Discrete Comput Geom (2013) 50:306–329
DOI 10.1007/s00454-013-9527-8

The Geodesic Diameter of Polygonal Domains

Sang Won Bae · **Matias Korman** · **Yoshio Okamoto**

Received: 10 July 2011 / Revised: 5 November 2012 / Accepted: 24 June 2013 /
Published online: 26 July 2013

Abstract This paper studies the geodesic diameter of polygonal domains having h holes and n corners. For simple polygons (i.e., $h = 0$), the geodesic diameter is determined by a pair of corners of a given polygon and can be computed in linear time, as shown by Hershberger and Suri. For general polygonal domains with $h \geq 1$, however, no algorithm for computing the geodesic diameter was known prior to this paper. In this paper, we present the first algorithms that compute the geodesic diameter of a given polygonal domain in worst-case time $O(n^{7.73})$ or $O(n^7(\log n + h))$. The main difficulty unlike the simple polygon case relies on the following observation revealed in this paper: two interior points can determine the geodesic diameter and in that case there exist at least five distinct shortest paths between the two.

Keywords Polygonal domain · Shortest path · Geodesic diameter · Exact algorithm · Convex function · Lower envelope

A preliminary version of this paper was presented at the 18th Annual European Symposium on Algorithms (ESA 2010).

S. W. Bae (✉)
Department of Computer Science, Kyonggi University, Suwon, Korea
e-mail: swbae@kgu.ac.kr

M. Korman
Department of Matemàtica Aplicada II, Universitat Politècnica de Catalunya, Barcelona, Spain
e-mail: matias.korman@upc.edu

Y. Okamoto
Department of Communication Engineering and Informatics, University of Electro-Communications, Chofu, Tokyo, Japan
e-mail: okamotoy@uec.ac.jp

1 Introduction

A *polygonal domain* $\mathcal{P}$ with h holes and n corners is a connected and closed subset of $\mathbb{R}^2$ having h holes whose boundary consists of $h+1$ simple closed polygonal chains of n total line segments. Given a polygonal domain $\mathcal{P}$, the geodesic distance $\mathrm{d}(p,q)$ between two points p and q of $\mathcal{P}$ is defined as the length of a shortest path that connects p and q and stays within $\mathcal{P}$.

This paper addresses the geodesic diameter problem in polygonal domains having one or more holes. The geodesic diameter $\mathrm{diam}(\mathcal{P})$ of domain $\mathcal{P}$ is defined as the largest possible geodesic distance between any two points of $\mathcal{P}$, that is, $\mathrm{diam}(\mathcal{P}) = \max_{s,t\in\mathcal{P}} \mathrm{d}(s,t)$.

For simple polygons (i.e., domains with no hole), the geodesic diameter has been extensively studied. Chazelle [7] provided the first $O(n^2)$-time algorithm computing the geodesic diameter of a simple polygon. Afterwards, Suri [20] presented an $O(n\log n)$-time algorithm that solves the all-geodesic-farthest neighbors problem, computing the farthest neighbor of every corner and thus finding the geodesic diameter. At last, Hershberger and Suri [12] showed that the diameter can be computed in linear time using fast matrix search techniques.

On the other hand, the geodesic diameter of a domain having one or more holes is less understood. Mitchell [16] has posed an open problem asking an algorithm for computing the geodesic diameter of a polygonal domain. However, even for the corner-to-corner diameter $\max_{u,v\in V} \mathrm{d}(u,v)$, where V denotes the set of corners of $\mathcal{P}$, we know nothing better than a brute-force algorithm that takes $O(n^2\log n)$ time, checking all the geodesic distances between every pair of corners.[1] Prior to our results, there was no known algorithm for computing the geodesic diameter in domains with holes. We should also mention that Koivisto and Polishchuk [14] had claimed an improved algorithm after a preliminary report of our work [6], but it was shown to be a failed trial through conversations with the authors.[2]

This fairly wide gap between simple polygons and polygonal domains with holes is seemingly due to the uniqueness of the shortest path between any two points. When a domain $\mathcal{P}$ has no hole, it is well known that there is a unique shortest path between any two points [10]. Using this uniqueness, one can show that the diameter $\mathrm{diam}(\mathcal{P})$ is realized by a pair of corners [12,20]. For general polygonal domains, however, this is not the case. In this paper, we exhibit several examples where the diameters are realized by non-corner points on $\partial\mathcal{P}$ or even by interior points of $\mathcal{P}$ (see Fig. 1). Such examples were constructed based on the multiplicity of shortest paths and, to our best knowledge, never known prior to this work. This observation also shows an immediate difficulty in devising any exhaustive algorithm since one sees no intuitive discretization of the search space.

The status of the geodesic center problem is also similar. A point in $\mathcal{P}$ is defined as a *geodesic center* if it minimizes the maximum geodesic distance from it to any other point of $\mathcal{P}$. Asano and Toussaint [3] introduced the first $O(n^4\log n)$-time algorithm

[1] Personal communication with Joseph S. B. Mitchell.

[2] Personal communication with Valentin Polishchuk.

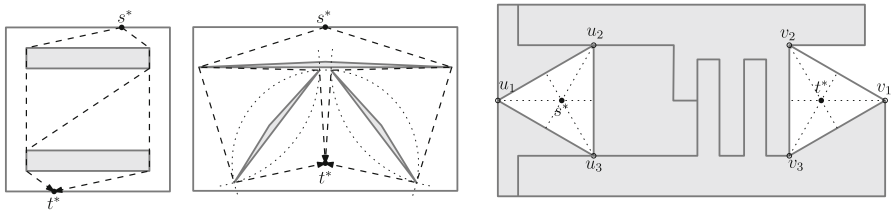

Fig. 1 Three polygonal domains where the geodesic diameter is determined by a pair (s^*, t^*) of non-corner points; *gray-shaded regions* depict the interior of the holes and *dark gray segments* depict the boundary $\partial\mathcal{P}$. Recall that $\mathcal{P}$, as a set, contains its boundary $\partial\mathcal{P}$. **a** Both s^* and t^* lie on $\partial\mathcal{P}$. There are three shortest paths between s^* and t^*. In this domain, there are two (symmetric) diametral pairs (only one is depicted for clarity). **b** $s^* \in \partial\mathcal{P}$ and $t^* \in \operatorname{int}\mathcal{P}$. Three triangular holes are placed in a symmetric way, obtaining four shortest paths between s^* and t^*. **c** Both s^* and t^* lie in the interior $\operatorname{int}\mathcal{P}$. Here, the five holes are packed like jigsaw puzzle pieces, forming narrow corridors (*dark gray paths*) and two empty, regular *triangles*. Observe that $\mathrm{d}(u_1, v_1) = \mathrm{d}(u_1, v_2) = \mathrm{d}(u_2, v_2) = \mathrm{d}(u_2, v_3) = \mathrm{d}(u_3, v_3) = \mathrm{d}(u_3, v_1)$. The points s^* and t^* lie at the centers of the triangles formed by the u_i and the v_i, respectively. There are six shortest paths between s^* and t^*.

for computing the geodesic center of a simple polygon (i.e., when $h = 0$), and Pollack et al. [19] improved it to $O(n \log n)$ time. As with the diameter problem, there is no known algorithm for domains with holes. See O'Rourke and Suri [18] and Mitchell [16] for more references on the geodesic diameter/center problem.

Since the geodesic diameter/center of a simple polygon is determined by its corners, one can exploit the *geodesic farthest-site Voronoi diagram* of the set V of corners to compute the diameter/center, which can be built in $O(n \log n)$ time [2]. Recently, Bae and Chwa [4] presented an $O(nk \log^3(n + k))$-time algorithm for computing the geodesic farthest-site Voronoi diagram of k sites in polygonal domains with holes. This result can be used to compute the geodesic diameter $\max_{p,q\in S} \mathrm{d}(p, q)$ of a *finite* set S of points in $\mathcal{P}$. However, this approach cannot be directly used for computing $\operatorname{diam}(\mathcal{P})$ without any characterization of the diameter. Moreover, when $S = V$, this approach is no better than the brute-force $O(n^2 \log n)$-time algorithm for computing the corner-to-corner diameter $\max_{u,v\in V} \mathrm{d}(u, v)$.

In this paper, we present the first algorithms that compute the geodesic diameter of a given polygonal domain in $O(n^{7.73})$ or $O(n^7(\log n + h))$ time in the worst case. Our new geometric results underlying the algorithms show that the existence of any diametral pair consisting of non-corner points implies multiple shortest paths between the pair; among other results, we show that *if (s, t) is a diametral pair and both s and t lie in the interior of $\mathcal{P}$, then there are at least five shortest paths between s and t.*

Some analogies between polygonal domains and convex polytopes in $\mathbb{R}^3$ can be seen. O'Rourke and Schevon [17] proved that if the geodesic diameter on a convex 3-polytope is realized by two non-corner points, then at least five shortest paths exist between the two; see also Zalgaller [21] for simpler arguments. Based on this observation, they presented an $O(n^{14} \log n)$-time algorithm for computing the geodesic diameter on a convex 3-polytope. Afterwards, the time bound was improved to $O(n^8 \log n)$ by Agarwal et al. [1] and recently to $O(n^7 \log n)$ by Cook IV and Wenk [9].

The rest of the paper is organized as follows: after introducing preliminary definitions and concepts in Sect. 2, we investigate local maxima of the lower envelope of convex functions in Sect. 3, resulting in Theorem 1. Section 4 extensively exploits the intermediate result to show lower bounds on the number of shortest paths between a diametral pair for every possible case, and then Sect. 5 describes our algorithms for the geodesic diameter. We finally conclude the paper with a summary, some remarks, and open issues in Sect. 6. Also, we exhibit several examples that cover all possible combinatorial cases in Appendix 1.

2 Preliminaries

Throughout the paper, we frequently use several topological concepts such as open and closed subsets, neighborhoods, and the boundary ∂A and the interior $\operatorname{int} A$ of a set A; unless stated otherwise, all of them are supposed to be derived with respect to the standard topology on $\mathbb{R}^d$ with the Euclidean norm $\|\cdot\|$ for fixed $d \geq 1$. We also denote the straight line segment joining two points a, b by $\overline{ab}$.

A *polygonal domain* $\mathcal{P}$ with h holes and n corners[3] is a connected and closed subset of $\mathbb{R}^2$ with h holes whose boundary $\partial\mathcal{P}$ consists of $h+1$ simple closed polygonal chains of n total line segments. The boundary $\partial\mathcal{P}$ of a polygonal domain $\mathcal{P}$ is regarded as a series of *obstacles* so that any feasible path in $\mathcal{P}$ is not allowed to cross $\partial\mathcal{P}$. The *geodesic distance* $\mathrm{d}(p, q)$ between any two points p, q in a polygonal domain $\mathcal{P}$ is defined as the length of a shortest feasible path between p and q, where the *length* of a path is the sum of the Euclidean lengths of its segments. It is well known from earlier work [15] that there always exists a shortest feasible path between any two points $p, q \in \mathcal{P}$, and thus the geodesic distance function $\mathrm{d}(\cdot,\cdot)$ is well defined. The geodesic diameter $\mathrm{diam}(\mathcal{P})$ of a polygonal domain $\mathcal{P}$ is defined as the largest geodesic distance between any two points of $\mathcal{P}$, that is,

$$\mathrm{diam}(\mathcal{P}) = \max_{s,t\in\mathcal{P}} \mathrm{d}(s, t).$$

A pair (s, t) of points in $\mathcal{P}$ that realizes the geodesic diameter $\mathrm{diam}(\mathcal{P})$ is called a *diametral pair*.

Shortest Path Map. Let V be the set of all corners of $\mathcal{P}$ and $\pi(s, t)$ be a shortest path between $s \in \mathcal{P}$ and $t \in \mathcal{P}$. Such a path $\pi(s, t)$ is represented as a sequence $\pi(s, t) = (s, v_1, \ldots, v_k, t)$ for some $v_1, \ldots, v_k \in V$; that is, a polygonal chain through a sequence of corners [15]. Note that we can have $k = 0$ when $\mathrm{d}(s, t) = \|s - t\|$. If two paths (with possibly different endpoints) have the same sequence of corners $(v_1, \ldots, v_k)$, then they are said to have the same *combinatorial structure*.

