The geodesic diameter of polygonal domains

Sang Won Bae* Matias Korman† Yoshio Okamoto‡

Abstract

This paper studies the geodesic diameter of polygonal domains having h holes and n corners. For simple polygons (i.e., $h = 0$), it is known that the geodesic diameter is determined by a pair of corners of a given polygon and can be computed in linear time. For general polygonal domains with $h \geq 1$, however, no algorithm for computing the geodesic diameter was known prior to this paper. We present 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 algorithms are based on new geometric observations, part of which states as follows: the geodesic diameter of a polygonal domain can be determined by two points in its interior, and in that case there are at least five shortest paths between the two points.

1 Introduction

In this paper, we address the geodesic diameter problem in polygonal domains. The geodesic distance $\mathrm{d}(p, q)$ between any two points p, q in a polygonal domain $\mathcal{P}$ is defined as the (Euclidean) length of a shortest obstacle-avoiding path between p and q. The *geodesic diameter* $\mathrm{diam}(\mathcal{P})$ of a polygonal domain $\mathcal{P}$ is defined as $\mathrm{diam}(\mathcal{P}) := \max_{s,t\in\mathcal{P}} \mathrm{d}(s, t)$. A pair (s, t) of points in $\mathcal{P}$ that realizes the geodesic diameter $\mathrm{diam}(\mathcal{P})$ is called a *diametral pair*. The geodesic diameter problem is to find the value of $\mathrm{diam}(\mathcal{P})$ and a diametral pair.

For simple polygons (i.e., $h = 0$), the geodesic diameter has been extensively studied and fully understood. Chazelle [2] provided the first $O(n^2)$-time algorithm computing the geodesic diameter of a simple polygon, and Suri [9] 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 [5] showed that the diameter can be computed in linear time using their fast matrix search technique. On the other hand, to the best of our knowledge, no algorithm for computing $\mathrm{diam}(\mathcal{P})$ has yet been discovered when $\mathcal{P}$ is a polygonal domain having one or more holes ($h \geq 1$).

This fairly wide gap between simple polygons and polygonal domains is seemingly due to the uniqueness of the shortest path between any two points; it is well known that there is a unique shortest path between any two points in a simple polygon [4]. Using this uniqueness, one can show that the diameter is indeed realized by a pair of corners in V; that is, $\mathrm{diam}(\mathcal{P}) = \max_{u,v\in V} \mathrm{d}(u, v)$ if $h = 0$ [5, 9]. For general polygonal domains with $h \geq 1$, however, this is not the case. In this paper, we exhibit several examples where the diameter is realized by non-corner points on $\partial\mathcal{P}$ or even by interior points of $\mathcal{P}$ (see Figure 1). This observation also shows an immediate difficulty in devising any exhaustive algorithm since the search space like $\partial\mathcal{P}$ or the whole domain $\mathcal{P}$ is not discrete.

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. We also show that for small constant h the diameter can be computed much faster.

2 Preliminaries

We are given as input a polygonal domain $\mathcal{P}$ with h holes and n corners. More precisely, $\mathcal{P}$ consists of an outer simple polygon in the plane $\mathbb{R}^2$ and a set of h (≥ 0) disjoint simple polygons inside the outer polygon. As a subset of $\mathbb{R}^2$, $\mathcal{P}$ is the region contained in its outer polygon *excluding* the interior of the holes; thus $\mathcal{P}$ is a bounded, closed subset of $\mathbb{R}^2$. The boundary $\partial\mathcal{P}$ of $\mathcal{P}$ is regarded as a series of *obstacles* so that any feasible path inside $\mathcal{P}$ is not allowed to cross $\partial\mathcal{P}$. Note that some portion or the whole of a feasible path may go along the boundary $\partial\mathcal{P}$. The *length* of a path is the sum of the Euclidean lengths of its segments. It is well known from earlier work that there always exists a *shortest (feasible) path* between any two points $p, q \in \mathcal{P}$ [7]. The *geodesic distance*, denoted by $\mathrm{d}(p, q)$, is then defined to be the length of a shortest path between $p \in \mathcal{P}$ and $q \in \mathcal{P}$.

