

Almost Optimal Distance Oracles for Planar Graphs*

Panagiotis Charalampopoulos
King's College London
UK

IDC Herzliya
Israel
panagiotis.charalampopoulos@kcl.ac.uk

Shay Mozes
IDC Herzliya
Israel
smozes@idc.ac.il

Paweł Gawrychowski
University of Wrocław
Poland

gawry@cs.uni.wroc.pl

Oren Weimann
University of Haifa
Israel
oren@cs.haifa.ac.il

ABSTRACT

We present new tradeoffs between space and query-time for exact distance oracles in directed weighted planar graphs. These tradeoffs are almost optimal in the sense that they are within polylogarithmic, subpolynomial or arbitrarily small polynomial factors from the naïve linear space, constant query-time lower bound. These tradeoffs include: (i) an oracle with space $O(n^{1+\epsilon})$ and query-time $\tilde{O}(1)$ for any constant $\epsilon > 0$, (ii) an oracle with space $\tilde{O}(n)$ and query-time $O(n^\epsilon)$ for any constant $\epsilon > 0$, and (iii) an oracle with space $n^{1+o(1)}$ and query-time $n^{o(1)}$.

CCS CONCEPTS

• Theory of computation → Shortest paths.

KEYWORDS

planar graphs, distance oracles, shortest paths, Voronoi diagrams

ACM Reference Format:

Panagiotis Charalampopoulos, Paweł Gawrychowski, Shay Mozes, and Oren Weimann. 2019. Almost Optimal Distance Oracles for Planar Graphs. In *Proceedings of the 51st Annual ACM SIGACT Symposium on the Theory of Computing (STOC '19)*, June 23–26, 2019, Phoenix, AZ, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3313276.3316316>

1 INTRODUCTION

Computing shortest paths is one of the most fundamental and well-studied algorithmic problems, with numerous applications in various fields. In the data structure version of the problem, the goal is to preprocess a graph into a compact representation such that the distance (or a shortest path) between any pair of vertices can be retrieved efficiently. Such data structures are called *distance oracles*.

*This work was partially supported by the Israel Science Foundation under grant number 592/17.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

STOC '19, June 23–26, 2019, Phoenix, AZ, USA
© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6705-9/19/06...\$15.00
<https://doi.org/10.1145/3313276.3316316>

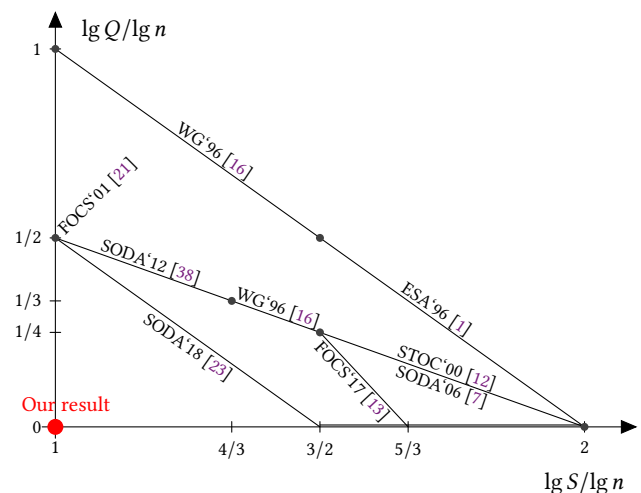


Figure 1: Tradeoff of the space (S) vs. the query time (Q) for exact distance oracles in planar graphs on a doubly logarithmic scale, hiding subpolynomial factors. The previous tradeoffs are indicated by solid black lines and points, while our new tradeoff is indicated by the red point at the bottom left.

Distance oracles are useful in applications ranging from navigation, geographic information systems and logistics, to computer games, databases, packet routing, web search, computational biology, and social networks. The topic has been studied extensively; see for example the survey by Sommer [43] for a comprehensive overview and references.

The two main measures of efficiency of a distance oracle are the space it occupies and the time it requires to answer a distance query. To appreciate the tradeoff between these two quantities consider two naïve oracles. The first stores all $\Theta(n^2)$ pairwise distances in a table, and answers each query in constant time using table lookup. The second only stores the input graph, and runs a shortest path algorithm over the entire graph upon each query. Both of these oracles are not adequate when working with mildly large graphs. The first consumes too much space, and the second is too slow in answering queries. A third quantity of interest is the preprocessing time required to construct the oracle. Since computing the data

structure is done offline, this quantity is often considered less important. However, one approach to dealing with dynamic networks is to recompute the entire data structure quite frequently, which is only feasible when the preprocessing time is reasonably small.

One of the ways to design oracles with small space is to consider *approximate distance oracles* (allowing a small stretch in the distance output). However, it turns out that one cannot get both small stretch and small space. In their seminal paper, Thorup and Zwick [45] showed that, assuming the girth conjecture of Erdős [18], there exist dense graphs for which no oracle with size less than $n^{1+1/k}$ and stretch less than $2k - 1$ exists. Pătraşcu and Roditty [42] showed that even sparse graphs with $\tilde{O}(n)$ edges do not have distance oracles with stretch better than 2 and subquadratic space, conditioned on a widely believed conjecture on the hardness of set intersection. To bypass these impossibility results one can impose additional structure on the graph. In this work we follow this approach and focus on distance oracles for planar graphs.

1.1 Distance Oracles for Planar Graphs

The importance of studying distance oracles for planar graphs stems from several reasons. First, distance oracles for planar graphs are ubiquitous in real-world applications such as geographical navigation on road networks [43] (road networks are often theoretically modeled as planar graphs even though they are not quite planar due to tunnels and overpasses). Second, shortest paths in planar graphs exhibit a remarkable wealth of structural properties that can be exploited to obtain efficient oracles. Third, techniques developed for shortest paths problems in planar graphs often carry over (because of intricate and elegant connections) to maximum flow problems. Fourth, planar graphs have proven to be an excellent sandbox for the development of algorithms and techniques that extend to broader families of graphs.

As such, distance oracles for planar graphs have been extensively studied. Works on exact distance oracles for planar graphs started in the 1990's with oracles requiring $\tilde{O}(n^2/Q)$ space and $\tilde{O}(Q)$ query-time for any $Q \in [1, n]$ [1, 16]. Note that this includes the two trivial approaches mentioned above. Over the past three decades, many other works presented exact distance oracles for planar graphs with increasingly better space to query-time tradeoffs [1, 7, 12, 13, 16, 21, 23, 38, 40]. Figure 1 illustrates the advancements in the space/query-time tradeoffs over the years. Until recently, no distance oracles with subquadratic space and polylogarithmic query time were known. Cohen-Addad et al. [13], inspired by Cabello's use of Voronoi diagrams for the diameter problem in planar graphs [8], provided the first such oracle. The currently best known tradeoff [23] is an oracle with $\tilde{O}(n^{3/2}/Q)$ space and $\tilde{O}(Q)$ query-time for any $Q \in [1, n^{1/2}]$. Note that all known oracles with nearly linear (i.e. $\tilde{O}(n)$) space require $\Omega(\sqrt{n})$ time to answer queries.

The holy grail in this area is to design an exact distance oracle for planar graphs with both linear space and constant query-time. It is not known whether this goal can be achieved. We do know, however, (for nearly twenty years now) that approximate distance oracles can get very close. For any fixed $\epsilon > 0$ there are $(1 + \epsilon)$ -approximate distance oracles that occupy nearly-linear space and answer queries

in polylogarithmic, or even constant time [10, 24, 28–30, 44, 46]. However, the main question of whether an approximate answer is the best one can hope for, or whether exact distance oracles for planar graphs with linear space and constant query time exist, remained a wide open important and interesting problem.

1.2 Our Results and Techniques

In this paper we approach the optimal tradeoff between space and query time for reporting exact distances. We design exact distance oracles that require almost-linear space and answer distance queries in polylogarithmic time. Specifically, given a planar graph of size n , we show how to construct in roughly $O(n^{3/2})$ time a distance oracle admitting any of the following (space, query-time) tradeoffs (see Theorem 7 and Corollary 8 for the exact statements).

- (i) $\langle \tilde{O}(n^{1+\epsilon}), O(\log^{1/\epsilon} n) \rangle$, for any constant $1/2 \geq \epsilon > 0$;
- (ii) $\langle O(n \log^{2+1/\epsilon} n), O(n^{2\epsilon}) \rangle$, for any constant $\epsilon > 0$;
- (iii) $\langle n^{1+o(1)}, n^{o(1)} \rangle$.

1.2.1 Voronoi Diagrams. The main tool we use to obtain this result is point location in *Voronoi diagrams on planar graphs*. The concept of Voronoi diagrams has been used in computational geometry for many years (cf. [2, 14]). We consider *graphical* (or network) Voronoi diagrams [20, 35]. At a high level, a graphical Voronoi diagram with respect to a set S of sites is a partition of the vertices into $|S|$ parts, called Voronoi cells, where the cell of site $s \in S$ contains all vertices that are closer (in the shortest path metric) to s than to any other site in S . Graphical Voronoi diagrams have been studied and used quite extensively, most notably in applications in road networks (e.g., [15, 17, 26, 41]).

Perhaps the most elementary operation on Voronoi diagrams is *point location*. Given a point (vertex) v , one wishes to efficiently determine the site s such that v belongs to the Voronoi cell of s . Cohen-Addad et al. [13], inspired by Cabello's [8] breakthrough use of Voronoi diagrams in planar graphs, suggested a way to perform point location that led to the first exact distance oracle for planar graphs with subquadratic space and polylogarithmic query time. A simpler and more efficient point location mechanism for Voronoi diagrams in planar graphs was subsequently developed in [23]. In both oracles, the space efficiency is obtained from the fact that the size of the representation of a Voronoi diagram is proportional to the number of sites and not to the total number of vertices.

To obtain our result, we add two new ingredients to the point location mechanism of [23]. The first is the use of what might be called *external Voronoi diagrams*. Unlike previous constructions, instead of working with Voronoi diagrams for every piece in some partition (r -division) of the graph, we work with many overlapping Voronoi diagrams, representing the complements of such pieces. This is analogous to the use of external dense distance graphs in [5, 38]. This approach alone leads to an oracle with space $\tilde{O}(n^{4/3})$ and query time $\tilde{O}(\log^2 n)$ (see Section 3). The obstacle with pushing this approach further is that the point location mechanism consists of auxiliary data for each piece, whose size is proportional to the size of the complement of the piece, which is $\Theta(n)$ rather than the much smaller size of the piece. We show that this problem can be mitigated by using recursion, and storing much less auxiliary data at a coarser level of granularity. This approach is made possible by

¹The $\tilde{O}(\cdot)$ notation hides polylogarithmic factors.

a more modular view of the point location mechanism which we present in Section 2, along with other preliminaries. The proof of our main space/query-time tradeoffs is given in Section 4 and the algorithm to efficiently construct these oracles is given in Section 5.

2 PRELIMINARIES

In this section we review the main techniques required for describing our result. Throughout the paper we consider as input a weighted directed planar graph $G = (V(G), E(G))$, embedded in the plane. We use the term length for edges and paths throughout the paper. We use $|G|$ to denote the number of vertices in G . Since planar graphs are sparse, $|E(G)| = O(|G|)$ as well. The dual of a planar graph G is another planar graph G^* whose vertices correspond to faces of G and vice versa. Arcs of G^* are in one-to-one correspondence with arcs of G ; there is an arc e^* from vertex f^* to vertex g^* of G^* if and only if the corresponding faces f and g of G are to the left and right of the arc e , respectively.

We assume that the input graph has no negative length cycles. We can transform the graph in $O(n \frac{\log^2 n}{\log \log n})$ time in a standard way [39] so that all edge lengths are non-negative and distances are preserved. With another standard transformation we can guarantee that each vertex has constant degree and that the graph is triangulated, while distances are preserved and without increasing asymptotically the size of the graph. We further assume that shortest paths are unique; this can be ensured in $O(n)$ time by a deterministic perturbation of the edge lengths [19]. Each original distance can be recovered from the corresponding distance in the transformed graph in constant time. Let $d_G(u, v)$ denote the distance from a vertex u to a vertex v in G .