The *shortest path map* $\mathsf{SPM}(s)$ for a fixed $s \in \mathcal{P}$ is a decomposition of $\mathcal{P}$ into cells such that every point in a common cell can be reached from s by shortest paths of the same combinatorial structure. Each cell $\sigma_s(v)$ of $\mathsf{SPM}(s)$ is associated with a corner $v \in V$ which is the last corner of $\pi(s, t)$ for any t in the cell $\sigma_s(v)$. We also define the cell $\sigma_s(s)$ as the set of points $t \in \mathcal{P}$ such that $\pi(s, t)$ passes through no corner of $\mathcal{P}$, so

[3] We reserve the term "vertex" for 0-dimensional faces of subdivisions of a certain space.

$\pi(s,t) = \overline{st}$. Each edge of $\mathsf{SPM}(s)$ either belongs to $\partial\mathcal{P}$ or is an arc on the boundary of two incident cells $\sigma_s(v_1)$ and $\sigma_s(v_2)$ determined by two corners $v_1, v_2 \in V \cup \{s\}$. Similarly, each vertex of $\mathsf{SPM}(s)$ is a vertex of $\mathcal{P}$ or is determined by at least three distinct corners $v_1, v_2, v_3 \in V \cup \{s\}$.

Note that, for fixed $s \in \mathcal{P}$, a point farthest from s lies at either (1) a vertex of $\mathsf{SPM}(s)$, (2) an intersection between the boundary $\partial\mathcal{P}$ and an edge of $\mathsf{SPM}(s)$, or (3) a corner in V. The shortest path map $\mathsf{SPM}(s)$ has $O(n)$ total number of cells, edges, and vertices and can be computed in $O(n \log n)$ time using $O(n \log n)$ working space [13]. For more details on shortest path maps, see [13,15,16].

Path-Length Function. If $\pi(s,t) \neq \overline{st}$, then there are two corners $u, v \in V$ such that u and v are the first and last corners along $\pi(s,t)$ from s to t, respectively. Here, the path $\pi(s,t)$ is formed as the union of $\overline{su}$, $\overline{vt}$ and a shortest path $\pi(u,v)$ from u to v. Note that u and v are not necessarily distinct. In order to realize such a path, we assert that s is visible from u and t is visible from v. That is, $s \in \mathsf{VR}(u)$ and $t \in \mathsf{VR}(v)$, where $\mathsf{VR}(p)$ for any $p \in \mathcal{P}$ is defined to be the set of all points $q \in \mathcal{P}$ such that $\overline{pq} \subset \mathcal{P}$, also called the *visibility region* of $p \in \mathcal{P}$.

We now define the *path-length function* $\mathrm{len}_{u,v} : \mathsf{VR}(u) \times \mathsf{VR}(v) \to \mathbb{R}$ for any fixed pair of corners $u, v \in V$ to be

$$\mathrm{len}_{u,v}(s,t) := \|s-u\| + \mathrm{d}(u,v) + \|v-t\|.$$

That is, $\mathrm{len}_{u,v}(s,t)$ represents the length of paths from s to t that have a common combinatorial structure; going straight from s to u, following a shortest path from u to v, and going straight to t. Also, unless $\mathrm{d}(s,t) = \|s-t\|$ (equivalently, $s \in \mathsf{VR}(t)$), the geodesic distance $\mathrm{d}(s,t)$ can be expressed as the pointwise minimum of some path-length functions:

$$\mathrm{d}(s,t) = \min_{u \in \mathsf{VR}(s),\ v \in \mathsf{VR}(t)} \mathrm{len}_{u,v}(s,t).$$

Consequently, we have two possibilities for a diametral pair (s^*, t^*); either we have $\mathrm{d}(s^*, t^*) = \|s^* - t^*\|$ or the pair (s^*, t^*) is a local maximum of the lower envelope of several path-length functions. In the following, we will mainly study the latter case, since the former can be easily handled.

3 Local Maxima of the Lower Envelope of Convex Functions

In this section, we give a property of the lower envelope of a family of convex functions which will afterwards be used in our geodesic diameter environment. We start with a basic observation on the intersection of hemispheres on a unit hypersphere in the d-dimensional space $\mathbb{R}^d$. For any fixed positive integer d, let $S^{d-1} := \{x \in \mathbb{R}^d \mid \|x\| = 1\}$ be the unit hypersphere in $\mathbb{R}^d$ centered at the origin. A *closed* (or *open*) *hemisphere* on S^{d-1} is defined to be the intersection of S^{d-1} and a closed (open, respectively) half-space of $\mathbb{R}^d$ bounded by a hyperplane that contains the origin.

We call a k-dimensional affine subspace of $\mathbb{R}^d$ a *k-flat*. Note that a hyperplane in $\mathbb{R}^d$ is a $(d-1)$-flat and a line in $\mathbb{R}^d$ is a 1-flat. Also, the intersection of S^{d-1} and a k-flat through the origin in $\mathbb{R}^d$ is called a *great $(k-1)$-sphere* on S^{d-1}. Note that a great 1-sphere is called a great circle and a great 0-sphere consists of two antipodal points.

Lemma 1 *For any two positive integers d and $m \le d$, a set of any m closed hemispheres on S^{d-1} has a nonempty common intersection. Moreover, if the intersection has an empty interior relative to S^{d-1}, then it includes a great $(d-m)$-sphere on S^{d-1}.*

Proof We only give a proof for the second statement, which implies the first. The case of $d = 1$ is trivial, so we assume $d > 1$. Let $H_1, \ldots, H_m$ be any m closed hemispheres on S^{d-1}, and h_i be the hyperplane through the origin in $\mathbb{R}^d$ such that H_i lies in a closed half-space supported by h_i. In this proof, we denote by $\widehat{H}_i$ the open hemisphere, defined to be $\widehat{H}_i = H_i \setminus h_i$. Also, let $\mathcal{H}_j := \bigcap_{1\le i\le j} H_i$ and $\widehat{\mathcal{H}}_j := \bigcap_{1\le i\le j} \widehat{H}_i$.

Suppose that $\widehat{\mathcal{H}}_m = \emptyset$. Let k be the smallest integer such that $\widehat{\mathcal{H}}_k = \emptyset$. By definition, $k \ge 2$ and $\widehat{\mathcal{H}}_{k-1} \ne \emptyset$. Note that the intersection of any $k-1$ non-parallel hyperplanes of $\mathbb{R}^d$ includes a $(d-k+1)$-flat and each h_i contains the origin. Hence, $\bigcap_{1\le i\le k-1} h_i$ includes a $(d-k+1)$-flat through the origin and thus $\mathcal{H}_{k-1}$ includes a great $(d-k)$-sphere G on S^{d-1}. Since $x \in G$ implies $-x \in G$ for any $x \in S^{d-1}$, we must have $G \subseteq h_k$, in order to have an empty intersection $\widehat{\mathcal{H}}_k$. This implies that $\bigcap_{1\le i\le k} h_i$ also includes a $(d-k+1)$-flat through the origin, and further that $\bigcap_{1\le i\le m} h_i$ includes a $(d-m+1)$-flat through the origin. We hence conclude that $\mathcal{H}_m = \bigcap_{1\le i\le m} H_i$ includes a great $(d-m)$-sphere on S^{d-1}. □

Using Lemma 1 we prove the following theorem.

Theorem 1 *For any fixed positive integer d, let $\mathcal{F}$ be a finite family of real-valued convex functions defined on a convex subset $C \subseteq \mathbb{R}^d$ and $g(x) := \min_{f\in\mathcal{F}} f(x)$ be their pointwise minimum. Suppose that g attains a local maximum at $x^* \in C$ and there are exactly $m \le d$ functions $f_1, \ldots, f_m \in \mathcal{F}$ such that $f_i(x^*) = g(x^*)$ for each $i = 1, \ldots, m$. Then there exists a $(d+1-m)$-flat $\varphi \subset \mathbb{R}^d$ through x^* such that g is constant on $\varphi \cap U$ for some neighborhood $U \subset \mathbb{R}^d$ of x^* with $U \subset C$.*

Proof First, we give an overview of our proof for the theorem. All functions $f \in \mathcal{F}$ other than $f_1, \ldots, f_m$ must satisfy $f(x) > g(x)$ in a small neighborhood of x^*. In particular, the function g is the lower envelope of the m convex functions f_i in a small neighborhood of x^*. By convexity, we will show that for each i, there is a hemisphere H_i of directions in S^{d-1} in which f_i does not decrease. (Note that the sphere S^{d-1} represents the space of all directions in $\mathbb{R}^d$.) This result combined with Lemma 1 gives that the intersection of hemispheres will be a $(d+1-m)$-flat in which neither of the m functions (nor g) can decrease. Since x^* is a local maximum of g, the only possibility is that g remains constant near x^* along the flat.

A more detailed proof is given as follows. Let $x^* \in C$ and $f_1, \ldots, f_m \in \mathcal{F}$ be as in the statement. For each i, consider the sublevel set $L_i := \{x \in C \mid f_i(x) \leq f_i(x^*)\}$. Here, we consider two cases: (i) x^* lies in the interior of L_i or (ii) on its boundary ∂L_i. Note that $L_i \subseteq C$ is convex since f_i is a convex function. For the latter case (ii), there exists a supporting hyperplane h_i to L_i at x^* since L_i is convex and $x^* \in \partial L_i$. Denote by $h_i^{\oplus}$ the closed half-space that is bounded by h_i and does not contain L_i. For the former case (i), we choose h_i to be any hyperplane of $\mathbb{R}^d$ through x^* and $h_i^{\oplus}$ to be any closed half-space supported by h_i. Then we have that $f_i(x^*) \leq f_i(x)$ for any $x \in h_i^{\oplus} \cap C$, regardless of the cases; in particular for Case (i), observe that $f_i(x) = f_i(x^*)$ for any $x \in L_i$ by convexity so that we can choose any hyperplane as h_i.

Now, we let

$$H_i := \{x - x^* \mid x \in h_i^{\oplus},\ \|x - x^*\| = 1\}$$

be a closed hemisphere of the unit sphere S^{d-1} centered at the origin. Note that f_i does not decrease if we move from x^* locally in any direction in H_i. Since $g(x^*) = f_i(x^*)$ for any $i \in \{1, \ldots, m\}$ and x^* is a local maximum of g, the intersection $\bigcap_{i=1}^m H_i$ has an empty interior relative to S^{d-1}; otherwise, there exists $y \in \operatorname{int} \bigcap_{i=1}^m H_i$ such that $f_i(x^* + \lambda y) \geq f_i(x^*)$ for any $i \in \{1, \ldots, m\}$ and any $\lambda > 0$ with $x^* + \lambda y \in C$.

Hence, by Lemma 1, $\bigcap_{i=1}^m H_i$ has a nonempty intersection including a great $(d - m)$-sphere G on S^{d-1}. Let φ be the corresponding $(d - m + 1)$-flat in $\mathbb{R}^d$ through x^* defined as

$$\varphi := \{x^* + \lambda y \in \mathbb{R}^d \mid y \in G \text{ and } \lambda \in \mathbb{R}\}.$$

Consider the restriction $f_i|_{\varphi \cap C}$ of f_i on $\varphi \cap C$. Since f_i is convex and φ is an affine subspace (thus, convex), $f_i|_{\varphi \cap C}$ is also convex and their pointwise minimum $g|_{\varphi \cap C}$ attains a local maximum at $x^* \in \varphi \cap C$. On the other hand, each $f_i|_{\varphi \cap C}$ attains a local minimum at x^*; since $\varphi \subseteq h_i^{\oplus}$, we have $f_i(x^*) \leq f_i(x)$ for any point $x \in \varphi \cap C$. Hence, $g|_{\varphi \cap C}$ also attains a local minimum at x^* since $g(x^*) = f_i(x^*)$ for any $i \in \{1, \ldots, m\}$. Consequently, g is locally constant at x^* on φ; more precisely, there is a sufficiently small neighborhood $U \subset \mathbb{R}^d$ of x^* with $U \subset C$ such that g is constant on $U \cap \varphi$, completing the proof. □

4 Properties of Geodesic-Maximal Pairs

We call a pair $(s^*, t^*) \in \mathcal{P} \times \mathcal{P}$ *maximal* if (s^*, t^*) is a local maximum of the geodesic distance function d. That is, (s^*, t^*) is maximal if and only if there are two neighborhoods $U_s, U_t \subset \mathbb{R}^2$ of s^* and of t^*, respectively, such that for any $s \in U_s \cap \mathcal{P}$ and any $t \in U_t \cap \mathcal{P}$ we have $\mathrm{d}(s^*, t^*) \geq \mathrm{d}(s, t)$. Clearly, any diametral pair is maximal.

