

Planar Distance Oracles with Better Time-Space Tradeoffs*

Yaowei Long[†]

Seth Pettie[‡]

Abstract

In a recent breakthrough, Charalampopoulos, Gawrychowski, Mozes, and Weimann [9] showed that exact distance queries on planar graphs could be answered in $n^{o(1)}$ time by a data structure occupying $n^{1+o(1)}$ space, i.e., up to $o(1)$ terms, optimal exponents in time (0) and space (1) can be achieved *simultaneously*. Their distance query algorithm is recursive: it makes successive calls to a point-location algorithm for planar Voronoi diagrams, which involves many recursive distance queries. The depth of this recursion is non-constant and the branching factor logarithmic, leading to $(\log n)^{\omega(1)} = n^{o(1)}$ query times.

In this paper we present a new way to do point-location in planar Voronoi diagrams, which leads to a new exact distance oracle. At the two extremes of our space-time tradeoff curve we can achieve either

$n^{1+o(1)}$ space and $\log^{2+o(1)} n$ query time, or

$n \log^{2+o(1)} n$ space and $n^{o(1)}$ query time.

All previous oracles with $\tilde{O}(1)$ query time occupy space $n^{1+\Omega(1)}$, and all previous oracles with space $\tilde{O}(n)$ answer queries in $n^{\Omega(1)}$ time.

1 Introduction

A *distance oracle* is a data structure that answers distance queries (or approximate distance queries) w.r.t. some underlying graph or metric space. On general graphs there are many well known distance oracles that pit space against multiplicative approximation [40], space against mixed multiplicative/additive approximation [34, 1], and, in sparse graphs, space against query time [2, 38]. Refer to Sommer [37] for a survey on distance oracles.

Whereas *approximation* seems to be a necessary ingredient to achieve any reasonable space/query time on general graphs, structured graph classes may admit *exact* distance oracles with attractive time-space tradeoffs. In this paper we continue a long line of work [3, 14, 10, 17, 27, 41, 33, 7, 32, 11, 19, 9] focused on exact distance oracles for weighted, directed planar graphs.

*This work was supported by NSF grants CCF-1637546 and CCF-1815316, and a grant from IIS, Tsinghua University. A full version of this paper is available at <https://arxiv.org/abs/2007.08585>.

[†]Tsinghua University, longyw17@mails.tsinghua.edu.cn.

[‡]University of Michigan, pettie@umich.edu.

1.1 History. Between 1996-2012, work of Arikati et al. [3], Djidjev [14], Chen and Xu [10], Fakcharoenphol and Rao [17], Klein [27], Wulff-Nilsen [41], Nussbaum [33], Cabello [7], and Mozes and Sommer [32] achieved space $\tilde{O}(S)$ and query time $\tilde{O}(n/\sqrt{S})$, for various ranges of S that ultimately covered the full range $[n, n^2]$.

In 2017, Cabello [8] introduced *planar* Voronoi diagrams as a tool for solving metric problems in planar graphs, such as diameter and sum-of-distances. This idea was incorporated into new planar distance oracles, leading to $\tilde{O}(n^{5/2}/S^{3/2})$ query time [11] for $S \in [n^{3/2}, n^{5/3}]$ and $\tilde{O}(n^{3/2}/S)$ query time [19] for $S \in [n, n^{3/2}]$. Finally, in a major breakthrough Charalampopoulos, Gawrychowski, Mozes, and Weimann [9] demonstrated that up to $n^{o(1)}$ factors, *there is no tradeoff* between space and query time, i.e., space $n^{1+o(1)}$ and query time $n^{o(1)}$ can be achieved simultaneously. In more detail, they proved that space $O(n^{4/3} \sqrt{\log n})$ allows for query time $O(\log^2 n)$, space $\tilde{O}(n^{1+\epsilon})$ allows for query time $O(\log n)^{1/\epsilon-1}$, and space $O(n \log^{2+1/\epsilon} n)$ allows for query time $O(n^{2\epsilon})$.

The Charalampopoulos et al. structure is based on a hierarchical \vec{r} -decomposition of the graph, $\vec{r} = (n, n^{(m-1)/m}, \dots, n^{1/m})$. (See Section 2.) Given u, v , it iteratively finds the last boundary vertex u_i on the shortest u - v path that lies on the boundary of the level- i region containing u . Given u_{i-1} , finding u_i amounts to solving a *point location* problem on an *external* Voronoi diagram, i.e., a Voronoi diagram of the *complement* of a region in the hierarchy. Each point location query is solved via a kind of binary search, and each step of the binary search involves 3 *recursive* distance queries that begin at a “higher” level in the hierarchy. This leads to a tradeoff between space $\tilde{O}(n^{1+1/m})$ and query time $O(\log n)^{m-1}$.

See Table 1 for a summary of the space-time tradeoffs for exact and approximate planar distance oracles.

1.2 New Results. In this paper we develop a more direct and more efficient way to do point location in external Voronoi diagrams. It uses a new persistent data structure for maintaining sets of non-crossing systems of *chords*, which are paths that begin and end at the boundary vertices of a region, but are internally vertex

REFERENCE		SPACE	QUERY TIME
Arikati, Chen, Chew Das, Smid & Zaroliagis	1996	$S \in [n^{3/2}, n^2]$	$O\left(\frac{n^2}{S}\right)$
Djidjev	1996	$S \in [n, n^2]$	$O\left(\frac{n^2}{S}\right)$
		$S \in [n^{4/3}, n^{3/2}]$	$O\left(\frac{n}{\sqrt{S}} \log n\right)$
Chen & Xu	2000	$S \in [n^{4/3}, n^2]$	$O\left(\frac{n}{\sqrt{S}} \log\left(\frac{n}{\sqrt{S}}\right)\right)$
Fakcharoenphol & Rao	2006	$O(n \log n)$	$O(\sqrt{n} \log^2 n)$
Wulff-Nilsen	2010	$O\left(n^2 \frac{\log^4 \log n}{\log n}\right)$	$O(1)$
Nussbaum	2011	$O(n)$	$O(n^{1/2+\epsilon})$
		$S \in [n^{4/3}, n^2]$	$O\left(\frac{n}{\sqrt{S}}\right)$
Cabello	2012	$S \in [n^{4/3} \log^{1/3} n, n^2]$	$O\left(\frac{n}{\sqrt{S}} \log^{3/2} n\right)$
Mozes & Sommer	2012	$S \in [n \log \log n, n^2]$	$O\left(\frac{n}{\sqrt{S}} \log^2 n \log^{3/2} \log n\right)$
		$O(n)$	$O(n^{1/2+\epsilon})$
Cohen-Addad, Dahlgaard & Wulff-Nilsen	2017	$S \in [n^{3/2}, n^{5/3}]$	$O\left(\frac{n^{5/2}}{S^{3/2}} \log n\right)$
Gawrychowski, Mozes, Weimann & Wulff-Nilsen	2018	$\tilde{O}(S)$ for $S \in [n, n^{3/2}]$	$\tilde{O}\left(\frac{n^{3/2}}{S}\right)$
Charalampopoulos, Gawrychowski, Mozes & Weimann	2019	$n^{1+o(1)}$	$n^{o(1)}$
		$\tilde{O}(n^{1+\epsilon})$ for $0 < \epsilon \leq 1/3$	$O(\log^{\epsilon^{-1}-1} n)$
		$O(n \log^{2+1/\epsilon} n)$ for $\epsilon > 0$	$O(n^{2\epsilon})$
new	2020	$n^{1+o(1)}$	$\log^{2+o(1)} n$
		$n \log^{2+o(1)} n$	$n^{o(1)}$

$(1 + \epsilon)$ -APPROX. ORACLES		SPACE	QUERY TIME
Thorup	2001	$O(n\epsilon^{-1} \log^2 n)$	$O(\log \log n + \epsilon^{-1})$
		$O(n\epsilon^{-1} \log n)$	$O(\epsilon^{-1})$ (Unidir.)
Klein	2002	$O(n(\log n + \epsilon^{-1} \log \epsilon^{-1}))$	$O(\epsilon^{-1})$ (Unidir.)
Kawarabayashi, Klein & Sommer	2011	$O(n)$	$O(\epsilon^{-2} \log^2 n)$ (Unidir.)
Kawarabayashi, Sommer & Thorup	2013	$\bar{O}(n \log n)$	$\bar{O}(\epsilon^{-1})$ (Unidir.)
		$\bar{O}(n)$	$\bar{O}(\epsilon^{-1})$ (Unidir., Unweight.)

Table 1: Space-query time tradeoffs for exact and approximate planar distance oracles. \bar{O} hides $\log(\epsilon^{-1} \log n)$ factors.

disjoint from the boundary. By applying this point location method in the framework of Charalampopoulos et al. [9], we obtain a better time-space tradeoff, which is most noticeable at the “extremes” when $\tilde{O}(n)$ space or $\tilde{O}(1)$ query time is prioritized.

THEOREM 1.1. *Let G be an n -vertex weighted planar digraph with no negative cycles, and let $\kappa, m \geq 1$ be parameters. A distance oracle occupying space $O(m\kappa n^{1+1/m+1/\kappa})$ can be constructed in $\tilde{O}(n^{3/2+1/m} + n^{1+1/m+1/\kappa})$ time that answers exact distance queries in $O(2^m \kappa \log^2 n \log \log n)$ time. At the two extremes of the space-time tradeoff curve, we can construct oracles in $n^{3/2+o(1)}$ time with either*

- $n^{1+o(1)}$ space and $\log^{2+o(1)} n$ query time, or
- $n \log^{2+o(1)} n$ space and $n^{o(1)}$ query time.

Our new point-location routine suffices to get the query time down to $O(\log^3 n)$. In order to reduce it further to $O(\log^{2+o(1)} n)$, we develop a new dynamic tree data structure based on Euler-Tour trees [22] with $O(\kappa n^{1/\kappa})$ update time and $O(\kappa)$ query time. This allows us to generate MSSP (multiple-source shortest paths) structures with a similar space-query tradeoff, specifically, $O(\kappa n^{1+1/\kappa})$ space and $O(\kappa \log \log n)$ query time. Our MSSP construction follows Klein [27] (see also [19]), but uses our new dynamic tree in lieu of Sleator and Tarjan’s Link-Cut trees [36], and uses persistent arrays [12] in lieu of [15] to make the data structure persistent.

1.3 Organization. In Section 2 we review background on planar embeddings, planar separators, multiple-source shortest paths, and weighted Voronoi diagrams. In Section 3 we introduce key parts of the data structure and describe the query algorithm, assuming a certain point location problem can be solved. Section 4 introduces several more components of the data structure, and shows how they can be applied to solve this particular point location problem in near-logarithmic time. The space and query-time claims of Theorem 1.1 are proved in Section 5. We leave the following details and proofs in the full version [30] of this paper due to page limitations. The construction time claims of Theorem 1.1 are proved in [30, Appendix C]. The MSSP structure based on Euler Tour trees appear in [30, Appendix B]. In [30, Appendix A] it is explained how to remove a simplifying assumption made throughout the paper, that the boundary vertices of every region in the \vec{r} -decomposition lie on a *single* hole, which is bounded by a *simple* cycle.

2 Preliminaries

2.1 The Graph and Its Embedding. A weighted planar graph $G = (V, E, \ell)$ is represented by an abstract embedding: for each $v \in V(G)$ we list the edges incident to v according to a clockwise order around v . We assume the graph has no negative weight cycles and further assume the following, without loss of generality.

- All the edge-weights can be made non-negative ($\ell : E \rightarrow \mathbb{R}_{\geq 0}$) [23]. Furthermore, via randomized or deterministic perturbation [16], we can assume there are no zero weight edges, and that shortest paths are *unique* in *any* subgraph of G .
- The graph is connected and triangulated. Assign all artificial edges weight $n \cdot \max_{e \in E(G)} \{\ell(e)\}$ so as not to affect any finite distances.
- If $(u, v) \in E(G)$ then $(v, u) \in E(G)$ as well. (In the circular ordering around v , they are represented as a single element $\{u, v\}$.)

Suppose $P = (v_0, v_1, \dots, v_k)$ is a path oriented from v_0 to v_k , and $e = (v_i, u)$ is an edge not on P , $i \in [1, k-1]$. Then e is to the right of P if e appears between (v_i, v_{i+1}) and (v_{i-1}, v_i) in the clockwise order around v_i , and left of P otherwise.

2.2 Separators and Divisions. Lipton and Tarjan [29] proved that every planar graph contains a *separator* of $O(\sqrt{n})$ vertices that, once removed, breaks the graph into components of at most $2/3$ the size. Miller [31] showed that every triangulated planar graph has a $O(\sqrt{n})$ -size separator that consists of a simple cycle. Frederickson [18] defined a *division* to be a set of edge-induced subgraphs whose union is G . A vertex in more than one region is a *boundary* vertex; the boundary of a region R is denoted ∂R . Edges along the boundary between two regions appear in both regions. The r -divisions of [18] have $\Theta(n/r)$ regions each with $O(r)$ vertices and $O(\sqrt{r})$ boundary vertices.

We use a linear-time algorithm of Klein, Mozes, and Sommer [28] for computing a hierarchical \vec{r} -division, where $\vec{r} = (r_m, \dots, r_1)$ and $n = r_m > \dots > r_1 = \Omega(1)$. Such an \vec{r} -division has the following properties:

- (Division & Hierarchy) For each i , \mathcal{R}_i is the set of regions in an r_i -division of G , where $\mathcal{R}_m = \{G\}$ consists of the graph itself. For each $i < i' \leq m$ and $R_i \in \mathcal{R}_i$, there is a unique $R_{i'} \in \mathcal{R}_{i'}$ such that $E(R_i) \subseteq E(R_{i'})$. The \vec{r} -division is therefore represented as a rooted tree of regions.
- (Boundaries and Holes) The $O(\sqrt{r_i})$ boundary vertices of any $R_i \in \mathcal{R}_i$ lie on a constant number of

faces of R_i called *holes*, each bounded by a cycle (not necessarily simple).

