and  $r_- = r - r_+$  are deletions, the algorithm performs  $O(\varepsilon^{-1}(r_+ \log D + r_-(1 + \log \mu)))$  replacements and removes at most  $r_-$  vertices, each of which has degree at most  $\eta = O(\varepsilon^{-1})$  in the currently maintained tree.

Although our problem is similar to the one considered in [21], it does not seem clear if the analysis of incremental algorithm ([20], Theorem 4.5) is applicable to our fully dynamic scheme. The usage of this analysis in [21] strongly relies on the assumption that newly added terminals are new vertices in the graph, and their distance to previously deleted terminals is inferred from the triangle inequality. However, in our model of a fixed host graph  $G$  this assumption is no longer valid, hence we need to fall back to a different analysis here.

### 3.2 Dynamic vertex-color distance oracles

In this section we introduce dynamic vertex-color distance oracles that we later use to implement the online algorithms. The oracles' interface is tailored to the construction of the algorithms. In fact, the algorithms will access the graphs only by the oracles that we describe. We consider two variants of the oracle: an *incremental* and a more general *fully dynamic*. The former will be used for incremental algorithms, whereas the latter for decremental and fully dynamic ones.




Figure 3: merge and split operations in the fully dynamic oracle.

### 3.2.1 Oracle interface

Let  $G = (V, E)$  be a graph. Each oracle maintains a partition of  $V$  into sets  $C_1, \dots, C_k$ , i.e.,  $\bigcup_{i=1}^k C_i = V$  and  $C_i \cap C_j = \emptyset$  for all  $i \neq j, 1 \leq i, j \leq k$ . In other words we assign colors to the vertices of  $G$ . At any given point each color might be *active* or not. We say that a vertex is active if the associated color is active. The oracle allows us to find the distance from a given vertex to the nearest vertex of a given color or to the nearest active vertex. The updates supported by the oracle alter only the sets  $C_i$ , while the underlying graph  $G$  is never changed.

In the incremental variant of the oracle, the only operation that allows us to change colors is merging two colors. To be more precise, the incremental oracle supports the following set of operations:

- **distance**( $v, i$ ) – compute the approximate distance from  $v$  to the nearest vertex of color  $i$ ,
- **nearest**( $v, k$ ) – compute the approximate distance from  $v$  to the  $k$ -th nearest active color (for a constant  $k$ ),
- **activate**( $i$ ) – activate set (color)  $C_i$ ,
- **merge**( $i, j$ ) – merge two different active sets  $C_i$  and  $C_j$  into one active set  $C_l$ , where  $l \in \{i, j\}$ .

The fully dynamic variant also supports splitting sets  $C_i$ , but only in a restricted way. The oracle associates with every set  $C_i$  a tree  $T_i$  that spans  $C_i$ . Trees  $T_i$  may be arbitrary trees spanning vertices of  $C_i$ , and their edges are not necessarily present in  $G$ . Splitting a set is achieved by specifying an edge  $e$  in  $T_i$ . The edge  $e$  is removed from  $T_i$  and the two connected components that are created specify how the set should be split (see also Figure 3). Formally speaking, the fully dynamic oracle supports the following operations:

- **distance**( $v, i$ ) – compute the approximate distance from  $v$  to the nearest vertex or color  $i$ ,

- **nearest**( $v, k$ ) – compute the approximate distance from  $v$  to the  $k$ -th nearest active color (for a constant  $k$ ),
- **activate**( $i$ ) – activate set (color)  $C_i$ ,
- **deactivate**( $i$ ) – deactivate set (color)  $C_i$ ,
- **merge**( $i, j, u, v$ ) – merge active sets  $C_i$  and  $C_j$  associated with trees  $T_i$  and  $T_j$  into an active color set  $C_l$ ,  $l \in \{i, j\}$ , and associate it with  $T_l := T_i \cup T_j \cup \{uv\}$ ,
- **split**( $l, u, v$ ) – if  $uv$  is an edge of a spanning tree  $T_l$  of an active color set  $C_l$ , split  $C_l$  into active color sets  $C_i$  and  $C_j$  associated with the two connected components of  $T_l \setminus \{uv\}$ .

We would like to remark here that in our construction of the oracles, the **nearest**( $v, k$ ) query returns the distance to the  $k$ -th nearest color, but we use the distances measured by the **distance** query. That is, it returns the  $k$ -th smallest value of **distance**( $v, i$ ) over all active colors  $i$ .

### 3.2.2 Generic distance oracles

Our oracles are based on static distance oracles that may answer vertex-to-vertex distance queries. In addition to supporting the scheme listed above, our oracles have some structural properties important for our algorithms. These properties are summarized in the following definition and we explain why we need them right below. The definition is illustrated in Fig. 4.

**Definition 3.4.** *Let  $G = (V, E, d_G)$  be a weighted graph and  $\alpha \geq 1$ . An  $\alpha$ -approximate generic oracle for  $G$  associates with every  $v \in V$ :*

- a set of portals, denoted by  $portals(v)$ , which is a subset of  $V$ ,
- and a family of pieces, denoted by  $pieces(v)$ , where each element (referred to as piece) is a weighted planar<sup>3</sup> graph on a subset of  $V$ .

We say that a vertex  $w$  is piece-visible from  $v$  if there is a piece  $\mathfrak{r} \in pieces(v)$  such that  $w \in V(\mathfrak{r})$ . The piece distance from  $v$  to  $w$  is defined as  $R_{v,w} = \min_{\mathfrak{r} \in pieces(v)} \delta_{\mathfrak{r}}(v, w)$ , and  $R_{v,w} = +\infty$  if  $w$  is not piece-visible from  $v$ . The oracle additionally stores, for each  $v \in V$  and  $p \in portals(v)$ , an approximate distance  $D_{p,v} \geq \delta_G(p, v)$ <sup>4</sup>. For every pair  $v, w \in V$  we require that  $R_{v,w} = R_{w,v}$ , and either  $R_{v,w} = \delta_G(v, w)$  or there is a portal  $p \in portals(v) \cap portals(w)$ , such that there is a walk from  $v$  to  $w$  of length at most  $\alpha \cdot \delta_G(v, w)$  that goes through  $p$ , and  $D_{v,p} + D_{w,p} \leq \alpha \delta_G(v, w)$ . The portal distance between  $v$  and  $w$  is  $\min_{p \in portals(v) \cap portals(w)} D_{v,p} + D_{w,p}$ .

Let us now introduce the notation related to oracles. We denote the set of all portals as  $\mathcal{P} = \bigcup_{v \in V} portals(v)$ , and refer to its elements as portals. We denote the family of pieces as  $\mathcal{R} = \bigcup_{v \in V} pieces(v)$ . The pieces in  $pieces(v)$  are supposed to represent distances from  $v$  to some vertices visible from it. For every vertex  $w$ , either one of piece of  $v$  contains a path to  $w$  of length  $\delta_G(v, w)$ , or there is a walk from  $v$  to  $w$  through a portal of both  $v$  and  $w$ , whose length gives an approximation of  $\delta_G(v, w)$ . Consider a complete graph  $GD = (V, \binom{V}{2}, d_{GD})$ , where  $d_{GD}(v, w)$

<sup>3</sup>The requirement that pieces are planar is not essential for the definition, but we have it here for convenience. In our oracles, pieces are either planar graphs or simply trees.

<sup>4</sup>Ideally, we would like  $D_{p,v}$  to be equal to  $\delta_G(p, v)$  and  $R_{v,w}$  equal to  $\delta_G(v, w)$ . However, in some cases  $D_{p,v}$  and  $R_{v,w}$  might be greater.




Figure 4: A generic distance oracle. Panel (a) contains an unweighted graph with two distinguished portals (marked with squares). We have that  $portals(v) = \{b, j\}$  for all vertices  $v$ . In panel (b) there are two pieces. In this example we assume that a piece  $r_i \in pieces(v)$  if  $v \in V(r_i)$ . For every  $u, v \in V$ , the minimum of piece and portal distances between  $u$  and  $v$  is at most  $2\delta(u, v)$ . For example, the piece distance between  $e$  and  $g$  is infinite, but the portal distance is 4, which is twice as much as the exact distance.

is defined as the minimum over the piece distance and the portal distance between  $v$  and  $w$ . By Definition 3.4 it is a well defined simple graph with finite weights. This is the graph of distances seen by the oracle. Even though oracle's answers are approximate with respect to  $G$ , they are exact with respect to  $GD$ , i.e., the oracle always returns the cheapest edge in  $GD$  between a vertex and a set of vertices of interest. We need this property for our algorithms to work. In particular,  $\mathbf{distance}(v, i)$  returns the cheapest edge in  $GD$  between  $v$  and  $C_i$ , and  $\mathbf{nearest}(v, k)$  returns the cheapest edge in  $GD$  between  $v$  and an active vertex of  $V \setminus (C_{i_1} \cup \dots \cup C_{i_{k-1}})$ , where  $C_{i_1} \cup \dots \cup C_{i_{k-1}}$  are the  $k - 1$  active colors that are closest to  $v$  (w.r.t.  $GD$ ).

We define a *cluster* of a portal  $p \in \mathcal{P}$  to be set of all vertices, for which it is a portal, i.e.,  $cluster(p) = \{v \in V : p \in portals(v)\}$ . If  $p$  is not a portal,  $cluster(p) = \emptyset$ . One can view the connections between vertices  $V(G)$  and portals  $\mathcal{P}$  as a bipartite graph on  $V(G) \uplus \mathcal{P}$ , where there is an edge between  $p \in \mathcal{P}$  and  $v \in V(G)$  iff  $p \in portals(v)$  or equivalently  $v \in cluster(p)$ . The number of edges in this graph is denoted as  $p_{tot} = \sum_p |cluster(p)| = \sum_v |portals(v)|$ .

### 3.2.3 From a generic oracle to a dynamic vertex-color distance oracle

We show that given a generic oracle, we may obtain both an incremental and fully dynamic vertex-color distance oracles. Their efficiency is expressed in terms of oracle parameters, such as the total size of all pieces of each vertex or the total number of portals of all vertices.

We first show an incremental construction. For every portal  $p$  we maintain the distance to the nearest vertex of every color in  $cluster(p)$ . This requires only as much space as the oracle itself. When two colors  $i$  and  $j$  are merged, we find the color assigned to a smaller number of vertices. Assume it is color  $i$ . Then, we change colors of all vertices of color  $i$  to  $j$ . This way, a vertex changes colors only  $O(\log n)$  times and each time this happens, it only needs to update information in all its portals. To find the nearest vertex of color  $i$  from a vertex  $v$ , we first check distances to all vertices of color  $i$  that are piece-visible from  $v$  and then examine paths going through portals of  $v$ . For a given portal, we find the distance to the nearest vertex of color  $i$ , which immediately



yields a candidate path from  $v$  to a vertex of color  $i$  going through a portal. The minimum of all candidate lengths gives the desired approximate distance.

A fully dynamic construction is more involved. We describe a simplified version here. The ultimate goal is also to maintain, for every portal, the distance to the nearest vertex of every color. Recall that a fully dynamic oracle associates with every color  $C_i \subseteq V$  a tree  $T_i$  spanning  $C_i$ . For every portal  $p$  we maintain a dynamic tree  $ET_i^p$ , which has the same topology as  $T_i$  and associates with every vertex  $v \in V(T_i) \subseteq V$  a key equal to  $\delta_G(p, v)$ . The dynamic trees support link and cut operations and finding minimum of all keys in a tree. To handle an update, we iterate through all portals and link/cut the corresponding tree  $ET_i^p$ . This allows us to maintain, for every portal, the nearest vertex of every color. Answering queries can then be done as in the incremental oracle.

Next, we adapt some existing distance oracles for general and planar graphs in order to obtain generic oracles. Basing on an oracle by Thorup and Zwick [37], we construct a generic oracle of stretch 3. This requires enforcing some additional properties, that may be of independent interest. From the generic oracle, by applying our generic construction, we immediately obtain a fully dynamic 3-approximate vertex-color distance oracle that handles operations in  $\tilde{O}(\sqrt{n})$  expected time. In case of planar graphs, we adapt the construction by Thorup [36], and in the end obtain an incremental  $(1 + \varepsilon)$ -approximate vertex-color distance oracle handling operations in  $O(\varepsilon^{-1} \log^2 n \log D)$  expected amortized time, where  $D$  is the stretch of the metric induced by the input graph. We also get a fully dynamic  $(1 + \varepsilon)$ -approximate vertex-color distance oracle for planar graphs that handles operations in  $\tilde{O}(\varepsilon^{-1} \sqrt{n} \log D)$  time.

### 3.3 Implementing the online algorithms efficiently

We show how to use the vertex-color distance oracles to implement the online algorithms. A  $\mu$ -approximate nearest neighbor oracle for a graph  $G = (V, E, d_G)$  yields a  $\mu$ -near metric space, which approximates  $G$ . In other words, it works as if it was an exact oracle, but the distances between vertices were given by the edge lengths in  $GD$ . However, note that the edge lengths in  $GD$  may not satisfy triangle inequality.

#### 3.3.1 Efficient decremental algorithm

In the decremental scenario we maintain an approximate Steiner tree as terminals are deleted. The algorithm is based on vertex-color distance oracles, but it uses not only the operations provided by the oracle, but also the sets of portals and pieces of each vertex. In particular, the algorithm uses the concepts of portal and piece distances.

For a fixed  $\varepsilon > 0$  we plan to maintain a  $(1 + \varepsilon)$ -approximation of  $MST(GD[S])$ . Following the decremental scheme, we set  $\eta = 1 + \lceil \varepsilon^{-1} \rceil = O(\varepsilon^{-1})$ . We start with the tree  $T$  being equal to the MST on the terminals. When a vertex  $v$  of degree more than  $\eta$  is removed, we mark it as a nonterminal vertex, but do nothing more. Otherwise, we need to remove  $v$  and reconnect the connected components that emerge into a new tree. The main challenge lies in finding the reconnecting edges efficiently.

In order to do that, we use the dynamic MSF algorithm (see Theorem 2.2) on a graph  $H$  that we maintain. Since our goal is to maintain a tree  $T$  which is an MST of the set of terminal and some nonterminal vertices, we assure that  $H$  is a subgraph of  $GD$  that contains this MST (for simplicity, we assume here that the MST is unique). Moreover, the only nonempty (i.e., containing

edges) connected component of  $H$  is composed exactly of terminals and nonterminal vertices which used to be terminals. As a result, the dynamic MSF algorithm will find the tree  $T$ .

The simplest approach to maintaining  $H$  would be to add, for every two  $u, w \in S$ , an edge  $uw$  of length  $d_{GD}(u, w)$ . This, however, would obviously be inefficient, i.e., work in linear time. Instead of that, we use the structure of the vertex-color distance oracle. For every two  $u, w \in V(T)$  which are mutually piece-visible we add to  $H$  an edge  $uw$  of length being equal to the piece distance. Moreover, we will be adding some edges corresponding to portal distances, but they will be chosen in a careful way to ensure that the number of such edges is low.

We initialize the edge set of  $H$  to be the set of all edges of the initial  $T$  and all edges between mutually piece-visible terminals. Our algorithm uses a single instance  $\mathbb{D}$  of the fully dynamic vertex-color distance oracle. We update the vertex-color distance oracle  $\mathbb{D}$ , so that it has a single active color associated with a tree equal to  $T$ . All other vertices have distinct, inactive colors. Any change in the spanning forest of  $H$  results in a constant number of modifications to  $\mathbb{D}$ .

The main challenge is what happens if a vertex of degree  $s \leq \eta$  is removed. Then the tree  $T$  decomposes into trees  $T_1, \dots, T_s$  that we need to reconnect with a set of edges of minimum total cost. We rely on the dynamic MSF algorithm here and make sure that  $H$  contains a superset of the desired edges. First of all,  $H$  surely contains all edges corresponding to piece distances, as this is maintained as an invariant of our algorithm. Thus, we only need to assure that  $H$  also contains the necessary edges corresponding to portal distances. It is easy to observe that it suffices to add to  $H$  the MST of edges corresponding to portal distances between trees  $T_i$ . More formally, we consider a complete graph  $G_c$  over trees  $T_i$ , in which the edge length between two trees is the portal distance between the nearest vertices in these trees. We add to  $H$  only edges corresponding to  $MST(G_c)$ .

This MST can be simply computed by generating the graph  $G_c$  explicitly. We go through all portals and for each portal we generate edges corresponding to all portal distances that use this portal. Note that we use the assumption that  $\eta$  and the total number of portals are relatively small.

The main factors that influence the running time are: the number of pairs of mutually piece-visible terminals (all such edges are initially added to  $H$ ), the time for computing the reconnecting edges corresponding to portal distances, and the time for updating the oracle, when a vertex of degree at most  $\eta$  is removed. Note that  $r$  delete operations may cause at most  $r$  vertices to be removed.

In practice, the update time is dominated by the time needed by the oracles for executing  $\eta = O(\varepsilon^{-1})$  `merge` and `split` operations caused by a vertex removal. In general graphs we obtain a  $(6 + \varepsilon)$ -approximate algorithm handling updates in  $O(\varepsilon^{-1} \sqrt{n} \log n)$  expected amortized time. For planar graphs, the approximation ratio is  $2 + \varepsilon$ , whereas the amortized running time of a single update amounts to  $\tilde{O}(\varepsilon^{-1.5} \sqrt{n} \log D)$ .

### 3.3.2 Efficient incremental algorithm

In the incremental scenario, the goal is to approximate the Steiner tree, as the terminals are added. Denote the terminals that are added by  $t_1, t_2, t_3, \dots$ . Let  $S_i = \{t_1, \dots, t_i\}$  and let  $T_i$  denote the tree maintained by the algorithm after adding  $i$  terminals. We maintain an approximate minimum spanning tree of  $S_i$  in graph  $GD$ .

Fix a constant  $\varepsilon > 0$ . We assume that the edge lengths in  $GD$  are powers of  $1 + \varepsilon$ . This can be achieved by rounding up edge lengths. Assume that all edge lengths of  $GD$  belong to the interval  $[1, D]$ . We define the *level* of an edge  $uv$ , as  $\text{lvl}(uv) := \log_{1+\varepsilon} d_{GD}(uv)$ . The algorithm represents




Figure 5: Trees maintained by the oracles  $\mathbb{D}_j$  corresponding to layers  $L_j$  of the tree  $T_i$ . The small numbers above edges denote their levels.

the current tree  $T_i$  by maintaining  $h = \lceil \log_{1+\varepsilon} D \rceil$  layers  $L_1, \dots, L_h$ . Layer  $L_j$  contains edges of the current tree  $T_i$  whose level is at most  $j$ . In particular, layer  $L_h$  contains all edges of  $T_i$ .

For each layer  $L_j$ , we maintain an incremental vertex-color distance oracle  $\mathbb{D}_j$  on the graph  $G$  (see Figure 5). The partition of  $V$  into colors corresponds to the connected components of a graph  $(V, L_j)$ . A color is active if and only if its elements belong to  $S_i$ .

The process of adding a terminal  $t_i$  consists of two steps. First, we find the shortest edge in  $GD$  connecting  $t_i$  to any of  $t_1, \dots, t_{i-1}$  and add this edge to  $T_{i-1}$  to obtain tree  $T_i$ . Then, we apply a sequence of  $\varepsilon$ -efficient replacements  $(e, e')$  to  $T_i$  in order to decrease its weight. Recall that we may only consider replacements in which  $e$  is incident to the newly added terminal  $t_i$ .

In order to find the replacements, we use the vertex-color distance oracles. Fix a layer number  $j$  and assume that the colors of the oracle  $\mathbb{D}_j$  reflect the layer structure, including the newly added terminal  $t_i$ . We want to find a replacement pair  $(e, e')$  such that the  $\text{lvl}(e') > j$  and  $\text{lvl}(e) \leq j$ . Denote by  $C$  the connected component of the graph  $(V, L_j)$  that contains  $t_i$  (or a color of  $t_i$  in  $\mathbb{D}_j$ ). By definition of  $C$ , it consists of edges of level at most  $j$  and the path in  $T_i$  from  $t_i$  to every  $t \notin C$  contains an edge of level  $> j$ . We find the vertex  $t \notin C$  which is the nearest to  $t_i$  by querying  $\mathbb{D}_j$ . Since  $t_i$  has color  $C$ , we issue a  $\text{nearest}(t_i, 2)$  query to find the second nearest color from  $t_i$ . Assume that we find a vertex  $t'$ . If  $\text{lvl}(t_i t') \leq j$ , we have found a replacement pair. Let  $e'$  be the heaviest edge on the path from  $t_i$  to  $t'$  in  $T_i$  and  $e = t_i t'$ . Clearly,  $(e, e')$  is a replacement pair, and since the weight of  $e$  is lower than the weight of  $e'$  and the weights can differ at least by a factor of  $1 + \varepsilon$ , this replacement is  $\varepsilon$ -efficient.

The running time of the algorithm described above depends on  $\log D$ . In order to circumvent this dependency, we observe that we do not need to replace edges whose length is at most  $\varepsilon/(2n)$

fraction of the weight of the current tree, as their total weight is negligible. This allows us to limit the number of levels that we consider to roughly  $O(\log n)$ .

One of the factors influencing the running time is the number of replacements, but it depends only on  $\varepsilon$  and  $\mu$ . For each replacement, we issue a constant number of oracle operations on each level, so the running time is, roughly speaking, a product of the number of replacements, the number of layers (roughly  $O(\log n)$ ) and the running time of a single oracle operation. In the end, we obtain  $(6 + \varepsilon)$ -approximate algorithm for general graphs handling insertions in  $\tilde{O}(\varepsilon^{-1}\sqrt{n})$  expected amortized time and a  $(2 + \varepsilon)$ -approximate algorithm for planar graphs, which processes updates in  $O(\varepsilon^{-2} \log^2 n \log(n/\varepsilon) \log D)$  expected amortized time.

### 3.3.3 Efficient fully dynamic algorithm

The fully dynamic algorithm is obtained by appropriately merging the ideas of the decremental and incremental algorithms. In fact we maintain both the invariants of the incremental and decremental algorithms.

Fix  $\varepsilon > 0$  and set  $\eta = 1 + \lceil \varepsilon^{-1} \rceil$ . We assume that all edges' lengths in  $GD$  are powers of  $1 + \varepsilon$ , which may be assured by rounding their lengths up. We maintain a tree  $T$  that spans the set of terminals and nonterminals of degree more than  $\eta$ . The tree  $T$  is maintained by using dynamic MSF algorithm on a graph  $H$  that we update. It is an MST of the set of terminals and high-degree nonterminal vertices, and, since the length of every edge is a power of  $1 + \varepsilon$ , this is equivalent to the fact that it does not admit any  $\varepsilon$ -efficient replacements.

Assume that the edge lengths in  $GD$  belong to  $[1, D]$ . Let  $h = \lfloor \log_{1+\varepsilon} D \rfloor$ . We maintain a collection of fully dynamic nearest neighbor oracles  $\mathbb{D}_1, \dots, \mathbb{D}_h$ . Recall that the *level* of an edge  $uv$  is  $\text{lvl}(uv) := \log_{1+\varepsilon} d_{GD}(uv)$ . The oracle  $\mathbb{D}_h$  contains a single active color associated with a tree equal to  $T$ . All other colors are inactive and assigned to isolated vertices. In general, oracle  $\mathbb{D}_i$  reflects the forest consisting of edges of  $T$  of level at most  $i$ . Each active color of  $\mathbb{D}_i$  corresponds to one tree of this forest. In particular the trees assigned to colors are the same as the trees in the forest. The oracles  $\mathbb{D}_i$  are counterparts of the oracles from the incremental algorithm. Since  $\mathbb{D}_h$  contains all edges of the current tree, it may be used as the oracle  $\mathbb{D}$  used in the decremental algorithm.

When a terminal  $v$  is deleted, we use exactly the same procedure as in the decremental algorithm. The only difference is that we need to update all oracles  $\mathbb{D}_i$ , so that they reflect the changes made to the tree.

On the other hand, if a vertex  $v$  is added, we use a procedure similar to the incremental algorithm. Namely, we connect  $v$  to  $T$  using the shortest edge in  $GD$  and then apply all occurring  $\varepsilon$ -replacements. In order to find all replacements, we need to consider replacements at all possible  $O(\log D)$  levels. When a replacement  $(e, e')$  is detected, we add edge  $e$  to  $H$ . Since every replacement that we find is  $\varepsilon$ -efficient, the dynamic MSF algorithm will surely update the maintained tree accordingly. Note that  $e'$  may be left in  $H$ . The vertex-color distance oracles are updated to reflect the changes done by the algorithm.

After all replacements are made, the tree  $T$  is a MST of the set of terminals and nonterminal vertices of degree more than  $\eta$ . The last step is to add to  $H$  all edges  $vx$  for all vertices  $x$  which are piece-visible from  $v$ , as this is an invariant of the decremental algorithm.

Regarding efficiency, note that every edge that is added may be replaced by an edge that is shorter by a factor of  $(1 + \varepsilon)$ . In general, an edge insertion may trigger  $O(\log_{1+\varepsilon} D)$  replacements, in amortized sense. Each replacement requires us to update oracles on each of  $O(\log_{1+\varepsilon} D)$  levels.

Thus, for each added edge, we perform  $O(\varepsilon^{-2} \log^2 D)$  **merge** and **split** operations. Moreover, since for every vertex we need to maintain edges to all vertices piece-visible by it, we have to add or remove all such edges when a vertex is added to or removed from the tree.

For general graphs, we are able to maintain  $(6 + \varepsilon)$ -approximation of the Steiner tree, handling each update in  $\tilde{O}(\varepsilon^{-2} \sqrt{n} \log^2 D)$  expected amortized time. In case of planar graphs, the approximation ratio is  $(2 + \varepsilon)$ , whereas the amortized expected running time of a single operation amounts to  $O(\varepsilon^{-2} \sqrt{n} \log^{2.5} D)$ . To deal with the big dependency on  $\log D$ , we show a general technique for decreasing the dependency on  $\log D$  from  $\log^c D$  to  $\log D$ , at the cost of adding some additional  $\varepsilon^{-1}$  and  $\log n$  factors.

## 4 Online algorithms

In this section we describe our online algorithms for Steiner tree: we show how to maintain a low-weight tree spanning the terminals, using small number of changes to the tree. In other words, this section provides all details of the results mentioned in Section 3.1. For completeness, some parts of Section 3.1 are repeated here.

Let us now give a short overview of this section. First, we observe that the distance oracles, used with conjunction with our the online algorithms, give only approximate distances and, consequently, the distances seen by the online algorithm may satisfy the triangle inequality with some slackness. To cope with that, in Section 4.1 we introduce the notion of  $\mu$ -near metric space and prove its few properties that allow us to handle them almost as typical metric spaces.

Then, in Section 4.2 we identify which properties of the maintained tree are needed to guarantee low weight. Here we essentially follow the core ideas of Imase and Waxman [28], with a more modern improvements of [30, 20, 21]: to maintain good approximation ratio, it suffices to (a) as long as we can replace an edge  $e'$  in the tree with a new one  $e$  of significantly lower cost, proceed with the replacement; (b) defer deletion of large-degree nonterminal vertices from the tree. However, we need some technical work to formally state these properties, especially in the context of *near* metric spaces. We also improve the approximation ratio in comparison to previous works.

We proceed with the actual online algorithms in subsequent sections: the decremental scheme is treated in Section 4.3, the incremental in Section 4.4, and we merge the ideas of these two algorithms to obtain a fully dynamic one in Section 4.5.

Finally, in Section 4.5.1 we discuss how to bootstrap our algorithm to avoid superfluous dependency on the stretch of the metric in the fully dynamic setting.

In the previous works [28, 30, 20, 21] the main goal was to optimize the number of changes to the tree, regardless of the actual time needed to identify them. Here we take a different view: as our ultimate goal is to minimize the *total running time* of each operation, in amortized sense, we prefer simple and non obtrusive requirements for the online algorithms, at the cost of (slight) increase in the actual number of changes. Moreover, we put more stress on optimizing the final approximation ratio of the algorithm, and identify some tradeoffs that can be made.

Consequently, in all presented online algorithms we allow some additional slackness, comparing to the algorithms of [28, 30, 20, 21]. For example, the newly added terminal is not necessarily connected to the *closest* present terminal, but *approximately closest*; a similar slackness is allowed when performing edge replacements. This relaxation, although increases the number of possible steps of the algorithm by a constant factor, allows application of approximate distance oracles, and, consequently, yields fast implementations of the considered dynamic algorithms. Together with the

presence of near metric spaces, the introduced relaxations add some technical difficulties to our proofs.

We would like also to remark that, although the recent analysis of the incremental algorithm of Imase and Waxman [28] due to Gu, Gupta and Kumar [20] perfectly fits into our incremental scheme, we were not able to apply the same ideas to our fully dynamic scheme, as it is done in [21]. The reason is that our fully dynamic model differs in a few significant aspects from the one of [21]. In [21], a newly added terminal is treated as a new vertex in the graph, with new distances given *only* to all *currently active* (not yet deleted) terminals; the remaining distances are assumed implicitly by the triangle inequality. In this manner, the analysis of [20] is directly applicable, as the terminal closest to the newly added terminal is always an active one. However, in our case such an interpretation does not make sense, as we assume the entire algorithm runs on a fixed host graph, given at the beginning. It does not seem clear whether the analysis of [20] is applicable to our model at all.

#### 4.1 $\mu$ -near metric spaces

As already discussed, in our dynamic algorithms we do not access the underlying graph directly, but instead use approximate distance oracles that can answer distance queries. The oracles we use provide us with an *approximation* of the metric closure of the input graph, that itself may not be metric, but it is always close to being metric. In the following definition we formalize this notion of closeness.

**Definition 4.1** ( $\mu$ -near metric space). *We say that a complete graph  $GD = (V, \binom{V}{2}, d_{GD})$  is an  $\mu$ -near metric space if there exists a metric space  $\overline{G} = (V, \binom{V}{2}, d_{\overline{G}})$  on  $V$  such that for any  $u, v \in V$  we have*

$$d_{\overline{G}}(uv) \leq d_{GD}(uv) \leq \mu d_{\overline{G}}(uv).$$

*We additionally say that  $\overline{G}$  is approximated by  $GD$ .*

Intuitively, a  $\mu$ -near metric space behaves almost like a metric space, but a factor of  $\mu$  pops up if we use the triangle inequality.

**Lemma 4.2.** *Let  $GD = (V, \binom{V}{2}, d_{GD})$  be a  $\mu$ -near metric space. For any  $k \geq 1$  and any sequence  $v_0, v_1, v_2, \dots, v_k \in V$  it holds that*

$$d_{GD}(v_0 v_k) \leq \mu \sum_{i=1}^k d_{GD}(v_{i-1} v_i).$$

*Proof.* Assume  $GD$  approximates metric space  $\overline{G} = (V, \binom{V}{2}, d_{\overline{G}})$ . Then

$$d_{GD}(v_0 v_k) \leq \mu d_{\overline{G}}(v_0 v_k) \leq \mu \sum_{i=1}^k d_{\overline{G}}(v_{i-1} v_i) \leq \mu \sum_{i=1}^k d_{GD}(v_{i-1} v_i).$$

□

The analysis of the approximation factors we obtain relies on Lemma 2.1. So the ultimate goal is to maintain a good approximation of  $MST(\overline{G}[S])$ .

However, as the algorithm is given access only to a graph  $GD$ , which is an approximation of  $G$ , in subsequent sections we instead present how to maintain a good approximation of  $MST(GD[S])$  by maintaining some tree  $T$  spanning  $S$  in  $GD$ . Note that, as  $T$  is allowed to use some non-terminal vertices, it may happen that actually  $d_{GD}(T) < d_{GD}(MST(GD[S]))$ . In what follows, we say that  $T$  is a  $\alpha$ -approximation of  $MST(GD[S])$  if  $d_{GD}(T) \leq \alpha d_{GD}(MST(GD[S]))$ .

Let us now formally show that approximating  $MST(GD[S])$  is sufficient for our needs.