Consider any maximal pair (s^*, t^*) in $\mathcal{P}$. Let $\Pi(s^*, t^*)$ be the set of all shortest paths from s^* to t^*. Then each path $\pi \in \Pi(s^*, t^*)$ is associated with a pair of corners (u, v) that are its first and last corners as discussed in Sect. 2. Note that such a pair (u, v) of corners always exists for any $\pi \in \Pi(s^*, t^*)$; even if $\mathrm{d}(s^*, t^*) = \|s^* - t^*\|$, then both endpoints s^* and t^* must be corners in V by its maximality. We now focus on the set of such pairs of the first and last corners, defined to be

$$\mathcal{V}(s^*, t^*) := \{(u, v) \mid \exists \pi \in \Pi(s^*, t^*) \text{ s.t. } u, v \in V \text{ are the first and last corners along } \pi, \text{ resp.}\}.$$

We set $\mathcal{V}(s^*, t^*) = \{(u_1, v_1), \ldots, (u_m, v_m)\}$, where m is the cardinality of $\mathcal{V}(s^*, t^*)$, and the pairs of vertices are sorted in no particular order. Also, we let

$$V_{s^*} := \{u_1, \ldots, u_m\}, \quad V_{t^*} := \{v_1, \ldots, v_m\}.$$

Some immediate bounds are $|\Pi(s^*, t^*)| \geq m$, $|V_{s^*}| \leq m$, and $|V_{t^*}| \leq m$. Observe that it is not true that we always have the equality $|\Pi(s^*, t^*)| = m$; in some cases, there can be multiple shortest paths between a pair of corners. In the following, we show a tight bound on the cardinality m of the set $\mathcal{V}(s^*, t^*)$, provided that (s^*, t^*) is maximal.

Let E be the set of all sides of $\mathcal{P}$ without their endpoints and $\mathcal{B}$ be their union. Note that $\mathcal{B} = \partial\mathcal{P} \setminus V$ is the boundary of $\mathcal{P}$ except the corners V. The goal of this section is to prove the following theorem, which is the main combinatorial result of this paper.

Theorem 2 *Suppose that (s^*, t^*) is a maximal pair in $\mathcal{P}$, and that $\mathcal{V}(s^*, t^*)$, V_{s^*}, and V_{t^*} are defined as above. Then we have the following implications.*

(V-V)	$s^* \in V$,	$t^* \in V$	*implies*	$\lvert\mathcal{V}(s^*, t^*)\rvert \geq 1, \lvert V_{s^*}\rvert \geq 1, \lvert V_{t^*}\rvert \geq 1$;
(V-B)	$s^* \in V$,	$t^* \in \mathcal{B}$	*implies*	$\lvert\mathcal{V}(s^*, t^*)\rvert \geq 2, \lvert V_{s^*}\rvert \geq 1, \lvert V_{t^*}\rvert \geq 2$;
(V-I)	$s^* \in V$,	$t^* \in int\,\mathcal{P}$	*implies*	$\lvert\mathcal{V}(s^*, t^*)\rvert \geq 3, \lvert V_{s^*}\rvert \geq 1, \lvert V_{t^*}\rvert \geq 3$;
(B-B)	$s^* \in \mathcal{B}$,	$t^* \in \mathcal{B}$	*implies*	$\lvert\mathcal{V}(s^*, t^*)\rvert \geq 3, \lvert V_{s^*}\rvert \geq 2, \lvert V_{t^*}\rvert \geq 2$;
(B-I)	$s^* \in \mathcal{B}$,	$t^* \in int\,\mathcal{P}$	*implies*	$\lvert\mathcal{V}(s^*, t^*)\rvert \geq 4, \lvert V_{s^*}\rvert \geq 2, \lvert V_{t^*}\rvert \geq 3$;
(I-I)	$s^* \in int\,\mathcal{P}$,	$t^* \in int\,\mathcal{P}$	*implies*	$\lvert\mathcal{V}(s^*, t^*)\rvert \geq 5, \lvert V_{s^*}\rvert \geq 3, \lvert V_{t^*}\rvert \geq 3$.

Moreover, each of the above bounds is tight.

Together with the bound $|\Pi(s^*, t^*)| \geq |\mathcal{V}(s^*, t^*)|$, Theorem 2 immediately implies tight lower bounds on the number of shortest paths between any maximal pair.

Corollary 1 *For any $p \in \mathcal{P}$, let $\delta(p) := 0$ if $p \in V$; $\delta(p) := 1$ if $p \in \mathcal{B}$; $\delta(p) := 2$ if $p \in int\mathcal{P}$. If (s^*, t^*) is a maximal pair in $\mathcal{P}$, then we have*

$$|\Pi(s^*, t^*)| \geq \delta(s^*) + \delta(t^*) + 1.$$

Moreover, the above bound is tight.

To see the tightness of the bounds, we present examples with remarks in Fig. 1 and Appendix 1. In particular, one can easily see the tightness of the bounds on $|V_{s^*}|$ and $|V_{t^*}|$ from shortest path maps $\mathsf{SPM}(s^*)$ and $\mathsf{SPM}(t^*)$, when $V \cup \{s^*, t^*\}$ is in general position.

We first give an overview of the proof. The general reasoning is roughly the same for all the different scenarios, and we thus focus on the case in which (s^*, t^*) is a maximal pair and both s^* and t^* are interior points (Case (I-I)). Regard the geodesic distance function d as a four-variate function in a small convex neighborhood of (s^*, t^*). As mentioned in Sect. 2, the geodesic distance is the pointwise minimum of a finite number of path-length functions. Since the pair (s^*, t^*) is maximal, we will apply Theorem 1 and obtain that the geodesic distance is constant in a flat of dimension $d+1-m = 5-m$, where $m = |\mathcal{V}(s^*, t^*)|$. On the other hand, we will also show that the geodesic distance function can only remain constant in a zero-dimensional flat (i.e., at a point), hence $m \geq 5$. In the other cases (boundary–interior, boundary–boundary, etc.) the boundary of $\mathcal{P}$ introduces additional constraints that reduce the degrees of freedom of the geodesic distance function. Hence, fewer paths are enough to *pin* the solution.

The main technical difficulty of the proof is the fact that the path-length functions $\mathrm{len}_{u,v}$ are not globally defined. Thus, we must properly extend them in a way that all conditions of Theorem 1 are satisfied.

4.1 Proof of Theorem 2

We start with several basic observations. The proof of Theorem 2 will be done separately for each case.

The following lemma proves the bounds on $|V_{s^*}|$ and $|V_{t^*}|$ of Theorem 2.

Lemma 2 *Let* (s^*, t^*) *be a maximal pair.*

1. *If* $t^* \in \mathcal{B}$, *then* $|V_{t^*}| \geq 2$. *Moreover, if* $t^* \in e \in E$, *then there exists* $v \in V_{t^*}$ *such that* v *is off the line supporting* e.
2. *If* $t^* \in \operatorname{int}\mathcal{P}$, *then* $|V_{t^*}| \geq 3$ *and* t^* *lies in the interior of the convex hull of* V_{t^*}.

Proof Since (s^*, t^*) is a maximal pair, the function $d_{s^*}(t) := \mathrm{d}(s^*, t)$ is maximized at $t = t^*$ on a sufficiently small subset $U \subset \mathcal{P}$ with $t^* \in U$. As discussed in Sect. 2, if $t^* \notin V$, then t^* must be either a vertex of $\mathsf{SPM}(s^*)$ or an intersection point between an edge of $\mathsf{SPM}(s^*)$ and $\partial\mathcal{P}$. If $t^* \in \operatorname{int}\mathcal{P}$, then t^* should fall into the former case and hence we have at least three corners $v_1, v_2, v_3 \in V$ determining the vertex t^* of $\mathsf{SPM}(s^*)$. If $t^* \in \mathcal{B}$, then t^* may also occur at the latter case. In that case, t^* lies on an edge of $\mathsf{SPM}(s^*)$ and thus we have at least two corners $v_1, v_2 \in V$ determining an edge of $\mathsf{SPM}(s^*)$.

The other claims of the lemma can be shown as follows. If $t^* \in \operatorname{int}\mathcal{P}$ but t^* lies out of the interior of the convex hull of V_{t^*}, then we can find another point $t \in \mathcal{P}$ arbitrarily close to t^* such that $\|t - v_i\| > \|t^* - v_i\|$ for every $v_i \in V_{t^*}$. This implies that $\mathrm{d}(s^*, t) > \mathrm{d}(s^*, t^*)$, contradicting the maximality of (s^*, t^*). If $t^* \in e \in E$ but every $v_i \in V_{t^*}$ lies on the supporting line ℓ of e, then we obtain a strictly larger distance than $\mathrm{d}(s^*, t^*)$, as moving t^* in a perpendicular direction to ℓ. (Notice that a similar argument can be also found in [17, Lemma 2.2]). □

Lemma 2 immediately implies the lower bound on $|\mathcal{V}(s^*, t^*)|$ when $s^* \in V$ or $t^* \in V$ since $|\mathcal{V}(s^*, t^*)| \geq \max\{|V_{s^*}|, |V_{t^*}|\}$. This completes Cases (V-*). Note that the bounds for Case (V-V) are trivial.

From now on, we assume that neither s^* nor t^* is a corner of $\mathcal{P}$. This assumption, together with Lemma 2, implies multiple shortest paths between s^* and t^*, and thus $\mathrm{d}(s^*, t^*) > \|s^* - t^*\|$. Hence, as discussed in Sect. 2, any maximal pair falling into one of Cases (B-B), (B-I), and (I-I) appears as a local maximum of the lower envelope of some path-length functions.

Case (I-I): *When both s^* and t^* lie in* $\mathrm{int}\,\mathcal{P}$. We will apply Theorem 1 to prove Theorem 2 for Case (I-I). Recall the definition of $\mathcal{V}(s^*, t^*) = \{(u_1, v_1), \ldots, (u_m, v_m)\}$ and $m = |\mathcal{V}(s^*, t^*)|$. For each $(u_i, v_i) \in \mathcal{V}(s^*, t^*)$, we have the corresponding shortest path π_i between s^* and t^* and $\mathrm{len}_{u_i,v_i}(s^*, t^*) = \mathrm{d}(s^*, t^*)$. Thus, we have at least m functions f among$\{\mathrm{len}_{u,v} \mid u, v \in V\}$ such that $f(s^*, t^*) = \mathrm{d}(s^*, t^*)$. If the number of such path-length functions are exactly m, we can apply Theorem 1 directly.

Unfortunately, this is not always the case. A single shortest path $\pi_i \in \Pi(s^*, t^*)$ may give additional pairs (u, v) of corners with $u, v \in \pi_i$ such that $(u, v) \neq (u_i, v_i)$ and $\mathrm{len}_{u,v}(s^*, t^*) = \mathrm{d}(s^*, t^*)$. This situation can occur even when the corners of $\mathcal{P}$ are in general position. Observe that this happens only when u, u_i, s^* or v, v_i, t^* are collinear. In order to resolve this problem, we define the *merged* path-length functions that satisfy all the requirements of Theorem 1 even in degenerate cases.