We supplement the \bar{r} -division with a zeroth level. The layer-0 $\mathcal{R}_0 = \{\{v\} \mid v \in V(G)\}$ consists of singleton sets, and each $\{v\}$ is attached as a (leaf) child of an arbitrary $R \in \mathcal{R}_1$ for which $v \in R$.

Suppose f is one of the $O(1)$ holes of region R and C_f the cycle around f . The cycle C_f partitions $E(G) - C_f$ into two parts. Let $R^{f,\text{out}}$ be the graph induced by the part disjoint from R , together with C_f , i.e., C_f appears in both R and $R^{f,\text{out}}$. To keep the description of the algorithm as simple as possible, we will assume that ∂R lies on a single simple cycle (hole) f_R , and let R^{out} be short for $R^{f_R,\text{out}}$. The modifications necessary to deal with multiple holes and non-simple boundary cycles are explained in Appendix B in the full version.

2.3 Multiple-source Shortest Paths. Suppose H is a weighted planar graph with a distinguished face f on vertices S . Klein’s MSSP algorithm takes $O(|H| \log |H|)$ time and produces an $O(|H| \log |H|)$ -size data structure such that given $s \in S$ and $v \in V(H)$, returns $\text{dist}_H(s, v)$ in $O(\log |H|)$ time. Klein’s algorithm can be viewed as continuously moving the source vertex around the boundary face f , recording all changes to the SSSP tree in a dynamic tree data structure [36]. It is shown [27] that each edge in H enters and leaves the SSSP tree exactly once, meaning the number of changes is $O(|H|)$. Each change to the tree is effected in $O(\log |H|)$ time [36], and the generic persistence method of [15] allows for querying any state of the SSSP tree. The important point is that the total space is linear in the number of updates to the structure ($O(|H|)$) times the update time ($O(\log |H|)$). As observed in [19], this structure can also answer other useful queries in $O(\log |H|)$ time. Lemma 2.1 is similar to [27, 19] except that we use a dynamic tree data structure based on Euler Tour trees [22] rather than Link-Cut trees [36], which allows for a more flexible tradeoff between update and query time. Because our data structure does not satisfy the criteria of Driscoll et al.’s [15] persistence method for pointer-based data structures, we use the folklore implementation of persistent arrays¹ to make any RAM data structure persistent, with doubly-logarithmic slowdown in the query time. See Appendix A in the full version for a proof of Lemma 2.1.

LEMMA 2.1. (Cf. Klein [27], Gawrychowski et al. [19]) Let H be a planar graph, S be the vertices on some

¹Dietz [12] credits this method to an oral presentation of Dietzfelbinger et al. [13], which highlighted it as an application of dynamic perfect hashing.




Figure 1: The clockwise order of e_x, e_u, e_v around x tells us whether the shortest $s-u$ path branches from the shortest $s-v$ path to the right or left.

distinguished face f , and $\kappa \geq 1$ be a parameter. An $O(\kappa |H|^{1+1/\kappa})$ -space data structure can be computed in $O(\kappa |H|^{1+1/\kappa})$ time that answers the following queries in $O(\kappa \log \log |H|)$ time.

- Given $s \in S, v \in V(H)$, return $\text{dist}_H(s, v)$.
- Given $s \in S, u, v \in V(H)$, return (x, e_u, e_v) , where x is the least common ancestor of u and v in the SSSP tree rooted at s and e_z is the edge on the path from x to z (if $x \neq z$), $z \in \{u, v\}$.

The purpose of the second query is to tell whether u lies on the shortest $s-v$ path ($x = u$) or vice versa, or to tell which direction the $s-u$ path branches from the $s-v$ path. Once we retrieve the LCA x and edges e_u, e_v , we get the edge e_x from x to its parent. The clockwise order of e_x, e_u, e_v around x tells us whether $s-u$ branches from $s-v$ to the left or right. See Figure 1.

2.4 Additively Weighted Voronoi Diagrams.

Let H be a weighted planar graph, f a distinguished face whose vertices S are called *sites*, and $\omega : S \rightarrow \mathbb{R}_{\geq 0}$ be a weight function on sites. We augment H with large-weight edges so that it is triangulated, except for f . For $s \in S, v \in V(H)$, define

$$d^\omega(s, v) \stackrel{\text{def}}{=} \omega(s) + \text{dist}_H(s, v).$$

The *Voronoi diagram* $\text{VD}[H, S, \omega]$ is a partition of $V(H)$ into *Voronoi cells*, where for $s \in S$, $\text{Vor}(s) \stackrel{\text{def}}{=} \{v \in V(H) \mid \forall s' \neq s. (d^\omega(s, v), -\omega(s)) < (d^\omega(s', v), -\omega(s'))\}$. In other words, $\text{Vor}(s)$ is the set of vertices that are closer to s than any other site, breaking ties in favor of larger ω -values. We usually work with the dual representation of a Voronoi diagram. It is constructed as follows.

- Define \hat{S} to be the set of sites with nonempty Voronoi cells, i.e., $\hat{S} = \{s \in S \mid s \in \text{Vor}(s)\}$. The case $|\hat{S}| = 1$ is trivial, so assume $|\hat{S}| \geq 2$.

- Add large-weight dummy edges to H so that \hat{S} appear on the boundary of a single face \hat{f} , but is otherwise triangulated. Observe that this has no effect on the Voronoi cells.
- An edge is *bichromatic* if its endpoints are in different cells. In particular, the edges bounding \hat{f} are all bichromatic. Define VD_0^* to be the (undirected) subgraph of H^* consisting of the duals of bichromatic edges.
- Obtain VD_1^* from VD_0^* by repeatedly contracting edges incident to a degree-2 vertex, terminating when there are no degree-2 vertices, or when it becomes a self-loop.² Observe that in VD_1^* , \hat{f}^* has degree $|\hat{S}|$ and all other vertices have degree 3; moreover, the faces of VD_1^* are in one-to-one correspondence with the Voronoi cells.
- We obtain $\text{VD}^* = \text{VD}^*[H, S, \omega]$ by splitting \hat{f}^* into $|\hat{S}|$ degree-1 vertices, each taking an edge formerly incident to \hat{f}^* . It was proved in [19, Lemma 4.1] that VD^* is a single tree.³
- We store with VD^* supplementary information useful for point location. Each degree-3 vertex g^* in VD^* corresponds a *trichromatic* face g whose three vertices, say y_0, y_1, y_2 , belong to different Voronoi cells. We store in VD^* the sites $s_0, s_1, s_2 \in S$ such that $y_i \in \text{Vor}(s_i)$. We also store a *centroid decomposition* of VD^* . A centroid of a tree T is a vertex c that partitions the edge set of T into disjoint subtrees $T_1, \dots, T_{\deg(c)}$, each containing at most $(|E(T)| + 1)/2$ edges, and each containing c as a leaf. The decomposition is a tree rooted at c , whose subtrees are the centroid decompositions of $T_1, \dots, T_{\deg(c)}$. The recursion bottoms out when T consists of a single edge, which is represented as a single (leaf) node in the centroid decomposition.⁴

The most important query on Voronoi diagrams is *point location*.

LEMMA 2.2. (Gawrychowski et al. [19]) *The $\text{PointLocate}(\text{VD}^*[H, S, \omega], v)$ function is given the dual representation of a Voronoi diagram $\text{VD}^*[H, S, \omega]$ and a vertex $v \in V(H)$ and reports the $s \in S$ for which $v \in \text{Vor}(s)$. Given access to an MSSP data structure for H with source-set S and query time τ , we can answer*

²The latter case only occurs when $|\hat{S}| = 2$.

³If we skipped the step of forming the face \hat{f} on the site-set \hat{S} and triangulating the rest, VD^* would still be acyclic, but perhaps disconnected. See [19, 9].

⁴I.e., internal nodes correspond to vertices of T ; leaf nodes correspond to edges of T .

$\text{PointLocate}(\text{VD}^*[H, S, \omega], v)$ queries in $O(\tau \cdot \log |H|)$ time.

The challenge in our data structure (as in [9]) is to do point location when our space budget precludes storing all the relevant MSSP structures. Nonetheless, we do make use of **PointLocate** when the MSSP data structures are available.

3 The Distance Oracle

As in [9], the distance oracle is based on an \vec{r} -decomposition, $\vec{r} = (r_m, \dots, r_1)$, where $r_i = n^{i/m}$ and m is a parameter. Suppose we want to compute $\text{dist}_G(u, v)$. Let $R_0 = \{u\}$ be the artificial level-0 region containing u and $R_i \in \mathcal{R}_i$ be the level- i ancestor of R_0 . (Throughout the paper, we will use “ R_i ” to refer specifically to the level- i ancestor of $R_0 = \{u\}$, as well as to a *generic* region at level- i . Surprisingly, this will cause no confusion.) Let t be the smallest index for which $v \notin R_t$ but $v \in R_{t+1}$. Define u_i to be the *last* vertex on ∂R_i encountered on the shortest path from u to v . The main task of the distance query algorithm is to compute the sequence $(u = u_0, \dots, u_t)$. Suppose that we know the identity of u_i and $t > i$. Finding u_{i+1} now amounts to a point location problem in $\text{VD}^*[R_{i+1}^{\text{out}}, \partial R_{i+1}, \omega]$, where $\omega(s)$ is the distance from u_i to $s \in \partial R_{i+1}$. However, we cannot apply the fast **PointLocate** routine because we cannot afford to store an MSSP structure for every $(R_{i+1}^{\text{out}}, \partial R_{i+1})$, since $|R_{i+1}^{\text{out}}| = \Omega(|G|)$. Our point location routine narrows down the number of possibilities for u_{i+1} to at most two candidates in $O(\kappa \log^{2+o(1)} n)$ time, then decides between them using two recursive distance queries, but starting at a higher level in the hierarchy. There are about 2^m recursive calls in total, leading to a $O(2^m \kappa \log^{2+o(1)} n)$ query time.

The data structure is composed of several parts. Parts (A) and (B) are explained below⁵ while parts (C)–(E) will be revealed in Section 4.2.

- (A) (**MSSP Structures**) For each $i \in [0, m - 1]$ and each region $R_i \in \mathcal{R}_i$ with parent $R_{i+1} \in \mathcal{R}_{i+1}$, we store an MSSP data structure (Lemma 2.1) for the graph R_i^{out} , and source set ∂R_i . However, the structure only answers queries for $s \in \partial R_i$ and $u, v \in R_i^{\text{out}} \cap R_{i+1}$. Rather than represent the full SSSP tree from each root on $s \in \partial R_i$, the MSSP data structure only stores the tree induced by $R_i^{\text{out}} \cap R_{i+1}$, i.e., the parent of any vertex $v \in R_i^{\text{out}} \cap R_{i+1}$ is its nearest ancestor v' in the SSSP tree such that $v' \in R_i^{\text{out}} \cap R_{i+1}$. If (v', v) is

⁵They are similar to corresponding parts of the [9] structure, but the MSSP data structures in (A) use “shortcut” edges to simplify the distance query algorithm.

