

Approximate Euclidean Shortest Paths amid Convex Obstacles*

Pankaj K. Agarwal R. Sharathkumar Hai Yu

Department of Computer Science
Duke University
{pankaj, sharath, fishhai}@cs.duke.edu

Abstract

We develop algorithms and data structures for the approximate Euclidean shortest path problem amid a set \mathcal{P} of k convex obstacles in \mathbb{R}^2 and \mathbb{R}^3 , with a total of n faces. The running time of our algorithms is linear in n , and the size and query time of our data structure are independent of n . We follow a “core-set” based approach, i.e., we quickly compute a small sketch \mathcal{Q} of \mathcal{P} whose size is independent of n and then compute approximate shortest paths with respect to \mathcal{Q} .

1 Introduction

The Euclidean shortest-path problem is defined as follows: Given a set of pairwise-disjoint polyhedral obstacles in \mathbb{R}^2 or \mathbb{R}^3 and two points s and t in the free space, compute the shortest path between s and t that avoids the interior of the obstacles. This problem has received much attention in computational geometry and robotics; see, e.g., the survey paper [16] for a comprehensive review. Let n denote the total complexity of the obstacles. In \mathbb{R}^2 , Hershberger and Suri [15] presented an optimal $O(n \log n)$ -time algorithm for computing the exact Euclidean shortest path between two points amid polygonal obstacles. In \mathbb{R}^3 , Canny and Reif [6] proved that the problem is NP-Hard, and Mitchell and Sharir [17] showed that the problem remains NP-Hard even for pairwise-disjoint convex obstacles. The best known exact algorithm runs in singly exponential time [22]; the running time can be improved to $O(n^{O(k)})$ [24] if the obstacles consist of k disjoint convex polytopes with a total of n faces.

The known hardness results in \mathbb{R}^3 have motivated researchers to develop fast approximation algorithms and exact algorithms for special cases. Papadimitriou [19] gave an $O(n^4(L + \log(n/\varepsilon))/\varepsilon^2)$ -time algo-

rithm for computing an ε -short path, i.e., a path whose length is at most $(1 + \varepsilon)$ times the length of the shortest path. Here L is the number of bits of precision in the model of computation. A rigorous analysis of Papadimitriou’s algorithm was later given by Choi *et al.* [9]; see also [5]. Clarkson [10] proposed a different algorithm for computing an ε -short path; the running time of his algorithm is roughly $O(n^2 \log^{O(1)} n/\varepsilon^4)$ (the running time also depends on the geometry of obstacles).

The special case of computing a shortest path between two points along the surface of a single convex polytope has been widely studied. After an initial $O(n^3 \log n)$ algorithm by Sharir and Schorr [25], the bound was improved to $O(n^2)$ by Chen and Han [8] (see also [18]). A major open problem was whether the running time can be improved to $O(n \log n)$. Such an algorithm is recently developed by Shreiber and Sharir [23]. Hershberger and Suri [14] proposed a simple $O(n)$ algorithm to compute a 1-short path. Later, Agarwal *et al.* [3] developed an $O(n \log(1/\varepsilon) + 1/\varepsilon^3)$ algorithm to compute an ε -short path (see also [2]). Their algorithm computes a convex polytope of size $O(1/\varepsilon^{3/2})$ that approximates the original polytope and runs a quadratic shortest-path algorithm on the simplified polytope. For a fixed point s , the aforementioned exact shortest-path algorithms [25, 8, 18] can also construct a data structure for a given polytope, so that the shortest distance from a fixed source s to a query point can be answered quickly. Har-Peled [13] described a data structure of size $O((n/\varepsilon) \log(1/\varepsilon))$ that can compute an ε -short distance from s to a query point in $O(\log(n/\varepsilon))$ time. His technique applies to the more general case of polyhedral obstacles, albeit with much worse preprocessing time (roughly $O(n^4/\varepsilon^6)$) and space complexity ($O(n^2/\varepsilon^{4+\delta})$ for any $\delta > 0$). There is also extensive work on computing shortest paths on a nonconvex surface: see [4, 16] and the references therein.

In this paper we study the problem of computing approximate shortest paths amid convex obstacles. Let

*Work on this paper is supported by NSF under grants CNS-05-40347, CFF-06-35000, and DEB-04-25465, by ARO grants W911NF-04-1-0278 and W911NF-07-1-0376, by an NIH grant 1P50-GM-08183-01, by a DOE grant OEG-P200A070505, and by a grant from the U.S.–Israel Binational Science Foundation.

$\mathcal{P} = \{P_1, \dots, P_k\}$ be a set of k pairwise-disjoint¹ convex obstacles in \mathbb{R}^d ($d = 2, 3$), where each P_i has n_i facets. Set $n = \sum_{i=1}^k n_i$, which denotes the total complexity of \mathcal{P} . The *free space* of \mathcal{P} , denoted by $\mathcal{F}(\mathcal{P})$, is defined as the closure of $\mathbb{R}^3 \setminus \bigcup \mathcal{P}$. Given two points $s, t \in \mathcal{F}(\mathcal{P})$, let $\pi_{\mathcal{P}}(s, t)$ denote the shortest path between s and t in $\mathcal{F}(\mathcal{P})$, and let $d_{\mathcal{P}}(s, t)$ denote the length of $\pi_{\mathcal{P}}(s, t)$. Let $\varepsilon > 0$ be a fixed parameter. The ε -*approximate shortest-path problem* is to compute a path $\pi \subseteq \mathcal{F}(\mathcal{P})$ between s and t whose length is at most $(1 + \varepsilon)d_{\mathcal{P}}(s, t)$. Such a path π is called an ε -*short path*, and its length is called an ε -*short distance*. For a fixed source $s \in \mathcal{F}(\mathcal{P})$, the *approximate shortest-path query problem* is to preprocess \mathcal{P} into a data structure so that for any query point $t \in \mathcal{F}(\mathcal{P})$, an ε -short distance (or an ε -short path) between s and t can be reported quickly.

We propose algorithms for computing approximate shortest paths between two points whose running time depends linearly in n , and data structures for answering approximate shortest-path queries to a fixed source whose size is independent of n . As far as we know, our results are the first to achieve this kind of performance. More specifically, we obtain the following:

- In \mathbb{R}^2 , for any two points $s, t \in \mathcal{F}(\mathcal{P})$ and a parameter $0 < \varepsilon \leq 1$, an ε -short path between s and t can be computed in $O(n + (k/\sqrt{\varepsilon}) \log(k/\varepsilon))$ time (Section 2).
- In \mathbb{R}^3 , for any two points $s, t \in \mathcal{F}(\mathcal{P})$ and a parameter $0 < \varepsilon \leq 1$, an ε -short path between s and t can be computed in $O(n + (k^4/\varepsilon^7) \log^3(k/\varepsilon))$ time (Section 3).
- In \mathbb{R}^3 , for a fixed source $s \in \mathcal{F}(\mathcal{P})$ and a parameter $0 < \varepsilon \leq 1$, a data structure of size $O(k^3 \text{poly}(\log k, 1/\varepsilon))$ can be constructed in $O(n \log k + k^7 \text{poly}(\log k, 1/\varepsilon))$ so that an ε -short distance between s and a query point $t \in \mathcal{F}(\mathcal{P})$ can be reported in $O(\log^2(k/\varepsilon) \log(1/\varepsilon))$ time (Section 4).

As can be seen, when $k \ll n$, our algorithms and data structures perform much better in terms of space and running time than previously known results. Our algorithms quickly compute a small “sketch” \mathcal{Q} of polytopes in \mathcal{P} whose size is independent of n and then compute an approximate shortest path in $\mathcal{F}(\mathcal{Q})$.

¹For our purpose it suffices to assume that the interiors of polytopes in \mathcal{P} are pairwise disjoint. The boundaries may touch each other, so nonconvex obstacles can be decomposed into convex obstacles. However, for simplicity of presentation, we assume that the boundaries also do not intersect.

2 Approximate Shortest Paths in \mathbb{R}^2

For a parameter $\varepsilon > 0$, a set $\mathcal{Q} = \{Q_1, \dots, Q_k\}$ of k pairwise-disjoint convex polygons is called an ε -*sketch* of \mathcal{P} if

- (i) $P_i \subseteq Q_i$, for $i = 1, \dots, k$;
- (ii) for any $s, t \in \mathcal{F}(\mathcal{Q})$, $d_{\mathcal{Q}}(s, t) \leq (1 + \varepsilon)d_{\mathcal{P}}(s, t)$.