**Lemma 4.3.** *Assume  $GD = (V, \binom{V}{2}, d_{GD})$  is a  $\mu$ -near metric space approximating a metric space  $\bar{G} = (V, \binom{V}{2}, d_{\bar{G}})$ . Let  $S \subseteq V$  be a terminal set, and let  $T$  be a tree spanning  $S$  and being a  $\alpha$ -approximation of the minimum spanning tree of  $GD[S]$  in the complete graph  $GD$  (i.e., with regards to weights  $d_{GD}$ ). Then  $T$  is a  $2\alpha\mu$ -approximation of a minimum Steiner tree connecting  $S$  in  $\bar{G}$  (i.e., with regards to weights  $d_{\bar{G}}$ ).*

*Proof.* Let  $\bar{T}_{MST}$  be a minimum spanning tree of  $\bar{G}[S]$  and let  $\bar{T}$  be a minimum Steiner tree connecting  $S$  in  $\bar{G}$  (i.e., both  $\bar{T}_{MST}$  and  $\bar{T}$  are minimum with regards to weights  $d_{\bar{G}}$ ). As  $\bar{G}$  is a metric, by Lemma 2.1, we have  $d_{\bar{G}}(\bar{T}_{MST}) \leq 2d_{\bar{G}}(\bar{T})$ . Let  $T_{MST}$  be a minimum spanning tree of  $GD[S]$  (i.e., with regards to weights  $d_{GD}$ ). Using the definition of a  $\mu$ -near metric space we obtain:

$$d_{\bar{G}}(T) \leq d_{GD}(T) \leq \alpha d_{GD}(T_{MST}) \leq \alpha d_{GD}(\bar{T}_{MST}) \leq \alpha \mu d_{\bar{G}}(\bar{T}_{MST}) \leq 2\alpha \mu d_{\bar{G}}(\bar{T}).$$

□

Lemma 4.3 allows us to henceforth focus on a weighted complete graph  $GD$  and approximate  $MST(GD[S])$ , in most cases forgetting about the underlying approximated metric space  $\bar{G}$ .

#### 4.1.1 Analysis of the greedy approach in near metric spaces

As discussed at the beginning of this section, we would like to apply the recent analysis of [20] of the greedy approach for incremental algorithm to our case. However, the analysis of [20] is performed in the presence of a metric space, and our algorithm perceives the input graph via distance oracles, and hence operates on a *near* metric space. In this short section, we adapt the results of [20] to near metric spaces.

In the metric case, the following is proven in [20].

**Theorem 4.4** ([20]). *Consider a set of terminals  $S = \{t_1, t_2, \dots, t_n\}$  in a metric space  $\bar{G} = (V, \binom{V}{2}, d_{\bar{G}})$ . For any  $2 \leq k \leq n$ , let  $g_k$  be the minimum-cost (w.r.t.  $d_{\bar{G}}$ ) edge connecting  $t_k$  with  $\{t_1, t_2, \dots, t_{k-1}\}$ . Then*

$$\prod_{k=2}^n d_{\bar{G}}(g_k) \leq 4^{n-1} \prod_{e \in E(MST(\bar{G}[S]))} d_{\bar{G}}(e).$$

As shown below, in the near-metric setting, the constant 4 increases to  $4\mu$ .

**Theorem 4.5.** *Consider a set of terminals  $S = \{t_1, t_2, \dots, t_n\}$  in a  $\mu$ -near metric space  $GD = (V, \binom{V}{2}, d_{GD})$ . For any  $2 \leq k \leq n$ , let  $f_k$  be the minimum-cost (w.r.t.  $d_{GD}$ ) edge connecting  $t_k$  with  $\{t_1, t_2, \dots, t_{k-1}\}$ . Then*

$$\prod_{k=2}^n d_{GD}(f_k) \leq (4\mu)^{n-1} \prod_{e \in E(MST(GD[S]))} d_{GD}(e).$$

*Proof.* Assume  $GD$  approximates metric space  $\bar{G} = (V, \binom{V}{2}, d_{\bar{G}})$ . For any  $2 \leq k \leq n$ , let  $g_k$  be defined as in Theorem 4.4, that is,  $g_k$  is the edge connecting  $t_k$  with  $\{t_1, t_2, \dots, t_{k-1}\}$  of minimum possible cost in the metric  $d_{\bar{G}}$ . By the definition of a  $\mu$ -near metric space and the choice of  $f_k$  we have

$$d_{GD}(f_k) \leq d_{GD}(g_k) \leq \mu d_{\bar{G}}(g_k). \quad (1)$$

Moreover, again by the definition of a  $\mu$ -near metric space we obtain:

$$\prod_{e \in E(MST(GD[S]))} d_{GD}(e) \geq \prod_{e \in E(MST(GD[S]))} d_{\bar{G}}(e) \geq \prod_{e \in E(MST(\bar{G}[S]))} d_{\bar{G}}(e). \quad (2)$$

The theorem follows by combining (1) and (2) with Theorem 4.4.  $\square$

## 4.2 Properties guaranteeing good approximation factors

Our goal is to maintain a tree  $T$  that spans the current set of terminals  $S$  in the near-metric space  $GD$ . In this section we identify the properties of the tree  $T$  that guarantee good approximation factor.

Following [28], when a terminal  $t$  is added to  $S$ , we can update the current tree by connecting  $t$  with  $T$  using the cheapest connecting edge. This does not lead to a good approximation of  $MST(GD[S \cup \{t\}])$  because some other edges incident to  $t$  could possibly be used to replace expensive edges in  $T$ . In our algorithms, we repeatedly replace tree edges with non-tree edges (assuring that after the replacement, we obtain a tree spanning  $S$ ) until we reach a point at which we can be sure that the current tree is a good approximation of  $MST(GD[S])$ . We now formalize the notion of replacement.

Let  $T$  be a tree in  $GD$ . For any edge  $e$  with both endpoints in  $V(T)$ , we say that an edge  $e_T \in E(T)$  is a *friend* of  $e$  (with respect to  $T$ ) if  $e_T$  lies on the unique path in  $T$  between the endpoints of  $e$ . For a friend  $e_T$  of  $e$  with regards to  $T$ ,  $T' := (T \setminus \{e_T\}) \cup \{e\}$  is a tree that spans  $V(T)$  as well, with cost  $d_{GD}(T) - (d_{GD}(e_T) - d_{GD}(e))$ . We say that  $T'$  is created from  $T$  by *replacing* the edge  $e_T$  with  $e$ , and  $(e, e_T)$  is a *replacement pair* in  $T$ .

**Definition 4.6** (heavy, efficient and good replacement). *Let  $c > \tau \geq 0$  be constants and let  $T$ ,  $e$  and  $e_T$  be defined as above. We say that the  $(e, e_T)$  pair is a*

1.  $\tau$ -heavy replacement if  $d_{GD}(e_T) > \tau d_{GD}(T)/|V|$ ;
2.  $\tau$ -efficient replacement if  $(1 + \tau)d_{GD}(e) < d_{GD}(e_T)$ ;
3.  $(\tau, c)$ -good replacement if it is both  $\tau$ -efficient and  $(\tau/c)$ -heavy.

Lemmas 4.8 and 4.9 that follow will later be used to expose the properties of the tree we need to maintain in order to obtain a constant approximation of Steiner tree. For their proof, we need the following variant of the classical Hall's theorem.

**Theorem 4.7** (Hall's theorem). *Let  $H$  be a bipartite graph with bipartition classes  $A \uplus B = V(H)$ . Assume that there exist two positive integers  $p, q$  such that for each  $X \subseteq A$  it holds that  $|\Gamma_H(X)| \geq \frac{p}{q}|X|$ . Then there exists a function  $\pi : A \times \{1, 2, \dots, p\} \rightarrow B$  such that (i)  $a\pi(a, i) \in E(H)$  for every  $a \in A$  and  $1 \leq i \leq p$ , and (ii)  $|\pi^{-1}(b)| \leq q$  for each  $b \in B$ .*



**Lemma 4.8.** *Let  $c > \tau \geq 0$  be constants and let  $T$  be a tree in a complete graph  $GD$ . If  $T$  does not admit any  $(\tau, c)$ -good replacement, then it is a  $c(1 + \tau)/(c - \tau)$ -approximation of minimum spanning tree of  $V(T)$ . In particular, if  $T$  does not admit any  $\tau$ -efficient replacement, then it is a  $(1 + \tau)$ -approximation of minimum spanning tree on  $V(T)$ .*

*Proof.* Let  $T_{MST}$  be a minimum spanning tree of  $V(T)$ . We claim that there exists a bijection  $\pi : E(T) \rightarrow E(T_{MST})$  such that  $e$  is a friend of  $\pi(e)$  with regards to  $T$ , for any  $e \in E(T)$ . Indeed, for any  $X \subseteq E(T)$ , there are  $|X| + 1$  connected components of  $T \setminus X$  and, as  $T$  and  $T_{MST}$  spans the same set of vertices, at least  $|X|$  edges of  $T_{MST}$  connect two distinct connected components of  $T \setminus X$ . As each such edge has a friend in  $X$ , the claim follows from Hall's theorem.

Since  $T$  does not admit a  $(\tau, c)$ -good replacement, for any  $e \in E(T)$ , either  $d_{GD}(e) \leq \tau d_{GD}(T)/(c|V|)$  or  $d_{GD}(e) \leq (1 + \tau)d_{GD}(\pi(e))$ . The cost of the edges of the first type sum up to at most  $\tau d_{GD}(T)/c$ , whereas the cost of the edges of the second type sum up to at most  $(1 + \tau)d_{GD}(T_{MST})$ . We infer

$$d_{GD}(T) \leq \frac{\tau}{c}d_{GD}(T) + (1 + \tau)d_{GD}(T_{MST})$$

and the lemma follows.  $\square$

We now move to the core arguments needed for a deletion step. Imase and Waxman [28] proved that deferring the deletion of non-terminal vertices from the tree  $T$  as long as they have degree at least three incurs only an additional factor of two in the approximation guarantee. We present here a new, simplified proof of a generalization of this fact to arbitrary degree threshold.

**Lemma 4.9.** *Let  $\eta \geq 2$  be an integer and let  $T$  be a tree spanning  $N \cup S$  where  $\deg_T(v) > \eta$  for any  $v \in N$ . Let  $T_S$  be any tree spanning  $S$ . Then there exists a function  $\pi : E(T) \times \{1, 2, \dots, \eta - 1\} \rightarrow E(T_S)$  such that*

1.  $e$  is a friend of  $\pi(e, i)$  with respect to  $T$ , for any  $e \in E(T)$  and  $1 \leq i \leq \eta - 1$ ; and
2.  $|\pi^{-1}(e')| \leq \eta$  for any edge  $e'$  of  $T_S$ .

*Proof.* We use Hall's theorem. Let  $X \subseteq E(T)$ ; we need to prove that there are at least  $\frac{\eta-1}{\eta}|X|$  edges of  $E(T_S)$  that have a friend in  $X$ . Let  $\mathcal{T}$  be the family of connected components of the forest  $T \setminus X$ , and  $\mathcal{T}_{-S} \subseteq \mathcal{T}$  be the family of these components of  $\mathcal{T}$  that do not contain any vertex of  $S$ . By the properties of the set  $N$ , each component of  $\mathcal{T}_{-S}$  is adjacent to at least  $\eta + 1$  edges of  $X$ . Consequently, since  $T$  is a tree, we have

$$|\mathcal{T}_{-S}| < |\mathcal{T}|/\eta = (|X| + 1)/\eta.$$

Hence, at least  $\frac{\eta-1}{\eta}|X| + 1$  components of  $\mathcal{T}$  contain a vertex of  $S$ . To connect these vertices,  $T_S$  needs to contain at least  $\frac{\eta-1}{\eta}|X|$  edges with endpoints in different components of  $T \setminus X$ , and each such edge has at least one friend in  $X$ . The lemma follows.  $\square$

With a similar strategy, we provide a proof Lemma 7.3, used later on in Section 7, but proven here for convenience.

*Proof of Lemma 7.3.* Recall that  $T_{MST} = MST(\tilde{G}[S])$ . Define a tree  $\hat{T}$  as follows: for each edge  $uw \in E(T_{MST})$ , take its corresponding edges  $ux, wx \in E(B)$ ,  $d_B(ux) + d_B(wx) = \tilde{\delta}_B(uw)$ , and add

them to  $\widehat{T}$ . In this way we obtain a tree  $\widehat{T}$  in  $B$  of length at most  $\widetilde{\delta}_B(T_{MST})$  that spans a superset of  $S$ . Denote  $\widehat{S} = V(\widehat{T})$ . It suffices to prove that  $d_B(T) \leq 2d_B(\widehat{T})$ .

Let  $\mathcal{L}_{N \setminus \widehat{S}}$  be the set of leaf vertices of  $T'$  that are contained in  $N \setminus \widehat{S}$ , and let  $T_{\widehat{S}} = T' \setminus \mathcal{L}_{N \setminus \widehat{S}}$ . Clearly,  $T$  is a subtree of  $T_{\widehat{S}}$ , so it suffices to prove  $d_B(T_{\widehat{S}}) \leq 2d_B(\widehat{T})$ .

To this end, we show that there exists a function  $\pi : E(T_{\widehat{S}}) \rightarrow E(\widehat{T})$  such that

1.  $e$  is a friend of  $\pi(e)$  with respect to  $T_{\widehat{S}}$ , for any  $e \in E(T_{\widehat{S}})$ ;
2.  $|\pi^{-1}(e')| \leq 2$  for every  $e' \in E(\widehat{T})$ .

Observe that  $T_{\widehat{S}}$  is a minimum spanning tree of  $B[\widehat{S}]$ , since  $T'$  is a minimum spanning tree as well. Hence, for such a function  $\pi$  we would have  $d_G(e) \leq d_B(\pi(e))$  for any  $e \in E(T_{\widehat{S}})$  and, consequently,  $d_B(T_{\widehat{S}}) \leq 2d_B(\widehat{T})$ , concluding the proof of the lemma.

By Hall's theorem, it suffices to prove that for any  $X \subseteq E(T_{\widehat{S}})$ , at least  $|X|/2$  edges of  $E(\widehat{T})$  have a friend in  $X$  (with respect to  $T_{\widehat{S}}$ ). To this end, consider any  $X \subseteq E(T_{\widehat{S}})$  and let  $\mathcal{T}$  be the family of connected components of the forest  $T_{\widehat{S}} \setminus X$ , and  $\mathcal{T}_{-\widehat{S}} \subseteq \mathcal{T}$  be the family of these components of  $\mathcal{T}$  that do not contain any vertex of  $\widehat{S}$ . Observe that each component  $Q$  of  $\mathcal{T}_{-\widehat{S}}$  is either adjacent to at least three edges of  $X$ , or consists of a single vertex of  $N$  and is adjacent in  $T$  to exactly two vertices of  $S \subseteq \widehat{S}$ . In the latter case, we contract  $Q$  together with its two neighboring components, decreasing  $|\mathcal{T}|$  by 2 and  $|\mathcal{T}_{-\widehat{S}}|$  by one. After  $k$  contractions, we have as in Lemma 4.9

$$|\mathcal{T}_{-\widehat{S}}| - k < (|\mathcal{T}| - 2k)/2.$$

Consequently,  $|\mathcal{T}_{-\widehat{S}}| < |\mathcal{T}|/2 = (|X| + 1)/2$ . Thus, at least  $|X|/2$  edges of  $\widehat{T}$  are needed to connect the components of  $\mathcal{T} \setminus \mathcal{T}_{-\widehat{S}}$ , and all such edges have some friend in  $X$ . This finishes the proof of an existence of the function  $\pi$ , and concludes the proof of the lemma.  $\square$

Lemma 4.10 that follows summarizes the properties we need to maintain a good tree. With respect to terminal addition, it says that it suffices to perform good replacements. With respect to terminal removal, it allows us to postpone the deletion of a terminal from the approximate Steiner tree as long as it has degree larger than some fixed integer threshold  $\eta$ .

**Lemma 4.10.** *Let  $\tau \geq 0$  be a constant,  $\eta \geq 2$  be an integer,  $GD = (V, \binom{V}{2}, d_{GD})$  be a complete weighted graph and  $S \subseteq V$ . Let  $T_{MST} = MST(GD[S])$  be a minimum spanning tree of  $S$  and let  $T$  be a tree spanning  $S \cup N$  such that for any  $v \in N$  we have  $\deg_T(v) > \eta$ . Then the following hold.*

1. *If  $N = \emptyset$  and  $T$  does not admit a  $(\tau/2, 1 + \tau)$ -good replacement, then  $d_{GD}(T) \leq (1 + \tau)d_{GD}(T_{MST})$ .*
2. *If  $T$  is a minimum spanning tree of  $S \cup N$ , then  $d_{GD}(T) \leq \frac{\eta}{\eta-1}d_{GD}(T_{MST})$ .*
3. *If  $T$  does not admit a  $\tau$ -efficient replacement then  $d_{GD}(T) \leq \frac{\eta}{\eta-1}(1 + \tau)d_{GD}(T_{MST})$ .*

*Proof.* Point 1 follows from Lemma 4.8. To prove the remaining points, invoke Lemma 4.9 on trees  $T$  and  $T_{MST}$  with threshold  $\eta$  to obtain a mapping  $\pi : E(T) \times \{1, 2, \dots, \eta - 1\} \rightarrow E(T_{MST})$ . If  $T$  is a minimum spanning tree of  $S \cup N$ ,  $d_{GD}(e) \leq d_{GD}(\pi(e, i))$  for any  $e \in E(T)$  and  $1 \leq i \leq \eta - 1$ ; Point 2 follows by summation over the entire domain of  $\pi$ . If  $T$  does not admit a  $\tau$ -efficient replacement,  $d_{GD}(e) \leq (1 + \tau)d_{GD}(\pi(e, i))$  for any  $e \in E(T)$  and  $1 \leq i \leq \eta - 1$  and Point 3 follows similarly.  $\square$

Knowing which properties we need our tree to satisfy, we may now investigate how to maintain these properties when confronted with addition and deletion of a terminal to and from the terminal set. In the next subsections we show how to modify the tree to preserve the desired properties.

### 4.3 Decremental online algorithm

In the decremental scheme, we study the case when terminals are removed from the terminal set. The main idea is to use Lemma 4.9 and to maintain the minimum spanning tree on the terminal set, but to postpone the deletion of terminals that are of degree above some fixed threshold  $\eta \geq 2$  in this spanning tree.

When a vertex  $v$  of degree  $s \leq \eta$  is deleted from the tree  $T$ , the tree breaks into  $s$  components  $\mathcal{T}^v = \{T_1, T_2, \dots, T_s\}$  that need to be reconnected. The natural idea is use for this task a set of edges of  $GD$  of minimum possible total weight. That is, we first for each  $1 \leq i < j \leq s$  identify an edge  $e_{ij}$  of minimum possible cost in  $GD$  among edges between  $T_i$  and  $T_j$ . Then, we construct an auxiliary complete graph  $H_v$  with vertex set  $\mathcal{T}^v$  and edge cost  $d(T_i T_j) := d_{GD}(e_{ij})$  and find a minimum spanning tree  $MST(H_v)$  of this graph. We define  $F^v := \{e_{ij} : T_i T_j \in E(MST(H_v))\}$  to be the set of reconnecting edges and  $T' := (T \setminus \{v\}) \cup F^v$  to be the reconnected tree  $T$ . We now proceed with a formal argumentation that this natural idea indeed works as expected.

**Lemma 4.11.** *Let  $T$  be a tree in a complete graph  $GD = (V, \binom{V}{2}, d_{GD})$ , let  $\mathcal{T}$  be a family of pairwise vertex-disjoint subtrees of  $T$  such that there exists a subtree  $T^c$  of  $T$  with  $|V(T^c) \cap V(\widehat{T})| = 1$  for any  $\widehat{T} \in \mathcal{T}$ , and let  $F$  be a set of edges of  $GD$  such that (a) each edge of  $F$  connects two different trees of  $\mathcal{T}$ , and (b)  $T' := F \cup \bigcup \mathcal{T}$  is a tree. Let  $(e, e')$  be a replacement pair in  $T'$ . Then one of the following holds:*

1.  $(e, e')$  is a replacement pair in  $T$  as well;
2.  $e' \in F$ ; or
3. there exists  $f \in F$  such that  $(e, f)$  is a replacement pair in  $T'$  and  $(f, e')$  is a replacement pair in  $T$ .

*Proof.* Let  $P$  be the unique path between the endpoints of  $e$  in the tree  $T'$ . By the construction of  $T'$ ,  $P$  consists of subpaths  $P_1, P_2, \dots, P_s$ , where  $P_i$  is a path in  $T_i \in \mathcal{T}$ , and edges  $e_2, e_3, \dots, e_s \in F$ , where each  $e_i$  connects the endpoint of  $P_{i-1}$  with the starting point of  $P_i$ .

If  $s = 1$  then  $P = P_1$  and hence  $(e, e')$  is a replacement pair in  $T$  as well. If  $e' = e_i$  for some  $2 \leq i \leq s$  then  $e' \in F$  and the lemma is proven. Otherwise, as  $e'$  lies on  $P$ ,  $e' \in E(P_j)$  for some  $1 \leq j \leq s$ . Observe that for any  $2 \leq i \leq s$ , the pair  $(e, e_i)$  is a replacement pair in  $T'$ . To finish the proof of the lemma it suffices to show that either  $(e, e')$  or  $(e_i, e')$  for some  $2 \leq i \leq s$  is a replacement pair in  $T$ .

Let  $x$  be the first vertex of  $P_j$ ,  $y$  be the last vertex of  $P_j$ , and let  $z$  be the unique vertex of  $V(T^c) \cap V(T_j)$ . Since  $e' \in E(P_j)$ , the edge  $e'$  lies either on the unique path between  $x$  and  $z$  in  $T_j$ , or on the unique path between  $y$  and  $z$ . In the first case, observe that  $(e_j, e')$  is a replacement pair in  $T$  if  $j > 1$ , and  $(e, e')$  is a replacement pair in  $T$  if  $j = 1$ . In the second case, symmetrically, observe that  $(e_{j+1}, e')$  is a replacement pair in  $T$  if  $j < s$  and  $(e, e')$  is a replacement pair in  $T$  if  $j = s$ .  $\square$

**Lemma 4.12.** *Let  $T$  be a tree in a complete graph  $GD = (V, \binom{V}{2}, d_{GD})$  that does not admit a  $\tau$ -efficient replacement and let  $v \in V(T)$ . Let  $\mathcal{T}$  be the family of connected components of  $T \setminus \{v\}$  and define the reconnecting edges  $F^v$  and the reconnected tree  $T' = (T \setminus \{v\}) \cup F^v$  as at the beginning of this section. Then  $T'$  does not admit a  $\tau$ -efficient replacement as well.*

*Proof.* Let  $(e, e')$  be a replacement pair in  $T'$ ; our goal is to prove that it is not  $\tau$ -efficient. Apply Lemma 4.11 to the pair  $(e, e')$ , the family  $\mathcal{T}$ , tree  $T^c = T[\{v\} \cup \Gamma_T(v)]$  and set  $F^v$ , and consider the three possible outcomes.

If  $(e, e')$  is a replacement pair in  $T$  as well, then clearly it is not  $\tau$ -efficient. Otherwise, observe first that for every  $g \in F^v$  such that  $(e, g)$  is a replacement pair in  $T'$ , it follows from the choice of  $F^v$  that  $d_{GD}(e) \geq d_{GD}(g)$ . Hence, if  $e' \in F^v$  then  $d_{GD}(e) \geq d_{GD}(e')$  and  $(e, e')$  is not  $\tau$ -efficient.

Consider now an edge  $f \in F^v$  promised in the remaining case of Lemma 4.11. As  $(e, f)$  is a replacement pair in  $T'$ , we have from the previous reasoning that  $d_{GD}(e) \geq d_{GD}(f)$ . As  $(f, e')$  is a replacement pair in  $T$ , it is not  $\tau$ -efficient. Therefore,  $(e, e')$  is also not  $\tau$ -efficient and the lemma is proven.  $\square$

We are ready to describe our dynamic algorithm. Fix an accuracy parameter  $\varepsilon > 0$  and denote  $\eta := 1 + \lceil \varepsilon^{-1} \rceil$ , that is,  $\eta$  is the smallest positive integer for which  $\frac{\eta}{\eta-1} \leq 1 + \varepsilon$ . Thus, by Lemma 4.10, Point 2, to obtain a  $(1 + \varepsilon)$ -approximation of  $MST(GD[S])$  where  $S$  is the set of currently active terminals, it is sufficient to maintain a tree  $T$  such that

1.  $S \subseteq V(T)$ ;
2.  $T$  is a minimum spanning tree of  $V(T)$  (i.e.,  $T$  does not admit any 0-efficient replacement);
3. any non-terminal vertex of  $T$  is of degree larger than  $\eta$  in  $T$ .

To achieve this goal, we describe a procedure  $\text{REMOVE}_\eta(v)$  that checks if a vertex  $v$  should be removed from the currently maintained tree  $T$  and, if this is the case, performs the removal.

1. If  $v$  is a terminal,  $v \notin V(T)$  or  $\text{deg}_T(v) > \eta$ , do nothing.
2. Otherwise, delete  $v$  and its incident edges from  $T$ , and compute the reconnecting edges  $F^v$  and the reconnected tree  $T' = (T \setminus \{v\}) \cup F^v$  as described in the beginning of this section. Replace  $T$  with  $T'$ . Then, recursively invoke  $\text{REMOVE}_\eta$  on all ex-neighbors of  $v$  in  $T$ .

If we delete  $v$  from the terminal set, we first mark  $v$  as non-terminal and then call  $\text{REMOVE}_\eta(v)$ . Lemma 4.12 for  $\tau = 0$  ensures that, if we initialize  $T$  as the minimum spanning tree on the terminal set, then  $T$  satisfies all aforementioned three properties and is a  $(1 + \varepsilon)$ -approximation of  $MST(GD[S])$ , where  $S$  is the current set of terminals. By Lemma 4.3, we obtain in this manner a  $2(1 + \varepsilon)\mu$ -approximation of minimum Steiner tree on  $S$  in the metric  $\overline{G}$  that is approximated by  $GD$ .

**Lemma 4.13.** *Assume that an algorithm implements the decremental scheme. Then, it maintains a  $2(1 + \varepsilon)\mu$ -approximation of minimum Steiner tree of  $S$  in  $G$ . In the course of the first  $r$  deletions, the algorithm performs at most  $r$  calls to the procedure  $\text{REMOVE}_\eta$  that result in a modification of the maintained tree. Moreover, each vertex removed by the procedure  $\text{REMOVE}_\eta$  is of degree at most  $\eta = O(\varepsilon^{-1})$  in the currently maintained tree.*

*Proof.* The approximation ratio follows from the discussion above. The claim about efficiency follows from an observation that  $\text{REMOVE}_\eta$  modifies the maintained tree only when called upon an non-terminal vertex in the tree of degree at most  $\eta$ , and each deletion operation introduces only one non-terminal vertex.  $\square$

#### 4.4 Incremental online algorithm

We consider here the case when new terminals are added to the terminal set and we are to update the currently maintained tree. Following Imase and Waxman [28] and Megow et al [30], we would like to first connect the new terminal to tree we already have, and then try to apply all occurring  $(\tau/2, 1+\tau)$ -good replacement pairs, for some  $\tau > 0$ . The main technical contribution of this section is a proof that it is sufficient to look only at replacement pairs  $(e, e')$  where  $e$  is incident to the newly added vertex.

However, as we would like to use our scheme in conjunction with approximate distance oracles, we need to introduce some additional relaxation in terms of efficiency of conducted replacements. We would like to require that the algorithm does not leave any  $(\tau/2, 1+\tau)$ -good replacement pair, but any replacement made by the algorithm is only  $\varsigma$ -efficient for some  $0 < \varsigma \leq \tau/2$ . The first requirement yields the approximation guarantee, while the latter controls the number of performed replacements.

We now proceed with formal description of the incremental scheme. Assume the algorithm operates on a  $\mu$ -near metric space  $GD = (V, \binom{V}{2}, d_{GD})$  and we are additionally equipped with parameters  $\tau \geq 2\varsigma > 0$ . Define  $c = 2\mu(1+\tau)^2$ .

At one step, given a tree  $T$  and a new terminal vertex  $v \in V \setminus V(T)$ , we:

1. add an edge to  $T$  connecting  $v$  with  $V(T)$  of cost  $\min_{u \in V(T)} d_{GD}(uv)$ ;
2. apply a sequence of  $\varsigma$ -efficient replacement pairs  $(e, e')$  that satisfy the following additional property:  $e'$  is a friend of  $e$  in the currently maintained tree of maximum possible cost;
3. after all the replacements, we require that there does not exist a  $(\tau/2, c)$ -good replacement pair  $(e, e')$  with  $e$  incident to  $v$ .

To argue about the efficiency of incremental scheme we need to use the fact that graph  $GD$  we work on is an  $\mu$ -near metric space. The next lemma shows that the cost of the optimal tree cannot decrease much during the course of the algorithm.

**Lemma 4.14.** *Let  $T, T'$  be two trees in a  $\mu$ -near metric space  $GD = (V, \binom{V}{2}, d_{GD})$ , such that  $V(T) \subseteq V(T')$  and  $T$  is a  $\alpha$ -approximation of a minimum spanning tree on  $V(T)$ , for some  $\alpha \geq 1$ . Then  $d_{GD}(T') \geq d_{GD}(T)/(2\alpha\mu)$ .*

*Proof.* Let  $\bar{G} = (V, \binom{V}{2}, d_{\bar{G}})$  be a metric space approximated by  $GD$ . Let  $\bar{T}_{MST} = MST(\bar{G}[V(T)])$  be the minimum spanning tree of  $V(T)$  in  $\bar{G}$  (i.e., with regards to weights  $d_{\bar{G}}$ ) and let  $\bar{T}_S$  be the minimum Steiner tree on the terminal set  $S = V(T)$  in  $\bar{G}$ . Let  $T_{MST} = MST(GD[V(T)])$  be the minimum spanning tree of  $V(T)$  in  $GD$ . Note that  $\bar{T}_{MST}$  is a 2-approximation of the minimum Steiner tree on  $V(T)$ , whereas  $T'$  is a (not necessarily optimal) Steiner tree of  $V(T)$ . Hence

$$\begin{aligned} d_{GD}(T') &\geq d_{\bar{G}}(T') \geq d_{\bar{G}}(\bar{T}_S) \geq d_{\bar{G}}(\bar{T}_{MST})/2 \geq d_{GD}(\bar{T}_{MST})/(2\mu) \\ &\geq d_{GD}(T_{MST})/(2\mu) \geq d_{GD}(T)/(2\alpha\mu). \end{aligned}$$

$\square$

In what follows we prove that it is sufficient to find replacement edges adjacent to the new terminal.

**Lemma 4.15.** *If an incremental algorithm satisfies the scheme above in an  $\mu$ -near metric space  $GD$  then after each step tree  $T$  does not admit a  $(\tau/2, 1 + \tau)$ -good replacement pair and, consequently, is a  $(1 + \tau)$ -approximation of  $MST(GD[S])$  and a  $2(1 + \tau)\mu$ -approximation of minimum Steiner tree spanning  $S$  in  $G$ .*

*Proof.* Let  $t_1, t_2, t_3, \dots$  be a sequence of terminals added to the terminal set,  $S_k = \{t_1, t_2, \dots, t_k\}$  and  $T_k$  be the tree maintained by the algorithm after  $k$ -th terminal addition, i.e.,  $V(T_k) = S_k$ . We inductively prove that  $T_k$  does not admit a  $(\tau/2, 1 + \tau)$ -good replacement pair; the second part of the lemma follows directly from Lemmas 4.10 and 4.3. The claim for  $k \leq 1$  is obvious; let us assume that  $k \geq 2$  and that the statement is true for the trees  $T_1, T_2, \dots, T_{k-1}$ .

First, note that by Lemma 4.8, any tree  $T_i$  for  $i < k$  is a  $(1 + \tau)$ -approximation of a minimum spanning tree of  $GD[S_i]$ . By Lemma 4.14,  $d_{GD}(T_k) \geq d_{GD}(T_i)/(2\mu(1 + \tau))$  for any  $i < k$ .