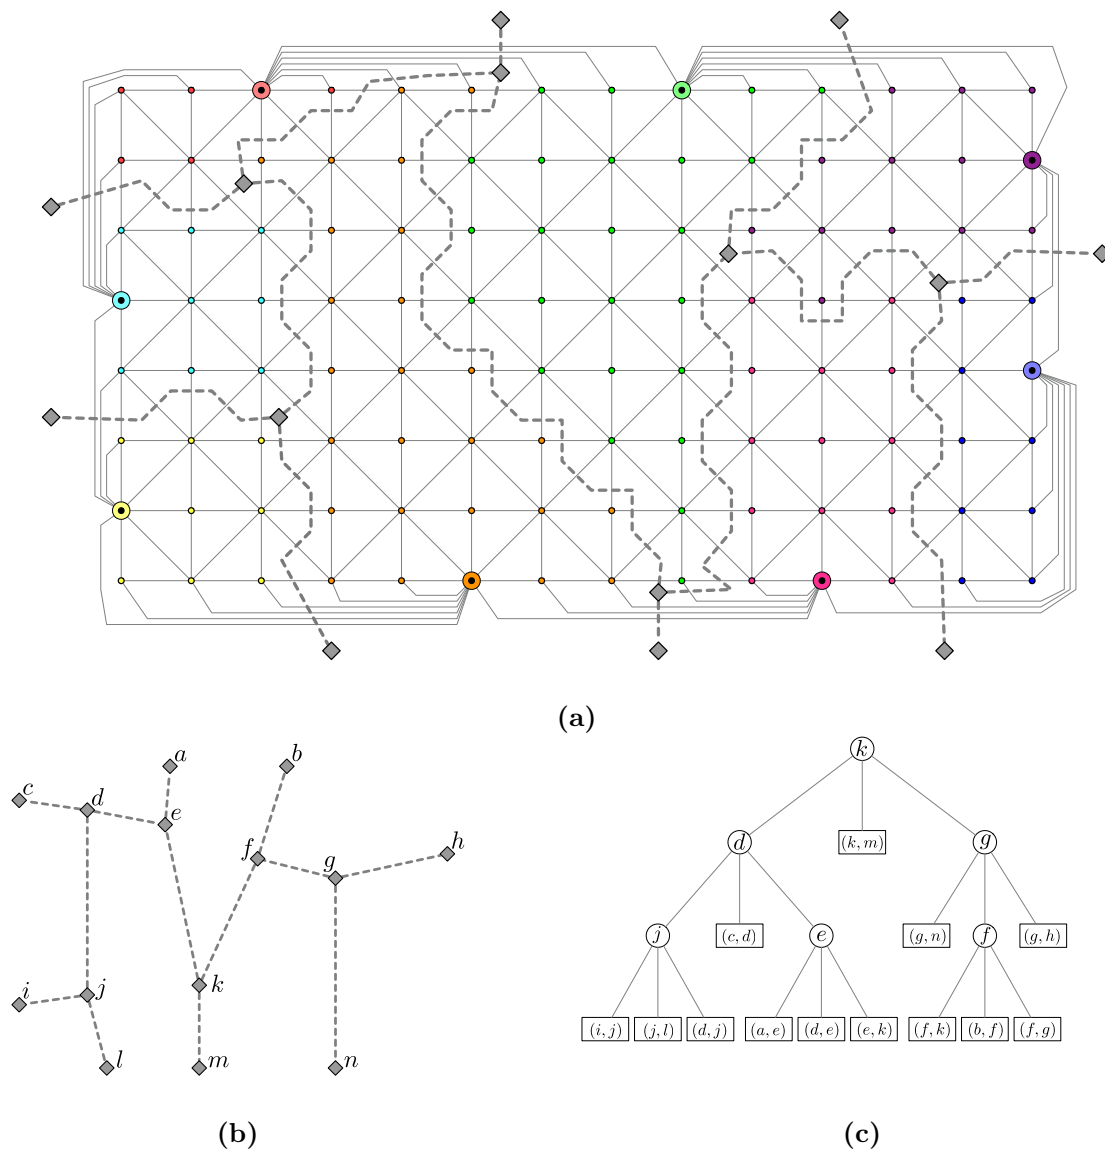


Figure 2: **(a)** The original H is a triangulated grid, with f being the exterior face. The boundary vertices \hat{S} with non-empty Voronoi cells are marked with colored halos. Edges are added so that \hat{S} are on the exterior face \hat{f} . The vertices of VD^* are the duals of trichromatic faces, and those derived by splitting \hat{f}^* into $|\hat{S}|$ vertices. The edges of VD^* correspond to paths of duals of bichromatic edges. **(b)** The dual representation VD^* . **(c)** A centroid decomposition of VD^* .

a “shortcut” edge corresponding to a path in R_{i+1}^{out} , it has weight $\text{dist}_{R_{i+1}^{\text{out}}}(v', v)$.

We fix a κ and let the update time in the dynamic tree data structure be $O(\kappa n^{1/\kappa})$ time. Thus, the space for this structure is $O((|R_i^{\text{out}} \cap R_{i+1}| + |\partial R_i| \cdot |\partial R_{i+1}|) \cdot \kappa n^{1/\kappa}) = O(r_{i+1} \cdot \kappa n^{1/\kappa})$ since each edge in $R_i^{\text{out}} \cap R_{i+1}$ is swapped into and out of the SSSP tree once [27], and the number of shortcut edges on ∂R_{i+1} swapped into and out of the SSSP is at most $|\partial R_{i+1}|$ for each of the $|\partial R_i|$ sources. Over all i and $\Theta(n/r_i)$ choices of R_i , the space is $O(m\kappa n^{1+1/m+1/\kappa})$ since $r_{i+1}/r_i = n^{1/m}$.

- (B) (**Voronoi Diagrams**) For each $i \in [0, m - 2]$ and $R_i \in \mathcal{R}_i$ with parent $R_{i+1} \in \mathcal{R}_{i+1}$, and each $q \in \partial R_i$, define $\text{VD}_{\text{out}}^*(q, R_{i+1})$ to be $\text{VD}^*[R_{i+1}^{\text{out}}, \partial R_{i+1}, \omega]$, with $\omega(s) = \text{dist}_G(q, s)$. The space to store the dual diagram and its centroid decomposition is $O(|\partial R_{i+1}|) = O(\sqrt{r_{i+1}})$. Over all choices for i, R_i , and q , the space is $O(mn^{1+1/(2m)})$ since $\sqrt{r_{i+1}/r_i} = n^{1/(2m)}$.

Due to our tie-breaking rule in the definition of $\text{Vor}(\cdot)$, locating u_{i+1} ($t \geq i + 1$) is tantamount to performing a point location on a Voronoi diagram in part (B) of the data structure.

LEMMA 3.1. *Suppose that $q \in \partial R_i$ and $v \notin R_{i+1}$. Consider the Voronoi diagram associated with $\text{VD}_{\text{out}}^*(q, R_{i+1})$ with sites ∂R_{i+1} and additive weights defined by distances from q in G . Then $v \in \text{Vor}(s)$ if and only if s is the *last* ∂R_{i+1} -vertex on the shortest path from q to v in G , and $d^\omega(s, v) = \text{dist}_G(q, v)$.*

Proof. By definition, $d^\omega(s, v)$ is the length of the shortest path from q to v that passes through s and whose s - v suffix does not leave R_{i+1}^{out} . Thus, $d^\omega(s, v) \geq \text{dist}_G(q, v)$ for every s , and $d^\omega(s, v) = \text{dist}_G(q, v)$ for some s . Because of our assumption that all edges are strictly positive, and our tie-breaking rule for preferring larger ω -values in the definition of $\text{Vor}(\cdot)$, if $v \in \text{Vor}(s)$ then s must be the *last* ∂R_{i+1} -vertex on the shortest q - v path. \square

3.1 The Query Algorithm. A distance query is given $u, v \in V(G)$. We begin by identifying the level-0 region $R_0 = \{u\} \in \mathcal{R}_0$ and call the function $\text{Dist}(u, v, R_0)$. In general, the function $\text{Dist}(u_i, v, R_i)$ takes as arguments a region R_i , a source vertex u_i on the boundary ∂R_i , and a target vertex $v \notin R_i$. It returns a value d such that

$$(3.1) \quad \text{dist}_G(u_i, v) \leq d \leq \text{dist}_{R_i^{\text{out}}}(u_i, v).$$

Note that $R_0^{\text{out}} = G$, so the initial call to this function correctly computes $\text{dist}_G(u, v)$. When v is “close” to u_i ($v \in R_i^{\text{out}} \cap R_{i+1}$) it computes $\text{dist}_{R_i^{\text{out}}}(u_i, v)$ without recursion, using part (A) of the data structure. When $v \in R_{i+1}^{\text{out}}$ it performs point location using the function **CentroidSearch**, which culminates in up to two recursive calls to **Dist** on the level- $(i + 1)$ region R_{i+1} . Thus, the correctness of **Dist** hinges on whether **CentroidSearch** correctly computes distances when $v \in R_{i+1}^{\text{out}}$.

Algorithm 1 $\text{Dist}(u_i, v, R_i)$

Input: A region R_i , a source $u_i \in \partial R_i$ and a destination $v \in R_i^{\text{out}}$.

Output: A value d such that $\text{dist}_G(u_i, v) \leq d \leq \text{dist}_{R_i^{\text{out}}}(u_i, v)$.

- 1: **if** $v \in R_i^{\text{out}} \cap R_{i+1}$ **then**
 - 2: **return** $d \leftarrow \text{dist}_{R_i^{\text{out}}}(u_i, v)$
 - 3: **end if**
 - 4: $f^* \leftarrow$ root of the centroid decomposition of $\text{VD}_{\text{out}}^*(u_i, R_{i+1})$
 - 5: **return** $d \leftarrow \text{CentroidSearch}(\text{VD}_{\text{out}}^*(u_i, R_{i+1}), v, f^*)$
-

Algorithm 2 $\text{CentroidSearch}(\text{VD}_{\text{out}}^*(u_i, R_{i+1}), v, f^*)$

Input: The dual representation $\text{VD}_{\text{out}}^* = \text{VD}_{\text{out}}^*(u_i, R_{i+1})$ of a Voronoi diagram with additive weights $\omega(s) = \text{dist}_G(u_i, s)$, a vertex $v \in R_{i+1}^{\text{out}}$, and a node f^* in the centroid decomposition of VD_{out}^* .

Output: The distance $\text{dist}_G(u_i, v)$.

- 1: **if** f^* is a leaf in the centroid decomposition (an edge in VD_{out}^*) **then**
 - 2: $s_1, s_2 \leftarrow$ sites whose Voronoi cells are bounded by f^*
 - 3: $d_1 \leftarrow \omega(s_1) + \text{Dist}(s_1, v, R_{i+1})$
 - 4: $d_2 \leftarrow \omega(s_2) + \text{Dist}(s_2, v, R_{i+1})$
 - 5: **return** $\min(d_1, d_2)$
 - 6: **end if**
 - 7: $(\text{flag}, a^*) \leftarrow \text{Navigation}(\text{VD}_{\text{out}}^*(u_i, R_{i+1}), v, f^*)$
 - 8: **if** $\text{flag} = \text{terminal}$ **then**
 - 9: **return** $\omega(a^*) + \text{Dist}(a^*, v, R_{i+1})$
 - 10: **else** (i.e., $\text{flag} = \text{nonterminal}$)
 - 11: **return** $\text{CentroidSearch}(\text{VD}_{\text{out}}^*(u_i, R_{i+1}), v, a^*)$
 - 12: **end if**
-

The procedure **CentroidSearch** is given $u_i \in \partial R_i$, $v \in R_{i+1}^{\text{out}}$, $\text{VD}_{\text{out}}^* = \text{VD}_{\text{out}}^*(u_i, R_{i+1})$ and a node f^* on the centroid decomposition of VD_{out}^* . It ultimately computes $u_{i+1} \in \partial R_{i+1}$ for which $v \in \text{Vor}(u_{i+1})$ and returns $\omega(u_{i+1}) + \text{Dist}(u_{i+1}, v, R_{i+1}) \leq \text{dist}_G(u_i, u_{i+1}) + \text{dist}_{R_{i+1}^{\text{out}}}(u_{i+1}, v) = \text{dist}_G(u_i, v)$. The algorithm is recursive, and bottoms out in one of two base cases (Line 5

or Line 9). The first way the recursion can end is if we reach the bottom of the centroid decomposition. If f^* is a leaf of the decomposition, it corresponds to an edge in VD_{out}^* separating the Voronoi cells of two sites, say s_1 and s_2 . At this point we know that either $u_{i+1} = s_1$ or $u_{i+1} = s_2$, and determine which case is true with two recursive calls to $\mathbf{Dist}(s_j, v, R_{i+1})$, $j \in \{1, 2\}$ (Lines 2–5). In general, f^* is dual to a trichromatic face f composed of three vertices y_0, y_1, y_2 in clockwise order, which are, respectively, in distinct Voronoi cells of s_0, s_1, s_2 . The three shortest s_j - y_j paths and f partition the vertices of R_{i+1}^{out} into six parts, namely the shortest s_j - y_j paths themselves, and the interiors of the regions bounded by ∂R_{i+1} , two of the s_j - y_j paths and an edge of f . See Figure 3. The **Navigation** function returns a pair (flag, a^*) that identifies which part v is in. If $\text{flag} = \text{terminal}$ then $a^* \in \{s_0, s_1, s_2\}$ is interpreted as a site, indicating that v lies on the shortest path from a^* to its f -vertex. In this case we return $\omega(a^*) + \mathbf{Dist}(a^*, v, R_{i+1}) = \text{dist}_G(u_i, v)$ with just one call to \mathbf{Dist} . If $\text{flag} = \text{nonterminal}$ then a^* is the correct child of f^* in the centroid decomposition. In particular, f^* is incident to three edges e_0^*, e_1^*, e_2^* dual to $\{y_0, y_2\}, \{y_1, y_0\}, \{y_2, y_1\}$. The children of f^* in the centroid decomposition are f_0^*, f_1^*, f_2^* , with f_j^* ancestral to e_j^* . We have $a^* = f_j^*$ if v lies to the right of the chord $(s_j, \dots, y_j, y_{j-1}, \dots, s_{j-1})$ in R_{i+1}^{out} . For example, in Figure 3, v lies to the right of the chord $(s_0, \dots, y_0, y_2, \dots, s_2)$. In this case we continue the search recursively from $a^* = f_0^*$.

