

An Improved Algorithm for Shortest Paths in Weighted Unit-Disk Graphs*

Bruce W. Brewer 

Kahlert School of Computing, University of Utah, Salt Lake City, UT 84112, USA

Haitao Wang 

Kahlert School of Computing, University of Utah, Salt Lake City, UT 84112, USA

Abstract

Let V be a set of n points in the plane. The unit-disk graph $G = (V, E)$ has vertex set V and an edge $e_{uv} \in E$ between vertices $u, v \in V$ if the Euclidean distance between u and v is at most 1. The weight of each edge e_{uv} is the Euclidean distance between u and v . Given V and a source point $s \in V$, we consider the problem of computing shortest paths in G from s to all other vertices. The previously best algorithm for this problem runs in $O(n \log^2 n)$ time [Wang and Xue, SoCG'19]. The problem has an $\Omega(n \log n)$ lower bound under the algebraic decision tree model. In this paper, we present an improved algorithm of $O(n \log^2 n / \log \log n)$ time (under the standard real-RAM model). Furthermore, we show that the problem can be solved using $O(n \log n)$ comparisons under the algebraic decision tree model, matching the $\Omega(n \log n)$ lower bound.

Keywords and phrases Unit-disk graphs, shortest paths, dynamic nearest neighbor search, additively-weighted Voronoi diagrams

Digital Object Identifier 10.57717/cgt.v5i2.69

Funding This research was supported in part by NSF under Grant CCF-2300356.

1 Introduction

Let V be a set of n points in the plane. The unit-disk graph $G = (V, E)$ has vertex set V and an edge $e_{uv} \in E$ between vertices $u, v \in V$ if the Euclidean distance between u and v is at most 1. Alternatively, G can be seen as the intersection graph of disks with radius $\frac{1}{2}$ centered at the points in V (i.e., two disks have an edge in the graph if they intersect). In the *weighted graph*, the weight of each edge $e_{uv} \in E$ is the Euclidean distance between u and v . In the *unweighted graph*, all edges have the same weight.

Given V and a source point $s \in V$, we study the single source shortest path (SSSP) problem where the goal is to compute shortest paths from s to all other vertices in G . Like in general graphs, the algorithm usually returns a shortest path tree rooted at s . The problem in the unweighted graph has an $\Omega(n \log n)$ lower bound in the algebraic decision tree model since even deciding if G is connected requires that much time by a reduction from the max-gap [3]. The unweighted problem has been solved optimally in $O(n \log n)$ time by Cabello and Jejčić [3], or in $O(n)$ time by Chan and Skrepetos [5] if the points of V are pre-sorted (by both the x - and y -coordinates). Several algorithms for the weighted case are also known [3, 9, 12, 14, 16]. Roditty and Segal [14] first solved the problem in $(n^{4/3+\delta})$ time, where $\delta > 0$ is an arbitrarily small constant. Cabello and Jejčić [3] improved it to $O(n^{1+\delta})$ time. Subsequent improvements were made by Kaplan, Mulzer, Roditty, Seiferth, and Sharir [9] and also by Liu [12] by developing more efficient dynamic bichromatic closest pair data structures and plugging them into the algorithm of [3]. Both algorithms are randomized, and specifically, the algorithm of [9] takes roughly $O(n \log^9 n)$ expected time

* A preliminary version of this paper appeared in *Proceedings of the 36th Canadian Conference on Computational Geometry (CCCG 2024)* [2].



and the algorithm of [12] runs in $O(n \log^4 n)$ expected time. Wang and Xue [16] proposed a new method that solves the problem in $O(n \log^2 n)$ deterministic time without using dynamic bichromatic closest pair data structures. It is currently the best algorithm for the problem.

1.1 Our result

We present a new algorithm of $O(n \log^2 n / \log \log n)$ time for the weighted case and, therefore, slightly improve the result of [16]. Our algorithm follows the framework of Wang and Xue [16] but provides a more efficient solution to a bottleneck subproblem in their algorithm, called the *offline insertion-only additively-weighted nearest neighbor problem with a separating line* (or IOAWNN-SL for short). Specifically, we are given a sequence of n operations of the following two types: (1) Insertion: Insert a weighted point to P (which is \emptyset initially); (2) Query: given a query point q , find the *additively-weighted nearest neighbor* of q in P , where the distance between q to any point $p \in P$ is defined to be their Euclidean distance plus the weight of p . The points of P and all the query points are required to be separated by a given line (say the x -axis). The goal of the problem is to answer all queries.

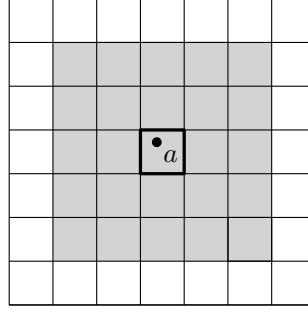
Wang and Xue [16] solved the IOAWNN-SL problem in $O(n \log^2 n)$ time using the traditional logarithmic method of Bentley [1]. This is the bottleneck of their overall shortest path algorithm; all other parts of the algorithm take $O(n \log n)$ time. We derive a more efficient algorithm that solves IOAWNN-SL in $O(n \log^2 n / \log \log n)$ time (see Theorem 1 for details). Plugging this result into the algorithm framework of Wang and Xue [16] solves the shortest path problem in $O(n \log^2 n / \log \log n)$ time.

► **Theorem 1.** *Let P be an initially empty set of n weighted points in the plane such that all points of P lie below the x -axis ℓ . There exists a data structure $\mathcal{D}(P)$ of $O(n)$ space supporting the following operations:*

1. *Insertion: Insert a weighted point p below ℓ to P in amortized $O(\log^2 n / \log \log n)$ time.*
2. *Query: Given a query point q above ℓ , find the additively-weighted nearest neighbor to q in P in worst-case $O(\log^2 n / \log \log n)$ time.*

Our algorithm for Theorem 1 needs to solve a subproblem about merging two additively-weighted Voronoi diagrams. Specifically, let S_a and S_b each be a subset of n weighted points in the plane such that all points of $S_a \cup S_b$ are below the x -axis ℓ . Let $\mathcal{VD}(S_a)$ denote the additively-weighted Voronoi diagram of S_a , and $\mathcal{VD}_+(S_a)$ denote the portion of $\mathcal{VD}(S_a)$ above ℓ . Similarly, define $\mathcal{VD}(S_b)$ and $\mathcal{VD}_+(S_b)$ for S_b , and define $\mathcal{VD}(S_a \cup S_b)$ and $\mathcal{VD}_+(S_a \cup S_b)$ for $S_a \cup S_b$. Given $\mathcal{VD}_+(S_a)$ and $\mathcal{VD}_+(S_b)$, our problem is to compute $\mathcal{VD}_+(S_a \cup S_b)$. We solve the problem in $O(n)$ time by modifying Kirkpatrick's algorithm for merging two standard Voronoi diagrams [10] and by making use of the property that $\mathcal{VD}_+(S_a \cup S_b)$ and all points of $S_a \cup S_b$ are separated by ℓ . Note that directly applying Kirkpatrick's algorithm does not work (see Section 3 for more details). It would be more interesting to have a linear time algorithm to compute the complete diagram $\mathcal{VD}(S_a \cup S_b)$ by merging $\mathcal{VD}(S_a)$ and $\mathcal{VD}(S_b)$. Our technique, however, does not immediately work because it relies on the separating line ℓ . Nevertheless, we hope our result will serve as a stepping stone towards achieving that goal. We summarize our result in the following theorem.