*Department of Computer Science and Engineering, POSTECH, swbae@postech.ac.kr. Work supported by the Brain Korea 21 Project.

†Computer Science Department, Université Libre de Bruxelles, mkormanc@ulb.ac.be. Work supported by the Communauté française de Belgique - Actions de Recherche Concertées (ARC).

‡Graduate School of Information Science and Engineering, Tokyo Institute of Technology, okamoto@is.titech.ac.jp. Work 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

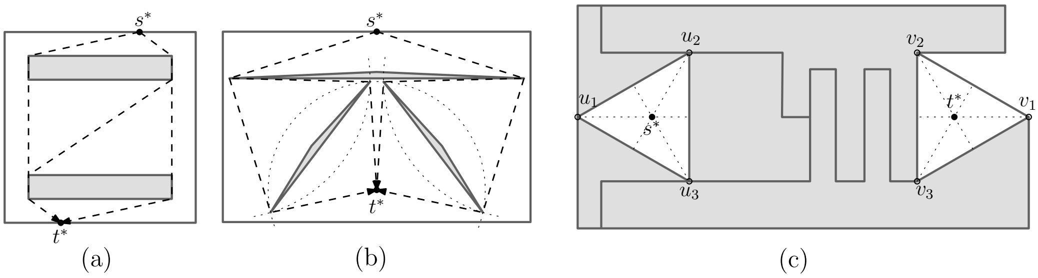

Figure 1: Three polygonal domains where the geodesic diameter is determined by a pair (s^*, t^*) of non-corner points; Gray-shaded regions depict the interior of the holes and dark gray segments depict the boundary $\partial\mathcal{P}$. Recall that $\mathcal{P}$, as a set, contains its boundary $\partial\mathcal{P}$. (a) Both s^* and t^* lie on $\partial\mathcal{P}$. There are three shortest paths between s^* and t^*. In this polygonal domain, there are two (symmetric) diametral pairs. (b) $s^* \in \partial\mathcal{P} \setminus V$ and $t^* \in \mathrm{int}\mathcal{P}$. Three triangular holes are placed in a symmetric way. There are four shortest paths between s^* and t^*. (c) Both s^* and t^* lie in the interior $\mathrm{int}\mathcal{P}$. Here, the five holes are packed like jigsaw puzzle pieces, forming narrow corridors (dark gray paths) and two empty, regular triangles. Observe that $\mathrm{d}(u_1, v_1) = \mathrm{d}(u_1, v_2) = \mathrm{d}(u_2, v_2) = \mathrm{d}(u_2, v_3) = \mathrm{d}(u_3, v_3) = \mathrm{d}(u_3, v_1)$. 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^*. More details on this example can be found in the extended version of this paper [1].

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 $t \in \mathcal{P}$. Then, it 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 [7]. Note that possibly we may have $k = 0$ when $\mathrm{d}(s, t) = \|s - t\|$. If two paths (with possibly different endpoints) induce the same sequence of corners, 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 $v \in V$ or s itself, which is the last corner of $\pi(s, t)$ for any t in the cell $\sigma_s(v)$. In particular, the cell $\sigma_s(s)$ is the set of points t such that $\pi(s, t)$ passes through no corner in V and thus $\mathrm{d}(s, t) = \|s - t\|$. Each edge of $\mathsf{SPM}(s)$ is an arc on the boundary of two incident cells $\sigma_s(v_1)$ and $\sigma_s(v_2)$ and thus determined by two corners $v_1, v_2 \in V \cup \{s\}$. Similarly, each vertex of $\mathsf{SPM}(s)$ is determined by at least three corners $v_1, v_2, v_3 \in V \cup \{s\}$. Note that for fixed $s \in \mathcal{P}$ a point t that locally maximizes $\mathrm{d}_s(t) := \mathrm{d}(s, t)$ 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 $O(n)$ complexity can be computed in $O(n \log n)$ time using $O(n \log n)$ working space [6]. For more details on shortest path maps, see [7, 6, 8].