2.1 Multiple-Source Shortest Paths

The multiple-source shortest paths (MSSP) data structure [31] represents all shortest path trees rooted at the vertices of a single face f in a planar graph using a persistent dynamic tree. It can be constructed in $O(n \log n)$ time, requires $O(n \log n)$ space, and can report any distance between a vertex of f and any other vertex in the graph in $O(\log n)$ time. We note that it can be augmented to also return the first edge of this path within the same complexities. The authors of [23] show that it can be further augmented—within the same complexities—such that given two vertices $u, v \in G$ and a vertex x of f it can return, in $O(\log n)$ time, whether u is an ancestor of v in the shortest path tree rooted at x as well as whether u occurs before v in a preorder traversal of this tree. We consider shortest path trees as ordered trees with the order inherited from the planar embedding.

2.2 Separators and Recursive Decompositions

Miller [36] showed how to compute, in a biconnected triangulated planar graph with n vertices, a simple cycle of size $O(\sqrt{n})$ that separates the graph into two subgraphs, each with at most $2n/3$ vertices. Simple cycle separators can be used to recursively separate a planar graph until pieces have constant size. The authors of [33] show how to obtain a complete recursive decomposition tree \mathcal{A} of G in $O(n)$ time. \mathcal{A} is a binary tree whose nodes correspond to subgraphs of G (pieces), with the root being all of G and the leaves being pieces of constant size. We identify each piece P with the

node representing it in \mathcal{A} . We can thus abuse notation and write $P \in \mathcal{A}$. An r -division [22] of a planar graph, for $r \in [1, n]$, is a decomposition of the graph into $O(n/r)$ pieces, each of size $O(r)$, such that each piece has $O(\sqrt{r})$ boundary vertices, i.e. vertices that belong to some separator along the recursive decomposition used to obtain P . Another desired property of an r -division is that the boundary vertices lie on a constant number of faces of the piece (holes). For every r larger than some constant, an r -division with few holes is represented in the decomposition tree \mathcal{A} of [33]. In fact, it is not hard to see that if the original graph G is triangulated then all vertices of each hole of a piece are boundary vertices. Throughout the paper, to avoid confusion, we use “nodes” when referring to \mathcal{A} and “vertices” when referring to G . We denote the boundary vertices of a piece P by ∂P . We refer to non-boundary vertices as internal. We assume for simplicity that each hole is a simple cycle. Non-simple cycles do not pose a significant obstacle, as we discuss at the end of Section 4.

It is shown in [33, Theorem 3] that, given a geometrically decreasing sequence of numbers $(r_m, r_{m-1}, \dots, r_1)$, where r_1 is a sufficiently large constant, $r_{i+1}/r_i = b$ for all i for some $b > 1$, and $r_m = n$, we can obtain the r_i -divisions for all i in time $O(n)$ in total. For convenience we define the only piece in the r_m division to be G itself. These r -divisions satisfy the property that a piece in the r_i -division is a—not necessarily strict—descendant (in \mathcal{A}) of a piece in the r_j -division for each $j > i$. This ancestry relation between the pieces of an r -division can be captured by a tree \mathcal{T} called the recursive r -division tree.

The boundary vertices of a piece $P \in \mathcal{A}$ that lie on a hole h of P separate the graph G into two subgraphs G_1 and G_2 (the cycle is in both subgraphs). One of these two subgraphs is enclosed by the cycle and the other is not. Moreover, P is a subgraph of one of these two subgraphs, say G_1 . We then call G_2 the outside of hole h with respect to P and denote it by $P^{h, out}$. In the sections where we assume that the boundary vertices of each piece lie on a single hole that is a simple cycle, the outside of this hole with respect to P is $G \setminus (P \setminus \partial P)$ and to simplify notation we denote it by P^{out} .

2.3 Additively Weighted Voronoi Diagrams

Let H be a directed planar graph with real edge-lengths, and no negative-length cycles. Assume that all faces of H are triangles except, perhaps, a single face h . Let S be the set of vertices that lie on h . The vertices of S are called *sites*. Each site $s \in S$ has a weight $\omega(s) \geq 0$ associated with it. The additively weighted distance between a site $s \in S$ and a vertex $v \in V$, denoted by $d^\omega(s, v)$ is defined as $\omega(s)$ plus the length of the s -to- v shortest path in H .

Definition 1. *The additively weighted Voronoi diagram of (S, ω) ($VD(S, \omega)$) within H is a partition of $V(H)$ into pairwise disjoint sets, one set $Vor(s)$ for each site $s \in S$. The set $Vor(s)$, which is called the Voronoi cell of s , contains all vertices in $V(H)$ that are closer (w.r.t. $d^\omega(\cdot, \cdot)$) to s than to any other site in S .*

There is a dual representation $VD^*(S, \omega)$ (or simply VD^*) of a Voronoi diagram $VD(S, \omega)$. Let H^* be the planar dual of H . Let VD_0^* be the subgraph of H^* consisting of the duals of edges uv of H such that u and v are in different Voronoi cells. Let VD_1^* be the graph obtained from VD_0^* by contracting edges incident to degree-2 vertices one after another until no degree-2 vertices

remain. The vertices of VD_1^* are called Voronoi vertices. A Voronoi vertex $f^* \neq h^*$ is dual to a face f such that the vertices of H incident to f belong to three different Voronoi cells. We call such a face *trichromatic*. Each Voronoi vertex f^* stores for each vertex u incident to f the site s such that $u \in \text{Vor}(s)$. Note that h^* (i.e. the dual vertex corresponding to the face h to which all the sites are incident) is a Voronoi vertex. Each face of VD_1^* corresponds to a cell $\text{Vor}(s)$. Hence there are at most $|S|$ faces in VD_1^* . By sparsity of planar graphs, and by the fact that the minimum degree of a vertex in VD_1^* is 3, the complexity (i.e., the number of vertices, edges and faces) of VD_1^* is $O(|S|)$. Finally, we define VD^* to be the graph obtained from VD_1^* after replacing the node h^* by multiple copies, one for each occurrence of h as an endpoint of an edge in VD_1^* . It was shown in [23] that VD^* is a forest, and that, if all vertices of h are sites and if the additive weights are such that each site is in its own nonempty Voronoi cell, then VD^* is a ternary tree. See Fig. 2 (also used in [23]) for an illustration.

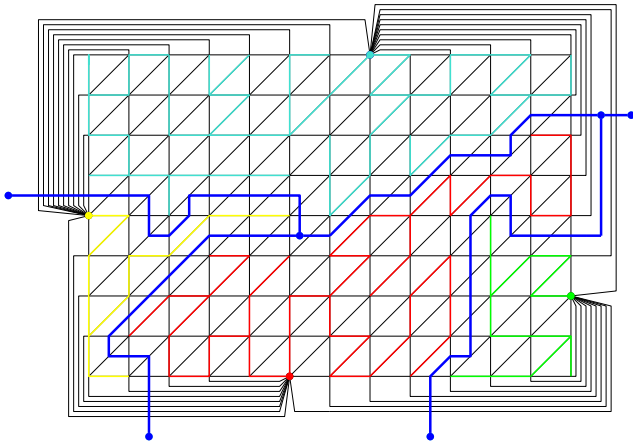


Figure 2: A planar graph (black edges) with four sites on the infinite face together with the dual Voronoi diagram VD^* (in blue). The sites are shown together with their corresponding shortest path trees (in turquoise, red, yellow, and green).

2.4 Point Location in Voronoi Diagrams

A *point location query* for a node v in a Voronoi diagram VD asks for the site s of VD such that $v \in \text{Vor}(s)$ and for the additive distance from s to v . Gawrychowski et al. [23] described a data structure supporting efficient point location, which is captured by the following theorem.

Theorem 2 ([23]). *Given an $O(|H| \log |H|)$ -sized MSSP data structure for H with sources S , and some VD^* for H with sources S , after an $O(|S|)$ -time preprocessing, point location queries can be answered in time $O(\log^2 |H|)$.*

We now describe this data structure. The data structure is essentially the same as in [23], but the presentation is a bit more modular. We will later adapt the implementation in Section 4.

Recall that H is triangulated (except the face h). The main idea is as follows. In order to find the Voronoi cell $\text{Vor}(s)$ to which a

query vertex v belongs, it suffices to identify an edge e^* of VD^* that is adjacent to $\text{Vor}(s)$. Given e^* we can simply check which of its two adjacent cells contains v by comparing the distances from the corresponding two sites to v . The point location structure is based on a *centroid decomposition* of the tree VD^* into connected subtrees, and on the ability to determine which of the subtrees is the one that contains the desired edge e^* .

The preprocessing consists of just computing a centroid decomposition of VD^* . A *centroid* of an n -node tree T is a node $u \in T$ such that removing u and replacing it with copies, one for each edge incident to u , results in a set of trees, each with at most $\frac{n+1}{2}$ edges. A centroid always exists in a tree with at least one edge. In every step of the centroid decomposition of VD^* , we work with a connected subtree T^* of VD^* . Recall that there are no nodes of degree 2 in VD^* . If there are no nodes of degree 3, then T^* consists of a single edge of VD^* , and the decomposition terminates. Otherwise, we choose a centroid c^* , and partition T^* into the three subtrees T_0^*, T_1^*, T_2^* obtained by splitting c^* into three copies, one for each edge incident to c^* . Clearly, the depth of the recursive decomposition is $O(\log |S|)$. The decomposition can be computed in $O(|S|)$ time [6] and can be represented as a ternary tree which we call the *centroid decomposition tree*, in $O(|S|)$ space. Each non-leaf node of the centroid decomposition tree corresponds to a centroid vertex c^* , which is stored explicitly. We will refer to nodes of the centroid decomposition tree by their associated centroid. Each node also corresponds implicitly to the subtree of VD^* of which c^* is the centroid. The leaves of the centroid decomposition tree correspond to single edges of VD^* , which are stored explicitly.

Point location queries for a vertex v in the Voronoi diagram VD are answered by employing procedure `POINTLOCATE`(VD^*, v) (Algorithm 1), which takes as input the Voronoi diagram, represented by the centroid decomposition of VD^* , and the vertex v . This in turn calls the recursive procedure `HANDLECENTROID`(VD^*, c^*, v) (Algorithm 2), where c^* is the root centroid in the centroid decomposition tree of VD^* .

Algorithm 1 `POINTLOCATE`(VD^*, v)

Input: The centroid decomposition of the dual representation VD^* of a Voronoi diagram VD and a vertex v .

Output: The site s of VD such that $v \in \text{Vor}(s)$ and the additive distance from s to v .

- 1: $c^* \leftarrow$ root centroid in the centroid decomposition of VD^*
 - 2: **return** `HANDLECENTROID`(VD^*, c^*, v)
-

We now describe the procedure `HANDLECENTROID`. It gets as input the Voronoi diagram, represented by its centroid decomposition tree, the centroid node c^* in the centroid decomposition tree that should be processed, and the vertex v to be located. It returns the site s such that $v \in \text{Vor}(s)$, and the additive distance to v . If c^* is a leaf of the centroid decomposition, then its corresponding subtree of VD^* is the single edge e^* we are looking for (Lines 1-6). Otherwise, c^* is a non-leaf node of the centroid decomposition tree, so it corresponds to a node in the tree VD^* , which is also a vertex of the dual H^* of H . Thus, c is a face of H . Let the vertices of c be y_1, y_2 and y_3 . We obtain s_i , the site such that $y_i \in \text{Vor}(s_i)$, from the representation of VD^* (Line 7). (Recall that $s_i \neq s_k$ if $i \neq k$ since

Algorithm 2 HANDLECENTROID(VD^*, c^*, v)

Input: The centroid decomposition of the dual representation VD^* of a Voronoi diagram, a centroid c^* and a vertex v such that the subtree of which c^* is a centroid contains an edge e^* incident to the Voronoi cell to which v belongs.