► **Theorem 2.** *Let S_a and S_b each be a set of n weighted points in the plane such that all the points of $S_a \cup S_b$ are below the x -axis ℓ . Given $\mathcal{VD}_+(S_a)$ and $\mathcal{VD}_+(S_b)$, $\mathcal{VD}_+(S_a \cup S_b)$ can be constructed in $O(n)$ time.*



■ **Figure 1** Illustrating \square_a (the central highlighted square) and \boxplus_a (the gray area).

Algebraic decision tree model. The above result holds for the standard real-RAM model. Under the algebraic decision tree model in which we only count comparisons toward the time complexity, using a technique recently developed by Chan and Zheng [6], we show that the problem IOAWNN-SL can be solved using $O(n \log n)$ comparisons. This leads to an $O(n \log n)$ time algorithm for the shortest path problem in weighted unit-disk graphs under the algebraic decision tree model, matching the $\Omega(n \log n)$ lower bound [3].

Outline. The rest of the paper is organized as follows. We describe the shortest path algorithm framework in Section 2, mainly by reviewing Wang and Xue’s algorithm [16]. In Section 3, we introduce our data structure for IOAWNN-SL and thus prove Theorem 1. Section 4 presents our Voronoi diagram merging algorithm for Theorem 2. The algebraic decision tree algorithm is described in Section 5. Section 6 concludes.

2 The shortest path algorithm

In this section, we describe the shortest path algorithm. We begin with reviewing Wang and Xue’s algorithm [16] and explain why the IOAWNN-SL problem is a bottleneck (we only state their algorithm and refer the interested reader to their paper [16] for the correctness analysis). We will show how our solution to IOAWNN-SL in Theorem 1 can lead to an $O(n \log^2 n / \log \log n)$ time algorithm for the shortest path problem.

Given a set V of n points in \mathbb{R}^2 and a source point $s \in V$, we wish to compute shortest paths from s to all vertices in the weighted unit-disk graph $G = (V, E)$. We use $e_{uv} \in E$ to denote the edge between two points $u, v \in V$ and $w(e_{uv})$ to denote the weight of the edge. Recall that $w(e_{uv}) = \|u - v\| \leq 1$, where $\|u - v\|$ denotes the Euclidean distance between u and v . The algorithm will compute a table $\text{dist}[\cdot]$ such that after the algorithm finishes, $\text{dist}[v]$ is the length of a shortest path from s to v for all $v \in V$. Using a predecessor table, we could also maintain a shortest path tree, but we will omit the discussion about it.

We overlay the plane with a grid Γ of square cells with side lengths $1/2$. For any point $a \in \mathbb{R}^2$, denote by \square_a the cell of Γ such that $a \in \square_a$, and \boxplus_a the 5×5 patch of cells in Γ centered around \square_a (see Figure 1). For a set of points $A \subseteq \mathbb{R}^2$ and $a \in A$, we use $A_{\square_a} = A \cap \square_a$ and $A_{\boxplus_a} = A \cap \boxplus_a$. The algorithm makes use of the following properties: (1) For any two points a, b in the same cell of Γ , $\|a - b\| \leq 1$ holds; (2) if $\|a - b\| \leq 1$, then b is in \boxplus_a and a is in \boxplus_b .

Wang and Xue’s algorithm is summarized in Algorithm 1. It can be understood by contrasting with Dijkstra’s algorithm, which we write in Algorithm 2 using similar notation. In particular, a subroutine $\text{UPDATE}(A, B)$ is used to “push” the current candidate shortest

path information from A to B where $A, B \subseteq V$. Specifically, for each point $b \in B$, we find:

$$p_b = \arg \min_{\{a \in A : e_{ab} \in E\}} \text{dist}[a] + w(e_{ab}). \quad (1)$$

We then update $\text{dist}[b]$ to $\min\{\text{dist}[b], \text{dist}[p_b] + w(e_{p_b b})\}$.

■ **Algorithm 1** Wang and Xue's algorithm [16]

```

1  $\text{dist}[a] \leftarrow \infty$  for all  $a \in V$ 
2  $\text{dist}[s] \leftarrow 0$ 
3  $A \leftarrow V$ 
4 while  $A \neq \emptyset$  do
5    $c \leftarrow \arg \min_{a \in A} \{\text{dist}[a]\}$ 
6   UPDATE( $A_{\square_c}, A_{\square_c}$ ) // First Update
7   UPDATE( $A_{\square_c}, A_{\square_c}$ ) // Second Update
8    $A \leftarrow A \setminus A_{\square_c}$ 
9 return  $\text{dist}[\cdot]$ 

```

■ **Algorithm 2** Dijkstra's algorithm

```

1  $\text{dist}[a] \leftarrow \infty$  for all  $a \in V$ 
2  $\text{dist}[s] \leftarrow 0$ 
3  $A \leftarrow V$ 
4 while  $A \neq \emptyset$  do
5    $c \leftarrow \arg \min_{a \in A} \{\text{dist}[a]\}$ 
6   UPDATE( $\{c\}, A$ )
7    $A \leftarrow A \setminus \{c\}$ 
8 return  $\text{dist}[\cdot]$ 

```

The main difference between Wang and Xue's algorithm and Dijkstra's is that instead of operating on single vertices, Wang and Xue's algorithm operates on cells of Γ . Generally speaking, the first update (Line 6) in Algorithm 1 is to update the shortest path information for the points in A_{\square_c} using the shortest path information of their neighbors. The second update is to use the shortest path information for the points in V_{\square_c} to update the shortest path information of their neighbors. Wang and Xue prove that after the first update, the shortest path information for all points of V_{\square_c} is correctly computed.

Wang and Xue give an $O(n \log^2 n)$ time solution for the second update, i.e., Line 7. The rest of Algorithm 1 takes $O(n \log n)$ time. We will improve the runtime for the second update to $O(n \log^2 n / \log \log n)$ using Theorem 1, which improves the runtime for Algorithm 1 to $O(n \log^2 n / \log \log n)$. The details are discussed in the following.