Since $\mathcal{F}(\mathcal{Q}) \subseteq \mathcal{F}(\mathcal{P})$, $\pi_{\mathcal{Q}}(s, t) \subseteq \mathcal{F}(\mathcal{P})$ for any $s, t \in \mathcal{F}(\mathcal{Q})$. Therefore (ii) implies that $\pi_{\mathcal{Q}}(s, t)$ is an ε -short path between s and t . Next we describe an algorithm to construct an ε -sketch \mathcal{Q} of small complexity.

Set $r = \lceil \sqrt{2\pi/\sqrt{\varepsilon}} \rceil$. Let $u_j = (\cos 2j\pi/r, \sin 2j\pi/r)$ for $0 \leq j < r$ and set $\mathcal{N} = \{u_j \mid 0 \leq j < r\}$. \mathcal{N} is a uniform sample of directions in \mathbb{S}^1 ; for simplicity, we assume that r is divisible by 4 so that $0, \pi/2, \pi, 3\pi/2 \in \mathcal{N}$. For a convex polygon P and a direction $u_j \in \mathcal{N}$, let $\ell_j(P)$ denote the line passing through the extreme point of P in direction u_j , and let $h_j(P)$ denote the (closed) halfplane bounded by $\ell_j(P)$ and containing P . Set $H_i = \{h_j(P_i) \mid 0 \leq j < r\}$.

We call a pair $\{P_i, P_j\}$ *vertically visible* if there is a vertical segment e connecting ∂P_i to ∂P_j whose relative interior does not intersect any polygon of \mathcal{P} (see Figure 1). Let $\ell_{ij}(= \ell_{ji})$ be the line that separates P_i and P_j . Let Φ be the set of vertically visible pairs. We note that $|\Phi| = O(k)$ and it can be computed in $O(n + k \log n) = O(n + k \log k)$ time by a sweep-line algorithm.

For each P_i , set $\mathcal{P}_i = \{P_j \mid \{P_i, P_j\} \in \Phi\}$. For each $P_j \in \mathcal{P}_i$, We translate the separating line ℓ_{ij} until it supports P_i . Let g_{ij} be the halfplane bounded by this line and containing P_i . Set $G_i = \{g_{ij} \mid P_j \in \mathcal{P}_i\}$. We define

$$Q_i = \left(\bigcap_{g \in G_i} g \right) \cap \left(\bigcap_{h \in H_i} h \right).$$

Set $\mathcal{Q} = \{Q_1, \dots, Q_k\}$. We next prove that \mathcal{Q} is an ε -sketch of \mathcal{P} . We call a vertex $v \in \partial Q_i$ *new* if $v \notin \partial P_i$. Each edge of Q_i touches a vertex of P_i . For each new vertex v of Q_i , let l_v (resp., r_v) be the vertex of P_i lying on the edge adjacent to v in counterclockwise (resp., clockwise) direction. We denote by Δ_v the triangle formed by l_v, v , and r_v . Using the fact that the internal angle of each new vertex in Q_i is at least $\pi - \sqrt{2\varepsilon}$, we can prove the following.

LEMMA 2.1. *For any pair of points $p \in \overline{vl_v}$ and $q \in \overline{vr_v}$,*

$$\|pv\| + \|vq\| \leq (1 + \varepsilon)\|pq\| \leq (1 + \varepsilon)d_{\mathcal{P}}(p, q).$$

Proof. We note that

$$\|pv\| + \|qv\| \leq \|pq\| / \sin(\angle pvq/2)$$




Figure 1. The pairs $\{P_1, P_2\}$ and $\{P_1, P_3\}$ are vertically visible, but $\{P_2, P_3\}$ is not. A path $\pi \subseteq \mathcal{F}(\mathcal{P})$ can be modified into $\pi' \subseteq \mathcal{F}(\mathcal{Q})$ so that $|\pi'| \leq (1 + \varepsilon)|\pi|$.

$$\begin{aligned} &\leq \|pq\| / \sin(\pi/2 - \sqrt{\varepsilon/2}) \\ &\leq \|pq\| / (1 - \varepsilon/2) \\ &\leq (1 + \varepsilon)\|pq\| \leq (1 + \varepsilon)d_{\mathcal{P}}(p, q), \end{aligned}$$

as claimed. \square

LEMMA 2.2. \mathcal{Q} is an ε -sketch of \mathcal{P} .

Proof. By construction of H_i , the orthogonal bounding boxes of P_i and Q_i are identical. Therefore a pair of polygons $\{Q_i, Q_j\}$ intersect if and only if there exists a vertically visible pair of polygons (P_i, P_k) such that $\{Q_i, Q_k\}$ intersect. Since g_{ij} and g_{ji} are disjoint, $Q_i \subseteq g_{ij}$ and $Q_j \subseteq g_{ji}$ are disjoint. Furthermore, since each halfplane in $G_i \cup H_i$ contains P_i , $P_i \subseteq Q_i$. It thus remains to prove that for any $s, t \in \mathcal{F}(\mathcal{Q})$, $d_{\mathcal{Q}}(s, t) \leq (1 + \varepsilon)d_{\mathcal{P}}(s, t)$.

Set $\Sigma = \{\Delta_v \mid v \text{ is a new vertex of some } Q_i \in \mathcal{Q}\}$. The set Σ consists of obtuse-angled triangles whose interiors are pairwise-disjoint and which cover the region $\mathcal{F}(\mathcal{Q}) \setminus \mathcal{F}(\mathcal{P})$. For any pair of points $s, t \in \mathcal{F}(\mathcal{Q})$, let $\pi = \pi_{\mathcal{P}}(s, t)$. If π does not intersect any triangle in Σ , then $\pi \subseteq \mathcal{F}(\mathcal{Q})$ and $d_{\mathcal{Q}}(s, t) = d_{\mathcal{P}}(s, t)$. Let $\Sigma_{st} = \langle \Delta_{v_1}, \dots, \Delta_{v_m} \rangle \subseteq \Sigma$ be the sequence of triangles that π intersects and let $\langle (p_1, q_1), \dots, (p_m, q_m) \rangle$ be the sequence of pairs of intersection points of π with the boundaries of triangles in Σ_{st} . Set $s = q_0$ and $t = p_{m+1}$. We obtain a path π' from π by replacing each segment $p_i q_i$ with $p_i v_i \cup v_i q_i$ (see Figure 1). Clearly, $\pi' \subseteq \mathcal{F}(\mathcal{Q})$. In addition,

$$\begin{aligned} |\pi'| &= \sum_{i=0}^m d_{\mathcal{Q}}(q_i, p_{i+1}) + \sum_{i=1}^m (\|p_i v_i\| + \|v_i q_i\|) \\ &\leq \sum_{i=0}^m d_{\mathcal{P}}(q_i, p_{i+1}) + (1 + \varepsilon) \sum_{i=1}^m d_{\mathcal{P}}(p_i, q_i) \\ &\leq (1 + \varepsilon)|\pi|, \end{aligned}$$

thereby implying that $d_{\mathcal{Q}}(s, t) \leq (1 + \varepsilon)d_{\mathcal{P}}(s, t)$. \square

THEOREM 2.1. Given a set \mathcal{P} of k pairwise-disjoint convex polygons of total complexity n in \mathbb{R}^2 , an ε -sketch of \mathcal{P} with size $O(k/\sqrt{\varepsilon})$ can be computed in $O(n + k \log k)$ time.

REMARK. If we assume that the vertices of the input polygons in \mathcal{P} are given sorted in an array, we can compute an ε -sketch \mathcal{Q} of \mathcal{P} in $O((k/\sqrt{\varepsilon}) \log n)$ time.