Recall that the combinatorial structure of each shortest path π_i can be represented by a sequence $(u_i = u_{i,1}, \ldots, u_{i,k} = v_i)$ of corners in V. We define u'_i to be one of the $u_{i,j}$ as follows. If s^* does not lie on the line ℓ through u_i and $u_{i,2}$, then $u'_i := u_i$; otherwise, if $s^* \in \ell$, then $u'_i := u_{i,j}$, where j is the largest index such that for any open neighborhood $U \subset \mathbb{R}^2$ of s^* there exists a point $s \in (U \cap \mathsf{VR}(u_{i,j})) \setminus \ell$. Intuitively speaking, u'_i is the last corner of π that sees a neighborhood of s^* (although it need not see s^*). Note that such u'_i always exists, and if no three of V are collinear, then we always have either $u'_i = u_i$ or $u'_i = u_{i,2}$. Figure 2a illustrates how to determine u'_i. Also, we define v'_i in an analogous way. Let $\alpha_i : \mathsf{VR}(u'_i) \cup \mathsf{VR}(u_i) \to \mathbb{R}$ and $\omega_i : \mathsf{VR}(v'_i) \cup \mathsf{VR}(v_i) \to \mathbb{R}$ be two functions defined as

$$\alpha_i(s) := \begin{cases} \|s - u'_i\| & \text{if } s \in \mathsf{VR}(u'_i), \\ \|s - u_i\| + \|u_i - u'_i\| & \text{if } s \in \mathsf{VR}(u_i) \setminus \mathsf{VR}(u'_i); \end{cases}$$

$$\omega_i(t) := \begin{cases} \|t - v'_i\| & \text{if } t \in \mathsf{VR}(v'_i), \\ \|t - v_i\| + \|v_i - v'_i\| & \text{if } t \in \mathsf{VR}(v_i) \setminus \mathsf{VR}(v'_i). \end{cases}$$

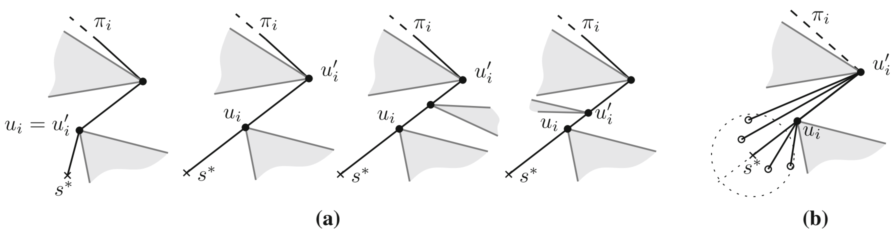

Fig. 2 **a** How to determine u'_i. (*left to right*) $u'_i = u_{i,1}, u_{i,2}, u_{i,3}$, and $u_{i,2}$; Observe that in the last two examples we have three collinear points. **b** For points in a small disk B centered at s^* with $B \subset \mathsf{VR}(u'_i) \cup \mathsf{VR}(u_i)$, the function α_i measures the length of the shortest path from u'_i to the depicted points

This allows us to define the merged path-length function $f_i : D_i \to \mathbb{R}$ as

$$f_i(s,t) := \alpha_i(s) + \mathrm{d}(u_i', v_i') + \omega_i(t),$$

where $D_i := (\mathsf{VR}(u_i') \cup \mathsf{VR}(u_i)) \times (\mathsf{VR}(v_i') \cup \mathsf{VR}(v_i)) \subseteq \mathcal{P} \times \mathcal{P}$ (see Fig. 2b). We consider $\mathcal{P} \times \mathcal{P}$ as a subset of $\mathbb{R}^4$ and each pair $(s,t) \in \mathcal{P} \times \mathcal{P}$ as a point in $\mathbb{R}^4$. Also, we denote by (s_x, s_y) the coordinates of a point $s \in \mathcal{P}$ and we write $s = (s_x, s_y)$ or $(s,t) = (s_x, s_y, t_x, t_y)$ by an abuse of notation. Observe that

$$f_i(s,t) = \min\{\mathrm{len}_{u_i,v_i}(s,t), \mathrm{len}_{u_i',v_i}(s,t), \mathrm{len}_{u_i,v_i'}(s,t), \mathrm{len}_{u_i',v_i'}(s,t)\}$$

for any $(s,t) \in D_i$ if we define $\mathrm{len}_{u,v}(s,t) = \infty$ when $s \notin \mathsf{VR}(u)$ or $t \notin \mathsf{VR}(v)$.

Lemma 3 *The following properties hold for the functions f_i.*

(i) *$f_i(s^*, t^*) = d(s^*, t^*)$ for any $i \in \{1, \dots, m\}$.*
(ii) *There exists a convex neighborhood $C \subset \mathbb{R}^4$ of (s^*, t^*) with $C \subseteq \bigcap_{i=1}^m D_i$ such that $d(s,t) = \min_{i \in \{1,\dots,m\}} f_i(s,t)$ for any $(s,t) \in C$.*
(iii) *Each of the functions f_i for $i \in \{1, \dots, m\}$ is convex on C.*
(iv) *For any $i \in \{1, \dots, m\}$, there exists a unique line $\ell_i \subset \mathbb{R}^4$ through $(s^*, t^*) \in \mathbb{R}^4$ such that f_i is constant on $\ell_i \cap C$. Moreover, there exists at most one index $j \neq i$ such that $\ell_i = \ell_j$.*
(v) *For any $i, j \in \{1, \dots, m\}$, any $(s,t) \in C$, and any neighborhood $U \subseteq C$ of (s,t), there exists $(s', t') \in U$ such that $f_i(s,t) < f_i(s', t')$ and $f_j(s,t) < f_j(s', t')$.*

Proof (i) This immediately follows from the fact that $f_i(s^*, t^*) = \mathrm{len}_{u_i,v_i}(s^*, t^*)$.
(ii) In this proof, we extend $\mathrm{len}_{u,v}$ to any $(s,t) \in \mathcal{P} \times \mathcal{P}$ where $\mathrm{len}_{u,v}(s,t) = \infty$ if $s \notin \mathsf{VR}(u)$ or $t \notin \mathsf{VR}(v)$. By the definition of f_i, there exists a small neighborhood $U_i \subset D_i$ of (s^*, t^*) such that $f_i(s,t) = \min\{\mathrm{len}_{u_i,v_i}(s,t), \mathrm{len}_{u_i',v_i}(s,t), \mathrm{len}_{u_i,v_i'}(s,t), \mathrm{len}_{u_i',v_i'}(s,t)\} = \min_{u,v \in \pi_i \cap V} \mathrm{len}_{u,v}(s,t)$ for all $(s,t) \in U_i$. We claim that there exists an open convex neighborhood $C \subset \bigcap_i U_i$ such that for any $(s,t) \in C$

$$\mathrm{d}(s,t) = \min_{1 \le i \le m} f_i(s,t).$$

To prove our claim, assume to the contrary that for every open convex neighborhood $C \subset \mathbb{R}^4$ of $(s^*, t^*) \in \mathbb{R}^4$ there exist a pair (u, v) of corners and $(s,t) \in C$ such that $\mathrm{d}(s,t) = \mathrm{len}_{u,v}(s,t) < \min_i f_i(s,t)$. Note that none of the shortest paths $\pi_i \in \Pi(s^*, t^*)$ between s^* and t^* can have u and v as the first and last corners along π_i (since, otherwise, we must have $(u,v) \in \mathcal{V}(s^*, t^*)$ and thus $(u,v) = (u_j, v_j)$ for some $1 \le j \le m$). This implies that $\mathrm{d}(s,t) = \mathrm{len}_{u_j,v_j}(s,t) < \min_i f_i(s,t) = \mathrm{d}(s,t)$, a contradiction.

Consider a sequence $C_1, C_2, \dots$ of neighborhoods of $(s^*, t^*) \in \mathbb{R}^4$ that converges to the singleton $\{(s^*, t^*)\}$. Since there are only n^2 pairs of corners, there exist a fixed pair (u_0, v_0) of corners and a subsequence $C_{k_1}, C_{k_2}, \dots$ converging to the singleton $\{(s^*, t^*)\}$ such that none of the π_i pass through both u_0 and v_0, and for any integer $j > 0$ there exists $(s_j, t_j) \in C_{k_j}$ with

$$\mathrm{d}(s_j, t_j) = \mathrm{len}_{u_0,v_0}(s_j, t_j) < \min_{1\le i\le m} f_i(s_j, t_j).$$

Since $\lim_{j\to\infty}(s_j, t_j) = (s^*, t^*)$, it holds that $\lim_{j\to\infty} \mathrm{d}(s_j, t_j) = \lim_{j\to\infty} \min_i f_i(s_j, t_j) = \mathrm{d}(s^*, t^*)$ by Property (i). By the sandwich theorem, we have

$$\lim_{j\to\infty} \mathrm{len}_{u_0,v_0}(s_j, t_j) = \mathrm{len}_{u_0,v_0}(s^*, t^*) = \mathrm{d}(s^*, t^*).$$

This implies the existence of the $(m+1)$-st shortest path between s^* and t^* since none of the $\pi_i \in \Pi(s^*, t^*)$ contains both u_0 and v_0, a contradiction.
(iii) Since the sum of convex functions is a convex function, it suffices to show that α_i and ω_i are convex. More precisely, for any $(s_1, t_1), (s_2, t_2) \in C$ and $0 \le \lambda \le 1$, we have

$$\begin{aligned} f_i(\lambda(s_1, t_1) + (1-\lambda)(s_2, t_2)) &= \alpha_i(\lambda s_1 + (1-\lambda)s_2) + \mathrm{d}(u_i', v_i') \\ &\quad + \omega_i(\lambda t_1 + (1-\lambda)t_2) \\ &\le \lambda\alpha_i(s_1) + (1-\lambda)\alpha_i(s_2) + \mathrm{d}(u_i', v_i') + \lambda\omega_i(t_1) \\ &\quad + (1-\lambda)\omega_i(t_2) \\ &\le \lambda f_i(s_1, t_1) + (1-\lambda) f_i(s_2, t_2) \end{aligned}$$

if α_i and ω_i are convex.

We now show the convexity of α_i on any convex subset $C \subset \mathsf{VR}(u_i') \cup \mathsf{VR}(u_i)$. Note that the convexity of ω_i can be shown in the same way. There are two cases: $u_i' = u_i$ or $u_i' \neq u_i$. For the former case, α_i is convex on C since it measures the Euclidean distance between u_i and a given point in C. For the latter case, let ℓ_0 be the line through u_i, u_i', and also s^*. Then C may be partitioned by ℓ_0 into two regions A_1 and A_2, where $A_1 = C \cap \mathsf{VR}(u_i')$ and $A_2 = C \setminus A_1$. Note that α_i is convex on A_1 and on A_2. Thus, we are done by checking every point on $\ell_0 \cap C$.

Pick any $s \in \ell_0 \cap C$ and any line $\ell \subset \mathbb{R}^2$ through s. Let θ be the angle between ℓ_0 and ℓ. If we restrict the domain of α_i on $\ell \cap C$, then one can check with elementary calculus that both the derivatives of $\|s - u_i\| + \|u_i - u_i'\|$ and of $\|s - u_i'\|$ are equal to $c\cos\theta$ at s for some constant c. Hence, α_i is smooth and convex along ℓ. Since we have taken any line ℓ through any point on $\ell_0 \cap C$, this suffices to prove the convexity of α_i on C.
(iv) Fix any $i \in \{1, \dots, m\}$. Any ray $\gamma \subset \mathbb{R}^4$ with endpoint $(s^*, t^*) \in \mathbb{R}^4$ can be determined by three parameters $(\theta_{s^*}, \theta_{t^*}, \lambda)$ with $0 \le \theta_{s^*}, \theta_{t^*} \le \pi$ and $\lambda \ge 0$ as follows: Let γ_{s^*} and γ_{t^*} be the projections of γ onto the (s_x, s_y)-plane and the (t_x, t_y)-plane, respectively. Note that γ_{s^*} is a ray in the (s_x, s_y)-plane with endpoint s^* and γ_{t^*} is a ray in the (t_x, t_y)-plane with endpoint t^*. Let θ_{s^*} be the smaller angle at s^* made by γ_{s^*} and another ray starting from s^* in the direction away from u_i. Define θ_{t^*} analogously with γ_{t^*}, t^*, and v_i. The derivative of f_i at (s^*, t^*) along γ is represented as $c(\cos\theta_{s^*} + \lambda\cos\theta_{t^*})$ for some constants $\lambda \ge 0$ and $c > 0$ depending only on i. Note that the second derivative of f_i at (s^*, t^*) along γ is derived as