2.1 The second update

To implement the second update $\text{UPDATE}(A_{\square_c}, A_{\square_c})$, since A_{\square_c} has $O(1)$ cells, it suffices to perform $\text{UPDATE}(A_{\square_c}, A_{\square})$ for each cell $\square \in \square_c$ individually.

If \square is \square_c , then $A_{\square_c} = A_{\square}$. Since the distance between two points in \square_c is at most 1, $\text{UPDATE}(A_{\square_c}, A_{\square})$ can be performed in $O(|A_{\square_c}| \log |A_{\square_c}|)$ time (and $O(|A_{\square_c}|)$ space) by constructing the additively-weighted Voronoi diagram for A_{\square_c} [8].

If \square is not \square_c , a useful property is that \square and \square_c are separated by an axis-parallel line. To perform $\text{UPDATE}(A_{\square_c}, A_{\square})$, Wang and Xue [16] proposed Algorithm 3 below.

Algorithm 3 $\text{UPDATE}(A, B)$ from [16]

```

1 dist'[a]  $\leftarrow$  dist[a] for  $a \in A$ 
2 Sort the points in  $A = \{a_1, \dots, a_{|A|}\}$  so that  $\text{dist}'[a_1] \leq \dots \leq \text{dist}'[a_{|A|}]$ 
3 for  $i = 1, \dots, |A|$  do
4    $B_i \leftarrow \{b \in B : e_{a_i b} \in E \text{ and } e_{a_j b} \notin E \text{ for all } j < i\}$ 
5    $U \leftarrow \emptyset$ 
6   for  $i = |A|, \dots, 1$  do
7      $U \leftarrow U \cup \{a_i\}$ 
8     for  $b \in B_i$  do
9        $p = \arg \min_{u \in U} \{\text{dist}'[u] + w(e_{ub})\}$ 
10       $\text{dist}[b] \leftarrow \min\{\text{dist}[b], \text{dist}'[p] + w(e_{pb})\}$ 

```

The correctness of Algorithm 3 hinges on the fact that p found by Line 9 is the same as p_b in Equation (1). This is seen by arguing that $p_b \in U$ and $e_{pb} \in E$.

We now analyze the runtime of Algorithm 3. Sorting A takes $O(|A| \log |A|)$ time. Computing the subsets B_i , $1 \leq i \leq |A|$, can be done in $O((|A| + |B|) \log(|A| + |B|))$ time (and $O(|A| + |B|)$ space) [16]. The for loop (Lines 6–10) is an instance of the IOAWNN-SL problem introduced in Section 1. Indeed, if we assign each point u in U a weight equal to $\text{dist}'[u]$, then p in Line 9 is essentially the additively-weighted nearest neighbor of b in U . The set U is dynamically changed with point insertions in Line 7. As such, by Theorem 1, the for loop can be implemented in $O(k \log^2 k / \log \log k)$ time (and $O(k)$ space) with $k = |A| + |B|$. Therefore, $\text{UPDATE}(A_{\square_c}, A_{\square})$ can be performed in $O(k \log^2 k / \log \log k)$ time and $O(k)$ space, with $k = |A_{\square_c}| + |A_{\square}|$.

In summary, since A_{\square_c} has $O(1)$ cells, the second update $\text{UPDATE}(A_{\square_c}, A_{\square_c})$ can be implemented in $O(|A_{\square_c}| \log^2 |A_{\square_c}| / \log \log |A_{\square_c}|)$ time and $O(|A_{\square_c}|)$ space as $A_{\square_c} \subseteq A_{\square_c}$. As analyzed in [16], the total sum of $|A_{\square_c}|$ in the entire Algorithm 1 is $O(n)$. This leads to the following result.

► **Theorem 3.** *Given a set V of n points in the plane and a source point s , shortest paths from s to all other vertices in the weighted unit-disk graph $G = (V, E)$ can be computed in $O(n \log^2 n / \log \log n)$ time and $O(n)$ space.*

3 The offline insertion-only additively-weighted nearest neighbor problem with a separating line (IOAWNN-SL)

In this section, we prove Theorem 1. We follow the notation in Section 1. In particular, for any subset $P' \subseteq P$, $\mathcal{VD}_+(P')$ denotes the portion of the additively-weighted Voronoi diagram of P' above the x -axis ℓ .

Our data structure $\mathcal{D}(P)$ for Theorem 1 consists of two components: $\mathcal{D}(P')$ and $\mathcal{VD}_+(P \setminus P')$ for some subset $P' \subseteq P$; we maintain the invariant $|P'| \leq |P| / \log |P|$. We also build a point location data structure on $\mathcal{VD}_+(P \setminus P')$ so that, given a query point, the cell of $\mathcal{VD}_+(P \setminus P')$ containing the point can be found in $O(\log |P \setminus P'|)$ time [7, 11]. As such, $\mathcal{D}(P)$ is a recursive structure: $\mathcal{D}(P)$ is defined in terms of $\mathcal{D}(P')$ which in turn is defined in terms of $\mathcal{D}(P'')$ and so on. As the base case, if $|P| \leq c$ for some constant c , then we simply let

$\mathcal{D}(P) = \mathcal{VD}_+(P)$. Similar recursive data structures have been used before in the literature, e.g., [4, 13].

In the following, we discuss how to handle the two operations: insertions and queries.

Queries. Given a query point q above ℓ , the query is to find the (weighted) nearest neighbor of q in P , denoted by p_q . Note that searching for p_q is a decomposable searching problem: If p_q^1 is the nearest neighbor of q in P' and p_q^2 is the nearest neighbor of q in $P \setminus P'$, then p_q is whichever of p_q^1 and p_q^2 is nearer to q . We can find p_q^1 by recursively querying $\mathcal{D}(P')$. We can find p_q^2 using a point location in $\mathcal{VD}_+(P \setminus P')$.

For the time analysis, let $Q(n)$ denote the runtime of this query algorithm. By the invariant, $|P'| \leq |P| / \log |P|$ holds, and therefore computing p_q^1 takes $Q(n / \log n)$ time. Finding p_q^2 with point location can be done $O(\log |P \setminus P'|) = O(\log n)$ time. Therefore, we have the following recurrence: $Q(n) = Q(n / \log n) + O(\log n)$. This solves to $Q(n) = O(\log^2 n / \log \log n)$. As such, each query operation takes worst-case $O(\log^2 n / \log \log n)$ time.