Output: The site s of VD such that $v \in \text{Vor}(s)$ and the additive distance from s to v .

```

1: if  $c^*$  is a leaf then
2:    $s_1, s_2 \leftarrow$  sites corresponding to  $c^*$ 
3:   for  $k = 1, 2$  do
4:      $d_k \leftarrow \omega(s_k) + d_H(s_k, v)$ 
5:    $j \leftarrow \text{argmin}_k(d_k)$ 
6:   return  $(s_j, d_j)$ 
7:  $s_1, s_2, s_3 \leftarrow$  sites corresponding to  $c^*$ 
8: for  $k = 1, 2, 3$  do
9:    $d_k \leftarrow \omega(s_k) + d_H(s_k, v)$ 
10:  $j \leftarrow \text{argmin}_k(d_k)$ 
11:  $a \leftarrow$  left/right/ancestor relationship of  $v$  to  $v_j^c$  in the shortest
    path tree of  $s_j$  in  $H$ 
12: if  $v$  is an ancestor of  $v_j^c$  then
13:   return  $(s_j, d_j)$ 
14: else
15:    $c^* \leftarrow \text{NEXTCENTROID}(c^*, a)$ 
16: return HANDLECENTROID( $\text{VD}^*, c^*, v$ )
    
```

c is a trichromatic face.) Next, for each i , we retrieve the additive distance from s_i to v . Let s_j be the site among them with minimum additive distance to v .

If v is an ancestor of the node y_j in the shortest path tree rooted at s_j , then $v \in \text{Vor}(s_j)$ and we are done. Otherwise, handling c consists of finding the child c'^* of c^* in the centroid decomposition tree whose corresponding subtree contains e^* .

Definition 3. For a tree T and two nodes x and y , such that none is an ancestor of the other, we say that x is left of y if the preorder number of x is smaller than the preorder number of y . Otherwise, we say that x is right of y .

For every face f , whose vertices are x_1, x_2, x_3 , let v_j^f for $j = 1, 2, 3$ be artificial vertices embedded in f , each with a single zero-length incident edge (x_j, v_j^f) . It is shown in [23] that identifying c'^* amounts to determining whether v is left or right of v_j^c in the shortest path tree rooted at s_j . We call the procedure that computes c'^* from c^* NEXTCENTROID. Having found the next centroid in Line 15, HANDLECENTROID moves on to handle it recursively.

The efficiency of procedure HANDLECENTROID depends on the time to compute distances in H (Lines 4 and 9) and the left/right/ancestor relationship (Line 11). Given an MSSP data structure for H , with sources S , each of these operations can be performed in time $O(\log |H|)$ and hence Theorem 2 follows.

3 AN $\tilde{O}(n^{4/3})$ -SPACE $O(\log^2 n)$ -QUERY ORACLE

In this section we present a simple oracle that illustrates the use of external Voronoi diagrams. Understanding this oracle may be

helpful, but is not crucial, to understanding the main result in Section 4. We start by stating the following result from [23] that we use in the oracle presented in this section.

Theorem 4 ([23]). For a planar graph G of size n , there is an $O(n^{3/2})$ -sized data structure that answers distance queries in time $O(\log n)$.

For clarity of presentation, we first describe our oracle under the assumption that the boundary vertices of each piece P in the r -division of the graph lie on a single hole and that each such hole is a simple cycle. Multiple holes and non-simple cycles do not pose any significant complications; we explain how to treat pieces with multiple holes that are not necessarily simple cycles, separately.

3.1 Data Structure

We obtain an r -division for $r = n^{2/3} \sqrt{\log n}$. The data structure consists of the following for each piece P of the r -division:

- (1) The $O(|P|^{3/2})$ -space, $O(\log |P|)$ -query-time distance oracle of Theorem 4. These occupy $O(n\sqrt{r})$ space in total.
- (2) Two MSSP data structures, one for P and one for P^{out} , both with sources the nodes of ∂P . The MSSP data structure for P requires space $O(r \log r)$, while the one for P^{out} requires space $O(n \log n)$. The total space required for the MSSP data structures is $O(\frac{n^2}{r} \log n)$, since there are $O(\frac{n}{r})$ pieces.
- (3) For each node u of P :
 - $\text{VD}_{in}^*(u, P)$, the dual representation of the Voronoi diagram for P with sites the nodes of ∂P , and additive weights the distances from u to these nodes in G ;
 - $\text{VD}_{out}^*(u, P)$, the dual representation of the Voronoi diagram for P^{out} with sites the nodes of ∂P , and additive weights the distances from u to these nodes in G .
 The representation of each Voronoi diagram occupies $O(\sqrt{r})$ space and hence, since each vertex belongs to a constant number of pieces, all Voronoi diagrams require space $O(n\sqrt{r})$.

3.2 Query

We obtain a piece P of the r -division that contains u . Let us first suppose that $v \in P$. We have to consider both the case that the shortest u -to- v path crosses ∂P and the case that it does not. If it does cross, we retrieve this distance by performing a point location query for v in the Voronoi diagram $\text{VD}_{in}(u, P)$. If the shortest u -to- v path does not cross ∂P , the path lies entirely within P . We thus retrieve the distance by querying the exact distance oracle of Theorem 4 stored for P . The answer is the minimum of the two returned distances. This requires $O(\log^2 n)$ time by Theorems 2 and 4. Else, $v \notin P$ and the shortest path from u to v must cross ∂P . The answer can be thus obtained by a point location query for v in the Voronoi diagram $\text{VD}_{out}(u, P)$ in time $O(\log^2 n)$ by Theorem 2. The pseudocode of the query algorithm is presented below as procedure SIMPLEDIST(u, v) (Algorithm 3).

Algorithm 3 SIMPLEDIST(u, v)

Input: Two nodes u and v .

Output: $d_G(u, v)$.

```

1:  $P \leftarrow$  a piece of the  $r$ -division containing  $u$ 
2: if  $v \in P$  then
3:    $d_1 \leftarrow d_P(u, v)$ 
4:    $d_2 \leftarrow \text{POINTLOCATE}(\text{VD}_{in}^*(u, P), v)$ 
5:   return  $\min(d_1, d_2)$ 
6: else
7:   return  $\text{POINTLOCATE}(\text{VD}_{out}^*(u, P), v)$ 
    
```

3.3 Dealing with Holes

The data structure has to be modified as follows.

- (2) For each hole h of P , two MSSP data structures, one for P and one for $P^{h, out}$, both with sources the nodes of ∂P that lie on h .
- (3) For each node u of P , for each hole h of P :
 - $\text{VD}_{in}^*(u, P, h)$, the dual representation of the Voronoi diagram for P with sites the nodes of ∂P that lie on h , and additive weights the distances from u to these nodes in G ;
 - $\text{VD}_{out}^*(u, P, h)$, the dual representation of the Voronoi diagram for $P^{h, out}$ with sites the nodes of ∂P that lie on h , and additive weights the distances from u to these nodes in G .

As for the query, if $v \in P$ we have to perform a point location query in $\text{VD}_{in}(u, P, h)$ for each hole h of P . Else $v \notin P$ and we have to perform a point location query in $\text{VD}_{out}(u, P, h)$ for the hole h of P such that $v \in P^{h, out}$. We can afford to store the required information to identify this hole explicitly in balanced search trees.

We thus obtain the following result.

Theorem 5. *For a planar graph G of size n , there is an $O(n^{4/3} \sqrt{\log n})$ -sized data structure that answers distance queries in time $O(\log^2 n)$.*

4 AN ORACLE WITH ALMOST OPTIMAL TRADEOFFS

In this section we describe how to obtain an oracle with almost optimal space to query-time tradeoffs. Consider the oracle in the previous section. The size of the representation we store for each Voronoi diagram is proportional to the number of sites, while the size of an MSSP data structure is roughly proportional to the size of the graph in which the Voronoi diagram is defined. Thus, the MSSP data structures that we store for the outside of pieces of the r -division are the reason that the oracle in the previous section requires $\Omega(n^{4/3})$ space. However, storing the Voronoi diagrams for the outside of each piece is not a problem. For instance, if we could somehow afford to store MSSP data structures for P and P^{out} of each piece P of an n^ϵ -division using just $O(n^\epsilon)$ space, then plugging $r = n^\epsilon$ into the data structure from the previous sections would yield an oracle with space $O(n^{1+\epsilon})$ and query-time $O(\log^2 n)$.

We cannot hope to have such a compact MSSP representation. However, we can use recursion to get around this difficulty. We compute a recursive r -division, represented by a recursive r -division tree \mathcal{T} . We store, for each piece $P \in \mathcal{T}$, the Voronoi diagram for P^{out} . However, instead of storing the costly MSSP for the entire

P^{out} , we store the MSSP data structure (and some additional information) just for the portion of P^{out} that belongs to the parent Q of P in \mathcal{T} . Roughly speaking, when we need to perform point location on a vertex of P^{out} that belongs to Q we can use this MSSP information. When we need to perform point location on a vertex of P^{out} that does not belong to Q (i.e., it is also in Q^{out}), we recursively invoke the point location mechanism for Q^{out} .

We next describe the details. For clarity of presentation, we assume that the boundary vertices of each piece P in \mathcal{T} lie on a single hole which is a simple cycle. We later explain how to remove these assumptions. In what follows, if a vertex in a Voronoi diagram can be assigned to more than one Voronoi cell, we assign it to the Voronoi cell of the site with largest additive weight. In other words, since we define the additive weights as distances from some vertex u , and since shortest paths are unique, we assign each vertex v to the Voronoi cell of the last site on the shortest u -to- v path. In particular, this implies that there are no empty Voronoi cells as every site belongs to its own cell. Thus VD^* is a ternary tree (see [23]). We can make such an assignment by perturbing the additive weights to slightly favor sites with larger distances from u at the time that the Voronoi diagram is constructed.

4.1 Data Structure

Consider a recursive (r_m, \dots, r_1) -division of G for some $n = r_m > \dots > r_1 = O(1)$ to be specified later. Recall that our convention is that the r_m -division consists of G itself. For convenience, we define each vertex v to be a boundary vertex of a singleton piece at level 0 of the recursive division. Denote the set of pieces of the r_i -division by \mathcal{R}_i . Let \mathcal{T} denote the tree representing this recursive (r_m, \dots, r_0) -division (each singleton piece v at level 0 is attached as the child of an arbitrary piece P at level 1 such that $v \in P$).

We will handle distance queries between vertices u, v that have the same parent P in \mathcal{R}_1 separately by storing these distances explicitly (this takes $O(n)$ space because pieces at level 1 have constant size).

The oracle consists of the following for each $0 \leq i \leq m-1$, for each piece $R \in \mathcal{R}_i$ whose parent in \mathcal{T} is $Q \in \mathcal{R}_{i+1}$:

- (1) If $i > 0$, two MSSP data structures for $Q \setminus (R \setminus \partial R)$, with sources ∂R and ∂Q , respectively.
- (2) If $i < m-1$, for each vertex $u \in \partial R$:
 - $\text{VD}_{near}^*(u, Q)$, the dual representation of the Voronoi diagram for $Q \setminus (R \setminus \partial R)$ with sites the nodes of ∂Q , and additive weights the distances from u to these nodes in R^{out} ;²
 - $\text{VD}_{far}^*(u, Q)$, the dual representation of the Voronoi diagram for Q^{out} with sites the nodes of ∂Q , and additive weights the distances from u to these nodes in R^{out} ;
 - if $i > 0$, the coarse tree T_u^R , which is the tree obtained from the shortest path tree rooted at u in R^{out} (the fine tree) by recursively contracting a rootmost edge that is incident to a vertex not in $\partial Q \cup \partial R$. Note that the left/right/ancestor relationship between vertices of $\partial Q \cup \partial R$ in the coarse tree is the same as in the fine tree. Also note that each (coarse) edge in T_u^R originates from a contracted path in the fine

²In the query algorithm we will use $\text{VD}_{near}^*(u, Q)$ (and $\text{VD}_{far}^*(u, Q)$ which will be defined immediately) only to compute distances in R^{out} . Hence, the additive weights are with respect to R^{out} , not to G .

tree. We preprocess the coarse tree T_u^R in time proportional to its size to allow for $O(1)$ -time lowest common ancestor (LCA)³ and level ancestor⁴ queries [3, 4]. In addition, for every (coarse) edge of T_u^R , we store the first and last edges of the underlying path in the fine tree. We also store the preorder numbers of the vertices of T_u^R .

