

A linear-time algorithm for the geodesic center of a simple polygon

Hee-Kap Ahn* Luis Barba^{†,‡} Prosenjit Bose[†] Jean-Lou De Carufel[†]
Matias Korman[§] Eunjin Oh*

August 13, 2015

Abstract

Let P be a closed simple polygon with n vertices. For any two points in P , the geodesic distance between them is the length of the shortest path that connects them among all paths contained in P . The geodesic center of P is the unique point in P that minimizes the largest geodesic distance to all other points of P . In 1989, Pollack, Sharir and Rote [Disc. & Comput. Geom. 89] showed an $O(n \log n)$ -time algorithm that computes the geodesic center of P . Since then, a longstanding question has been whether this running time can be improved. In this paper we affirmatively answer this question and present a linear time algorithm to solve this problem.

1 Introduction

Let P be a simple polygon with n vertices. Given two points x, y in P (either on the boundary or in the interior), the *geodesic path* $\pi(x, y)$ is the shortest path contained in P connecting x with y . If the straight-line segment connecting x with y is contained in P , then $\pi(x, y)$ is a straight-line segment. Otherwise, $\pi(x, y)$ is a polygonal chain whose vertices (other than its endpoints) are reflex vertices of P . We refer the reader to [22] for more information on geodesic paths.

The *geodesic distance* between x and y , denoted by $|\pi(x, y)|$, is the sum of the Euclidean lengths of each segment in $\pi(x, y)$. Throughout this paper, when referring to the distance between two points in P , we mean the geodesic distance between them.

Given a point $x \in P$, a (geodesic) *farthest neighbor* of x , is a point $f_P(x)$ (or simply $f(x)$) of P whose geodesic distance to x is maximized. To ease the description, we assume that each vertex of P has a unique farthest neighbor. This *general position* condition was also assumed by Aronov et al. [2] and can be obtained by applying a slight perturbation to the positions of the vertices [10].

Let $F_P(x)$ be the function that maps each $x \in P$ to the distance to a farthest neighbor of x (i.e., $F_P(x) = |\pi(x, f(x))|$). A point $c_P \in P$ that minimizes $F_P(x)$ is called the *geodesic center*

*Department of Computer Science and Engineering, POSTECH, 77 Cheongam-Ro, Nam-Gu, Pohang, Gyeongbuk, Korea. {heekap@postech.ac.kr, jin9082@postech.ac.kr}. Supported by the NRF grant 2011-0030044 (SRC-GAIA) funded by the Korea government (MSIP).

[†]School of Computer Science, Carleton University, Ottawa, Canada. jit@scs.carleton.ca, jdecaruf@csg.scs.carleton.ca

[‡]Département d'Informatique, Université Libre de Bruxelles, Brussels, Belgium. lbarbaf1@ulb.ac.be

[§]Tohoku University, mati@dais.is.tohoku.ac.jp. Partially supported by the ELC project (MEXT KAKENHI No. 24106008).

31 of P . Similarly, a point $s \in P$ that maximizes $F_P(x)$ (together with $f(s)$) is called a *geodesic*
 32 *diametral pair* and their distance is known as the *geodesic diameter*. Asano and Toussaint [3]
 33 showed that the geodesic center is unique (whereas it is easy to see that several geodesic diametral
 34 pairs may exist).

35 In this paper we show how to compute the geodesic center of P in $O(n)$ time.

36 1.1 Previous Work

37 Since the early 80s the problem of computing the geodesic center (and its counterpart, the
 38 geodesic diameter) has received a lot of attention from the computational geometry community.
 39 Chazelle [7] gave the first algorithm for computing the geodesic diameter (which runs in $O(n^2)$
 40 time using linear space). Afterwards, Suri [27] reduced it to $O(n \log n)$ -time without increasing
 41 the space constraints. Finally, Hershberger and Suri [15] presented a fast matrix search technique,
 42 one application of which is a linear-time algorithm for computing the diameter.

43 The first algorithm for computing the geodesic center was given by Asano and Toussaint [3],
 44 and runs in $O(n^4 \log n)$ -time. In 1989, Pollack, Sharir, and Rote [25] improved it to $O(n \log n)$
 45 time. Since then, it has been an open problem whether the geodesic center can be computed
 46 in linear time (indeed, this problem was explicitly posed by Pollack et al. [25] and later by
 47 Mitchell [22, Chapter 27]).

48 Several variations of these two problems have been considered. Indeed, the same problem
 49 has been studied under different metrics. For example, the L_1 geodesic distance [6], the link
 50 distance [26, 16, 9] (where we look for the path with the minimum possible number of bends or
 51 *links*), or even rectilinear link distance [23, 24] (a variation of the link distance in which only
 52 isothetic segments are allowed). The diameter and center of a simple polygon for both the L_1
 53 and rectilinear link metrics can be computed in linear time (whereas $O(n \log n)$ time is needed
 54 for the link distance).

55 Another natural extension is the computation of the diameter and center in polygonal domains
 56 (i.e., polygons with one or more holes). Polynomial time algorithms are known for both the
 57 diameter [4] and center [5], although the running times are significantly larger (i.e., $O(n^{7.73})$ and
 58 $O(n^{12+\epsilon})$, respectively).

59 1.2 Outline

60 In order to compute the geodesic center, c_P , Pollack et al. [25] introduced a linear time *chord-*
 61 *oracle*. Given a chord C that splits P into two sub-polygons, the oracle determines which
 62 sub-polygon contains c_P . Combining this operation with an efficient search on a triangulation
 63 of P , Pollack et al. narrow the search of c_P within a triangle (and afterwards find the center
 64 using optimization techniques). Their approach however, does not allow them to reduce the
 65 complexity of the problem in each iteration, and hence it runs in $\Theta(n \log n)$ time.

66 The general approach of our algorithm described in Section 6 is similar: partition P into $O(1)$
 67 cells, use an oracle to determine which cell contains c_P , and recurse within the cell. Our approach
 68 differs however in two important aspects that allows us to speed-up the algorithm. First, we do
 69 not use the chords of a triangulation of P to partition the problem into cells. We use instead
 70 a cutting of a suitable set of chords. Secondly, we compute a set Σ of $O(n)$ functions, each
 71 defined in a triangular domain contained in P , such that their upper envelope, $\phi(x)$, coincides
 72 with $F_P(x)$. Thus, we can “ignore” the polygon P and focus only on finding the minimum of the
 73 function $\phi(x)$.

74 The search itself uses ε -nets and cutting techniques, which guarantee that both the size of the
 75 cell containing c_P and the number of functions of Σ defined in it decrease by a constant fraction

(and thus leads to an overall linear time algorithm). This search has however two stopping conditions, (1) reach a subproblem of constant size, or (2) find a triangle containing c_P . In the latter case, we show that $\phi(x)$ is a convex function when restricted to this triangle. Thus, finding its minimum becomes an optimization problem that we solve in Section 7 using cuttings in \mathbb{R}^3 .

The key of this approach lies in the computation of the functions of Σ and their triangular domains. Each function $g(x)$ of Σ is defined in a triangular domain Δ contained in P and is associated to a particular vertex w of P . Intuitively speaking, $g(x)$ maps points in Δ to their (geodesic) distance to w . We guarantee that, for each point $x \in P$, there is one function g defined in a triangle containing x , such that $g(x) = F_P(x)$. To compute these triangles and their corresponding functions, we proceed as follows.

In Section 2, we use the matrix search technique introduced by Hershberger and Suri [15] to decompose the boundary of P , denoted by ∂P , into connected edge-disjoint chains. Each chain is defined by either (1) a consecutive list of vertices that have the same farthest neighbor v (we say that v is *marked* if it has such a chain associated to it), or (2) an edge whose endpoints have different farthest neighbors (such edge is called a *transition edge*).