Using the above theorem, we show how to compute an ε -short path between two points $s, t \in \mathcal{F}(\mathcal{P})$. We treat s and t as two additional (degenerate) obstacles and compute an ε -sketch \mathcal{Q} of $\mathcal{P} \cup \{s, t\}$. This ensures that $s, t \in \mathcal{F}(\mathcal{Q})$. We then apply the the algorithm of Hershberger and Suri [15] to obtain $\pi_{\mathcal{Q}}(s, t)$; the running time is $O((k/\sqrt{\varepsilon}) \log(k/\sqrt{\varepsilon}))$. Since \mathcal{Q} is an ε -sketch, $\pi_{\mathcal{Q}}(s, t)$ is an ε -short path of s, t in $\mathcal{F}(\mathcal{P})$. Moreover, the path consists of $O(k/\sqrt{\varepsilon})$ edges.

COROLLARY 2.1. Given a set \mathcal{P} of k pairwise-disjoint convex polygons of total complexity n in \mathbb{R}^2 and two points $s, t \in \mathcal{F}(\mathcal{P})$, an ε -short path between s and t which consists of $O(k/\sqrt{\varepsilon})$ edges can be computed in $O(n + (k/\sqrt{\varepsilon}) \log(k/\varepsilon))$ time.

Other immediate applications of Theorem 2.1 include constructing small-sized spanners or shortest-path maps using ε -sketches.

3 Approximate Shortest Paths in \mathbb{R}^3

In this section we present an efficient algorithm for computing approximate shortest paths amid a set of convex polytopes in \mathbb{R}^3 . The basic idea of our algorithm is the same as in the preceding section, i.e., we first compute a sketch of small size for the convex obstacles and then compute a path amid the sketch. However, a simple example shows that one cannot hope for a small-sized sketch that works for *all* pairs of points $s, t \in \mathcal{F}(\mathcal{P})$ simultaneously. Nonetheless, we show that a small-sized sketch can indeed be computed for any *fixed* pair of points $s, t \in \mathcal{F}(\mathcal{P})$, which suffices for our purpose.

Outer approximations. For a set $U \subseteq \mathbb{R}^3$ and a real number $r > 0$, let $U_r = U \oplus \mathbb{B}_r$ where \oplus refers to the Minkowski sum and \mathbb{B}_r denotes a ball of radius r centered at the origin. For a parameter $r > 0$ and a convex polytope P of n vertices, an *outer r -approximation* of P is a convex polytope $P(r)$ such that $P \subseteq P(r) \subseteq P_r$. Dudley [11] has shown that there is a polytope $P(r)$ of size $O(1/\delta)$, where $\delta = r/\text{diam}(P)$. Agarwal *et al.* [3] have shown that such a polytope can be computed in $O(n + (1/\delta) \log(1/\delta))$ time. The concept of outer approximation can be generalized to k pairwise-disjoint convex polytopes as follows. Given a parameter

$r > 0$ and a set $\mathcal{P} = \{P_1, \dots, P_k\}$ of k pairwise-disjoint convex polytopes in \mathbb{R}^3 , we call $\mathcal{Q} = \{Q_1, \dots, Q_k\}$ an *outer r -approximation* of \mathcal{P} if the convex polytopes in \mathcal{Q} are pairwise-disjoint and $P_i \subseteq Q_i \subseteq (P_i)_r$ for $i = 1, \dots, k$.

Let D be the maximum diameter of a polytope in \mathcal{P} . Let $\delta = r/D$. An outer r -approximation \mathcal{Q} of \mathcal{P} of total complexity $O(k^2 + k/\delta)$ can be computed as follows. For each $P_i \in \mathcal{P}$, we first compute Dudley's outer r -approximation $P_i(r)$ of P_i in time $O(n_i + (1/\delta) \log(1/\delta))$. Next, for $j \neq i$, we compute a supporting plane $h_{i,j}$ of P_i that separates P_i and P_j . Let $h_{i,j}^+$ denote the halfspace bounded by $h_{i,j}$ and containing P_i . We set

$$Q_i = P_i(r) \cap \bigcap_{j \neq i} h_{i,j}^+.$$

The resulting set $\mathcal{Q} = \{Q_1, \dots, Q_k\}$ is a set of pairwise-disjoint convex polytopes such that $P_i \subseteq Q_i \subseteq P_i(r) \subseteq (P_i)_r$. Hence \mathcal{Q} is an outer approximation of \mathcal{P} . Since the complexity of each $Q_i \in \mathcal{Q}$ is $O(k + 1/\delta)$, the total complexity of \mathcal{Q} is $O(k^2 + k/\delta)$. Each supporting hyperplane $h_{i,j}$ can be computed by using the Dobkin-Kirkpatrick hierarchies of P_i and P_j in $O(\log |P_i| \cdot \log |P_j|) = O(\log^2 n)$ time. Hence, the time spent in computing \mathcal{Q} is $O(n + k^2 \log^2 n + (k/\delta) \log(1/\delta))$.

ε -Sketches. For two points $s, t \in \mathcal{F}(\mathcal{P})$ and a value such that $d \geq \|st\|$, we show how to construct a set \mathcal{Q} of at most k pairwise-disjoint convex obstacles of total complexity $O(k^2/\varepsilon^{3/2})$ such that

- (C1) $s, t \in \mathcal{F}(\mathcal{Q})$,
- (C2) $d_{\mathcal{Q}}(s, t) \leq (1 + \varepsilon/2)d_{\mathcal{P}}(s, t) + \varepsilon d/8$, and
- (C3) If $d_{\mathcal{Q}}(s, t) \leq 2d$, then $d_{\mathcal{P}}(s, t) \leq d_{\mathcal{Q}}(s, t)$.

Let C_{4d} be a cube centered at s of side length $4d$. Note that $t \in C_{4d}$ because $d \geq \|st\|$. We clip every polytope of \mathcal{P} within C_{4d} and obtain a set \mathcal{P}' of at most k pairwise-disjoint convex obstacles, each of diameter at most $4\sqrt{3}d$. We treat s, t as two additional (degenerate) obstacles and compute an outer r -approximation \mathcal{Q} of $\mathcal{P}' \cup \{s, t\}$ with $r = \varepsilon^{3/2}d/c_k$, where c is a sufficiently large constant. Observe that $s, t \in \mathcal{F}(\mathcal{Q})$. The resulting set \mathcal{Q} has total complexity $O(k^2/\varepsilon^{3/2})$ and can be constructed in time $O(n + k^2 \log^2 k + (k^2/\varepsilon^{3/2}) \log(k/\varepsilon))$. In fact, if we precompute the Dobkin-Kirkpatrick hierarchy of each polytope in \mathcal{P} as well as their pairwise separating supporting planes in a total of $O(n + k^2 \log^2 n)$ time, then \mathcal{Q} can be computed in an additional $O((k^2/\varepsilon^{3/2}) \log(k/\varepsilon))$ time for any $s, t \in \mathcal{F}(\mathcal{P})$ and $d \geq \|st\|$.

We need the following lemma to prove (C2) and (C3).

LEMMA 3.1. *Let P and Q be two convex polytopes such that $P \subseteq Q \subseteq P_r$. For a parameter $0 < \varepsilon \leq 1$ and for any pair of points $p, q \in \partial Q$,*

$$d_{\mathcal{P}}(p, q) \leq d_{\mathcal{Q}}(p, q) \leq (1 + \varepsilon/2)d_{\mathcal{P}}(p, q) + (13 + 100/\sqrt{\varepsilon})r. \quad (3.1)$$

Proof. Let p' and q' be the closest points of p and q on ∂P respectively. Let p'' (resp., q'') be the intersection of ∂P_r with the ray emanating from p' (resp., q') in direction $p'p$ (resp., $q'q$). See Figure 2. Observe that $p''p$ (resp., $q''q$) is the normal of ∂P at p (resp., q). Therefore $\|p'p''\| = \|q'q''\| = r$. Furthermore, since $P \subseteq Q \subseteq P_r$, the segment $\overline{p'p''}$ contains p and the segment $\overline{q'q''}$ contains q . It was shown in [3] that

$$d_{\mathcal{Q}}(p'', q'') \leq (1 + \varepsilon/2)d_{\mathcal{P}}(p', q') + 2\pi r + 2r + 100r/\sqrt{\varepsilon}.$$

Furthermore, observe that

$$d_{\mathcal{Q}}(p, q) \leq \|pp''\| + d_{\mathcal{Q}}(p'', q'') + \|q''q\|$$

and

$$d_{\mathcal{P}}(p', q') \leq \|p'p\| + d_{\mathcal{P}}(p, q) + \|qq'\|.$$

Putting these inequalities together, we obtain (3.1) as claimed. \square




Figure 2. Illustration for the proof of Lemma 3.1.