Space. The space required to store the described data structure is $O(n \sum_i \frac{r_{i+1}}{r_i} \log r_{i+1})$. Part 1 of the data structure: At the r_i -division, we have $O(\frac{n}{r_i})$ pieces and for each of them, we store two MSSP data structures, each of size $O(r_{i+1} \log r_{i+1})$. Thus, the total space required for the MSSP data structures is $O(n \sum_i \frac{r_{i+1}}{r_i} \log r_{i+1})$. Part 2 of the data structure: At the r_i -division, we store, for each of the $O(\frac{n}{\sqrt{r_i}})$ boundary nodes, the representation of two Voronoi diagrams and a coarse tree, each of size $O(\sqrt{r_{i+1}})$. The space for this part is thus $O(n \sum_i (\frac{r_{i+1}}{r_i})^{1/2})$.

4.2 Query

Upon a query $d_G(u, v)$ for the distance between vertices u and v in G , we pick the singleton piece $Q_0 = \{u\} \in \mathcal{R}_0$ and call the procedure $\text{DIST}(u, v, Q_0)$ presented below. In what follows we denote by Q_i the ancestor of Q_0 in \mathcal{T} that is in \mathcal{R}_j .

The procedure $\text{DIST}(u, v, Q_i)$ (Algorithm 4) gets as input a piece Q_i , a source vertex $u \in \partial Q_i$ and a destination vertex $v \in Q_i^{\text{out}}$. The output of $\text{DIST}(u, v, Q_i)$ is a tuple (β, d) , where d is the distance from u to v in Q_i^{out} and β is a sequence of certain boundary vertices of the recursive division that occur along the recursive query. (Note that $Q_0^{\text{out}} = G$ and thus $\text{DIST}(u, v, Q_0)$ returns $d_G(u, v)$.) Let i^* be the smallest index such that $v \in Q_{i^*}$. For every $1 \leq i < i^*$, the i 'th element in the list β is the last boundary vertex on the shortest u -to- v path that belongs to ∂Q_i . As we shall see, the list β enables us to obtain sufficient information to determine the required left/right/ancestor relationships during the recursion. $\text{DIST}(u, v, Q_i)$ works as follows.

- (1) If $v \in Q_{i+1} \setminus Q_i$, it computes the required distance by considering the minimum of the distance from u to v returned by a query in the MSSP data structure for $Q_{i+1} \setminus (Q_i \setminus \partial Q_i)$ with sources ∂Q_i and the distance returned by a (vanilla) point location query for v in $\text{VD}_{\text{near}}(u, Q_{i+1})$ using the procedure POINTLOCATE . The first query covers the case where the shortest path from u to v lies entirely within Q_{i+1} , while the second one covers the complementary case. The time required in this case is $O(\log^2 n)$.
- (2) If $v \notin Q_{i+1}$, DIST performs a recursive point location query on $\text{VD}_{\text{far}}(u, Q_{i+1})$ by calling the procedure $\text{MODIFIEDPOINTLOCATE}$.

Since DIST returns a list of boundary vertices as well as the distance, POINTLOCATE and HANDLECENTROID must pass and augment these lists as well. The pseudocode for $\text{MODIFIEDPOINTLOCATE}$ is identical to that of POINTLOCATE , except it calls $\text{MODIFIEDHANDLECENTROID}$ instead of HANDLECENTROID ; see Algorithm 5. In what follows, when any of the three procedures is called with respect

³An LCA query takes as input two nodes of a rooted tree and returns the deepest node of the tree that is an ancestor of both.

⁴A level ancestor query takes as input a node v at depth d and an integer $\ell \leq d$ and returns the ancestor of v that is at depth ℓ .

Algorithm 4 $\text{DIST}(u, v, Q_i)$

Input: A piece Q_i , a vertex $u \in \partial Q_i$ and a vertex $v \in Q_i^{\text{out}}$ (v may belong to Q_i if $i = m$).

Output: (β, d) , where β is a list of (some) boundary vertices on the u -to- v shortest path in Q_i^{out} , and d is the u -to- v distance in Q_i^{out} .

- 1: **if** $v \in Q_{i+1}$ **then**
 - 2: $(\beta, d) \leftarrow (\text{null}, \text{distance returned by querying the MSSP data structure for } Q_{i+1} \setminus (Q_i \setminus \partial Q_i))$
 - 3: $(\beta', d') \leftarrow \text{POINTLOCATE}(\text{VD}_{\text{near}}^*(u, Q_{i+1}), v)$
 - 4: **if** $d' < d$ **then**
 - 5: $(\beta, d) \leftarrow (\beta', d')$
 - 6: **else if** $i = m - 2$ **then**
 - 7: $(\beta, d) \leftarrow \text{POINTLOCATE}(\text{VD}_{\text{far}}^*(u, Q_{m-1}), v)$
 - 8: **else**
 - 9: $(\beta, d) \leftarrow \text{MODIFIEDPOINTLOCATE}(\text{VD}_{\text{far}}^*(u, Q_{i+1}), v, i + 1)$
 - 10: **return** (β, d)
-

Algorithm 5 $\text{MODIFIEDPOINTLOCATE}(\text{VD}^*, v, i)$

Input: The centroid decomposition of the dual representation VD^* of a Voronoi diagram VD for Q_i^{out} with sites ∂Q_i and a vertex $v \in Q_i^{\text{out}}$.

Output: The site s of VD such that $v \in \text{Vor}(s)$ and the additive distance from s to v .

- 1: $c^* \leftarrow$ root centroid in the centroid decomposition of VD^*
 - 2: **return** $\text{MODIFIEDHANDLECENTROID}(\text{VD}^*, c^*, v, i)$
-

to a piece Q_i , we refer to this as an invocation of this procedure at level i .

The pseudocode of $\text{MODIFIEDHANDLECENTROID}$ (Algorithm 6) is similar to that of the procedure HANDLECENTROID , except that when the site s such that $v \in \text{Vor}(s)$ is found, s is prepended to the list. A more significant change in $\text{MODIFIEDHANDLECENTROID}$ stems from the fact that, since we no longer have an MSSP data structure for all of Q_i^{out} , we use recursive calls to DIST to obtain the distances from the sites s_k to v in Q_i^{out} in Lines 4 and 9. We highlight these changes in red in the pseudocode provided below (Algorithm 6). We next discuss how to determine the required left/right/ancestor relationships in VD_{far} (Line 12) in the absence of MSSP information for the entire Q_i^{out} .

4.2.1 Left/Right/Ancessor Relationships in VD_{far} . Let s_j be the site among the three sites corresponding to a centroid c^* such that v is closest to s_j with respect to the additive distances (see Lines 7–10 of Algorithm 6). Recall that y_j is the vertex of the centroid face c that belongs to $\text{Vor}(s_j)$ and that v_j^c is the artificial vertex connected to y_j and embedded inside the face c . In Line 12 of $\text{MODIFIEDHANDLECENTROID}$ we have to determine whether v is an ancestor of y_j in the shortest path tree rooted at s_j in Q_i^{out} , and if not, whether the s_j -to- v path is left or right of the s_j -to- v_j^c path. To avoid clutter, we will omit the subscript j in the following, and refer to s_j as s , to v_j^c as v^c , and to the sequence of boundary vertices β_j returned by the recursion as β . To infer the relationship between the two paths we use the sites (boundary vertices) stored in the list

Algorithm 6 MODIFIEDHANDLECENTROID(VD^*, c^*, v, i)

Input: The centroid decomposition of the dual representation VD^* of a Voronoi diagram VD for Q_i^{out} with sites ∂Q_i , a centroid c^* and a vertex $v \in Q_i^{out}$ such that the subtree of which c^* is a centroid contains an edge e^* incident to the Voronoi cell to which v belongs.

Output: The site s of VD such that $v \in \text{Vor}(s)$ and the additive distance from s to v .

```

1: if  $c^*$  is a leaf then
2:    $s_1, s_2 \leftarrow$  sites corresponding to  $c^*$ 
3:   for  $k = 1, 2$  do
4:      $(\beta_k, d_k) \leftarrow \text{DIST}(s_k, v, Q_i)$ 
5:      $j \leftarrow \text{argmin}_k(\omega(s_k) + d_k)$ 
6:     return ( $\text{prepend}(s_j, \beta_j)$ ),  $\omega(s_j) + d_j$ )
7:  $s_1, s_2, s_3 \leftarrow$  sites corresponding to  $c^*$ 
8: for  $k = 1, 2, 3$  do
9:    $(\beta_k, d_k) \leftarrow \text{DIST}(s_k, v, Q_i)$ 
10:  $j \leftarrow \text{argmin}_k(\omega(s_k) + d_k)$ 
11:  $(\gamma, d) \leftarrow \text{DIST}(s_j, v_j^c, Q_i)$ 
12:  $a \leftarrow$  left/right/ancestor relationship of  $v$  to  $v_j^c$  in the shortest
    path tree of  $s_j$  in  $Q_i^{out}$  using  $\beta$  and  $\gamma$ 
13: if  $v$  is an ancestor of  $v_j^c$  then
14:   return ( $\text{prepend}(s_j, \beta_j)$ ),  $\omega(s_j) + d_j$ )
15: else
16:    $c^* \leftarrow \text{NEXTCENTROID}(c^*, a)$ 
17:   return MODIFIEDHANDLECENTROID( $VD^*, c^*, v, i$ )
    
```

β . We prepend s to β , and, if v is not already the last element of β , we append v to β and use a flag to denote that v was appended.

To be able to compare the s -to- v path with the s -to- v^c path, we perform another recursive call $\text{DIST}(s, v^c, Q_i)$ in Line 11. Let γ be the list of sites returned by this call. As above, we prepend s to γ , and append v^c if it is not already the last element. The intuition is that the lists β and γ are a coarse representation of the s -to- v and s -to- v^c shortest paths, and that we can use this coarse representation to roughly locate the vertex where the two paths diverge. The left/right relationship between the two paths is determined solely by the left/right relationship between the two paths at that vertex (the divergence point). We can use the local coarse tree information or the local MSSP information to infer this relationship.

Recall that s is a boundary vertex of the piece Q_i at level i of the recursive division. More generally, for any $k \geq 0$, $\beta[k]$ is a boundary vertex of Q_{i+k} (except, possibly, when $\beta[k] = v$). To avoid this shift between the index of a site in the list and the level of the corresponding piece in the recursive r -division, we prepend i empty elements to both lists, so that now, for any $k \geq i$, $\beta[k]$ is a boundary vertex of Q_k . Let k be the largest integer such that $\beta[k] = \gamma[k]$. Note that k exists because s is the first vertex in both β and γ .

Observation 6. *The restriction of the shortest path from $\beta[i]$ to $\beta[i+1]$ in Q_i^{out} to the nodes of $\partial Q_i \cup \partial Q_{i+1}$ is identical to the path from $\beta[i]$ to $\beta[i+1]$ in the coarse tree $T_{\beta[i]}^{Q_i}$.*

We next analyze the different cases that might arise and describe how to correctly infer the left/right relationship in each case. An illustration is provided as Fig. 3.