Path-length function. If $\pi(s, t) \neq \overline{st}$, then there are two corners $u, v \in V$ such that $\pi(s, t)$ is formed as the union of a shortest path from u to v and two segments $\overline{su}$ and $\overline{vt}$. 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; thus, $s \in \mathsf{VP}(u)$ and $t \in \mathsf{VP}(v)$, where $\mathsf{VP}(p)$ for any $p \in \mathcal{P}$ is defined to be the set of all points $q \in \mathcal{P}$ such that $\overline{pq} \subset \mathcal{P}$. The set $\mathsf{VP}(p)$ is also called the *visibility profile* of $p \in \mathcal{P}$ [3].

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

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

Then, $\mathrm{len}_{u,v}(s, t)$ represents the length of the path from s to t that has the fixed combinatorial structure, entering u from s and exiting v to t. Also, unless $\mathrm{d}(s, t) = \|s - t\|$ (equivalently, $s \in \mathsf{VP}(t)$), the geodesic distance $\mathrm{d}(s, t)$ can be expressed as the pointwise minimum of some path-length functions:

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

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.

3 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 $\mathrm{d}(s^*, t^*) \geq \mathrm{d}(s, t)$. For any pair (s, t), let $\Pi(s, t) = \{\pi_1, \ldots, \pi_m\}$ be the set of all distinct shortest paths from s to t, where m denotes the

(VV)	$s^* \in V$,	$t^* \in V$	implies	$\lvert\Pi(s^*,t^*)\rvert \geq 1, \lvert V_{s^*}\rvert \geq 1, \lvert V_{t^*}\rvert \geq 1$;
(VB)	$s^* \in V$,	$t^* \in \mathcal{B}$	implies	$\lvert\Pi(s^*,t^*)\rvert \geq 2, \lvert V_{s^*}\rvert \geq 1, \lvert V_{t^*}\rvert \geq 2$;
(VI)	$s^* \in V$,	$t^* \in \text{int}\mathcal{P}$	implies	$\lvert\Pi(s^*,t^*)\rvert \geq 3, \lvert V_{s^*}\rvert \geq 1, \lvert V_{t^*}\rvert \geq 3$;
(BB)	$s^* \in \mathcal{B}$,	$t^* \in \mathcal{B}$	implies	$\lvert\Pi(s^*,t^*)\rvert \geq 3, \lvert V_{s^*}\rvert \geq 2, \lvert V_{t^*}\rvert \geq 2$;
(BI)	$s^* \in \mathcal{B}$,	$t^* \in \text{int}\mathcal{P}$	implies	$\lvert\Pi(s^*,t^*)\rvert \geq 4, \lvert V_{s^*}\rvert \geq 2, \lvert V_{t^*}\rvert \geq 3$;
(II)	$s^* \in \text{int}\mathcal{P}$,	$t^* \in \text{int}\mathcal{P}$	implies	$\lvert\Pi(s^*,t^*)\rvert \geq 5, \lvert V_{s^*}\rvert \geq 3, \lvert V_{t^*}\rvert \geq 3$.

Figure 2: Necessary conditions for a pair of points to be maximal.

.

number of shortest paths. Let u_i and v_i be the first and the last corners in V along π_i from s to t, and let $V_s := \{u_1, \ldots, u_m\}$ and $V_t := \{v_1, \ldots, v_m\}$.

Let E be the set of all sides of $\mathcal{P}$ without their endpoints and $\mathcal{B}$ be their union. Note that $\mathcal{B} = \partial\mathcal{P} \setminus V$, the boundary of $\mathcal{P}$ except the corners V.

Theorem 1 *Suppose that (s^*, t^*) is a maximal pair in $\mathcal{P}$ and $\Pi(s^*, t^*)$, V_{s^*}, and V_{t^*} be defined as above. The implications of Figure 2 hold. Moreover, each of the above bounds is best possible by examples.*

Due to space constraints proofs of this theorem is omitted (and can be found in the extended version [1]).