In Section 3, we consider each transition edge ab of ∂P independently and compute its *hourglass*. Intuitively, the hourglass of ab , H_{ab} , is the region of P between two chains, the edge ab and the chain of ∂P that contains the farthest neighbors of all points in ab . Inspired by a result of Suri [27], we show that the sum of the complexities of all hourglasses defined on a transition edge is $O(n)$. (The *combinatorial complexity*, or simply complexity of a geometric object is the total number of vertices and edges that define it.) In addition, we provide a new technique to compute all these hourglasses in linear time.

While the hourglasses cover a big part of P , to complete the coverage we need to consider *funnels* from each marked vertex. Intuitively, the funnel of a marked vertex v is the region of P that contains all the paths from v to every point of ∂P that is farther from v than from any other vertex of P . In Section 4, we prove that the total complexity of all the funnels of all marked vertices is $O(n)$. Moreover, we provide an algorithm to compute all these funnels in linear time.

In Section 5 we show how to compute the functions in Σ and their respective triangles. We distinguish two cases: (1) Inside each hourglass H_{ab} of a transition edge, we use a technique introduced by Aronov et al. [2] that uses the shortest-path trees of a and b in H_{ab} to construct $O(|H_{ab}|)$ triangles with their respective functions (for more information on shortest-path trees refer to [11]). (2) Inside the funnel of each marked vertex v , we compute triangles that encode the distance from v . Moreover, we guarantee that these triangles cover every point of P whose farthest neighbor is v . Overall, we compute the $O(n)$ functions of Σ in linear time.

2 Decomposing the boundary

In this section, we decompose the boundary of P into chains of consecutive vertices that share the same farthest neighbor and edges of P whose endpoints have distinct farthest neighbors.

Using a result from Hershberger and Suri [15], in $O(n)$ time we can compute the farthest neighbor of each vertex of P . Recall that the farthest neighbor of each vertex of P is always a convex vertex of P [3] and is unique by our general position assumption. The (farthest) *Voronoi region* of a vertex v of P is the set of points $R(v) = \{x \in P : F_P(x) = |\pi(x, v)|\}$ (including boundary points).

We mark the vertices of P that are farthest neighbors of at least one vertex of P . Let M denote the set of marked vertices of P (clearly this set can be computed in $O(n)$ time after applying the result of Hershberger and Suri). In other words, M contains all vertices of P whose Voronoi region contains at least one vertex of P .




Figure 1: Each vertex of the boundary of P is assigned with a farthest neighbor which is then marked. The boundary is then decomposed into vertex-disjoint chains, each associated with a marked vertex, joined by transition edges (blue) whose endpoints have different farthest neighbors.

122 Given a vertex v of P , the vertices of P whose farthest neighbor is v appear contiguously
 123 along ∂P [2]. Therefore, after computing all these farthest neighbors, we effectively split the
 124 boundary into subchains, each associated with a different vertex of M ; see Figure 1.

125 Given two points x and y on ∂P , let $\partial P(x, y)$ be the polygonal chain that starts at x and
 126 follows the boundary of P clockwise until reaching y . We say that three (non-empty) disjoint
 127 sets A, B and C contained in ∂P are in *clockwise order* if $B \subset \partial P(a, c)$ for any $a \in A$ and any
 128 $c \in C$. (To ease notation, we say that three points $x, y, z \in \partial P$ are in clockwise order if $\{x\}, \{y\}$
 129 and $\{z\}$ are in clockwise order).

130 Let a and b be the endpoints of an edge of ∂P such that b is the clockwise neighbor of a along
 131 ∂P and $f(a) \neq f(b)$. Recall that we have computed $f(a)$ and $f(b)$ in the previous step and note
 132 that $a, b, f(a), f(b)$ are in clockwise order. For any vertex $v \in \partial P$ such that $f(a), v, f(b)$ are in
 133 clockwise order, we know that there cannot be a vertex u of P such that $f(u) = v$. As proved by
 134 Aronov et al. [2, Corollary 2.7.4], if there is a point x on ∂P whose farthest neighbor is v , then
 135 x must lie on the open segment (a, b) . In other words, the Voronoi region $R(v)$ restricted to ∂P
 136 is contained in (a, b) .

137 3 Hourglasses

138 For any polygonal chain $C = \partial P(p_0, p_k)$, the *hourglass* of C , denoted by H_C , is the simple polygon
 139 contained in P bounded by C , $\pi(p_k, f(p_0))$, $\partial P(f(p_0), f(p_k))$ and $\pi(f(p_k), p_0)$; see Figure 2. We
 140 call C and $\partial P(f(p_0), f(p_k))$ the *top* and *bottom* chains of H_C , respectively, while $\pi(p_k, f(p_0))$
 141 and $\pi(f(p_k), p_0)$ are referred to as the *walls* of H_C .

142 We say that the hourglass H_C is *open* if its walls are vertex-disjoint. We say that C is a
 143 *transition chain* if $f(p_0) \neq f(p_k)$ and neither $f(p_0)$ nor $f(p_k)$ are interior vertices of C . In
 144 particular, if an edge ab of ∂P is a transition chain, we say that it is a *transition edge* (see
 145 Figure 2).

146 **Lemma 3.1.** [Restatement of Lemma 3.1.3 of [2]] *If C is a transition chain of ∂P , then the*
 147 *hourglass H_C is an open hourglass.*

148 Note that by Lemma 3.1, the hourglass of each transition chain is open. In the remainder
 149 of the paper, all the hourglasses considered are defined by a transition chain. That is, they are
 150 open and their top and bottom chains are edge-disjoint.




Figure 2: Given two edge-disjoint transition chains, their hourglasses are open and the bottom chains of their hourglasses are also edge-disjoint. Moreover, these bottom chains appear in the same cyclic order as the top chains along ∂P .

151 The following lemma is depicted in Figure 2 and is a direct consequence of the Ordering
 152 Lemma proved by Aronov et al. [2, Corollary 2.7.4].

153 **Lemma 3.2.** *Let C_1, C_2, C_3 be three edge-disjoint transition chains of ∂P in clockwise order.*
 154 *Then, the bottom chains of H_{C_1}, H_{C_2} and H_{C_3} are also edge-disjoint and are in clockwise order.*

155 Let γ be a geodesic path joining two points on the boundary of P . We say that γ *separates*
 156 two points x_1 and x_2 of ∂P if the points of $X = \{x_1, x_2\}$ and the endpoints of γ alternate along
 157 the boundary of P . For consistency, we also say that γ separates x_1 and x_2 if either of them
 158 coincides with an endpoint of γ . We say that a geodesic path γ *separates* an hourglass H if it
 159 separates the points of its top chain from those of its bottom chain.

160 **Lemma 3.3.** *Let C_1, \dots, C_r be edge-disjoint transition chains of ∂P . Then, there is a set of*
 161 *$t \leq 10$ geodesic paths $\gamma_1, \dots, \gamma_t$ with endpoints on ∂P such that for each $1 \leq i \leq r$ there exists*
 162 *$1 \leq j \leq t$ such that γ_j separates H_{C_i} . Moreover, this set can be computed in $O(n)$ time.*

163 *Proof.* Aronov et al. showed that there exist four vertices v_1, \dots, v_4 of P and geodesic paths
 164 $\pi(v_1, v_2), \pi(v_2, v_3), \pi(v_3, v_4)$ such that for any point $x \in \partial P$, one of these paths separates x from
 165 $f(x)$ [2, Lemma 2.7.6]. Moreover, they show how to compute this set in $O(n)$ time.