- (1) $(\beta[k+1] = v \notin \partial Q_{k+1})$ or $(\gamma[k+1] = c \notin \partial Q_{k+1})$. This corresponds to the case that in at least one of the lists there are no boundary vertices after the divergence point. We have a few cases depending on whether $\beta[k+1]$ and $\gamma[k+1]$ are in Q_k or not.
 - a. If $\beta[k+1], \gamma[k+1] \in Q_k$, then it must be that $\beta[k+1] = v$ and $\gamma[k+1] = c$ were appended to the sequences manually. In this case we can query the MSSP data structure for $Q_k \setminus (Q_{k-1} \setminus \partial Q_{k-1})$ with sources the boundary vertices of ∂Q_k to determine the left/right/ancestor relationship.
 - b. If $\beta[k+1], \gamma[k+1] \notin Q_k$, we can use the MSSP data structure stored for $Q_{k+1} \setminus (Q_k \setminus \partial Q_k)$ with sources the nodes of ∂Q_k . Let us assume without loss of generality that $\beta[k+1] = v \notin \partial Q_{k+1}$. We then define w to be $\gamma[k+1]$ if $\gamma[k+1] = c \notin \partial Q_{k+1}$ or the child w of the root $\beta[k]$ of the coarse $T_{\beta[k]}^{Q_k}$ that is an ancestor of $\gamma[k+1]$ otherwise. (In the latter case we can compute w in $O(1)$ time with a level ancestor query.) We then query the MSSP data structure for the relation between v and the w . Note that the relation between the s -to- v and s -to- c paths is the same as the relation between the s -to- v and s -to- w paths.
 - c. Else, one of $\beta[k+1], \gamma[k+1]$ is in Q_k and the other is not. Let us assume without loss of generality that $\beta[k+1] = v \in Q_k$. We can infer the left/right relation by looking at the circular order of the following edges: (i) the last fine edge in the root-to- $\beta[k]$ path in $T_{\beta[k-1]}^{Q_{k-1}}$, (ii) the first edge in the $\beta[k]$ -to- $\beta[k+1]$ path, and (iii) the first edge in the $\beta[k]$ -to- w path, where w is defined as in Case 1(b). Edge (i) is stored in $T_{\beta[k-1]}^{Q_{k-1}}$ —note that $\beta[k-1]$ exists since $\beta[k+1] \in Q_k$. Edge (ii) can be retrieved from the MSSP data structure for $Q_k \setminus (Q_{k-1} \setminus \partial Q_{k-1})$ with sources the boundary vertices of ∂Q_k , and edge (iii) can be retrieved from the MSSP data structure for $Q_{k+1} \setminus (Q_k \setminus \partial Q_k)$ with sources the boundary vertices of ∂Q_k .
- (2) $\beta[k+1], \gamma[k+1] \in \partial Q_{k+1}$. We first compute the LCA x of $\beta[k+1]$ and $\gamma[k+1]$ in $T_{\beta[k]}^{Q_k}$.
 - a. If neither of $\beta[k+1], \gamma[k+1]$ is an ancestor of the other in $T_{\beta[k]}^{Q_k}$, we are done by utilising the preorder numbers stored in $T_{\beta[k]}^{Q_k}$.
 - b. Else, x is one of $\beta[k+1], \gamma[k+1]$. We can assume without loss of generality that $x = \beta[k+1]$. Using a level ancestor query we find the child z of $\beta[k+1]$ in $T_{\beta[k]}^{Q_k}$ that is a (not necessarily strict) ancestor of $\gamma[k+1]$. The x -to- z shortest path in Q_k^{out} is internally disjoint from ∂Q_{k+1} ; i.e. it starts and ends at ∂Q_{k+1} , but either entirely lies inside Q_{k+1} or entirely outside Q_{k+1} . If $\beta[k+2] = v \notin \partial Q_{k+2}$ let w be v ; otherwise let w be the child of the root in $T_{\beta[k+1]}^{Q_{k+1}}$ that is an ancestor of $\beta[k+2]$. The x -to- w shortest path also lies either entirely inside or entirely outside Q_{k+1} . It remains

to determine the relationship of the x -to- w and the x -to- z paths.

- i. If the shortest x -to- z path and the shortest x -to- w path lie on different sides of ∂Q_{k+1} , we can infer the left/right relation by looking at the circular order of the following edges: (i) the last edge on the coarse edge from $\beta[k]$ to x in $T_{\beta[k]}^{Q_k}$, (ii) the first edge on the x -to- z shortest path in Q_k^{out} , and (iii) the first edge on the x -to- w shortest path in Q_k^{out} . Edges (ii) and (iii) can be retrieved from the MSSP data structure.
- ii. Else, the x -to- z and the x -to- w shortest paths lie on the same side, and we compute the required relation using the MSSP data structure for this side, with sources the sites of ∂Q_{k+1} .

Query Correctness. Let us consider an invocation of `DIST` at level i . We first show that if such a query is answered without invoking `MODIFIEDPOINTLOCATE` (Line 9 of Algorithm 4) then the distance it returns is correct. Indeed, let us recall that the input to `DIST` is a piece Q_i , a vertex $u \in \partial Q_i$ and a vertex $v \in Q_i^{out}$ and the desired output is the u -to- v distance in Q_i^{out} (along with some vertices on the u -to- v shortest path in Q_i^{out}). Let us now look at the cases in which `DIST` does not call `MODIFIEDPOINTLOCATE` (Lines 1-7). If $v \in Q_{i+1}$ then we have two cases: either the shortest u -to- v path in Q_i^{out} crosses ∂Q_{i+1} or it does not. In the first case the answer is obtained by a point location query in the graph $Q_{i+1} \setminus (Q_i \setminus \partial Q_i)$ with sites the vertices of ∂Q_{i+1} and additive weights the distances from u to these vertices in Q_i^{out} , while in the second case by a query on the MSSP data structure for this graph with sources the vertices of ∂Q_i . Else, if $v \notin Q_{i+1}$ and $i = m-2$, the shortest path must cross ∂Q_{i+1} . We can thus obtain the distance by performing a point location query in the Voronoi diagram for Q_{m-1}^{out} with sites the vertices of ∂Q_{i+1} and additive weights the distances to these vertices from u in Q_{m-2}^{out} . Note that in both the cases that we perform point location queries, we have MSSP information for the underlying graph on which the Voronoi diagram is defined and hence the correctness of the answer follows from the correctness of the point location mechanism designed in [23].

We now observe that if invocations of `DIST` at level $i+1$ return the correct answer, then so does the invocation of `MODIFIEDPOINTLOCATE` at line 9 of the invocation of `DIST` at level i , since the only other requirement for the correctness of `MODIFIEDPOINTLOCATE` is the ability to resolve the left/right/ancestor relationship; this follows from inspecting the cases analyzed above. It follows that invocations of `DIST` at level i also return the correct answer. The correctness of the query algorithm then follows by induction.

Query Time. We initially call `DIST` at level 0. A call to `DIST` at level i , either returns the answer in $O(\log^2 n)$ time by using MSSP information and/or invoking `POINTLOCATE` (Lines 1-7), or makes a call to `MODIFIEDPOINTLOCATE` at level $i+1$ (Line 9). For the latter to happen it must be that $i \leq m-3$. This call to `MODIFIEDPOINTLOCATE` at level $i+1$ invokes a call to `MODIFIEDHANDLECENTROID` at level $i+1$. Then, `MODIFIEDHANDLECENTROID` makes $O(\log r_{i+1})$ calls to `DIST` at level $i+1$, while going down the centroid decomposition of a Voronoi diagram with $O(\sqrt{r_{i+1}})$ sites.

Let $D[i]$ be the time required by a call to `DIST` at level i and $C[i, t]$ be the time required by a call to `MODIFIEDHANDLECENTROID` at

level i , where the size of the subtree of VD^* corresponding to the input centroid is t . We then have the following:

$$D[i] = \begin{cases} C[i+1, \sqrt{r_{i+1}}] + O(\log n + \log^2 r_{i+1}) & \text{If } i < m-2 \\ O(\log^2 n) & \text{If } i = m-2 \end{cases}$$

(The additive $O(\log n)$ factor comes from the possible existence of multiple holes, see Section 4.3.)

$$C[i, k] = \begin{cases} O(D[i] + \log n) + C[i, k/2] & \text{If } k > 1 \\ 2D[i] & \text{If } k = 1 \end{cases}$$

(The additive $O(\log n)$ factor comes from resolving the left/right/ancestor relationship.)

Hence, for $i < m-2$, $D[i] = O(\log r_{i+1}(D[i+1] + \log n))$. The total time complexity of the query is thus $D[0] = O(\log^2 n \prod_{i=1}^{m-3} c \log r_{i+1}) = O(c^m \log^{m-1} n)$, where $c > 0$ is a constant.

Retrieving the Shortest Path. First of all, let us recall that given an MSSP data structure for a planar graph H of size n with sources the nodes of a face f , we can construct the shortest path from a node x in f to any node y of H , backwards from y in $O(\log \log n)$ per reported edge (cf. [32]). We can easily retrieve one by one the edges of the coarse trees $T_{\beta[i]}^{Q_i}$ corresponding to paths whose concatenation would give us the shortest u -to- v path. The shortest path is in fact the concatenation of the $\beta[i]$ -to- $\beta[i+1]$ paths in $T_{\beta[i]}^{Q_i}$, plus perhaps an extra final leg, which lies entirely in $Q_j \setminus (Q_{j-1} \setminus \partial Q_{j-1})$ for some j , if v was appended to β . Our MSSP data structures for each i allow us to refine those coarse edges of $T_{\beta[i]}^{Q_i}$ when their underlying shortest path lies entirely within $Q_{i+1} \setminus (Q_i \setminus \partial Q_i)$. For an edge (x, y) of $T_{\beta[i]}^{Q_i}$ that corresponds to a path lying entirely in Q_{i+1}^{out} , we use the x -to- y path in $T_x^{Q_{i+1}}$, which we can refine using the MSSP information we have for $Q_{i+2} \setminus (Q_{i+1} \setminus \partial Q_{i+1})$. We then proceed inductively on coarse edges whose underlying path is entirely in Q_{i+2}^{out} . In each step we obtain and refine an edge from a node in ∂Q_i to a node in ∂Q_j , $j \in \{i-1, i, i+1\}$ using the local MSSP information. To summarize, we retrieve one by one the edges of the coarse trees $T_{\beta[i]}^{Q_i}$, corresponding to portions of the actual shortest path that lie entirely in $Q_i \setminus (Q_{i-1} \setminus \partial Q_{i-1})$ in $O(1)$ time each. We then use the local MSSP information to obtain the actual edges of each coarse edge in $O(\log \log n)$ time per actual edge.

We prove in Section 5 that the time bound for the construction of the data structure is $\tilde{O}(n \sum_i \frac{r_{i+1}}{\sqrt{r_i}})$. We summarize our main result in the following theorem.