Insertions. For insertion, we are to insert a point p below ℓ to P . Recall that our data structure $\mathcal{D}(P)$ consists of two parts: $\mathcal{D}(P')$ and $\mathcal{VD}_+(P \setminus P')$. We start by inserting p into P' . This involves a recursive insertion call to $\mathcal{D}(P')$. Second, we check if the invariant $|P'| \leq |P| / \log |P|$ still holds. If it holds, then we are done. Otherwise, we set $P' = \emptyset$. This means that we must compute $\mathcal{VD}_+(P)$ and a point location data structure for it. We do this as follows. Before setting $P' = \emptyset$, P' contains p and we have $\mathcal{VD}_+(P \setminus P')$ available. First, we construct $\mathcal{VD}_+(P')$ recursively. Then, we compute $\mathcal{VD}_+(P)$ by merging $\mathcal{VD}_+(P')$ and $\mathcal{VD}_+(P \setminus P')$. Finally, we compute a point location data structure for $\mathcal{VD}_+(P)$.

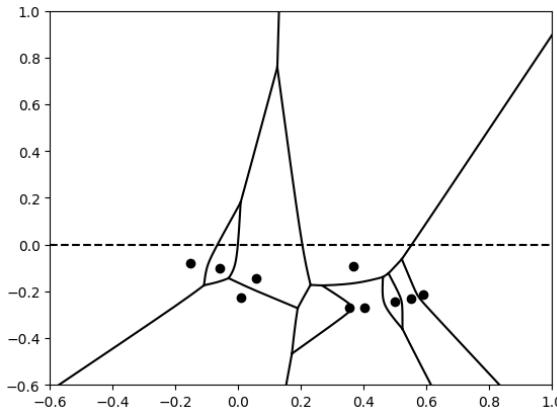
For the time analysis, let $T(n)$ denote the time to construct $\mathcal{VD}_+(P)$ and a point location data structure for it. By the invariant, constructing $\mathcal{VD}_+(P')$ takes $T(n / \log n)$ time. By Theorem 2, merging $\mathcal{VD}_+(P')$ and $\mathcal{VD}_+(P \setminus P')$ takes $O(n)$ time. Computing a point location data structure for $\mathcal{VD}_+(P)$ can be done in $O(n)$ time [7, 11]. Therefore, we have the following recurrence: $T(n) = T(n / \log n) + O(n)$. This solves to $T(n) = O(n)$.

Let $I(n)$ be amortized time it takes to perform an insertion. By the invariant, the recursive call to $\mathcal{D}(P')$ for inserting p into P' takes $I(n / \log n)$ time. Because we set $P' = \emptyset$ once $\mathcal{VD}_+(P)$ is constructed, we only need to construct $\mathcal{VD}_+(P)$ every $\Theta(n / \log n)$ insertions. Because constructing $\mathcal{VD}_+(P)$ takes $O(n)$ time, the amortized time for constructing $\mathcal{VD}_+(P)$ per insertion is $O(\log n)$. Therefore, we have $I(n) = I(n / \log n) + O(\log n)$. This solves to $I(n) = O(\log^2 n / \log \log n)$. We conclude that each insertion takes $O(\log^2 n / \log \log n)$ amortized time.

Space. Let $S(n)$ be the space of $\mathcal{D}(P)$. By the invariant, $\mathcal{D}(P')$ takes $S(n / \log n)$ space. The size of $\mathcal{VD}_+(P \setminus P')$ is $O(|P \setminus P'|) = O(n)$ [8] and the space of the point location data structure for it is the same [7, 11]. Therefore, the recurrence for $S(n)$ is $S(n) = S(n / \log n) + O(n)$. This solves to $S(n) = O(n)$. This concludes the proof of Theorem 1.

4 Merging two additively-weighted Voronoi diagrams

In this section, we prove Theorem 2. For completeness, we first introduce the formal definition of additively-weighted Voronoi diagrams and then present our merging algorithm.



■ **Figure 2** Illustrating an additively-weighted Voronoi diagram. The dashed horizontal line is the x -axis ℓ .

4.1 Additively-weighted Voronoi diagrams

Let $S = \{s_1, s_2, \dots, s_n\}$ be a set of n points in the plane such that each point s_i has a weight w_i that can be positive, zero, or negative. Following the literature, we refer to points of S as *sites*. We define the additively-weighted Euclidean distance (or *weighted distance* for short) of a point $p \in \mathbb{R}^2$ to a site s_i as $d(s_i, p) = \|s_i - p\| + w_i$.

The additively-weighted Voronoi diagram of S , denoted by $\mathcal{VD}(S)$, partitions the plane into Voronoi regions, Voronoi edges, and Voronoi vertices; see Figure 2. Each Voronoi region R_i is associated with a site s_i and is defined to be the set of points that are closer to s_i than to any other site measured by the weighted distances:

$$R_i = \{p \in \mathbb{R}^2 : d(s_i, p) < d(s_j, p), \forall j \neq i\}.$$

Each Voronoi edge E_{ij} is associated with two distinct sites s_i and s_j and is defined to be the set of points that are equidistant to s_i and s_j and closer to these sites than any other sites:

$$E_{ij} = \{p \in \mathbb{R}^2 : d(s_i, p) = d(s_j, p) < d(s_k, p), \forall k \neq i, j\}.$$

Each Voronoi vertex is associated with three or more distinct sites and is defined to be the point that is equidistant to these sites and closer to these sites than any other site.

We will also talk about the *bisector* between two sites, which is defined to be the set of points in the plane that are equidistant to the two sites:

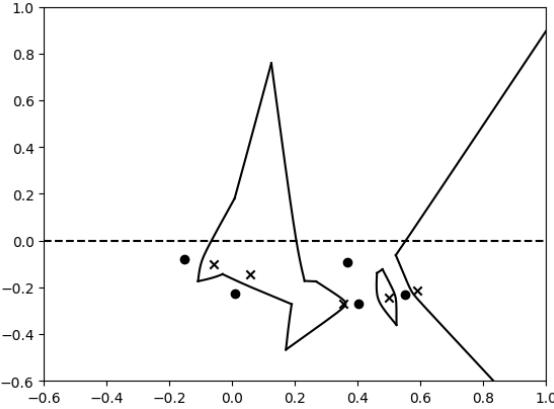
$$B(s_i, s_j) = \{p \in \mathbb{R}^2 : d(s_i, p) = d(s_j, p)\}.$$

$B(s_i, s_j)$ is a hyperbolic arc whose foci are s_i and s_j . Note that a Voronoi edge associated with two sites is a subset of their bisector.

Observation 4 states some properties about $\mathcal{VD}(S)$ that are well known in the literature; we will use these properties in our algorithm.

► **Observation 4.** ([8])

1. Every Voronoi region of $\mathcal{VD}(S)$ must contain its associated site.
2. Each Voronoi region R_i of $\mathcal{VD}(S)$ is star-shaped with respect to its site s_i , that is, the line segment $\overline{s_i p}$ is inside R_i for any point $p \in R_i$.
3. The combinatorial size of $\mathcal{VD}(S)$ is $O(|S|)$.