If the path $\pi_{\mathcal{P}}(s, t)$ does not intersect any polytope in \mathcal{Q} , then $\pi_{\mathcal{Q}}(s, t) = \pi_{\mathcal{P}}(s, t)$ and therefore (C2) holds. So assume that $\pi_{\mathcal{P}}(s, t)$ intersects a polytope of \mathcal{Q} . It is well known that for any $Q_i \in \mathcal{Q}$, the intersection $\pi_{\mathcal{P}}(s, t) \cap Q_i$ consists of at most one connected component [20]. For each polytope $Q_i \in \mathcal{Q}$ intersected by $\pi_{\mathcal{P}}(s, t)$, let $p_i, q_i \in \partial Q_i$ be the corresponding entry and exit points of $\pi_{\mathcal{P}}(s, t)$. We obtain a new path π from $\pi_{\mathcal{P}}(s, t)$ by replacing its subpath $\pi_{\mathcal{P}}(p_i, q_i)$ with $\pi_{\mathcal{Q}}(p_i, q_i)$, for each pair $\{p_i, q_i\}$. Clearly, $\pi \subseteq \mathcal{F}(\mathcal{Q})$. Furthermore, for each pair $\{p_i, q_i\}$, applying Lemma 3.1 on P_i, Q_i with $p_i, q_i \in \partial Q_i$ and $r = \varepsilon^{3/2}d/c_k$, we obtain

$$d_{\mathcal{Q}}(p_i, q_i) \leq (1 + \varepsilon/2)d_{\mathcal{P}}(p_i, q_i) + \varepsilon d/8k,$$

provided c is a sufficiently large constant. Hence

$$|\pi| \leq (1 + \varepsilon/2)d_{\mathcal{P}}(p_i, q_i) + \varepsilon d/8,$$

implying (C2) as desired.

To prove (C3), observe that if $d_{\mathcal{Q}}(s, t) \leq 2d$, then $\pi_{\mathcal{Q}}(s, t)$ is contained in the interior of C_{4d} and hence, by construction, $\pi_{\mathcal{Q}}(s, t) \subseteq \mathcal{F}(\mathcal{P})$ implying $d_{\mathcal{Q}}(s, t) \geq d_{\mathcal{P}}(s, t)$ as desired.

LEMMA 3.2. *Let \mathcal{P} be a set of k pairwise-disjoint convex polytopes with a total of n faces. Suppose the Dobkin-Kirkpatrick hierarchy of each polytope in \mathcal{P} as well as their pairwise separating supporting planes have been computed in $O(n + k^2 \log^2 n)$ time. Let $s, t \in \mathcal{F}(\mathcal{P})$, and let $d \geq \|st\|$ be a real value. A set \mathcal{Q} of at most k pairwise-disjoint convex polytopes, with a total complexity $O(k^2/\varepsilon^{3/2})$, can be computed in $O((k^2/\varepsilon^{3/2}) \log(k/\varepsilon))$ time such that (C1)–(C3) hold.*

Computing ε -short paths. Let $d^* = d_{\mathcal{P}}(s, t)$. We start by computing a $2k$ -factor approximation \tilde{d} to d^* in $O(n)$ time using the algorithm of Hershberger and Suri [14]. We have $d^* \leq \tilde{d} \leq 2kd^*$. Set $m = \log(4k)$, and let $d_0 = \tilde{d}/2k$ and $d_i = 2^i \cdot d_0$, for $i = 1, \dots, m$. Note that $d_0 \leq d^*$ and $d_m \geq 2d^*$.

For each $i = 1, \dots, m$, we run the algorithm of Lemma 3.2 with $d = d_i$ and error parameter $\varepsilon/4$ to compute a set \mathcal{Q}_i . We then apply the “refinement step” of the Clarkson’s algorithm² on \mathcal{Q}_i , with parameter $\varepsilon/12$ and d_i , which computes a path π_i between s and t in $\mathcal{F}(\mathcal{Q}_i)$ such that $|\pi_i| \leq d_{\mathcal{Q}_i}(s, t) + \varepsilon d_i/24$. Combining this with Lemma 3.2, we have

$$(3.2) \quad |\pi_i| \leq (1 + \varepsilon/6)d^* + \varepsilon d_i/4.$$

Our algorithm terminates when $|\pi_i| \leq 2d_i$ and reports π_i as an ε -short path. Let j be the index such that $d^* \leq d_j \leq 2d^*$. When $i = j$, since $d^* \leq d_j$, from (3.2), we obtain $|\pi_j| \leq (1 + \varepsilon)d_j \leq 2d_j$. This implies that the algorithm terminates on or before iteration j . Let the algorithm terminate in some iteration $k \leq j$. Since $d_k \leq d_j \leq 2d^*$, from (3.2) and Lemma 3.2, it follows that $d^* \leq |\pi_k| \leq (1 + \varepsilon)d^*$. Hence the reported path π_k is indeed an ε -short path.

We spend $O(n + k^2 \log^2 n)$ time for precomputing the Dobkin-Kirkpatrick hierarchy of each polytope in \mathcal{P} and their pairwise separating supporting planes, as

²Clarkson’s algorithm is divided into two steps. The estimation step computes a 1-short distance d and the refinement step with parameters ε and d computes a path π whose length exceeds the shortest-path distance by at most $\varepsilon d/2$; as such, π is an ε -short path. See [10] for details.

well as computing the value of \tilde{d} . Then, in each iteration, computing \mathcal{Q}_i takes $O((k^2/\varepsilon^{3/2}) \log(k/\varepsilon))$ time (Lemma 3.2), and running the “refinement step” of Clarkson’s algorithm takes $O((k^4/\varepsilon^7) \log^2(k/\varepsilon))$ time. Since the total number of iterations is $O(\log k)$, the total running time is $O(n + (k^4/\varepsilon^7) \log^3(k/\varepsilon))$ (the $O(k^2 \log^2 n)$ term is absorbed). The output path has $O(k^2/\varepsilon^{3/2})$ edges, because each \mathcal{Q}_i is of total complexity $O(k^2/\varepsilon^{3/2})$. Hence, we conclude the following.

THEOREM 3.1. *Let \mathcal{P} be a set of k pairwise-disjoint convex polytopes in \mathbb{R}^3 of total complexity n . For any two points $s, t \in \mathcal{F}(\mathcal{P})$ and a parameter $0 < \varepsilon \leq 1$, an ε -short path between s and t which consists of $O(k^2/\varepsilon^{3/2})$ edges can be computed in $O(n + (k^4/\varepsilon^7) \log^3(k/\varepsilon))$ time.*

By combining the algorithms of Har-Peled [12] and Hershberger and Suri [14], it is possible to use the Dobkin-Kirkpatrick hierarchies of the polytopes in \mathcal{P} to compute the value of \tilde{d} in $O(k \log n)$ time. We then obtain the following result.

COROLLARY 3.1. *Let \mathcal{P} be a set of k pairwise-disjoint convex polytopes in \mathbb{R}^3 of total complexity n . Suppose the Dobkin-Kirkpatrick hierarchies of each polytope in \mathcal{P} as well as their pairwise separating supporting planes have been computed in $O(n + k^2 \log^2 n)$ time. For any $s, t \in \mathcal{F}(\mathcal{P})$ and for any $0 < \varepsilon \leq 1$, an ε -short path between s and t , which consists of $O(k^2/\varepsilon^{3/2})$ edges, can be computed in $O(k \log n + (k^4/\varepsilon^7) \log^3(k/\varepsilon))$ time.*

4 Approximate Shortest-Path Queries in \mathbb{R}^3

In this section we study the approximate shortest-path query problem amid \mathcal{P} for a fixed source $s \in \mathcal{F}(\mathcal{P})$ and a fixed parameter $0 < \varepsilon \leq 1$. The main result is a data structure whose size and query time are independent of n . In Sections 4.1 and 4.2, we prove a few key geometric lemmas that our data structure relies on. In Section 4.3, we present the details of the data structure.