Theorem 7. *Given a planar graph G of size n , for any decreasing sequence $n = r_m > r_{m-1} > \dots > r_1$, where r_1 is a sufficiently large constant and $m \geq 3$, there is a data structure of size $O(n \sum_i \frac{r_{i+1}}{r_i} \log r_{i+1})$ that answers distance queries in time $O(c^m \log^{m-1} n)$, where $c > 0$ is a constant. The construction time is $\tilde{O}(n \sum_i (\frac{r_{i+1}}{\sqrt{r_i}} + \frac{\sqrt{nr_{i+1}}}{\sqrt{r_i}}))$. The shortest path can be retrieved backwards from v to u , at a cost of $O(\log \log n)$ time per edge.*

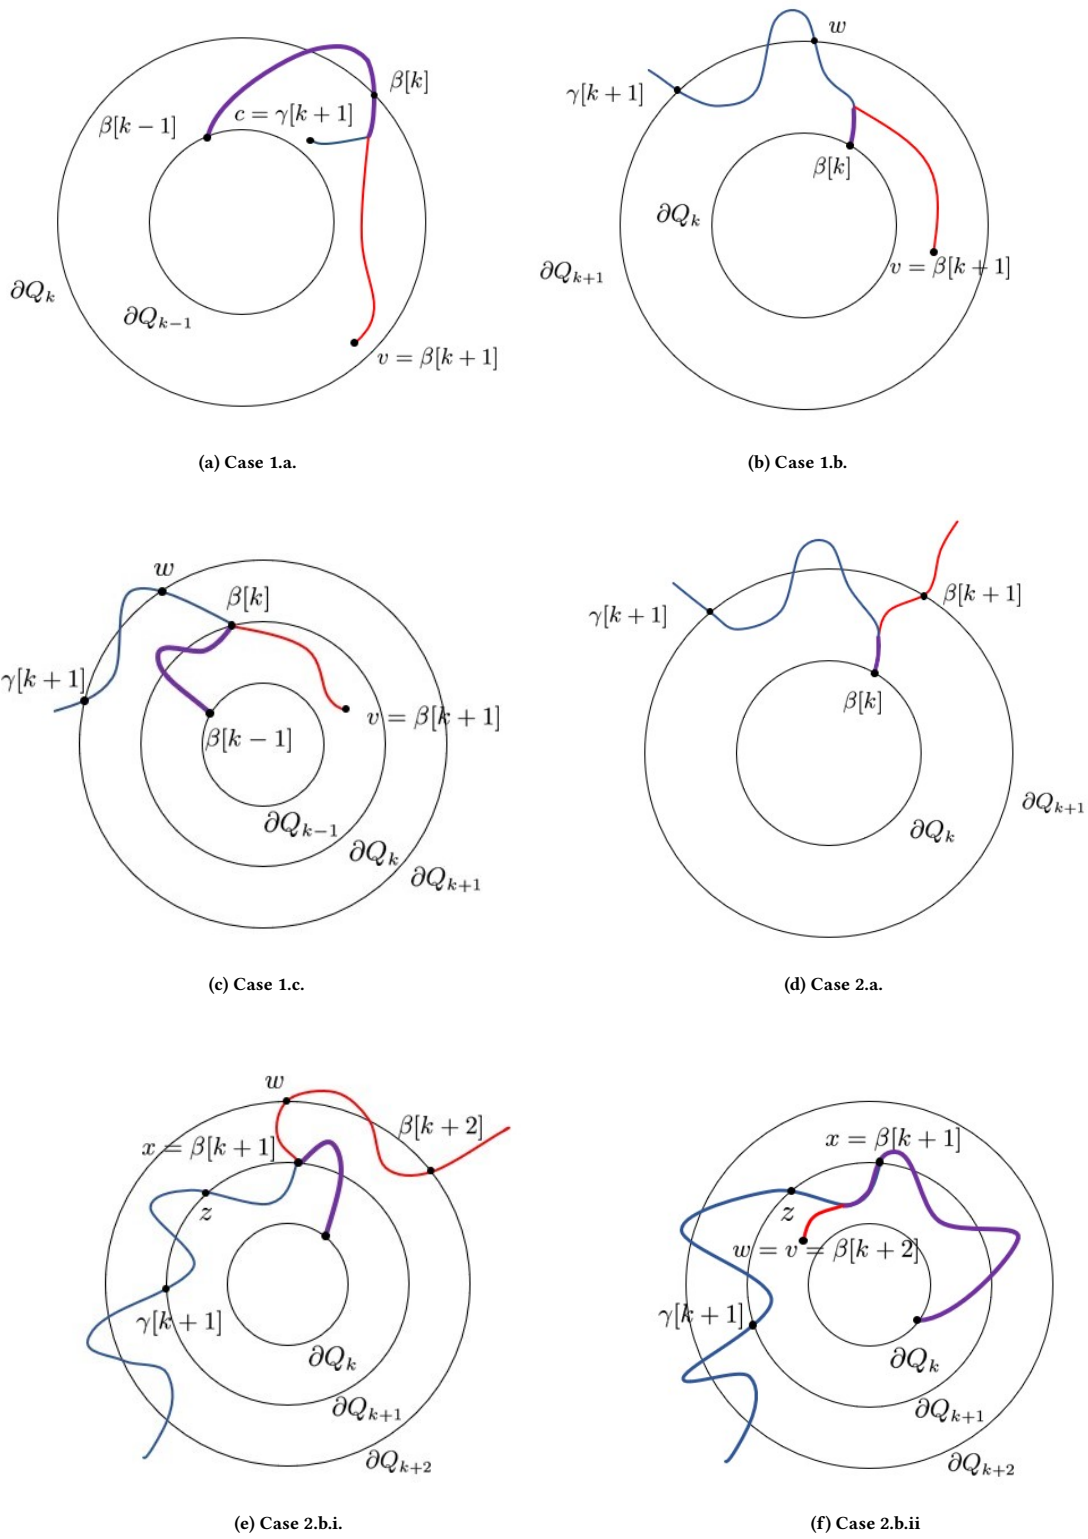


Figure 3: An illustration of the different cases that arise in determining the left/right/ancestor relationship of $\beta[k+1]$ and $\gamma[k+1]$ in the shortest path tree rooted at s in R^{out} . The s -to- v path is red, the s -to- c path is blue, and their common prefix is purple.

We next present three specific tradeoffs that follow from this theorem.

Corollary 8. *Given a planar graph G of size n , we can construct a distance oracle admitting any of the following (space, query-time) tradeoffs:*

- (a) $\langle O(n \log^{2+1/\epsilon} n), O(n^{2\epsilon}) \rangle$, for any constant $\epsilon > 0$;
- (b) $\langle n^{1+o(1)}, n^{o(1)} \rangle$;
- (c) $\langle O(n^{1+\epsilon} \log n), O(\log^{1/\epsilon} n) \rangle$, for any constant $1/2 \geq \epsilon > 0$.

The construction time is $\tilde{O}(n^{3/2} \log^{1/2\epsilon} n)$, $n^{3/2+o(1)}$ and $\tilde{O}(n^{3/2+\epsilon/2})$, respectively.

PROOF. Let ρ_i denote $\frac{r_{i+1}}{r_i}$; we will set all ρ_i to be equal to a value ρ . We get tradeoffs (a), (b) and (c) by setting $\rho = \log^{1/\epsilon} n$, $\rho = 2^{\sqrt{\log n}}$ and $\rho = n^\epsilon$, respectively.

We provide the arithmetic below.

- (a) If we set $\rho = \log^{1/\epsilon} n$, then the depth of the recursive r -division is $m = \log_\rho n + 1 = \frac{\log n}{\log \rho} + 1 = \epsilon \log n / \log \log n + 1$. Then, the space required is $O(nm\rho \log n) = O(n \log^{2+1/\epsilon} n)$, and the query time is $O((c \log n)^m) = O(2^{\log(c \log n)\epsilon \log n / \log \log n}) = O(n^{2\epsilon})$.
- (b) If we set $\rho = 2^{\sqrt{\log n}}$, then the depth of the recursive r -division is $m = \log_\rho n + 1 = \frac{\log n}{\log \rho} + 1 = \sqrt{\log n} + 1$. Then, the space required is $O(nm\rho \log n) = n^{1+o(1)}$ and the query time is $O((c \log n)^m) = O((c \log n)^{\sqrt{\log n}}) = O(2^{2 \log(c \log n) \sqrt{\log n}}) = n^{o(1)}$.
- (c) If we set $\rho = n^\epsilon$, then the depth of the recursive r -division is $m = \log_\rho n + 1 = 1/\epsilon + 1$. Then, the space required is $O(nm\rho \log n) = O(n^{1+\epsilon} \log n)$ and the query time is $O((c \log n)^m) = O((c \log n)^{1/\epsilon}) = O(\log^{1/\epsilon} n)$.

As for the construction time, we note that $\frac{r_{i+1}}{\sqrt{r_i}} = \sqrt{r_{i+1}\rho}$ and $\frac{r_{i+1}^{1/4}}{\sqrt{r_i}} = \frac{\sqrt{\rho}}{r_i^{1/4}}$. Hence $\tilde{O}(n \sum_i (\frac{r_{i+1}}{\sqrt{r_i}} + \frac{\sqrt{n} r_{i+1}^{1/4}}{\sqrt{r_i}}))$ is $\tilde{O}(nm\sqrt{n\rho})$. The stated bounds follow. \square

4.3 Dealing with Multiple Holes

We now remove the assumption that all boundary vertices of each piece lie on a single hole.

4.3.1 Data Structure. The oracle consists of the following for each $0 \leq i \leq m-1$, for each piece $R \in \mathcal{R}_i$ whose parent in \mathcal{T} is $Q \in \mathcal{R}_{i+1}$:

- (1) If $i > 0$, for each pair of holes h of R and g of Q , such that g lies in $R^{h,out}$, two MSSP data structures for $R^{h,out} \cap Q$, one with sources the vertices of ∂R that lie on h , and one with sources the vertices of ∂Q that lie on g .
- (2) If $i < m-1$, for each boundary vertex u of R that lies on a hole h of R :
 - For each hole g of Q that lies in $R^{h,out}$:
 - $\text{VD}_{near}^*(u, Q, g)$, the dual representation of the Voronoi diagram for $R^{h,out} \cap Q$ with sites the nodes of ∂Q that lie on g , and additive weights the distances from u to these nodes in $R^{h,out}$;

- $\text{VD}_{far}^*(u, Q, g)$, the dual representation of the Voronoi diagram for $Q^{g,out}$ with sites the nodes of ∂Q that lie on g , and additive weights the distances from u to these nodes in $R^{h,out}$;
- If $i > 0$, the coarse tree $T_u^{R,h}$, which is the tree obtained from the shortest path tree rooted at u in $R^{h,out}$ by recursively contracting a rootmost edge that is incident to a vertex that is neither in h nor in ∂Q , preprocessed as in the single hole case.

4.3.2 Query. If $v \in Q_{i+1} \setminus Q_i$ then in Line 3 of `DIST`, we need to consider $O(1)$ Voronoi diagrams $\text{VD}_{near}^*(u, Q_{i+1}, g)$, one for each hole of Q_{i+1} , instead of just one. If not, we find the correct hole h of Q_{i+1} such that $v \in Q_{i+1}^{h,out}$, and recurse on $\text{VD}_{far}^*(u, Q_{i+1}, h)$.

We can find the correct hole by storing some information for each hole of each piece. It is not hard to see that each separator of the $O(\log n)$ ancestors of a piece $P \in \mathcal{A}$ lies in $P^{h,out}$ for some hole h of P . For each piece $P \in \mathcal{A}$, we store, for each separator of an ancestor of P in \mathcal{A} , the hole h of P such that $P^{h,out}$ contains that separator. Then, by performing an LCA query for the constant size piece of \mathcal{A} containing v and Q_{i+1} in \mathcal{A} we find the separator that separated v from Q_{i+1} and we can thus find the appropriate hole of Q_{i+1} in $O(1)$ time.

Obtaining the left/right/ancestor information is essentially the same as in the single hole case. There is one additional case that needs to be considered, in which the vertex at which the s -to- v path and the s -to- v^c path diverge is located inside a hole. In this case we may have to look into finer and finer coarse trees until we get to a level of resolution where we can either deduce the left/right relationship from the coarse tree at that level, or from the MSSP data structure at that level. This refining process has $O(\log n)$ levels and takes constant time per level, except for the last level where we use the MSSP data structure in $O(\log n)$ time as in the single hole case. Therefore, the asymptotic query time is unchanged.

As for retrieving the edges of the shortest path in $O(\log \log n)$ time each, the presence of multiple holes does not pose any further complications.