166 Let $\Gamma = \{\pi(v_i, v_j) : 1 \leq i < j \leq 4\}$ and note that v_1, \dots, v_4 split the boundary of P
 167 into at most four connected components. If a chain C_i is completely contained in one of these
 168 components, then one path of Γ separates the top and bottom chain of H_{C_i} . Otherwise, some
 169 vertex v_j is an interior vertex of C_i . However, because the chains C_1, \dots, C_r are edge-disjoint,
 170 there are at most four chains in this situation. For each chain C_i containing a vertex v_j , we
 171 add the geodesic path connecting the endpoints of C_i to Γ . Therefore, Γ consists of at most 10
 172 geodesic paths and each hourglass H_{C_i} has its top and bottom chain separated by some path of
 173 Γ . Since only $O(1)$ additional paths are computed, this can be done in linear time. \square

174 A *chord* of P is an edge joining two non-adjacent vertices a and b of P such that $ab \subseteq P$.
 175 Therefore, a chord splits P into two sub-polygons.

176 **Lemma 3.4.** *[Restatement of Lemma 3.4.3 of [2]] Let C_1, \dots, C_r be a set of edge-disjoint tran-*
 177 *sition chains of ∂P in clockwise order. Then each chord of P appears in $O(1)$ hourglasses among*
 178 *H_{C_1}, \dots, H_{C_r} .*

179 *Proof.* Note that chords can only appear on walls of hourglasses. Because hourglasses are open,
 180 any chord must be an edge on exactly one wall of each of these hourglasses. Assume, for the

181 sake of contradiction, that there exist two points $s, t \in P$ whose chord st is in three hourglasses
 182 H_{C_i}, H_{C_j} and H_{C_k} (for some $1 \leq i < j < k \leq r$) such that s is visited before t when going from
 183 the top chain to the bottom one along the walls of the three hourglasses. Let s_i and t_i be the
 184 points in the top and bottom chains of H_{C_i} , respectively, such that $\pi(s_i, t_i)$ is the wall of
 185 H_{C_i} that contains st (analogously, we define s_k and t_k).

186 Because C_i, C_j, C_k are in clockwise order, Lemma 3.2 implies that the bottom chains of C_i, C_j
 187 and C_k are also in clockwise order. Therefore, C_j lies between s_i and s_k and the bottom chain of
 188 H_{C_j} lies between t_i and t_k . That is, for each $x \in C_j$ and each y in the bottom chain of H_{C_j} , the
 189 geodesic path $\pi(x, y)$ is “sandwiched” by the paths $\pi(s_i, t_i)$ and $\pi(s_k, t_k)$. In particular, $\pi(x, y)$
 190 contains st for each pair of points in the top and bottom chain of H_{C_j} . However, this implies
 191 that the hourglass H_{C_j} is not open—a contradiction that comes from assuming that st lies in
 192 the wall of three open hourglasses, when this wall is traversed from the top chain to the bottom
 193 chain. Analogous arguments can be used to bound the total number of walls that contain the
 194 edge st (when traversed in any direction) to $O(1)$. \square

195 **Lemma 3.5.** *Let x, u, y, v be four vertices of P in clockwise order. Given the shortest-path trees*
 196 *T_x and T_y of x and y in P , respectively, such that T_x and T_y can answer lowest common ancestor*
 197 *(LCA) queries in $O(1)$ time, we can compute the path $\pi(u, v)$ in $O(|\pi(u, v)|)$ time. Moreover,*
 198 *all edges of $\pi(u, v)$, except perhaps one, belong to $T_x \cup T_y$.*

199 *Proof.* Let X (resp. Y) be the set containing the LCA in T_x (resp. T_y) of u, y , and of v, y (resp.
 200 u, x and x, y). Note that the points of $X \cup Y$ lie on the path $\pi(x, y)$ and can be computed in
 201 $O(1)$ time by hypothesis. Moreover, using LCA queries, we can decide their order along the path
 202 $\pi(x, y)$ when traversing it from x to y . (Both X and Y could consist of a single vertex in some
 203 degenerate situations.) Two cases arise:

204 **Case 1.** If there is a vertex $x^* \in X$ lying after a vertex $y^* \in Y$ along $\pi(x, y)$, then the
 205 path $\pi(u, v)$ contains the path $\pi(y^*, x^*)$. In this case, the path $\pi(u, v)$ is the concatenation of
 206 the paths $\pi(u, y^*)$, $\pi(y^*, x^*)$, and $\pi(x^*, v)$ and that the three paths are contained in $T_x \cup T_y$.
 207 Moreover, $\pi(u, v)$ can be computed in time proportional to its length by traversing along the
 208 corresponding tree; see Figure 3 (top).

209 **Case 2.** In this case the vertices of X appear before the vertices of Y along $\pi(x, y)$. Let x'
 210 (resp. y') be the vertex of X (resp. Y) closest to x (resp. y).

211 Let u' be the last vertex of $\pi(u, x)$ that is also in $\pi(u, y)$. Note that u' can be constructed
 212 by walking from u' towards x until the path towards y diverges. Thus, u' can be computed in
 213 $O(|\pi(u, u')|)$ time. Define v' analogously and compute it in $O(|\pi(v, v')|)$ time.

214 Let P' be the polygon bounded by the geodesic paths $\pi(x', u')$, $\pi(u', y')$, $\pi(y', v')$ and $\pi(v', x')$.
 215 Because the vertices of X appear before those of Y along $\pi(x, y)$, P' is a simple polygon; see
 216 Figure 3 (bottom).

217 In this case the path $\pi(u, y)$ is the union of $\pi(u, u')$, $\pi(u', v')$ and $\pi(v', y)$. Because $\pi(u, u')$
 218 and $\pi(v', y)$ can be computed in time proportional to their length, it suffices to compute $\pi(u', v')$
 219 in $O(|\pi(u', v')|)$ time.

220 Note that P' is a simple polygon with only four convex vertices x', u', y' and v' , which are
 221 connected by chains of reflex vertices. Thus, the shortest path from x' to y' can have at most one
 222 diagonal edge connecting distinct reflex chains of P' . Since the rest of the points in $\pi(u', v')$ lie
 223 on the boundary of P' and from the fact that each edge of P' is an edge of $T_x \cup T_y$, we conclude
 224 all edges of $\pi(u, v)$, except perhaps one, belong to $T_x \cup T_y$.

225 We want to find the common tangent between the reflex paths $\pi(u', x')$ and $\pi(v', y')$, or the
 226 common tangent of $\pi(u', y')$ and $\pi(v', x')$ as one of them belongs to the shortest path $\pi(u', v')$.
 227 Assume that the desired tangent lies between the paths $\pi(u', x')$ and $\pi(v', y')$. Since these paths
 228 consist only of reflex vertices, the problem can be reduced to finding the common tangent of two




Figure 3: (top) Case 1 of the proof of Lemma 3.5 where the path $\pi(u, v)$ contains a portion of the path $\pi(x, y)$. (bottom) Case 2 of the proof of Lemma 3.5 where the path $\pi(u, v)$ has exactly one edge being the tangent of the paths $\pi(u', y')$ and $\pi(v', x')$.

229 convex polygons. By slightly modifying the linear time algorithm to compute this tangents, we
 230 can make it run in $O(|\pi(u', v')|)$ time.

231 Since we do not know if the tangent lies between the paths $\pi(u', x')$ and $\pi(v', y')$, we process
 232 the chains $\pi(u', y')$ and $\pi(v', x')$ in parallel and stop when finding the desired tangent. Conse-
 233 quently, we can compute the path $\pi(u, v)$ in time proportional to its length. \square

Lemma 3.6. *Let P be a simple polygon with n vertices. Given k disjoint transition chains C_1, \dots, C_k of ∂P , it holds that*

$$\sum_{i=1}^k |H_{C_i}| = O(n).$$

234 *Proof.* Because the given transition chains are disjoint, Lemma 3.2 implies that the bottom
 235 chains of their respective hourglasses are also disjoint. Therefore, the sum of the complexities
 236 of all the top and bottom chains of these hourglasses is $O(n)$. To bound the complexity of their
 237 walls we use Lemma 3.4. Since no chord is used more than a constant number of times, it suffices
 238 to show that the total number of chords used by all these hourglasses is $O(n)$.