4.1 Pseudoconvex subdivisions

For a convex polytope P , an ε -pseudoface of P is a region $F \subseteq \partial P$ such that for any $s, t \in F$, there exist outward normals u_s and u_t of P at s and t respectively such that $\angle u_s, u_t \leq \sqrt{\varepsilon}/2$. Note that an ε -pseudoface is not necessarily a connected region.

LEMMA 4.1. *The boundary of a convex polytope P in \mathbb{R}^3 (resp., \mathbb{R}^2) can be decomposed into a collection \mathcal{S} of $O(1/\varepsilon)$ (resp., $O(1/\sqrt{\varepsilon})$) ε -pseudofaces, each of which is the union of a subset of faces of P . The decomposition can be computed in $O(|P|)$ time.*

Proof. We only prove the lemma for \mathbb{R}^3 ; the case for \mathbb{R}^2 is simpler. We draw a uniform grid \mathcal{G} on \mathbb{S}^2 of longitudes

and latitudes so that the geodesic diameter of each grid cell is at most $\sqrt{\varepsilon}/2$. For a cell $\sigma \in \mathcal{G}$, let $S(\sigma)$ be the collection of faces of P whose outward normals fall inside σ . (To avoid ambiguity, we choose \mathcal{G} such that the outward normal of each face of P is contained in the interior of a unique cell of \mathcal{G} .) Let $F(\sigma) = \bigcup_{f \in S(\sigma)} f$. Clearly, $F(\sigma)$ is an ε -pseudoface of P . We conclude that $\mathcal{S} = \{F(\sigma) \mid \sigma \in \mathcal{G}\}$ is the desired collection of ε -pseudofaces of P , and clearly it can be computed in $O(|P|)$ time. \square

For a set \mathcal{P} of convex polytopes, we define an ε -pseudoconvex region in $\mathcal{F}(\mathcal{P})$ as a region $\sigma \subseteq \mathcal{F}(\mathcal{P})$ such that for any $s, t \in \sigma$, $d_{\mathcal{P}}(s, t) \leq (1 + \varepsilon)\|st\|$. Again, an ε -pseudoconvex region is not necessarily connected. We define an ε -pseudoconvex decomposition Ξ of $\mathcal{F}(\mathcal{P})$ as a decomposition of $\mathcal{F}(\mathcal{P})$ into ε -pseudoconvex regions.

A wedge W is the intersection of two halfspaces $h_1^+ \cap h_2^+$, and the angle of W is $\angle u_1, u_2$, where u_1, u_2 are the outward normals of h_1^+ and h_2^+ respectively. Let W be a wedge of angle at most $\sqrt{\varepsilon}/2$. A standard calculation shows that $d_W(s, t) \leq (1 + \varepsilon)\|st\|$ for any $s, t \in \mathcal{F}(W)$.

LEMMA 4.2. *Let P be a convex polytope in \mathbb{R}^3 . An ε -pseudoconvex decomposition $\Xi(P)$ of $\mathcal{F}(P)$ of size $O(1/\varepsilon)$ can be computed in $O(|P|)$ time.*

Proof. Let Q be the vertical projection of P onto the xy -plane, and let C be the (infinite) vertical prism $Q \times \mathbb{R}$. We produce an ε -pseudoconvex decomposition of $\mathcal{F}(P)$ by first constructing an ε -pseudoconvex decomposition Ξ' of size $O(1/\varepsilon)$ within $\mathcal{F}(P) \cap C = C \setminus \text{int } P$, and then constructing an ε -pseudoconvex decomposition Ξ'' of $\mathcal{F}(P) \setminus \text{int } C = \mathcal{F}(C)$ of size $O(1/\sqrt{\varepsilon})$.

The region $C \setminus \text{int } P$ consists of two components: the upper component C^+ extending infinitely in $(+z)$ -direction, and the lower component C^- extending infinitely in $(-z)$ -direction. For a face f of P in C^+ , let σ_f be the vertical prism with base f and extending infinitely in $(+z)$ -direction. Let \mathcal{S} be a set of $O(1/\varepsilon)$ ε -pseudofaces of P , by Lemma 4.1. For an ε -pseudoface $F \in \mathcal{S}$, let $\sigma_F = \bigcup_f \sigma_f$ where the union is taken over all faces of P that lie in $C^+ \cap F$. We claim that σ_F is an ε -pseudoconvex region. To see this, for any $s, t \in \sigma_F$, let s', t' be their projections onto P in $(-z)$ -direction. Since s', t' lie in the same ε -pseudoface F , there are supporting planes $h_{s'}$ and $h_{t'}$ of P at s' and t' with outer normals $u_{s'}$ and $u_{t'}$, respectively, such that $\angle u_{s'}, u_{t'} \leq \sqrt{\varepsilon}/2$. Let W be the wedge $h_{s'}^+ \cap h_{t'}^+$, where $h_{s'}^+$ (resp., $h_{t'}^+$) denotes the halfspace bounded by $h_{s'}$ (resp., $h_{t'}$) and containing P . Observe that $P \subseteq W$ and $s, t \in \mathcal{F}(W)$. Since the angle of W is at most $\sqrt{\varepsilon}/2$, we then have $d_P(s, t) \leq d_W(s, t) \leq (1 + \varepsilon)\|st\|$ as desired. It fol-

lows that $\{\sigma_F \mid F \in \mathcal{S}\}$ is an ε -pseudoconvex decomposition of C^+ . Symmetrically, we can construct an ε -pseudoconvex decomposition of C^- . They together form an ε -pseudoconvex decomposition Ξ' of $C \setminus \text{int } P$ of size $O(1/\varepsilon)$.

In the xy -plane, using a similar method, we can construct an ε -pseudoconvex decomposition $\Xi(Q)$ of $\mathcal{F}(Q)$ of size $O(1/\sqrt{\varepsilon})$; in particular, for any $\sigma \in \Xi(Q)$ and any $s, t \in \sigma$, there exists a wedge $W_{s,t}$ of angle at most $\sqrt{\varepsilon}/2$ in the xy -plane such that $Q \subseteq W_{s,t}$ and $s, t \in \mathcal{F}(W_{s,t})$. For any $s, t \in \sigma \times \mathbb{R}$ with $\sigma \in \Xi(Q)$, let s', t' be the projection of s, t onto σ . Then $W = W_{s',t'} \times \mathbb{R}$ is a wedge in \mathbb{R}^3 of angle at most $\sqrt{\varepsilon}/2$ such that $P \subseteq W$ and $s, t \in \mathcal{F}(W)$, implying $d_P(s, t) \leq d_W(s, t) \leq (1 + \varepsilon)\|st\|$. As such, $\sigma \times \mathbb{R}$ is an ε -pseudoconvex region, and the set $\Xi'' = \{\sigma \times \mathbb{R} \mid \sigma \in \Xi(Q)\}$ is an ε -pseudoconvex decomposition of $\mathcal{F}(C)$. We conclude that $\Xi(P) = \Xi' \cup \Xi''$ is an ε -pseudoconvex decomposition of $\mathcal{F}(P)$ of size $O(1/\varepsilon)$. \square

Using Lemma 4.2, we construct an ε -pseudoconvex decomposition of $\mathcal{F}(\mathcal{P})$ as follows. For $i = 1, \dots, k$, set $P'_i = \bigcap_{j \neq i} h_{i,j}^+$, where $h_{i,j}$ is a plane separating P_i and P_j and $h_{i,j}^+$ is the halfspace bounded by $h_{i,j}$ and containing P_i . Clearly, P'_i is a convex polytope of complexity $O(k)$ with $P_i \subseteq P'_i$, and $\mathcal{P}' = \{P'_1, P'_2, \dots, P'_k\}$ is a set of pairwise-disjoint convex polytopes. We decompose $\mathcal{F}(\mathcal{P}')$ into a set Ξ_0 of $O(k^3 \log k)$ tetrahedra as described in [1]; each tetrahedron is clearly an ε -pseudoconvex region. Next, for each polytope P_i , we obtain an ε -pseudoconvex decomposition $\Xi(P_i)$ of $\mathcal{F}(P_i)$ from Lemma 4.2, and clip each region $\sigma \in \Xi(P_i)$ with P'_i . Let $\Xi_i = \{\sigma \cap P'_i \mid \sigma \in \Xi(P_i)\}$ denote the resulting decomposition of $P'_i \setminus \text{int } P_i$. (We remark that, for our purpose, there is no need to represent each cell $\sigma \cap P'_i$ explicitly.) Each region $\sigma \cap P'_i \in \Xi_i$ is an ε -pseudoconvex region. This is because, for any pair of points $s, t \in \sigma \cap P'_i$, $\pi_{\mathcal{P}}(s, t) \subseteq P'_i \setminus \text{int } P_i$, implying $d_{\mathcal{P}}(s, t) = d_P(s, t) \leq (1 + \varepsilon)\|st\|$. Setting $\Xi(\mathcal{P}) = \Xi_0 \cup \Xi_1 \cup \dots \cup \Xi_k$, we obtain the following.