4.3.3 Dealing with Holes That Are Non-simple Cycles. Non-simple holes do not introduce any significant difficulties. The technique to handle non-simple holes is the one described in [27, Section 5.1, pp. 27]. We make an incision along the non-simple boundary of a hole to turn it into a simple cycle. Making this incision creates multiple copies of vertices that are visited multiple times by the non-simple cycle. However, the total number of copies is within a constant of the original number of boundary vertices along the hole. These copies do not pose any problems because, for our purposes, they can be treated as distinct sites of the Voronoi diagram.

5 CONSTRUCTION TIME

In this section we show how to construct our oracles in $\tilde{O}(n \sum_i \frac{r_{i+1}}{\sqrt{r_i}})$ time, as stated in Theorem 7. Before doing so, we give some preliminaries on dense distance graphs.

5.1 Dense Distance Graphs and FR-Dijkstra

The *dense distance graph* of a piece P , denoted DDG_P is a complete directed graph on the boundary vertices of P . Each edge (u, v) has length $d_P(u, v)$, equal to the length of the shortest u -to- v path in P . DDG_P can be computed in time $O((|\partial P|^2 + |P|) \log |P|)$ using the multiple source shortest paths (MSSP) algorithm [9, 31]. Over all pieces of the recursive decomposition this takes time $O(n \log^2 n)$ in total and requires space $O(n \log n)$. We refer to the aforementioned (standard) DDGs as *internal*; the *external* DDG of a piece P is the complete directed graph on the vertices of ∂P , with the edge (u, v) having length equal to $d_{P_{out}}(u, v)$. There is an efficient implementation of Dijkstra's algorithm (nicknamed FR-Dijkstra [21]) that runs on any union of DDGs. The algorithm exploits the fact that, due to planarity, a DDG's adjacency matrix M satisfies a Monge property (namely, for any $i \neq j$ we have that $M[i + 1, j] + M[i, j + 1] \leq M[i + 1, j + 1] + M[i, j]$). We next give a —convenient for our purposes— interface for FR-Dijkstra which was essentially proved in [21], with some additional components and details from [27, 39].

Theorem 9 ([21, 27, 39]). *A set of DDGs with $O(M)$ vertices in total (with multiplicities), each having at most m vertices, can be preprocessed in time and extra space $O(M \log m)$ in total. After this preprocessing, Dijkstra's algorithm can be run on the union of any subset of these DDGs with $O(N)$ vertices in total (with multiplicities) in time $O(N \log N \log m)$.*

5.2 Construction

For ease of presentation we present the construction algorithm assuming pieces have only one hole. Generalizing to multiple holes does not pose any new obstacles. Recall that for each piece R in the recursive decomposition whose parent in \mathcal{T} is Q we need to compute two MSSP data structures, and for each boundary vertex u of R we need to compute $VD_{near}^*(u, Q)$, $VD_{far}^*(u, Q)$, and the coarse tree T_u^R .

Computing the MSSP data structures takes nearly linear time in the total size of these data structures. To be able to compute the distances between vertices in the same piece of the r_1 division, additive distances for all Voronoi diagrams, and the coarse trees T_u^R , we compute the internal and external dense distance graphs of every piece in the recursive r -division. This can be done in nearly linear time [38]. We also compute the dense distance graph of $\partial R \cup \partial Q$. This is done by running MSSP twice in $Q \setminus (R \setminus \partial R)$, and then querying all pairwise distances. For a piece R of the r_i division, this takes $\tilde{O}(r_{i+1})$ time, so it takes a total of $\tilde{O}(\sum_i n \frac{r_{i+1}}{r_i})$ time overall.

We then compute the distances from each vertex in G to all other vertices in the same piece R of the r_1 -division in time $O(1)$, using Dijkstra on R and the external DDG of R . To compute the additive weights for the Voronoi diagrams $VD_{near}^*(u, Q)$ and $VD_{far}^*(u, Q)$, as well as the coarse tree T_u^R , we need to compute distances from u to $\partial Q \cup \partial R$ in R^{out} . This can be done in time $O(|\partial Q| \log^2 |\partial Q|)$ by running FR-Dijkstra from u on the union of the DDG of $Q \setminus (R \setminus \partial R)$ and the external DDG of Q . The total time to compute all such additive weights is thus $O(\sum_i \frac{n}{\sqrt{r_i}} \sqrt{r_{i+1}} \log^2 r_{i+1})$.

To compute each Voronoi diagram $VD_{near}^*(u, Q)$ we compute the primal Voronoi diagram $VD_{near}(u, Q)$ with a single-source shortest path computation in Q from an artificial super-source s connected to all vertices ∂Q with edges whose lengths are the additive weights. This can be done in linear $O(|Q|)$ time [25]. In $O(|Q|)$ time it is then easy to obtain the dual $VD_{near}^*(u, Q)$ from the primal $VD_{near}(u, Q)$. The total time to compute all such diagrams is therefore $O(\sum_i n \frac{r_{i+1}}{\sqrt{r_i}})$.

It remains to compute the diagram $VD_{far}^*(u, Q)$. Here we cannot afford to make an explicit computation of the primal Voronoi diagram $VD_{far}(u, Q)$. Instead, we will compute just the tree structure and the trichromatic faces of $VD_{far}(u, Q)$. These suffice as the representation of $VD_{far}^*(u, Q)$ since all we need for point location is the centroid decomposition of $VD_{far}^*(u, Q)$. The main idea is to use FR-Dijkstra to locate the trichromatic faces. Let K be a star with center u and leaves ∂Q , where for each vertex $w \in \partial Q$, the length of the arc uw is the additive distance of w in $VD_{far}^*(u, Q)$. We start by computing a shortest path tree T rooted at u in the union of K and the DDGs of all siblings of pieces in the complete recursive decomposition tree \mathcal{A} that contain Q . We shall prove that, by inspecting the restriction of T to the DDG of each piece P , we can infer whether P contains a trichromatic face or not. If P contains a trichromatic face we refine T inside P by replacing the DDG of P with the DDGs of its two children in \mathcal{A} , and recomputing the part of T inside P using FR-Dijkstra. We continue doing so until we locate all $O(|\partial Q|)$ trichromatic faces. The tree structure of $VD_{far}(u, Q)$ is captured by the structure of the shortest path tree in the DDGs of all the pieces at the end of this process. The total time to locate all the trichromatic faces is proportional, up to polylogarithmic factors, to the total number of vertices in all of the DDGs involved in all these computations. Since the total number of vertices in all DDGs containing a particular face is $O(\sqrt{n})$, the total time for finding all trichromatic faces as well as the tree structure is $\tilde{O}(\sqrt{n} \cdot \sqrt{|Q|})$. A more careful analysis that takes into account double counting of large pieces (cf. [11]) bounds this time by $\tilde{O}(\sqrt{n \sqrt{|Q|}})$. Since there are $n/\sqrt{r_i}$ boundary vertices at level i of the recursive r -division, the total time to compute all Voronoi diagrams $VD_{far}^*(u, Q)$ for a single level is $\tilde{O}(\frac{n}{\sqrt{r_i}} \cdot \sqrt{nr_{i+1}})$. Thus, the total time to construct all

Voronoi diagrams $VD_{far}^*(u, Q)$ for all levels is $\tilde{O}(n^{3/2} \sum_i \frac{r_{i+1}^{1/4}}{\sqrt{r_i}})$.

Putting everything together, the total construction time sums up to $\tilde{O}(n \sum_i (\frac{r_{i+1}}{\sqrt{r_i}} + \frac{\sqrt{nr_{i+1}}}{\sqrt{r_i}}))$.

Remark 10. *We could use the described algorithm to also compute $VD_{near}^*(u, Q)$, bringing the complexity down to $\tilde{O}(n \sum_i \frac{r_{i+1}^{3/4}}{\sqrt{r_i}})$. We omit the description due to technicalities arising from having to consider DDGs of ancestors of R , excluding R . If we did that, the construction time in Theorem 7 would be $\tilde{O}(n^{3/2} \sum_i \frac{r_{i+1}^{1/4}}{\sqrt{r_i}})$.*

The only missing part in the explanation above is showing we can infer whether a piece P contains a trichromatic face just by inspecting the restriction of T to P . Our choice of additive distance guarantees that each vertex of ∂Q is a child of u in T . We label

each vertex w of T by its unique ancestor in T that belongs to ∂Q . Note that the label of a vertex w corresponds to the Voronoi cell containing w in $\text{VD}_{far}^*(u, Q)$. Consider the restriction of T to the DDG of P . We use a representation of size $O(|\partial P|)$ of the edges of T embedded as curves in P , such that each edge of T is homologous (w.r.t. the holes of P) to its underlying shortest path in P . See [34, 37] for details on such a representation. We make incisions in the embedding of P along the edges of T (the endpoints of edges of T are duplicated in this process). Let \mathcal{P} be the set of connected components of P after all incisions are made.

Lemma 11. *P contains a trichromatic face if and only if some connected component C in \mathcal{P} contains boundary vertices of P with more than two distinct labels.*

Intuitively, for each connected component C in \mathcal{P} , each label appears as the label of boundary vertices along at most a single sequence of consecutive boundary vertices along the boundary of C . Consider the component C' obtained from C by connecting an artificial new vertex to all the boundary vertices with the same label. Since these boundary vertices form a single consecutive interval on the boundary of C , C' is also a planar graph when the new vertices are embedded in the infinite face of C . Now C' has a new infinite face where every vertex on that face has a distinct label, and there are more than two labels. The Voronoi diagram of C' necessarily contains a trichromatic face.

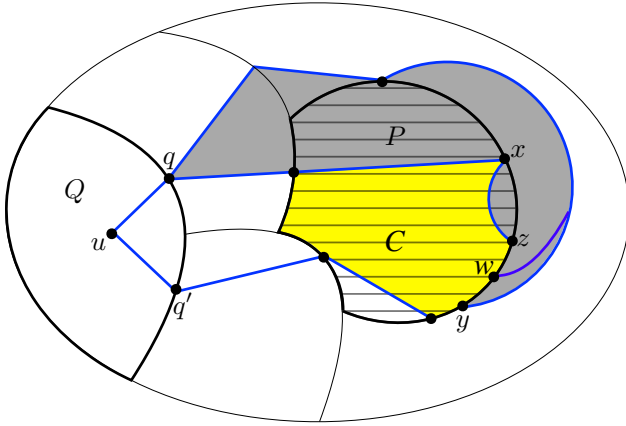


Figure 4: Illustration for the proof of Lemma 11. Some pieces in a graph G are shown. The piece Q (in bold) is a piece of some r_i -division in the recursive r -division of G . The piece P (bold boundary, horizontal stripes) is some piece in the complete recursive decomposition of G that lies outside Q . The shortest path tree T is shown in blue, the connected component C in yellow, and the cycle D in gray. Vertices x and y have the same label q . The vertices w, z between x and y (on the cyclic walk F along the infinite face of C) must also be labeled q .

PROOF. Let C be a connected component in \mathcal{P} . Note that the vertices of ∂Q either do not belong to C or they are incident to a single face f of C . In the former case, let f be the face of C such that

∂Q is embedded in f . We think of f as the infinite face of C . Note that, because any path from ∂Q to any vertex of C must intersect f , the set of labels of the vertices of f is identical to the set of labels of all of C .

We first claim that the vertices of f that have the same label are consecutive in the cyclic order of f . To see this, consider any two distinct vertices x, y of f that have the same label q . Note that, by the incision process defining C , if x is an ancestor of y in T (or vice versa) then all the vertices on one of the x -to- y subpaths of f are on the x -to- y path in T , and hence all have the label q . Assume, therefore, that neither is an ancestor of the other (as illustrated in Figure 4), and consider the (not necessarily simple) cycle D (in Q^{out}) formed by the unique x -to- y path in T , and the x -to- y path F of f such that D does not enclose C . Observe that D is non self crossing, and that D encloses no vertices of ∂Q except, perhaps, q itself. Suppose some vertex of F has a label $q' \neq q$, and let w be a rootmost such vertex in T . Since D does not enclose q' , the q' -to- w path in T is not enclosed by D , and hence must intersect C . But then C should have been further dissected when the incisions along T were performed, a contradiction.