Assume that there exists a  $(\tau/2)$ -efficient replacement pair  $(e, e')$  in  $T_k$ . Assume  $e = t_i t_j$  where  $1 \leq i < j \leq k$ . Let  $e_j$  be an edge of maximum cost on the unique path between  $t_i$  and  $t_j$  in  $T_j$ . We now prove the following.

**Claim 4.16.** *For any  $j \leq j' \leq k$ , the unique path between  $t_i$  and  $t_j$  in  $T_{j'}$  does not contain edges of cost greater or equal to  $d_{GD}(e_j)$ .*

*Proof.* We use induction on  $j'$ . In the base case  $j' = j$  the claim follows from the choice of  $e_j$ . In the inductive step, consider a single replacement pair  $(f, f')$  applied by the algorithm. Denote by  $T^1$  (resp.  $T^2$ ) the maintained tree before (resp. after) the replacement pair was applied, and by  $P_1$  (resp.  $P_2$ ) the unique path connecting  $t_i$  and  $t_j$  in  $T^1$  (resp.  $T^2$ ). Assume that no edge on  $P_1$  is of cost greater or equal to  $d_{GD}(e_j)$ ; we would like to prove the same statement for  $P_2$ . If  $P_1 = P_2$  there is nothing to prove so assume otherwise. We have  $f \in E(P_2) \setminus E(P_1)$ ,  $f' \in E(P_1) \setminus E(P_2)$  and, by the efficiency of  $(f, f')$ , we have  $d_{GD}(f) \leq d_{GD}(f') \leq d_{GD}(e_j)$ . Moreover, observe that each edge  $f'' \in E(P_2) \setminus (\{f\} \cup E(P_1))$  is a friend of  $f$  in the tree  $T^1$ . By the choice of  $f'$ , we have  $d_{GD}(f'') \leq d_{GD}(f') \leq d_{GD}(e_j)$  and the claim follows.  $\square$

By Claim 4.16, we infer that  $d_{GD}(e') \leq d_{GD}(e_j)$ , as  $e'$  lies on the unique path between  $t_i$  and  $t_j$  in  $T_k$ . Hence,  $(e, e_j)$  is a  $(\tau/2)$ -efficient replacement pair in  $T_j$  and, moreover,  $e$  is incident with  $t_j$ . By the requirements we impose on the algorithm,  $(e, e_j)$  is not  $(\tau/(2c))$ -heavy in  $T_j$ , that is:

$$d_{GD}(e_j) \leq \frac{\tau/2}{c} \cdot \frac{d_{GD}(T_j)}{|V|} \leq \frac{\tau/2}{2\mu(1 + \tau)^2} \cdot \frac{2\mu(1 + \tau)d_{GD}(T_k)}{|V|} = \frac{\tau/2}{1 + \tau} \cdot \frac{d_{GD}(T_k)}{|V|}.$$

As  $d_{GD}(e') \leq d_{GD}(e_j)$ , we infer that  $(e, e')$  is not  $(\tau/2, 1 + \tau)$ -good in  $T_k$ , and the lemma is proven.  $\square$

We now apply the analysis of [20], i.e. Theorem 4.5, to obtain a bound on the number of replacements performed by the incremental scheme.

**Lemma 4.17.** *Assume that an algorithm implements the incremental scheme for some choice of parameters  $\tau \geq 2\varsigma > 0$  with  $\varsigma = O(1)$ . Then, in the course of the first  $r$  additions, the algorithm performs  $O(r\varsigma^{-1}(1 + \log \mu))$  replacements.*

*Proof.* Let  $t_1, t_2, t_3, \dots, t_r$  be a sequence of terminals added to the terminal set,  $S_k = \{t_1, t_2, \dots, t_k\}$  and  $T_k$  be the tree maintained by the algorithm after  $k$ -th terminal addition, i.e.,  $V(T_k) = S_k$ . The set of replacement pairs applied by the algorithm can be arranged into a set of  $r - 1$  sequences of the form  $(e_k^2, e_k^1), (e_k^3, e_k^2), \dots, (e_k^{s(k)}, e_k^{s(k)-1})$ ,  $2 \leq k \leq r$ , where  $e_k^{s(k)} \in E(T_r)$  and  $e_k^1$  is the edge added to connect  $t_k$  to  $T_{k-1}$  at the beginning of the  $k$ -th addition step. As each replacement is  $\varsigma$ -efficient, we have

$$s(k) - 1 \leq \log_{1+\varsigma} \left( d_{GD}(e_k^1) / d_{GD}(e_k^{s(k)}) \right). \quad (3)$$

Consider now  $MST(GD[S_r])$  and let  $e_2, e_3, \dots, e_r$  be such a permutation of  $E(MST(GD[S_r]))$  such that  $e_k$  is a friend of  $e_k^{s(k)}$  w.r.t.  $MST(GD[S_r])$ , for every  $2 \leq k \leq r$ . By the properties of a minimum spanning tree,

$$d_{GD}(e_k^{s(k)}) \geq d_{GD}(e_k). \quad (4)$$

Altogether, we can bound the number of replacements, equal to  $\sum_{k=2}^r s(k) - 1$ , as follows.

$$\begin{aligned} \sum_{k=2}^r s(k) - 1 &\leq \sum_{k=2}^r \log_{1+\varsigma} \frac{d_{GD}(e_k^1)}{d_{GD}(e_k^{s(k)})} && \text{by (3)} \\ &\leq \sum_{k=2}^r \log_{1+\varsigma} \frac{d_{GD}(e_k^1)}{d_{GD}(e_k)} && \text{by (4)} \\ &\leq O(\varsigma^{-1}) \left( r + \log \frac{\prod_{k=2}^r d_{GD}(e_k^1)}{\prod_{k=2}^r d_{GD}(e_k)} \right) && \text{as } \varsigma = O(1) \\ &\leq O(\varsigma^{-1}) (r + \log((4\mu)^{r-1})) && \text{by Theorem 4.5} \\ &\leq O(r\varsigma^{-1}(1 + \log \mu)). \end{aligned}$$

□

## 4.5 Fully dynamic online algorithm

We now merge the ideas of two previous sections to obtain a scheme for a fully dynamic algorithm. We fix three accuracy parameters  $\tau, \varsigma, \varepsilon > 0$ , where  $\tau$  controls the efficiency of replacement pairs that are allowed to remain in the tree,  $\varsigma \leq \tau$  controls the efficiency of the replacement pairs actually made, and hence the total number of replacements made by the algorithm, whereas  $\varepsilon$  controls the loss caused by postponing the deletion of high-degree non-terminal vertices from the tree, and hence the degree threshold at which we delete non-terminal vertices from the tree. We aim at  $(1 + \tau)(1 + \varepsilon)$ -approximation of  $MST(GD[S])$  for  $S$  being the current set of terminals. As in the decremental step, we define  $\eta = 1 + \lceil \varepsilon^{-1} \rceil = O(\varepsilon^{-1})$ , that is,  $\eta$  is the minimum positive integer with  $\frac{\eta}{\eta-1} \leq 1 + \varepsilon$ . In the algorithm, we maintain a tree  $T$  spanning a set of terminals and nonterminals. The degree of every nonterminal is greater than  $\eta$  and the tree does not admit any  $\tau$ -efficient replacements.

In a deletion step, we behave in exactly the same manner as in the decremental scheme in Section 4.3. If we want to delete  $v$  from the terminal set, we mark it as non-terminal and call the procedure  $\text{REMOVE}_\eta(v)$  that tries to remove vertex  $v$  from the tree  $T$  if  $v$  is not a terminal and its degree is at most  $\eta$ .

In an addition step, we perform similarly as in the incremental scheme, but there are two significant differences. First, we do not have the guarantee that the cost of the tree will not

decrease much in the future (as in Lemma 4.14), so we cannot stop replacing edges at some cost threshold. Second, we need to watch out for non-terminal vertices whose degree may drop to the threshold  $\eta$  as a consequence of a replacement. Formally, to add a vertex  $v$  to the terminal set  $S$  we perform the following operations on the currently maintained tree  $T$ .

1. If  $v \in V(T)$ , mark  $v$  as a terminal and finish.
2. Otherwise, connect  $v$  to any vertex of  $T$ , initialize an auxiliary set  $R$  as  $R = \emptyset$ ; and
3. apply a sequence of  $\varsigma$ -efficient replacement pairs, where for each such pair  $(e, e')$  we require that  $e'$  has the maximum possible cost among the friends of  $e$ ; moreover, for each such pair we insert both endpoints of  $e'$  into the set  $R$ ;
4. once all replacement pairs are applied, we require that there exists no  $\tau$ -efficient replacement pair  $(e, e')$  with  $e$  incident to  $v$ ;
5. at the end, invoke  $\text{REMOVE}_\eta$  on all vertices in  $R$  in an arbitrary order.

We claim the following.

**Lemma 4.18.** *If a fully dynamic algorithm satisfies the scheme above, after each step the tree  $T$  does not admit a  $\tau$ -efficient replacement pair and each non-terminal vertex of  $T$  has degree larger than  $\eta$  in  $T$ . Consequently,  $T$  is a  $(1 + \tau)(1 + \varepsilon)$ -approximation of  $\text{MST}(GD[S])$  and a  $2(1 + \tau)(1 + \varepsilon)\mu$ -approximation of minimum Steiner tree spanning  $S$  in  $G$ , where  $S$  is the current set of terminals.*

*Proof.* The second claim of the lemma follows directly from the first claim and lemmas 4.10 and 4.3. The claim that each non-terminal vertex of  $T$  has degree larger than  $\eta$  in  $T$  is obvious, as we invoke the procedure  $\text{REMOVE}_\eta$  on a vertex whenever its degree drops in  $T$ .

Consider now a step of adding a vertex  $v$  to the terminal set. Let  $T$  be the tree before this step, and  $T'$  be the tree obtained after connecting  $v$  to the  $T$  and performing all replacements, but before the first call to the procedure  $\text{REMOVE}_\eta$  on an element of  $R$ . We claim that if  $T$  does not admit a  $\tau$ -efficient replacement pair, neither does  $T'$ .

Assume the contrary. By the definition of  $T'$ , if  $T'$  admits a  $\tau$ -efficient replacement pair  $(e, e')$ , then  $e$  is not incident with  $v$  (such a pair is henceforth called *forbidden replacement pair*). Let  $T^0$  be the tree  $T$  with  $v$  connected to  $T$ , and  $T^1, T^2, \dots, T^s = T'$  be the sequence of trees computed by the algorithm;  $T^j$  is constructed from  $T^{j-1}$  by applying a  $\varsigma$ -efficient replacement pair  $(e_j, e'_j)$ . Note that  $\text{deg}_{T^0}(v) = 1$  and it follows from Lemma 4.12 for the vertex  $v$  that  $T^0$  does not admit a forbidden replacement pair. By our assumption,  $T^s = T'$  contains a forbidden replacement pair; let  $i$  be the smallest integer for which  $T^i$  contains a forbidden replacement pair, and let  $(e, e')$  be any one of them. Note that  $i > 0$ .

We apply Lemma 4.11 to the tree  $T^{i-1}$ , the family  $\mathcal{T}$  being the two trees of  $T^{i-1} \setminus \{e'_i\}$ , the tree  $T^c$  consisting of a single edge  $e'_i$ , and set  $F = \{e_i\}$ . By the choice of  $i$ ,  $(e, e')$  is not a replacement pair in  $T^{i-1}$ , as no forbidden replacement pair exists in  $T^{i-1}$ .

Consider now the case  $e' \in F$ , that is,  $e' = e_i$ . As  $(e, e')$  is  $\tau$ -efficient and  $(e_i, e'_i)$  is  $\varsigma$ -efficient, we have that  $(1 + \tau)d_{GD}(e) < d_{GD}(e'_i)$  and, in particular,  $e \neq e'_i$ . As  $e$  connects the two connected components of  $T^{i-1} \setminus \{e'_i\}$ ,  $(e, e'_i)$  is a  $\tau$ -efficient replacement pair in  $T^{i-1}$ . Since  $e$  is not incident to  $v$ ,  $(e, e'_i)$  is a forbidden replacement pair in  $T^{i-1}$ , a contradiction to the choice of  $i$ .



In the remaining case,  $(e_i, e')$  is a replacement pair in  $T^{i-1}$ . As we require  $e'_i$  to be an edge of maximum cost among the friends of  $e_i$  in  $T^{i-1}$ , we infer that  $d_{GD}(e'_i) \geq d_{GD}(e') > d_{GD}(e)/(1 + \tau)$ , where the last inequality follows from the  $\tau$ -efficiency of  $(e, e')$ . In particular,  $e \neq e'_i$ , and  $(e, e'_i)$  is a  $\tau$ -efficient replacement pair in  $T^{i-1}$ . Since  $e$  is not incident to  $v$ ,  $(e, e'_i)$  is a forbidden replacement pair, a contradiction to the choice of  $i$ .

To finish the proof of the lemma, observe that Lemma 4.12 ensures that a call to the procedure `REMOVE $_{\eta}$`  cannot introduce a  $\tau$ -efficient replacement pair, if it was not present prior to the call.  $\square$

We conclude this section with a note on the efficiency of our fully dynamic scheme.

**Lemma 4.19.** *Assume that an algorithm implements the fully dynamic scheme for some choice of parameters  $\tau, \varsigma, \varepsilon > 0$  with  $\varsigma \leq \tau$  and  $\varsigma = O(1)$ , and is run on an  $\mu$ -near metric space  $GD = (V, \binom{V}{2}, d_{GD})$  with stretch  $D$ , and a sequence of  $r$  operations, where  $r_+$  of this operations are additions and  $r_- = r - r_+$  are deletions. Then the algorithm performs  $O(\varsigma^{-1}(r_+ \log D + r_-(1 + \log \mu)))$  replacements, and at most  $r_-$  calls to `REMOVE $_{\eta}$`  procedure that result in a modification of the maintained tree. Moreover, each vertex removed by the procedure `REMOVE $_{\eta}$`  is of degree at most  $\eta = O(\varepsilon^{-1})$  in the currently maintained tree.*

*Proof.* The second and third bounds follow directly from the fact that `REMOVE $_{\eta}$`  modifies the tree only if it is called upon a non-terminal vertex of degree at most  $\eta$  in the tree, and each deletion operation introduces only one non-terminal vertex. For the first claim, let  $d_{min}$  be the minimum cost of an edge in  $GD$ , and consider a potential  $\Phi(T) = \sum_{e \in E(T)} \log_{1+\varsigma}(d_{GD}(e)/d_{min})$ . Clearly, this potential is zero at the beginning and never becomes negative. Moreover, each  $\varsigma$ -efficient replacement drops the potential by at least one.

If we introduce a new edge, we increase the potential by at most  $\lceil \log_{1+\varsigma} D \rceil = O(\varsigma^{-1} \log D)$ . Consider now a call to the procedure `REMOVE $_{\eta}$`  that replaces a vertex  $v$  of degree  $s \leq \eta$  in  $T$  with a set  $F^v$  of  $s-1$  edges, obtaining the reconnected tree  $T'$ . Let  $\mathcal{T}^v = \{T_1, T_2, \dots, T_s\}$  be the connected components of  $T \setminus \{v\}$  and for every  $1 \leq i \leq s$ , let  $e_i$  be the unique edge connecting  $V(T_i)$  with  $v$  in the tree  $T$ , and let  $u_i$  be the endpoint of  $e_i$  that belongs to  $V(T_i)$ . Without loss of generality assume  $d_{GD}(e_1) \geq d_{GD}(e_2) \geq \dots \geq d_{GD}(e_s)$ . Observe that, by Lemma 4.2, for every  $1 \leq i < s$  we have  $d_{GD}(u_i u_{i+1}) \leq \mu(d_{GD}(e_i) + d_{GD}(e_{i+1})) \leq 2\mu d_{GD}(e_i)$ . Moreover, the set  $F' := \{u_i u_{i+1} : 1 \leq i < s\}$  is a valid candidate for the set  $F^v$  and hence there exists a bijective mapping  $\pi : F^v \rightarrow F'$  such that  $d_{GD}(\pi(e)) \geq d_{GD}(e)$  for any  $e \in F^v$ . We infer that

$$\Phi(T') - \Phi(T) \leq \sum_{i=1}^{s-1} \log_{1+\varsigma}(d_{GD}(u_i u_{i+1})/d_{GD}(e_i)) \leq \log_{1+\varsigma}(2\mu) = O(\varsigma^{-1}(1 + \log \mu)).$$

The lemma follows.  $\square$

As discussed at the beginning of this section, it does not seem clear if the analysis of incremental algorithm of [20] (i.e., Theorem 4.5) is applicable to our fully dynamic scheme. The usage of this analysis in [21] strongly relies on the assumption that newly added terminals are new vertices in the graph, with distances to terminals already deleted being implicitly defined by the triangle inequality. However, in our model of a fixed host graph  $G$  this assumption is no longer valid, hence we need to fall back to the more straightforward analysis of Lemma 4.19.

### 4.5.1 Decreasing the dependence on the stretch of the metric

The fully dynamic scheme presented above yields algorithms depending on the stretch of the metric. The dependence in Lemma 4.19 is proportional to  $\log D$ . However, in some implementations the dependence on  $\log D$  may be bigger. In this section we present a generic method of avoiding such a dependency, at the cost of additional polylogarithmic factors in  $n = |V|$ .

The main idea is similar as in the incremental algorithm: the edges of cost  $\varepsilon d_{GD}(T)/n$  are almost for free, and, consequently, the window of interesting edge costs is of stretch  $O(n/\varepsilon)$ . However, in the fully dynamic algorithm we cannot simply ignore what happens to the tree on low-cost edges, as they may become important in the future, when the cost of the tree decreases significantly. Thus, we adapt a different strategy: for each possible window, we maintain a solution that considers only edges of levels from this window (and assumes the cheaper ones are for free). Formally, we prove the following.

**Theorem 4.20.** *Let  $\mathcal{G}$  be a minor-closed graph class and  $\alpha, \varepsilon > 0$  be constants. Assume we are given a fully dynamic algorithm that, given a graph  $G \in \mathcal{G}$  with nonnegative edge weights,  $n$  vertices and metric stretch  $D$ , initializes in  $f_{ini}(n, \log D)$  time, handles updates (terminal additions and deletions) in  $f_{upd}(n, \log D)$  time, handles queries on the cost of minimum Steiner tree in  $f_q(n, \log D)$  time with approximation guarantee  $\alpha$ , and uses  $f_{mem}(n, \log D)$  memory. Assume additionally that all  $f_{ini}, f_{upd}, f_q$  are non-decreasing with respect to their parameters and that  $f_{ini}$  and  $f_{mem}$  are convex and at least linear in  $n$ .*

*Then there for some  $\Delta = O(\varepsilon^{-1} \log(\varepsilon^{-1}n))$  there exists a fully dynamic algorithm that initializes in  $O(f_{ini}(n, \Delta)\varepsilon^{-1} \log D)$  time, handles updates in  $O(f_{upd}(n, \Delta)\varepsilon^{-1} \log D)$  time, handles queries in  $f_q(n, \Delta) + O(1)$  time with approximation guarantee  $(\alpha + \varepsilon)$  and uses  $O(f_{mem}(n, \Delta)\varepsilon^{-1} \log D)$  memory.*

Before we dive into the depths of the proof of Theorem 4.20, we would like to clarify that this theorem operates on the actual graph  $G$  and the metric space  $\overline{G}$  being the closure of the edge weights of  $G$ . In other words, all distance oracles and subsequent  $\mu$ -near metric spaces approximating  $\overline{G}$  are hidden inside the assumed fully dynamic algorithm.

*Proof.* Let  $\overline{G} = (V, \binom{V}{2}, d_{\overline{G}})$  be the metric closure of  $G$ , recall that  $n = |V|$ . Let us introduce levels of edge weights in  $\overline{G}$  as follows. Let  $\beta = \beta(\varepsilon) = 1 + \varepsilon$  and for any  $u, v \in V$ , we define the level of  $uv$ ,  $\text{lvl}(uv)$ , as  $\lfloor \log_{\beta} d(uv) \rfloor$ , i.e.,  $\beta^{\text{lvl}(uv)} \leq d_G(uv) < \beta^{\text{lvl}(uv)+1}$ . Let  $\mathfrak{a}$  be the smallest integer such that if for any two edges  $e, e'$  we have  $\text{lvl}(e') - \text{lvl}(e) \geq \mathfrak{a}$  then  $d_G(e) < d_G(e')/n$ . Similarly we define  $\mathfrak{b}$  to be the minimum difference in levels that guarantee a ratio of costs of at least  $\varepsilon/n$ . Note that  $\mathfrak{a} = O(\varepsilon^{-1} \log n)$  and  $\mathfrak{b} = O(\varepsilon^{-1} \log(\varepsilon^{-1}n))$ . We fix  $\Delta = \mathfrak{a} + \mathfrak{b}$ .

For an integer  $i$ , we define a binary relation  $\mathcal{R}^i$  on  $V$  as follows:  $(u, v) \in \mathcal{R}^i$  iff there exists a path between  $u$  and  $v$  in  $\overline{G}$  whose each edge is of level at most  $i$ . Note that, equivalently, we may require that this path resides in  $G$ , as the weights of edges in  $G$  are nonnegative. It is easy to see that  $\mathcal{R}^i$  is an equivalence relation; let  $\mathcal{V}^i$  be the set of its equivalence classes and, for any  $v \in V$ ,  $\mathbf{v}^i(v)$  is the equivalence class of  $v$  in  $\mathcal{R}^i$ .

For a fixed level  $i$ , let  $G^i$  be a graph constructed from  $G$  as follows: we contract all edges of  $G$  of level at most  $i - \mathfrak{b}$ , and we delete all edges  $uv$ , where  $(u, v) \notin \mathcal{R}^{i+\mathfrak{a}}$ . Note that  $V(G^i) = \mathcal{V}^{i-\mathfrak{b}}$  and each connected component of  $G^i$  corresponds to one equivalence class of  $\mathcal{R}^{i+\mathfrak{a}}$ . Moreover,  $G^i$  is a minor of  $G$  and can be constructed in linear time.

Let  $d^i$  be the metric on  $V(G^i) = \mathcal{V}^{i-b}$  induced by the weights of the edges of  $G^i$ . Let  $\mathbf{v}$  and  $\mathbf{v}'$  be two distinct vertices of the same connected component of  $G^i$ . As  $\mathbf{v}$  and  $\mathbf{v}'$  can be connected with a path with edges of level at most  $i + \mathbf{a}$ , we have  $d^i(\mathbf{v}\mathbf{v}') < (n-1)\beta^{i+1+\mathbf{a}}$  and therefore  $\text{lvl}(\mathbf{v}\mathbf{v}') \leq i + O(\varepsilon^{-1} \log n)$  in the metric  $d^i$ . As we contracted all edges of level at most  $i - \mathbf{b}$  in the process of obtaining the graph  $G^i$ , we have  $\text{lvl}(\mathbf{v}\mathbf{v}') > i - \mathbf{b}$ . Since  $\mathbf{b} = O(\varepsilon^{-1} \log(\varepsilon^{-1}n))$ , we infer that the logarithm of the stretch of the metric  $d^i$  is at most  $\Delta = O(\varepsilon^{-1} \log(\varepsilon^{-1}n))$ .

Intuitively, the graph  $G^i$  contains all the information we need to construct the approximate Steiner tree for a given set of terminals, if the most expensive edge of such a tree would have level roughly  $i$ .

We run several copies of the assumed fully dynamic algorithm: for each level  $i$  out of  $\log_{1+\varepsilon} D = O(\varepsilon^{-1} \log D)$  levels and for each  $\mathbf{c} \in \mathcal{V}^{i+\mathbf{a}}$  we create a copy of the fully dynamic algorithm  $\mathcal{A}_{\mathbf{c}}^i$  created on the connected component  $\mathbf{c}$  of the graph  $G^i$ . At any time, if  $S$  is the current set of terminals in  $G$ , a vertex  $\mathbf{v} \in V(G^i)$  is a terminal in  $\mathcal{A}_{\mathbf{c}}^i$  where  $\mathbf{v} \subseteq \mathbf{c}$  if  $\mathbf{v} \cap S \neq \emptyset$ . Denote  $S^i = \{\mathbf{v} \in V(G^i) : S \cap \mathbf{v} \neq \emptyset\}$ . Note that the memory usage of the structures is as promised due to convexity of  $f_{mem}$ .

At initialization, note that the construction of the graph  $G^i$  takes time linear in  $|V(G)| + |E(G)|$ , and, therefore is subsumed by the time taken to initialize the fully dynamic algorithms. The initialization time follows from the assumptions on the convexity of  $f_{ini}$ .

To handle updates (additions and deletions from the terminal set), for each level  $i$  and for each  $\mathbf{v} \in V(G^i) = \mathcal{V}^{i-b}$  we keep a counter how many terminals of  $S$  are contained in  $\mathbf{v}$ . An addition or deletion of a vertex  $v$  from the terminal set results in a change of the counter at each of the  $O(\varepsilon^{-1} \log D)$  levels; if at any level the counter changes from 0 to 1 or from 1 to 0, an operation (addition or deletion) is performed on  $\mathcal{A}_{\mathbf{c}}^i$ , where  $\mathbf{v} \subseteq \mathbf{c} \in \mathcal{V}^{i+\mathbf{a}}$ . The promised update time follows.

To handle queries, first, for each level  $i$ , we define  $F_i$  to be an (arbitrarily chosen) spanning forest of the subgraph of  $G$  that contains exactly those edges that have level at most  $i$ . Note that  $F_i$  can be constructed in linear time for each  $i$  at the time of initialization. Let us define  $i(S)$  to be the minimum level where all terminals are contained in the same equivalence class  $\mathbf{c}(S) \in \mathcal{V}^{i(S)}$ . Assume  $\mathbf{c}(S) \subseteq \mathbf{c}'(S) \in \mathcal{V}^{i(S)+\mathbf{a}}$ , that is, let  $\mathbf{c}'(S)$  be the equivalence class of  $\mathcal{R}^{i(S)+\mathbf{a}}$  that contains  $S$ . Let  $T$  be the tree maintained by  $\mathcal{A}_{\mathbf{c}'(S)}^{i(S)}$  and let  $H$  be the edge-induced subgraph of  $G$  that contains all edges of  $T$  as well as  $F_{i-b}$ . Clearly,  $H$  contains a connected component that spans  $S$  in  $G$ . We claim that the total cost of all edges of  $H$  is a good approximation of the cost of the minimum spanning tree of  $S$  in  $G$  and, consequently, a good approximation of the cost of the minimum Steiner tree of  $S$  in  $G$ . Let  $T_{OPT}$  be any such minimum Steiner tree of  $S$  in  $G$ .

By the definition of  $i(S)$ , any tree spanning  $S$  in  $G$  needs to contain at least one edge of level at least  $i$ . Hence,  $d_G(T_{OPT}) \geq \beta^i$ . On the other hand, by the definition of  $\mathbf{b}$ ,  $d_G(F_{i-b}) \leq (n-1) \cdot \varepsilon \beta^i / n < \varepsilon \beta^i \leq \varepsilon d_G(T_{OPT})$ .

As  $S \subseteq \mathbf{c}(S) \in \mathcal{V}^{i(S)}$ ,  $F_i$  contains a connected component that spans  $S$  and, consequently,  $d_G(T_{OPT}) \leq d_G(F_i) < (n-1)\beta^{i+1}$ . By the definition of  $\mathbf{a}$ ,  $d_G(T_{OPT}) < \beta^{i+\mathbf{a}}$ , so  $T_{OPT}$  does not contain any edge of level  $i + \mathbf{a}$  or higher. We infer that  $V(T_{OPT}) \subseteq \mathbf{c}'(S)$  and, by construction of  $G^i$ , the cost of a minimum Steiner tree of  $S^i$  in  $G^i$  is at most  $d_G(T_{OPT})$ . Therefore, the tree maintained by  $\mathcal{A}_{\mathbf{c}'(S)}^{i(S)}$  has cost at most  $\alpha d_G(T_{OPT})$ . We infer that  $d_G(H) \leq (\alpha + \varepsilon) d_G(T_{OPT})$ .

We note that the value of  $i(S)$  can be updated after each addition or deletion operation by counting the number of terminals in each connected component of each graph  $G^i$ . Thus, to handle a query we need to query one algorithm  $\mathcal{A}_{\mathbf{c}'(S)}^{i(S)}$  and to look up the precomputed cost of the forest  $F_{i(S)-\mathbf{b}}$ . The query time follows.  $\square$

## 5 Dynamic vertex-color distance oracles

In this section we construct dynamic vertex-color distance oracles that we later use to implement algorithms from Section 4. Their interface, as well as some basic definitions have been introduced in Section 3.2. This section covers the construction that allows us to extend any generic oracle to a vertex-color distance oracle. Moreover, we show how to build generic oracles for general and planar graphs. In the end, we obtain an incremental vertex-color distance oracle for planar graphs, and fully dynamic oracles, both for planar and general graphs.

### 5.1 Generic construction

In the Section 3.2 we defined a list of queries and update operations that our oracles should support. In this section we present generic implementations of the two schemes introduced in the previous section. The constructions presented here are generic in a sense that no specific choice of portals or pieces is imposed. Instead, the running times and space requirements are measured in terms of a few general parameters:

- $p_{\text{tot}} = \sum_{v \in V} |\text{portals}(v)|$  is the (expected) total number of vertex-portal connections,
- $p_{\text{num}} = |\mathcal{P}|$  is the (expected) total number of portals,
- $p_{\text{max}} = \max_{v \in V} |\text{portals}(v)|$  is an upper bound on the (expected) number of portals assigned to a vertex,
- $r_{\text{tot}} = \sum_{r \in \mathcal{R}} |r|$  is the (expected) total size of all pieces,
- $r_{\text{num}} = |\mathcal{R}|$  is the (expected) total number of different pieces,
- $r_{\text{max}}^{\#} = \max_{r \in \mathcal{R}} |r|$  is an upper bound on the (expected) size of a piece,
- $r_{\text{max}} = \max_{v \in V} |\text{pieces}(v)|$  is an upper bound on the (expected) number of pieces assigned to a vertex,
- $a_{\text{tot}} = \sum_{v \in A} |\text{portals}(v)|$ , where  $A$  is a set of vertices activated by the oracle, is the (expected) total number of vertex-portal connections over vertices that were activated by the oracle.

Note that some of our constructions are randomized, and then we measure the quality of the oracle in expectation. We also assume that all these values are bounded polynomially in  $n$ , and hence the logarithm of any of it can be bounded by  $O(\log n)$ .

The problem one deals with in our oracles is finding the vertex-set portal distances. The piece distances are easy to find: upon a query it suffices to run single source shortest paths [24] algorithm from  $v$  in every piece of  $\text{pieces}(v)$  and compare the distances. This section is mostly devoted to designing a data structure able to find portal distances. To that end, for a portal  $p$  and color  $i$  we define  $N_i^p$  to be the distance between  $p$  and the nearest vertex in  $\text{cluster}(p)$  that has color  $i$ . Every portal  $p$  holds a heap  $H_p$ , where it stores the values  $N_i^p$  that correspond to active colors in  $\text{cluster}(p)$ . The main goal of this section is to describe how to update the values  $N_i^p$  and heaps  $H_p$  as colors of vertices are modified.

**Lemma 5.1.** *Assume we are given an  $\alpha$ -approximate generic distance oracle for  $G$ . One can maintain expected constant-time access to the values  $N_i^p$  and deterministic constant-time access to heaps  $H_p$  in  $O(a_{\text{tot}} \log n)$  total expected time under an arbitrary sequence of `activate`( $i$ ) and `merge`( $i, j$ ) operations. The data structure requires  $O(p_{\text{tot}})$  space and can be initialized in  $O(p_{\text{tot}} \log n)$  time.*

*The data structure can be derandomized, at the cost of increase to total  $O(a_{\text{tot}} \log n \log \log n)$  running time, and  $O(\log \log n)$  time to access a single value of  $N_i^p$ .*

*Proof.* Every update operation affects vertices of some color and all their portals. Thus, for every color  $i$  we define a set of pairs  $L(i) = \{(v, p) | v \in C_i, p \in \text{portals}(v)\}$  and maintain it as a list. Note that the total size of all those lists is exactly  $p_{\text{tot}}$  and at the beginning all the colors are inactive. Recall that colors have to be activated first in order to be merged. The total size of the lists of colors that will be activated during the runtime of the oracle is  $a_{\text{tot}}$ . On the remaining lists no operations will ever be executed.