LEMMA 4.3. *$\Xi(\mathcal{P})$ is an ε -pseudoconvex decomposition of $\mathcal{F}(\mathcal{P})$ of size $O(k^3 \log k + k/\varepsilon)$.*

4.2 Critical distance values

Let s be a fixed source in $\mathcal{F}(\mathcal{P})$. For a region $U \subseteq \mathcal{F}(\mathcal{P})$, we call a set of distance values $d_1 < \dots < d_m$ critical if for any $t \in U$, one of the following holds:

- (i) $d_{\mathcal{P}}(s, t) \leq (1 + \varepsilon)\|st\|$; or
- (ii) there exists an index i such that $d_i \leq d_{\mathcal{P}}(s, t) \leq d_{i+1} \leq 2d_i$.

Intuitively, the critical distance values of U are what we need to focus on when answering approximate shortest-path queries for points in U ; Euclidean distances are good approximations for other cases. We next describe an algorithm to compute a set $\Sigma(P_i)$ of $O((1/\varepsilon)\log(1/\varepsilon))$ critical distance values for the region $P'_i \setminus \text{int } P_i$, for each $i = 1, \dots, k$.

Let Ξ_i be an $(\varepsilon/4)$ -pseudoconvex decomposition of $P'_i \setminus \text{int } P_i$ of size $O(1/\varepsilon)$. We will compute a set Σ_σ of critical distance values for each $(\varepsilon/4)$ -pseudoconvex region $\sigma \in \Xi_i$ and then set $\Sigma(P_i) = \bigcup_{\sigma \in \Xi_i} \Sigma_\sigma$. The set Σ_σ is computed as follows. We first find the Euclidean nearest neighbor v_e of s in σ (i.e., $v_e = \arg \min_{p \in \Xi_i} \|sp\|$) using a method to be explained shortly. Using Corollary 3.1, we compute a value \tilde{r} such that $d_{\mathcal{P}}(s, v_e) \leq \tilde{r} \leq 2d_{\mathcal{P}}(s, v_e)$. We then set

$$\Sigma_\sigma = \{\tilde{r}/8, 2\tilde{r}/8, 2^2\tilde{r}/8, \dots, 2^{m+3}\tilde{r}/8\},$$

where $m = \lceil \log_2(4 + 4/\varepsilon) \rceil$.

To compute the Euclidean nearest neighbor v_e of s in σ , recall that $\sigma = \sigma' \cap P'_i$ for some $\sigma' \in \Xi(P_i)$, where $\Xi(P_i)$ is an $(\varepsilon/4)$ -pseudoconvex decomposition of $\mathcal{F}(P_i)$ from Lemma 4.2. Using the Dobkin-Kirkpatrick hierarchy of P'_i , one can compute the Euclidean nearest neighbor of s in $\sigma' \cap P'_i$ (i.e., v_e) in $O(|\sigma'| \log |P'_i|) = O(|\sigma'| \log k)$ time, where $|\sigma'|$ denotes the complexity of the cell σ' .

Once v_e is identified, the value \tilde{r} and thus Σ_σ can be computed in $O(k \log n + k^4 \log^3 k + |\sigma'| \log k)$ time using Corollary 3.1 (after preprocessing). Since $\sum_{\sigma' \in \Xi_i} |\sigma'| = O(|P_i|)$ and $|\Xi_i| = O(1/\varepsilon)$, $\Sigma(P_i)$ can then be computed in $O((k/\varepsilon) \log n + (k^4/\varepsilon) \log^3 k + |P_i| \log k)$ time.

It remains to prove that Σ_σ is indeed a set of critical distance values for the region σ . Let $v_g \in \sigma$ be the geodesic nearest neighbor of s in σ (i.e., $v_g = \arg \min_{p \in \sigma} d_{\mathcal{P}}(s, p)$), and let $r = d_{\mathcal{P}}(s, v_g)$. We claim that $r \leq \tilde{r} \leq 8r$. Indeed, since σ is an $(\varepsilon/4)$ -pseudoconvex region and $v_e, v_g \in \sigma$, we have

$$\begin{aligned} d_{\mathcal{P}}(v_e, v_g) &\leq (1 + \varepsilon/4) \|v_e v_g\| \\ &\leq (1 + \varepsilon/4) (\|s v_e\| + \|s v_g\|) \\ &\leq 2(1 + \varepsilon/4) d_{\mathcal{P}}(s, v_g). \end{aligned}$$

It follows that

$$d_{\mathcal{P}}(s, v_e) \leq d_{\mathcal{P}}(s, v_g) + d_{\mathcal{P}}(v_g, v_e) \leq 4d_{\mathcal{P}}(s, v_g),$$

implying $r \leq \tilde{r} \leq 8r$.

Next, we partition the region σ into two subsets:

$$\begin{aligned} \sigma_1 &= \{t \in \sigma \mid \|v_g t\| \geq r(1 + 4/\varepsilon)\}, \\ \sigma_2 &= \{t \in \sigma \mid \|v_g t\| \leq r(1 + 4/\varepsilon)\}. \end{aligned}$$

For any point $t \in \sigma_1$, we have $r \leq \|v_g t\|/(1 + 4/\varepsilon)$. Hence,

$$\begin{aligned} \|st\| &\geq \|v_g t\| - \|s v_g\| \geq \|v_g t\| - r \\ &\geq \|v_g t\| (1 - 1/(1 + 4/\varepsilon)) = \|v_g t\|/(1 + \varepsilon/4). \end{aligned}$$

Furthermore,

$$\begin{aligned} d_{\mathcal{P}}(s, t) &\leq d_{\mathcal{P}}(s, v_g) + d_{\mathcal{P}}(v_g, t) \\ &\leq r + (1 + \varepsilon/4) \|v_g t\| \\ &\leq (1 + \varepsilon) \|v_g t\|/(1 + \varepsilon/4). \end{aligned}$$

Therefore, $d_{\mathcal{P}}(s, t) \leq (1 + \varepsilon) \|st\|$. On the other hand, for any point $t \in \sigma_2$,

$$\begin{aligned} r &\leq d_{\mathcal{P}}(s, t) \leq r + (1 + \varepsilon/4) \|v_g t\| \\ &\leq (1 + (1 + \varepsilon/4)(1 + 4/\varepsilon)) r \leq (4 + 4/\varepsilon) r. \end{aligned}$$

Together with $r \leq \tilde{r} \leq 8r$, we obtain

$$\tilde{r}/8 \leq d_{\mathcal{P}}(s, t) \leq (4 + 4/\varepsilon) \tilde{r}.$$

Therefore, there exists an index $0 \leq i < m + 3$ such that

$$2^i \tilde{r}/8 \leq d_{\mathcal{P}}(s, t) \leq 2^{i+1} \tilde{r}/8,$$

as desired.

LEMMA 4.4. *A set $\Sigma(P_i)$ of $O((1/\varepsilon)\log(1/\varepsilon))$ critical distance values for the region $P'_i \setminus \text{int } P_i$ can be computed in $O((k/\varepsilon) \log n + (k^4/\varepsilon) \log^3 k + |P_i| \log k)$ time.*

4.3 Shortest path queries

We are now ready to describe the data structure for approximate shortest-path queries with respect to a fixed source $s \in \mathcal{F}(\mathcal{P})$.