LEMMA 3.2. *CentroidSearch correctly computes $\text{dist}_G(u_i, v)$.*

Proof. Define $f, y_j, s_j, e_j^*, f_j^*$ as usual, and let u_{i+1} be such that $v \in \text{Vor}(u_{i+1})$. The loop invariant is that in the subtree of the centroid decomposition rooted at f^* , there is *some* leaf edge on the boundary of the cell $\text{Vor}(u_{i+1})$. This is clearly true in the initial recursive call, when f^* is the root of the centroid decomposition. Suppose that **Navigation** tells us that v lies to the right of the oriented chord $C^* = (s_j, \dots, y_j, y_{j-1}, \dots, s_{j-1})$. Observe that since the s_j - y_j and s_{j-1} - y_{j-1} shortest paths are monochromatic, all edges of the centroid decomposition correspond to paths in G^* that lie strictly to the left or right of C^* , with the exception of e_j^* . Moreover, since $v \in \text{Vor}(u_{i+1})$, $\text{Vor}(u_{i+1})$ must be bounded by *some* edge that is either e_j^* or one entirely to the right of C^* , from which it follows that $f_j^* = a^*$ is ancestral to at least one edge bounding $\text{Vor}(u_{i+1})$. When f^* is a single edge on the boundary of $\text{Vor}(s_1), \text{Vor}(s_2)$ the loop invariant guarantees that either $u_{i+1} = s_1$ or $u_{i+1} = s_2$; suppose that $u_{i+1} = s_1$. It follows from the

specification of **Dist** (Eqn. (3.1)) and Lemma 3.1 that

$$\begin{aligned} d_1 &= \omega(s_1) + \mathbf{Dist}(s_1, v, R_{i+1}) \\ &\leq \text{dist}_G(u_i, s_1) + \text{dist}_{R_{i+1}^{\text{out}}}(s_1, v) \\ &= \text{dist}_G(u_i, v). \end{aligned}$$

Furthermore,

$$\begin{aligned} d_2 &= \omega(s_2) + \mathbf{Dist}(s_2, v, R_{i+1}) \\ &\geq \text{dist}_G(u_i, s_2) + \text{dist}_G(s_2, v) \\ &\geq \text{dist}_G(u_i, v), \end{aligned}$$

so in this base case **CentroidSearch** correctly returns $d_1 = \text{dist}_G(u_i, v)$. If **Navigation** ever reports that v is on an s_j - y_j path, then by definition $v \in \text{Vor}(s_j)$. By the specification of **Dist** (Eqn. (3.1)) and Lemma 3.1 we have

$$\begin{aligned} &\omega(s_j) + \mathbf{Dist}(s_j, v, R_{i+1}) \\ &\leq \text{dist}_G(u_i, s_j) + \text{dist}_{R_{i+1}^{\text{out}}}(s_j, v) \\ &= \text{dist}_G(u_i, v) \end{aligned}$$

and the base case on Lines 8–9 also works correctly. \square

Thus, the main challenge is to design an efficient **Navigation** function, i.e., to solve the restricted point location problem in R_{i+1}^{out} depicted in Figure 3. Whereas Charalampopoulos et al. [9] solve this problem using several *more* recursive calls to **Dist**, we give a new method to do this point location directly, in $O(\kappa \log^{1+o(1)} n)$ time per call to **Navigation**.

4 The Navigation Oracle

The input to **Navigation** is the same as **CentroidSearch**, except that f^* is guaranteed to correspond to a trichromatic face f . Define y_j, s_j, e_j, f_j , $j \in \{0, 1, 2\}$ as in the discussion of **CentroidSearch**. The **Navigation** function determines the location of v relative to f and the shortest s_j - y_j paths. It delegates nearly all the actual computation to two functions: **SitePathIndicator**, which returns a boolean indicating whether v is on the shortest s_j - y_j path, and **ChordIndicator**, which indicates whether v lies strictly to the right of the oriented chord $(s_j, \dots, y_j, y_{j-1}, \dots, s_{j-1})$. If so, we return the centroid child f_j^* of f^* in this region. Three calls each to **SitePathIndicator** and **ChordIndicator** suffice to determine the location of v .




Figure 3: Here f^* is a degree-3 vertex in $\text{VD}_{\text{out}}^*(u_i, R_{i+1})$, corresponding to a trichromatic face f on vertices y_0, y_1, y_2 , which are in the Voronoi cells of s_0, s_1, s_2 on the boundary $\partial R_{i+1}^{\text{out}}$. The shortest s_j - y_j paths partition $V(R_{i+1}^{\text{out}})$ into six parts: the three shortest paths and the three regions bounded by them and f . Let e_0^*, e_1^*, e_2^* be the edges in VD_{out}^* dual to $\{y_0, y_2\}, \{y_1, y_0\}, \{y_2, y_1\}$. In the centroid decomposition e_0^*, e_1^*, e_2^* are in separate subtrees of f^* . Let f_j^* be the child of f^* ancestral to e_j^* , which is either e_j^* itself, or a trichromatic face to the right of the “chord” $(s_j, \dots, y_j, y_{j-1}, \dots, s_{j-1})$. **CentroidSearch** locates the site whose Voronoi cell contains v via recursion. It calls **Navigation**, a function that finds which of the 6 parts contains v . If v lies on an s_j - y_j path the **CentroidSearch** recursion terminates; otherwise it recurses on the correct child f_j^* of f^* .

Algorithm 3 Navigation($\text{VD}_{\text{out}}^*(u_i, R_{i+1}), v, f^*$)

Input: The dual representation $\text{VD}_{\text{out}}^*(u_i, R_{i+1})$ of a Voronoi diagram, a vertex $v \in R_{i+1}^{\text{out}}$, and a centroid f^* in the centroid decomposition. The face f is on y_0, y_1, y_2 , which are in the Voronoi cells of s_0, s_1, s_2 , and f_j^* is the child of f^* containing the edge dual to $\{y_j, y_{j-1}\}$.

Output: (terminal, s_j) if v is on the shortest s_j - y_j path, or (nonterminal, f_j^*) where f_j^* is the child of f^* ancestral to an edge bounding v 's Voronoi cell.

- 1: $s_0, s_1, s_2 \leftarrow$ sites corresponding to f^*
 - 2: **for** $j = 0, 1, 2$ **do**
 - 3: **if** **SitePathIndicator**($\text{VD}_{\text{out}}^*(u_i, R_{i+1}), v, f^*, j$) **returns True** **then**
 - 4: **return** (terminal, s_j)
 - 5: **end if**
 - 6: **end for**
 - 7: **for** $j = 0, 1, 2$ **do**
 - 8: **if** **ChordIndicator**($\text{VD}_{\text{out}}^*(u_i, R_{i+1}), v, f^*, j$) **returns True** **then**
 - 9: **return** (nonterminal, f_j^*)
 - 10: **end if**
 - 11: **end for**
-

In Section 4.1 we formally introduce the notion of *chords* used informally above, as well as some related concepts like *laminar* sets of chords and *maximal* chords. In Section 4.2 we introduce parts (C)-(E) of the data structure used to support **Navigation**. The functions **SitePathIndicator** and **ChordIndicator** are presented in Sections 4.3 and 4.4.

4.1 Chords and Pieces. We begin by defining the key concepts of our point location method: *chords*, *laminar chord sets*, *pieces*, and the *occludes* relation.

DEFINITION 4.1. (Chords) Fix an R in the \vec{r} -decomposition and two vertices $c_0, c_1 \in \partial R$. An oriented simple path $\overrightarrow{c_0 c_1}$ is a chord of R^{out} if it is contained in R^{out} and is internally vertex-disjoint from ∂R . When the orientation is irrelevant we write it as $\overline{c_0 c_1}$.

DEFINITION 4.2. (Laminar Chord Sets) A set of chords \mathcal{C} for R^{out} is laminar (non-crossing) if for any two such chords $C = \overrightarrow{c_0 c_1}, C' = \overrightarrow{c_2 c_3}$, if there exists a $v \in (C \cap C') - \partial R$ then the subpaths from c_0 to v and from c_2 to v are identical; in particular $c_0 = c_2$.

The orientation of chords does not always coincide with a natural orientation of paths defined by the algorithm. For example, in Figure 3, the oriented chord $\overrightarrow{s_0 s_2} = (s_0, \dots, y_0, y_2, \dots, s_2)$ is composed of three parts: a shortest s_0 - y_0 path (whose natural orientation coincides with that of $\overrightarrow{s_0 s_2}$), the edge $\{y_0, y_2\}$ (which has




Figure 4: A laminar set of chords partition R^{out} into pieces. Observe that the chords separating pieces P_5 – P_9 overlap in certain prefixes. The piece tree is indicated by diamond vertices and pink edges. Note two pieces (e.g. P_5 and P_9) may share a boundary, but not be adjacent.

no natural orientation in this context), and the shortest s_2 - y_2 path (whose natural orientation is the reverse of its orientation in $\overrightarrow{s_0s_2}$). The orientation serves two purposes. In Definition 4.1 we can speak unambiguously about the parts of R^{out} to the right and left of $\overrightarrow{s_0s_2}$. In Definition 4.2 the role of the orientation is to ensure that the partition of R^{out} into pieces induced by \mathcal{C} can be represented by a tree, as we show in Lemma 4.1.

DEFINITION 4.3. (Pieces) A laminar chord set \mathcal{C} for R^{out} partitions the faces of R^{out} into pieces, excluding the face on ∂R . Two faces f, g are in the same piece iff f^* and g^* are connected by a path in $(R^{\text{out}})^*$ that avoids the duals of edges in \mathcal{C} and edges along the boundary cycle on ∂R . A piece is regarded as the subgraph induced by its faces, i.e., it includes their constituent vertices and edges. Two pieces P_1, P_2 are adjacent if there is an edge e on the boundary of P_1 and P_2 and e is in a unique chord of \mathcal{C} . See Figure 4.

LEMMA 4.1. Suppose \mathcal{C} is a laminar chord set for R^{out} , $\mathcal{P} = \mathcal{P}(\mathcal{C})$ is the corresponding piece set and \mathcal{E} are pairs of adjacent pieces. Then $\mathcal{T} = (\mathcal{P}, \mathcal{E})$ is a tree, called the piece tree induced by \mathcal{C} .

Proof. The claim is clearly true when \mathcal{C} contains zero or one chords, so we will try to reduce the general case to

this case via a peeling argument. We will find a piece P with degree 1 in \mathcal{T} , remove it and the chord bounding it, and conclude by induction that the truncated instance is a tree. Reattaching P implies \mathcal{T} is a tree.

Let $C = \overrightarrow{c_0c_1} \in \mathcal{C}$ be a chord such that no edge of any other chord appears strictly to one side of C , say to the right of C . Let P be the piece to the right of C . (In Figure 4 the chords bounding P_1, P_2, P_{11}, P_{12} would be eligible to be C .) Let $C = (c_0 = v_0, v_1, v_2, \dots, v_k = c_1)$ and v_{j^*} be such that the edges of the suffix (v_{j^*}, \dots, v_k) are on no other chord, meaning the vertices $\{v_{j^*+1}, \dots, v_{k-1}\}$ are on no other chord. Let g_j be the face to the left of (v_j, v_{j+1}) . It follows that there is a path from $g_{j^*}^*$ to g_{k-1}^* in $(R^{\text{out}})^*$ that avoids the duals of all edges in \mathcal{C} and along ∂R . All pieces adjacent to P contain some face among $\{g_{j^*}, \dots, g_{k-1}\}$, but these are in a single piece, hence P corresponds to a degree-1 vertex in \mathcal{T} . Let P be bounded by C and an interval B of the boundary cycle on ∂R . Obtain the “new” R^{out} by cutting along C and removing P , the new ∂R by substituting C for B , and the new chord set \mathcal{C} by removing C and trimming any chords that shared a non-empty prefix with C . By induction the resulting piece-adjacency graph is a tree; reattaching P as a degree-1 vertex shows \mathcal{T} is a tree. \square

DEFINITION 4.4. (Occludes Relation) Fix R^{out} , chord C , and two faces f, g , neither of which is the hole defined by ∂R . If f and g are on opposite sides of C , we say that from vantage f , C occludes g . Let \mathcal{C} be a set of chords. We say $C \in \mathcal{C}$ is maximal in \mathcal{C} with respect to a vantage f if there is no $C' \in \mathcal{C}$ such that C' occludes a strict superset of the faces that C occludes. (Note that the orientation of chords is irrelevant to the occludes relation.)

It follows from Definition 4.4 that if \mathcal{C} is laminar, the set of maximal chords with respect to f are exactly those chords whose own suffixes are on the boundary of f 's piece in $\mathcal{P}(\mathcal{C})$.

We can also speak unambiguously about a chord C occluding a vertex or edge not on C , from a certain vantage. Specifically, we can say that from some vantage, C occludes an interval of the boundary cycle on ∂R , say according to a clockwise traversal around the hole on ∂R in R^{out} .⁶ This will be used in the **ChordIndicator** procedure of Section 4.4.2.

4.2 Data Structures for Navigation. Parts (C)–(E) of the data structure are used to implement the **SitePathIndicator** and **ChordIndicator** functions.

⁶This is one place where we use the assumption that all boundary holes are simple cycles.

- (C) (**More Voronoi Diagrams**) For each $i \in [1, m - 1]$, each $R_i \in \mathcal{R}_i$, and each $q \in \partial R_i$, we store $\text{VD}_{\text{out}}^*(q, R_i)$, which is $\text{VD}^*[R_i^{\text{out}}, \partial R_i, \omega]$, where $\omega(s) = \text{dist}_G(q, s)$. The total space for these diagrams is $\tilde{O}(n)$ and dominated by part (B).
- (D) (**Chord Trees; Piece Trees**) For each $i \in [1, m - 1]$, each $R_i \in \mathcal{R}_i$, and source $q \in \partial R_i$, we store the SSSP tree from q with respect to G induced by ∂R_i as a *chord tree* $T_q^{R_i}$. In particular, the parent of $x \in \partial R_i$ in $T_q^{R_i}$ is the nearest ancestor in the SSSP tree from q that lies on ∂R_i . Every edge of $T_q^{R_i}$ is designated a *chord* if the corresponding path is contained in R_i^{out} but not in R_i , or a *non-chord* otherwise. Define $\mathcal{C}_q^{R_i}$ to be the set of all chords in $T_q^{R_i}$, oriented away from q ; this is clearly a laminar set since shortest paths are unique and all prefixes of shortest paths are shortest paths. Define $\mathcal{P}_q^{R_i}$ to be the corresponding partition of R_i^{out} into pieces, and $\mathcal{T}_q^{R_i}$ the corresponding piece tree. Define $T_q^{R_i}[x]$ to be the path from q to x in $T_q^{R_i}$, $\mathcal{C}_q^{R_i}[x]$ the corresponding chord set, and $\mathcal{P}_q^{R_i}[x]$ the corresponding piece set.

The data structure answers the following queries

MaximalChord(R_i, q, x, P, P'): We are given R_i , $q, x \in \partial R_i$, a piece $P \in \mathcal{P}_q^{R_i}$, and possibly another piece $P' \in \mathcal{P}_q^{R_i}$ (which may be **Null**). If P' is **Null**, return any maximal chord in $\mathcal{C}_q^{R_i}[x]$ from vantage P . If P' is not **Null**, return the maximal chord in $\mathcal{C}_q^{R_i}[x]$ (if any) that occludes P' from vantage P .