As mentioned before, each portal  $p \in P$  holds a heap  $H_p$  storing distances  $N_i^p$  to all active colors  $i$  in  $\text{cluster}(p)$ . Together with  $N_i^p$  we store on  $H_p$  the vertex that provides  $N_i^p$ , i.e., the vertex  $v$  in  $\text{cluster}(p)$  of color  $i$  whose distance from  $p$  is  $N_i^p$ . We define  $\bar{N}_i^p = (N_i^p, v)$  to be such a pair. We store pairs  $\bar{N}_i^p$  in a heap, sorted according to the first coordinate only.

At the beginning, all colors are inactive. Every portal  $p$  initializes  $H_p$  to be an empty heap. In addition to  $H_p$ , every portal maintains two dictionaries, where one can add/remove (*key, value*) pairs and look up keys in time  $O(t_{\text{dict}})$  per operation; we address the choice of the dictionary, and henceforth determine the value of  $t_{\text{dict}}$ , at the end of the proof. A dictionary of inactive distances,  $I_p$ , stores pairs  $(i, \bar{N}_i^p)$  for every inactive color  $i$  which has a nonempty intersection with  $\text{cluster}(p)$ . Initially, every portal iterates over its cluster and computes pairs  $(i, \bar{N}_i^p)$  based on stored distances  $D_{p,v}$ . Every portal also holds a dictionary of active colors  $A_p$ , which stores pairs  $(i, q)$  for every active color  $i$  in  $\text{cluster}(p)$ , where  $q$  is a pointer to a position on  $H_p$  where  $\bar{N}_i^p$  is stored. Initially  $A_p$  is empty.

To perform `activate`( $i$ ) it suffices to update all relevant heaps  $H_p$  with information about color  $i$ . We iterate over elements of  $L(i)$  and for each pair  $(v, p)$  we remove  $(i, \bar{N}_i^p)$  from  $I_p$  (if it is there) and add  $\bar{N}_i^p$  to  $H_p$  and  $(i, q)$  to  $A_p$ , where  $q$  is a pointer to  $\bar{N}_i^p$  on  $H_p$ . This is done at most once for each of the  $O(a_{\text{tot}})$  pairs, so the total time of this operation is  $O(a_{\text{tot}}(\log n + t_{\text{dict}}))$ .

Now, consider a `merge`( $i, j$ ) operation, and keep in mind that both  $i$  and  $j$  are active. Without loss of generality, we may assume that  $|L(i)| \leq |L(j)|$ . We change the color of all vertices in  $C_i$  to  $j$ . To be more precise, we iterate over  $L(i)$  and for each pair  $(v, p)$  we:

remove $i, j$ from $A_p$ , set $\bar{N}_j^p = \min(\bar{N}_i^p, \bar{N}_j^p)$ and insert $\bar{N}_j^p$ to $H_p$ and $A_p$	if $i, j \in A_p$
we remove $(i, q)$ from $A_p$ and add $(j, q)$ to $A_p$ ,	if $i \in A_p, j \notin A_p$
do nothing	if $i \notin A_p$

Finally, we append list  $L(i)$  to  $L(j)$  in  $O(|L(i)|)$  time.

Note that for each appended pair, the length of the list it belongs to doubles. Since the active list length is bounded by  $a_{\text{tot}}$ , each of  $a_{\text{tot}}$  pairs will change its list at most  $O(\log a_{\text{tot}}) = O(\log n)$  times. Each such list change of a pair may trigger a dictionary update and a heap update. The total time for list merging, iterating and dictionary modifications is therefore  $O(a_{\text{tot}} t_{\text{dict}} \log n)$ . We will bound the number of operations on heaps separately. Note that only when  $i, j \in A_p$  the merge is effective in  $\text{cluster}(p)$  and  $O(\log n)$  operation on a heap is triggered. Otherwise, if  $j \notin A_p$ , only  $A_p$  is modified and not  $H_p$ . For every portal  $p$  we effectively merge colors in its cluster at most

$|cluster(p) \cap A| - 1$  times, where  $A$  is the set of activated vertices. So the number of times we modify a heap is  $\sum_{v \in A} |portals(v)| = O(a_{tot})$ . This gives the total running time of  $O(a_{tot} t_{dict} \log n)$ .

Observe that each heap  $H_p$  can be accessed in constant time, and each value  $N_i^p$  in  $O(t_{dict})$  time.

Finally, we address the choice of the dictionaries. If use hash maps to implement dictionaries, then we obtain  $t_{dict} = O(1)$  in expectation. As an alternative, observe that every used dictionary uses color numbers as keys. Hence, one might use Y-fast-tries [38] yielding a slightly worse but deterministic  $t_{dict} = O(\log \log n)$  worst case time.  $\square$

**Lemma 5.2.** *Assume we are given an  $\alpha$ -approximate generic distance oracle for a weighted graph  $G = (V, E, d)$ . Let  $n = |V|$ . One can maintain constant-time access to the values  $N_i^p$  and heaps  $H_p$ , under a sequence of `activate( $i$ )`, `deactivate( $i$ )`, `merge( $i, j$ )` or `split( $l, i, j$ )` operations. Each operation requires  $O(p_{num} \log n)$  worst case time. The data structure uses  $O(p_{tot} \log n)$  space and can be initialized in  $O(p_{tot} \log n)$  time with arbitrary forest  $T_1, \dots, T_k$  spanning  $V(G)$  corresponding to colors  $C_1, \dots, C_k$ .*

*Proof.* In the proof we use Euler tour trees introduced in [35], called further the ET-trees. An ET-tree represents an Euler tour of an arbitrary unrooted tree  $T$ , i.e., a tour that traverses each edge of the tree twice, once in each direction. More formally, each edge of  $T$  is replaced with two oppositely directed edges and a directed self loop is added in every vertex. We represent the Euler tour of the resulting graph as a balanced binary tree that allows join and split operations. The tree has a single node for every edge of the tour, and the nodes are arranged in the same order as the tour.

The trees represented in this way can be linked by adding an edge or split by removing some edge (i.e., they support link and cut operations), and the corresponding ET-trees may be updated by joining and splitting the representing binary trees. Each of these operations can be implemented in  $O(\log n)$  time. ET-trees can also support typical extensions of binary search trees, including associating keys with vertices and finding the minimum over all the keys in the tree. Note that the keys of vertices of  $T$  are associated with the self loops in the Euler tour representation. Thus, for every node of  $T$ , the ET-tree has a single node holding its key. On the other hand, some nodes of the ET-tree may not correspond to any node of  $T$ .

In our data structure, for each color  $C_i$  we maintain the corresponding tree  $T_i$  as an ET-tree  $ET_i$ . We associate a key of size  $O(p_{num})$  with every node of  $ET_i$ . Hence, each operation on  $ET_i$  requires  $O(p_{num} \log n)$  time.

Let us now describe the keys in detail. Consider a color  $i$ , a node  $x$  of  $ET_i$  and a portal  $p$ . The subtree of  $ET_i$  rooted at  $x$  corresponds naturally to some set of vertices of color  $i$ . Denote this set by  $C_i(x)$ . In particular, the entire  $ET_i$  corresponds to  $C_i$ . We define  $N_i^p(x)$  to be the distance between  $p$  and the nearest vertex that belongs both to  $C_i(x)$  and  $cluster(p)$ .

A key associated with  $x$  is a list  $L(x)$  of pairs  $(p, N_i^p(x))$ . The list is sorted according to some (arbitrary) linear ordering of portals. This allows us to compute the key of any node in  $ET_i$  in  $O(p_{num})$  time, using the keys stored in the child nodes.

In order to save space, if a subtree of  $ET_i$  rooted at a node  $x$  does not contain any nodes from  $cluster(p)$ , we do not store the list element for portal  $p$  in the list  $L(x)$ . Let us bound the total length of all lists  $L(x)$  in all ET-trees. Consider a vertex  $v$ . It belongs to exactly one node  $x$  of some ET-tree and may cause that  $O(|portals(v)|)$  elements appear in every list on the path from  $x$  to the root of the corresponding ET-tree. Since the depth of each tree representing an

ET-tree is  $O(\log n)$ , the total length of all lists  $L(x)$ , as well as the space usage of all ET-trees is  $O(\sum_{v \in V} |\text{portals}(v)| \log n) = O(p_{\text{tot}} \log n)$ . They can also be initialized, given any forest  $T_1 \dots T_k$  spanning  $V(G)$  and representing colors  $C_1 \dots C_k$ , in  $O(p_{\text{tot}} \log n)$  time.

Clearly, each value  $N_i^p$  can be extracted from the root of the tree  $ET_i$ . In order to enable constant lookup of values  $N_i^p$ , useful for efficient distance queries, for each color  $i$  we store a deterministic dictionary of Hagerup et al. [23]. A dictionary with  $k$  entries can be initialized in  $O(k \log k)$ , and allows constant time access. It uses  $O(k)$  space. Since every operation affects at most two colors, we can afford to rebuild the affected dictionaries after every operation in  $O(p_{\text{num}} \log n)$  time.

A  $\text{merge}(i, j)$  or  $\text{split}(l, i, j)$  operation can be handled with a single operation on one of the ET-trees. Each such operation may affect at most two elements of each heap  $H_p$ . Thus, updating both the ET-trees and the heaps  $H_p$  requires  $O(p_{\text{num}} \log n)$  time.

The  $\text{activate}(i)$  and  $\text{deactivate}(i)$  operations do not affect the values  $N_i^p$ , but they may affect the heaps  $H_p$ . Again, it takes  $O(p_{\text{num}} \log n)$  time to update them.

We remark here that, contrary to the previous lemma, we have  $O(p_{\text{num}} \log n)$  deterministic worst case time for each operation, as we can afford to use always Y-fast-tries to maintain pointers to all colors present in a heap  $H_p$ .  $\square$

In our data structures we also need to find vertices of a given color in a particular piece. Using the data structures from the above lemma, for a given vertex, we may determine its color in  $O(\log n)$  time. This can be achieved by simply maintaining a pointer from each vertex to the corresponding node in an ET-tree. Once we find this node, we need to find the root of its ET-tree which requires  $O(\log n)$  time. However, if the total number of distinct pieces is small, we may use a slightly faster approach.

**Lemma 5.3.** *Assume we are given an  $\alpha$ -approximate generic distance oracle for a graph  $G = (V, E, d)$ . Let  $n = |V|$ . There exists a data structure that allows to list all vertices of a given color in a given piece, and list all active vertices in a given piece. Both lists are computed in time that is linear in their size.*

*The data structure can be initialized in  $O(r_{\text{tot}} \log n)$  time and updated after every  $\text{activate}(i)$ ,  $\text{deactivate}(i)$ ,  $\text{merge}(i, j)$  or  $\text{split}(l, i, j)$  operation in  $O(r_{\text{num}} \log n)$  worst case time. It requires  $O(r_{\text{tot}} \log n)$  space.*

*Proof.* For each color  $i$  we maintain yet another ET-tree  $\widetilde{ET}_i$  corresponding to tree  $T_i$  associated with  $C_i$ . Consider a piece  $\tau \in \mathcal{R}$  and a node  $x$  of  $\widetilde{ET}_i$ . We define  $x$  to be  $\tau$ -heavy if either  $x$  corresponds to a vertex  $v \in \tau$ ,  $x$  is a root of  $\widetilde{ET}_i$ , or  $x$  has  $\tau$ -heavy nodes in both subtrees.

For a node  $y$  of  $\widetilde{ET}_i$  (not necessarily an  $\tau$ -heavy one) we extend the key associated with  $y$  with a list  $L(y)$  of triples  $(\tau, \text{left}_\tau(y), \text{right}_\tau(y))$ , where  $\tau$  is a piece. Let  $z$  and  $z'$  be the highest  $\tau$ -heavy vertices in left and right subtree of  $y$  respectively. Then,  $\text{left}_\tau(y)$  and  $\text{right}_\tau(y)$  are pointers to the position of  $\tau$  in  $L(z)$  and  $L(z')$  respectively. We do not store such a triple in  $L(y)$  if none of the subtrees contains an  $\tau$ -heavy node. The list  $L(y)$  is sorted according to some arbitrary linear order on pieces.

Note that we are able to calculate values of  $\text{left}_\tau(y)$  and  $\text{right}_\tau(y)$  using the keys stored in the sons of  $y$ . This allows us to compute the key of a node of  $\widetilde{ET}_i$  in  $O(r_{\text{num}})$  time, so each ET-tree operation requires additional  $O(r_{\text{num}} \log n)$  time.

Consider  $\tau$ -heavy vertices of  $\widetilde{ET}_i$ . Together with  $\text{left}_\tau$  and  $\text{right}_\tau$  pointers they form a binary tree, whose leaves correspond to vertices of piece  $\tau$ . Since the size of a binary tree is linear in the

number of its leaves, the pointers allow us to iterate over all vertices of color  $i$  in a piece  $\tau$  in linear time.

We store information about an arbitrary vertex  $v$  of a color  $i$  in a piece  $\tau$  only on a path between a corresponding vertex  $x$  of  $\widehat{ET}_i$  and a root of  $\widehat{ET}_i$  using  $O(\log n)$  space. Summing over all vertices of all pieces, we get a space usage bounded by  $O(r_{\text{tot}} \log n)$ .

To allow listing all active vertices in a piece for each piece  $\tau$  we maintain a set (represented as a balanced tree) containing all active colors of vertices in piece  $\tau$ . After each update operation we update all sets in  $O(r_{\text{num}} \log n)$  time. To list all active vertices in  $\tau$  we iterate over active colors in  $\tau$  and list all their vertices in  $\tau$ .

Lastly, we note that we need to access elements of lists  $L(y)$  in the roots of all ET-trees in constant time. In order to achieve that, after some ET-tree is updated, we build a deterministic dictionary [23] in  $O(r_{\text{num}} \log r_{\text{num}}) = O(r_{\text{num}} \log n)$  time.  $\square$

In the next two lemmas we show that our template for distance oracles is sufficient to implement the schemes of Section 4.

**Lemma 5.4.** *Assume we are given an  $\alpha$ -approximate generic distance oracle for a weighted graph  $G = (V, E, d)$ . Let  $n = |V|$ . One can implement the incremental scheme in  $O(p_{\text{max}} \log n)$  expected amortized time per update and  $O(r_{\text{max}} r_{\text{max}}^{\#} + p_{\text{max}})$  expected time per query operation. The data structure requires  $O(p_{\text{tot}})$  space and can be initialized in  $O(p_{\text{tot}} \log n)$  time.*

*The algorithm can be derandomized at the cost of increase to  $O(p_{\text{max}} \log n \log \log n)$  amortized time per update and  $O(r_{\text{max}} r_{\text{max}}^{\#} + p_{\text{max}} \log \log n)$  worst case time per query operation.*

*Proof.* According to Lemma 5.1, there is a data structure that can maintain heaps  $H_p$  with distances  $N_i^p$  for every portal  $p \in \mathcal{P}$  in total expected time  $O(a_{\text{tot}} \log n)$  (or  $O(a_{\text{tot}} \log n \log \log n)$  deterministic time), using  $O(p_{\text{tot}})$  space and  $O(p_{\text{tot}} \log n)$  time for initialization. The incremental oracle starts with all vertices assigned to distinct inactive colors. Hence, the number of update operations is at least  $|A|$ , where  $A$  is the set of vertices active at the end. Clearly,  $a_{\text{tot}} \log n = O(|A| p_{\text{max}} \log n)$  and hence the amortized time per update is  $O(p_{\text{max}} \log n)$ .

Now we move on to handling distance queries. Note that in the proof of Lemma 5.1 we iterated through all vertices that changed color upon every **merge** operation. So it is possible to maintain a flag in each vertex which allows checking the color of a vertex in constant time.

Consider **distance**( $v, i$ ) query. In order to find a piece distance from  $v$  to the closest vertex of color  $i$ , we run the linear time single source shortest paths algorithm for planar graphs [24] in every piece in  $\text{pieces}(v)$  with a source at  $v$ . We then check the colors of all the vertices visible from  $v$  and choose a minimum distance among the vertices of color  $i$ . In order to find a portal distance we iterate through  $\text{portals}(v)$  and compute  $\min_{p \in \text{portals}(v)} (D_{p,v} + N_i^p)$ . Recall that the dynamic dictionary we use allows portals to access values  $N_i^p$  in expected constant time (or  $O(\log \log n)$  deterministic). Therefore answering **distance**( $v, i$ ) query takes expected  $O(r_{\text{max}} r_{\text{max}}^{\#} + p_{\text{max}})$  time (or  $O(r_{\text{max}} r_{\text{max}}^{\#} + p_{\text{max}} \log \log n)$  deterministic time).

Consider now **nearest**( $v, k$ ) query. In order to find the shortest piece distance from  $v$  to some active color, we again iterate through all vertices that are piece-visible from  $v$ , and among these choose the closest to  $v$  active vertex. To find the shortest portal distance, we compute  $\min_{p \in \text{portals}(v)} (D_{p,v} + \text{top}(H_p))$ , where  $\text{top}(H_p)$  is the minimum distance stored on  $H_p$ . We take the minimum over piece and portal distance, and based on the vertex associated with this distance (recall that on heaps  $H_p$  we store pairs  $(N_i^p, v)$ ) we decide which active color is the closest to  $v$ .



This takes  $O(r_{\max}r_{\max}^{\#} + p_{\max})$  time. Next, we repeat this procedure  $k - 1$  times, and in  $l$ 'th iteration we ignore the distances to  $l - 1$  closest colors we found already. So if at  $l$ 'th iteration  $i_1, \dots, i_{l-1}$  are the first  $l - 1$  closest colors, then we look in this iteration for an active vertex in set  $\bigcup_i C_i \setminus (C_{i_1} \cup \dots \cup C_{i_{l-1}})$ . To that end, we iterate through vertices that are piece-visible from  $v$  and find the closest to  $v$  vertex of a color we are interested in. Then we iterate through portals and for every  $p \in \text{portals}(v)$  we find  $l$  smallest elements in a heap  $H_p$ . Since  $l \leq k = O(1)$ , this can be done in constant time. From these  $l$  elements, we select the minimum element that does not correspond to one of  $i_1, \dots, i_{l-1}$ . The procedure is repeated  $k$  times and, if  $k = O(1)$ , every step takes as much time as finding the nearest active vertex. Hence the total time to answer  $\text{nearest}(v, k)$  is  $O(r_{\max}r_{\max}^{\#} + p_{\max})$  if  $k$  is constant.  $\square$

**Lemma 5.5.** *Assume we are given an  $\alpha$ -approximate generic distance oracle for  $G = (V, E, d)$  and let  $n = |V|$ . One can implement the fully dynamic scheme in  $O((p_{\text{num}} + r_{\text{num}}) \log n)$  worst case time per update and  $O(r_{\max}r_{\max}^{\#} + p_{\max})$  worst case time per query operation. The data structure requires  $O((p_{\text{tot}} + r_{\text{tot}}) \log n)$  space and can be initialized in  $O((p_{\text{tot}} + r_{\text{tot}}) \log n)$  time.*

*Proof.* Here we employ the data structure of Lemma 5.2 to maintain heaps  $H_p$  and for a constant time access to  $N_i^p$  for each portal  $p \in \mathcal{P}$ . Recall, that this data structure requires  $O(p_{\text{tot}} \log n)$  space and the same initialization time. It allows updates in  $O(p_{\text{num}} \log n)$  worst case time.

Hence, the only thing left to do is to handle distance queries. This is done in the same manner as in the proof of Lemma 5.4 with one difference. The update operations do not iterate through all vertices that change color, so we cannot maintain a color flag in each vertex so that it knows what color it is. We could access the color of a vertex via ET-tree  $ET_i$  containing it, but that takes  $O(\log n)$  time per vertex. Lemma 5.3 comes in useful here. We maintain a structure which requires  $O(r_{\text{tot}} \log n)$  space and the same initialization time, and can be updated in  $O(r_{\text{num}} \log n)$  worst case time. This structure allows us to list all vertices active colors within a given piece. Recall that in order to answer distance queries we first find the related piece distance and compare it with a portal distance. In order to find the piece distance, for each piece  $\tau \in \text{pieces}(v)$  we list all active vertices in this piece and mark their colors in some global array. We run the linear time single source shortest paths algorithm [24] on  $\tau$  with the source at  $v$  in order to find the distances from  $v$  to vertices in  $\tau$ . Then we iterate through vertices  $w \in \tau$ , and knowing their color from the global array, we choose the closest vertex we are interested in. This takes linear time per piece by Lemma 5.3. Finding the portal distances is essentially the same as in the proof of Lemma 5.4, with one small difference: the heaps  $H_p$  can be easily modified to store triples  $(N_i^p, w, i)$ , where  $w$  is the vertex corresponding to distance  $N_i^p$  and  $i$  is its color, so that when we know the colors of vertices in the heaps  $H_p$ . The Query time of  $O(r_{\max}r_{\max}^{\#} + p_{\max})$  follows.  $\square$

**Corollary 5.6.** *Assume we are given an  $\alpha$ -approximate generic distance oracle for  $G = (V, E, d)$  and let  $n = |V|$ . One can implement the fully dynamic scheme in  $O(p_{\text{num}} \log n)$  worst case time per update and  $O(r_{\max}r_{\max}^{\#} \log n + p_{\max})$  worst case time per query operation. The data structure requires  $O(p_{\text{tot}} \log n)$  space and can be initialized in  $O(p_{\text{tot}} \log n)$  time.*

*Proof.* Instead of using structure of Lemma 5.3, we pay an additional logarithmic factor per distance query, as while iterating through a vertex seen by the queried vertex  $v$ , we additionally check its color using the ET-trees.  $\square$

	Initialization	Query	Update
Incremental	$O(p_{\text{tot}} \log n)$	$O(r_{\text{max}} r_{\text{max}}^{\#} + p_{\text{max}})$ exp. $O(r_{\text{max}} r_{\text{max}}^{\#} + p_{\text{max}} \log \log n)$ det.	$O(p_{\text{max}} \log n)$
Fully dynamic 1	linear	$O(r_{\text{max}} r_{\text{max}}^{\#} + p_{\text{max}})$	$O((p_{\text{num}} + r_{\text{num}}) \log n)$
Fully dynamic 2	linear	$O(r_{\text{max}} r_{\text{max}}^{\#} \log n + p_{\text{max}})$	$O(p_{\text{num}} \log n)$

Table 2: Summary of our generic constructions of vertex-color oracles.

## 5.2 3-approximate oracles for general graphs

In the previous section we showed how to efficiently implement the incremental and fully dynamic vertex-color distance oracle scheme. These constructions were based on generic oracles, which provide an choice and assignment of portals and pieces to vertices in the graph and store partial information about the distances in the graph. In other words, we showed how to implement both schemes given portals and pieces.

Here we show how to choose portals and pieces, i.e., how to construct generic oracles. In principle we want to keep the right balance between the parameters introduced in the previous section in order to obtain the best running times.

The first generic oracle we present here is a 3-approximate generic oracle for general weighted graphs. It is based on ideas by Thorup and Zwick [37]. In their construction, however, the procedure that measures distance between  $u$  and  $v$  may give different results when the two vertices are swapped. This does not comply with Definition 3.4, which requires symmetry in answering queries, that is, if the portal distance between two vertices  $u$  and  $w$  is finite, it has to be the same in  $\text{pieces}(u)$  and  $\text{pieces}(w)$ . Also, our data structures measure distance between some vertex and the nearest vertex of some given color. We extend the construction, so that the oracle works on portal and piece distances, which are invariant under swapping vertices. Some additional properties that we obtain are summarized in Section 5.4.

**Lemma 5.7.** *Let  $G = (V, E, d)$  be a weighted graph,  $n = |V|$  and  $m = |E|$ . In  $O(\sqrt{n}(m + n \log n))$  expected time we can build a 3-approximate generic distance oracle for  $G$  that uses expected  $O(n\sqrt{n})$  space. The set of portals  $\mathcal{P} \subseteq V$  of size  $O(\sqrt{n})$  is the same for every vertex, i.e.,  $\text{portals}(v) = \mathcal{P}$  for every  $v \in V$ . Moreover, for each  $v \in V$ ,  $\text{pieces}(v)$  contains a single graph that is a star on  $O(\sqrt{n})$  vertices (in expectation), whose center is  $v$ . The parameters of the oracle are as in Table 3.*

*Proof.* We choose  $\mathcal{P} \subseteq V$  uniformly at random with probability  $1/\sqrt{n}$ , so in expectation it is a subset of  $V$  of size  $\sqrt{n}$ . For each  $v \in V$  we set  $\text{portals}(v) = \mathcal{P}$ , so the number of portals is  $p_{\text{num}} = O(\sqrt{n})$  in expectation and the expected number of portals assigned to a vertex is  $p_{\text{max}} = O(\sqrt{n})$ . Also, the sum of sizes of all clusters is  $p_{\text{tot}} = O(n\sqrt{n})$  in expectation. For each  $w \in V$  we define  $p_w$  to be the portal which is closest to  $w$ , i.e.  $\delta(w, p_w) = \min_{p \in \mathcal{P}} \delta(w, p)$ . By  $B_w$  we denote the set of vertices that are closer to  $w$  than any portal, i.e.  $B_w = \{v \in V : \delta(w, v) < \delta(w, p_w)\}$ . In principle our goal is to create a single piece per vertex (i.e.,  $r_{\text{max}} = 1$ ). A star  $w \times (B_w \setminus \{w\})$  satisfies all the requirements but the symmetry, that is the piece distance between  $u$  and  $w$  may be different when measured in  $\text{pieces}(u)$  from the one given by  $\text{pieces}(w)$ . In order to ensure the symmetry, we reduce  $B_w$  to  $B_w^{\text{red}} \subseteq B_w$ . The reduction, for every vertex  $w$ , removes from  $B_w$  vertices  $u \in B_w$  such that  $w \notin B_u$ . For every  $w \in V$  we declare  $\text{pieces}(w)$  to contain a single graph  $r$  with the

parameter	$p_{\text{tot}}$	$r_{\text{max}}$	$r_{\text{max}}^{\#}$	$p_{\text{num}}$	$p_{\text{max}}$
value	$O(n\sqrt{n})$	1	$O(\sqrt{n})$	$O(\sqrt{n})$	$O(\sqrt{n})$

Table 3: Summary of parameters: all parameters but  $r_{\text{max}}$  are in expectation.

vertex set  $V(r) = B_w^{\text{red}}$ , the edge set  $E(r) = \{uw : u \in B_w^{\text{red}} \setminus \{w\}\}$  and the weights on the edges  $d_r(w, u) = \delta_G(w, u)$  for every  $u \in B_w^{\text{red}} \setminus \{w\}$ . As shown in [37], the expected size of each  $B_w$ , and consequently of each piece  $r_{\text{max}}^{\#} = O(\sqrt{n})$ .

To build the oracle we pick  $\mathcal{P}$ , compute shortest path distances from every vertex of  $\mathcal{P}$  using Dijkstra’s algorithm, and store the distances. To obtain distances from  $w$  to vertices of  $B_w^{\text{red}}$ , we also run a modified version of Dijkstra’s algorithm starting from each  $w \in V$  to find shortest distances inside  $B_w$ . We break the execution of the algorithm as soon as we visit a vertex from  $\mathcal{P}$ ,<sup>5</sup> and store this vertex as  $p_w$ .

The total time of running shortest paths algorithms from all portals is  $|\mathcal{P}| \cdot O(m + n \log n) = O(\sqrt{n}(m + n \log n))$ . Computing the distances for sets  $B_w$  requires  $n \cdot O(m/\sqrt{n} + n \log n/\sqrt{n}) = O(\sqrt{n}(m + n \log n))$  expected time. To store resulting distances and graphs we require  $O(n\sqrt{n})$  expected space.

Let us now prove that we obtain a generic oracle. Consider two vertices  $u$  and  $w$ . If  $u \in B_w^{\text{red}}$ , then, by the construction,  $w \in B_u^{\text{red}}$  and the shortest distance between them can be measured exactly both by inspecting  $\text{pieces}(u)$  and  $\text{pieces}(w)$ . Now, assume  $u \notin B_w^{\text{red}}$ . This means that  $u \notin B_w$  or  $w \notin B_u$ . Without loss of generality, assume  $u \notin B_w$ . This implies  $\delta(u, w) \geq \delta(w, p_w)$ . Then,  $\delta(w, p_w) + \delta(p_w, u) \leq \delta(w, p_w) + (\delta(p_w, w) + \delta(w, u)) \leq 3\delta(u, w)$ . Obviously  $p_w \in \text{portals}(v) \cap \text{portals}(w)$ .  $\square$

We now turn the oracle of Lemma 5.7 into a vertex-color distance oracle.

**Theorem 5.8.** *Let  $G = (V, E)$  be a graph,  $n = |V|$ . There exists a 3-approximate vertex-color distance oracle for  $G$ , which:*

- may answer  $\text{distance}(v, i)$  and  $\text{nearest}(v, k)$  queries in  $O(\sqrt{n} \log n)$  expected time, provided that  $k$  is a constant,
- handles  $\text{merge}(i, j, u, v)$ ,  $\text{split}(l, u, v)$ ,  $\text{activate}(i)$  and  $\text{deactivate}(i)$  operations in  $O(\sqrt{n} \log n)$  expected time,
- can be constructed in  $O(\sqrt{n}(m + n \log n))$  expected time,
- uses expected  $O(n\sqrt{n} \log n)$  space.

*Proof.* We use Lemma 5.7 to construct a generic 3-approximate oracle in  $O(\sqrt{n}(m + n \log n))$  expected time using expected  $O(n\sqrt{n})$  space. The parameters of this oracle are shown in Table 3. We plug this parameters into Corollary 5.6 and obtain what is claimed.  $\square$

We also observe that, using the ideas of [37], we can improve the running time at the cost of worse — but still constant — approximation ratio.

<sup>5</sup>This modification is described in detail in [37].

**Theorem 5.9.** *Let  $G = (V, E)$  be a graph,  $n = |V|$  and  $l$  be a parameter. There exists a  $(2l - 1)$ -approximate vertex-color distance oracle for  $G$ , which:*

- *answers  $\text{distance}(v, i)$  and  $\text{nearest}(v, k)$  queries in  $O(\ln^{1/l})$  expected time, provided that  $k$  is a constant,*
- *handles  $\text{merge}(i, j)$  and  $\text{activate}(i)$  operations in  $O(\ln^{1/l} \log n)$  amortized expected time,*
- *can be constructed in  $O(lmn^{1/l} + \ln^{1+1/l} \log n)$  expected time,*
- *uses expected  $O(\ln^{1+1/l})$  space.*

*Proof.* We again follow the ideas of [37]. In particular the construction is the same as in Lemma 7.5. In  $O(lmn^{1/l})$  we compute for each  $v \in V$  a set  $B(v)$  called a bunch. The expected size of each bunch is  $O(\ln^{1/l})$ . From the query algorithm in [37] it follows that for any two  $u, v \in V$  there exists  $w \in B(u) \cap B(v)$  such that  $\delta(u, w) + \delta(w, v) \leq (2l - 1)\delta(u, v)$ . For every vertex  $v$ , we set  $\text{portals}(v) = B(v)$  and  $\text{pieces}(v) = \emptyset$ . We obtain the expected total number of portals  $p_{\text{tot}} = O(kn^{1+1/l})$  and the expected number of portals per vertex  $p_{\text{max}} = O(\ln^{1/l})$ . The number of pieces assigned to a vertex is  $r_{\text{max}} = 0$ . By Lemma 5.4 (see Table 2), we obtain an incremental oracle using total  $O(\ln^{1+1/l})$  expected space, which is initialized in  $O(lmn^{1/l} + \ln^{1+1/l} \log n)$  expected time, handles updates in  $O(\ln^{1/l} \log n)$  amortized expected time and handles queries in  $O(\ln^{1/l})$  expected time.  $\square$

### 5.3 $(1 + \varepsilon)$ -approximate oracles for planar graphs

In this section we show how to employ generic oracle constructions from Section 5.1 in order to obtain dynamic  $(1 + \varepsilon)$ -approximate vertex-color distance oracles for planar graphs. We are able to obtain a better approximation and, for the incremental variant, better running times, due to planar separators. The separators are used to construct small sets of portals. We follow here the construction of Thorup [36], extending it slightly in order to support vertex colors.