To prove this, we use Lemma 3.3 to construct $O(1)$ paths $\gamma_1, \dots, \gamma_t$ such that for each $1 \leq i \leq k$, there is a path γ_j that separates the top and bottom chains of H_{C_i} . For each $1 \leq j \leq t$, let

$$\mathcal{H}^j = \{H_{C_i} : \text{the top and bottom chain of } H_{C_i} \text{ are separated by } \gamma_j\}.$$

Since the complexity of the shortest-path trees of the endpoints of γ_j is $O(n)$ [11], and from the fact that the chains C_1, \dots, C_k are disjoint, Lemma 3.5 implies that the total number of edges in all the hourglasses of \mathcal{H}^j is $O(n)$. Moreover, because each of these edges appears in $O(1)$ hourglasses among C_1, \dots, C_k , we conclude that

$$\sum_{H \in \mathcal{H}^j} |H| = O(n).$$

239 Since we have only $O(1)$ separating paths, our result follows. \square

240 3.1 Building hourglasses

241 Let E be the set of transition edges of ∂P . Given a transition edge $ab \in E$, we say that H_{ab} is
 242 a *transition hourglass*. In order to construct the triangle cover of P , we construct the transition
 243 hourglass of each transition edge of E . By Lemma 3.6, we know that $\sum_{ab \in E} |H_{ab}| = O(n)$.
 244 Therefore, our aim is to compute the cover in time proportional to the size of H_{ab} .

245 By Lemma 3.3 we can compute a set Γ of $O(1)$ paths such that for any transition edge
 246 ab , the transition hourglass H_{ab} is separated by one (or more) paths in this set. For each
 247 endpoint of the $O(1)$ paths of Γ , we compute its shortest-path tree in linear time [8, 11]. In
 248 addition, we preprocess these trees in linear time to support constant time LCA queries [14].
 249 Both computations need linear time per endpoint and use $O(n)$ space. Since we do this process
 250 for a constant number of endpoints, overall this preprocessing takes $O(n)$ time.

251 Let $\gamma \in \Gamma$ be a separating path. Note that γ separates the boundary of P into two chains S_γ
 252 and S'_γ such that $S_\gamma \cup S'_\gamma = \partial P$. Let $\mathcal{H}(\gamma)$ be the set of transition hourglasses separated by γ
 253 whose transition edge is contained in S_γ (whenever an hourglass is separated by more than one
 254 path, we pick one arbitrarily). Note that we can classify all transition hourglasses into the sets
 255 $\mathcal{H}(\gamma)$ in $O(n)$ time (since $|\Gamma| = O(1)$).

256 **Lemma 3.7.** *Given a separating path $\gamma \in \Gamma$, it holds that $\sum_{H \in \mathcal{H}(\gamma)} |H| = O(n)$. Moreover, we
 257 can compute all transition hourglasses of $\mathcal{H}(\gamma)$ in $O(n)$ time.*

258 *Proof.* Since all transition hourglasses in \mathcal{H} are defined from disjoint transition edges, Lemma 3.6
 259 implies that $\sum_{H \in \mathcal{H}(\gamma)} |H| = O(n)$.

260 By construction, the wall of each of these hourglasses consists of a (geodesic) path that
 261 connects a point in S_γ with a point in S'_γ . Let $u \in S_\gamma$ and $v \in S'_\gamma$ be two vertices such that
 262 $\pi(u, v)$ is the wall of a hourglass in $\mathcal{H}(\gamma)$. Because LCA queries can be answered in $O(1)$ time,
 263 Lemma 3.5 allows us to compute this path in $O(|\pi(u, v)|)$ time. Therefore, we can compute all
 264 hourglasses of $\mathcal{H}(\gamma)$ in $O(\sum_{H \in \mathcal{H}(\gamma)} |H| + n) = O(n)$. \square

265 Because $|\Gamma| = O(1)$ and since every transition hourglass is separated by at least one path
 266 in Γ , we obtain the following result.

267 **Corollary 3.8.** *The added complexity of the transition hourglasses of all transition edges of P
 268 is $O(n)$. Moreover, all these hourglasses can be constructed in $O(n)$ time.*

269 4 Funnels

270 Let $C = (p_0, \dots, p_k)$ be a chain of ∂P and let v be a vertex of P not in C . The *funnel* of v to C ,
 271 denoted by $S_v(C)$, is the simple polygon bounded by C , $\pi(p_k, v)$ and $\pi(v, p_0)$; see Figure 4 (a).
 272 Note that the paths $\pi(v, p_k)$ and $\pi(v, p_0)$ may coincide for a while before splitting into disjoint
 273 chains. We call C the *main chain* of $S_v(C)$ while $\pi(p_k, v)$ and $\pi(v, p_0)$ are referred to as the
 274 *walls* of the funnel. See Lee and Preparata [17] or Guibas et al. [11] for more details on funnels.

275 A subset $R \subset P$ is *geodesically convex* if for every $x, y \in R$, the path $\pi(x, y)$ is contained in
 276 R . This funnel $S_v(C)$ is then the minimum geodesically convex set that contains v and C .

277 Given two points $x, y \in P$, the (geodesic) *bisector* of x and y is the set of points contained
 278 in P that are equidistant from x and y . This bisector is a curve, contained in P , that consists
 279 of straight-line segments and hyperbolic arcs. Moreover, this curve intersects ∂P only at its
 280 endpoints [1, Lemma 3.22].




Figure 4: a) The funnel $S_v(C)$ of a vertex v and a chain C contained in ∂P are depicted. b) The decomposition of $S_v(C)$ into apexed triangles produced by the shortest-path map of v .

281 **Lemma 4.1.** *Let v be a vertex of P and let C be a transition chain such $R(v) \cap \partial P \subseteq C$ and*
 282 *$v \notin C$. Then, $R(v)$ is contained in the funnel $S_v(C)$*

283 *Proof.* Let a and b be the endpoints of C such that $a, b, f(a)$ and $f(b)$ are in clockwise order.
 284 Because $R(v) \cap \partial P \subset C$, we know that $f(a), v$ and $f(b)$ are in clockwise order.

285 Let α (resp. β) be the bisector of v and $f(a)$ (resp. $f(b)$). Let h_a (resp. h_b) be the set of
 286 points of P that are farther from v than from $f(a)$ (resp. $f(b)$). Note that α is the boundary of
 287 h_a while β bounds h_b .

288 By definition, we know that $R(v) \subseteq h_a \cap h_b$. Therefore, it suffices to show that $h_a \cap h_b \subset S_v(C)$.
 289 Assume for a contradiction that there is a point of $h_a \cap h_b$ lying outside of $S_v(C)$. By continuity
 290 of the geodesic distance, the boundaries of $h_a \cap h_b$ and $S_v(C)$ must intersect. Because $a \notin h_a$
 291 and $b \notin h_b$, both bisectors α and β must have an endpoint on the edge ab . Since the boundaries
 292 of $h_a \cap h_b$ and $S_v(C)$ intersect, we infer that $\beta \cap \pi(v, b) \neq \emptyset$ or $\alpha \cap \pi(v, a) \neq \emptyset$. Without loss of
 293 generality, assume that there is a point $w \in \beta \cap \pi(v, b)$, the case where w lies in $\alpha \cap \pi(v, a)$ is
 294 analogous.