AdjacentPiece(R_i, q, e): Here e is an edge on the boundary cycle on ∂R_i . Return the unique piece in $\mathcal{P}_q^{R_i}$ with e on its boundary.⁷

- (E) (**Site Tables; Side Tables**) For each i and Voronoi diagram $\text{VD}_{\text{out}}^* = \text{VD}_{\text{out}}^*(u', R_i)$ from part (B) or (C). Let f^* be any node in the centroid decomposition of VD_{out}^* , with $y_j, s_j, j \in \{0, 1, 2\}$ defined as usual, and let $R_{i'} \in \mathcal{R}_{i'}$ be any ancestor of $R_i, i' \geq i$. Fix $j \in \{0, 1, 2\}$, define q and x to be the *first* and *last* vertices on the shortest s_j - y_j path that lie on $\partial R_{i'}$. For all $\text{VD}_{\text{out}}^*, f^*, j$ and i' , we store (q, x) and $\text{dist}_G(u', x)$ with respect to them.

For all $\text{VD}_{\text{out}}^*, f^*, j$ and i' , we also store whether $R_{i'}^{\text{out}}$ lies to the left or right of the site-centroid-site chord $\overrightarrow{s_j y_j y_{j-1} s_{j-1}}$ in R_i^{out} , or **Null** if the relationship cannot be determined, i.e., if the chord

crosses $\partial R_{i'}$. These tables increase the space of (B) and (C) by a small $O(m)$ factor.

Part (D) of the data structure is the only one that is non-trivial to store compactly. Our strategy is as follows. We fix R_i and $q \in \partial R_i$ and build a dynamic data structure for these operations relative to a dynamic subset $\hat{\mathcal{C}} \subseteq \mathcal{C}_q^{R_i}$ subject to the insertion and deletion of chords in $O(\log |\partial R_i|)$ time. By inserting/deleting $O(|\partial R_i|)$ chords in the correct order, we can arrange that $\hat{\mathcal{C}} = \mathcal{C}_q^{R_i}[x]$ at some point in time, for every $x \in \partial R_i$. Using the generic persistence technique for RAM data structures (see [12]) we can answer **MaximalChord** queries relative to $\mathcal{C}_q^{R_i}[x]$ in $O(\log |\partial R_i| \log \log |\partial R_i|)$ time.⁸

LEMMA 4.2. *Part (D) of the data structure can be stored in $O(mn \log n)$ total space and answer **MaximalChord** queries in $O(\log n \log \log n)$ time and **AdjacentPiece** queries in $O(1)$ time.*

Proof. We first address **MaximalChord**. Let $\mathcal{T} = \mathcal{T}_q^{R_i}$ be the piece tree. The edges of \mathcal{T} are in 1-1 correspondence with the chords of $\mathcal{C} = \mathcal{C}_q^{R_i}$ and if $P, P' \in \mathcal{P} = \mathcal{P}_q^{R_i}$ are two pieces, the path from P to P' in \mathcal{T} crosses exactly those chords that occlude P' from vantage P (and vice versa). We will argue that to implement **MaximalChord** it suffices to design an efficient dynamic data structure for the following problem; initially all edges are *unmarked*.

Mark(e) Mark an edge $e \in E(\mathcal{T})$.

Unmark(e) Unmark e .

LastMarked(P', P) Return the *last* marked edge on the path from P' to P , or **Null** if all are unmarked.

By doing a depth-first traversal of the chord tree $T_q^{R_i}$, marking/unmarking chords as they are encountered, the set $\{e \in E(\mathcal{T}) \mid e \text{ is marked}\}$ will be equal to $\mathcal{C}_q^{R_i}[x]$ precisely when x is first encountered in DFS. To answer a **MaximalChord**(R_i, q, x, P, P') query we interact with the state of the data structure when the marked set is $\mathcal{C}_q^{R_i}[x]$. If P' is not null we return **LastMarked**(P', P). Otherwise we pick an arbitrary (marked) chord $C \in \mathcal{C}_q^{R_i}[x]$, get the adjacent pieces P'_1, P'_2 on either side of C , then query **LastMarked**(P'_1, P) and **LastMarked**(P'_2, P). At least one of these queries will return a chord and that

⁸Our data structure works in the pointer machine model, but it has unbounded in-degrees so the theorem of Driscoll et al. [15, 35] cannot be applied directly. It is probably possible to improve the bound to $O(\log |\partial R_i|)$ but this is not a bottleneck in our algorithm.

⁷This is another place where we use the assumption that holes are bounded by simple cycles.

chord is maximal w.r.t. vantage P . (Note that C must separate P from either P'_1 or P'_2 .)

The full version [30] shows how to support these operations in $O(\log n)$ worst case time. A more sophisticated data structure of Brodal et al. [6] improves this to optimal $O(\log n / \log \log n)$ time.

For fixed $R_i, q \in \partial R_i$ there are $O(|\partial R_i|)$ **Mark** and **Unmark** operations, each taking $O(\log n)$ time. Over all choices of i, R_i , and q the total update time is $O(mn \log n)$. After applying generic persistence for RAM data structures (see [12]) the space becomes $O(mn \log n)$ and the query time for **LastMarked** becomes $O(\log n \log \log n)$.

Turning to **AdjacentPiece**(R_i, q, e), there are $|\partial R_i|^2$ choices of (q, e) . Hence all answers can be precomputed in a lookup table in $O(mn)$ space. \square

4.3 The SitePathIndicator Function. The **SitePathIndicator** function is relatively simple. We are given $\text{VD}_{\text{out}}^*(u_i, R_{i+1})$, $v \in R_{i+1}^{\text{out}}$, a centroid $f^* \in R_{i+1}^{\text{out}}$, f being a trichromatic face on y_0, y_1, y_2 , which are, respectively, in the Voronoi cells of $s_0, s_1, s_2 \in \partial R_{i+1}$, and an index $j \in \{0, 1, 2\}$. We would like to know if v is on the shortest s_j -to- y_j path. Recall that t is such that $v \notin R_t$ but $v \in R_{t+1}$.

Using the lookup tables in part (E) of the data structure, we find the first and last vertices (q and x) of ∂R_t on the s_j - y_j path. If q, x do not exist then v is certainly not on the s_j - y_j path (Line 4). Using parts (A,C,E) of the data structure, we invoke **PointLocate** to find the last point z of ∂R_t on the shortest path (in G) from q to v . (See Lemma 3.1.) If z is not on the path from q to x in $T_q^{R_t}$, stored in Part (D)), then once again v is certainly not on the s_j - y_j path (Line 8). So we may assume z lies on the q - x path. If $z = x$ then there are three cases to consider, depending on whether the destination y_j of the path is in $R_t^{\text{out}} \cap R_{t+1}$, or in R_{t+1}^{out} , or in R_t . If $y_j \in R_t^{\text{out}} \cap R_{t+1}$ we let $x' = y_j$; if $y_j \notin R_{t+1}$ we let x' be the last vertex of ∂R_{t+1} encountered on the shortest s_j - y_j path (part (E)); and if $y_j \notin R_t^{\text{out}}$ we let $x' = x$. In all cases, x' is the last vertex of the shortest s_j - y_j path that is contained in the relevant subgraph $R_t^{\text{out}} \cap R_{t+1}$. (Figure 5(a,b) illustrates the first two possibilities for x' .) Now v is on the s_j - y_j path iff it is on the x - x' shortest path, which can be answered using part (A) of the data structure (Lines 13, 15). (Figure 5(b) illustrates one way for v to appear on the x - x' path.) In the remaining case z is on the shortest q - x path but is not x , meaning the child z' of z on $T_q^{R_t}[x]$ is well defined. If $\overrightarrow{zz'}$ is a chord (corresponding to a path in R_t^{out}) then v is on the shortest s_j - y_j path iff it is on the shortest z - z' path

in R_t^{out} , which, once again, can be answered with part (A) of the data structure (Lines 19, 21). See Figure 5(a) for an illustration of this case.

REMARK 4.1. *Strictly speaking we cannot apply Lemma 2.2 (Gawrychowski et al. [19]) since we do not have an MSSP structure for all of R_t^{out} . Part (A) only handles distance/LCA queries when the query vertices are in $R_t^{\text{out}} \cap R_{t+1}$. It is easy to make Gawrychowski et al.'s algorithm work using parts (A) and (E) of the data structure. See the discussion at the end of Section 4.4.3.*

Algorithm 4 SitePathIndicator($\text{VD}_{\text{out}}^*(u_i, R_{i+1}), v, f^*, j$)

Input: The dual representation $\text{VD}_{\text{out}}^*(u_i, R_{i+1})$ of a Voronoi diagram, a vertex $v \in R_{i+1}^{\text{out}}$, and an s_j -to- y_j site-centroid shortest path (s_j, y_j are with respect to f^*) in VD^* .

Output: **True** if v is on s_j -to- y_j shortest path, or **False** otherwise.

- 1: $R_t \leftarrow$ the ancestor of R_i s.t. $v \notin R_t, v \in R_{t+1}$.
 - 2: $(q, x) \leftarrow$ first and last ∂R_t vertices on the shortest s_j - y_j path.
 - 3: **if** q, x are **Null** **then**
 - 4: **return False**
 - 5: **end if**
 - 6: $z \leftarrow$ **PointLocate**($\text{VD}_{\text{out}}^*(q, R_t), v$)
 - 7: **if** z is not on $T_q^{R_t}[x]$ **then**
 - 8: **return False**
 - 9: **end if**
 - 10: **if** $z = x$ **then**
 - 11: Let x' be

$$\begin{cases} y_j, & \text{if } y_j \in R_t^{\text{out}} \cap R_{t+1}; \\ \text{last } \partial R_{t+1} \text{ vertex on } s_j\text{-}y_j \text{ path,} & \text{if } y_j \notin R_{t+1}; \\ x, & \text{if } y_j \notin R_t^{\text{out}}. \end{cases}$$
 - 12: **if** v is on the shortest x - x' path **then**
 - 13: **return True**
 - 14: **end if**
 - 15: **return False**
 - 16: **end if**
 - 17: $z' \leftarrow$ the child of z on $T_q^{R_t}[x]$
 - 18: **if** $\overrightarrow{zz'}$ is a chord in $\mathcal{C}_q^{R_t}[x]$ and v is on the shortest z - z' path in R_t^{out} **then**
 - 19: **return True**
 - 20: **end if**
 - 21: **return False**
-

4.4 The ChordIndicator Function. The **ChordIndicator** function is given $\text{VD}_{\text{out}}^*(u_i, R_{i+1})$, $v \in R_{i+1}^{\text{out}}$, a centroid f^* , with $\{y_j, s_j\}$ defined as usual,




Figure 5: **(a)** If $z = x$ and y_j is not in R_{t+1} , x' is the last boundary vertex of ∂R_{t+1} on the s_j - y_j path. **(b)** If $z = x$ and y_j is in $R_t^{\text{out}} \cap R_{t+1}$ then $x' = y_j$. (Not depicted: if $y_j \in R_t$ then $x' = x$.) We test whether v is on the shortest x - x' path. If $z \neq x$ then z' is well defined and the position of y_j is immaterial; we test whether v is on the shortest z - z' path (depicted in **(a)**).

and an index $j \in \{0, 1, 2\}$. The goal is to report whether v lies to right of the oriented *site-centroid-site* chord

$$C^* = \overrightarrow{s_j y_j y_{j-1} s_{j-1}},$$

which is composed of the shortest s_j - y_j and s_{j-1} - y_{j-1} paths, and the single edge $\{y_j, y_{j-1}\}$. Note that C^* is a simple path since the shortest s_j - y_j and s_{j-1} - y_{j-1} paths belong to different Voronoi cells. See Figure 3 as an illustration. It is guaranteed that v does not lie on C^* , as this case is already handled by the **SitePathIndicator** function.

Figure 6 illustrates why this point location problem is so difficult. Since we know $v \in R_{t+1}$ but not in R_t , we can narrow our attention to $R_t^{\text{out}} \cap R_{t+1}$. However the projection of C^* onto R_t^{out} can touch the boundary ∂R_t an arbitrary number of times. Define \mathcal{C} to be the set of oriented chords of R_t^{out} obtained by projecting C^* onto R_t^{out} .

Luckily \mathcal{C} has some structure. Let (q_j, x_j) and (q_{j-1}, x_{j-1}) be the first and last ∂R_t vertices on the shortest s_j - y_j and s_{j-1} - y_{j-1} paths, respectively. (One or both of these pairs may not exist.) The chords of \mathcal{C} are in one-to-one correspondence with the chords of $\mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3$, defined below, but as we will see, sometimes with their orientation reversed.

\mathcal{C}_1 : By definition $\mathcal{C}_1 = C_{q_j}^{R_t}[x_j]$ contains all the chords on the path from q_j to x_j , stored in part (D) of the data structure. Moreover, the orientation of \mathcal{C}_1

agrees with the orientation of C^* . The blue chords of Figure 6(a) are isolated as \mathcal{C}_1 in Figure 6(b).

\mathcal{C}_2 : By definition $\mathcal{C}_2 = C_{q_{j-1}}^{R_t}[x_{j-1}]$ contains all the chords on the path from q_{j-1} to x_{j-1} . The red chords of \mathcal{C} in Figure 6(a) are *represented* by chords \mathcal{C}_2 , but with reversed orientation. Figure 6(c) depicts \mathcal{C}_2 .