The structure. Sections 4.1 and 4.2 imply that we can compute:

- (i) a set $\mathcal{P}' = \{P'_1, \dots, P'_k\}$ of k pairwise-disjoint convex polytopes such that $P_i \subseteq P'_i$ and $|P'_i| = O(k)$,
- (ii) a set $\Sigma(P_i)$ of $m = O((1/\varepsilon)\log(1/\varepsilon))$ critical distances for the region $P'_i \setminus P_i$, and
- (iii) a decomposition Ξ_0 of $\mathcal{F}(\mathcal{P}')$ into $O(k^3 \log k)$ tetrahedra.

The overall structure consists of three components. First, for each tetrahedron $\Delta \in \Xi_0$, we construct a data structure $\mathbb{D}(\Delta)$ of Har-Peled [13] of size $O(1/\varepsilon^5)$ in $O(1/\varepsilon^5)$ time so that for any point $t \in \Delta$, an ε -short distance between s and t amid \mathcal{P} can be reported in $O(\log(1/\varepsilon))$ time. The data structure of [13] in a

tetrahedron Δ is constructed by sprinkling a set X_Δ of $O((1/\varepsilon^2)\log(1/\varepsilon))$ weighted points in Δ and then computing a weighted Voronoi diagram $\text{Vor}(X_\Delta)$ of X_Δ , and preprocessing $\text{Vor}(X_\Delta)$ into a point-location structure; the weight w_p of each point $p \in X_\Delta$ is an ε -short distance between s and p amid \mathcal{P} .

Next, for each $P \in \mathcal{P}$ and each $d \in \Sigma(P)$, we construct a data structure $\mathbb{D}(P, d)$ as follows. Set $r = \varepsilon d/64$. Let C_{4d} be a cube of side-length $4d$ centered at s . We compute an inner r -approximation I of $P \cap C_{4d}$ ($I \subseteq P \cap C_{4d} \subseteq I_r$ and $|I| = O(1/\varepsilon)$) and an outer r -approximation O of $P \cap C_{4d}$ ($|O| = O(1/\varepsilon)$). We also compute the Dobkin-Kirkpatrick hierarchies of I and O . We decompose the region $(P' \cap C_{4d}) \setminus \text{int } O$ into $O(|P' \cap C_{4d}| + |O|) = O(k + 1/\varepsilon)$ tetrahedra using an algorithm in [7]. Let $\Xi(P, d)$ be this decomposition. Note that each $\Delta \in \Xi(P, d)$ lies in $\mathcal{F}(\mathcal{P})$. We process $\Xi(P, d)$ into a point-location structure of size $O((k + 1/\varepsilon)\log^2(k/\varepsilon))$ with query time $O(\log^2(k/\varepsilon))$ [21]. For each $\Delta \in \Xi(P, d)$, we can construct the aforementioned data structure of Har-Peled [13] so that for any point $t \in \Delta$, an $(\varepsilon/4)$ -short distance between s and t amid \mathcal{P} can be reported in $O(\log(1/\varepsilon))$ time.

Using Lemma 4.4, we can compute the critical distance values in $\Sigma(P_1) \cup \dots \cup \Sigma(P_k)$ in $O(n \log k + (k^5/\varepsilon)\log^3 k)$ time. The weight of each sprinkled point in Har-Peled's data structure can be computed in $O(k \log n + (k^4/\varepsilon^7)\log^3(k/\varepsilon))$ time using Corollary 3.1 (after preprocessing). Therefore, the total time for constructing the entire data structure, summed over all $P \in \mathcal{P}$ and $d \in \Sigma(P)$, is $O(n \log k + k^7 \text{poly}(\log k, 1/\varepsilon))$. The size of the entire data structure is $O(k^3 \text{poly}(\log k, 1/\varepsilon))$.

Finally, we preprocess Ξ_0 and \mathcal{P}' into a point-location data structure $\mathbb{D}(\Xi_0, \mathcal{P}')$ of size $O(k^3 \log^3 k)$ so that given a point $t \in \mathcal{F}(\mathcal{P})$, one of $\Delta \in \Xi_0$ or $P' \in \mathcal{P}'$ that contains t can be located in $O(\log^2 k)$ time [21].

Query algorithm. Let $t \in \mathcal{F}(\mathcal{P})$ be a query point. We compute a value $\theta(t)$ such that

$$d_{\mathcal{P}}(s, t) \leq \theta(t) \leq (1 + \varepsilon)d_{\mathcal{P}}(s, t).$$

Using the point location data structure $\mathbb{D}(\Xi_0, \mathcal{P}')$, we compute the tetrahedron $\Delta \in \Xi_0$ or the polytope $P' \in \mathcal{P}'$ that contains t . If $t \in \Delta \in \Xi_0$, we compute the (weighted) nearest neighbor $q \in X_\Delta$ of t and return the value $w_q + \|qt\|$ as $\theta(t)$.

Next suppose $t \in P'$. Let $P \in \mathcal{P}$ be the polytope such that $P \subseteq P'$. By a binary search on the values in $\Sigma(P)$, we find the smallest value d_i such that $d_i > \|st\|$. If no such d_i exists, we simply return $(1 + \varepsilon)\|st\|$ as $\theta(t)$. Otherwise, We query the data structure $\mathbb{D}(P, d_i)$ as follows. Using the Dobkin-Kirkpatrick heirarchy of O_i , we determine in $O(\log(1/\varepsilon))$ time whether $t \in \text{int } O_i$. Suppose $t \notin \text{int } O_i$, i.e., $t \in P' \setminus \text{int } O_i$, we find

in $O(\log^2 k)$ time the tetrahedron $\Delta \in \Xi(P, d_i)$ that contains t , and then query Har-Peled's structure built on Δ . Let $\theta_i(t)$ be the value returned by this procedure. We return $\theta(t) = \theta_i(t)$.

Finally, if $t \in O_i$, we first compute the nearest neighbor $t' \in I_i$ of t using the Dobkin-Kirkpatrick heirarchy, and then compute the intersection point t'' of ∂O_i with the ray $t't$, using the Dobkin-Kirkpatrick heirarchy built on O_i . This step takes $O(\log(1/\varepsilon))$ time. Note that $tt'' \subset \mathcal{F}(\mathcal{P})$. (If we simply take t'' as the nearest neighbor of t on ∂O_i , then the segment tt'' may intersect $\text{int } P$; that is why we had to use I_i to find the appropriate t'' .) We compute $\theta_i(t'')$ as described above. We set $\theta_i(t) = \theta_i(t'') + \|tt''\|$. If $d_i \leq 4\|st\|$ or $\theta_i(t) - \varepsilon d_i/8 \geq (1 + 2\varepsilon)\|st\|$, we return $\theta_i(t)$ as the value of $\theta(t)$, else return $(1 + \varepsilon)\|st\|$ as the value of $\theta(t)$.

This completes the description of the query procedure. The total query time is $O(\log^2(k/\varepsilon))$.

LEMMA 4.5. For any $t \in \mathcal{F}(\mathcal{P})$,

$$(4.3) \quad d_{\mathcal{P}}(s, t) \leq \theta(t) \leq (1 + \varepsilon)d_{\mathcal{P}}(s, t).$$

Proof. If $t \in \Xi_0$, then (4.3) follows from the correctness of Har-Peled's structure, so assume $t \in P'$. Let $d_i \in \Sigma(P)$ be the smallest value such that $d_i \geq \|st\|$. If no such d_i exists, then by the definition of critical distance values, $\theta(t) = (1 + \varepsilon)\|st\|$ satisfies (4.3). So assume that d_i exists. We have $t \in C_{4d_i} \cap P'$. If $t \in P' \setminus \text{int } O_i$, then a tetrahedron $\Delta \in \Xi(P, d_i)$ contains t , and Har-Peled's data structure on Δ returns a value that satisfies (4.3). Hence assume that $t \in P' \cap O'$.