Since $w \in \beta$, we know that $|\pi(w, v)| = |\pi(w, f(b))|$. By the triangle inequality and since w
 cannot be a vertex of P as w intersects ∂P only at its endpoints, we get that

$$|\pi(b, f(b))| < |\pi(b, w)| + |\pi(w, f(b))| = |\pi(b, w)| + |\pi(w, v)| = |\pi(b, v)|.$$

295 Which implies that b is farther from v than from $f(b)$ —a contradiction that comes from assuming
 296 that $h_a \cap h_b$ is not contained in $S_v(C)$. □

297 4.1 Funnels of marked vertices

298 Recall that for each marked vertex $v \in M$, we know at least of one vertex on ∂P such that v is
 299 its farthest neighbor.

300 For any marked vertex v , let u_1, \dots, u_{k-1} be the vertices of P such that $v = f(u_i)$ and assume
 301 that u_1, \dots, u_{k-1} are in clockwise order. Let u_0 and u_k be the neighbors of u_1 and u_{k-1} other
 302 than u_2 and u_{k-2} , respectively. Note that both $u_0 u_1$ and $u_{k-1} u_k$ are transition edges of P . Let
 303 $C_v = (u_0, \dots, u_k)$ and consider the funnel $S_v(C_v)$ defined by v .

304 **Lemma 4.2.** *Let x be a point in P . If $f(x) = v$ for some marked vertex $v \in M$, then $x \in S_v(C_v)$.*

305 *Proof.* Since $f(u_0) \neq f(u_k)$, C_v is a transition chain. Moreover, C_v contains $R(v) \cap \partial P$ by
 306 definition. Therefore, Lemma 4.1 implies that $R(v) \subset S_v(C_v)$. Since $v = f(x)$, we know that
 307 $x \in R(v)$ and hence that $x \in S_v(C_v)$. \square

308 We focus now on computing the funnels defined by the marked vertices of P .

309 **Lemma 4.3.** *Given a marked vertex v such that $C_v = (u_0, \dots, u_k)$, it holds that $|S_v(C_v)| =$
 310 $O(k + |H_{u_0u_1}| + |H_{u_{k-1}u_k}|)$. Moreover, $S_v(C_v)$ can be computed in $O(|S_v(C_v)|)$ time.*

311 *Proof.* Since $H_{u_0u_1}$ and $H_{u_{k-1}u_k}$ are transition hourglasses, we can assume that they have been
 312 computed using Corollary 3.8.

313 Because $v = f(u_1) = f(u_{k-1})$, we know that v is a vertex of both $H_{u_0u_1}$ and $H_{u_{k-1}u_k}$. By
 314 definition, we have $\pi(v, u_0) \subset H_{u_0u_1}$ and $\pi(v, u_k) \subset H_{u_{k-1}u_k}$. Thus, we can compute both paths
 315 $\pi(v, u_0)$ and $\pi(v, u_k)$ in $O(|H_{u_0u_1}| + |H_{u_{k-1}u_k}|)$ time [11] and their complexities are $O(|H_{u_0u_1}|)$ and
 316 $O(|H_{u_{k-1}u_k}|)$, respectively. So, overall, the funnel $S_v(C_v)$ has total complexity $O(k + |H_{u_0u_1}| +$
 317 $|H_{u_{k-1}u_k}|)$ and can be constructed in linear time in its size. \square

Recall that, by Corollary 3.8, the total sum of the complexities of all transition hourglasses
 is $O(n)$. Therefore, Lemma 4.3 implies that

$$\sum_{v \in M} |S_v(C_v)| = O\left(n + \sum_{ab \in E} |H_{ab}|\right) = O(n).$$

318 Because the funnel of each marked vertex can be computed in linear time in its size, we obtain
 319 the following result.

320 **Corollary 4.4.** *The added complexity of the funnels of all marked vertices of P is $O(n)$. More-*
 321 *over, all these funnels can be constructed in $O(n)$ time.*

322 5 Covering the polygon with apexed triangles

An *apexed triangle* $\Delta = (a, b, c)$ with *apex* a is a triangle contained in P with an associated
 distance function $g_\Delta(x)$, called the *apex function* of Δ , such that (1) a is a vertex of P , (2) there
 is an edge of ∂P containing both b and c , and (3) there is a vertex w of P , called the *definer* of
 Δ , such that

$$g_\Delta(x) = \begin{cases} -\infty & \text{if } x \notin \Delta \\ |xa| + |\pi(a, w)| = |\pi(x, w)| & \text{if } x \in \Delta \end{cases}$$

323 Intuitively, Δ bounds a constant complexity region where the geodesic distance function from
 324 w is explicitly stored by our algorithm.

325 In this section, we show how to find a set of $O(n)$ apexed triangles of P such that the upper
 326 envelope of their apex functions coincides with $F_P(x)$. To this end, we first decompose the
 327 transition hourglasses into apexed triangles that encode all the geodesic distance information
 328 inside them. For each marked vertex $v \in M$ we construct a funnel that contains the Voronoi
 329 region of v . We then decompose this funnel into apexed triangles that encode the distance from v .

330 The same approach was already used by Pollack et al. in [25, Section 3]. Given a segment
 331 contained in the interior of P , they show how to compute a linear number of apexed triangles
 332 such that $F_P(x)$ coincides with the upper envelope of the corresponding apex functions in the
 333 given segment.

334 While the construction we follow is analogous, we use it in the transition hourglass H_{ab}
 335 instead of the full polygon P . Therefore, we have to specify what is the relation between the




Figure 5: (left) A vertex v visible from the segment ab lying on the bottom chain of H_{ab} , and the triangle Δ_v which contains the portion of ab visible from v . (right) The children u_1 and u_2 of v are visible from ab while u_3 is not. The triangle Δ_v is split into apexed triangles by the rays going from u_1 and u_2 to v .

336 upper envelope of the computed functions and $F_P(x)$. We will show that the upper envelope
 337 of the apex functions computed in H_{ab} coincides with $F_P(x)$ inside the Voronoi region $R(v)$ of
 338 every vertex $v \in B_{ab}$.

339 5.1 Inside a transition hourglass

340 Let ab be a transition edge of P such that b is the clockwise neighbor of a along ∂P . Let B_{ab}
 341 denote the open bottom chain of H_{ab} . As noticed above, a point on ∂P can be farthest from a
 342 vertex in B_{ab} only if it lies in the open segment ab . That is, if v is a vertex of B_{ab} such that
 343 $R(v) \neq \emptyset$, then $R(v) \cap \partial P \subset ab$.

344 In fact, not only this Voronoi region is inside H_{ab} when restricted to the boundary of P , but
 345 we can further bound its location and show that $R(v) \subset H_{ab}$. The next result follows trivially
 346 from Lemma 4.1.

347 **Corollary 5.1.** *Let v be a vertex of B_{ab} . Then $R(v) \subset H_{ab}$.*

348 Our objective is to compute $O(|H_{ab}|)$ apexed triangles contained in H_{ab} , each with its distance
 349 function, such that the upper envelope of these apex functions coincides with $F_P(x)$ restricted
 350 to H_{ab} inside the Voronoi region of every vertex in B_{ab} .

351 Let T_a and T_b be the shortest-path trees in H_{ab} rooted at a and b , respectively. We can
 352 compute these trees in $O(|H_{ab}|)$ time [11]. For each vertex v such that $f(a), v$ and $f(b)$ are in
 353 clockwise order, let v_a and v_b be the neighbors of v in the paths $\pi(v, a)$ and $\pi(v, b)$, respectively.
 354 We say that a vertex v is *visible* from ab if $v_a \neq v_b$. Note that if a vertex v is visible, then
 355 the extension of the segments vv_a and vv_b must intersect the top segment ab at points
 356 x_a and x_b , respectively. Therefore, for each visible vertex v , we obtain a triangle $\Delta_v = \Delta(v, x_a, x_b)$ as
 357 shown in Figure 5.