\mathcal{C}_3 : This is the singleton set containing the oriented chord $\overrightarrow{x_j x_{j-1}}$ consisting of the shortest x_j - y_j and x_{j-1} - y_{j-1} paths and the edge $\{y_j, y_{j-1}\}$.

The chord set \mathcal{C} partitions R_t^{out} into a piece set \mathcal{P} , with one such piece $P \in \mathcal{P}$ containing v . (Remember that v is not on C^* .) We can also consider the piece sets $\mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3$ generated by $\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3$. Let $P_1 \in \mathcal{P}_1, P_2 \in \mathcal{P}_2, P_3 \in \mathcal{P}_3$ be the pieces containing v . Since, ignoring orientation, $\mathcal{C} = \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3$, it must be that $P = P_1 \cap P_2 \cap P_3$. In order to determine whether v is to the right of C^* , it suffices to find some chord $C \in \mathcal{C}$ bounding P and ask whether v is to the right of C . Thus, C must also be on the boundary of one of P_1, P_2 , or P_3 .

The high-level strategy of **ChordIndicator** is as follows. First, we will find some piece $P'_1 \in \mathcal{P}_{q_j}^{R_t}$ that is contained in P_1 using the procedure **PieceSearch** described below, in Section 4.4.1. The chords of \mathcal{C}_1 bounding P_1 are precisely the *maximal* chords in \mathcal{C}_1 from vantage P'_1 . Using **MaximalChord** (part (D)) we will

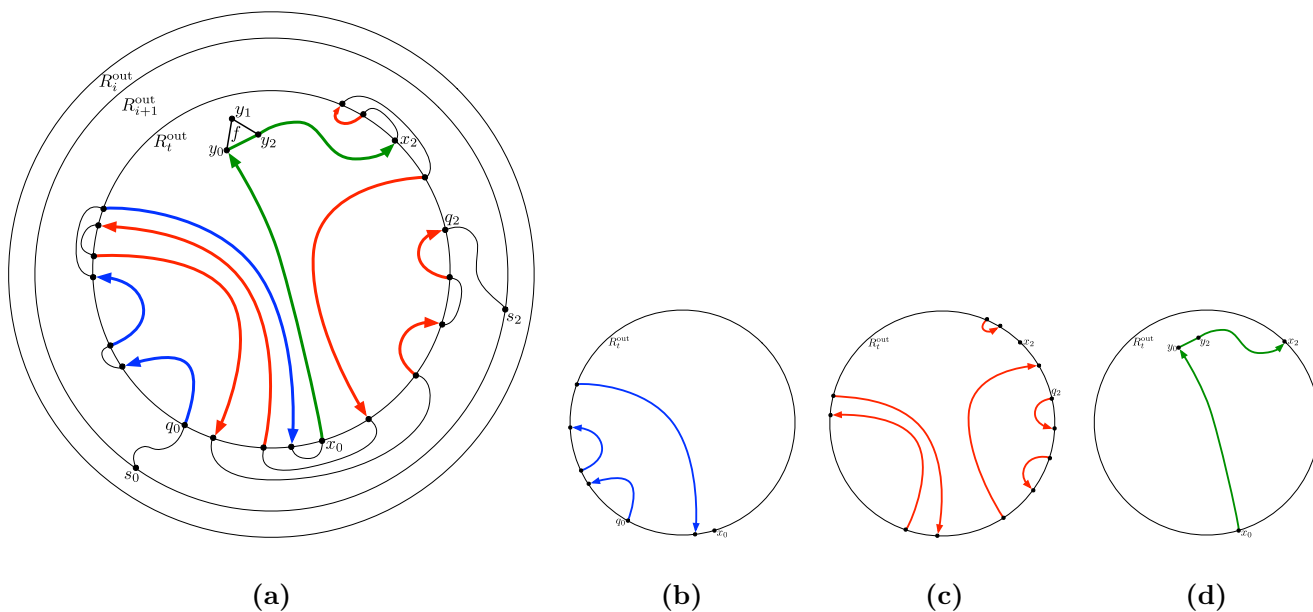


Figure 6: (a) The projection of a site-centroid-site chord $C^* = \overrightarrow{s_j y_j y_{j-1} s_{j-1}}$ of R_{i+1}^{out} onto R_t^{out} yields a set \mathcal{C} of chords of R_t^{out} , partitioned into three classes. Let q_j, x_j and q_{j-1}, x_{j-1} be the first and last ∂R_t -vertices on the s_j - y_j and s_{j-1} - y_{j-1} paths. (b) \mathcal{C}_1 : all chords in $T_{q_j}^{R_t}[x_j]$. (c) \mathcal{C}_2 : all chords in $T_{q_{j-1}}^{R_t}[x_{j-1}]$. Their orientation is the reverse of their counterparts in C^* . (d) \mathcal{C}_3 : the single chord $\overrightarrow{x_j y_j y_{j-1} x_{j-1}}$.

find a candidate chord $C_1 \in \mathcal{C}_1$, and one edge e on the boundary cycle of ∂R_t occluded by C_1 from vantage P'_1 . Turning to \mathcal{C}_2 , we use **AdjacentPiece** to find the piece $P_e \in \mathcal{P}_{q_{j-1}}^{R_t}$ adjacent to e . Then, using **PieceSearch** and **MaximalChord** again, we find a $P'_2 \in \mathcal{P}_{q_{j-1}}^{R_t}$ contained in P_2 and the maximal chord C_2 occluding P_e from vantage P'_2 . Let C_3 be the singleton chord in \mathcal{C}_3 . We determine an “eligible” chord $C_\ell \in \{C_1, C_2, C_3\}$, decide whether v lies to the right of C_ℓ , and return this answer if $\ell \in \{1, 3\}$ or reverse it if $\ell = 2$. Recall that chords in \mathcal{C}_2 have the opposite orientation as their counterparts in \mathcal{C} .

PieceSearch is presented in Section 4.4.1 and **ChordIndicator** in Section 4.4.2.

4.4.1 PieceSearch. We are given a region R_t , a vertex $v \in R_t^{\text{out}} \cap R_{t+1}$, and two vertices $q, x \in \partial R_t$. We must locate *any* piece $P' \in \mathcal{P}_q^{R_t}$ that is contained in the unique piece $P \in \mathcal{P}_q^{R_t}[x]$ containing v . The first thing we do is find the *last* ∂R_t vertex z on the shortest path from q to v , which can be found with a call to **PointLocate** on $\text{VD}_{\text{out}}^*(q, R_t)$. (This uses parts (A,C,E) of the data structure.) The shortest path from z to v cannot cross any chord in $\mathcal{C}_q^{R_t}[x]$ (since they are part of a shortest path tree), but it can coincide with a prefix of some chord in $\mathcal{C}_q^{R_t}[x]$. Thus, if no chord of $\mathcal{C}_q^{R_t}[x]$ is incident to z , then we are free to return *any* piece containing z . (There may be multiple options if z is an endpoint of a chord in $\mathcal{C}_q^{R_t}$. This case is depicted in Fig-

ure 7. When $z = z_0$, we know that $v \in P_5 \cup \dots \cup P_9$ and return any piece containing z .) In general z may be incident to up to two chords $C_1, C_2 \in \mathcal{C}_q^{R_t}[x]$. (This occurs when the shortest q - x path touches ∂R_t at z without leaving R_t^{out} .) In this case we determine which side of C_1 and C_2 v is on (using parts (A) and (E) of the data structure; see Lemma 4.3 in Section 4.4.3 for details) and return the appropriate piece adjacent to C_1 or C_2 . This case is depicted in Figure 7 with $z = z_1$; the three possible answers coincide with $v \in \{v_1, v_2, v_3\}$.

Algorithm 5 PieceSearch(R_t, q, x, v)

Input: A region R_t , two vertices $q, x \in \partial R_t$, and a vertex v not on the q -to- x shortest path in G .

Output: A piece $P' \in \mathcal{P}_q^{R_t}$, which is a subpiece of the unique piece $P \in \mathcal{P}_q^{R_t}[x]$ containing v .

- 1: $z \leftarrow \text{PointLocate}(\text{VD}_{\text{out}}^*(q, R_t), v)$
 - 2: **if** z is not the endpoint of any chord in $\mathcal{C}_q^{R_t}[x]$ **then**
 - 3: **return** any piece in $\mathcal{P}_q^{R_t}$ containing z .
 - 4: **end if**
 - 5: $C_1, C_2 \leftarrow$ two chords in $\mathcal{C}_q^{R_t}[x]$ adjacent to z (C_2 may be **Null**)
 - 6: Determine whether v is to the left or right of C_1 and C_2 .
 - 7: **return** a piece adjacent to C_1 or C_2 that respects the queries of Line 6.
-




Figure 7: Solid chords are in $\mathcal{C}_q^{R_t}[x]$. Dashed chords are in $\mathcal{C}_q^{R_t}$ but not $\mathcal{C}_q^{R_t}[x]$. When $z = z_0, v = v_0$, the piece in $\mathcal{P}_q^{R_t}[x]$ containing v is the union of P_5 – P_9 . **PieceSearch** reports any piece containing z_0 . When $z = z_1, v \in \{v_1, v_2, v_3\}$, z is incident to two chords C_1, C_2 . **PieceSearch** decides which side of C_1, C_2 v is on (see Lemma 4.3), and returns the appropriate piece adjacent to C_1 or C_2 .

4.4.2 ChordIndicator. Let us walk through the **ChordIndicator** function. If $C^* = \overrightarrow{s_j y_j y_{j-1} s_{j-1}}$ does not touch the interior of R_t^{out} then the left-right relationship between C^* and $v \notin R_t$ is known, and stored in part (E) of the data structure. If this is the case the answer is returned immediately, at Line 3. A relatively simple case is when \mathcal{C}_1 and \mathcal{C}_2 are empty, and $\mathcal{C} = \mathcal{C}_3$ consists of just one chord $C_3 = \overrightarrow{x_j y_j y_{j-1} x_{j-1}}$. We determine whether v is to the right or left of C_3 and return this answer (Line 8). (Lemma 4.3 in Section 4.4.3 explains how to test whether v is to one side of a chord.) Thus, without loss of generality we can assume $\mathcal{C}_1 \neq \emptyset$ and \mathcal{C}_2 may or may not be empty.

Recall that P_1 is v 's piece in $\mathcal{P}_{q_j}^{R_t}[x_j]$. Using **PieceSearch** we find a piece $P'_1 \subseteq P_1$ in the more refined partition $\mathcal{P}_{q_j}^{R_t}$ and find a maximal chord $C_1 \in \mathcal{C}_1$ from vantage P'_1 , and hence from vantage v as well. We regard ∂R_t as circularly ordered according to a clockwise walk around the hole on ∂R_t in R_t^{out} . The chord C_1 occludes an interval I_1 of ∂R_t from vantage v . If C_1 is *not* one of the chords bounding P , then C_3 or some $C_2 \in \mathcal{C}_2$ must occlude a superset I_2 of I_1 , so we will attempt to find such a C_2 , as follows.

Let e be the first edge on the boundary cycle occluded by C_1 , i.e., e joins the first two elements of I_1 . Using **AdjacentPiece** we find the unique piece

Algorithm 6 ChordIndicator($\text{VD}_{\text{out}}^*(u_i, R_{i+1}), v, f^*, j$)

Input: The dual representation $\text{VD}_{\text{out}}^* = \text{VD}_{\text{out}}^*(u_i, R_{i+1})$ of a Voronoi diagram, a centroid f^* in VD_{out}^* with face f on vertices y_0, y_1, y_2 , which are in the Voronoi cells of s_0, s_1, s_2 , an index $j \in \{0, 1, 2\}$, and a vertex $v \in R_{i+1}^{\text{out}}$ that does not lie on the site-centroid-site chord $C^* = \overrightarrow{s_j y_j y_{j-1} s_{j-1}}$.

Output: **True** if v lies to the right of C^* , and **False** otherwise.