■ **Figure 3** Illustrating the contour between two sets of points. The dashed horizontal line is the x -axis ℓ .

4.2 Merging algorithm for Theorem 2

We follow the notation introduced in Section 1, e.g., ℓ , n , S_a , S_b , $\mathcal{VD}(S_a)$, $\mathcal{VD}(S_b)$, $\mathcal{VD}_+(S_a)$, $\mathcal{VD}_+(S_b)$, etc. Let $S = S_a \cup S_b$. Given $\mathcal{VD}_+(S_a)$ and $\mathcal{VD}_+(S_b)$, our goal is to compute $\mathcal{VD}_+(S)$ in $O(n)$ time. For ease of exposition, we make a general position assumption that no point in the plane is equidistant to four points of S .

Our strategy is to identify the *contour* which consists of edges in the complete Voronoi diagram $\mathcal{VD}(S)$ that are associated with a site in S_a and a site in S_b ; see Figure 3. Note that the contour may have multiple connected components. The contour partitions the plane into regions CR_i such that $\mathcal{VD}(S) \cap CR_i$ is either $\mathcal{VD}(S_a) \cap CR_i$ or $\mathcal{VD}(S_b) \cap CR_i$ (we show in Lemma 6 later that each contour component has the topology of a line or a circle). As such, once we have identified the contour, computing $\mathcal{VD}(S)$ is straightforward. To compute the contour, the idea is to first find a point on each contour component and then trace the component by traversing $\mathcal{VD}(S_a)$ and $\mathcal{VD}(S_b)$ simultaneously. This strategy follows Kirkpatrick's algorithm [10] for merging two standard Voronoi diagrams. However, we cannot directly apply Kirkpatrick's algorithm because his method for finding a point in each contour component is not applicable to the weighted case. More specifically, his method relies on the property that the Euclidean minimum spanning tree of a point set in the plane must be a subgraph of the dual graph of its standard Voronoi diagram. However, this is not true anymore for the additively-weighted Voronoi diagrams. We make it formally as an observation below.

► **Observation 5.** *The Euclidean minimum spanning tree of a set of points in the plane is not necessarily a subgraph of the dual graph of the additively-weighted Voronoi diagram of the point set.*

Proof. Figure 4 gives an example for the observation with $S = \{p_1, p_2, p_3, p_4\}$. It is obtained by setting $p_1 = (0, 12)$, $p_2 = (5, 0)$, $p_3 = (0, -12)$, and $p_4 = (-5, 0)$ with weights $w_1 = -10$, $w_2 = 0$, $w_3 = -10$, and $w_4 = 0$. We find that the unweighted distances are $d(p_1, p_2) = d(p_1, p_4) = d(p_3, p_2) = d(p_3, p_4) = 13$, $d(p_1, p_3) = 24$, and $d(p_2, p_4) = 10$. Since (p_2, p_4) is the closest pair among the four points of S , $\overline{p_2 p_4}$ must be an edge in the Euclidean minimum spanning tree of S . However, there is no edge between p_2 and p_4 in the dual graph of the additively-weighted Voronoi diagram of S because their Voronoi regions are not adjacent. ◀

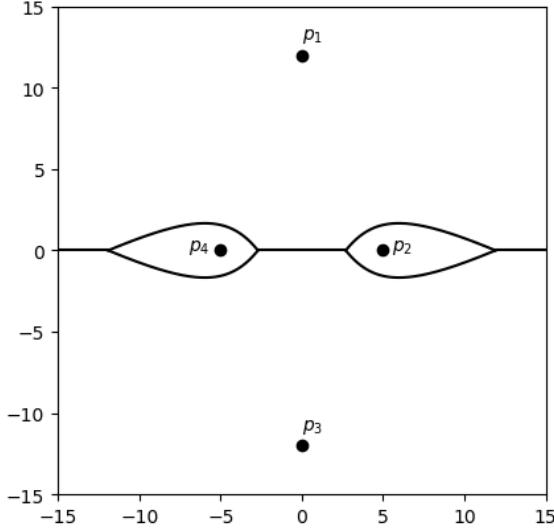


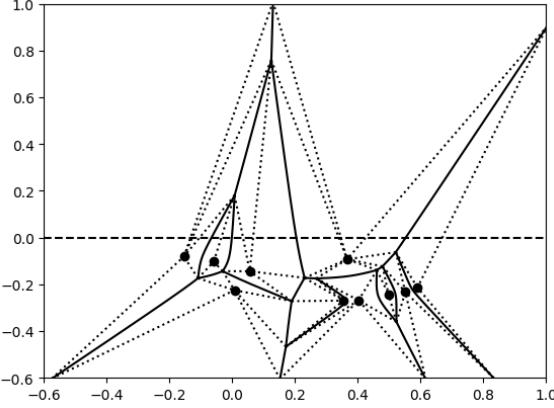
Figure 4 Illustrating the additively-weighted Voronoi diagram of four points $\{p_1, p_2, p_3, p_4\}$ for Observation 5.

In our problem, we are interested in merging $\mathcal{VD}_+(S_a)$ and $\mathcal{VD}_+(S_b)$ into $\mathcal{VD}_+(S)$, so it suffices to compute the portions of the contour above ℓ . With the help of ℓ , it is relatively easy to find a point on each contour component using the following property proved in Lemma 7: Every contour component above ℓ must intersect ℓ .

At a high level, our algorithm has two main procedures. The first one is to identify the intersections between the contour and ℓ . The second procedure is to start at these intersection points and trace each component of the contour above ℓ .

The first main procedure: Finding intersections between the contour and ℓ . By definition, ℓ is divided into segments by its intersections with $\mathcal{VD}_+(S_a)$, which we call ℓ -edges of $\mathcal{VD}_+(S_a)$; similarly, we define ℓ -edges for $\mathcal{VD}_+(S_b)$. We sweep ℓ from left to right, looking for places where the contour intersects ℓ . We start with the leftmost ℓ -edge of $\mathcal{VD}_+(S_a)$ and the leftmost ℓ -edge of $\mathcal{VD}_+(S_b)$. At each step, we are on some ℓ -edge e_a of $\mathcal{VD}_+(S_a)$ and some ℓ -edge e_b of $\mathcal{VD}_+(S_b)$. Let $s_a \in S_a$ be the site associated with the cell of $\mathcal{VD}_+(S_a)$ containing e_a ; define $s_b \in S_b$ similarly. We compute the bisector $B(s_a, s_b)$ and determine where it intersects ℓ . The bisector is a hyperbolic arc and ℓ is a straight line, so they have at most two intersections p_1 and p_2 . If $p_i \in e_a \cap e_b$, then p_i is a point of intersection between the contour and ℓ . In this way, we can compute all intersections between ℓ and the contour. Since the combinatorial sizes of $\mathcal{VD}_+(S_a)$ and $\mathcal{VD}_+(S_b)$ are $O(n)$, this procedure computes $O(n)$ intersections between ℓ and the contour in $O(n)$ time.