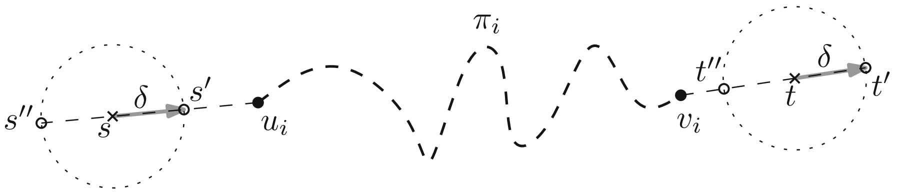

Fig. 3 Illustration to Lemma 3(iv); for any $(s,t) \in D_i$ and any sufficiently small δ if we pick (s',t') such that s' is δ closer to u_i than s and t' is δ farther from v_i than t, then we have $f_i(s',t') = f_i(s,t)$. Symmetrically, $f_i(s'',t'') = f_i(s,t)$ with $\|s''-u_i\| = \|s-u_i\| + \delta$ and $\|t''-v_i\| = \|t-v_i\| - \delta$

$$c\left(\frac{\sin^2\theta_{s^*}}{\|s^*-u_i\|} + \lambda\frac{\sin^2\theta_{t^*}}{\|t^*-v_i\|}\right).$$

Suppose that f_i is constant along γ locally around (s^*,t^*). Then, its first and second derivatives along γ should be zero in a small neighborhood $U \subset \mathbb{R}^4$ of (s^*,t^*) with $U \subset D_i$. First, we observe that λ should be positive; if $\lambda = 0$, then $t = t^*$ is fixed while s moves from s^* along γ_{s^*}, and hence f_i does not stay constant. Since every term of the second derivative is nonnegative and $\lambda > 0$, we only obtain two solutions $(\theta_{s^*}, \theta_{t^*}) = (0,\pi)$ or $(\pi,0)$. Consequently, we have two such rays $\gamma = (0,\pi,1)$ or $(\pi,0,1)$ such that f_i remains constant along γ. These two rays form a unique line $\ell_i \subset \mathbb{R}^4$ through (s,t) such that f_i is constant on $\ell_i \cap U$. See Fig. 3 for a more intuitive and geometric description of ℓ_i.

The projections of ℓ_i onto the (s_x,s_y)-plane and the (t_x,t_y)-plane appear the lines through s^* and u_i and through t^* and v_i, respectively. Hence, one can easily check that f_i remains constant on $\ell_i \cap D_i$, which completes the proof of the first part of the claim.

We now show the second part of the claim. As observed above, we have that the projection of ℓ_i onto the (s_x,s_y)-plane is the line through s^* and u_i. Also, the projection of ℓ_i onto the (t_x,t_y)-plane is the line through t^* and v_i. Hence, $\ell_i = \ell_j$ implies that u_i, u_j, s^* are collinear and v_i, v_j, t^* are collinear. First, since the pairs (u_i,v_i) are all distinct, we have $u_i \neq u_j$ or $v_i \neq v_j$. If $u_i = u_j$ and $v_i \neq v_j$, then one can easily check that $\ell_i \neq \ell_j$ from the geometric interpretation of ℓ_i as shown in Fig. 3. We hence have $u_i \neq u_j$ and $v_i \neq v_j$. Moreover, s^* must lie in between u_i and u_j and t^* must lie in between v_i and v_j by definition; if u_j lies in between u_i and s^*, then the first corner of π_i from s^* becomes u_j since the three are collinear. Therefore, for each $i \in \{1,\ldots,m\}$, there is at most one index $j \in \{1,\ldots,m\}$ such that $i \neq j$ and $\ell_i = \ell_j$.

(v) Pick any $i, j \in \{1,\ldots,m\}$ and consider the sublevel sets $L_i = \{(\tilde{s},\tilde{t}) \in \mathbb{R}^4 \mid f_i(\tilde{s},\tilde{t}) \le f_i(s,t)\}$ and $L_j = \{(\tilde{s},\tilde{t}) \in \mathbb{R}^4 \mid f_j(\tilde{s},\tilde{t}) \le f_j(s,t)\}$. Since f_i and f_j are convex and non-constant functions, L_i and L_j are closed convex sets that have (s,t) on their boundaries. Therefore, there exist hyperplanes h_i and h_j tangent to L_i and L_j, respectively, at (s,t). Let $h_i^{\oplus}$ be a closed half-space bounded by h_i that avoids L_i and $H_i' := \{(s',t') \in h_i^{\oplus} \mid \|(s',t')-(s,t)\| = 1\}$ be a closed hemisphere on the unit sphere centered at (s,t). Define H_j' analogously for h_j.

Since H_i' and H_j' are closed hemispheres with a common center, $H_i' \cap H_j' \neq \emptyset$. By construction, we have $f_i(s,t) \le f_i(s',t')$ for any $(s',t') \in H_i'$, and $f_j(s,t) \le$

$f_j(s', t')$ for any $(s', t') \in H'_j$. On the other hand, by Property (iv) of the lemma, the equality holds only when (s', t') lies on line ℓ_i or ℓ_j, respectively. Therefore, for any $(s', t') \in (H'_i \cap H'_j) \setminus (\ell_i \cup \ell_j)$, the claimed inequalities $f_i(s, t) < f_i(s', t')$ and $f_j(s, t) < f_j(s', t')$ hold strictly. The last task is to check that $(H'_i \cap H'_j) \setminus (\ell_i \cup \ell_j) \neq \emptyset$, which follows by Lemma 1. □

Returning to the proof of Theorem 2, we take a convex neighborhood C of (s^*, t^*) satisfying Property (ii) of Lemma 3 and apply Theorem 1. Note that Properties (i)–(iii) of Lemma 3 ensure that the preconditions of Theorem 1 are satisfied.

Suppose that $m < 5$. Then, by Theorem 1, there exists at least one line $\ell \in \mathbb{R}^4$ through (s^*, t^*) such that d is constant on $\ell \cap C$. Since (s^*, t^*) is a local maximum, there exists a small neighborhood $U \subset C$ of (s^*, t^*) such that $\mathrm{d}(s, t) \leq \mathrm{d}(s^*, t^*)$ for all $(s, t) \in U$. By Property (iv) of Lemma 3, at most two functions f_i are constant on $\ell \cap U$. Without loss of generality, we can assume that functions $f_3, \dots, f_m$ are not constant. Since the geodesic distance function d is constant on $\ell \cap U$ and $\mathrm{d}(s, t) = \min_{i \in \{1,\dots,m\}} f_i(s, t)$, any of $f_3, \dots, f_m$ must strictly increase in both directions along ℓ. That is, for any $(s', t') \in \ell \cap U$ with $(s', t') \neq (s^*, t^*)$ and for all $i \geq 3$, we have $\min\{f_1(s', t'), f_2(s', t')\} < f_i(s', t')$. Thus, there exists a small neighborhood $U' \subseteq U$ of (s', t') such that $\mathrm{d}(s, t) = \min\{f_1(s, t), f_2(s, t)\}$ for all $(s, t) \in U'$. However, by Property (v) of Lemma 3, there exists a pair $(s'', t'') \in U'$ such that $f_1(s', t') < f_1(s'', t'')$ and $f_2(s', t') < f_2(s'', t'')$, contradicting the maximality of (s^*, t^*) (see Fig. 4). Hence, we achieve a bound $m = |\mathcal{V}(s^*, t^*)| \geq 5$, as claimed in Case (I-I) of Theorem 2.

Case (B-B): *When Both s^* and t^* Lie on $\mathcal{B}$*. In this case, we assume that $s^* \in e_s \in E$ and $t^* \in e_t \in E$. The outline of the proof is analogous to the above discussion for Case (I-I); the only difference is that the search space has a lower dimension.

Let p be an endpoint of e_s and l_s be the length of e_s. We denote by $s(\zeta_s)$ the unique point on e_s such that $\|s(\zeta_s) - p\| = \zeta_s$ for any $0 < \zeta_s < l_s$. Thus, $s\colon (0, l_s) \to e_s$ establishes a bijection between the open interval $(0, l_s) \subset \mathbb{R}$ and the segment $e_s \subset \mathbb{R}^2$ except its endpoints. We also define $t(\zeta_t)$, analogously. Then, we let $\bar{f}_i\colon D_i \to \mathbb{R}$ be a function defined as the composition of f_i and the two bijections:

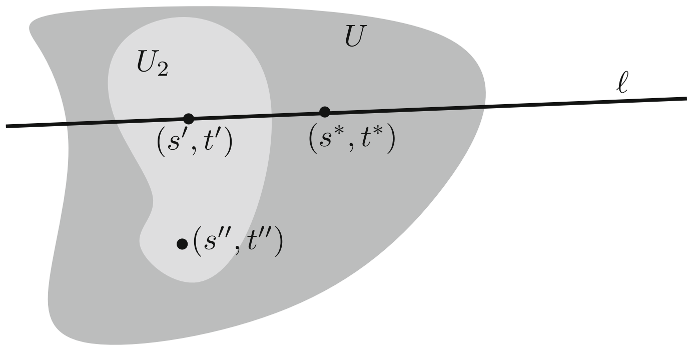

Fig. 4 Proof of Theorem 2 for the (I-I) case; If $\mathrm{d}(s, t)$ is constant on ℓ, we can find pairs of points (s'', t'') arbitrarily close to (s^*, t^*) whose geodesic distance is larger than $\mathrm{d}(s^*, t^*)$

$$\bar{f}_i(\zeta_s, \zeta_t) := \alpha_i(s(\zeta_s)) + \mathrm{d}(u_i', v_i') + \omega_i(t(\zeta_t)),$$

where the domain of $\bar{f}_i$ is $D_i := s^{-1}((\mathsf{VR}(u_i') \cup \mathsf{VR}(u_i)) \cap e_s) \times t^{-1}((\mathsf{VR}(v_i') \cup \mathsf{VR}(v_i)) \cap e_t)$. We consider D_i as a subset of $\mathbb{R}^2$ and each pair $(\zeta_s, \zeta_t) \in D_i$ as a point in $\mathbb{R}^2$. Let ζ_s^* and ζ_t^* be real numbers such that $s^* = s(\zeta_s^*)$ and $t^* = t(\zeta_t^*)$. We obtain the analogue of Lemma 3.

Lemma 4 *The following properties hold for the functions $\bar{f}_i$.*

(i) $\bar{f}_i(\zeta_s^*, \zeta_t^*) = d(s(\zeta_s^*), t(\zeta_t^*))$ *for any* $i \in \{1, \dots, m\}$.
(ii) *There exists a convex neighborhood* $C \subset \mathbb{R}^2$ *of* (ζ_s^*, ζ_t^*) *with* $C \subseteq \bigcap_{i=1}^m D_i$ *such that* $d(s(\zeta_s), t(\zeta_t)) = \min_{i\in\{1,\dots,m\}} \bar{f}_i(\zeta_s, \zeta_t)$ *for any* $(\zeta_s, \zeta_t) \in C$.
(iii) *Each of the functions* $\bar{f}_i$ *for* $i \in \{1, \dots, m\}$ *is convex on* C.
(iv) *If there exists a line* $\ell_i \subset \mathbb{R}^2$ *such that* $\bar{f}_i$ *is constant on* $\ell_i \cap C$, *then* u_i *lies on the line supporting* e_s *and* v_i *lies on the line supporting* e_t.
(v) *For any* $i \in \{1, \dots, m\}$, *any* $(\zeta_s, \zeta_t) \in C$, *and any neighborhood* $U \subseteq C$ *of* (ζ_s, ζ_t), *there exists* $(\zeta_s', \zeta_t') \in U$ *such that* $\bar{f}_i(\zeta_s, \zeta_t) < \bar{f}_i(\zeta_s', \zeta_t')$.

Note that the above claims are almost identical to those of Lemma 3. The results have been adapted taking into account that $\bar{f}_i$ is the composition of f_i and both ζ_s and ζ_t. Proofs follow verbatim, thus we omit them. Property (v) is the only exception: since the degrees of freedom have decreased, we cannot certify the existence of points arbitrarily close that increase two functions $\bar{f}_i$. Instead, we will use the second property of Lemma 2 to lead to a contradiction.