- 1: $R_t \leftarrow$ the ancestor of R_i s.t. $v \notin R_t, v \in R_{t+1}$. \mathcal{C} is the projection of C^* onto R_t^{out} .
 - 2: **if** the left/right relationship between R_t^{out} and $C^* = \overrightarrow{s_j y_j y_{j-1} s_{j-1}}$ is known **then**
 - 3: **return** stored **True/False** answer.
 - 4: **end if**
 - 5: $(q_j, x_j) \leftarrow$ first and last ∂R_t -vertices on shortest s_j - y_j path.
 - 6: $(q_{j-1}, x_{j-1}) \leftarrow$ first and last ∂R_t -vertices on shortest s_{j-1} - y_{j-1} path.
 - 7: **if** $\mathcal{C}_1 = \mathcal{C}_2 = \emptyset$ **then**
 - 8: **return True** if v is to the right of the \mathcal{C}_3 -chord $\overrightarrow{x_j y_j y_{j-1} x_{j-1}}$, or **False** otherwise.
 - 9: **end if**
 - 10: $P'_1 \leftarrow$ **PieceSearch**(R_t, q_j, x_j, v)
 - 11: $C_1 \leftarrow$ **MaximalChord**($R_t, q_j, x_j, P'_1, \perp$)
 - 12: $I_1 \leftarrow$ the clockwise interval of hole ∂R_t occluded by C_1 from vantage v .
 - 13: $e \leftarrow$ edge joining first two elements of I_1 .
 - 14: $P_e \leftarrow$ **AdjacentPiece**(R_t, q_{j-1}, e)
 - 15: $P'_2 \leftarrow$ **PieceSearch**(R_t, q_{j-1}, x_{j-1}, v)
 - 16: $C_2 \leftarrow$ **MaximalChord**($R_t, q_{j-1}, x_{j-1}, P'_2, P_e$)
 - 17: $I_2 \leftarrow$ the clockwise interval of hole ∂R_t occluded by C_2 from vantage v .
 - 18: $C_3 \leftarrow$ single chord in \mathcal{C}_3 , if any.
 - 19: $I_3 \leftarrow$ the clockwise interval of hole ∂R_t occluded by C_3 from vantage v .
 - 20: $\ell \leftarrow$ index such that I_ℓ covers e , and $|I_\ell|$ is maximum.
 - 21: **if** v is to the right of C_ℓ and $\ell \in \{1, 3\}$ or v is to the left of C_ℓ and $\ell = 2$ **then**
 - 22: **return True**
 - 23: **end if**
 - 24: **return False**
-

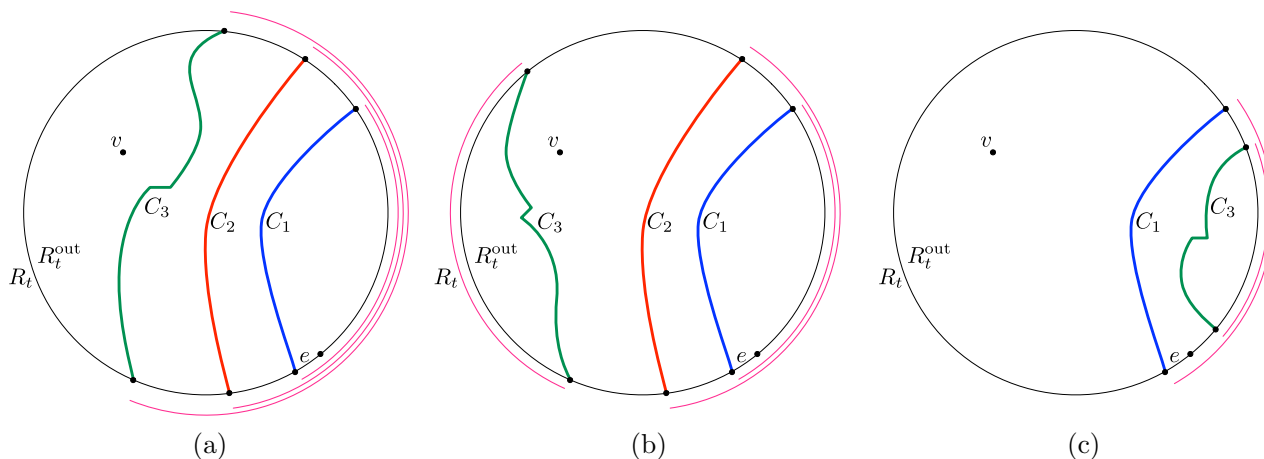


Figure 8: The intervals I_1, I_2, I_3 are represented as pink circular arcs. The edge e is the first edge of I_1 in a clockwise walk around the hole bounded by ∂R_t in R_t^{out} . (Note that in this drawing the hole on ∂R_t is the infinite face. Thus, a clockwise walk around ∂R_t *looks like* a counter-clockwise walk in the plane.) In (a) C_2 exists and C_3 is eligible since $I_3 \supset I_2 \supset I_1$. In (b) C_2 exists, but C_3 occludes an interval I_3 that does not contain e , so C_2 is an eligible chord. In (c) C_2 is **Null**, and C_3 does not occlude e from v , so C_1 is the only eligible chord. (In the figure $I_3 \subset I_1$ but it could also be as in (b), with I_3 disjoint from I_1 .)

$P_e \in \mathcal{P}_{q_{j-1}}^{R_t}$ with e on its boundary. Using **PieceSearch** again we find $P'_2 \in \mathcal{P}_{q_{j-1}}^{R_t}$ contained in P_2 , and using **MaximalChord** again, we find the maximal chord $C_2 \in \mathcal{C}_2$ that occludes P_e from vantage P'_2 , and hence from vantage v as well. Observe that since all chords in \mathcal{C}_2 are vertex-disjoint from C_1 , if $C_2 \neq \mathbf{Null}$ then C_2 must occlude a strictly larger interval $I_2 \supset I_1$ of ∂R_t . (If C_2 is **Null** then $I_2 = \emptyset$.) It may be that C_1 and C_2 are both not on the boundary of P , but the only way that could occur is if $C_3 \in \mathcal{C}_3$ occludes a superset of I_1 and I_2 on the boundary ∂R_t . We check whether v lies to the right or left of C_3 and let I_3 be the interval of ∂R_t occluded by C_3 from vantage v . If I_3 does not cover e , then we cannot conclude that C_3 is superior than C_1/C_2 . Thus, we find the chord $C_\ell \in \{C_1, C_2, C_3\}$ that covers e and maximizes $|I_\ell|$. C_ℓ must be on the boundary of P , so the left-right relationship between v and C^* is exactly the same as the left-right relationship between v and C_ℓ , if $\ell \in \{1, 3\}$, and the reverse of this relationship if $\ell = 2$ since chords in \mathcal{C}_2 have the opposite orientation as their subpath counterparts in C^* . Figure 8 illustrates how ℓ could take on all three values.

4.4.3 Side Queries. Lemma 4.3 explains how we test whether v is to the right or left of a chord, which is used in both **PieceSearch** and **ChordIndicator**.

LEMMA 4.3. *For any $C \in \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3$ and v not on C , we can test whether v lies to the right or left of C in $O(\kappa \log \log n)$ time, using parts (A) and (E) of the data structure.*

Proof. There are several cases.

Case 1. Suppose that $C = \overrightarrow{c_0 c_1} \in \mathcal{C}_1 \cup \mathcal{C}_2$ corresponds to the shortest path from c_0 to c_1 in R_t^{out} , $c_0, c_1 \in \partial R_t$. Let c'_0, c'_1 be pendant vertices attached to c_0, c_1 embedded inside the face of R_t^{out} bounded by ∂R_t . The shortest c'_0-v and $c'_0-c'_1$ paths branch at some point. We ask the MSSP structure (part (A)) for the least common ancestor, w , of v and c'_1 in the shortcutted SSSP tree rooted at c'_0 . This query also returns the two tree edges $e_v, e_{c'_1}$ leading to v and c'_1 , respectively. Let e_w be the edge connecting w to its parent.⁹ If the clockwise order around w is $e_w, e_{c'_1}, e_v$ then v lies to the right of $\overrightarrow{c_0 c_1}$; otherwise it lies to the left. Note that if the shortest $c'_0-c'_1$ and c'_0-v paths in G branch at a point in R_{t+1}^{out} , then w will be the nearest ancestor of the branchpoint on ∂R_{t+1} and one or both of $e_v, e_{c'_1}$ may be “shortcut” edges in the MSSP structure. See Figure 9(a) for a depiction of this case.

Case 2. Now suppose $C = \overrightarrow{x_j y_j y_{j-1} x_{j-1}}$ is the one chord in \mathcal{C}_3 . Consider the following distance function \hat{d} for vertices in $z \in R_t^{\text{out}}$: $\hat{d}(z) = \min\{\text{dist}_G(u_i, x_j) + \text{dist}_{R_t^{\text{out}}}(x_j, z), \text{dist}_G(u_i, x_{j-1}) + \text{dist}_{R_t^{\text{out}}}(x_{j-1}, z)\}$. Observe that the terms involving u_i are stored in part (E) and, if $z \in R_t^{\text{out}} \cap R_{t+1}$, the other terms can be queried in $O(\kappa \log \log n)$ time using part (A). It follows that the

⁹The purpose of adding c'_0, c'_1 is to make sure all three edges $e_w, e_v, e_{c'_1}$ exist. The vertices c'_0, c'_1 are not represented in the MSSP structure. The edges (c'_0, c_0) and (c_1, c'_1) can be simulated by inserting them between the two boundary edges on ∂R_t adjacent to c_0 and c_1 , respectively.

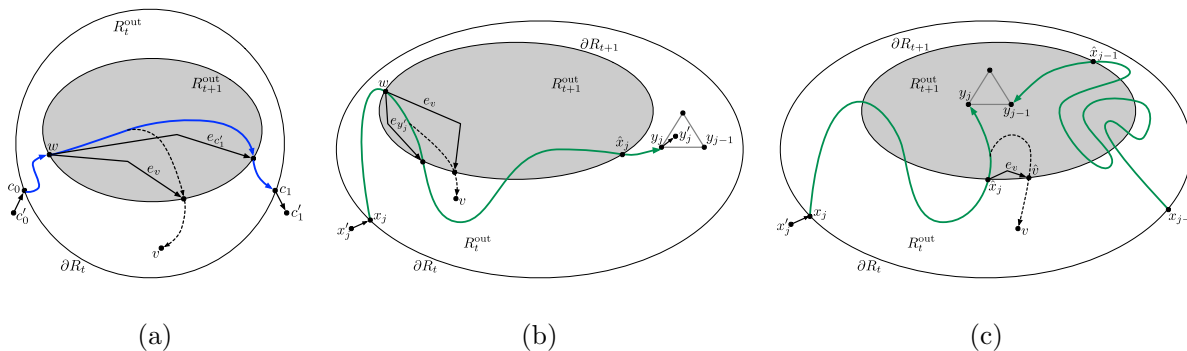


Figure 9: (a) The chord $C \in \mathcal{C}_1 \cup \mathcal{C}_2$ corresponds to a shortest path, which may pass through R_{t+1}^{out} , in which case it is represented in the MSSP structure with shortcut edges (solid, angular edges). (b) The chord $C = \overrightarrow{x_j y_j y_{j-1} x_{j-1}}$ is in \mathcal{C}_3 , and f lies in $R_t^{\text{out}} \cap R_{t+1}$. This is handled similarly to (a). (c) Here f lies in R_{t+1}^{out} , \hat{x}_j, \hat{x}_{j-1} are the last ∂R_{t+1} vertices on the $s_j - y_j$ and $s_{j-1} - y_{j-1}$ paths. If the shortest $x'_j - \hat{x}_j$ and $x'_j - v$ paths branch, we can answer the query as in (b). If $x'_j - \hat{x}_j$ is a prefix of $x'_j - v$, $e_v = (\hat{x}_j, \hat{v})$, and $\hat{v} \in \partial R_{t+1}$, then we can use the clockwise order of $\hat{x}_j, \hat{v}, \hat{x}_{j-1}$ around the hole on ∂R_{t+1} to determine whether v lies to the right of C . (Not depicted: the case when $\hat{v} \notin \partial R_{t+1}$.)

shortest path forest w.r.t. \hat{d} has two trees, rooted at x_j and x_{j-1} . Using part (A) of the data structure we compute $\hat{d}(v)$, which reveals the $j^* \in \{j, j-1\}$ such that v is in x_{j^*} 's tree. At this point we break into two cases, depending on whether f is in $R_t^{\text{out}} \cap R_{t+1}$, or in R_{t+1}^{out} . We assume $j^* = j$ without loss of generality and depict only this case in Figure 9(b,c).

Case 2a. Suppose that f is in $R_t^{\text{out}} \cap R_{t+1}$. Let y'_j be a pendant vertex attached to y_j embedded inside f and let x'_j be a pendant attached to x_j embedded in the face on ∂R_t . The shortest $x'_j - y'_j$ and $x'_j - v$ paths diverge at some point. We query the MSSP structure (part (A)) to get the least common ancestor w of y'_j and v and the three edges $e_{y'_j}, e_v, e_w$ around w , then determine the left/right relationship as in Case 1. (If $j^* = j-1$ then we would reverse the answer due to the reversed orientation of the $x_{j-1} - y_{j-1}$ subpath w.r.t. C .) Once again, some of $e_{y'_j}, e_v, e_w$ may be shortcut edges between ∂R_{t+1} -vertices or artificial pendant edges. See Figure 9(b)

Case 2b. Now suppose f lies in R_{t+1}^{out} . We get from part (E) the last vertices $\hat{x}_j, \hat{x}_{j-1} \in \partial R_{t+1}$ that lie on the $s_j - y_j$ and $s_{j-1} - y_{j-1}$ shortest paths. We ask the MSSP structure of part (A) for the least common ancestor w of \hat{x}_j and v in the shortcutted SSSP tree rooted at x'_j , and also get the three incident edges $e_{\hat{x}_j}, e_v, e_w$. The edges e_v and e_w exist and are different, but $e_{\hat{x}_j}$ may not exist if $w = \hat{x}_j$, i.e., if v is a descendant of \hat{x}_j . If all three edges $\{e_{\hat{x}_j}, e_v, e_w\}$ exist we can determine whether v lies to the right of C as in Case 1 or 2a.

Case 2b(i). Suppose $w = \hat{x}_j$ and $e_{\hat{x}_j}$ does not exist. Let $e_v = (\hat{x}_j, \hat{v})$. If $\hat{v} \in \partial R_{t+1}$ then e_v represents

a path that is completely contained in R_{t+1}^{out} . Thus, if we walk clockwise around the hole of R_{t+1}^{out} on ∂R_{t+1} and encounter $\hat{x}_j, \hat{v}, \hat{x}_{j-1}$ in that order then v lies to the right of C , and if we encounter them in the reverse order then v lies to the left of C . See Figure 9(c).