The second main procedure: Tracing the contour. We trace the contour components from the intersection points computed above. Specifically, for each intersection p , we trace the contour component containing p as follows. Suppose that p is on an ℓ -edge e_a of $\mathcal{VD}_+(S_a)$ and an ℓ -edge e_b of ℓ in $\mathcal{VD}_+(S_b)$. These edges are associated with sites $s_a \in S_a$ and $s_b \in S_b$. Our trace begins at p and continues above ℓ along the bisector $B(s_a, s_b)$. This bisector enters a Voronoi region R_a of $\mathcal{VD}_+(S_a)$ and a region R_b of $\mathcal{VD}_+(S_b)$. We find which edge of R_a or R_b the bisector intersects first. If no intersection exists or the bisector first intersects ℓ ,



■ **Figure 5** The dotted segments are spokes. Our algorithm only uses the portions of these spokes above ℓ , the dashed line.

then we finish the trace by reporting that the portion of $B(s_a, s_b)$ past p is an edge of the contour. Otherwise, assume that we intersect an edge e'_a of R_a before an edge of R_b (the case where we intersect an edge of R_b first is handled the same way) and denote this point of intersection by p' . We rule out the case where $B(s_a, s_b)$ intersects a vertex instead of an edge because if we were to intersect a vertex, this vertex would be equidistant to three sites in S_a and one site in S_b , which would contradict our general position assumption that no point is equidistant to four sites of $S = S_a \cup S_b$. We report that the portion of $B(s_a, s_b)$ between p and p' is an edge of the contour. Then, we rename R_a to be the Voronoi region of $\mathcal{VD}_+(S_a)$ on the other side of e'_a and update $p \leftarrow p'$. We then continue the tracing from p following the same process as above.

Our tracing algorithm is similar to the well-known algorithm for merging the standard Voronoi diagrams of two sets of points separated by a line [15]. One difficulty with our algorithm is efficiently determining which edges of R_a and R_b the contour intersects first. This may not be a constant time operation since R_a and R_b may have many edges. The merge algorithm by Shamos and Hoey [15] takes advantage of the fact that the contour in their problem is monotone so that they can find all contour edges in a region by a single scan of the boundary of that region. In our problem, the contour may not be monotone. To resolve the issue, we follow the same technique used by Kirkpatrick [10] for merging standard Voronoi diagrams of two arbitrary sets of points. Specifically, before our tracing algorithm, we subdivide Voronoi regions of $\mathcal{VD}_+(S_a)$ and $\mathcal{VD}_+(S_b)$ each into sub-regions of at most four edges by drawing segments between each site and each vertex of the Voronoi region of the site (see Figure 5; we can do this because each Voronoi region is star-shaped by Observation 4); as in [10], we refer to these segments as *spokes*. Because each sub-region only has at most four edges, finding where a bisector intersects a sub-region can be done in $O(1)$ time. We then apply our above tracing algorithm using these subdivisions of $\mathcal{VD}_+(S_a)$ and $\mathcal{VD}_+(S_b)$. Each tracing step now finds an intersection between the contour and either a spoke or a Voronoi edge in constant time. As such, the total time of the tracing procedure is linear in the number of such intersections. By Lemma 8, the number of such intersections, and hence the runtime of the tracing procedure, is $O(n)$. Therefore, the total time of the algorithm for merging $\mathcal{VD}_+(S_a)$ and $\mathcal{VD}_+(S_b)$ is $O(n)$. This proves Theorem 2.

4.3 Useful lemmas

It remains to prove Lemmas 6, 7 and 8, which our algorithm relies on.

Recall that the contour also includes its portions below ℓ , i.e., it is defined with respect to the complete Voronoi diagram $\mathcal{VD}(S)$. We first have the following lemma, which is also needed in the proof of Lemma 7; a similar result on the standard Voronoi diagrams is already used in [10].

► **Lemma 6.** *Each contour component never terminates or splits; that is, it has the topology of an infinite line or a circle.*

Proof. A contour component is made up of edges in $\mathcal{VD}(S)$, so if it were to terminate or split, it would be at a Voronoi vertex of $\mathcal{VD}(S)$. Due to our general position assumption that no point is equidistant to four sites, each Voronoi vertex in $\mathcal{VD}(S)$ is adjacent to three sites in S . If the contour hits a Voronoi vertex v , then at least one of these sites must be in S_a and at least one must be in S_b . Without loss of generality, let these sites be s_1, s_2 , and s_3 with $s_1, s_2 \in S_a$ and $s_3 \in S_b$. The Voronoi edge between s_1 and s_3 and the Voronoi edge between s_2 and s_3 will be on the contour, so the contour will not terminate at v . The edge between s_1 and s_2 will not be on the contour, so the contour will not split at v . ◀

► **Lemma 7.** *If a contour component contains a point above ℓ , then the contour component must intersect ℓ .*

Proof. Lemma 6 establishes that a contour component divides the plane into two regions, called *contour regions*. Notice that because a contour component is made up of edges in $\mathcal{VD}(S)$, each contour region must contain at least one Voronoi region of $\mathcal{VD}(S)$ and thus contains at least one site of S by Observation 4.

Now assume to the contrary that a contour component C contains a point above ℓ but C does not intersect ℓ . Then, the entire C is above ℓ . As such, one of the contour regions divided by C must be entirely in the halfplane above ℓ ; let R be the region. This implies that the sites of S contained in R must be above ℓ , but this contradicts the fact that all sites of S are below ℓ . ◀

► **Lemma 8.** 1. *The total number of intersections between the contour and the Voronoi edges of the complete Voronoi diagrams $\mathcal{VD}(S_a)$ and $\mathcal{VD}(S_b)$ is at most $O(n)$.*
 2. *The total number of intersections between the contour and the spokes of the complete Voronoi diagrams $\mathcal{VD}(S_a)$ and $\mathcal{VD}(S_b)$ is at most $O(n)$.*

Proof. We adapt the proof from [10] for a similar lemma on standard Voronoi diagrams.

Notice that the intersection between the contour and a Voronoi edge in $\mathcal{VD}(S_a)$ or $\mathcal{VD}(S_b)$ is a vertex in $\mathcal{VD}(S)$. There are $O(n)$ vertices in $\mathcal{VD}(S)$, so the total number of intersections between the contour and the Voronoi edges of $\mathcal{VD}(S_a)$ and $\mathcal{VD}(S_b)$ is at most $O(n)$. This proves the first lemma statement.