Recall that by the first claim of Lemma 2 we have $m \geq 2$. Thus, we are done by showing that the case $m = 2$ is not possible. Suppose that $m = 2$. Then, by Theorem 1, there exists a line $\ell \subset \mathbb{R}^2$ through $(\zeta_s^*, \zeta_t^*) \in \mathbb{R}^2$ such that d is constant on $\ell \cap C$. By the second claim of Lemma 2, there exists a vertex $v \in V_{s^*}$ off the line supporting e_s. Without loss of generality, we assume that $v = v_2$. By Property (iv) of Lemma 4, function $\bar{f}_2$ cannot remain constant in any line.

Now, we proceed as in Case (I-I). Consider any small neighborhood $U \subseteq C$ of (ζ_s^*, ζ_t^*). Any point $(\zeta_s', \zeta_t') \in \ell \cap C$ with $(\zeta_s', \zeta_t') \neq (\zeta_s^*, \zeta_t^*)$ satisfies the strict inequality $\mathrm{d}(s(\zeta_s'), t(\zeta_t')) = \bar{f}_1(\zeta_s', \zeta_t') < \bar{f}_2(\zeta_s', \zeta_t')$, since $\bar{f}_2$ cannot remain constant and d is a local maximum. Thus, there exists a sufficiently small neighborhood $U' \subseteq U$ of (ζ_s', ζ_t') such that $\mathrm{d}(s(\zeta_s), t(\zeta_t)) = \bar{f}_1(\zeta_s, \zeta_t)$ for all $(\zeta_s, \zeta_t) \in U'$.

Now, we apply Property (v) of Lemma 4 to obtain a point (ζ_s'', ζ_t'') arbitrarily close to (ζ_s^*, ζ_t^*) with strict inequality $\mathrm{d}(s(\zeta_s''), t(\zeta_t'')) > \mathrm{d}(s(\zeta_s^*), t(\zeta_t^*)) = \mathrm{d}(s^*, t^*)$, contradicting the maximality of (s^*, t^*). We hence conclude that $m = |\mathcal{V}(s^*, t^*)| \geq 3$ for Case (B-B) when both s^* and t^* lie on $\mathcal{B}$.

Case (B-I): *When* $s^* \in \mathcal{B}$ *and* $t^* \in \operatorname{int} \mathcal{P}$. This case is a mixture of the two previous cases. Without loss of generality, we can also assume that $s^* \in e_s \in E$ and $t^* \in \operatorname{int} \mathcal{P}$. We define $s(\zeta_s)$ as in Case (B-B) with $s(\zeta_s^*) = s^*$. We now define function $\hat{f}_i : D_i \to \mathbb{R}$ as $\hat{f}_i(\zeta_s, t_x, t_y) := \alpha_i(s(\zeta_s)) + \mathrm{d}(u_i', v_i') + \omega_i(t_x, t_y)$, where $D_i := s^{-1}((\mathsf{VR}(u_i') \cup \mathsf{VR}(u_i)) \cap e_s) \times (\mathsf{VR}(v_i') \cup \mathsf{VR}(v_i))$ is a subset of $\mathbb{R}^3$.

We obtain another analogue of Lemmas 3 and 4.

Lemma 5 *The following properties hold for the functions $\hat{f}_i$.*

(i) $\hat{f}_i(\zeta_s^*, t^*) = d(s(\zeta_s^*), t)$ *for any* $i \in \{1, \dots, m\}$.
(ii) *There exists a convex neighborhood* $C \subset \mathbb{R}^3$ *of* (ζ_s^*, t^*) *with* $C \subseteq \bigcap_{i=1}^m D_i$ *such that* $d(s(\zeta_s), t) = \min_{i \in \{1,\dots,m\}} \hat{f}_i(\zeta_s, t)$ *for any* $(\zeta_s, t) \in C$.
(iii) *Each of the functions* $\hat{f}_i$ *for* $i \in \{1, \dots, m\}$ *is convex on* C.
(iv) *For any* $i \in \{1, \dots, m\}$, *there exists a unique line* $\ell_i \subset \mathbb{R}^3$ *through* $(\zeta_s^*, t^*) \in \mathbb{R}^3$ *such that* $\hat{f}_i$ *is constant on* $\ell_i \cap C$. *Moreover, there is at most one index* $j \neq i$ *such that* $\ell_i = \ell_j$.
(v) *For any* $i, j \in \{1, \dots, m\}$, *any* $(\zeta_s, t) \in C$, *and any neighborhood* $U \subseteq C$ *of* (ζ_s, t), *there exists* $(\zeta_s', t') \in U$ *such that* $\hat{f}_i(\zeta_s, t) < \hat{f}_i(\zeta_s', t')$ *and* $\hat{f}_j(\zeta_s, t) < \hat{f}_j(\zeta_s', t')$.

We proceed as in Case (B-B). Suppose $m \leq 3$ and apply Theorem 1. Then we obtain a line ℓ such that the geodesic distance (composed with ζ_s) is constant on $\ell \cap C$. However, since at most two functions f_i can remain constant on ℓ by Property (iv) of Lemma 5, there must exist a point arbitrarily close to (ζ_s^*, t^*) with strictly larger function value. Details are almost identical to the previous cases, and we get the claimed bound $m = |\mathcal{V}(s^*, t^*)| \geq 4$ for Case (B-I).

The claimed bounds on $|V_{s^*}|$ and $|V_{t^*}|$ are shown by Lemma 2, which completes the proof of Theorem 2. □

5 Computing the Geodesic Diameter

Since a diametral pair is in fact maximal, it falls into one of the cases shown in Theorem 2. In order to find a diametral pair we examine all possible scenarios accordingly.

Cases (V-*), where at least one point is a corner in V, can be handled in $O(n^2 \log n)$ time by computing $\mathsf{SPM}(v)$ for every $v \in V$ and traversing it to find the farthest point from v, as discussed in Sect. 2. We thus focus on Cases (B-B), (B-I), and (I-I), where a diametral pair consists of two non-corner points.

From the computational point of view, the most difficult case corresponds to Case (I-I) of Theorem 2. In particular, if $|V_{s^*}| = |V_{t^*}| = 5$, ten corners of V are involved and thus any exhaustive method would check $O(n^{10})$ possibilities to find maximal pairs of this case. Observe that such a case can happen even under a general position assumption as shown in Appendix 1(c). By Theorem 2, in Case (I-I), it is guaranteed that there are at least five distinct pairs $(u_1, v_1), \dots, (u_5, v_5)$ of corners in V such that $\mathrm{len}_{u_i,v_i}(s^*, t^*) = \mathrm{d}(s^*, t^*)$ for any $i \in \{1, \dots, 5\}$ and the system of equations $\mathrm{len}_{u_1,v_1}(s, t) = \cdots = \mathrm{len}_{u_5,v_5}(s, t)$ determines a 0-dimensional zero set, corresponding to a constant number of candidate pairs in $\operatorname{int}\mathcal{P} \times \operatorname{int}\mathcal{P}$. On the other hand, each path-length function $\mathrm{len}_{u,v}$ is an algebraic function of degree at most 4. Thus, given five distinct pairs (u_i, v_i) of corners, we can compute all candidate pairs (s, t) in $O(1)$ time by solving the system.[4] For each candidate pair we compute the geodesic distance between the pair to check its validity. Since the geodesic distance between

[4] Here, we assume that fundamental operations on a constant number of polynomials of constant degree with a constant number of variables can be performed in constant time.

any two points $s, t \in \mathcal{P}$ can be computed in $O(n \log n)$ time [13], we obtain a brute-force $O(n^{11} \log n)$-time algorithm, checking $O(n^{10})$ candidate pairs obtained from all possible combinations of 10 corners in V.

As a different approach, one can exploit the SPM-*equivalence decomposition* of $\mathcal{P}$, which subdivides $\mathcal{P}$ into regions such that the shortest path map of any two points in a common region are *topologically equivalent* [8]. It is not difficult to see that if (s, t) is a pair of points that equalizes any five path-length functions, then both s and t appear as vertices of the decomposition. However, the current best upper bound on the complexity of the SPM-equivalence decomposition is $O(n^{10})$ [8], and thus this approach hardly leads to a remarkable improvement.

Instead, we do the following for Case (I-I) with $|V_{s^*}| = 5$. We choose any five corners $u_1, \ldots, u_5 \in V$ (as a candidate for the set V_{s^*}) and overlay their shortest path maps $\mathsf{SPM}(u_i)$. Since each $\mathsf{SPM}(u_i)$ has $O(n)$ complexity, the overlay consists of $O(n^2)$ cells. Any cell of the overlay is the intersection of five cells associated with $v_1, \ldots, v_5 \in V$ in $\mathsf{SPM}(u_1), \ldots, \mathsf{SPM}(u_5)$, respectively. Choosing a cell of the overlay, we get five (possibly, not distinct) corners $v_1, \ldots, v_5$ and a constant number of candidate pairs by solving the system $\mathrm{len}_{u_1,v_1}(s, t) = \cdots = \mathrm{len}_{u_5,v_5}(s, t)$. We iterate this process for all possible tuples of five corners $u_1, \ldots, u_5$, to obtain a total of $O(n^7)$ candidate pairs, spending $O(n^7 \log n)$ time. Note that the other subcases with $|V_{s^*}| \leq 4$ can be handled similarly, resulting in $O(n^6)$ candidate pairs.

The validity of each candidate pair (s, t) is examined by checking if the paths from s through u_i and v_i to t are indeed shortest. For the purpose, we evaluate its geodesic distance $\mathrm{d}(s, t)$ using a two-point query structure of Chiang and Mitchell [8]. For a fixed parameter $0 < \delta \leq 1$ and any fixed $\varepsilon > 0$, one can construct, in $O(n^{5+10\delta+\varepsilon})$ time, a data structure that supports $O(n^{1-\delta} \log n)$-time two-point shortest path queries. The total running time is $O(n^7 \log n) + O(n^{5+10\delta+\varepsilon}) + O(n^7) \times O(n^{1-\delta} \log n)$. We set $\delta = \frac{3}{11}$ to optimize the running time to $O(n^{7+\frac{8}{11}+\varepsilon})$.

Also, we can use an alternative two-point query data structure whose performance is sensitive to the number h of holes [8]: after $O(n^5)$ preprocessing time using $O(n^5)$ storage, two-point queries can be answered in $O(\log n + h)$ time.[5] Using this alternative structure, the total running time of our algorithm amounts to $O(n^7(\log n + h))$. Note that this method gives a better bound than the previous one when $h = O(n^{\frac{8}{11}})$.

The other cases can be handled analogously with strictly better time bound. For Case (B-I), by Theorem 2, we have $|\mathcal{V}(s^*, t^*)| \geq 4$ and thus there are at least four distinct pairs (u_i, v_i) of corners with $\mathrm{len}_{u_i,v_i}(s^*, t^*) = \mathrm{d}(s^*, t^*)$. Here, we handle only the case of $|V_{t^*}| = 3$ or 4. For the subcase with $|V_{t^*}| = 4$, we choose any four corners from V as $v_1, \ldots, v_4$ as a candidate for V_{t^*} and overlay their shortest path maps $\mathsf{SPM}(v_i)$. The overlay, together with V, decomposes $\partial\mathcal{P}$ into $O(n)$ intervals. Each such interval determines $u_1, \ldots, u_4$ as above, and the side $e_s \in E$ on which s^* should lie. Now, we have a system of four equations in four variables: three from the corresponding path-length functions len_{u_i,v_i} with $1 \leq i \leq 4$ which should be equalized at (s^*, t^*),

[5] If h is relatively small, one could use the structure of Guo et al. [11] which answers a two-point query in $O(h \log n)$ time after $O(n^2 \log n)$ preprocessing time using $O(n^2)$ storage, or another structure by Chiang and Mitchell [8] that supports a two-point query in $O(h \log n)$ time, spending $O(n + h^5)$ preprocessing time and storage.

and the fourth from the supporting line of e_s. Solving the system, we get a constant number of candidate maximal pairs, again by Theorem 2. In total, we obtain $O(n^5)$ candidate pairs. The other subcase with $|V_{t^*}| = 3$ can be handled similarly, resulting in $O(n^4)$ candidate pairs. As above, we can exploit two different structures for two-point queries. Consequently, we can handle Case (B-I) in $O(n^{5+\frac{10}{11}+\varepsilon})$ or $O(n^5(\log n + h))$ time.