The argument above established that the vertices of f that have the same label are consecutive in the cyclic order of f . For every unique label q of a vertex of f we embed inside f an artificial vertex and connect it to every vertex w of f that has label q with an edge whose length is the q -to- w distance in T . By the consecutiveness property above, this can be done without violating planarity. We connect the artificial vertices by infinite length edges, according to the cyclic order of the labels along f , so that resulting graph C' has a new infinite face containing just the artificial vertices. Consider now the additively weighted Voronoi diagram $\text{VD}(C')$ with the artificial vertices as sites, where the additive weight of an artificial vertex with label q is equal to the additive weight of the corresponding site q in $\text{VD}_{far}^*(u, Q)$. Note that, by construction, the additive distances in $\text{VD}(C')$ and in $\text{VD}_{far}^*(u, Q)$ are the same, so the restrictions of both diagrams to C are identical. It therefore suffices to prove that $\text{VD}(C')$ has a trichromatic face in C .

Since C has boundary vertices of P with more than two distinct labels, C' has more than two sites. Since all the vertices on the infinite face of C' are sites, and since every site is in its own Voronoi cell, $\text{VD}(C')$ has at least two trichromatic faces [8, 23] (one being the infinite face of C'). Observe that all the newly introduced faces of C' (those created by the edges connecting artificial vertices to vertices of f or to each other) have either a single label (in case the face is a triangle formed by an artificial vertex and two consecutive vertices of f with the same label), or two labels (in case the face has size 4, and is formed by two artificial vertices and two consecutive vertices of f with two distinct labels). It follows that $\text{VD}(C')$ must have a trichromatic face of C , and so does $\text{VD}_{far}^*(u, Q)$. \square

The entire discussion in this section on efficient computation of dual Voronoi Diagrams can be generalized to the following statement.

Theorem 12. *Given a face f on a graph G , and a complete recursive decomposition of G , such that every piece has been preprocessed for FR-Dijkstra as in Theorem 9, we can compute VD^* with sites the vertices on f and arbitrary input additive weights in time $\tilde{O}(\sqrt{|G|} \cdot |f|)$.*

6 FINAL REMARKS

The main open question is whether there exists a distance oracle occupying linear (or nearly linear) space and requiring constant (or polylogarithmic) time to answer queries. Note that currently, any oracle with $\tilde{O}(n)$ space requires polynomial query time. In particular, the fastest known oracles with strictly linear space [38, 40] require $\Omega(n^{1/2+\epsilon})$ query time. Another important question concerns the construction time. Is there a nearly-linear time algorithm to construct our oracle, or oracles with comparable space to query-time tradeoffs?

REFERENCES

- [1] Srinivasa Rao Arikati, Danny Z. Chen, L. Paul Chew, Gautam Das, Michiel H. M. Smid, and Christos D. Zaroliagis. 1996. Planar Spanners and Approximate Shortest Path Queries among Obstacles in the Plane. In *4th ESA*. 514–528.
- [2] Franz Aurenhammer. 1991. Voronoi Diagrams: a Survey of a Fundamental Geometric Data Structure. *ACM Comput. Surv.* 23, 3 (1991), 345–405.
- [3] Michael A. Bender and Martin Farach-Colton. 2000. The LCA Problem Revisited. In *4th LATIN*. 88–94.
- [4] Michael A. Bender and Martin Farach-Colton. 2004. The Level Ancestor Problem simplified. *Theor. Comput. Sci.* 321, 1 (2004), 5–12.
- [5] Glencora Borradaile, Piotr Sankowski, and Christian Wulff-Nilsen. 2010. Min st-cut Oracle for Planar Graphs with Near-Linear Preprocessing Time. In *51th FOCS*. 601–610.
- [6] Gerth Stølting Brodal, Rolf Fagerberg, Christian N. S. Pedersen, and Anna Östlin. 2001. The Complexity of Constructing Evolutionary Trees Using Experiments. In *Automata, Languages and Programming, 28th International Colloquium, ICALP 2001, Crete, Greece, July 8–12, 2001, Proceedings*. 140–151. https://doi.org/10.1007/3-540-48224-5_12
- [7] Sergio Cabello. 2012. Many Distances in Planar Graphs. *Algorithmica* 62, 1-2 (2012), 361–381.
- [8] Sergio Cabello. 2017. Subquadratic Algorithms for the Diameter and the Sum of Pairwise Distances in Planar Graphs. In *28th SODA*. 2143–2152.
- [9] Sergio Cabello, Erin W. Chambers, and Jeff Erickson. 2013. Multiple-Source Shortest Paths in Embedded Graphs. *SIAM J. Comput.* 42, 4 (2013), 1542–1571.
- [10] Timothy M. Chan and Dimitrios Skrepetos. 2017. Faster Approximate Diameter and Distance Oracles in Planar Graphs. In *25th ESA*. 25:1–25:13.
- [11] Panagiotis Charalampopoulos, Shay Mozes, and Benjamin Tebeka. 2019. Exact Distance Oracles for Planar Graphs with Failing Vertices. In *30th SODA*. 2110–2123.
- [12] Danny Z. Chen and Jinhui Xu. 2000. Shortest path queries in planar graphs. In *32nd STOC*. 469–478.
- [13] Vincent Cohen-Addad, Søren Dahlgaard, and Christian Wulff-Nilsen. 2017. Fast and Compact Exact Distance Oracle for Planar Graphs. In *58th FOCS*. 962–973.
- [14] Mark de Berg, Otfried Cheong, Marc J. van Kreveld, and Mark H. Overmars. 2008. *Computational geometry: algorithms and applications, 3rd Edition*. Springer. <http://www.worldcat.org/oclc/227584184>
- [15] Ugur Demiryurek and Cyrus Shahabi. 2012. Indexing Network Voronoi Diagrams. In *17th DASFAA*. 526–543.
- [16] Hristo Djidjev. 1996. On-Line Algorithms for Shortest Path Problems on Planar Digraphs. In *22nd WG*. 151–165.
- [17] David Eppstein and Michael T. Goodrich. 2008. Studying (non-planar) road networks through an algorithmic lens. In *16th ACM-GIS*. 16.
- [18] Paul Erdős. 1963. Extremal problems in graph theory. In *Proc. Symp. Theory of Graphs and its Applications*. 29–36.
- [19] Jeff Erickson, Kyle Fox, and Luvsandondov Lkhamsuren. 2018. Holiest minimum-cost paths and flows in surface graphs. In *50th STOC*. 1319–1332.
- [20] Martin Erwig. 2000. The graph Voronoi diagram with applications. *Networks* 36, 3 (2000), 156–163.
- [21] Jittat Fakcharoenphol and Satish Rao. 2006. Planar graphs, negative weight edges, shortest paths, and near linear time. *J. Comput. Syst. Sci.* 72, 5 (2006), 868–889.
- [22] Greg N. Frederickson. 1987. Fast Algorithms for Shortest Paths in Planar Graphs, with Applications. *SIAM J. Comput.* 16, 6 (1987), 1004–1022.
- [23] Paweł Gawrychowski, Shay Mozes, Oren Weimann, and Christian Wulff-Nilsen. 2018. Better Tradeoffs for Exact Distance Oracles in Planar Graphs. In *29th SODA*. 515–529.
- [24] Qian-Ping Gu and Gengchun Xu. 2015. Constant Query Time $(1+\epsilon)$ -Approximate Distance Oracle for Planar Graphs. In *26th ISAAC*. 625–636.
- [25] Monika Rauch Henzinger, Philip N. Klein, Satish Rao, and Sairam Subramanian. 1997. Faster Shortest-Path Algorithms for Planar Graphs. *J. Comput. Syst. Sci.* 55, 1 (1997), 3–23.
- [26] Shinichi Honiden, Michael E. Houle, Christian Sommer, and Martin Wolff. 2010. Approximate Shortest Path Queries Using Voronoi Duals. *Trans. Computational Science* 9 (2010), 28–53.
- [27] Haim Kaplan, Shay Mozes, Yahav Nussbaum, and Micha Sharir. 2017. Submatrix Maximum Queries in Monge Matrices and Partial Monge Matrices, and Their Applications. *ACM Trans. Algorithms* 13, 2 (2017), 26:1–26:42.
- [28] Ken-ichi Kawarabayashi, Philip N. Klein, and Christian Sommer. 2011. Linear-Space Approximate Distance Oracles for Planar, Bounded-Genus and Minor-Free Graphs. In *37th ICALP*. 135–146.
- [29] Ken-ichi Kawarabayashi, Christian Sommer, and Mikkel Thorup. 2013. More Compact Oracles for Approximate Distances in Undirected Planar Graphs. In *23rd SODA*. 550–563.
- [30] Philip N. Klein. 2002. Preprocessing an Undirected Planar Network to Enable Fast approximate distance queries. In *12th SODA*. 820–827.
- [31] Philip N. Klein. 2005. Multiple-source shortest paths in planar graphs. In *16th SODA*. 146–155.
- [32] Philip N. Klein and Shay Mozes. [n. d.]. Optimization Algorithms for Planar Graphs. <http://planarity.org>. Book draft.
- [33] Philip N. Klein, Shay Mozes, and Christian Sommer. 2013. Structured recursive separator decompositions for planar graphs in linear time. In *45th STOC*. 505–514.
- [34] Jakub Lacki and Piotr Sankowski. 2011. Min-Cuts and Shortest Cycles in Planar Graphs in $O(n \log \log n)$ Time. In *19th ESA*. 155–166.
- [35] Kurt Mehlhorn. 1988. A Faster Approximation Algorithm for the Steiner Problem in Graphs. *Inf. Process. Lett.* 27, 3 (1988), 125–128.
- [36] Gary L. Miller. 1984. Finding Small Simple Cycle Separators for 2-Connected Planar Graphs. In *16th STOC*. 376–382.
- [37] Shay Mozes, Kirill Nikolaev, Yahav Nussbaum, and Oren Weimann. 2018. Minimum Cut of Directed Planar Graphs in $O(n \log \log n)$ Time. In *29th SODA*. 477–494.
- [38] Shay Mozes and Christian Sommer. 2012. Exact distance oracles for planar graphs. In *23rd SODA*. 209–222.
- [39] Shay Mozes and Christian Wulff-Nilsen. 2010. Shortest Paths in Planar Graphs with Real Lengths in $O(n \log^2 n / \log \log n)$ Time. In *18th ESA*. 206–217.
- [40] Yahav Nussbaum. 2011. Improved Distance Queries in Planar Graphs. In *12th WADS*. 642–653.
- [41] Atsuyuki Okabe, Toshiaki Satoh, Takehiro Furuta, Atsuo Suzuki, and K. Okano. 2008. Generalized network Voronoi diagrams: Concepts, computational methods, and applications. *International Journal of Geographical Information Science* 22, 9 (2008), 965–994.
- [42] Mihai Patrascu and Liam Roditty. 2014. Distance Oracles beyond the Thorup-Zwick Bound. *SIAM J. Comput.* 43, 1 (2014), 300–311.
- [43] Christian Sommer. 2014. Shortest-path queries in static networks. *ACM Comput. Surv.* 46, 4 (2014), 45:1–45:31.
- [44] Mikkel Thorup. 2004. Compact oracles for reachability and approximate distances in planar digraphs. *J. ACM* 51, 6 (2004), 993–1024.
- [45] Mikkel Thorup and Uri Zwick. 2005. Approximate distance oracles. *J. ACM* 52, 1 (2005), 1–24.
- [46] Christian Wulff-Nilsen. 2016. Approximate Distance Oracles for Planar Graphs with Improved Query Time-Space Tradeoff. In *26th SODA*. 351–362.