The base building block of the construction is the following:

**Lemma 5.10.** *Let  $G = (V, E)$  be a planar graph,  $Q$  be a shortest path in  $G$  of length at most  $\alpha$ ,  $n = |V|$  and  $\varepsilon > 0$ . There exists a set  $P \subseteq V$  of size  $O(\varepsilon^{-1})$  such that the following holds. Let  $u, v \in V$ . Assume that the shortest path between  $u$  and  $v$  crosses  $Q$  and  $\delta_G(u, v) \in [\frac{\alpha}{2}, \alpha]$ . Then, there exists  $p \in P$  such that  $\delta_G(u, v) \leq \delta_G(u, p) + \delta_G(p, v) \leq (1 + \varepsilon)\delta_G(u, v)$ . The set  $P$  can be computed in  $O(n)$  time.*

*Proof.* The set  $P$  is constructed by putting portals on the path  $Q$ . We let the first portal be the first vertex of  $Q$ , and then put portals one by one, each at a distance  $\frac{\varepsilon}{4}d(Q)$  from the previous one. If this would require us to put a portal on an edge  $e$  of  $Q$ , we put two portals in the two endpoints of  $e$ . Clearly,  $|P| = O(\varepsilon^{-1})$ .

It remains to show that the oracle correctly approximates distance between  $u$  and  $v$ . Let  $X_{v,u}$  be the shortest path connecting  $u$  to  $v$ . Assume that it crosses  $Q$  in vertex  $q \in G$ . Let  $p_q \in P$  be the portal closest to  $q$ , that is,  $\delta(q, p_q) \leq \frac{\varepsilon}{4}d(Q)$ . Then

$$\delta(u, v) \leq \delta(v, p_q) + \delta(p_q, u) \leq d(X_{v,u}) + \frac{\varepsilon}{2}d(Q) \leq d(X_{v,u}) + \frac{\varepsilon}{2}\alpha \leq (1 + \varepsilon)d(X_{v,u}),$$

as requested.  $\square$

We now show how to use the above Lemma to construct an oracle that may answer queries about any two vertices in the graph. This construction was done by Thorup [36], but we need to extend it slightly in order to support vertex colors.

First, we fix  $\alpha$  and show how to build a generic oracle that finds all distances in range  $[\frac{\alpha}{2}, \alpha]$ . The construction begins by replacing  $G$  with a set of graphs  $G_1^\alpha \dots G_k^\alpha$  of diameter  $O(\alpha)$ , which carry all necessary distance information, and then building a recursive subdivision for each  $G_i^\alpha$ . The division is obtained by using paths of length  $\leq \alpha$  as separators. Consider a recursive call for a subgraph  $G'$  and one of the separator paths  $Q$ . We run the construction of Lemma 5.10 on  $G'$  and  $Q$  and then add all vertices of  $P$  to the portal sets of all vertices of  $G'$ .

Finally, we repeat the construction for  $\log D$  different values of  $\alpha$ , so that the intervals  $[\frac{\alpha}{2}, \alpha]$  altogether cover the entire range of possible distances.

**Lemma 5.11** ([36]). *Given a planar  $G$  and a scale  $\alpha$  we can construct a series of graphs  $G_1^\alpha, \dots, G_k^\alpha$  satisfying the following properties:*

- *the total number of vertices and edges in all  $G_i^\alpha$  is linear in the size of  $G$ ,*
- *$G_i^\alpha$  is a minor of  $G$  (in particular, every  $G_i^\alpha$  is planar),*
- *each  $G_i^\alpha$  has exactly one vertex obtained by contractions, which we call a supervertex,*
- *each  $G_i^\alpha$  has a spanning tree  $T_i^\alpha$  of depth at most  $3\alpha$  rooted in the supervertex,*
- *each vertex  $v$  has index  $\tau(v)$  such that a vertex  $w$  is at distance  $d \leq \alpha$  from  $v$  in  $G$  if and only if  $d$  is the shortest distance from  $v$  to  $w$  in  $G_{\tau(v)}^\alpha$ ; this also holds after removing the supervertex,*
- *each vertex  $v$  is in exactly three graphs:  $G_{\tau(v)-1}^\alpha$ ,  $G_{\tau(v)}^\alpha$  and  $G_{\tau(v)+1}^\alpha$*

*Proof sketch.* For the full proof we refer to [36]. We construct a partition of  $V(G)$  into layers  $L_i$ . Let  $L_{<i} = \bigcup_{j=0}^{i-1} L_j$ . We arbitrarily choose vertex  $v_0$  and let  $L_0 = \{v_0\}$ . Layer  $L_{i+1}$  contains new vertices reachable from vertices of  $L_i$  within the distance of  $\alpha$ , i.e., all reachable vertices not in previous layers. To create  $G_i^\alpha$  we take a subgraph of  $G$  induced by  $L_{<i+2}$  and contract vertices of  $L_{<i-1}$ . For a vertex  $v$  in  $G$  we let  $\tau(v)$  be the index of its layer. Note that any path starting in  $v$  of length at most  $\alpha$  is contained in  $G_{\tau(v)}^\alpha$  and does not cross the contracted supervertex.  $\square$

As a consequence of the above, each distance query concerning  $v$  can be answered by accessing  $G_{\tau(v)}^\alpha$ , provided that the corresponding distance is at most  $\alpha$ . In our construction we also use the following Lemma for finding separators.

**Lemma 5.12** (Lemma 2.3 in [36]). *In linear time, given an undirected planar graph  $G$  with a rooted spanning tree  $T$  and nonnegative vertex weights, we can find three vertices  $u, v$  and  $w$  such that each component of  $G \setminus (V(T(u) \cup T(v) \cup T(w)))$  has at most half the weight of  $G$ , where  $T(x)$  denotes the path from  $x$  to the root in  $T$ .*

We are now ready to show the construction.

**Lemma 5.13.** *Let  $G$  be a planar graph,  $n = |V|$ ,  $\varepsilon > 0$  and  $D$  be the stretch of the metric induced by  $G$ . A  $(1 + \varepsilon)$ -approximate generic distance oracle for  $G$  can be built in  $O(\varepsilon^{-1} n \log n \log D)$  time. For every  $v \in V$ ,  $|\text{portals}(v)| = O(\varepsilon^{-1} \log n \log D)$  and  $\text{pieces}(v) = \emptyset$ .*

*Proof.* Initially, we set  $portals(v) = \emptyset$  for every  $v \in V$ . During the construction, the algorithm adds new portals.

We first show the construction for a fixed value of  $\alpha$ . It starts with computing the sequence  $G_1^\alpha, \dots, G_k^\alpha$  by applying Lemma 5.11. Then, we recursively decompose each  $G_i^\alpha$ .

We use Lemma 5.12 on  $G_i^\alpha$  and  $T_i^\alpha$  setting the weight of the supervertex to 0 and the weights of other vertices to 1. Lemma 5.12 guarantees that  $G_i^\alpha$  has a separator  $Q$  composed of three paths of  $T_i^\alpha$ . The first step is to represent all distances that correspond to paths crossing  $Q$  in a graph obtained from  $G_i^\alpha$  by removing the supervertex. Denote this graph by  $\tilde{G}_i^\alpha$ . Observe that although  $G_i^\alpha$  is a minor of  $G$ ,  $\tilde{G}_i^\alpha$  is a *subgraph* of  $G$ , as we have removed the only vertex obtained by contractions.

We remove the supervertex from  $Q$  and obtain a separator  $\tilde{Q}$  of  $\tilde{G}_i^\alpha$ . Note that the three paths that form  $Q$  start in the supervertex, so  $\tilde{Q}$  also consists of three paths. Since these paths have length  $\leq 3\alpha$ , we may split each of them into three paths of length  $\leq \alpha$ , thus decomposing  $\tilde{Q}$  into nine paths  $Q_1, \dots, Q_9$ , each of length  $\leq \alpha$ .

To represent all distances in  $\tilde{G}_i^\alpha$  crossing  $\tilde{Q}$ , we use the construction of Lemma 5.10 for each of the nine paths  $Q_j$ ,  $j \in \{1 \dots 9\}$ . In this way for a path  $Q_j$  we obtain a set of vertices  $P_j$ . We add those vertices to the set  $portals(v)$  for every vertex  $v$  of  $\tilde{G}_i^\alpha$ . For every new portal, we compute (and store) distances to every vertex of  $\tilde{G}_i^\alpha$  using linear time algorithm for shortest paths in planar graphs [24].

Next in  $G_i^\alpha$  we contract separator  $Q$  and obtain a planar graph  $G'$ . Since the separator  $Q$  contains the supervertex of  $G_i^\alpha$ , the contracted  $Q$  in  $G'$  becomes a supervertex  $r'$ , and tree  $T_i^\alpha$  becomes a tree  $T'$  rooted at  $r'$ . We note that  $T'$  is a tree spanning  $G'$  rooted in  $r'$  and its depth is at most  $3\alpha$ . Moreover, removing  $r'$  splits  $G'$  into at least two components (cut out in  $G_i^\alpha$  by the separator  $Q$ ), each containing at most half vertices of  $V(G')$ . Each of these components has a spanning tree which is a subtree of  $T'$  rooted at the supervertex  $r'$ . Hence, we can recurse on each of these components (each component is taken together with  $r'$ ).

The construction allows us to approximate distances in range  $[\frac{\alpha}{2}, \alpha]$ , so we repeat it  $\log D$  times for different values of  $\alpha$ , where  $D$  is the stretch of the metric induced by  $G$ .

We now show that we obtain a  $(1 + \varepsilon)$ -approximate generic oracle. Consider two vertices  $u$  and  $v$ . Let  $\alpha$  be the scale such that  $\delta_G(u, v) \in [\frac{\alpha}{2}, \alpha]$ . There exists a graph  $G_i^\alpha$  such that  $\delta_{G_i^\alpha}(u, v) = \delta_G(u, v)$ . Consider the shortest path  $X_{u,v}$  in  $G_i^\alpha$  between  $u$  and  $v$ . In some recursive call of the decomposition algorithm, the path crosses the separator  $Q$ . It is easy to see that in that step, we add both to  $portals(u)$  and  $portals(v)$  a vertex  $p$  such that  $\delta_G(u, p) + \delta_G(p, v) \leq (1 + \varepsilon)\delta_G(u, v)$ . Moreover, if we denote by  $G'$  the subgraph in the recursion in exactly that step, then  $G'$  contains  $Q'$ , so  $\delta_{G'}(u, p) = \delta_G(u, p)$  and  $\delta_{G'}(p, v) = \delta_G(p, v)$ .

It remains to bound the running time and the number of portals of each vertex. Consider a vertex  $v$  of  $G$ . For each of  $O(\log D)$  values of  $\alpha$ ,  $v$  belongs to three layers  $G_i^\alpha$ . Let  $n_i = |V(G_i^\alpha)|$ . In a recursive subdivision of  $G_i^\alpha$ ,  $v$  can appear once on each of the  $O(\log n_i)$  levels. On a single level, we add  $O(\varepsilon^{-1})$  vertices to  $portals(v)$ . Thus, since  $\log n_i \leq \log n$ ,  $|portals(v)| = O(\varepsilon^{-1} \log n \log D)$ . Observe that the distances inside each cluster are computed in linear time, so the initialization time is equal to the total size of the clusters. The lemma follows.  $\square$

We combine the incremental oracle construction from Lemma 5.4 (see Table 2) with the above lemma to obtain the following result.

**Theorem 5.14.** *Let  $G = (V, E)$  be a planar graph,  $n = |V|$  and  $D$  be the stretch of the metric induced by  $G$ . For every  $\varepsilon > 0$  there exists an  $(1 + \varepsilon)$ -approximate vertex-color distance oracle for  $G$ , which:*

- *answers  $\text{distance}(v, i)$  and  $\text{nearest}(v, k)$  queries in  $O(\varepsilon^{-1} \log n \log D)$  expected time or  $O(\varepsilon^{-1} \log n \log \log n \log D)$  deterministic time, provided that  $k$  is a constant,*
- *handles  $\text{merge}(i, j)$  and  $\text{activate}(i)$  operations in  $O(\varepsilon^{-1} \log^2 n \log D)$  expected amortized time or  $O(\varepsilon^{-1} \log^2 n \log \log n \log D)$  deterministic amortized time,*
- *can be constructed in  $O(\varepsilon^{-1} n \log^2 n \log D)$  time,*
- *uses  $O(\varepsilon^{-1} n \log n \log D)$  space.*

*Proof.* Lemma 5.13 gives a generic oracle with  $p_{\max} = \max |portals(v)| = O(\varepsilon^{-1} \log n \log D)$ , and hence the total size of all clusters  $p_{\text{tot}} = O(\varepsilon^{-1} \log n \log D)$ . In addition to that every vertex is assigned an empty set of pieces so  $r_{\max} = 0$ . We employ Lemma 5.4 (see Table 2) to obtain what is claimed.  $\square$

We now describe how to build a fully dynamic vertex-color distance oracle, which also supports `split` and `deactivate` operations. Its running time depends on two main factors. By Lemma 5.2 the running time of an update operation is proportional to the number of portals, whereas the query time depends on the size of pieces. Thus, in our construction we try to keep both these values low.

**Lemma 5.15.** *Let  $G = (V, E)$  be a planar graph,  $n = |V|$  and  $D$  be the stretch of the metric induced by  $G$ . For every  $\rho$  such that  $1 \leq \rho \leq n$ , in  $O(\varepsilon^{-1} n \log n \log D)$  time we can build an  $(1 + \varepsilon)$ -approximate generic distance oracle for  $G$ . The total number of portals is  $O(\varepsilon^{-1} \frac{n}{\rho} \log D)$ , and the total size of all of clusters is  $O(\varepsilon^{-1} n \log n \log D)$ . For every  $v \in V$ ,  $\text{pieces}(v)$  is a set consisting of  $O(\log D)$  planar graphs of size  $O(\rho)$ . There are  $O(\frac{n}{\rho} \log D)$  distinct pieces in total. The parameters are summarized in Table 4.*

*Proof.* For the purpose of this construction a graph is called *big* if it contains at least  $\frac{\rho}{3}$  vertices and *small* otherwise. Moreover, a graph with more than  $\rho$  vertices is called *too big*. We aim to alter the construction of a generic oracle from Lemma 5.13 in such a way that the recursive division of graphs  $G_i^\alpha$  produces subgraphs which are neither small nor too big.

First, we use Lemma 5.11 to obtain a sequence  $G_i^\alpha$ . As long as the sequence contains two consecutive small layers  $G_i^\alpha, G_{i+1}^\alpha$ , we merge them into one layer. Since the layers are minors of  $G$ , we first need to remove the supervertex from both of them. After that, they become subgraphs of  $G$ , so we can merge them in a natural way.

Note that after this process every small layer is a neighbor of a big layer, so the number of small layers is  $O(\frac{n}{\rho})$ . Additionally, each too big layer is not a result of merging smaller layers and satisfies conditions of Lemma 5.12.

We now deal with too big layers  $G_i^\alpha$ . For each such layer, we perform a recursive subdivision similar to the one in the proof of Lemma 5.13. However, we introduce two changes. First, we assure that at each step, the graph is divided into exactly two subgraphs of roughly even size. Recall that in the construction, we contract the separator in a currently handled graph  $G'$  into a supervertex  $r'$ , and each connected component of  $G' \setminus r'$  is of size at most  $\frac{1}{2}$  of the size of  $G'$ . By standard

$p_{\text{tot}}$	$p_{\text{max}}$	$p_{\text{num}}$	$r_{\text{num}}$	$r_{\text{max}}$	$r_{\text{max}}^{\#}$	$r_{\text{tot}}$
$O(\varepsilon^{-1}n \log n \log D)$	$O(\varepsilon^{-1} \log n \log D)$	$O(\varepsilon^{-1} \frac{n}{\rho} \log D)$	$\frac{n}{\rho} \log D$	$O(\log D)$	$\Theta(\rho)$	$O(n \log D)$

Table 4: Summary of parameters.

arguments, we can group the components of  $G' \setminus r'$  into exactly two groups, each of total size at most  $\frac{2}{3}$  of the size of  $G'$ . We recurse on the union of the graphs in each group separately.

The second modification to the algorithm from Lemma 5.13 is that we stop the recursion as soon as we reach a graph which is not too big. As a result, leaves of the recursive subdivision correspond to big graphs, which are not too big.

Consider a layer  $G_i^\alpha$  (after the merging step). If  $G_i^\alpha$  is not too big we add  $G_i^\alpha$  to  $\text{pieces}(v)$  for every  $v \in V(G_i^\alpha)$  (except for the supervertex). Otherwise, we decompose the graph recursively and for every separator that we use, we apply Lemma 5.10 and add the obtained set to  $\text{portals}(v)$ . Moreover, once we reach a graph  $G'$  which is not too big and the recursion stops, we add this graph to  $\text{pieces}(v)$  for every  $v \in V(G')$ .

The entire construction is repeated for  $O(\log D)$  values of  $\alpha$ , so that the intervals  $[\frac{\alpha}{2}, \alpha]$  altogether cover all possible distances.

We now show that we obtain a  $(1 + \varepsilon)$ -approximate generic oracle. Consider two vertices  $u$  and  $v$  and the value of  $\alpha$  such that  $\delta_G(u, v) \in [\frac{\alpha}{2}, \alpha]$ . There exists a graph  $G_i^\alpha$  such that  $\delta_{G_i^\alpha}(u, v) = \delta_G(u, v)$ . Consider the shortest path  $X_{u,v}$  in  $G_i^\alpha$  between  $u$  and  $v$ . If layer  $G_i^\alpha$  is not too big then it is added both to  $\text{pieces}(u)$  and  $\text{pieces}(v)$  and the property follows trivially. Otherwise, there are two cases to consider. One possibility is that in some recursive call of the decomposition algorithm, the path crosses the separator  $Q$ . In this case we add both to  $\text{portals}(u)$  and  $\text{portals}(v)$  a vertex  $p$  such that  $\delta_G(u, p) + \delta_G(p, v) \leq (1 + \varepsilon)\delta_G(u, v)$ . Otherwise, we reach a subgraph which is not too big and contains the entire path  $X_{u,v}$ . This subgraph is then added to  $\text{pieces}(u)$  and  $\text{pieces}(v)$ .

It remains to bound the size of the oracle. From the construction it follows that the size of every piece is  $\Theta(\rho)$  so the maximal size of a piece is clearly  $r_{\text{max}}^{\#} \in \Theta(\rho)$ . For every  $\alpha$  during the recursive decomposition of  $G_i^\alpha$  we create pieces that are nonoverlapping subgraphs of  $G_i^\alpha$ . Thus, for a fixed  $\alpha$ , the sum of sizes of all pieces that we create is  $O(n)$ . We use  $O(\log D)$  values of  $\alpha$ , so the total size of all pieces is  $r_{\text{tot}} = O(n \log D)$ .

Moreover, for a fixed  $\alpha$ , because all but  $O(\frac{n}{\rho})$  pieces have size  $\Theta(\rho)$ , there are  $O(\frac{n}{\rho})$  pieces in total. So the total number of pieces is  $r_{\text{num}} = O(\frac{n}{\rho} \log D)$  as there are  $\log D$  values of  $\alpha$ . For every layer  $G_i^\alpha$  and a fixed vertex  $v$ , at most one piece may be added to  $\text{pieces}(v)$ . Each vertex belongs to  $O(\log D)$  layers  $|\text{pieces}(v)| = O(\log D)$  for every  $v \in G$  and hence  $r_{\text{max}} = O(\log D)$ .

Let us now bound the total number of portals. Consider the recursive subdivision algorithm. We add  $O(\varepsilon^{-1})$  new portals every time when we divide a graph into two subgraphs with a separator. For a fixed  $\alpha$ , in the end all layers are divided into  $O(\frac{n}{\rho})$  subgraphs (which become pieces). Thus, the number of division steps is also  $O(\frac{n}{\rho})$ , and they create  $O(\varepsilon^{-1} \frac{n}{\rho})$  new portals. This gives  $O(\varepsilon^{-1} \frac{n}{\rho} \log D)$  portals for all values of  $\alpha$ , in other words  $p_{\text{num}} = O(\varepsilon^{-1} \frac{n}{\rho} \log D)$ .

The total size of clusters is bounded by the total size of all clusters in the construction of Lemma 5.13, which is  $p_{\text{tot}} = O(\varepsilon^{-1}n \log n \log D)$ . Also after Lemma 5.13, the maximum number of portals assigned to a vertex is  $p_{\text{max}} = O(\varepsilon^{-1} \log n \log D)$ . The lemma follows.



□

**Theorem 5.16.** *Let  $G = (V, E)$  be a planar graph,  $n = |V|$  and let  $D$  be the stretch of the metric induced by  $G$ . For every  $\rho, \varepsilon > 0$ , there exists an  $(1 + \varepsilon)$ -approximate vertex-color distance oracle for  $G$ , which:*

- *may answer  $\mathbf{distance}(v, i)$  and  $\mathbf{nearest}(v, k)$  queries in  $O((\varepsilon^{-1} \log n + \rho) \log D)$  time, provided that  $k$  is a constant,*
- *handles  $\mathbf{merge}(i, j, u, v)$ ,  $\mathbf{split}(l, u, v)$ ,  $\mathbf{activate}(i)$  and  $\mathbf{deactivate}(i)$  operations in  $O(\varepsilon^{-1} \frac{n}{\rho} \log n \log D)$  time,*
- *can be constructed in  $O(\varepsilon^{-1} n \log^2 n \log D)$  time,*
- *uses  $O(\varepsilon^{-1} n \log^2 n \log D)$  space.*

*In addition to that, the oracles admits parameters as in Table 4.*

*Proof.* The theorem follows from combining the oracle of Lemma 5.15 (the parameters are given in Table 4) with Corollary 5.6 (see Fully Dynamic 1 oracle in Table 2).

□

## 5.4 Additional operations and properties

**Near-metric spaces** The generic oracles we build measure distance between  $u$  and  $v$  by taking the minimum of portal and piece distance (see Definition 3.4) between  $u$  and  $v$ . Let us call this distance an *oracle distance*. From the construction of our vertex-color distance oracles it is easy to observe that they answer  $\mathbf{distance}$  queries by choosing the minimum oracle distance to the vertex of respective color. A similar property holds for  $\mathbf{nearest}$  queries. Roughly speaking, the oracles behave as if they were exact and the real distance between two vertices was the oracle distance (though the oracle distance does not satisfy triangle inequality). This property, which is used in our algorithms, is summarized in the following.

**Corollary 5.17.** *Consider one of nearest neighbor oracles from Section 3.2. Assume it provides  $\mu$ -approximate distances between vertices of a graph  $G = (V, E, d_G)$ . Then, the oracle yields a  $\mu$ -near metric space  $GD$  that approximates  $G$ . In particular:*

- *The answer to a  $\mathbf{distance}(v, i)$  query is equal to  $\min\{d_{GD}(v, w) : w \text{ is active and has color } i\}$ .*
- *The answer to a  $\mathbf{nearest}(v, k)$  query is the  $k$ -th smallest value of  $\mathbf{distance}(v, i)$  over all active colors  $i$ .*

*Proof.* Consider a clique  $GD = (V, \binom{V}{2}, d_{GD})$ , where  $d_{GD}(u, v)$  is defined as the minimum over the piece distance and the portal distance between  $u$  and  $v$  (i.e., the oracle distance). By Definition 3.4 it is a well defined simple graph with finite weights. This graph certifies the claim stated in the corollary. □

**Reconnecting edges.** In one of our algorithms, at some point we encounter the following problem. We are given a set of active colors  $c_1, \dots, c_k$  that we need to connect together in a minimum spanning tree fashion. That is, we would like to compute the MST of a complete graph  $G_c$  over colors  $c_1, \dots, c_k$ , where the length of an edge connecting two colors is the distance between the nearest vertices of these colors. Solving this problem efficiently seems hard, even for the case when  $k = 2$ . However, we are able to show a solution for its simpler variant. Namely, we define the length of the edge in  $G_c$  to be the length of the shortest path that connects two vertices of respective colors and goes through a portal.

**Lemma 5.18.** *Assume we are given a vertex-color distance oracle with  $p_{num}$  portals and a set of active colors  $c_1, \dots, c_k$ . Consider a complete graph  $G_c$  on the set of colors, such that the length of an edge connecting two colors  $c_i$  and  $c_j$  is the length of a shortest path that connects vertices  $v_i$  and  $v_j$  of colors  $c_i$  and  $c_j$  and goes through a vertex in  $portals(v_i) \cap portals(v_j)$ . Then,  $MST(G_c)$  can be computed in  $O(kp_{num} + k \log k)$  time.*

*Proof.* We first build a graph  $G'_c$ , and then compute its minimum spanning tree.  $G'_c$  is a subgraph of the graph  $G_c$  defined above. For efficiency reasons it may not contain some edges of  $G_c$  that surely do not belong to  $MST(G_c)$ .

We iterate through all portals of the oracle and add edges to  $G'_c$ . Consider a portal  $p$  and assume that vertices of colors  $c_1, \dots, c_k$  are at distances  $d_1 \leq \dots \leq d_s$  from  $p$ . In particular,  $c_1$  is the color that is nearest from  $p$ . The shortest path connecting colors  $i$  and  $j$  going through  $p$  has length  $d_i + d_j$ . Observe that we may add to  $G'_c$  only edges corresponding to pairs of colors  $c_1, c_i$  (for  $i \geq 2$ ). Indeed, if  $i, j \geq 2$ , then  $d_i + d_j \geq d_1 + d_j$  and  $d_i + d_j \geq d_i + d_1$ . Thus, any other edge that we could add would be the heaviest on the cycle formed by colors  $c_1, c_i$  and  $c_j$  and as such it is not necessary for computing the MST.

Thus, we end up with a graph with  $k$  vertices and  $O(kp_{num})$  edges. We compute its MST with Prim's algorithm in  $O(kp_{num} + k \log k)$  time.  $\square$

**Efficient Steiner tree computation.** We also observe that a vertex-color distance oracle gives us an efficient way of computing an approximate Steiner tree for a query consisting of an entire set  $S$  of terminals.

**Lemma 5.19.** *Assume we are given a graph  $G = (V, E, d_G)$ ,  $n = |V|$  and an instance  $\mathbb{D}$  of a vertex-color distance oracle for  $G$ , which yields a near-metric space  $GD$ . Then, given a set  $S \subseteq V(G)$ , one can compute  $MST(GD[S])$  in  $O(|S|(p_{max} + r_{max}r_{max}^\# \log n))$  time, where  $p_{max}$  is the maximum number of portals per vertex,  $r_{max}$  is maximum number of pieces assigned to a vertex and  $r_{max}^\#$  bounds the size of a piece.*

*Proof.* We compute a minimum spanning tree of  $S$  in  $GD$  using Prim's algorithm. We start with  $T$  being an isolated vertex (an arbitrarily chosen element of  $S$ ) and, at each step, we choose an edge connecting  $S \setminus V(T)$  and  $V(T)$  in  $GD$  of minimum possible weight and add it to  $T$ . To choose such an edge, it suffices to choose the minimum among the portal and piece distance between  $S \setminus V(T)$  and  $V(T)$ .

First, consider piece distances. For each  $v \in S \setminus V(T)$  we maintain the piece distances between  $v$  and the nearest vertex (w.r.t piece distance) of  $V(T)$ . All these piece distances are kept in a heap  $H_v$ . In addition to that, we maintain a global heap  $H$  with the minimas of the heaps  $H_v$  for all vertices  $v \in S \setminus V(T)$ . The minimum piece distance between  $V(T)$  and  $S \setminus V(T)$  is the minimum

value on  $H$ . If a vertex  $u$  is added to  $V(T)$ , to update this information it suffices to run Dijkstra's shortest-path algorithm from  $u$  in each of  $pieces(u)$ , update information in heaps  $H_v$  for all vertices of  $S \setminus V(T)$  that were found, each time updating the value on  $H$  as well. For all vertices of  $S$  this requires  $O(|S|r_{\max}r_{\max}^{\#} \log n)$  time.

To find the minimum portal distance, for each portal  $p$  of a vertex in  $S$  we consider  $cluster(p)$  and store the distance  $d_T^p$  to the nearest vertex of  $V(T)$  and a heap of distances to vertices of  $S \setminus V(T)$  (both in  $cluster(p)$ ). The portal distance is the minimum over portals of the sum of  $d_T^p$  and the minimum on this portals heap. We again maintain a global heap which has access to minimum portal distance. Thus, for every vertex  $v \in S$ , we need to add distances to all  $portals(v)$  to a heap in  $O(|S| \cdot |portals(v)| \log n)$  time, each time updating the global heap as well. Whenever we connect a vertex  $u$  to the tree  $T$ , we update the heaps of all its portals and the global heap. This again takes a total  $O(|S| \cdot |portals(v)| \log n)$  time. The time for removing the elements is at most as big as the time for inserting them. The lemma follows.  $\square$

By Lemma 5.19, we obtain the following two corollaries, depending on which oracle we use.

**Corollary 5.20.** *Assume we are given a graph  $G = (V, E, d_G)$ ,  $n = |V|$ . After  $O(\sqrt{n}(m + n \log n))$  preprocessing, using  $O(n\sqrt{n} \log n)$  space, we can repeatedly compute a 3-approximation of  $MST(\overline{G}[S])$ , and consequently a 6-approximation of the minimum Steiner tree spanning  $S$  in  $G$ , in  $O(|S|\sqrt{n} \log n)$  time for any subset  $S \subseteq V$ .*

*Proof.* We employ as  $\mathbb{D}$  in Lemma 5.19 the fully dynamic oracle of Theorem 5.8. By Theorem 5.8,  $\mathbb{D}$  can be constructed in  $O(\sqrt{n}(m + n \log n))$  expected time and uses expected  $O(n\sqrt{n} \log n)$  space. By Lemma 5.19, we can query for  $MST(GD[S])$  in  $O(|S|\sqrt{n} \log n)$  time. By Lemma 2.1, it is a 6-approximation of the Steiner tree.  $\square$

**Corollary 5.21.** *Assume we are given a planar graph  $G = (V, E, d_G)$ ,  $n = |V|$ . After  $O(\varepsilon^{-1}n \log^2 n \log D)$  preprocessing, using  $O(\varepsilon^{-1}n \log n \log D)$  space, we can repeatedly compute a  $(1 + \varepsilon)$ -approximation of  $MST(\overline{G}[S])$ , and consequently a  $2 + \varepsilon$ -approximation of the minimum Steiner tree spanning  $S$  in  $G$ , in  $O(|S|\varepsilon^{-1} \log^2 n \log D)$  time for any subset  $S \subseteq V$ .*

*Proof.* We employ as  $\mathbb{D}$  in Lemma 5.19 the fully dynamic oracle of Theorem 5.14. By Theorem 5.14,  $\mathbb{D}$  can be constructed in  $O(\varepsilon^{-1}n \log^2 n \log D)$  time and uses  $O(\varepsilon^{-1}n \log n \log D)$  space. By Theorem 5.14,  $\mathbb{D}$  assigns at maximum  $p_{\max} = \max |portals(v)| = O(\varepsilon^{-1} \log n \log D)$  portals to a vertex and the number of pieces assigned to a vertex is  $r_{\max} = 0$ . By Lemma 5.19, we can query for  $MST(GD[S])$  in  $O(|S|\varepsilon^{-1} \log^2 n \log D)$  time. By Lemma 2.1, it is  $2(1 + \varepsilon)$ -approximation of the Steiner tree, which can be turned into  $(2 + \varepsilon)$ -approximation by using  $\varepsilon' = \varepsilon/2$  instead of  $\varepsilon$ .  $\square$

**Faster implementation of the Imase-Waxman incremental algorithm.** Imase and Waxman [28] proved that the natural online algorithm for Steiner tree — the one that attaches the new terminal to the closest previous one — achieves  $O(\log n)$  approximation guarantee. We remark here that our oracles give a different way to obtain this guarantee; our implementation is faster if the number of terminals is large in terms of the size of host graph.