To prove the second lemma statement, we show that the contour can intersect each spoke at most once. We exploit the fact that Voronoi regions are star-shaped (Observation 4). If the contour intersects a spoke of the Voronoi region for site s in $\mathcal{VD}(S_a)$ or $\mathcal{VD}(S_b)$, then the open segment between s and this intersection will lie in the Voronoi region of s in $\mathcal{VD}(S)$. Because this segment is in the Voronoi region for s in $\mathcal{VD}(S)$, the contour cannot intersect this segment.

Now, assume for the sake of contradiction that the contour were to intersect a spoke twice. This would mean the closer to s of the two intersections would lie on the segment between

s and the further of the two intersections, which we have shown above to be impossible. Therefore, the contour can only intersect each spoke at most once, and there are $O(n)$ spokes in $\mathcal{VD}(S_a)$ and $\mathcal{VD}(S_b)$, so the total number of intersections between the contour and the spokes is at most $O(n)$. \blacktriangleleft

5 Algebraic decision tree algorithm

Under the algebraic decision tree model, where the time complexity is measured only by the number of comparisons, we show that the IOAWNN-SL problem can be solved using $O(n \log n)$ comparisons. Consequently, we can solve the shortest path problem in weighted unit-disk graphs in $O(n \log n)$ time under the algebraic decision tree model. In the following, we first describe an $O(n \log^2 n)$ time algorithm under the conventional computational model and then show how to improve it to $O(n \log n)$ time under the algebraic decision tree model.

5.1 Algorithm under the conventional computational model

Let p_1, p_2, \dots, p_n be the points to be inserted in this order; each point has a weight. Let P denote the set of all these points. Let Q be a set of $O(n)$ query points, such that all points of P are above the x -axis ℓ while all points of Q are below ℓ . For each query point $q \in Q$, we know the timer when the query is conducted, i.e., we know the index i such that the query looks for the nearest neighbor of q among the first i points of P . Our goal is to answer all queries for the points of Q .

We construct a complete binary tree T whose leaves from left to right correspond to points p_1, p_2, \dots, p_n in this order. For each node $v \in T$, let P_v denote the set of points that are in the leaves of the subtree rooted at v . Let $\mathcal{VD}(P_v)$ be the additively-weighted Voronoi diagram for the weighted points of P_v ; let $\mathcal{VD}_+(P_v)$ be the portion of $\mathcal{VD}(P_v)$ above ℓ . We construct $\mathcal{VD}_+(P_v)$. If we construct $\mathcal{VD}_+(P_v)$ for all nodes v of T in a bottom-up manner and use our linear time merge algorithm in Theorem 2, constructing the diagrams $\mathcal{VD}_+(P_v)$ for all nodes $v \in T$ can be done in $O(n \log n)$ time. In addition, we construct Kirkpatrick's point location data structure [11] on $\mathcal{VD}_+(P_v)$ for each node $v \in T$, which takes $O(|P_v|)$ time.¹ Note that we use Kirkpatrick's point location data structure instead of others such as the one in [7] because we will need to apply a technique from [6] that requires Kirkpatrick's data structure. Constructing the point location data structures for all nodes of T takes $O(n \log n)$ time.

Consider a query point $q \in Q$. Suppose we are looking for the nearest neighbor of q among the first i points p_1, p_2, \dots, p_i of P . Let v_i be the leaf of T corresponding to p_i . Following the path in T from the root to v_i , we can find a set V_q of nodes of T such that the union of P_v for all $v \in V_q$ is exactly $\{p_1, p_2, \dots, p_i\}$. As such, the query can be answered after performing $O(\log n)$ point location queries on $\mathcal{VD}_+(P_v)$ for all $v \in V_q$. As each point

¹ Note that Kirkpatrick's data structure is originally for planar subdivisions in which each edge is a straight line segment. However, as discussed in [11], the algorithm also works for additively-weighted Voronoi diagrams (and other types of Voronoi diagrams) since each cell of the diagram is star-shaped. A subtle issue in our problem is that $\mathcal{VD}_+(P_v)$ is only the portion of the complete diagram $\mathcal{VD}(P_v)$ above ℓ , and each cell of $\mathcal{VD}_+(P_v)$ does not contain its site. To circumvent the issue, we can enlarge each cell of $\mathcal{VD}_+(P_v)$ by including its site, as follows. For each cell $R \in \mathcal{VD}_+(P_v)$, if \overline{ab} is a maximal segment of $R \cap \ell$, then we add the triangle Δpab to R , where p is the site of R . Note that Δpab must be inside the cell of p in $\mathcal{VD}(P_v)$, denoted by R' . As such, the enlarged region R is still star-shaped, contains its site p , and is a subset of R' . We can then construct Kirkpatrick's point location data structure on the subdivision of all these enlarged regions R .

location query takes $O(\log n)$ time, answering the nearest neighbor query for q can be done in $O(\log^2 n)$ time. Therefore, the total time for answering the queries for all points of Q is $O(n \log^2 n)$.

5.2 Algorithm under the algebraic decision tree model

The above solves the problem in $O(n \log^2 n)$ time. To improve the time to $O(n \log n)$, the bottleneck is to solve all $O(n \log n)$ point location queries. For this, we resort to a technique recently developed by Chan and Zheng [6] under the algebraic decision tree model.

For our problem, we can simply apply [6, Theorem 7.2] as a black box to solve all our point location queries using $O(n \log n)$ comparisons. Roughly speaking, the theorem says the following: Given a collection of t planar triangulated subdivisions S_i of the plane, $1 \leq i \leq t$, of total size L , and a collection of t point sets Q_i , $1 \leq i \leq t$, of total size M , we can locate the cell of S_i that contains each query point $q \in Q_i$, for all $1 \leq i \leq t$, using a total of $O(L + M + N \log N)$ comparisons, where N is the input size.

In our problem, for each node $v \in T$, we have $\mathcal{VD}_+(P_v)$ as a planar subdivision. Note that the theorem statement requires the input planar subdivisions to be triangulated. The triangulation is mainly used to construct Kirkpatrick's point location data structure [11] on each planar subdivision. Since we already have Kirkpatrick's point location data structure for each $\mathcal{VD}_+(P_v)$ as discussed above, we can simply follow the same algorithm of the theorem. For each node $v \in T$, we also have a subset Q_v of points of Q that need to locate the cells of $\mathcal{VD}_+(P_v)$ containing them. Hence, we have $t = O(n)$, the number of nodes of T . Also, the total size of $\mathcal{VD}_+(P_v)$ for all nodes $v \in T$ is $O(n \log n)$ and the total size of Q_v for all $v \in T$ is $O(n \log n)$. Thus, we have $L = O(n \log n)$ and $M = O(n \log n)$. Finally, since the input of our problem consists of the $O(n)$ points of P and Q , we have $N = O(n)$. Consequently, according to the theorem, all point location queries can be solved using $O(L + M + N \log N)$ comparisons, which is $O(n \log n)$.