358 We further split Δ_v into a series of triangles with apex at v as follows: Let u be a child of v
 359 in either T_a or T_b . As noted by Pollack et al., v can be of three types, either (1) u is not visible
 360 from ab (and is hence a child of v in both T_a and T_b); or (2) u is visible from ab , is a child of v
 361 only in T_b , and v_bvu is a left turn; or (3) u is visible from ab , is a child of v only in T_a , and v_avu
 362 is a right turn.

363 Let u_1, \dots, u_{k-1} be the children of v of type (2) sorted in clockwise order around v . Let $c(v)$
 364 be the maximum distance from v to any invisible vertex in the subtrees of T_a and T_b rooted at
 365 v ; if no such vertex exists, then $c(v) = 0$. Define a function $d_l(v)$ on each vertex v of H_{ab} in a
 366 recursive fashion as follows: If v is invisible from ab , then $d_l(v) = c(v)$. Otherwise, let $d_l(v)$ be

367 the maximum of $c(v)$ and $\max\{d_l(u_i) + |u_i v| : u_i \text{ is a child of } v \text{ of type (2)}\}$. Symmetrically, we
 368 define a function $d_r(v)$ using the children of type (3) of v .

For each $1 \leq i \leq k-1$, extend the segment $u_i v$ past v until it intersects ab at a point s_i . Let
 s_0 and s_k be the intersections of the extensions of vv_a and vv_b with the segment ab . We define
 then k apexed triangles contained in Δ_v as follows. For each $0 \leq i \leq k-1$, consider the triangle
 $\Delta(s_i, v, s_{i+1})$ whose associated apexed (left) function is

$$f_i(x) = \begin{cases} |xv| + \max_{j>i}\{c(v), |vu_j| + d_l(u_j)\} & \text{if } x \in \Delta(s_i, v, s_{i+1}) \\ -\infty & \text{otherwise} \end{cases}$$

369 In a symmetric manner, we define a set of apexed triangles induced by the type (3) children of v
 370 and their respective apexed (right) functions.

371 Let g_1, \dots, g_r and $\Delta_1, \dots, \Delta_r$ respectively be an enumeration of all the generated apex func-
 372 tions and apexed triangles such that g_i is defined in the triangle Δ_i . Because each function is
 373 determined uniquely by a pair of adjacent vertices in T_a or in T_b , and since these trees have
 374 $O(|H_{ab}|)$ vertices each, we conclude that $r = O(|H_{ab}|)$.

375 Note that for each $1 \leq i \leq r$, the apexed triangle Δ_i has two vertices on the segment ab
 376 and a third vertex, say a_i , being its apex such that for each $x \in \Delta_i$, $g_i(x) = |\pi(x, w_i)|$ for some
 377 vertex w_i of H_{ab} . Recall that w_i is called the definer of Δ_i .

378 **Lemma 5.2.** *Given a transition edge ab of P , we can compute a set \mathcal{A}_{ab} of $O(|H_{ab}|)$ apexed
 379 triangles in $O(|H_{ab}|)$ time with the property that for any point $p \in P$ such that $f(p) \in B_{ab}$, there
 380 is an apexed triangle $\Delta \in \mathcal{A}_{ab}$ with apex function g and definer equal to $f(p)$ such that*

- 381 1. $p \in \Delta$ and
- 382 2. $g(p) = F_P(p)$.

383 *Proof.* Because $p \in R(f(p))$, Lemma 5.1 implies that $p \in H_{ab}$. Consider the path $\pi(p, f(p))$ and
 384 let v be the successor of p along this path. By construction of \mathcal{A}_{ab} , there is a triangle $\Delta \in \mathcal{A}_{ab}$
 385 apexed at v with definer w that contains p . The apex function $g(x)$ of Δ encodes the geodesic
 386 distance from x to w . Because $F_P(x)$ is the upper envelope of all the geodesic functions, we know
 387 that $g(p) \leq F_P(p)$.

388 To prove the other inequality, note that if $v = f(p)$, then trivially $g(p) = |pv| + |\pi(v, w)| \geq$
 389 $|pv| = |\pi(p, f(p))| = F_P(p)$. Otherwise, let z be the next vertex after v in the path $\pi(p, f(p))$.
 390 Three cases arise:

(a) If z is invisible from ab , then so is $f(p)$ and hence,

$$|\pi(p, f(p))| = |pv| + |\pi(v, f(p))| \leq |pv| + c(v) \leq g(p).$$

(b) If z is a child of type (2), then z plays the role of some child u_j of v in the notation used
 during the construction. In this case:

$$|\pi(p, f(p))| = |pv| + |vz| + |\pi(z, f(p))| \leq |pv| + |vu_j| + d_l(u_j) \leq g(p).$$

391 (c) If z is a child of type (3), then analogous arguments hold using the (right) distance d_r .

392 Therefore, regardless of the case $F_P(p) = |\pi(p, f(p))| \leq g(p)$.

393 To bound the running time, note that the recursive functions d_l, d_r and c can be computed in
 394 $O(|T_a| + |T_b|)$ time. Then, for each vertex visible from ab , we can process it in time proportional
 395 to its degree in T_a and T_b . Because the sum of the degrees of all vertices in T_a and T_b is
 396 $O(|T_a| + |T_b|)$ and from the fact that both $|T_a|$ and $|T_b|$ are $O(|H_{ab}|)$, we conclude that the total
 397 running time to construct \mathcal{A}_{ab} is $O(|H_{ab}|)$. \square

398 In other words, Lemma 5.2 says that no information on farthest neighbors is lost if we only
 399 consider the functions of \mathcal{A}_{ab} within H_{ab} . In the next section we use a similar approach to
 400 construct a set of apexed triangles (and their corresponding apex functions), so as to encode the
 401 distance from the vertices of M .

402 5.2 Inside the funnels of marked vertices

We now proceed to split a given funnel into $O(|S_v(C_v)|)$ apexed triangles that encode the distance
 function from v . To this end, we use the algorithm described by Guibas et al. [12, Section 2] to
 compute the shortest-path map of v in $S_v(C_v)$ in $O(|S_v(C_v)|)$ time. This algorithm produces a
 partition of $S_v(C_v)$ into $O(|S_v(C_v)|)$ interior disjoint triangles with vertices on ∂P , such that each
 triangle consists of all points in $S_v(C_v)$ whose shortest path to v consists of the same sequence
 of vertices; see Figure 4 (b). Let Δ be a triangle in this partition and let a be its apex, i.e., the
 first vertex found along each path $\pi(x, v)$, where $x \in \Delta$. We define the apex function $g_\Delta(x)$ of
 Δ as follows:

$$g_\Delta(x) = \begin{cases} |xa| + |\pi(a, v)| & \text{if } x \in \Delta \\ -\infty & \text{otherwise} \end{cases}$$

403 Notice that for each $x \in \Delta$, $g_\Delta(x) = |\pi(x, v)|$.

404 **Lemma 5.3.** *The shortest-path map of v in $S_v(C_v)$ can be computed in $O(|S_v(C_v)|)$ time and pro-*
 405 *duces $O(|S_v(C_v)|)$ interior disjoint apexed triangles such that their union covers $S_v(C_v)$. More-*
 406 *over, for each point $x \in R(v)$, there is an apexed triangle Δ with apex function $g(x)$ such that*
 407 *(1) $x \in \Delta$ and (2) $g(x) = F_P(x)$.*

408 *Proof.* The above procedure splits $S_v(C_v)$ into $O(|S_v(C_v)|)$ apexed triangles, such that the apex
 409 function in each of them is defined as the geodesic distance to v . By Lemma 4.2, if $x \in R(v)$,
 410 then $x \in S_v(C_v)$. Therefore, there is an apexed triangle Δ with apex function $g(x)$ such that
 411 $x \in \Delta$ and $g(x) = |\pi(x, v)| = F_P(x)$. Thus, we obtain properties (1) and (2). \square

412 6 Prune and search

413 With the tools introduced in the previous sections, we can describe the prune and search algo-
 414 rithm to compute the geodesic center. The idea of the algorithm is to partition P into $O(1)$ cells,
 415 determine in which cell of P the center lies and recurse on that cell as a new subproblem with
 416 smaller complexity.