4 Computing the Geodesic Diameter

Since a diametral pair is in fact maximal, it falls into one of the cases shown in Theorem 1. 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 Section 2. We thus focus on Cases **(BB)**, **(BI)**, and **(II)**, where a diametral pair consists of two non-corner points.

From the computational point of view, the most difficult case corresponds to Case **(II)** of Theorem 1; in particular, the case in which $|\Pi(s^*, t^*)| = |V_{s^*}| = |V_{t^*}| = 5$. For such a case we do the following: 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. Then, 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) $v_1, \ldots, v_5$ and thus a constant number of candidate pairs by solving the system $\text{len}_{u_1,v_1}(s,t) = \cdots = \text{len}_{u_5,v_5}(s,t)$. We iterate this process for all possible tuples of five corners $u_1, \ldots, u_5$, obtaining a total of $O(n^7)$ candidate pairs in $O(n^7 \log n)$ time. Note that this method also covers the case of $|\Pi(s^*, t^*)| > 5$. Recall that each path-length function $\text{len}_{u,v}$ 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[1]. Indeed when five distinct pairs $(u_1, v_1), \ldots, (u_5, v_5)$ of corners in V such that $\text{len}_{u_i,v_i}(s^*, t^*) = \text{d}(s^*, t^*)$ for any $i \in \{1, \ldots, 5\}$ are known, their system of equations $\text{len}_{u_1,v_1}(s,t) = \cdots = \text{len}_{u_5,v_5}(s,t)$ determines a 0-dimensional zero set corresponding to a constant number of candidate pairs in $\text{int}\mathcal{P} \times \text{int}\mathcal{P}$. The **(II)** case (in which $|V_{s^*}| \leq 4$) can be handled similarly, resulting in $O(n^6)$ candidate pairs.

In order to test the validity of each candidate pair (s,t), we check the geodesic distance $\text{d}(s,t)$ using a two-point query structure of Chiang and Mitchell [3]: for a fixed parameter $0 < \delta \leq 1$ and any fixed $\epsilon > 0$, we can construct, in $O(n^{5+10\delta+\epsilon})$ time, a data structure that supports $O(n^{1-\delta} \log n)$-time two-point shortest path queries. Then, the total running time is $O(n^7 \log n) + O(n^{5+10\delta+\epsilon}) + O(n^7) \times O(n^{1-\delta} \log n)$. We set $\delta = \frac{3}{11}$ to optimize the running time to $O(n^{7+\frac{8}{11}+\epsilon})$.

Also, we can use an alternative two-point query data structure whose performance is sensitive to the number h of holes [3]: after $O(n^5)$ preprocessing time using $O(n^5)$ storage, two-point queries can be answered in $O(\log n + h)$ time.Using this alternative structure, the total running time of our algorithm becomes $O(n^7(\log n + h))$. Note that this method outperforms the previous one when $h = O(n^{\frac{8}{11}})$.

For Case **(BI)**, we handle only the case of $|\Pi(s^*, t^*)| = 4$ with $|V_{t^*}| = 3$ 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. Then, 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

[1] 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.

have a system of four equations on four variables: three from the corresponding path-length functions len_{u_i,v_i} which should be equalized at (s^*, t^*) 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 1 and its proof. 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 **(BI)** in $O(n^{5+\frac{10}{11}+\epsilon})$ or $O(n^5(\log n + h))$ time.

In Case **(BB)** when $s^*, t^* \in \mathcal{B}$, we handle the case of $|\Pi(s^*, t^*)| = 3$ with $|V_{s^*}| = 2$ 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. Then, 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 $\bigcap_{1 \le i \le 3} \mathsf{VP}(u_i)$ 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 **(BB)**; 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 performed as above, resulting in $O(n^{5+\frac{10}{11}+\epsilon})$ or $O(n^5(\log n + h))$ total running time.

As Case **(II)** being a bottleneck, we conclude the following.