Clearly, $\theta_i(t) \geq d_{\mathcal{P}}(s, t)$. Furthermore, Let t'' be the point as defined in the above query procedure. Then, by construction, $\|tt''\| \leq 2\varepsilon d_i/64 = \varepsilon d_i/32$. As such,

$$(4.4) \quad \begin{aligned} \theta_i(t) &\leq (1 + \varepsilon/4)d_{\mathcal{P}}(s, t'') + \varepsilon d_i/32, \\ &\leq (1 + \varepsilon/4)(d_{\mathcal{P}}(s, t) + \varepsilon d_i/32) + \varepsilon d_i/32, \\ &\leq (1 + \varepsilon/4)d_{\mathcal{P}}(s, t) + \varepsilon d_i/8. \end{aligned}$$

Now, if $d_i \leq 4\|st\|$, then

$$\begin{aligned} \theta(t) = \theta_i(t) &\leq (1 + \varepsilon/4)d_{\mathcal{P}}(s, t) + \varepsilon\|st\|/2, \\ &\leq (1 + \varepsilon)d_{\mathcal{P}}(s, t), \end{aligned}$$

as desired. So assume that $d_i > 4\|st\|$. There are two cases to consider:

(i) $d_{\mathcal{P}}(s, t) \leq (1 + \varepsilon)\|st\|$. By (4.4),

$$\begin{aligned} \theta_i(t) - \varepsilon d_i/8 &\leq (1 + \varepsilon/4)d_{\mathcal{P}}(s, t), \\ &< (1 + 2\varepsilon)\|st\|. \end{aligned}$$

Hence, the query procedure sets $\theta(t) = (1 + \varepsilon)\|st\|$, which satisfies (4.3).

(ii) $d_{\mathcal{P}}(s, t) > (1 + \varepsilon)\|st\|$. By the definition of critical distance values, there is some value $d_j \in \Sigma(P)$ such that

$$d_{j-1} \leq d_{\mathcal{P}}(s, t) \leq d_j \leq 2d_{j-1}.$$

Since $d_{i-1} \leq \|st\| < 4\|st\| < d_i$, we can deduce that $4\|st\| < d_i \leq d_{j-1} \leq d_{\mathcal{P}}(s, t)$. Thus, we have

$$\begin{aligned} \theta_i(t) - \varepsilon d_i/8 &> d_{\mathcal{P}}(s, t) - \varepsilon d_i/8, \\ &\geq d_{\mathcal{P}}(s, t)(1 - \varepsilon/8), \\ &> 4\|st\|(1 - \varepsilon/8), \\ &\geq (1 + 2\varepsilon)\|st\|. \end{aligned}$$

Hence the query procedure sets $\theta(t)$ to $\theta_i(t)$, and

$$\begin{aligned} \theta(t) = \theta_i(t) &\leq (1 + \varepsilon/4)d_{\mathcal{P}}(s, t) + \varepsilon d_i/8, \\ &\leq (1 + \varepsilon)d_{\mathcal{P}}(s, t). \end{aligned}$$

This completes the proof. \square

THEOREM 4.1. *Let \mathcal{P} be a set of k convex polytopes of total complexity n in \mathbb{R}^3 , and let s be a fixed source in $\mathcal{F}(\mathcal{P})$, and let $0 < \varepsilon \leq 1$. A data structure of size $O(k^3 \text{poly}(\log k, 1/\varepsilon))$ can be constructed in $O(n \log k + k^7 \text{poly}(\log k, 1/\varepsilon))$ time such that for any query point $t \in \mathcal{F}(\mathcal{P})$, an ε -short distance between s and t can be reported in $O(\log^2(k/\varepsilon))$ time.*

5 Conclusion

In this paper, we obtain algorithms and data structures for the approximate Euclidean shortest-path problem amid convex obstacles. Given a set \mathcal{P} of pairwise disjoint convex obstacles, we show how to quickly construct a sketch \mathcal{Q} of \mathcal{P} whose complexity is independent of the complexity of \mathcal{P} and then use \mathcal{Q} to compute approximate shortest paths or answer approximate shortest-path queries. We conclude by mentioning a related problem. Is there a pseudoconvex decomposition of $\mathcal{F}(\mathcal{P})$ of size $O(k^2)$? This problem is closely related to whether there is a binary space partition of size $O(k^2)$ of a set of k convex objects in \mathbb{R}^3 .

References

[1] P. K. Agarwal, B. Aronov, and S. Suri. Stabbing triangulations by lines in 3d. In *Proc. 11th Annu. Sympos. Comput. Geom.*, pages 267–276, 1995.
 [2] P. K. Agarwal, S. Har-Peled, and M. Karia. Computing approximate shortest paths on convex polytopes. In *Algorithmica*, 33:227–242, 2002.
 [3] P. K. Agarwal, S. Har-peled, M. Sharir, and K. R. Varadarajan. Approximating shortest paths on a convex polytope in three dimensions. *J. ACM*, 44:567–584, 1997.

[4] Lyudmil Aleksandrov, Anil Maheshwari, and Jörg-Rüdiger Sack. Determining approximate shortest paths on weighted polyhedral surfaces. *J. ACM*, 52(1):25–53, 2005.
 [5] T. Asano, D. G. Kirkpatrick, and C. Yap. Pseudo approximation algorithms with applications to optimal motion planning. *Discrete Comput. Geom.*, 31:139–171, 2004.
 [6] J. Canny and J. Reif. New lower bound techniques for robot motion planning problems. In *Proc. 28th IEEE Sympos. Foundat. Comput. Sci.*, pages 49–60, 1987.
 [7] B. Chazelle and N. Shouraboura. Bounds on the size of tetrahedralizations. *Discrete Comput. Geom.*, 14:429–444, 1995.
 [8] J. Chen and Y. Han. Shortest paths on a polyhedron, Part I: Computing shortest paths. *Internat. J. Comput. Geom. Appl.*, 6:127–144, 1996.
 [9] J. Choi, J. Sellen, and C.-K. Yap. Approximate Euclidean shortest paths in 3-space. *Internat. J. Comput. Geom. Appl.*, 7:271–295, 1997.
 [10] K. L. Clarkson. Approximation algorithms for shortest path motion planning. In *Proc. 19th Annu. ACM Sympos. Theory Comput.*, pages 56–65, 1987.
 [11] K. Dudley. Metric entropy of some classes of sets with differentiable boundaries. *J. Approx. Theory*, 10:227–236, 1974.
 [12] S. Har-Peled. Approximate shortest-path and geodesic diameter on convex polytopes in three dimensions. *Discrete Comput. Geom.*, 21:216–231, 1999.
 [13] S. Har-Peled. Constructing approximate shortest path maps in three dimensions. *SIAM J. Comput.*, 28:1182–1197, 1999.
 [14] J. Hershberger and S. Suri. Practical methods for approximating shortest path on a convex polytope in \mathbb{R}^3 . *Comput. Geom. Theory Appl.*, 10:31–46, 1998.
 [15] J. Hershberger and S. Suri. An optimal algorithm for Euclidean shortest paths in the plane. *SIAM J. Comput.*, 28:2215–2256, 1999.
 [16] J. Mitchell. Shortest paths and networks. In Jacob E. Goodman and Joseph O’Rourke, editors, *Handbook of Discrete and Computational Geometry*, pages 755–778. CRC Press, Inc., Boca Raton, FL, USA, 1997.
 [17] J. Mitchell and M. Sharir. New results on shortest paths in three dimensions. In *Proc. 20th Annu. Sympos. Comput. Geom.*, pages 124–133, 2004.
 [18] J.S.B. Mitchell, D. Mount, and C. Papadimitriou. The discrete geodesic problem. *SIAM J. Comput.*, 16:647–668, 1987.
 [19] C. Papadimitriou. An algorithm for shortest path motion planning in three dimension. *Inform. Process. Lett.*, 20:259–268, 1985.
 [20] A. Pogorelov. *Extrinsic Geometry of Convex Surfaces*, volume 35 of *Transactions of Mathematical Monographs*. American Mathematical Society, Providence, RI, 1973.
 [21] F. Preparata and R. Tamassia. Efficient point location in a convex spatial cell-complex. *SIAM J. Comput.*, 21:267–280, 1992.

- [22] J. Reif and J. Storer. A single-exponential upper bound for finding shortest paths in three dimensions. *J. ACM*, 41:1013–1019, 1994.
- [23] Y. Schreiber and M. Sharir. An optimal algorithm for shortest paths on a convex polytope in three dimensions. *Discrete Comput. Geom.*, 39:500–579, 2008.
- [24] M. Sharir. On shortest paths amidst convex polyhedra. *SIAM J. Comput.*, 16:561–572, 1987.
- [25] M. Sharir and A. Schorr. On shortest paths in polyhedral spaces. *SIAM J. Comput.*, 15:193–215, 1986.