417 We can discard all apexed triangles that do not intersect the new cell containing the center.
 418 Using cuttings to produce this partition of P , we can show that both the complexity of the cell
 419 containing the center, and the number of apexed triangles that intersect it decrease by a constant
 420 fraction in each iteration of the algorithm. This process is then repeated until either of the two
 421 objects has constant descriptive size.

422 Let τ be the set of all apexed triangles computed in previous sections. Corollary 3.8 and
 423 Lemma 5.2 bound the number of apexed triangles constructed inside the transition hourglasses,
 424 while Corollary 4.4 and Lemma 5.3 do so inside the funnels of the marked vertices. We obtain
 425 the following.

426 **Corollary 6.1.** *The set τ consists of $O(n)$ apexed triangles.*

427 Let $\phi(x)$ be the upper envelope of the apex functions of the triangles in τ (i.e., $\phi(x) =$
 428 $\max\{g(x) : \Delta \in \tau \text{ and } g(x) \text{ is the apex function of } \Delta\}$). The following result is a direct conse-
 429 quence of Lemmas 5.2 and 5.3, and shows that the $O(n)$ apexed triangles of τ not only cover P ,
 430 but their apex functions suffice to reconstruct the function $F_P(x)$.

431 **Corollary 6.2.** *The functions $\phi(x)$ and $F_P(x)$ coincide in the domain of points of P , i.e., for*
 432 *each $p \in P$, $\phi(p) = F_P(p)$.*

433 Given a chord C of P a *half-polygon* of P is one of the two simple polygons in which C splits
 434 P . A k -*cell* of P is a simple polygon obtained as the intersection of at most k half-polygons.
 435 Because a k -cell is the intersection of geodesically convex sets, it is also geodesically convex.

436 The recursive algorithm described in this section takes as input a 4-cell (initially equal to
 437 P) containing the geodesic center of P and the set of apexed triangles of τ that intersect it. In
 438 each iteration, it produces a new 4-cell of smaller complexity that intersects just a fraction of
 439 the apexed triangles and contains the geodesic center of P . By recursing on this new cell, the
 440 complexity of the problem is reduced in each iteration.

441 Let R be a 4-cell of P (initially equal to P) containing the geodesic center of P and let τ_R
 442 be the set of apexed triangles of τ that intersect R . Let $m_R = \max\{|R|, |\tau_R|\}$, where $|R|$ denotes
 443 the combinatorial complexity R . Recall that, by construction of the apexed triangles, for each
 444 triangle of τ_R at least one and at most two of its boundary segments are chords of P . Let \mathcal{C}_R
 445 be the set containing all chords that belong to the boundary of a triangle of τ_R . Therefore,
 446 $|\mathcal{C}_R| \leq 2|\tau_R| \leq 2m_R$.

447 To construct ε -nets, we need some definitions (for more information on ε -nets refer to [20]).
 448 Let φ be the set of all open 4-cells of P . For each $t \in \varphi$, let $\mathcal{C}_R(t) = \{C \in \mathcal{C}_R : C \cap t \neq \emptyset\}$ be the
 449 set of chords of \mathcal{C}_R induced by t . Finally, let $\varphi_{\mathcal{C}_R} = \{\mathcal{C}_R(t) : t \in \varphi\}$ be the family of subsets of
 450 \mathcal{C}_R induced by φ . Consider the set system $(\mathcal{C}_R, \varphi_{\mathcal{C}_R})$ (denoted by (\mathcal{C}_R, φ) for simplicity) defined
 451 by \mathcal{C}_R and φ . The proof of the next lemma is deferred to the Appendix for ease of readability.

452 **Lemma 6.3.** *The set system (\mathcal{C}_R, φ) has constant VC-dimension.*

453 Let $\varepsilon > 0$ be a sufficiently small constant whose exact value will be specified later. Because
 454 the VC-dimension of the set system (\mathcal{C}_R, φ) is finite by Lemma 6.3, we can compute an ε -net N
 455 of (\mathcal{C}_R, φ) in $O(|\mathcal{C}_R|/\varepsilon) = O(m_R)$ time [20]. The size of N is $O(\frac{1}{\varepsilon} \log \frac{1}{\varepsilon}) = O(1)$ and its main
 456 property is that any 4-cell that does not intersect a chord of N will intersect at most $\varepsilon|\mathcal{C}_R|$ chords
 457 of \mathcal{C}_R .

458 Observe that N partitions R into $O(1)$ sub-polygons (not necessarily 4-cells). We further
 459 refine this partition to obtain 4-cells. That is, we shoot vertical rays up and down from each
 460 endpoint of N , and from the intersection point of any two segments of N , see Figure 6. Overall,
 461 this partitions R into $O(1)$ 4-cells such that each either (i) is a convex polygon contained in P
 462 of at most four vertices, or otherwise (ii) contains some chain of ∂P . Since $|N| = O(1)$, the
 463 whole decomposition can be computed in $O(m_R)$ time (the intersections between segments of N
 464 are done in constant time, and for the ray shooting operations we walk along the boundary of R
 465 once).

466 In order to determine which 4-cell contains the geodesic center of P , we extend each edge of a
 467 4-cell to a chord C . This can be done with two ray-shooting queries (each of which takes $O(m_R)$
 468 time). We then use the chord-oracle from Pollack et al. [25, Section 3] to decide which side of
 469 C contains c_P . The only requirement of this technique is that the function $F_P(x)$ coincides with
 470 the upper envelope of the apex functions when restricted to C . This holds by Corollary 6.2 and
 471 from the fact that τ_R consists of all the apexed triangles of τ that intersect R .

472 Because the chord-oracle described by Pollack et al. [25, Section 3] runs in time linear in the
 473 number of functions defined on C , we can decide in total $O(m_R)$ time in which side of C the
 474 geodesic center of P lies. Since our decomposition into 4-cells has constant complexity, we need
 475 to perform $O(1)$ calls to the oracle before determining the 4-cell R' that contains the geodesic
 476 center of P .

477 Note that the chord-oracle computes the minimum of $F_P(x)$ restricted to the chord before
 478 determining the side containing the minimum. In particular, if c_P lies on any chord bounding




Figure 6: The ϵ -net N splits R into $O(1)$ sub-polygons that are further refined into a 4-cell decomposition using $O(1)$ ray-shooting queries from the vertices of the arrangement defined by N .

479 R' , then the chord-oracle will find it. Therefore, we can assume that c_P lies in the interior of R' .
 480 Moreover, since N is a ϵ -net, we know that at most $\epsilon|\mathcal{C}_R|$ chords of \mathcal{C}_R intersect R' .

481 Observe that both the complexity of R' and $\tau_{R'}$ decrease by a constant fraction. Indeed, by
 482 construction of the cutting at most $2\epsilon m_R$ apexed triangles can intersect R' (and thus $|\tau_{R'}| \leq$
 483 $2\epsilon m_R$).