**Theorem 5.22.** *Assume we are given a graph  $G = (V, E, d_G)$ ,  $n = |V|$ . Then there exists an algorithm that maintains an online Steiner tree upon a sequence of terminal addition (i.e., the algorithm needs to maintain a Steiner tree upon each addition, and is not allowed to alter the edges that are already in the tree) with  $O(\log n)$  approximation guarantee, using expected time*

$O(\sqrt{n} \log n)$  per terminal addition. The algorithm can be initialized in  $O(\sqrt{n}(m + n \log n))$  time and uses  $O(n\sqrt{n} \log n)$  space.

*Proof.* We use the vertex-color distance oracle of Theorem 5.8. In the oracle the current set of terminals is a single color, and only this color is active. To add a new terminal, it suffices to issue one **distance** and one **merge** operation. The new terminal is attached to not necessarily the closest previous terminal, but approximately closest.  $\square$

## 6 Sublinear time dynamic algorithms

We propose here dynamic algorithms for each of the dynamic scenarios we considered. The algorithms are obtained by implementing the online algorithms of Section 4 using vertex-color distance oracles from Section 3.2.

### 6.1 Decremental algorithm

We now consider the decremental scenario, in which the goal is to maintain an approximate Steiner tree subject to terminal deletions. The graph we work with is accessed with a fully dynamic vertex-color distance oracle  $\mathbb{D}$  (see Section 3.2), which yields a  $\mu$ -near metric space  $GD$  that approximates  $G$  (see Corollary 5.17). However, our algorithm not only uses the operations provided by the oracle, but also the sets of portals and pieces of each vertex. In particular, the algorithm uses the concepts of portal and piece distances, which are defined in Section 5.1.

We use the decremental scheme of Section 4.3, so for a fixed  $\varepsilon > 0$  we plan to maintain a  $(1 + \varepsilon)$ -approximation of  $MST(GD[S])$ . Following the decremental scheme, we set  $\eta = 1 + \lceil \varepsilon^{-1} \rceil$ .

Due to decremental scheme, when a vertex is removed from the terminal set, it is not necessary to remove it from the maintained tree  $T$  unless its degree in  $T$  is at most  $\eta$ . Thus an update may cause  $T$  to split into at most  $\eta$  trees. The main challenge of this section is to find in an efficient way the set of edges of minimum total cost that reconnects them.

In order to maintain the desired tree  $T$  we use the dynamic MSF algorithm (see Theorem 2.2). In what follows we construct a graph  $H$  on the vertex set  $V$  (the vertex set of both input graph  $G$  and  $GD$ ) with the property that there is a connected component  $H_0$  of  $H$  containing all current terminals. We declare  $T$  to be  $MST(H_0)$  (maintained by the dynamic MSF algorithm). All other connected components of  $H$  will be isolated vertices. Our goal now is to maintain  $H$  by adding and removing edges in such a way that maintaining the MSF of  $H$  with the dynamic MSF algorithm does in fact implement the decremental scheme of Section 4.3 on  $GD$ .

The simplest approach would be to add to  $H$ , for every two  $u, w \in S$ , an edge  $uw$  of length  $d_{GD}(u, w)$ . This approach, however, would obviously be inefficient. Instead of that, we use the structure of the vertex-color distance oracle. For every two  $u, w \in V(T)$  such that  $u$  is piece-visible from  $w$ , we add to  $H$  an edge  $uw$  of length being equal to the piece distance. Moreover, we will be adding some edges corresponding to portal distances, but they will be chosen in a careful way to ensure that the number of such edges is low.

Our algorithm uses a single instance  $\mathbb{D}$  of the fully dynamic vertex-color distance oracle. Initially all vertices are inactive and have distinct colors. Denote the terminal set by  $S$ . Using Lemma 5.19 we compute the initial minimum spanning tree  $T = MST(GD[S])$  and modify  $\mathbb{D}$  so that  $\mathbb{D}$  contains one active color, whose corresponding tree is equal to  $T$ . Other vertices remain inactive.

We initialize the edge set of  $H$  to be the set of all edges of the initial  $T$  and all edges between terminals corresponding to mutually piece-visible vertices. Note that we may break ties of weights in  $H$  in such a way that  $H$  has a unique minimum spanning forest. Consequently, the forest computed by the dynamic MSF algorithm contains the tree  $T$ .

We update the vertex-color distance oracle  $\mathbb{D}$  to maintain the following invariant: the colors of  $\mathbb{D}$  correspond to the connected components of the graph  $H$ , and the only active color is associated with a tree equal to  $T$ . Hence, any change to the spanning forest of  $H$  results in a constant number of modifications to  $\mathbb{D}$ .

According to the decremental scheme of Section 4.3, when a vertex of degree greater than  $\eta$  is deleted, we simply mark it as a non-terminal. Otherwise, if the procedure  $\text{REMOVE}_\eta(v)$  requires us to delete a vertex of degree at most  $\eta$  from the currently maintained tree, we update  $H$  and the distance oracle  $\mathbb{D}$ , as follows.

1. Split the active color  $V(T)$  in  $\mathbb{D}$  into  $s + 1$  pieces being  $\{v\}$  and the  $s$  connected components of  $T \setminus \{v\}$ , and then deactivate  $\{v\}$ .
2. Compute the set of reconnecting edges that correspond to portal distances (using Lemma 5.18) and add these edges to  $H$ .
3. Remove all edges incident to  $v$  from  $H$ .
4. At this point, the dynamic MSF algorithm may update the maintained tree. Update the vertex-color distance oracle  $\mathbb{D}$ , so that it reflects these changes of the MSF of  $H$ .

We now show that our algorithm indeed implements the decremental scheme of Section 4.3.

**Lemma 6.1.** *After each removal of a vertex from  $T$ ,  $T$  is a minimum spanning tree of  $GD[V(T)]$  and  $T$  implements the decremental scheme of Section 4.3.*

*Proof.* Clearly, at the initialization step for a given terminal set  $S$ , in  $H$  there is a component  $H_0$  containing  $T = \text{MST}(GD[S])$  and all vertices of  $V \setminus S$  are isolated. Moreover, the minimum spanning tree of  $H_0$  computed by the dynamic MSF algorithm equals  $T$ .

After initialization, for each pair of mutually piece-visible vertices of  $V(T)$ , the graph  $H$  contains the corresponding edge, as we added this edge explicitly to  $H$ . Clearly, this invariant is satisfied after every terminal deletion.

Let us now analyze the case when  $\text{REMOVE}_\eta$  removes a vertex  $v$  of degree  $s \leq \eta$ . Denote by  $T_1, \dots, T_s$  the trees which  $T$  decomposes into. We use Lemma 5.18 to compute a set of edges corresponding to portal distances that have minimum weight and connect trees  $T_1, \dots, T_s$ . These edges are added to  $H$  in Step 2.

Assume that before the operation of removing a vertex  $v$  from  $T$ , the tree  $T$  is a minimum spanning tree of  $GD[V(T)]$ . Then, the addition of the reconnecting edges in Step 2 does not modify  $T$ , as we use a tie breaking strategy for edges' weights and assume that out of two edges of the same weight, the lighter is the one that was added earlier (see Section 2). These edges may be used when the edges incident to  $v$  are deleted in Step 3.

We now prove that  $H$  contains a set of edges  $F^v$  of minimum total cost that reconnects trees  $T_1, \dots, T_s$  into a tree  $T'$ . If  $s = 1$ , the claim follows trivially (the set  $F^v$  is empty). Otherwise, consider an edge  $e$  of  $F^v$ . Removing  $e$  from  $T'$  yields two trees  $T'_1$  and  $T'_2$ . We show that  $H$  contains an edge connecting  $T'_1$  with  $T'_2$  of cost equal to  $e$ . If  $d_{GD}(e)$  is equal to the piece distance between

the endpoints of  $e$ , then the edge  $e$  is contained in  $H$ . Otherwise, the length of  $e$  corresponds to portal distance. But, since the set computed by Lemma 5.18 is a MST of a graph defined over portal distances, we have added to  $H$  the edge connecting  $T'_1$  with  $T'_2$  of minimum portal distance. The length of this edge is at most the length of  $e$  (in fact the lengths are equal).

The dynamic MSF algorithm guarantees that the new tree  $T'$  after the operation consists of the connected components of  $T \setminus v$  connected with a set of edges of minimum possible cost with respect to the weights in  $GD$ . Consequently,  $T'$  is a MST of  $GD[V(T')]$ .  $\square$

**Lemma 6.2.** *Let  $\varepsilon > 0$ . Assume we are given a fully dynamic  $\mu$ -approximate vertex-color distance oracle with  $p_{\text{num}}$  portals. For every  $v \in V$ ,  $|\text{portals}(v)| \leq p_{\text{max}}$ ,  $|\text{pieces}(v)| \leq r_{\text{max}}$  and an (expected) size of each piece is bounded by  $r_{\text{max}}^\#$ . Then, there exists a dynamic algorithm that maintains a  $2(1+\varepsilon)\mu$  approximation of a minimum cost Steiner tree under a sequence of terminal deletions. For every terminal deletion, in amortized sense, the algorithm issues  $O(\varepsilon^{-1})$  merge and split operations and uses  $O(\varepsilon^{-1} \log^4 n + \varepsilon^{-1} p_{\text{num}} + \varepsilon^{-1} \log \varepsilon^{-1})$  additional (expected) time. Preprocessing requires  $O(|S|)$  merge and split operations, as well as  $O(|S|(r_{\text{max}}^\# r_{\text{max}} \log^4 n + p_{\text{num}} \log n))$  (expected) time. The algorithm uses  $O(|S|(\varepsilon^{-1} + r_{\text{max}}^\# r_{\text{max}}))$  (expected) space.*

*Proof.* We use the decremental algorithm described in this section which, by Lemma 6.1, follows the decremental scheme. The approximation guarantee follows from Lemma 4.13.

Let us now analyze the efficiency of the algorithm. At the initialization step, we compute the minimum spanning tree of the terminal set. By Lemma 5.19, this requires  $O(|S|(p_{\text{num}} + r_{\text{max}}^\# r_{\text{max}}) \log n)$  (expected) time. Moreover, we add edges between mutually piece-visible terminals. For each vertex  $v$  we compute shortest paths in each element of  $\text{pieces}(v)$  and add respective edges to  $H$ . The computation of shortest paths in each piece takes  $O(|S| r_{\text{max}}^\# r_{\text{max}} \log n)$  (expected) time. Thus, during initialization, we add a total of  $O(|S| r_{\text{max}}^\# r_{\text{max}})$  edges to  $H$  (in expectation), which, by Theorem 2.2, requires  $O(|S| r_{\text{max}}^\# r_{\text{max}} \log^4 n)$  (expected) time.

When a vertex  $v$  is removed from the maintained spanning tree we need to compute the set of reconnecting edges, using Lemma 5.18 in  $O(\eta p_{\text{num}} + \eta \log \eta) = O(\varepsilon^{-1} p_{\text{num}} + \varepsilon^{-1} \log \varepsilon^{-1})$  (expected) time. We obtain a set of  $O(\eta) = O(\varepsilon^{-1})$  edges that we add to  $H$  in  $O(\varepsilon^{-1} \log^4 n)$  time. Moreover, with respect to the operations on the oracle  $\mathbb{D}$ , we merge and activate vertices of  $S$   $O(\eta)$  times per removal.

For the space bound, we need  $O(|S|(\varepsilon^{-1} + r_{\text{max}}^\# r_{\text{max}}))$  (expected) space to store the graph  $H$  and maintain the decremental MSF algorithm. In addition to that, space is used by the vertex-color distance oracle.  $\square$

By using the two oracles developed in Section 3.2, we obtain the following two dynamic algorithms.

**Theorem 6.3.** *Let  $G = (V, E, d_G)$  be a graph,  $n = |V|$ ,  $m = |E|$  and  $\varepsilon > 0$ . Let  $S \subseteq V$  be a dynamic set, subject to vertex removals. Then, after preprocessing in  $O(\sqrt{n}(m + n \log n) + |S|\sqrt{n} \log^4 n)$  expected time, we may maintain a  $(6 + \varepsilon)$ -approximate Steiner tree that spans  $S$ , handling each removal from  $S$  in  $O(\varepsilon^{-1} \sqrt{n} \log n)$  expected amortized time. The algorithm uses  $O(n\sqrt{n} \log n)$  expected space.*

*Proof.* We combine Lemma 6.2 with 3-approximate vertex-color distance oracles of Theorem 5.8 (based on Lemma 5.7). There are  $r_{\text{num}} = n$  pieces, each of expected size  $r_{\text{max}}^\# = O(\sqrt{n})$ . Each vertex is assigned a single piece ( $r_{\text{max}} = 1$ ). For every vertex the number of portals is  $p_{\text{max}} = O(\sqrt{n})$

and so is the total number of portals ( $p_{\text{num}} = O(\sqrt{n})$ ). Every **merge** and **split** operation requires  $T_{\text{merge}} = O(\sqrt{n} \log n)$  time. By plugging this parameters to Lemma 6.2 we obtain an algorithm which processes each removal in time

$$\begin{aligned} & O(\varepsilon^{-1} \log^4 n + \varepsilon^{-1} p_{\text{num}} + \varepsilon^{-1} \log \varepsilon^{-1} + \varepsilon^{-1} T_{\text{merge}}) \\ &= O(\varepsilon^{-1} \log^4 n + \varepsilon^{-1} \sqrt{n} + \varepsilon^{-1} \log \varepsilon^{-1} + \varepsilon^{-1} \sqrt{n} \log n) \\ &= O(\varepsilon^{-1} \sqrt{n} \log n). \end{aligned}$$

The last equality holds under a natural assumption that  $\log \varepsilon^{-1} = O(\sqrt{n})$ ; otherwise a brute-force approach works in  $O(\varepsilon^{-1})$  time.

It takes  $O(\sqrt{n}(m + n \log n))$  expected time to initialize the oracle and

$$\begin{aligned} & O(|S|(T_{\text{merge}} + r_{\text{max}}^{\#} r_{\text{max}} \log^4 n + p_{\text{num}} \log n)) \\ &= O(|S|(\sqrt{n} \log n + \sqrt{n} \log^4 n + \sqrt{n} \log n)) \\ &= O(|S| \sqrt{n} \log^4 n) \end{aligned}$$

expected time to initialize the algorithm. Altogether, the expected initialization time is  $O(\sqrt{n}(m + n \log n) + |S| \sqrt{n} \log^4 n)$ .

The space usage is  $O(|S|(\varepsilon^{-1} + r_{\text{max}}^{\#} r_{\text{max}}) + n \sqrt{n} \log n) = O(n \sqrt{n} \log n)$  expected space (again, we assume  $\log \varepsilon^{-1} = O(\sqrt{n})$ ).

We obtain approximation factor of  $2(1 + \varepsilon)3 = 6(1 + \varepsilon)$ , which can be adjusted to  $6 + \varepsilon$  by using  $\varepsilon' = \varepsilon/6$  instead of  $\varepsilon$ .  $\square$

**Theorem 6.4.** *Let  $G = (V, E, d_G)$  be a planar graph,  $n = |V|$  and  $\varepsilon > 0$ . Denote by  $D$  be the stretch of the metric induced by  $G$ . Let  $S \subseteq V$  be a dynamic set, subject to vertex removals. Then, after preprocessing in  $O(\varepsilon^{-1} n \log^2 n \log D + |S| \varepsilon^{-0.5} \sqrt{n} \log^{2.5} n \log D)$  time we may maintain  $(2 + \varepsilon)$ -approximate Steiner tree that spans  $S$ , handling each removal from  $S$  in  $\tilde{O}(\varepsilon^{-1.5} \sqrt{n} \log D)$  amortized time.*

*Proof.* We combine Lemma 6.2 with  $(1 + \varepsilon)$ -approximate vertex-color distance oracles of Theorem 5.16. There are  $r_{\text{num}} = O(n/\rho \log D)$  pieces, each of size  $r_{\text{max}}^{\#} = O(\rho)$ . For every vertex the number of pieces is  $r_{\text{max}} = O(\log D)$ . The total number of portals is  $p_{\text{num}} = O(\varepsilon^{-1} n/\rho \log D)$ , whereas the number of portals assigned to a vertex is at most  $p_{\text{max}} = O(\varepsilon^{-1} \log n \log D)$ . Every **merge** and **split** operation requires  $T_{\text{merge}} = O(\varepsilon^{-1} n/\rho \log n \log D)$  time.

We plug this parameters to Lemma 6.2 and obtain an algorithm which processes each removal in time

$$\begin{aligned} & O(\varepsilon^{-1} \log^4 n + \varepsilon^{-1} p_{\text{num}} + \varepsilon^{-1} \log \varepsilon^{-1} + \varepsilon^{-1} T_{\text{merge}}) \\ &= O(\varepsilon^{-1} \log^4 n + \varepsilon^{-2} n/\rho \log D + \varepsilon^{-1} \log \varepsilon^{-1} + \varepsilon^{-2} n/\rho \log n \log D) \\ &= O(\varepsilon^{-2} n/\rho \log n \log D). \end{aligned}$$

It takes  $O(\varepsilon^{-1} n \log^2 n \log D)$  time to initialize the oracle and

$$\begin{aligned} & O(|S|(T_{\text{merge}} + r_{\text{max}}^{\#} r_{\text{max}} \log^4 n + p_{\text{num}} \log n)) \\ &= O(|S|(\varepsilon^{-1} n/\rho \log n \log D + \rho \log D \log^4 n + \varepsilon^{-1} n/\rho \log n \log D)) \\ &= O(|S|(\varepsilon^{-1} n/\rho \log n \log D + \rho \log D \log^4 n)) \end{aligned}$$

time to initialize the algorithm. We balance the two terms from initialization time and get  $\rho = \sqrt{n}\varepsilon^{-0.5}/\log^{1.5} n$ . Thus, the initialization time becomes  $O(\varepsilon^{-1}n \log^2 n \log D + |S|\varepsilon^{-0.5}\sqrt{n} \log^{2.5} n \log D)$ .

The amortized time for removing one terminal is  $O(\varepsilon^{-2}n/\rho \log n \log D) = O(\varepsilon^{-1.5}\sqrt{n} \log^{2.5} n \log D)$ .

We obtain approximation factor of  $2(1 + \varepsilon)^2$ , which can be adjusted to  $2 + \varepsilon$  by using a value of  $\varepsilon$  which is a constant factor smaller.  $\square$

## 6.2 Incremental algorithm

The goal is to maintain an approximate Steiner tree on a set of terminals, subject to terminal additions. Fix a constant  $\tau > 0$  and set  $\varsigma = \tau/2$ . We assume that the input graph  $G$  may *only* be accessed by incremental vertex-color distance oracles (see Section 3.2). In particular, improving the oracles would immediately yield a better algorithm for our problem. We denote the approximation ratio of the oracles by  $\mu$  and assume  $\mu = O(1)$ . We also assume that all the oracles yield the same  $\mu$ -near metric space  $GD = (V, \binom{V}{2}, d_{GD})$  that approximates the metric closure of  $G$  (see Corollary 5.17). In particular, the tree that we maintain has edge lengths corresponding to the lengths in  $GD$ . Moreover, we assume that the distances returned by the oracles are discrete, meaning that every distance is a power of  $1 + \varsigma$ . If this is not the case, we may round up every distance returned by the oracle. As a result, the distances will correspond to a  $\mu(1 + \varsigma)$ -near metric space. However, for simplicity of presentation, let us assume that we work with a  $\mu$ -near metric space. We will include the additional  $1 + \varsigma$  factor in the final analysis.

Our first goal is to maintain a good approximation of the minimum spanning tree of  $S_i$  in graph  $GD$ . In order to achieve that, we use the incremental online algorithm from Section 4.4. The parameters  $\tau$  and  $\varsigma$  that we have chosen correspond to the parameters used in the description of the scheme. We denote the terminals that are added by  $t_1, t_2, t_3, \dots$ . Let  $S_i = \{t_1, \dots, t_i\}$  and let  $T_i$  denote the tree maintained by the algorithm after adding  $i$  terminals.

Assume that all edge lengths of  $GD$  belong to the interval  $[1, D]$ . Note that  $D$  is at most  $\mu$  times greater than the stretch of the metric induced by  $G$ . We define the *level* of an edge  $uv$ , as  $\text{lvl}(uv) := \log_{1+\varsigma} d_{GD}(uv)$ . Similarly, the level of a tree  $T$  is  $\text{lvl}(T) := \lfloor \log_{1+\varsigma} d_{GD}(T) \rfloor$ . Note that since the distances are powers of  $1 + \varsigma$ , the levels of edges are integers. The algorithm represents the current tree  $T_i$  by maintaining  $h = \lfloor \log_{1+\varsigma} D \rfloor$  layers  $L_1, \dots, L_h$ . Layer  $L_j$  contains edges of the current tree  $T_i$  whose level is at most  $j$ . In particular, layer  $L_h$  contains all edges of  $T_i$ . We also define  $\text{lvl}_\perp(T) = \lfloor \log_{1+\tau} \frac{\tau d_{GD}(T)}{8\mu^2(1+\tau)^3 n} \rfloor$ . Intuitively, edges of level at most  $\text{lvl}_\perp(T)$  are so short that they may sum up to at most  $\tau/2$  fraction of the weight of  $T$ .

For each layer  $L_j$ , we maintain an incremental vertex-color distance oracle  $\mathbb{D}_j$  on the graph  $G$ . The partition of  $V$  into colors corresponds to the connected components of a graph  $(V, L_j)$ . A color is active if and only if its elements belong to  $S_i$ .

We now describe the process of adding a terminal  $t_i$ . The pseudocode is given in Algorithm 1. It consists of two stages. First, we find the shortest edge in  $GD$  connecting  $t_i$  to any of  $t_1, \dots, t_{i-1}$  and add this edge to  $T_{i-1}$  to obtain tree  $T_i$ . Then, we apply a sequence of  $\varsigma$ -efficient replacements  $(e, e')$  to  $T_i$  in order to decrease its weight. According to the incremental scheme of Section 4.4 we only need to consider replacements in which  $e$  is incident to the newly added terminal  $t_i$ .

In order to find the replacements, we use the vertex-color distance oracles. Fix a layer number  $j$  and assume that the colors of the oracle  $\mathbb{D}_j$  reflect the layer structure, including the newly added terminal  $t_i$ . We want to find a replacement pair  $(e, e')$  such that the  $\text{lvl}(e') > j$  and  $\text{lvl}(e) \leq j$ . Denote by  $C$  the connected component of the graph  $(V, L_j)$  that contains  $t_i$  (observe that the vertex set of  $C$  is exactly the set of vertices with the same color as  $t_i$  in  $\mathbb{D}_j$ ). By definition of  $C$ , it consists



---

**Algorithm 1** Adding a new terminal  $t_i$ 

---

```
1: procedure FINDREPLACEMENTS( $j$ )
2:   while true do
3:      $t' :=$  active vertex in  $\mathbb{D}_j$  which is second nearest to  $t_i$ 
4:     if  $\text{lvl}(t_i t') > j$  then
5:       break
6:      $e' :=$  the heaviest edge in  $T_i$  on the path from  $t_i$  to  $t'$ 
7:     replace  $e'$  with  $t_i t'$  in  $T_i$ 
8:     for  $k = 1, \dots, h$  do
9:       if  $\text{lvl}(e') \leq k$  then
10:        replace  $e'$  with  $t_i t'$  in  $L_k$ 
11:       else if  $\text{lvl}(t_i t') \leq k$  then
12:        add  $t_i t'$  to layer  $L_k$ 
13:        merge colors of  $t_i$  and  $t'$  in  $\mathbb{D}_k$ 
14: procedure ADDTERMINAL( $t_i$ )
15:    $t' :=$  the active vertex in  $\mathbb{D}_h$  which is nearest to  $t_i$ 
16:    $T_i := T_{i-1} \cup \{t_i t'\}$ 
17:   for  $j = 1, \dots, h$  do
18:     activate color of  $t_i$  in  $\mathbb{D}_j$ 
19:     if  $\text{lvl}(t_i t') \leq j$  then
20:       add  $t_i t'$  to  $L_j$ 
21:       merge colors of  $t_i$  and  $t'$  in  $\mathbb{D}_j$ 
22:    $T_M :=$  maximum weight tree among  $T_0, \dots, T_{i-1}$ 
23:   for  $j = \text{lvl}_\perp(T_M), \dots, \text{lvl}(T_M)$  do
24:     FINDREPLACEMENTS( $j$ )
```

---

of edges of level at most  $j$  and the path in  $T_i$  from  $t_i$  to every  $t \notin C$  contains an edge of level  $> j$ . We find the vertex  $t \notin C$  which is the nearest to  $t_i$  by querying  $\mathbb{D}_j$ . Since  $t_i$  has color  $C$ , we issue a  $\text{nearest}(t_i, 2)$  query to find the second nearest color from  $t_i$ . Assume that we find a vertex  $t'$ . If  $\text{lvl}(t_i t') \leq j$ , we have found a replacement pair. Let  $e'$  be the heaviest edge on the path from  $t_i$  to  $t'$  in  $T_i$  and  $e = t_i t'$ . Clearly,  $(e, e')$  is a replacement pair, and since the weight of  $e$  is lower than the weight of  $e'$  and the weights can differ at least by a factor of  $1 + \varsigma$ , this replacement is  $\varsigma$ -efficient. Note that we choose  $e'$  to be the heaviest edge of the path from  $t_i$  to  $t'$  in  $T_i$ , as required by the incremental scheme.

Our algorithm iterates through layers in increasing order, and for each layer it repeatedly looks for replacements using the procedure described above. As we later show, in this way we will detect all replacements that we need. In particular, applying a replacement on a higher level may not create new  $(\tau/2, 2\mu(1 + \tau)^2)$ -good replacements on lower levels.

Let  $T_M$  be the tree with maximum weight among  $T_1, \dots, T_{i-1}$ . For efficiency, we only iterate through layers  $j = \text{lvl}_\perp(T_M), \dots, \text{lvl}(T_M)$  to perform replacements.<sup>6</sup> After this procedure, some  $\varsigma$ -efficient replacements may still be possible in  $T_i$ , but they would not be  $(\tau/2)/(2\mu(1 + \tau)^2)$ -heavy.

Every time we change the current tree, we update all layers accordingly. In case we add a new edge, this boils down to a single merge operation (in each layer that is supposed to contain the edge). A more difficult case is when we perform a replacement, as then we *remove* an edge that is replaced. Assume that we perform a replacement  $(e, e')$  and we need to update a layer  $L_j$  that contains the edge  $e'$ . From the algorithm we know that  $e'$  is the heaviest edge on the path connecting the endpoints of  $e$ , so all edges of this path are contained in  $L_j$  and the endpoints of  $e$  have the same color. Right after we remove  $e'$  from  $T_i$ , we add  $e$ . Since the level of  $e$  is lower than the level of  $e'$ , the edge  $e$  is added to layer  $L_j$ , which means that in the end the colors in  $L_j$  do not change at all.

We first show that one call to `FINDREPLACEMENTS` correctly finds the replacements.

**Lemma 6.5.** *After a call to `FINDREPLACEMENTS(j)`, there are no possible  $\varsigma$ -efficient replacements  $(e, e')$ , such that  $\text{lvl}(e') > j$ ,  $\text{lvl}(e) \leq j$  and  $e$  is incident to  $t_i$ . Every replacement made by `FINDREPLACEMENTS` is  $\varsigma$ -efficient. Every time a replacement  $(e, e')$  is made, the edge  $e'$  is the heaviest edge on the path between the endpoints of  $e$ .*

*Proof.* The second and third claims of the lemma follow directly from the discussion above. It remains to prove the first one. Note that once we identify a replacement pair, the loop in lines 8-13 correctly updates the data structures.

Assume that a desired replacement exists and  $e = t_i t''$ . The path from  $t_i$  to  $t''$  in the current spanning tree goes through  $e'$ . Consider the vertex-color distance oracle  $\mathbb{D}_j$ . Since its colors correspond to connected components of the subgraph consisting of edges of level at most  $j$  and  $\text{lvl}(e') > j$ , the color of  $t''$  is different from the color of  $t_i$ . Thus, the vertex  $t'$  found in 3rd line of the algorithm, is at distance at most  $d(t_i t'') \leq d(t_i t')$  from  $t_i$  and  $\text{lvl}(t_i t') \leq j$ , so the **while** loop does not terminate.  $\square$

We now show, that considering layers in the increasing order is correct. In particular, performing a replacement using one layer may not create replacements at lower levels.

---

<sup>6</sup>A simple solution would be to iterate through all layers, or through a range of layers defined in terms of  $T_{i-1}$ , but we use a slightly more sophisticated approach in preparation for more efficient algorithm that we describe later. The key property here is that the weight of  $T_M$  may be only slightly larger than the weight of  $T_i$ , but at the same time the weight of  $T_M$  may only increase after adding every terminal.

**Lemma 6.6.** *After  $\text{FINDREPLACEMENTS}(j)$  is called with  $j = a, \dots, b$ , there does not exist a  $\varsigma$ -efficient replacement  $(e, e')$ , such that  $\text{lvl}(e) \leq b$ ,  $\text{lvl}(e') > a$  and  $e$  is incident to  $t_i$ . Moreover, every time the algorithm identifies a replacement pair  $(e, e')$  in the iteration  $j$ , we either have  $\text{lvl}(e) = j$ , or  $j = a$  and  $\text{lvl}(e) \leq a$ .*

*Proof.* The proof is by induction on  $b - a$ . The base case  $b = a$  is provided by Lemma 6.5; observe that also any replacement pair  $(e, e')$  identified for  $j = a$  satisfies  $\text{lvl}(e) \leq a$ . Consider now an inductive step for  $b > a$ .

Throughout the proof we only consider replacements in which the replacement edge is incident to  $e$ . By induction hypothesis, we know that in the beginning of last iteration for  $j = b$  there are no  $\varsigma$ -efficient replacement  $(e, e')$ , such that  $\text{lvl}(e) \leq b - 1$  and  $\text{lvl}(e') > a$ . Moreover, by Lemma 6.5, when the iteration is complete, there are no  $\varsigma$ -efficient replacement  $(e, e')$ , such that  $\text{lvl}(e) \leq b$  and  $\text{lvl}(e') > b$ . We combine this two facts to show the induction step, but we need to prove that performing replacements may not introduce any new replacements  $(f, f')$  satisfying  $\text{lvl}(f) < b$  and  $\text{lvl}(f') > a$ .

Consider the first replacement  $(e, e')$  that we perform in the iteration  $j = b$ . Then, from the induction hypothesis, it follows immediately that  $\text{lvl}(e) = b$ . We show that performing the replacement may not introduce any new replacements  $(f, f')$ , such that  $\text{lvl}(f) < b$  and  $\text{lvl}(f') > a$ . Indeed, observe that such a replacement can be identified by querying  $\mathbb{D}_{\text{lvl}(f')}$ . But then the vertex-color distance oracles  $\mathbb{D}_{b'}$ , for  $b' < b$ , are not modified, so no replacement pair  $(f, f')$  such that  $\text{lvl}(f) < b$  and  $\text{lvl}(f') > a$  can be created.

Consequently, the next replacement  $(e, e')$  performed in the iteration  $j = b$  also satisfies  $\text{lvl}(e) = b$ , and the argument from the previous paragraph applies again. Consequently, as we perform replacements, we do not introduce any replacement pair  $(f, f')$  such that  $\text{lvl}(f) < b$  and  $\text{lvl}(f') > a$ . This concludes the proof of the lemma.  $\square$

**Lemma 6.7.** *The maintained tree  $T_k$  is a  $2(1 + \tau)\mu$ -approximation of a minimum Steiner tree in  $G$  that spans  $S_k$ .*

*Proof.* The proof proceeds by induction on the number of terminals  $i$ . Our goal is to use Lemma 4.15, so we have to show that the algorithm presented here implements the incremental scheme from Section 4.4.