In Case (B-B) when $s^*, t^* \in \mathcal{B}$, we have $|V_{s^*}| = 2$ or 3. For the subcase with $|V_{s^*}| = 3$, we choose three corners as a candidate of V_{s^*} and take the overlay of their shortest path maps $\mathsf{SPM}(u_i)$. It decomposes $\partial\mathcal{P}$ into $O(n)$ intervals. Each such interval determines three corners v_1, v_2, v_3 forming V_{t^*} and a side $e_t \in E$ on which t^* should lie. Note that we have only three equations so far; two from the three path-length functions and the third from the line supporting to e_t. Since s^* also should lie on a side $e_s \in E$ with $e_s \neq e_t$, we need to fix such a side e_s that $\bigcap_{1\leq i\leq 3} \mathsf{VR}(u_i)$ intersects e_s. In the worst case, the number of such sides e_s is $\Theta(n)$. Thus, we have $O(n^5)$ candidate pairs for Case (B-B); again, the other subcase with $|V_{s^*}| = 2$ contributes to a smaller number $O(n^4)$ of candidate pairs. Testing each candidate pair can be done as above, resulting in $O(n^{5+\frac{10}{11}+\varepsilon})$ or $O(n^5(\log n + h))$ total running time.

Alternatively, one can exploit a two-point query structure only for boundary points on $\partial\mathcal{P}$ for Case (B-B). The two-point query structure by Bae and Okamato [5] builds an explicit representation of the graph of the lower envelope of the path-length functions $\mathrm{len}_{u,v}$ restricted on $\partial\mathcal{P} \times \partial\mathcal{P}$ in $O(n^5 \log n \log^* n)$ time.[6] Since $|\mathcal{V}(s^*, t^*)| \geq 3$ in Case (B-B), such a pair appears as a vertex on the lower envelope. Hence, we are done by traversing all the vertices of the lower envelope.

The following table summarizes the discussion so far.

Case	Independent of h	Dependent on h
(V-*)	$O(n^2 \log n)$	
(B-B)	$O(n^5 \log n \log^* n)$	$O(n^5(\log n + h))$
(B-I)	$O(n^{5+\frac{10}{11}+\varepsilon})$	$O(n^5(\log n + h))$
(I-I)	$O(n^{7+\frac{8}{11}+\varepsilon})$	$O(n^7(\log n + h))$

As Case (I-I) is the bottleneck, we conclude the following.

Theorem 3 *Given a polygonal domain having n corners and h holes, the geodesic diameter and a diametral pair can be computed in* $O(n^{7+\frac{8}{11}+\varepsilon})$ *or* $O(n^7(\log n + h))$ *time in the worst case, where* ε *is any fixed positive number.*

6 Concluding Remarks

We have presented the first algorithms that compute the geodesic diameter of a given polygonal domain. As mentioned in the introduction, a similar result for convex

[6] More precisely, in $O(n^4 \lambda_{65}(n) \log n)$ time, where $\lambda_m(n)$ stands for the maximum length of a Davenport-Schinzel sequence of order m on n symbols.

3-polytopes was shown in [17]. We note that, although the main result of this paper is similar, the techniques used in the proof are quite different. Indeed, the key requirement for our proof is the fact that shortest paths in our environment are polygonal chains whose vertices are in V, a claim that does not hold in higher dimensions (even in 2.5-D surfaces). It would be interesting to find other environments in which a similar result holds.

Another interesting question would be finding out how many maximal pairs a polygonal domain can have. The analysis of Sect. 5 gives an $O(n^7)$ upper bound. On the other hand, one can easily construct a simple polygon in which the number of maximal pairs is $\Omega(n^2)$. Any improvement on the $O(n^7)$ upper bound would lead to an improvement in the running time of our algorithm.

Though in this paper we have focused on *exact* geodesic diameters only, an efficient algorithm for finding an *approximate* geodesic diameter would be also interesting. Notice that any point $s \in \mathcal{P}$ and its farthest point $t \in \mathcal{P}$ yield a 2-approximate diameter; that is, $\operatorname{diam}(\mathcal{P}) \leq 2 \max_{t \in \mathcal{P}} \mathrm{d}(s,t)$ for any $s \in \mathcal{P}$. Also, based on a standard technique using a rectangular grid with a specified parameter $0 < \varepsilon < 1$, one can obtain a $(1+\varepsilon)$-approximate diameter in $O((\frac{n}{\varepsilon^2} + \frac{n^2}{\varepsilon}) \log n)$ time as follows. Scale $\mathcal{P}$ so that $\mathcal{P}$ can fit into a unit square, and partition $\mathcal{P}$ with a grid of size $\varepsilon^{-1} \times \varepsilon^{-1}$. We define the set D as the point set that has the center of grid squares (that are inside $\mathcal{P}$) and intersection points between boundary edges and grid segments. It is easy to see that the distance between any two points s and t in $\mathcal{P}$ is within a $(1+\varepsilon)$ factor of the distance between two points of D. Hence, for each point of D we compute its furthest neighbor (in $O(n \log n)$ time), and keep the pair of points of largest distance.[7] Since $D = O(\frac{1}{\varepsilon^2} + \frac{n}{\varepsilon})$ we obtain the aforementioned bound.

Breaking the quadratic bound in n for the $(1+\varepsilon)$-approximate diameter seems a challenge at this stage. We conclude by posing the following problem: for any or some $0 < \varepsilon < 1$, is there any algorithm that finds a $(1+\varepsilon)$-approximate diametral pair in $O(n^{2-\delta} \cdot \operatorname{poly}(1/\varepsilon))$ time for some positive $\delta > 0$?

Acknowledgements We thank Hee-Kap Ahn, Jiongxin Jin, Christian Knauer, and Joseph Mitchell for fruitful discussion. We also thank Joseph O'Rourke for pointing out the reference [21]. Work by S.W. Bae was supported by National Research Foundation of Korea (NRF) Grant funded by the Korea government (MEST) (No. 2010-0005974 and NRF-2013R1A1A1A05006927). M. Korman received the support of the Secretary for Universities and Research of the Ministry of Economy and Knowledge of the Government of Catalonia, the European Union, and ESF EUROCORES Programme EuroGIGA-ComPoSe IP04-MICINN Project EUI-EURC-2011-4306. Work by Y. Okamoto was supported by Global COE Program "Computationism as a Foundation for the Sciences" and Grant-in-Aid for Scientific Research from Ministry of Education, Science and Culture, Japan, and Japan Society for the Promotion of Science.

Appendix: 1 More Examples and Remarks

In this section, we show more constructions of polygonal domains and their diametral pairs with remarks. In the figures, we keep the following rules: the boundary $\partial\mathcal{P}$ is depicted by dark gray segments and the interior of holes by light gray region. A

[7] The idea of this approximation algorithm is due to Hee-Kap Ahn.

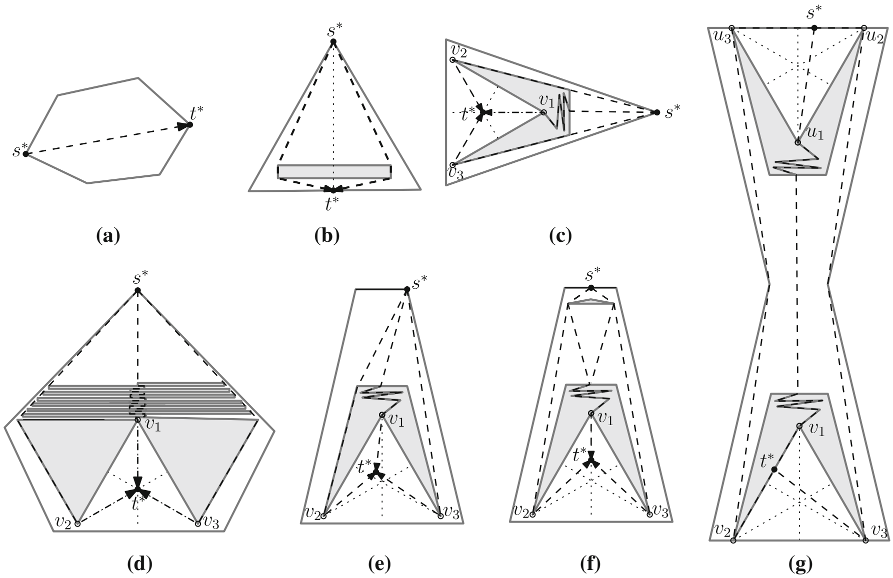

Fig. 5 **a**–**c** Polygonal domains whose geodesic diameter is determined by a corner s^* and (**d**–**g**) variations of the construction (**c**). **a** When both s^* and t^* are corners; **b** when t^* is a point on $\partial\mathcal{P}$; **c** when $t^* \in \mathrm{int}\,\mathcal{P}$. This polygonal domain consists of two holes, forming a narrow corridor and three shortest paths between s^* and t^*. Here, we have $\mathrm{d}(s^*, v_1) = \mathrm{d}(s^*, v_2) = \mathrm{d}(s^*, v_3)$ and t^* is indeed the vertex of $\mathsf{SPM}(s^*)$ defined by v_1, v_2, v_3; **d** variation of (c) with all convex holes; **e** three shortest paths are not enough to determine a boundary-interior diametral pair; **f** if we add one more hole, then the diameter is determined by $s^* \in \mathcal{B}$ and $t^* \in \mathrm{int}\,\mathcal{P}$ with four shortest paths; **g** a polygonal domain made by attaching two copies of (**e**) and modifying it to have $\mathrm{d}(u_1, v_1) = \mathrm{d}(u_2, v_2) = \mathrm{d}(u_3, v_3)$. Observe that, in this polygonal domain, the diameter is determined by two boundary points with three shortest paths.

diametral pair is given as (s^*, t^*) and shortest paths between s^* and t^* are described as black dashed polygonal chains.

1(a) Examples Where at Least One Point of a Diametral Pair Lies on $\partial\mathcal{P}$

Note that, as expected, every example in Fig. 5 obeys Theorem 2. An interesting construction is Fig. 5g, where neither of the two centers of $\triangle u_1u_2u_3$ and of $\triangle v_1v_2v_3$ appears in any diametral pair. Also note that Fig. 5d consists of *convex* holes only. We think that any complicated construction can be "convexified" in a similar fashion. This would suggest that computing the diameter in polygonal domains with convex holes only might be as difficult as the general case.

1(b) A Proof for Fig. 1c: Case (I-I) with 6 Shortest Paths

Claim 1 *In the polygonal domain described in Fig.* (1)c*,* (s^*, t^*) *is the unique diametral pair.*

Proof of Claim Recall that by construction of the problem instance, the triangles $\triangle u_1u_2u_3$ and $\triangle v_1v_2v_3$ are regular and $\mathrm{d}(u_1, v_1) = \mathrm{d}(u_1, v_2) = \mathrm{d}(u_2, v_2) = \mathrm{d}(u_2, v_3)$

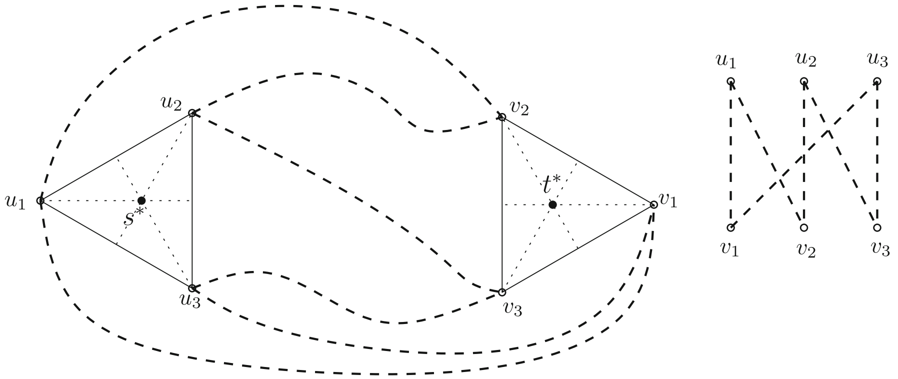

Fig. 6 A schematic diagram corresponding to the polygonal domain shown in Fig. 1c.

$= \mathrm{d}(u_3, v_3) = \mathrm{d}(u_3, v_1) = L$, for some arbitrarily large value $L > 0$. Also, s^* and t^* are the centers of $\triangle u_1u_2u_3$ and $\triangle v_1v_2v_3$, respectively.