Case 2b(ii). Finally, suppose $\hat{v} \notin \partial R_{t+1}$ and $e_v = (\hat{x}_j, \hat{v})$ is a normal edge in G . Redefine $e_{\hat{x}_j}$ to be the first edge on the path from \hat{x}_j to y_j .¹⁰ Now we can determine if v is to the right of C by looking at the clockwise order of $e_w, e_v, e_{\hat{x}_j}$ around \hat{x}_j . \square

As pointed out in Remark 4.1, Lemma 2.1 does not immediately imply that Line 6 of **SitePathIndicator** and Line 1 of **PieceSearch** can be implemented efficiently. Gawrychowski et al.'s [19] implementation of **PointLocate** requires MSSP access to R_t^{out} , whereas part (A) only lets us query vertices in $R_t^{\text{out}} \cap R_{t+1}$. Gawrychowski et al.'s algorithm is identical to **CentroidSearch**, except that **Navigation** is done directly with MSSP structures. Suppose we are currently at f^* in the centroid decomposition, with y_j, s_j defined as usual. Gawrychowski's algorithm finds j minimizing $\omega(s_j) + \text{dist}_{R_t^{\text{out}}}(s_j, v)$ using three distance queries to the MSSP structure, then decides whether the $s'_j - v$ shortest path is a prefix of the $s'_j - y'_j$ shortest path, and if not, which direction it branches in.¹¹ If f is in $R_t^{\text{out}} \cap R_{t+1}$ we can proceed exactly as in Gawrychowski et al. [19]. If

¹⁰ $\overline{\omega}$ could store $e_{\hat{x}_j}$ in part (E) of the data structure but that is not necessary. If e_0, e_1 are the edges adjacent to \hat{x}_j on the boundary cycle of ∂R_{t+1} , then we can use any member of $\{e_0, e_1\} \setminus \{e_w\}$ as a proxy for $e_{\hat{x}_j}$.

¹¹ s'_j, y'_j being pendant vertices attached to s_j, y_j , as in Lemma 4.3.

not, we retrieve from part (E) the last vertex \hat{x} of ∂R_{t+1} on the s_j - y_j shortest path, use \hat{x} in lieu of y'_j for the LCA queries, and tell whether the s'_j - v path branches to the right exactly as in Lemma 4.3, Case 2b.

5 Analysis

This section constitutes a proof of the claims of Theorem 1.1 concerning space complexity and query time; refer to Appendix C in the full version for an efficient construction algorithm.

Combining Lemmas 2.1 and 2.2 (see Section 4.4.3), **PointLocate** runs in $O(\kappa \log n \log \log n)$ time. Together with Lemma 4.3 it follows that **PieceSearch** also takes $O(\kappa \log n \log \log n)$ time. **SitePathIndicator** uses **PointLocate**, the MSSP structure, and $O(1)$ -time tree operations on $T_q^{R_i}$ and the \vec{r} -division like least common ancestors and level ancestors [21, 4, 5, 20]. Thus **SitePathIndicator** also takes $O(\kappa \log n \log \log n)$ time. The calls to **MaximalChord** and **AdjacentPiece** in **ChordIndicator** take $O(\log n \log \log n)$ time by Lemma 4.2, and testing which side of a chord v lies on takes $O(\kappa \log \log n)$ time by Lemma 4.3. The bottleneck in **ChordIndicator** is still **PieceSearch**, which takes $O(\kappa \log n \log \log n)$ time. The only non-trivial parts of **Navigation** are calls to **SitePathIndicator** and **ChordIndicator**, so it, too, takes $O(\kappa \log n \log \log n)$ time.

An initial call to **CentroidSearch** (Line 5 of **Dist**) generates at most $\log n$ recursive calls to **CentroidSearch**, culminating in the last recursive call making 1 or 2 calls to **Dist** with the “ i ” parameter incremented. Excluding the cost of recursive calls to **Dist**, the cost of **CentroidSearch** is dominated by calls to **Navigation**, i.e., an initial call to **CentroidSearch** costs $\log n \cdot O(\kappa \log n \log \log n) = O(\kappa \log^2 n \log \log n)$ time. Let $T(i)$ be the cost of a call to **Dist**(u_i, v, R_i). In the worst case that $t = m - 1$, we have

$$\begin{aligned} T(m-1) &= O(\kappa \log \log n) \\ T(i) &= 2T(i+1) + O(\kappa \log^2 n \log \log n) \end{aligned}$$

It follows that the time to answer a distance query is $T(0) = O(2^m \cdot \kappa \log^2 n \log \log n)$.

The space complexity of each part of the data structure is as follows. (A) is $O(\kappa mn^{1+1/m+1/\kappa})$ by Lemma 2.1 and the fact that $r_{i+1}/r_i = n^{1/m}$. (B) is $O(mn^{1+1/(2m)})$ since $\sqrt{r_{i+1}/r_i} = n^{1/(2m)}$. (C) is $O(mn)$ since $\sum_i n/r_i \cdot (\sqrt{r_i})^2 = O(mn)$. (D) is $O(mn \log n)$ by Lemma 4.2, and (E) is $O(m)$ times the space cost of (B) and (C), namely $O(m^2 n^{1+1/(2m)})$. The bottleneck is (A).

We now explain how m, κ can be selected to achieve the extreme space and query complexities claimed The-

orem 1.1. To optimize for query time, pick $\kappa = m$ to be any function of n that is $\omega(1)$ and $o(\log \log n)$. Then the query time is

$$O(2^m \kappa \log^2 n \log \log n) = \log^{2+o(1)} n$$

and the space is

$$O(m\kappa n^{1+1/m+1/\kappa}) = n^{1+o(1)}.$$

To optimize for space, choose $\kappa = \log n$ and m to be a function that is $\omega(\log n / \log \log n)$ and $o(\log n)$. Then the space is

$$\begin{aligned} O(m\kappa n^{1+1/m+1/\kappa}) &= o(n^{1+1/m} \log^2 n) \\ &= n \cdot 2^{o(\log \log n)} \cdot \log^2 n \\ &= n \log^{2+o(1)} n, \end{aligned}$$

and the query time

$$\begin{aligned} O(2^m \kappa \log^2 n \log \log n) &= 2^{o(\log n)} \log^3 n \log \log n \\ &= n^{o(1)}. \end{aligned}$$

References

- [1] Ittai Abraham and Cyril Gavoille. On approximate distance labels and routing schemes with affine stretch. In *Proceedings of the 25th International Symposium on Distributed Computing (DISC)*, volume 6950 of *Lecture Notes in Computer Science*, pages 404–415, 2011.
- [2] Rachit Agarwal. The space-stretch-time tradeoff in distance oracles. In *Proceedings of the 22nd European Symposium on Algorithms (ESA)*, volume 8737 of *Lecture Notes in Computer Science*, pages 49–60, 2014.
- [3] Srinivasa Rao Arikati, Danny Z. Chen, L. Paul Chew, Gautam Das, Michiel H. M. Smid, and Christos D. Zaroliagis. Planar spanners and approximate shortest path queries among obstacles in the plane. In *Proceedings 4th Annual European Symposium on Algorithms (ESA)*, volume 1136 of *Lecture Notes in Computer Science*, pages 514–528, 1996.
- [4] Michael A. Bender and Martin Farach-Colton. The LCA problem revisited. In *Proceedings of the 4th Latin American Symposium on Theoretical Informatics (LATIN)*, volume 1776 of *Lecture Notes in Computer Science*, pages 88–94. Springer, 2000.
- [5] Michael A. Bender and Martin Farach-Colton. The level ancestor problem simplified. *Theor. Comput. Sci.*, 321(1):5–12, 2004.
- [6] Gerth Stølting Brodal, Pooya Davoodi, and S Srinivasa Rao. Path minima queries in dynamic weighted trees. In *Proceedings 12th Int'l Symposium on Algorithms and Data Structures (WADS)*, volume 6844 of *Lecture Notes in Computer Science*, pages 290–301, 2011.

- [7] Sergio Cabello. Many distances in planar graphs. *Algorithmica*, 62(1-2):361–381, 2012.
- [8] Sergio Cabello. Subquadratic algorithms for the diameter and the sum of pairwise distances in planar graphs. *ACM Trans. Algorithms*, 15(2):21:1–21:38, 2019.
- [9] Panagiotis Charalampopoulos, Paweł Gawrychowski, Shay Mozes, and Oren Weimann. Almost optimal distance oracles for planar graphs. In *Proceedings of the 51st Annual ACM Symposium on Theory of Computing (STOC)*, pages 138–151, 2019.
- [10] Danny Z. Chen and Jinhui Xu. Shortest path queries in planar graphs. In *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 469–478, 2000.
- [11] Vincent Cohen-Addad, Søren Dahlgaard, and Christian Wulff-Nilsen. Fast and compact exact distance oracle for planar graphs. In *Proceedings 58th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 962–973, 2017.
- [12] Paul F. Dietz. Fully persistent arrays. In *Proceedings of the First Workshop on Algorithms and Data Structures (WADS)*, volume 382 of *Lecture Notes in Computer Science*, pages 67–74, 1989.
- [13] Martin Dietzfelbinger, Anna R. Karlin, Kurt Mehlhorn, Friedhelm Meyer auf der Heide, Hans Rohnert, and Robert Endre Tarjan. Dynamic perfect hashing: Upper and lower bounds. In *Proceedings of the 29th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 524–531, 1988.
- [14] Hristo Djidjev. On-line algorithms for shortest path problems on planar digraphs. In *Proceedings of the 22nd International Workshop on Graph-Theoretic Concepts in Computer Science (WG)*, volume 1197 of *Lecture Notes in Computer Science*, pages 151–165, 1996.
- [15] James R. Driscoll, Neil Sarnak, Daniel Dominic Sleator, and Robert Endre Tarjan. Making data structures persistent. *J. Comput. Syst. Sci.*, 38(1):86–124, 1989.
- [16] Jeff Erickson, Kyle Fox, and Luvsandondov Lkhamsuren. Holiest minimum-cost paths and flows in surface graphs. In *Proceedings of the 50th Annual ACM Symposium on Theory of Computing (STOC)*, pages 1319–1332, 2018.
- [17] Jittat Fakcharoenphol and Satish Rao. Planar graphs, negative weight edges, shortest paths, and near linear time. *J. Comput. Syst. Sci.*, 72(5):868–889, 2006.
- [18] Greg N. Frederickson. Fast algorithms for shortest paths in planar graphs, with applications. *SIAM J. Comput.*, 16(6):1004–1022, 1987.
- [19] Paweł Gawrychowski, Shay Mozes, Oren Weimann, and Christian Wulff-Nilsen. Better tradeoffs for exact distance oracles in planar graphs. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 515–529, 2018.
- [20] Torben Hagerup. Still simpler static level ancestors. *CoRR*, abs/2005.11188, 2020.
- [21] Dov Harel and Robert Endre Tarjan. Fast algorithms for finding nearest common ancestors. *SIAM J. Comput.*, 13(2):338–355, 1984.
- [22] Monika Rauch Henzinger and Valerie King. Randomized fully dynamic graph algorithms with polylogarithmic time per operation. *J. ACM*, 46(4):502–516, 1999.
- [23] Donald B. Johnson. Efficient algorithms for shortest paths in sparse networks. *J. ACM*, 24(1):1–13, 1977.
- [24] Ken-ichi Kawarabayashi, Philip N. Klein, and Christian Sommer. Linear-space approximate distance oracles for planar, bounded-genus and minor-free graphs. In *Proceedings of the 38th Int'l Colloquium on Automata, Languages and Programming (ICALP)*, volume 6755 of *Lecture Notes in Computer Science*, pages 135–146, 2011.
- [25] Ken-ichi Kawarabayashi, Christian Sommer, and Mikkel Thorup. More compact oracles for approximate distances in undirected planar graphs. In *Proceedings of the 24th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 550–563, 2013.
- [26] Philip Klein. Preprocessing an undirected planar network to enable fast approximate distance queries. In *Proceedings of the 13th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 820–827, 2002.
- [27] Philip N. Klein. Multiple-source shortest paths in planar graphs. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 146–155, 2005.
- [28] Philip N. Klein, Shay Mozes, and Christian Sommer. Structured recursive separator decompositions for planar graphs in linear time. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC)*, pages 505–514, 2013.
- [29] Richard J. Lipton and Robert Endre Tarjan. Applications of a planar separator theorem. *SIAM J. Comput.*, 9(3):615–627, 1980.
- [30] Yaowei Long and Seth Pettie. Planar distance oracles with better time-space tradeoffs. *CoRR*, abs/2007.08585, 2020.
- [31] Gary L. Miller. Finding small simple cycle separators for 2-connected planar graphs. *J. Comput. Syst. Sci.*, 32(3):265–279, 1986.
- [32] Shay Mozes and Christian Sommer. Exact distance oracles for planar graphs. In *Proceedings of the 23rd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 209–222, 2012.
- [33] Yahav Nussbaum. Improved distance queries in planar graphs. In *Proceedings 12th Int'l Workshop on Algorithms and Data Structures (WADS)*, pages 642–653, 2011.
- [34] Mihai Patrascu and Liam Roditty. Distance oracles beyond the Thorup-Zwick bound. *SIAM J. Comput.*, 43(1):300–311, 2014.
- [35] Neil Sarnak and Robert Endre Tarjan. Planar point location using persistent search trees. *Commun. ACM*, 29(7):669–679, 1986.
- [36] Daniel Dominic Sleator and Robert Endre Tarjan. A data structure for dynamic trees. *J. Comput. Syst. Sci.*, 26(3):362–391, 1983.

- [37] Christian Sommer. Shortest-path queries in static networks. *ACM Computing Surveys*, 46(4):1–31, 2014.
- [38] Christian Sommer, Elad Verbin, and Wei Yu. Distance oracles for sparse graphs. In *Proceedings of the 50th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 703–712, 2009.
- [39] Mikkel Thorup. Compact oracles for reachability and approximate distances in planar digraphs. *J. ACM*, 51(6):993–1024, 2004.
- [40] Mikkel Thorup and Uri Zwick. Approximate distance oracles. *J. ACM*, 52(1):1–24, 2005.
- [41] Christian Wulff-Nilsen. *Algorithms for planar graphs and graphs in metric spaces*. PhD thesis, PhD thesis, University of Copenhagen, 2010.