In order to show that we follow the incremental scheme from Section 4.4 we need to prove that there are no  $(\tau/2, 2\mu(1 + \tau)^2)$ -good replacements  $(e, e')$  with  $e$  incident to  $t_i$ . By Lemma 6.6, after all calls to  $\text{FINDREPLACEMENTS}$ , for every such replacement, we either have  $\text{lvl}(e) > \text{lvl}(T_M)$  or  $\text{lvl}(e') \leq \text{lvl}_\perp(T_M)$ . Since the edges that we replace either belong to  $T_{i-1}$  or have been added as a replacement for an edge in  $T_{i-1}$ , and  $d_{GD}(T_M) \geq d_{GD}(T_i)$  the first case may not hold. Let us then consider the second case. We show that the condition implies that  $(e, e')$  is not  $(\tau/2)/(2\mu(1 + \tau)^2)$ -heavy. We have

$$\begin{aligned} d(e') &\leq \frac{\tau d(T_M)}{8(1 + \tau)^3 \mu^2 n} && \text{from } \text{lvl}(e') \leq \text{lvl}_\perp(T_M) \\ &\leq \frac{\tau 2(1 + \tau)\mu d(T_i)}{8(1 + \tau)^3 \mu^2 n} && \text{by Lemma 4.14, } d(T_M) \leq 2(1 + \tau)\mu d(T_i) \\ &\leq \frac{\tau d(T_i)}{4(1 + \tau)^2 \mu n}. \end{aligned}$$

This shows that the update procedure correctly implements the incremental scheme from Section 4.4, so, by Lemma 4.15, it computes a  $2(1 + \tau)\mu$ -approximation of a minimum Steiner tree in  $G$  that spans  $S_k$ .  $\square$

The above algorithm can be speeded up by reducing the number of layers, so that it does not depend on the stretch of the metric induced by the input graph. Observe that although we maintain  $O(\log_{1+\varsigma} D)$  layers, we only use layers  $\text{lv}1_{\perp}(T_M), \dots, \text{lv}1(T_M)$  to find the replacements. Thus, we may modify our algorithm so that it only maintains the layers that are necessary. We will maintain  $h$  layers  $L_1, \dots, L_h$  that correspond to levels  $\text{lv}1_{\perp}(T_M), \dots, \text{lv}1_{\perp}(T_M) + h - 1$ . The value  $h$  is chosen such that  $\text{lv}1_{\perp}(T_M) + h - 1$  is at least  $\text{lv}1(T_M)$ , regardless of the weight of  $T_M$ . Thus, for some tree  $T$ ,

$$h = \lceil \log_{1+\varsigma} d(T) \rceil - \left\lfloor \log_{1+\varsigma} \frac{\tau d(T)}{8\mu^2(1+\tau)^3 n} \right\rfloor = O(\log_{1+\varsigma}(\mu^2(1+\tau)^3 n/\tau)) = O(\tau^{-1} \log(n/\tau))$$

Since the weight of  $T_M$  may increase over time, the layers will be *dynamic*, that is, the level corresponding to each layer will also increase.

The algorithm may be easily adapted to handle the dynamic layers. Once a new tree  $T_i$  has been computed and the weight of  $T_M$  increases, the layers may correspond to higher levels and may need to be updated by adding some edges to them. However, there are at most  $n - 1$  of these additions for every layer.

Let us now analyze the efficiency of the algorithm. Note that while the layers are used for the description of the algorithm, they may be omitted in the implementation, as they are never read. We first bound the running time of one call to `FINDREPLACEMENTS`.

**Lemma 6.8.** *Assume that a call to `FINDREPLACEMENTS` performs  $r$  replacements. Then, it issues  $r + 1$  `nearest` queries and at most  $r$  `merge` operations on every level. Moreover, it uses  $O(r \log n)$  additional time.*

*Proof.* We issue a `nearest` query to find replacements.  $r$  of these queries identify new replacements and the last one detects that there are no more replacements to be made. After a replacement is detected, we may execute a `merge` operation on every level. Moreover, after a replacement is detected, we find the heaviest edge on some path in the tree (see line 6 of Algorithm 1). To do that efficiently, we maintain the trees  $T_i$  as top trees [3]. This allows us to perform this operation in  $O(\log n)$  time.  $\square$

**Lemma 6.9.** *Let  $0 < \tau = O(1)$ . Given  $h = O(\tau^{-1} \log(n/\tau))$  instances of an incremental  $\mu$ -approximate vertex-color distance oracle ( $\mu = O(1)$ ), there exists a dynamic algorithm that maintains a  $2\mu(1 + \tau)$ -approximation of a minimum-cost Steiner tree under a sequence of terminal insertions. For every terminal addition, in amortized sense, the algorithm executes  $O(\tau^{-1} \log(n/\tau))$  `activate`, `merge` and `nearest` operations. The space usage is dominated by the space usage of the oracles.*

*Proof.* By Lemma 6.7, the incremental algorithm implements the incremental scheme. As a result, we may use Lemma 4.17 to bound the running time.

Assume that we have added  $k$  terminals. First, we need to bound the number of replacements. Since we round the distances up to a power of  $(1 + \varsigma)$ , we operate on a  $\mu(1 + \varsigma)$ -near metric space. Our algorithm implements the incremental scheme with parameters  $\varsigma = \tau/2$ . Thus, by Lemma 4.17,

adding  $k$  terminals may trigger  $O(k\tau^{-1}\log((1+\tau)\mu)) = O(k\tau^{-1})$  replacements. In particular, this shows that the update procedure terminates.

Let us now calculate the number of operations performed on the vertex-color distance oracles. We use one copy of the oracle for each layer. In the course of adding  $k$  terminals, in each of the  $h$  layers at most  $k$  vertices may be activated and colors may be merged  $k-1$  times. Note that this includes the merging of colors caused by maintaining dynamic layers. For every terminal insertion we also issue one nearest query to find the first edge connecting the new terminal with the current tree. Since a single terminal insertion triggers  $O(\tau^{-1})$  replacements, by Lemma 6.8, we issue  $O(h + \tau^{-1}) = O(\tau^{-1}\log(n/\tau))$  nearest queries, in amortized sense. We also spend  $O(\tau^{-1}\log n)$  additional time, but this time is dominated by the oracle operations.

Apart from the oracles, the algorithm uses  $O(n)$  space to represent the trees. Thus, the space usage is dominated by the oracles.  $\square$

By using the two oracles developed in Section 3.2, we obtain the following two dynamic algorithms.

**Theorem 6.10.** *Let  $G = (V, E, d_G)$  be a graph,  $n = |V|$  and  $\varepsilon > 0$ . Let  $S \subseteq V$  be a dynamic set, subject to vertex insertions (initially  $S = \emptyset$ ). Then, after preprocessing in  $O(\varepsilon^{-1}\sqrt{n}(m + n\log(n/\varepsilon)))$  expected time we may maintain  $(6+\varepsilon)$ -approximate Steiner tree that spans  $S$ , handling each insertion to  $S$  in  $\tilde{O}(\varepsilon^{-1}\sqrt{n})$  amortized time. The algorithm uses  $O(\varepsilon^{-1}n\sqrt{n}\log n\log(n/\varepsilon))$  space.*

*Proof.* We combine the 3-approximate vertex-color distance oracle from Theorem 5.8 (based on Lemma 5.7) with Lemma 6.9. Thus  $\mu = 3 = O(1)$ . We set  $\tau = \varepsilon/3$ .

We first build  $O(\tau^{-1}\log(n/\tau))$  copies of the oracle, by first building a single one and then copying it. Building the first oracle requires  $O(\sqrt{n}(m + n\log n))$  expected time. Since the space usage is  $O(n\sqrt{n}\log n)$ , the copying requires  $O(\tau^{-1}n\sqrt{n}\log n\log(n/\tau))$  time, and so is the space usage.

The oracle handles `activate` and `merge` operations in  $O(\sqrt{n}\log n)$  time, whereas `nearest` queries are answered in  $O(\sqrt{n})$  time. Since we perform  $O(\tau^{-1}\log(n/\tau))$  of each of these operations, they require  $O(\tau^{-1}\sqrt{n}\log n\log(n/\tau)) = \tilde{O}(\varepsilon^{-1}\sqrt{n})$  amortized time per terminal insertion.

We obtain an approximation factor of  $2(1+\varepsilon)3 = 6(1+\varepsilon)$ , which may be reduced to  $6+\varepsilon$  by manipulating the value of  $\varepsilon$ .  $\square$

**Theorem 6.11.** *Let  $G = (V, E, d_G)$  be planar a graph,  $n = |V|$  and  $D$  be the stretch of the metric induced by  $G$ . Fix  $\varepsilon > 0$ . Let  $S \subseteq V$  be a dynamic set, subject to vertex insertions (initially  $S = \emptyset$ ). Then, after preprocessing in  $O(\varepsilon^{-2}n\log n\log(n/\varepsilon)\log D)$  time we may maintain  $(2+\varepsilon)$ -approximate Steiner tree that spans  $S$ , handling every insertion to  $S$  in  $O(\varepsilon^{-2}\log^2 n\log(n/\varepsilon)\log D)$  expected amortized time or  $O(\varepsilon^{-2}\log^2 n\log\log n\log(n/\varepsilon)\log D)$  deterministic amortized time. The algorithm uses  $O(\varepsilon^{-2}n\log n\log(n/\varepsilon)\log D)$  space.*

*Proof.* We combine the  $(1+\varepsilon)$ -approximate vertex-color distance oracle of Theorem 5.14 with Lemma 6.9. We set  $\tau = \varepsilon$ .

We first build  $O(\tau^{-1}\log(n/\tau))$  copies of the oracle, by first building a single one and then copying it. Building the first oracle requires  $O(\tau^{-1}n\log^2 n\log D)$  time. Since the space usage is  $O(\tau^{-1}n\log n\log D)$ , the copying requires  $O(\tau^{-2}n\log n\log(n/\tau)\log D) = O(\varepsilon^{-2}n\log n\log(n/\varepsilon)\log D)$  time, and so is the space usage. The copying also dominates the preprocessing time.

The oracle handles **activate** and **merge** operations in  $O(\varepsilon^{-1} \log^2 n \log D)$  expected amortized time, whereas **nearest** queries are answered in  $O(\varepsilon^{-1} \log n \log D)$  time (and additional  $\log \log n$  factor is needed for deterministic algorithms). Since we perform  $O(\tau^{-1} \log(n/\tau))$  of each of these operations, they require  $O(\varepsilon^{-2} \log^2 n \log(n/\tau) \log D)$  expected amortized time (with an additional  $\log \log n$  factor for a deterministic algorithm).

We obtain approximation ratio of  $2(1 + \varepsilon)^2$ . It is easy to see that by manipulating the value of  $\varepsilon$ , we may obtain  $(2 + \varepsilon)$ -approximation.  $\square$

### 6.3 Fully dynamic algorithm

In this section we merge the ideas of the decremental and incremental algorithms to obtain a fully dynamic algorithm, which supports both terminal additions and deletions. The algorithm simply maintains the invariants of both the incremental and decremental algorithms. Let  $\varepsilon > 0$  and  $\varsigma = \tau > 0$ . We set  $\eta = 1 + \lceil \varepsilon^{-1} \rceil$ . We implement the fully dynamic scheme of Section 4.5 with the parameters  $\varepsilon$ ,  $\varsigma$  and  $\tau$ . Recall that  $\varepsilon$  (and  $\eta$ ) controls the degree threshold for nonterminals that are deleted, and  $\varsigma$  defines the efficiency of the replacements that we make. We maintain the invariant that the tree does not contain any  $\tau$ -efficient replacements. Since,  $\varsigma = \tau$ , in the following we only use  $\tau$ .

Similarly to the decremental algorithm, we work with a graph  $G = (V, E, d_G)$  that is accessed with approximate vertex-color distance oracles, that yield a  $\mu$ -near metric space  $GD$ , which approximates  $G$  (see Corollary 5.17). We assume that  $\mu = O(1)$  and that every edge length is  $GD$  has length being a power of  $(1 + \tau)$ . This can be achieved by rounding up the distances returned by the oracle. As a result, the graph  $GD$  yielded by the oracle becomes a  $\mu(1 + \tau)$ -near metric approximating  $G$ , but in the description we assume the approximation factor of  $\mu$ . The additional  $1 + \tau$  factor will be included in the final analysis.

As in the decremental algorithm, we do not maintain the tree explicitly, but instead we use the dynamic MSF algorithm (see Theorem 2.2) on top of a carefully chosen graph  $H$  with vertex set  $V$ . Again, each connected component of  $H$  is an isolated vertex, with the exception of one component, denoted  $H_0$ , that contains all terminals. We declare that the currently maintained Steiner tree,  $T$ , is the tree maintained by the dynamic MSF algorithm in the component  $H_0$ .

We define the level of an edge  $uv$  to be  $\text{lvl}(uv) := \log_{1+\tau} d_{GD}(uv)$ . Our algorithm maintains a tree  $T$  which spans the set of terminals and nonterminals of degree more than  $\eta$ . The tree  $T$  will not admit  $\tau$ -efficient replacements, but since the distances are powers of  $1 + \tau$ , this means that  $T$  does not admit any efficient replacements. Thus, by Lemma 4.8, it is the MST of its vertex set.

Denote by  $D$  the stretch of the metric induced by  $G$ . Thus, the longest edge in  $GD$  is at most  $\mu \cdot D = O(D)$  times longer than the shortest one. In the algorithm we maintain  $h = \lfloor \log_{1+\tau} D \rfloor = O(\tau^{-1} \log D)$  fully dynamic vertex-color distance oracles  $\mathbb{D}_1, \dots, \mathbb{D}_h$ .  $\mathbb{D}_h$  is the counterpart of the oracle from the decremental algorithm. The colors of  $\mathbb{D}_h$  correspond to connected components of  $H$ . The only active color is the one corresponding to component  $H_0$ . The tree associated with this color is equal to the tree  $T$ . The oracle  $\mathbb{D}_i$  is obtained from  $\mathbb{D}_h$  by splitting along edges of level more than  $i$ . In other words, we consider the forest consisting of edges of  $T$  of level at most  $i$  and for each tree  $F_j$  in this forest,  $\mathbb{D}_i$  contains an active color associated with the tree  $F_j$ . We update the oracles in order to maintain those properties.

Note that we may assume that all oracles have the same structure, i.e., the sets  $\text{portals}(v)$  and  $\text{pieces}(v)$  are the same in each oracle. This can be achieved by constructing one oracle and then

copying it to create all the others. Thus, referring to piece or portal distance between vertices is unambiguous.

As in the decremental algorithm of Section 6.1, we maintain an invariant that for every two vertices  $u, w \in V(T)$ , if  $u$  is piece-visible from  $w$ , the graph  $H$  contains an edge connecting them, whose length is equal to the piece distance.

With respect to the vertex-color distance oracles, note that each modification (edge addition, removal or replacement) in the tree  $T$  requires us to perform  $O(1)$  operations on each oracle  $\mathbb{D}_i$ .

**Vertex removal.** To use the fully dynamic scheme of Section 4.5, we first need to describe how the algorithm behaves if the procedure  $\text{REMOVE}_\eta$  requires us to remove a vertex  $v$  of degree at most  $\eta$  from the tree  $T$ . We resolve this situation in exactly the same manner as in the decremental algorithm of Section 6.1. For sake of completeness, let us recall the steps.

1. Split the active color  $V(T)$  in  $\mathbb{D}_h$  into  $s + 1$  pieces being  $\{v\}$  and the  $s$  connected components of  $T \setminus \{v\}$ , and then deactivate  $\{v\}$ . Update all oracles  $\mathbb{D}_j$  accordingly.
2. Compute the set of reconnecting edges that correspond to portal distances (using Lemma 5.18) and add these edges to  $H$ .
3. Remove all edges incident to  $v$  from  $H$ .
4. At this point, the dynamic MSF algorithm may update the maintained tree. Update the vertex-color distance oracles  $\mathbb{D}_i$ , so that they reflect these changes of the MSF of  $H$ .

**Vertex addition.** A second ingredient needed to implement the fully dynamic scheme is the behavior upon addition of a new terminal  $v$ . If  $v \in V(H_0)$ , we simply mark it as a terminal and finish.

Let us now assume that we are adding a new vertex  $v$  to the tree. We use a similar method as in the incremental scheme, but here we need to assure that after adding  $v$ , the tree does not admit any  $\tau$ -efficient replacements. We use the following algorithm. Note after every change to  $T$  done by the dynamic MSF algorithm we update the nearest oracles accordingly.

1. Connect  $v$  to the tree with the shortest possible edge in  $GD$ . Add this edge to  $H$ .
2. Iterate through levels in increasing order, and for each level find and apply all occurring  $\tau$ -efficient replacements. This is done with the procedure  $\text{FINDREPLACEMENTS}$  from the (see Algorithm 1, Section 6.2). Each time a replacement  $(e, e')$  is found, simply add edge  $e$  to  $H$ . The edge  $e'$  may safely remain in  $H$ .
3. For every vertex  $w$ , whose degree became at most  $\eta$  due to replacements, call  $\text{REMOVE}_\eta(w)$ .
4. Add edges corresponding to piece distances between  $v$  and vertices of  $T$ .

We claim that the dynamic MSF algorithm modifies the tree  $T$  exactly the way the fully dynamic scheme requires it to do.

**Lemma 6.12.** *The tree  $T$  implements the fully dynamic scheme of Section 4.5.*

*Proof.* The correctness of the implementation of the procedure  $\text{REMOVE}_\eta$  follows from the same arguments as in the decremental algorithm. Indeed, our algorithm maintains the same invariants about graph  $H$  and before the removal the tree  $T$  is the MST of the terminals. Since we use the same procedure for handling deletions, the proof is analogous to the proof of Lemma 6.1.

Let us now focus on a step when a vertex  $v \notin V(H_0)$  is added to the terminal set. By Lemma 6.5, every replacement  $(e, e')$  is  $\tau$ -efficient and  $e'$  is the friend of  $e$  with maximum possible cost. Since we iterate through all levels, by Lemma 6.6, we have that after handling the insertion there are no  $\tau$ -efficient replacements  $(e, e')$  with  $e$  incident to  $v$ . Now, by Lemma 4.18, if there are no  $\tau$ -efficient replacements  $(e, e')$  with  $e$  incident to  $v$ , there are no  $\tau$ -efficient replacements at all. In particular, the addition of edges incident to  $v$  corresponding to pieces distances does not cause any changes to  $T$ .

Finally, note that as the weights of the edges in the dynamic MSF algorithm equal their levels in  $GD$ , and we break ties by timestamps, if an addition of a new edge results in a replacement in  $T$ , such a replacement is  $\tau$ -efficient. This concludes the proof of the lemma.  $\square$

We are now ready to state the main result of this section.

**Lemma 6.13.** *Let  $G = (V, E, d_G)$  be a graph,  $\varepsilon > 0$ . Denote by  $D$  the stretch of the metric induced by  $G$ . Assume there exists a fully dynamic  $\mu$ -approximate vertex-color distance oracles, where  $\mu = O(1)$ . Let  $p_{\text{num}}$  be the (expected) total number of portals in the oracle. Moreover, assume that every vertex is assigned at most  $r_{\text{max}}$  pieces (in expectation), the (expected) size of each piece is bounded by  $r_{\text{max}}^\#$ , and the (expected) total number of pieces is  $r_{\text{tot}}$ .*

*Then, there exists a dynamic algorithm that maintains a  $2(1 + \varepsilon)^2 \mu$  approximation of a minimum cost Steiner tree under a sequence of terminal additions and deletions (initially, the terminal set is empty). For every terminal addition, in amortized sense, the algorithm executes  $O(\varepsilon^{-2} \log^2 D)$  **split** and **merge** and issues  $O(\varepsilon^{-1} \log D)$  **nearest** queries. Moreover, it uses  $O(\varepsilon^{-1} p_{\text{num}} + \varepsilon^{-2} \log n \log D + (r_{\text{max}} r_{\text{max}}^\# + \varepsilon^{-1}) \log^4 n)$  additional (expected) time.*

*The initialization time is dominated by the time needed to initialize  $\Theta(\varepsilon^{-1} \log D)$  fully dynamic vertex-color distance oracles. Similarly, the space usage is dominated by the oracles.*

*Proof.* Set  $\tau = \varepsilon$ . By Lemma 6.12, our algorithm implements the fully dynamic scheme. Thus, by Lemma 4.18, it maintains a  $2(1 + \varepsilon)^2 \mu$  approximation of a minimum cost Steiner tree.

The algorithm starts by constructing  $\Theta(\tau^{-1} \log D) = O(\varepsilon^{-1} \log D)$  oracles  $\mathbb{D}_i$ . Initially, the graph  $H$  is empty.

After each removal of a vertex in  $T$  we apply Lemma 5.18, which requires  $O(\varepsilon^{-1} p_{\text{num}} + \varepsilon^{-1} \log \varepsilon^{-1})$  time. Moreover, we issue  $O(\tau^{-1})$  **merge** and **split** operations in each of the  $O(\tau^{-1} \log D)$  oracles, that is, a total of  $O(\tau^{-2} \log D) = O(\varepsilon^{-2} \log D)$  operations.

Now consider a vertex addition. By Lemma 4.19, an addition may trigger  $O(\tau^{-1} \log D) = O(\varepsilon^{-1} \log D)$  replacements. When a new vertex is added we call  $\text{FINDREPLACEMENTS}$   $O(\tau^{-1} \log D)$  times, and each call issues one **nearest** query, regardless of the number of replacements. Moreover, for every replacement, all calls to  $\text{FINDREPLACEMENTS}$  issue an additional **nearest** query and  $O(\tau^{-1} \log D)$  **merge** operations. It also uses  $O(\tau^{-1} \log n)$  additional time. This gives a total of  $O(\tau^{-2} \log^2 D) = O(\varepsilon^{-2} \log^2 D)$  **merge** and **split** operations per insertion,  $O(\tau^{-1} \log D) = O(\varepsilon^{-1} \log D)$  **nearest** queries and  $O(\tau^{-2} \log n \log D) = O(\varepsilon^{-2} \log n \log D)$  additional time.

Whenever we add or remove a vertex from  $T$ , we add or remove all edges corresponding to piece distances. There are at most  $O(r_{\text{max}} r_{\text{max}}^\#)$  such edges for every vertex (in expectation). All these edges can be computed in  $O(r_{\text{max}} r_{\text{max}}^\# \log n)$  using Dijkstra's algorithm.



We are left with bounding the amortized cost of maintaining the dynamic MSF algorithm over the graph  $H$ . Each call to  $\text{REMOVE}_\eta$  that results in a modification of  $T$  as well as each addition of a new terminal adds to  $H$   $O(r_{\max} r_{\max}^\#)$  edges corresponding to piece distances. By Lemma 4.19 and Theorem 2.2 this amortizes to  $O(r_{\max} r_{\max}^\# \log^4 n)$  per operation. Moreover, each removal adds  $O(\varepsilon^{-1})$  new edges to  $H$ , which accounts for  $O(\varepsilon^{-1} \log^4 n)$  amortized time.

Altogether, the amortized cost of one operation is  $O(\varepsilon^{-2} \log n \log D + \varepsilon^{-1} p_{\text{num}} + \varepsilon^{-1} \log \varepsilon^{-1} + r_{\max} r_{\max}^\# \log^4 n + \varepsilon^{-1} \log^4 n) = O(\varepsilon^{-1} p_{\text{num}} + \varepsilon^{-2} \log n \log D + (r_{\max} r_{\max}^\# + \varepsilon^{-1}) \log^4 n)$ .

Apart from the vertex-color distance oracles, the space usage is  $O(n)$ . Thus, the oracles dominate the space usage.  $\square$

By using the two oracles developed in Section 3.2, we obtain the following dynamic algorithms.

**Theorem 6.14.** *Let  $G = (V, E, d_G)$  be a graph,  $n = |V|$ ,  $m = |E|$  and  $\varepsilon > 0$ . Denote by  $D$  the stretch of the metric induced by  $G$ . Let  $S \subseteq V$  be a dynamic set, subject to vertex additions and removals (initially  $S = \emptyset$ ). Then, after preprocessing in  $O(\sqrt{n}(m + \varepsilon^{-1} n \log n \log D))$  expected time, we may maintain a  $(6 + \varepsilon)$ -approximate Steiner tree that spans  $S$ , handling each change to  $S$  in  $\tilde{O}(\varepsilon^{-2} \sqrt{n} \log^2 D)$  expected amortized time. The algorithm uses  $O(n \sqrt{n} \log n \varepsilon^{-1} \log D)$  expected space.*

*Proof.* We combine Lemma 6.13 with vertex-color distance oracles for general graphs of Theorem 5.8 (based on Lemma 5.7). We have that  $\mu = 3 = O(1)$ . The expected total number of portals is  $p_{\text{num}} = \sqrt{n}$ . Every vertex is assigned a single piece  $r_{\max} = 1$  of expected size at most  $r_{\max}^\# = O(\sqrt{n})$ . The total number of pieces is  $r_{\text{tot}} = O(n)$ . Each **merge** and **split** operation requires  $T_{\text{merge}} = O(\sqrt{n} \log n)$  expected time, whereas a **nearest** query can be handled in  $T_{\text{query}} = O(\sqrt{n})$  expected time.

It takes  $O(\sqrt{n}(m + n \log n))$  expected time to initialize one oracle, and, since the expected space usage is  $O(n \sqrt{n} \log n)$ , copying the oracle  $\Theta(\varepsilon^{-1} \log D)$  times requires  $O(n \sqrt{n} \log n \varepsilon^{-1} \log D)$  time. Thus, the expected preprocessing time is  $O(\sqrt{n}(m + n \log n) + n \sqrt{n} \log n \varepsilon^{-1} \log D) = O(\sqrt{n}(m + \varepsilon^{-1} n \log n \log D))$ . The space usage is dominated by the oracles, and amounts to  $O(n \sqrt{n} \log n \varepsilon^{-1} \log D)$ .

Each terminal addition or deletion requires time

$$\begin{aligned} & O(\varepsilon^{-2} \log^2 D T_{\text{merge}} + \varepsilon^{-1} \log D T_{\text{query}} + \varepsilon^{-1} p_{\text{num}} + \varepsilon^{-2} \log n \log D + (r_{\max} r_{\max}^\# + \varepsilon^{-1}) \log^4 n) \\ &= O(\varepsilon^{-2} \log^2 D \sqrt{n} \log n + \varepsilon^{-1} \log D \sqrt{n} + \varepsilon^{-2} \log n \log D + (\sqrt{n} + \varepsilon^{-1}) \log^4 n) \\ &= O(\varepsilon^{-2} \log^2 D \sqrt{n} \log n + \sqrt{n} \log^4 n) \\ &= \tilde{O}(\varepsilon^{-2} \sqrt{n} \log^2 D). \end{aligned}$$

We obtain approximation factor of  $2(1 + \varepsilon)^2 3 = 6(1 + \varepsilon)^2$ , which can be reduced to  $6 + \varepsilon$  by manipulating the value of  $\varepsilon$ .  $\square$

Using Theorem 4.20, we may decrease the dependency on the stretch of the metric at the cost of increasing the polylogarithmic factors.

**Theorem 6.15.** *Let  $G = (V, E, d_G)$  be a graph,  $n = |V|$ ,  $m = |E|$  and  $\varepsilon > 0$ . Denote by  $D$  the stretch of the metric induced by  $G$ . Let  $S \subseteq V$  be a dynamic set, subject to vertex additions and removals (initially  $S = \emptyset$ ). Then, after preprocessing in  $O(\varepsilon^{-1} \sqrt{n}(m + \varepsilon^{-2} n \log n \log(\varepsilon^{-1} n) \log D))$  expected time, we may maintain a  $(6 + \varepsilon)$ -approximate Steiner tree that spans  $S$ , handling each change*

to  $S$  in  $\tilde{O}(\varepsilon^{-5}\sqrt{n}\log D)$  expected amortized time. The algorithm uses  $O(\varepsilon^{-2}n\sqrt{n}\log n\log(\varepsilon^{-1}n)\log D)$  expected space.

Finally, we show a fully dynamic algorithm for planar graphs.

**Theorem 6.16.** *Let  $G = (V, E, d_G)$  be a planar graph,  $n = |V|$  and  $\varepsilon > 0$ . Denote by  $D$  the stretch of the metric induced by  $G$ . Let  $S \subseteq V$  be a dynamic set, subject to vertex additions and removals (initially  $S = \emptyset$ ). Then, after preprocessing in  $O(\varepsilon^{-2}n\log^2 n\log^2 D)$  time, we may maintain a  $(2 + \varepsilon)$ -approximate Steiner tree that spans  $S$ , handling each change to  $S$  in  $\tilde{O}(\varepsilon^{-2}\log^{2.5} D\sqrt{n})$  amortized time. The algorithm uses  $O(\varepsilon^{-2}n\log^2 n\log^2 D)$  space.*

*Proof.* We combine Lemma 6.13 with  $(1 + \varepsilon)$ -approximate vertex-color distance oracles for general graphs of Theorem 5.8. We have that  $\mu = 1 + \varepsilon = O(1)$ . The total number of portals is  $p_{\text{num}} = O(\varepsilon^{-1}n/\rho\log D)$ . Every vertex belongs to  $r_{\text{max}} = O(\log D)$  pieces and each piece is a planar graph containing  $r_{\text{max}}^{\#} = O(\rho)$  vertices. The total number of pieces is  $r_{\text{tot}} = O(n/\rho\log D)$ . Each **merge** and **split** operation requires  $T_{\text{merge}} = O(\varepsilon^{-1}\frac{n}{\rho}\log n\log D)$  time, whereas a **nearest** query can be handled in  $T_{\text{query}} = O((\varepsilon^{-1}\log n + \rho)\log D)$  time.

It takes  $O(\varepsilon^{-1}n\log^2 n\log D)$  time to initialize one oracle, and, since we copy it  $\Theta(\varepsilon^{-1}\log D)$  and its space usage is  $O(\varepsilon^{-1}n\log^2 n\log D)$ , the copying requires  $O(\varepsilon^{-2}n\log^2 n\log^2 D)$  time. This also dominates the preprocessing time. The space usage of the oracles and the algorithm is also  $O(\varepsilon^{-2}n\log^2 n\log^2 D)$ .

Each terminal addition or deletion requires time

$$\begin{aligned} & O(\varepsilon^{-2}\log^2 DT_{\text{merge}} + \varepsilon^{-1}\log DT_{\text{query}} + \varepsilon^{-1}p_{\text{num}} + \varepsilon^{-2}\log n\log D + (r_{\text{max}}r_{\text{max}}^{\#} + \varepsilon^{-1})\log^4 n) \\ &= O(\varepsilon^{-3}\log^3 D(n/\rho)\log n + \varepsilon^{-1}\log^2 D(\varepsilon^{-1}\log n + \rho) + \varepsilon^{-2}(n/\rho)\log D + \varepsilon^{-2}\log n\log D + \\ &\quad (\rho\log D + \varepsilon^{-1})\log^4 n) \\ &= O(\varepsilon^{-3}\log^3 D(n/\rho)\log n + \varepsilon^{-1}\log^2 D(\varepsilon^{-1}\log n + \rho) + \varepsilon^{-2}(n/\rho)\log D + \varepsilon^{-2}\log n\log D + \\ &\quad (\rho\log D + \varepsilon^{-1})\log^4 n). \end{aligned}$$

We set  $\rho = \varepsilon^{-1}\sqrt{n\log D}$  and obtain a running time of  $O(\varepsilon^{-2}\log^{2.5} D\sqrt{n}\log n + \sqrt{n}\log^{2.5} D\log^4 n) = \tilde{O}(\varepsilon^{-2}\log^{2.5} D\sqrt{n})$ . We obtain approximation factor of  $2(1 + \varepsilon)^3$ , which can be reduced to  $2 + \varepsilon$  by manipulating the value of  $\varepsilon$ .  $\square$

We may again apply Theorem 4.20 to decrease the dependency on the stretch of the metric induced by  $G$ .