Theorem 2 *Given a polygonal domain having n corners and h holes, the geodesic diameter and a diametral pair can be computed in $O(n^{7+\frac{8}{11}+\epsilon})$ or $O(n^7(\log n + h))$ time in the worst case, where ϵ is any fixed positive number.*

We can avoid some difficult cases when h is a small constant based on a simple observation: if there are two distinct shortest paths between s and t in $\mathcal{P}$, then we know that there is at least one hole in the region closed by the two paths. In general, if $h < k-1$, there cannot exist two points that have k or more distinct shortest paths between them.

Theorem 3 *Given a polygonal domain having n corners and h holes, the geodesic diameter and a diametral pair can be computed in the following worst-case time bound, depending on h.*

- *$O(n)$ time, if $h = 0$ (by Hershberger and Suri [5]),*
- *$O(n^2 \log n)$ time, if $h = 1$,*
- *$O(n^5 \log n)$ time, if $h = 2$ or 3,*
- *$O(n^7(\log n + h))$ time, if $4 \le h = O(n^{\frac{8}{11}})$,*
- *$O(n^{7+\frac{8}{11}+\epsilon})$ time, otherwise.*

5 Concluding Remark

It is worth noting that with analysis in Section 4 the number of geodesic-maximal pairs is shown to be at most $O(n^7)$. On the other hand, one can easily construct a simple polygon in which the number of maximal pairs is $\Omega(n^2)$. An interesting question is how many maximal pairs are there in a polygonal domain in the worst case.

Though we, in this paper, have focused on exact diameters only, an efficient algorithm for finding an approximate diameter would be interesting. Notice that any point $s \in \mathcal{P}$ and its farthest point $t \in \mathcal{P}$ yield a $\frac{1}{2}$-approximate diameter; that is, $\mathrm{d}(s, t) \ge \frac{1}{2}\mathrm{diam}(\mathcal{P})$. Also, based on a standard technique using a rectangular grid with a specified parameter $0 < \epsilon < 1/2$, one can easily obtain a $(1 - \epsilon)$-approximate diameter in $O((\frac{n}{\epsilon^2} + \frac{n^2}{\epsilon}) \log n)$ time. However, breaking the quadratic bound (in n) for the $(1 - \epsilon)$-approximate diameter seems a challenge at this stage.

Acknowledgments

The authors thank Hee-Kap Ahn, Jiongxin Jin, Christian Knauer, and Joe Mitchell for fruitful discussion.

References

[1] S.W. Bae, M. Korman and Y. Okamoto The geodesic diameter of polygonal domains. *CoRR*, abs/1001.0695, 2010.

[2] B. Chazelle. A theorem on polygon cutting with applications. In *Proc. Annu. Sympos. Found. Comput. Sci. (FOCS)*, pages 339–349, 1982.

[3] Y.-J. Chiang and J. S. B. Mitchell. Two-point Euclidean shortest path queries in the plane. In *Proc. 10th ACM-SIAM Sympos. Discrete Algorithms (SODA)*, pages 215–224, 1999.

[4] L. J. Guibas and J. Hershberger. Optimal shortest path queries in a simple polygon. *J. Comput. Syst. Sci.*, 39(2):126–152, 1989.

[5] J. Hershberger and S. Suri. Matrix searching with the shortest path metric. *SIAM J. Comput.*, 26(6):1612–1634, 1997.

[6] J. Hershberger and S. Suri. An optimal algorithm for Euclidean shortest paths in the plane. *SIAM J. Comput.*, 28(6):2215–2256, 1999.

[7] J. S. B. Mitchell. Shortest paths among obstacles in the plane. *Internat. J. Comput. Geom. Appl.*, 6(3):309–331, 1996.

[8] J. S. B. Mitchell. Shortest paths and networks. In *Handbook of Discrete and Computational Geometry*, chapter 27, pages 607–641. CRC Press, Inc., 2nd edition, 2004.

[9] S. Suri. The all-geodesic-furthest neighbors problem for simple polygons. In *Proc. 3rd Annu. Sympos. Comput. Geom. (SoCG)*, page 64, 1987.