We assume that both triangles $\triangle u_1u_2u_3$ and $\triangle v_1v_2v_3$ are inscribed in a unit circle (and thus $\mathrm{d}(s^*, t^*) = 2 + L$). For any point s on any shortest path between u_i and v_j, it is easy to see that $\mathrm{d}(s, t) \le \sqrt{3} + L < \mathrm{d}(s^*, t^*)$ for every point $t \in \mathcal{P}$. In particular, no point on those paths can contribute to the diameter (see Fig. 6).

(1) First, observe that $\max_{t\in\triangle v_1v_2v_3} \mathrm{d}(s^*, t) = \max_{s\in\triangle u_1u_2u_3} \mathrm{d}(s, t^*) = \mathrm{d}(s^*, t^*)$.

(2) For any $s \in \triangle u_1u_2u_3$, its farthest point $t \in \triangle v_1v_2v_3$ is on the angle bisector of some v_i. Consider any $s \in \triangle u_1u_2u_3$. Without loss of generality we assume that $\|s - u_1\| \le \min_i\{\|s - u_i\|\}$. Both shortest paths to v_1 and to v_2 from s pass through u_1. We have $\mathrm{d}(s, v_1) = \mathrm{d}(s, v_2)$ by construction and its farthest point $t \in \triangle v_1v_2v_3$ must be in the angle bisector of v_3. By symmetry, the same property holds when the closest corner from s^* is either u_2 or u_3.
Conversely, for any t, its farthest point $s \in \triangle u_1u_2u_3$ must be on a bisector of some u_i. In any diametral pair (s, t), we have that t is the farthest point of s (and vice versa), so both must be on one of the angle bisectors.

(3) If (s, t) is a diametral pair, then $s \in \overline{u_is^*}$ and $t \in \overline{v_jt^*}$ for some i and j. Suppose that s lies on the bisector of u_1 but not in between u_1 and s^*. We then have $\|s-u_2\| = \|s-u_3\| < \|s-u_1\|$ and $\mathrm{d}(s, v_1) = \mathrm{d}(s, v_2) = \mathrm{d}(s, v_3) = \|s-u_2\|+L$ by construction. This implies that t^* is the farthest point of such s. Since $\|s-u_2\| < 1$ and $\mathrm{d}(s, t^*) < 2 + L$, $\overline{(s, t^*)}$ is not a diametral pair.

(4) Now, pick any point $s \in \overline{u_1s^*}$ with $s \ne s^*$. Suppose that $t \in \triangle v_1v_2v_3$ is the farthest point from s. We know that $t \in \overline{v_3t^*}$ by above discussions. In this case, we have four shortest paths between s and t through (u_1, v_1), (u_1, v_2), (u_2, v_3), and (u_3, v_3); the other two are strictly longer unless $s = s^*$. By Theorem 2, such $s \in \overline{u_1s^*}$ with $s \ne s^*$ and its farthest point t cannot form a maximal pair. By symmetry, the other cases where $s \in \overline{u_is^*}$ can be handled.

Hence, (s^*, t^*) is a unique diametral pair and the geodesic diameter is $2 + L$.

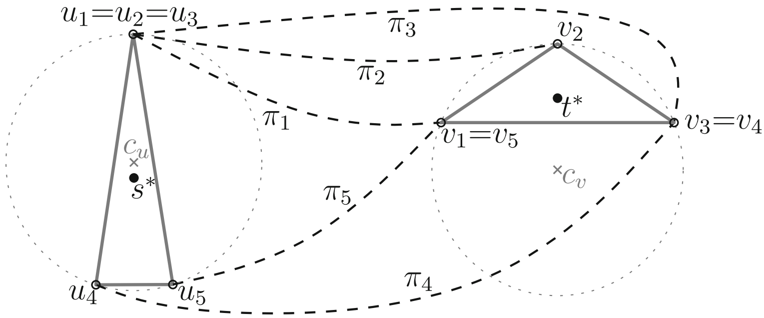

Fig. 7 A schematic diagram of a polygonal domain in which $|V_{s^*}| = |V_{t^*}| = 3$ and $|\Pi(s^*, t^*)| = 5$.

1(c) Diametral Pair of Case (I-I) with Exactly 5 Shortest Paths

Here, we present a polygonal domain in which the diameter is determined by two interior points and exactly five shortest paths between them. This proves the tightness of Case (I-I) in Theorem 2.

Figure 7 shows a schematic description of a polygonal domain $\mathcal{P}$. We assume that only the position of the vertices u_i and the v_i are geometrically precise. We construct the problem instance such that we have $u_1 = u_2 = u_3$, $v_1 = v_5$, and $v_3 = v_4$, and the convex hulls of the u_i and of the v_i form isosceles triangles $\triangle_u$ and $\triangle_v$. Each of $\triangle_u$ and $\triangle_v$ is inscribed in a unit circle centered at c_u and c_v. Moreover, the bases of both triangles are horizontal and the angles opposite to the bases are 18° and 112°, respectively. Note that the side lengths of the triangles $\triangle_u$ and $\triangle_v$ are as follows: $\|u_1 - u_4\| = 1.97537\cdots$ and $\|u_4 - u_5\| = 0.61803\cdots$; $\|v_2 - v_1\| = 1.11833\cdots$ and $\|v_1 - v_3\| = 1.85436\cdots$.

In this configuration, we set the constants as follows: letting $L := \mathrm{d}(u_1, v_1) = \mathrm{d}(u_3, v_3)$ be some sufficiently large number, we set $\mathrm{d}(u_2, v_2) = L + 0.5$ and $\mathrm{d}(u_4, v_4) = \mathrm{d}(u_5, v_5) = L + 0.2$. Note that this configuration can be realized with four obstacles in a similar way as Fig. 1c.

Since we have fixed all necessary parameters, we have a fully explicit description of the len_{u_i,v_i}. Due to the difficulty of finding an exact analytical solution, we used numerical methods to solve the system of equations $\mathrm{len}_{u_1,v_1}(s,t) = \cdots = \mathrm{len}_{u_5,v_5}(s,t)$. We have found that there is a unique solution (s^*, t^*) such that $s^* \in \triangle_u$ and $t^* \in \triangle_v$; we obtained $s^* = c_u + (0, -0.102795\cdots)$, $t^* = c_v + (0, 0.555361\cdots)$ and $\mathrm{d}(s^*, t^*) = 2.047433734\cdots + L$ (see Fig. 7).

We first checked that (s^*, t^*) is a maximal pair based on the following lemma, which can be shown using elementary linear algebra together with the convexity of the path-length functions.

Lemma 6 *Suppose that (s,t) is a solution to the system $len_{u_1,v_1}(s,t) = \cdots = len_{u_5,v_5}(s,t)$. If any four of the five gradients $\nabla len_{u_i,v_i}$ at (s,t) are linearly independent (as vectors in a 4-dimensional space) and one of them is represented as a linear combination of the other four with all "negative" coefficients, then (s,t) is a local maximum of the pointwise minimum of the five functions len_{u_i,v_i}.*

Next, to see that (s^*, t^*) is a diametral pair, we have run our algorithm for each of Cases (B-B), (B-I), and (I-I); as a result, there are 44 candidate pairs, including (s^*, t^*), falling into those cases among which at most 11 are maximal and only (s^*, t^*) is diametral. Note that the pair (s^*, t^*) is the only candidate pair of Case (I-I). Also, observe that any point on the shortest path between u_i and v_i cannot belong to a diametral pair. This implies that none of the u_i and the v_i belongs to a diametral pair. In particular, we have that none of the Cases (V-*) can happen. In addition, we also sampled about 350,000 points uniformly from each of $\triangle_u$ and $\triangle_v$, and evaluated the geodesic distances of the $350{,}000^2$ pairs.

Note that one can modify the construction to have $|V_{s^*}| = |V_{t^*}| = |\Pi(s^*, t^*)| = 5$. For the purpose, we can split u_1, u_2, u_3 into three close corners (analogously for corners, v_1, v_5 and v_3, v_4). The splitting process should preserve the differences between the distances $\mathrm{d}(u_i, v_i)$ for all $i = 1, \ldots, 5$ (and increase other distances). We have tested such an example in the same way as above and concluded that a solution equalizing the five path-length functions is indeed a diametral pair.

References

1. Agarwal, P.K., Aronov, B., O'Rourke, J., Schevon, C.A.: Star unfolding of a polytope with applications. SIAM J. Comput. **26**(6), 1689–1713 (1997)
2. Aronov, B., Fortune, S., Wilfong, G.: The furthest-site geodesic Voronoi diagram. Discrete Comput. Geom. **9**, 217–255 (1993)
3. Asano, T., Toussaint, G.: Computing the geodesic center of a simple polygon. Technical Report SOCS-85.32, McGill University, Montreal (1985)
4. Bae, S.W., Chwa, K.-Y.: The geodesic farthest-site Voronoi diagram in a polygonal domain with holes. In: Proceedings of 25th Annual Symposium Computational Geometry (SoCG), pp. 198–207 (2009)
5. Bae, S.W., Okamoto, Y.: Querying two boundary points for shortest paths in a polygonal domain. Comput. Geom. **45**(7), 284–293 (2012)
6. Bae, S.W., Korman, M., Okamoto, Y.: The geodesic diameter of polygonal domains. In: Proceedings of 18th Annual European Symposium on Algorithms. Part 1. Lecture Notes in Computer Science, vol. 6346, pp. 500–511 (2010)
7. Chazelle, B.: A theorem on polygon cutting with applications. In: Proceedings of 23rd Annual Symposium on Foundations of Computer Science (FOCS), pp. 339–349 (1982)
8. Chiang, Y.-J., Mitchell, J.S.B.: Two-point Euclidean shortest path queries in the plane. In: Proceedings of 10th ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 215–224 (1999)
9. Cook, A.F., IV, Wenk, C.: Shortest path problems on a polyhedral surface. In: Proceedings of 11th International Symposium on Algorithms and Data Structures (WADS), pp. 156–167 (2009)
10. Guibas, L.J., Hershberger, J.: Optimal shortest path queries in a simple polygon. J. Comput. Syst. Sci. **39**(2), 126–152 (1989)
11. Guo, H., Maheshwari, A., Sack, J.-R.: Shortest path queries in polygonal domains. In: Proceedings of the 4th International Conference on Algorithmic Aspects in Information and Management (AAIM). Lecture Notes in Computer Science, vol. 5034, pp. 200–211 (2008)
12. Hershberger, J., Suri, S.: Matrix searching with the shortest path metric. SIAM J. Comput. **26**(6), 1612–1634 (1997)
13. Hershberger, J., Suri, S.: An optimal algorithm for Euclidean shortest paths in the plane. SIAM J. Comput. **28**(6), 2215–2256 (1999)
14. Koivisto, M., Polishchuk, V.: Geodesic diameter of a polygonal domain in $O(n^4 \log n)$ time. CoRR, http://arxiv/abs/1006.1998 (2010)
15. Mitchell, J.S.B.: Shortest paths among obstacles in the plane. Int. J. Comput. Geom. Appl. **6**(3), 309–331 (1996)
16. Mitchell, J.S.B.: Shortest paths and networks. In: Handbook of Discrete and Computational Geometry, 2nd edn, Chap. 27, pp. 607–641. CRC Press, Boca Raton (2004)

17. O'Rourke, J., Schevon, C.: Computing the geodesic diameter of a 3-polytope. In: Proceedings of the 5th Annual Symposium on Computational Geometry (SoCG), pp. 370–379 (1989)
18. O'Rourke, J., Suri, S.: Polygons. In: Handbook of Discrete and Computational Geometry, 2nd edn, Chap. 26, pp. 583–606. CRC Press, Boca Raton (2004)
19. Pollack, R., Sharir, M., Rote, G.: Computing the geodesic center of a simple polygon. Discrete Comput. Geom. **4**(6), 611–626 (1989)
20. Suri, S.: The all-geodesic-furthest neighbors problem for simple polygons. In: Proceedings of 3rd Annual Symposium on Computational Geometry (SoCG), pp. 64–75 (1987)
21. Zalgaller, V.A.: An isoperimetric problem for tetrahedra. J. Math. Sci. **140**(4), 511–527 (2007)