**Theorem 6.17.** *Let  $G = (V, E, d_G)$  be a planar graph,  $n = |V|$  and  $\varepsilon > 0$ . Denote by  $D$  the stretch of the metric induced by  $G$ . Let  $S \subseteq V$  be a dynamic set, subject to vertex additions and removals (initially  $S = \emptyset$ ). Then, after preprocessing in  $\tilde{O}(\varepsilon^{-5}n\log D)$  time, we may maintain a  $(2 + \varepsilon)$ -approximate Steiner tree that spans  $S$ , handling each change to  $S$  in  $\tilde{O}(\varepsilon^{-5.5}\sqrt{n}\log D)$  amortized time.*

## 7 Steiner tree via bipartite emulators

In this section we present an alternative solution for maintaining approximate Steiner tree. This solution exposes a trade-off between the running time and the approximation ratio. The algorithms




Figure 6: An illustration of a bipartite emulator. The blue thick dashed path corresponds to the distance  $\tilde{\delta}_B(u, v) = d_B(ux) + d_B(xv)$ .

presented here are faster than the algorithms from previous sections, however their approximation ratio is much higher. We introduce the notion of *bipartite emulator*, which will be essential for the algorithms proposed here. In essence, a bipartite emulator is a sparse bipartite graph that stores the information about the approximate distances in the input graph. In this section we show how to use a bipartite emulator to maintain a good approximation of the minimum Steiner tree. In Section 7.1 we show how to construct the emulators themselves.

**Definition 7.1.** Let  $G$  be a graph and  $\alpha \geq 1$ . A bipartite emulator of  $G$  is a bipartite graph  $B = (V(G) \uplus N, E_B, d_B)$ , where the vertex set is partitioned into  $V(G)$  and a set of auxiliary vertices  $N$  representing the distances in  $G$ . For every  $u, v \in V(G)$ , it holds that  $\delta_G(u, v) \leq \delta_B(u, v) \leq \alpha \cdot \delta_G(u, v)$ . Moreover, for every  $u, v \in V(G)$ , there exists in  $B$  a vertex  $x \in N$ , such that  $d_B(ux) + d_B(xv) \leq \alpha \cdot \delta_G(u, v)$ . We say that  $\alpha$  is the stretch of emulator  $B$ . We denote  $\tilde{\delta}_B(u, v) = \min_{x \in N} \{d_B(ux) + d_B(xv)\}$  and we let  $\tilde{G} = (V(G), \binom{V(G)}{2}, \tilde{\delta}_B)$  be the approximation of the metric closure of  $G$  given by the emulator.

Note that a bipartite emulator is not an emulator according to the usual definition, as an emulator is required to have the same vertex set as the graph it emulates. Observe also that  $\tilde{G}$  may not be a metric space itself, as it may not satisfy the triangle inequality. See also Figure 6 for an illustration.

The definition shows that a bipartite emulator  $B$  contains estimates of all distances in  $G$  and that the distance estimate is given by a path of length two in  $B$ . Moreover, it easily follows that, to approximate  $ST(G, S)$ , it suffices to focus on  $MST(\tilde{G}[S])$ .

**Lemma 7.2.** Let  $B$  be a bipartite emulator of  $G$  with stretch  $\alpha$ . Then for any set of terminals  $S \subseteq V(G)$  it holds that

$$\tilde{\delta}_B(MST(\tilde{G}[S])) \leq 2\alpha d_G(ST(G, S)).$$

*Proof.* Denote  $T_{MST} = MST(\tilde{G}[S])$  and observe that

$$\begin{aligned} \tilde{\delta}_B(MST(\tilde{G}[S])) &\leq \tilde{\delta}_B(T_{MST}) && \text{by the definition of } MST \\ &\leq \alpha d_{\tilde{G}}(T_{MST}) && \text{by the stretch of } B \\ &\leq 2\alpha d_G(ST(G, S)) && \text{by Lemma 2.1.} \end{aligned}$$

□

To obtain the necessary distances (distances in  $\tilde{G}$  between vertices of  $S$ ) it is enough to consider the subgraph  $B[S \cup \Gamma(S)]$  of  $B$  induced by  $S$  and its neighborhood in  $B$ . We make sure that the




Figure 7: A minimum spanning tree  $T'$  of  $B[S \cup \Gamma(S)]$  (blue thick solid and dashed lines) and its subtree  $T$  created by removing all leaves that lie in  $N$  (blue thick solid lines).

degree of every vertex of  $V(G)$  in  $B$  is bounded. This will allow us to maintain graph  $B[S \cup \Gamma(S)]$  efficiently as  $S$  is undergoing changes. If a vertex of  $V(G)$  becomes a terminal, we add it and the adjacent edges to  $B[S \cup \Gamma(S)]$  in time proportional to its degree. We apply to  $B[S \cup \Gamma(S)]$  the dynamic MSF algorithm (see Theorem 2.2), which maintains the minimum spanning forest in a graph under edge additions or deletions. Each edge insertion/removal is handled in polylogarithmic time, so the time we need to add a vertex to  $B$  is bounded by its degree times logarithmic factors. The only problem that remains is that the maintained spanning tree may be too heavy, as  $\Gamma(S)$  may contain more vertices than needed. Note that every edge in  $\tilde{G}[S]$  corresponds to a vertex in  $\Gamma(S) \subseteq N$ . We need vertices corresponding to the edges of  $MST(\tilde{G}[S])$ , and the rest of  $\Gamma(S)$  unnecessarily increases the cost of the spanning tree. In the next lemma we show, that if we remove leaves of  $MST(B[S \cup \Gamma(S)])$  that belong to  $N$ , then the obtained tree costs at most twice the weight of  $MST(\tilde{G}[S])$ . The quite easy proof is provided in Section 4.2, where we develop a convenient language for such claims (see also Figure 7 for an illustration).

**Lemma 7.3.** *Let  $T' = MST(B[S \cup \Gamma(S)])$ , and let  $\mathcal{L}_N$  be the set of leaf vertices of  $T'$  contained in  $N$ . Let  $T$  be a tree obtained from  $T'$  by removing leaves in  $N$ , formally  $T = T'[V(T') \setminus \mathcal{L}_N]$ . Denote  $T_{MST} = MST(\tilde{G}[S])$ . Then  $d_B(T) \leq 2\delta_B(T_{MST})$ .*

Due to Lemma 7.3,  $MST(B[S \cup \Gamma(S)])$  without the leaves in  $N$  costs at most twice the weight of  $MST(\tilde{G}[S])$ , which is, by Lemma 7.2, at most  $4\alpha d_G(ST(G, S))$ . Clearly, the cost  $MST(B[S \cup \Gamma(S)])$  without the leaves in  $N$  cannot be smaller than the cost of the minimum Steiner tree in  $G$ . Thus, this approach gives us a  $4\alpha$ -approximation of the Steiner tree in the input graph  $G$ .

**Lemma 7.4.** *Let  $G$  be a graph and let  $B$  be its bipartite emulator of stretch  $\alpha$  (possibly obtained by a randomized algorithm). Moreover, assume that the (expected) degree in  $B$  of any vertex of  $V(G)$  is bounded by  $\Delta$ . Let  $S$  be a set subject to insertions and deletions, such that at any time  $S \subseteq V$ . Then, we can maintain a  $4\alpha$ -approximate Steiner tree of  $G$  that spans  $S$ , handling each update to  $S$  in (expected) amortized  $O(\Delta \log^4 n)$  time.*

Let us comment on the possible randomization of the construction of a bipartite emulator. In Section 7.1.1 we construct a randomized emulator, with only a bound on the expected degree of a vertex in  $V(G)$ . Then Lemma 7.4 gives us a bound on expected amortized time. That is, for any positive integer  $r$ , the expected time spent on the first  $r$  operations is bounded by  $O(r\Delta \log^4 n)$ .

In Section 7.1 we show two constructions of emulators. The construction for general graphs gives an emulator with  $\Delta = O(kn^{1/k})$  (in expectation) and stretch  $2k - 1$ , which yields an algorithm for maintaining  $(8k - 4)$ -approximate Steiner tree in  $O(kn^{1/k} \log^4 n)$  expected time per terminal addition or deletion. In case of planar graphs, we have an emulator with  $\Delta = O(\varepsilon^{-1} \log^2 n)$  and stretch  $1 + \varepsilon$ , so we obtain  $(4 + \varepsilon)$ -approximate algorithm, handling updates in  $O(\varepsilon^{-1} \log^6 n)$  time.

## 7.1 Constructing bipartite emulators

In this section we show how to construct emulators that can be plug into Lemma 7.4 to obtain dynamic algorithms for maintaining the Steiner tree.

### 7.1.1 General graphs

Our emulator for general graphs is based on an approximate distance oracle by Thorup and Zwick [37].

**Lemma 7.5** ([37]). *Let  $G = (V, E, d)$  be a graph,  $n = |V|$ ,  $m = |E|$ , and  $k \geq 1$  be an integer. Then, we can compute an emulator of  $G$  of stretch  $2k - 1$  in  $O(kmn^{1/k})$  expected time.*

We now modify the construction slightly in order to obtain a *bipartite* emulator of  $G$ . The emulator constructed in [37] is obtained by computing for each  $v \in V$  a set  $B(v)$  called a *bunch* and, for each  $w \in B(v)$ , adding an edge  $vw$  of length  $\delta(v, w)$ . The expected size of each bunch is  $O(kn^{1/k})$ . Moreover, from the query algorithm it follows that for any two  $u, v \in V$  there exists  $w \in B(u) \cap B(v)$  such that  $\delta(u, w) + \delta(w, v) \leq (2k - 1)\delta(u, v)$ .

Thus, the bipartite emulator  $G' = (V \cup V', E', d')$  can be constructed as follows. The vertex set of  $G'$  consists of two copies of  $V$ , i.e., for each  $v \in V$ , there is its copy  $v' \in V'$ . We connect each vertex  $v \in V$  to its copy with an edge of length 0. Moreover, for each  $v \in V$  and each  $w \in B(v)$  we add to  $G'$  an edge  $vw'$  of length  $\delta_G(v, w)$ .

It follows that  $G'$  is a bipartite graph, in which the degree of every  $v \in V$  is  $O(kn^{1/k})$  and for every two vertices  $u, w$  there exists a two-edge path of length  $\leq (2k - 1)\delta(u, v)$ . Thus, using Lemma 7.4 we obtain the following result.

**Theorem 7.6.** *Let  $G = (V, E, d)$  be a graph,  $n = |V|$ ,  $m = |E|$  and  $k \geq 1$  be an integer. Let  $S \subseteq V$  be a dynamic set, subject to vertex insertions and removals (initially  $S = \emptyset$ ). Then, after preprocessing in  $O(kmn^{1/k})$  expected time, we may maintain a  $(8k - 4)$ -approximate Steiner tree that spans  $S$ , handling each update to  $S$  in  $O(kn^{1/k} \log^4 n)$  expected amortized time.*

### 7.1.2 Planar graphs

In this section we show an emulator construction for planar graphs. As a result we obtain an algorithm that maintains a  $(4 + \varepsilon)$ -approximate Steiner tree in polylogarithmic time per update. In order to reach this goal, we use a construction by Thorup (section 3.8 in [36]) that we extend in order to construct an emulator (see Lemma 7.8).

Let  $G = (V, E)$  be a planar graph. The overall idea uses recursive division of  $G$  using balanced separators. A *balanced separator* of  $G$  is a set  $Q \subseteq V$  such that each connected component of  $G \setminus Q$  is a constant fraction smaller than  $G$ . We find a balanced separator of  $G$  that consists of a constant number of shortest paths  $P_1, \dots, P_k$  (the separator consists of vertices contained in these paths). For a shortest path  $P_i$ , we build an emulator that approximates all the shortest paths in  $G$  that

intersect  $P_i$ . Then, we recurse on each of the connected components of  $G \setminus (P_1 \cup \dots \cup P_k)$ . Hence, we now focus on the following problem. Given a planar graph  $G$  and a shortest path  $P$ , build an emulator that approximates all shortest paths intersecting  $P$ .

We define a *connection* to be an edge that connects a vertex  $v \in V$  with a vertex  $a \in P$  and has length  $d(va)$  at least  $\delta_G(v, a)$  (it would be convenient to assume that the length is *equal* to  $\delta_G(v, a)$ , but the algorithm we use may sometimes give longer connections). A connection  $vb$   $\varepsilon$ -covers  $x$  if  $d(vb) + \delta_G(b, x) \leq (1 + \varepsilon)\delta_G(v, x)$ ; observe that the distance  $\delta_G(b, x)$  can be equivalently measured along the path  $P$ . A set of connections  $C(v, P)$  is  $\varepsilon$ -covering if it  $\varepsilon$ -covers every  $x \in P$ .

**Lemma 7.7.** *Let  $G = (V, E, d)$  be a planar graph,  $n = |V|$ , and  $0 < \varepsilon \leq 1$ . Let  $P$  be a shortest path in  $G$ . For each  $v \in V(G)$  we can construct an  $\varepsilon$ -covering set  $C(v, P)$  of size  $O(\varepsilon^{-1})$  in  $O(\varepsilon^{-1}n \log n)$  total time.*

A very similar fact is shown in [36], but the definition of  $\varepsilon$ -covering used there is slightly different, so, for the sake of completeness, we rewrite the proof.

*Proof.* Let  $\varepsilon_0 = \varepsilon/2$ . We say that a connection  $vb$  *strongly- $\varepsilon$ -covers*  $a$  if  $d(vb) + (1 + \varepsilon)\delta_G(b, a) \leq (1 + \varepsilon)\delta_G(v, a)$ . By Lemma 3.18 in [36], for each  $v \in V(G)$  we can construct a strongly- $(\varepsilon_0/2)$ -covering set  $D(v, P)$  of size  $O(\varepsilon_0^{-1} \log n)$  in  $O(\varepsilon_0^{-1}n \log n)$  time.<sup>7</sup> We now show that we can use it to construct an  $\varepsilon$ -covering set  $C(v, P) \subseteq D(v, P)$  of size  $O(\varepsilon^{-1})$ .

Let  $vc$  be the shortest connection from  $D(v, P)$  and  $s$  be one of the two endpoints of  $P$ . We add  $vc$  to  $C(v, P)$ . Now, iterate through connections in  $D(v, P)$  starting from  $vc$  and going towards  $s$ . Let  $vb$  be the connection that was most recently added to  $C(v, P)$ . If for the current connection  $va$  we have  $d(vb) + \delta_G(b, a) > (1 + \varepsilon_0)d(va)$ , we add  $va$  to  $C(v, P)$ . Then, we run a similar procedure using the other endpoint of  $P$ .

To prove that  $C(v, P)$  covers every vertex between  $c$  and  $s$ , consider some vertex  $x \in P$ . There exists a connection  $va \in D(v, P)$  that strongly- $(\varepsilon_0/2)$ -covers  $vx$ , so  $d(va) + (1 + \varepsilon_0/2)\delta_G(a, x) \leq (1 + \varepsilon_0/2)\delta_G(v, x)$ . If this connection is in  $C(v, P)$  then  $vx$  is strongly- $\varepsilon_0/2$ -covered and obviously also  $\varepsilon$ -covered. Otherwise, there exists a connection  $vb$  such that  $d(vb) + \delta_G(b, a) \leq (1 + \varepsilon_0)d(va)$ . We have

$$\begin{aligned} d(vb) + \delta_G(b, x) &\leq d(vb) + \delta_G(b, a) + \delta_G(a, x) \\ &\leq (1 + \varepsilon_0)d(va) + \delta_G(a, x) \\ &\leq (1 + \varepsilon_0)d(va) + (1 + \varepsilon_0)(1 + \varepsilon_0/2)\delta_G(a, x) \\ &= (1 + \varepsilon_0)(d(va) + (1 + \varepsilon_0/2)\delta_G(a, x)) \\ &\leq (1 + \varepsilon_0)(1 + \varepsilon_0/2)\delta_G(v, x) \\ &\leq (1 + \varepsilon)\delta_G(v, x). \end{aligned}$$

The last inequality follows from  $\varepsilon_0 = \varepsilon/2 \leq 1$ . It remains to bound the size of  $C(v, P)$ . Let  $f(vb) = d(vb) + \delta_G(b, s)$ . As connections are added to  $C(v, P)$  we trace the value of  $f(vb)$ , where  $vb$  is the last connection that we have added. Every time we add a connection  $va$  we reduce the value of  $f$  by

$$f(vb) - f(va) = d(vb) + \delta_G(b, s) - d(va) - \delta_G(a, s) = d(vb) - d(va) + \delta_G(b, a) > \varepsilon_0 d(va) \geq \varepsilon_0 \delta_G(v, c).$$

<sup>7</sup>Strictly speaking, the strongly- $\varepsilon$ -covering set according to our definition is an  $\varepsilon/(\varepsilon + 1)$ -covering set according to the definition used in the statement of Lemma 3.18 of [36].

However, the total change equals

$$f(vc) - f(vs) = d(vc) + \delta_G(c, s) - d(vs) \leq (1 + \varepsilon_0)\delta_G(v, c) + \delta_G(c, s) - \delta_G(v, s) \leq (2 + \varepsilon_0)\delta_G(v, c).$$

Thus, at most  $O(\varepsilon_0^{-1}) = O(\varepsilon^{-1})$  connections are added to  $C(v, P)$ .

The same procedure is then repeated for the other endpoint of  $P$ , so we get a total of  $O(\varepsilon^{-1})$  connections.  $\square$

**Lemma 7.8.** *Let  $G = (V, E, d)$  be a planar graph,  $n = |V|$ ,  $0 < \varepsilon \leq 1$ . Let  $P$  be a shortest path in  $G$ . For each  $v \in V(G)$  we can construct a set of connections  $C'(v, P)$  of size  $O(\varepsilon^{-1} \log n)$ , which satisfies the following property. For any two vertices  $u, w \in V$ , if the shortest path between  $u$  and  $w$  intersects  $P$ , then for some  $x \in P$  there exist connections  $ux \in C'(u, P)$  and  $wx \in C'(w, P)$ , such that  $\delta(u, w) \leq d(ux) + d(wx) \leq (1 + \varepsilon)\delta(u, w)$ . The sets  $C'(v, P)$  can be constructed in  $O(\varepsilon^{-1} n \log n)$  time.*

*Proof.* First, using Lemma 7.7, for every  $v \in V$  we construct an  $\varepsilon$ -covering sets of connections  $C(v, P)$ . Assume that the shortest path  $Q$  between  $u$  and  $w$  intersects  $P$  in  $x \in P$ . There exists a path  $Q'$  between  $u$  and  $w$  which consists of a connection, subpath of  $P$ , denoted henceforth  $Q'_P$ , and another connection. Moreover,  $d(Q') \leq (1 + \varepsilon)d(Q)$ . We call each path of this form an *approximating path*. Our goal is to substitute every approximating path with an approximating path that consists solely of two connections from  $C'(v, P)$ .

The construction is done recursively. The parameter of the recursion is a subpath  $P'$  of  $P$ . Consider a single step, with a parameter  $P' = p_1 p_2 \dots p_k$ . Let  $p_m = p_{\lfloor k/2 \rfloor}$  be the middle vertex of  $P'$ . For any  $v \in V$  and  $p_i \in P'$ , if there is an connection  $vp_i \in C(v, P)$ , we add a connection  $vp_m$  of length  $d(vp_i) + \delta(p_i, p_m)$  to  $C'(v, P)$ . Then, we recurse on  $P_1 = p_1 p_2 \dots p_{m-1}$  and  $P_2 = p_{m+1} \dots p_k$ . Lastly, for each  $p_i \in P$  we add a connection  $p_i p_i$  of length 0.

To prove the correctness of the procedure, consider now the aforementioned approximating path  $Q'$ , and recall that  $Q'_P = P \cap Q'$ . Let  $p$  be the vertex that is taken as  $p_m$  in the closest to root node in the recursion tree of the algorithm, among all vertices of  $Q'_P$ . Observe that in the single recursive step when  $p_m = p$ , we add to  $C'(v, P)$  the connections  $up$  and  $wp$  of length exactly equal to the length of the part of  $Q'$  between  $u$  and  $p$ , and the part of  $q'$  between  $p$  and  $w$ , respectively. Also, the connections we add clearly do not cause any distances to be underestimated.

The running time of each step is proportional to the length of the subpath we consider and the number of connections incident to this subpath. Moreover, every connection may be considered in at most  $O(\log n)$  recursive calls, so we we add to  $C'(v, P)$  at most  $O(\varepsilon^{-1} \log n)$  connections. It follows that the total running time of the procedure is  $O(\varepsilon^{-1} n \log n)$ .  $\square$

**Lemma 7.9.** *Let  $G = (V, E, d)$  be a planar graph,  $n = |V|$ ,  $0 < \varepsilon \leq 1$ . We can construct a bipartite emulator  $B = (V \cup N, E_B, d_B)$  of  $G$  of stretch  $1 + \varepsilon$ . The degree of every  $v \in V$  in  $B$  is  $O(\varepsilon^{-1} \log^2 n)$ . The graph  $B$  can be constructed in  $O(\varepsilon^{-1} n \log^2 n)$ .*

*Proof.* We begin with  $B$  being a graph with vertex set  $V$  and no edges. The construction is done recursively. As it is shown, e.g., in [36], each planar graph admits a balanced separator that consists of a constant number of shortest paths  $P_1, \dots, P_k$ , and, moreover, such a separator can be found in  $O(n)$  time. For each path  $P_i$  we use Lemma 7.8 to construct a set of connections  $C'(v, P_i)$  for every  $v \in V$ . Next, we iterate through the vertices of the paths  $P_i$ . For each vertex  $w \in P_i$  we add a new auxiliary vertex  $w'$  to  $B$  and add an edge  $uw'$  for each connection  $uw$  from  $G$  (the length of

the edge is the length of the connection). After that, we recurse on each connected component of  $G \setminus (P_1 \cup \dots \cup P_k)$ .

Let us now prove the correctness of the construction. Consider any two  $v_1, v_2 \in V$  and the shortest path  $Q$  between them. At some step of the recursion, some vertex of  $Q$  belongs to the separator that we have found. From the construction, it follows that in this step we have added to  $B$  a vertex  $w'$  and edges  $v_1w'$  and  $v_2w'$  to  $B$  of total length at most  $(1 + \varepsilon)\delta_G(v_1, v_2)$ . Moreover, since the length of every connection is lower bounded by the length of the corresponding shortest path,  $B$  may not underestimate the distances between vertices of  $V$ .

Since every vertex  $v \in V$  takes part in  $O(\log n)$  recursive steps and in every step we add  $O(\varepsilon^{-1} \log n)$  edges incident to any  $v \in V$ , we have that the degree of any vertex of  $V$  in  $B$  is  $O(\varepsilon^{-1} \log^2 n)$ . As shown in [36], finding the separators requires  $O(n \log n)$  total time. The running time of every recursive step is dominated by the time from Lemma 7.8. Summing this over all recursive steps, we get that the construction can be done in  $O(\varepsilon^{-1} n \log^2 n)$  time.  $\square$

By constructing  $B$  according to Lemma 7.9 and applying Lemma 7.4 we obtain the following.

**Theorem 7.10.** *Let  $G = (V, E, d)$  be a planar graph and  $0 < \varepsilon \leq 1$ . Let  $S \subseteq V$  be a dynamic set, subject to vertex insertions and removals (initially  $S = \emptyset$ ). Then, after preprocessing in  $O(\varepsilon^{-1} n \log^2 n)$  time, we may maintain a  $(4 + \varepsilon)$ -approximate Steiner tree that spans  $S$ , handling each update to  $S$  in  $O(\varepsilon^{-1} \log^6 n)$  amortized time.*

## 8 Conclusions

In our work we have given the first sublinear dynamic algorithms for the Steiner tree problem and the subset TSP problem, where the set of terminals changes over time. We have exhibited a tight connection to vertex-color distance oracles and the replacement schemes previously developed for the online model.

We believe our result would inspire further work on the dynamic Steiner tree problem, focusing not only on the number of replacements required to maintain a good Steiner tree (as in the online model), but also on the entire efficiency of updates. In particular, we would like to emphasize here the following open questions.

1. The approximation guarantee of our dynamic algorithms, even for planar graphs, contains the term 2, inherited from the fact that we approximate Steiner tree by a minimum spanning tree in the metric closure of the input graph. We have chosen this approach for its simplicity. Can any of the approximation algorithms that guarantees better approximation than 2 be adapted to our dynamic setting? In particular, there exists a PTAS for the (static) Steiner tree problem in planar graphs [9]. Can we obtain a dynamic  $(1 + \varepsilon)$ -approximation algorithm for planar graphs? The other direction to approach this question is to devise simple and almost linear time approximation algorithm for Steiner tree that would beat the factor of 2 in non-planar graphs. This could be of some practical importance in big data applications.
2. A second question concerning better approximation ratios would be whether it is possible to obtain a  $c$ -approximate vertex-color distance oracle in general graphs for some constant  $c < 3$ .



3. In the deletion step, we relied on a somehow weak variant of a *color-to-color* distance query, where the returned distance takes into account only portal distances between the colors. Is it possible to enhance our oracles to perform a general color-to-color queries efficiently?
4. Our decremental and fully dynamic algorithms for general graphs require preprocessing time of roughly order  $O(\sqrt{nm})$  and space of roughly order  $O(n\sqrt{n})$ . Can it be improved to near-linear time and space?
5. The running time of our fully dynamic algorithms depends on the stretch of the metric. Can this dependency be avoided, or at least further reduced?
6. Which other combinatorial problems, e.g., facility location, allow for dynamic sublinear time approximation algorithms?

## References

- [1] I. Abraham, D. Malkhi, and D. Ratajczak. Compact multicast routing. In *23rd International Symposium on Distributed Computing (DISC 2009)*. Springer Verlag, September 2009.
- [2] E. Aharoni and R. Cohen. Restricted dynamic steiner trees for scalable multicast in datagram networks. In *INFOCOM '97. Sixteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Driving the Information Revolution., Proceedings IEEE*, volume 2, pages 876–883 vol.2, Apr 1997.
- [3] S. Alstrup, J. Holm, K. de Lichtenberg, and M. Thorup. Maintaining information in fully dynamic trees with top trees. *ACM Transactions on Algorithms*, 1(2):243–264, 2005.
- [4] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *In FOCS*, pages 295–304, 2004.
- [5] S. Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *J. ACM*, 45(5):753–782, 1998.
- [6] F. Bauer and A. Varma. ARIES: A rearrangeable inexpensive edge-based on-line Steiner algorithm. *IEEE Journal of Selected Areas in Communications*, 15:382–397, 1995.
- [7] M. W. Bern and P. E. Plassmann. The Steiner problem with edge lengths 1 and 2. *Inf. Process. Lett.*, 32(4):171–176, 1989.
- [8] V. Bilò, M. Flammini, and L. Moscardelli. The price of stability for undirected broadcast network design with fair cost allocation is constant. In *FOCS*, pages 638–647. IEEE Computer Society, 2013.
- [9] G. Borradaile, P. N. Klein, and C. Mathieu. An  $O(n \log n)$  approximation scheme for Steiner tree in planar graphs. *ACM Transactions on Algorithms*, 5(3), 2009.
- [10] J. Byrka, F. Grandoni, T. Rothvoss, and L. Sanità. Steiner tree approximation via iterative randomized rounding. *J. ACM*, 60(1):6:1–6:33, Feb. 2013.

- [11] T. M. Chan. Dynamic subgraph connectivity with geometric applications. In *Proceedings of the Thirty-fourth Annual ACM Symposium on Theory of Computing*, STOC '02, pages 7–13, New York, NY, USA, 2002. ACM.
- [12] T. M. Chan, M. Pătraşcu, and L. Roditty. Dynamic connectivity: Connecting to networks and geometry. *SIAM Journal on Computing*, 40(2):333–349, 2011. See also FOCS'08, arXiv:0808.1128.
- [13] S. Chechik. Improved distance oracles and spanners for vertex-labeled graphs. In L. Epstein and P. Ferragina, editors, *ESA*, volume 7501 of *Lecture Notes in Computer Science*, pages 325–336. Springer, 2012.
- [14] X. Cheng, Y. Li, D.-Z. Du, and H. Ngo. Steiner trees in industry. In D.-Z. Du and P. Pardalos, editors, *Handbook of Combinatorial Optimization*, pages 193–216. Springer US, 2005.
- [15] F. Chung. Graph theory in the information age. *Notices of the American Mathematical Society*, 57(06):726.
- [16] M. Cygan, L. Kowalik, M. Mucha, M. Pilipczuk, and P. Sankowski. Fast approximation in subspaces by doubling metric decomposition. In *ESA (1)*, pages 72–83, 2010.
- [17] S. E. Deering and D. R. Cheriton. Multicast routing in datagram internetworks and extended lans. *ACM Trans. Comput. Syst.*, 8(2):85–110, May 1990.
- [18] C. Demetrescu, D. Eppstein, Z. Galil, and G. F. Italiano. Algorithms and theory of computation handbook. chapter Dynamic Graph Algorithms, pages 9–9. Chapman & Hall/CRC, 2010.
- [19] N. Garg, A. Gupta, S. Leonardi, and P. Sankowski. Stochastic analyses for online combinatorial optimization problems. In *SODA*, pages 942–951, 2008.
- [20] A. Gu, A. Gupta, and A. Kumar. The power of deferral: maintaining a constant-competitive Steiner tree online. In *STOC*, pages 525–534, 2013.
- [21] A. Gupta and A. Kumar. Online Steiner tree with deletions. In C. Chekuri, editor, *SODA*, pages 455–467. SIAM, 2014.
- [22] A. Gupta, M. Pál, R. Ravi, and A. Sinha. Boosted sampling: approximation algorithms for stochastic optimization. In *STOC*, pages 417–426, 2004.
- [23] T. Hagerup, P. B. Miltersen, and R. Pagh. Deterministic dictionaries. *Journal of Algorithms*, 41(1):69 – 85, 2001.
- [24] M. R. Henzinger, P. N. Klein, S. Rao, and S. Subramanian. Faster shortest-path algorithms for planar graphs. *J. Comput. Syst. Sci.*, 55(1):3–23, 1997.
- [25] D. Hermelin, A. Levy, O. Weimann, and R. Yuster. Distance oracles for vertex-labeled graphs. In *ICALP*, pages 490–501, 2011.
- [26] J. Holm, K. de Lichtenberg, and M. Thorup. Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity. *J. ACM*, 48(4):723–760, 2001.

- [27] S.-P. Hong, H. Lee, and B. H. Park. An efficient multicast routing algorithm for delay-sensitive applications with dynamic membership. In *INFOCOM '98. Seventeenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 3, pages 1433–1440 vol.3, Mar 1998.
- [28] M. Imase and B. M. Waxman. Dynamic Steiner tree problem. *SIAM J. Discrete Math.*, 4(3):369–384, 1991.
- [29] M. Li, C. Ma, and L. Ning.  $(1 + \epsilon)$ -distance oracles for vertex-labeled planar graphs. In *Theory and Applications of Models of Computation*, volume 7876 of *Lecture Notes in Computer Science*, pages 42–51. 2013.
- [30] N. Megow, M. Skutella, J. Verschae, and A. Wiese. The power of recourse for online MST and TSP. In *ICALP*, pages 689–700. 2012.
- [31] K. Mehlhorn. A faster approximation algorithm for the Steiner problem in graphs. *Inf. Process. Lett.*, 27(3):125–128, 1988.
- [32] J. S. B. Mitchell. Guillotine subdivisions approximate polygonal subdivisions: A simple polynomial-time approximation scheme for geometric TSP, k-MST, and related problems. *SIAM J. Comput.*, 28:402–408, 1996.
- [33] S. Raghavan, G. Manimaran, C. Siva, and R. Murthy. A rearrangeable algorithm for the construction of delay-constrained dynamic multicast trees. *IEEE/ACM Transactions on Networking*, 7:514–529, 1999.
- [34] G. Robins and A. Zelikovsky. Tighter bounds for graph Steiner tree approximation. *SIAM J. Discret. Math.*, 19(1):122–134, May 2005.
- [35] R. E. Tarjan. Dynamic trees as search trees via euler tours, applied to the network simplex algorithm. *Mathematical Programming*, 78:169–177, 1997.
- [36] M. Thorup. Compact oracles for reachability and approximate distances in planar digraphs. *J. ACM*, 51:993–1024, 2004.
- [37] M. Thorup and U. Zwick. Approximate distance oracles. *J. ACM*, 52(1):1–24, 2005.
- [38] D. E. Willard. Log-logarithmic worst-case range queries are possible in space  $\Theta(n)$ . *Inf. Process. Lett.*, 17(2):81–84, 1983.