We remark that Chan and Zheng solved the offline insertion-only nearest neighbor problem as an example of [6, Theorem 7.2]. Our application of the theorem to the IOAWNN-SL problem mirrors their example. The main difference is that our algorithm for merging additively-weighted Voronoi diagrams above a separating line replaces their use of Kirkpatrick's [10] algorithm for merging standard Voronoi diagrams and that we modify the use of Kirkpatrick's [11] point location data structure as described in Footnote 1. Therefore, we refer readers interested in the details of Chan and Zheng's technique to their paper.

It should be noted that this algorithm is only applicable to the algebraic decision tree model of computation. The bound on the number of real-RAM operations given by Chan and Zheng for their technique is $n^{O(n)}$.

6 Conclusion

In this paper, we presented an $O(n \log^2 n / \log \log n)$ time algorithm for the SSSP problem in weighted unit-disk graphs in the standard real-RAM model of computation. The best known lower bound for this problem is $\Omega(n \log n)$, which is proven in the algebraic decision tree model and also applies to the real-RAM model. We have given an algorithm matching this lower bound in the algebraic decision tree model. Resolving the gap between $O(n \log^2 n / \log \log n)$ and $\Omega(n \log n)$ in the real-RAM model remains an open problem. Lower bounds in computational geometry are commonly proven in the algebraic decision tree model, so our latter algorithm is also significant because it rules out the possibility of the gap being closed by an improved lower bound in the algebraic decision tree model. In addition to our

algorithms for the SSSP problem in weighted unit-disk graphs, we derived a data structure for the IOAWNN-SL problem with $O(\log^2 n / \log \log n)$ query time and $O(\log^2 n / \log \log n)$ amortized insertion time. We also gave an $O(n)$ time algorithm for the problem of merging two additively-weighted Voronoi diagrams above a separating line. We leave it as an open problem to determine if these results can also be achieved in the case where the separating line constraint is relaxed.

References

- 1 Jon L. Bentley. Decomposable searching problems. *Information Processing Letters*, 8:244–251, 1979. doi:[10.1016/0020-0190\(79\)90117-0](https://doi.org/10.1016/0020-0190(79)90117-0).
- 2 Bruce W. Brewer and Haitao Wang. An improved algorithm for shortest paths in weighted unit-disk graphs. In *Proceedings of the 36th Canadian Conference on Computational Geometry (CCCG 2024)*, pages 57–64, 2024.
- 3 Sergio Cabello and Miha Jejčič. Shortest paths in intersection graphs of unit disks. *Computational Geometry: Theory and Applications*, 48:360–367, 2015. doi:[10.1016/j.comgeo.2014.12.003](https://doi.org/10.1016/j.comgeo.2014.12.003).
- 4 Timothy M. Chan. Dynamic geometric data structures via shallow cuttings. *Discrete and Computational Geometry*, 64:1235–1252, 2020. doi:[10.1007/s00454-020-00229-5](https://doi.org/10.1007/s00454-020-00229-5).
- 5 Timothy M. Chan and Dimitrios Skrepetos. All-pairs shortest paths in unit-disk graphs in slightly subquadratic time. In *Proceedings of the 27th International Symposium on Algorithms and Computation (ISAAC)*, pages 24:1–24:13, 2016. doi:[10.4230/LIPIcs.ISAAC.2016.24](https://doi.org/10.4230/LIPIcs.ISAAC.2016.24).
- 6 Timothy M. Chan and Da Wei Zheng. Hopcroft’s problem, log-star shaving, 2D fractional cascading, and decision trees. *ACM Transactions on Algorithms*, 20(3):24:1–24:27, 2023. doi:[10.1145/3591357](https://doi.org/10.1145/3591357).
- 7 Herbert Edelsbrunner, Leonidas J. Guibas, and Jorge Stolfi. Optimal point location in a monotone subdivision. *SIAM Journal on Computing*, 15(2):317–340, 1986. doi:[10.1137/0215023](https://doi.org/10.1137/0215023).
- 8 Steven Fortune. A sweepline algorithm for Voronoi diagrams. *Algorithmica*, 2:153–174, 1987. doi:[10.1007/BF01840357](https://doi.org/10.1007/BF01840357).
- 9 Haim Kaplan, Wolfgang Mulzer, Liam Roditty, Paul Seiferth, and Micha Sharir. Dynamic planar Voronoi diagrams for general distance functions and their algorithmic applications. *Discrete and Computational Geometry*, 64:838–904, 2020. doi:[10.1007/s00454-020-00243-7](https://doi.org/10.1007/s00454-020-00243-7).
- 10 David G. Kirkpatrick. Efficient computation of continuous skeletons. In *Proceedings of the 20th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 18–27, 1979. doi:[10.1109/SFCS.1979.15](https://doi.org/10.1109/SFCS.1979.15).
- 11 David G. Kirkpatrick. Optimal search in planar subdivisions. *SIAM Journal on Computing*, 12(1):28–35, 1983. doi:[10.1137/0212002](https://doi.org/10.1137/0212002).
- 12 Chih-Hung Liu. Nearly optimal planar k nearest neighbors queries under general distance functions. *SIAM Journal on Computing*, 51:723–765, 2022. doi:[10.1137/20M1388371](https://doi.org/10.1137/20M1388371).
- 13 Mark H. Overmars. *The Design of Dynamic Data Structures*, volume 156 of *Lecture Notes in Computer Science*. Springer, 1983. doi:[10.1007/BF0014927](https://doi.org/10.1007/BF0014927).
- 14 Liam Roditty and Michael Segal. On bounded leg shortest paths problems. *Algorithmica*, 59:583–600, 2011. doi:[10.1007/s00453-009-9322-3](https://doi.org/10.1007/s00453-009-9322-3).
- 15 Michael I. Shamos and Dan Hoey. Closest-point problems. In *Proceedings of the 16th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 151–162, 1975. doi:[10.1109/SFCS.1975.8](https://doi.org/10.1109/SFCS.1975.8).
- 16 Haitao Wang and Jie Xue. Near-optimal algorithms for shortest paths in weighted unit-disk graphs. *Discrete and Computational Geometry*, 64:1141–1166, 2020. doi:[10.1007/s00454-020-00219-7](https://doi.org/10.1007/s00454-020-00219-7).