484 In order to bound the complexity of R' we use Corollary 6.2: function $F_P(x)$ is defined in
 485 each point of R' , and in particular for each vertex v of R' there must exist an apexed triangle
 486 $\Delta \in \tau_{R'}$ such that $v \in \Delta$. By definition of apexed triangles, each such triangle can contain at
 487 most three vertices of R' . Combining this fact with the bound of $|\tau_{R'}|$ we obtain that R' has at
 488 most $3|\tau_{R'}| \leq 6\epsilon m_R$ vertices. Thus, if we choose $\epsilon = 1/12$, we guarantee that both the size of
 489 the 4-cell R' and the number of apexed triangles in $\tau_{R'}$ are at most $m_R/2$.

490 In order to proceed with the algorithm on R' recursively, we need to compute the set $\tau_{R'}$
 491 with the at most $\epsilon|\mathcal{C}_R|$ apexed triangles of τ_R that intersect R' (i.e., prune the apexed triangles
 492 that do not intersect with R'). For each apexed triangle $\Delta \in \tau_R$, we can determine in constant
 493 time if it intersects R' (either one of the endpoints is in $R' \cap \partial P$ or the two boundaries have
 494 non-empty intersection in the interior of P). Overall, we need $O(m_R)$ time to compute the at
 495 most $\epsilon|\mathcal{C}_R|$ triangles of τ_R that intersect R' .

496 By recursing on R' , we guarantee that after $O(\log m_R)$ iterations, we reduce the size of either
 497 τ_R or R' to constant. In the former case, the minimum of $F_P(x)$ can be found by explicitly
 498 constructing function ϕ in $O(1)$ time. In the latter case, we triangulate R' and apply the chord-
 499 oracle to determine which triangle will contain c_P . The details needed to find the minimum of
 500 $\phi(x)$ inside this triangle are given in the next section.

501 **Lemma 6.4.** *In $O(n)$ time we can find either the geodesic center of P or a triangle containing*
 502 *the geodesic center.*

503 7 Finding the center within a triangle

504 In order to complete the algorithm it remains to show how to find the geodesic center of P for the
 505 case in which R' is a triangle. If this triangle is in the interior of P , it may happen that several
 506 apexed triangles of τ (up to a linear number) fully contain R' . Thus, the pruning technique used
 507 in the previous section cannot be further applied. We solve this case with a different approach.

508 Recall that $\phi(x)$ denotes the upper envelope of the apex functions of the triangles in τ , and
 509 the geodesic center is the point that minimizes ϕ . The key observation is that, as it happened
 510 with chords, the function $\phi(x)$ restricted to R' is convex.

Let $\tau_{R'} = \{\Delta_1, \Delta_2, \dots, \Delta_m\}$ be the set of $m = O(n)$ apexed triangles of τ that intersect R' . Let a_i and w_i be the apex and the definer of Δ_i , respectively. Let $g_i(x)$ be the apex function of Δ_i such that

$$g(x) = \begin{cases} |xa_i| + \kappa_i & \text{if } x \in \Delta_i \\ -\infty & \text{otherwise} \end{cases},$$

511 where $\kappa_i = |\pi(a_i, w_i)|$ is a constant.

512 By Corollary 6.2, we have $\phi(x) = F_P(x)$. Therefore, the problem of finding the center is
513 equivalent to the following optimization problem in \mathbb{R}^3 :

(P1). Find a point $(x, r) \in \mathbb{R}^3$ minimizing r subject to $x \in R'$ and

$$g_i(x) \leq r, \text{ for } 1 \leq i \leq m.$$

514 Thus, we need to find the solution to (P1) to find the geodesic center of P . We use some
515 remarks described by Megiddo in order to simplify the description of (P1) [21].

To simplify the formulas, we square the equation $|xa_i| \leq r - \kappa_i$:

$$\|x\|^2 - 2x \cdot a_i + \|a_i\|^2 = |xa_i|^2 \leq (r - \kappa_i)^2 = r^2 - 2r\kappa_i + \kappa_i^2.$$

And finally for each $1 \leq i \leq m$, we define the function $h_i(x, r)$ as follows:

$$h_i(x, r) = \begin{cases} \|x\|^2 - 2x \cdot a_i + \|a_i\|^2 - r^2 + 2r\kappa_i - \kappa_i^2 & \text{if } x \in \Delta_i \\ -\infty & \text{otherwise} \end{cases}.$$

516 Therefore, our optimization problem can be reformulated as:

(P2). Find a point $(x, r) \in \mathbb{R}^3$ such that r is minimized subject to $x \in R'$ and

$$h_i(x, r) \leq 0 \text{ and } r > \max\{\kappa_i\}, \text{ for } 1 \leq i \leq m.$$

517 Let $h'_i(x, r) = \|x\|^2 - 2x \cdot a_i + \|a_i\|^2 - r^2 + 2r\kappa_i - \kappa_i^2$ be a function defined in the entire plane
518 and let (P2') be an optimization problem equivalent to (P2) where every instance of $h_i(x, r)$ is
519 replaced by $h'_i(x, r)$. The optimization (P2') was studied by Megiddo in [21]. We provide some
520 of the intuition used by Megiddo to solve this problem.

Although the functions $h'_i(x, r)$ are not linear, they all have the same non-linear terms. Therefore, for $i \neq j$, we get that equation $h'_i(x, r) = h'_j(x, r)$ defines a *separating plane*

$$\gamma_{i,j} = \{(x, r) \in \mathbb{R}^3 : 2(\kappa_i - \kappa_j)r - 2(a_i - a_j) \cdot x + \|a_i\|^2 - \|a_j\|^2 - \kappa_i^2 + \kappa_j^2 = 0\}$$

521 As noted by Megiddo [21], this separating plane has the following property: If the solution
522 (x, r) to (P2') is known to lie to one side of $\gamma_{i,j}$, then one of the constraints is redundant.

523 Thus, to solve (P2') it sufficed to have a *side-decision oracle* to determine in which side of
524 a plane $\gamma_{i,j}$ the solution lies. Megiddo showed how to implement this oracle in a way that the
525 running time is proportional to the number of constraints [21].

526 Once we have such an oracle, problem (P2') can be solved using a prune and search approach:
527 pair the functions arbitrarily, and consider the set of $m/2$ separating planes defined by these pairs.
528 For some constant t , compute a $1/t$ -cutting in \mathbb{R}^3 of the separating planes. A $1/t$ -cutting is a
529 partition of the plane into $O(t^3) = O(1)$ convex regions each of which is of constant complexity
530 and intersects at most $m/2t$ separating planes. A cutting of planes can be computed in linear
531 time in \mathbb{R}^3 for any $t = O(1)$ [18]. After computing the cutting, determine in which of the
532 regions the minimum lies by performing $O(1)$ calls to the side-decision oracle. Because at least
533 $(t-1)m/2t$ separating planes do not intersect this constant complexity region, for each of them
534 we can discard one of the constraints as it becomes redundant. By repeating this algorithm
535 recursively we obtain a linear running time.

536 To solve (P2) we follow a similar approach, but our set of separating planes needs to be
537 extended in order to handle apex functions as they are only defined in the same way as in (P2')
538 in a triangular domain. Note that no vertex of an apexed triangle can lie inside R' .

7.1 Optimization problem in a convex domain

In this section we describe our algorithm to solve the optimization problem (P2). We start by pairing the apexed triangles of $\tau_{R'}$ arbitrarily to obtain $m/2$ pairs. By identifying the plane where P lies with the plane $Z_0 = \{(x, y, z) : z = 0\}$, we can embed each apexed triangle in \mathbb{R}^3 . A *plane-set* is a set consisting of at most five planes in \mathbb{R}^3 . For each pair of apexed triangles (Δ_i, Δ_j) we define its plane-set as follows: For each chord of P bounding either Δ_i or Δ_j (at most two chords on each triangle), consider the line extending this chord and the vertical extrusion of this line in \mathbb{R}^3 , i.e., the plane containing this chord orthogonal to Z_0 . The set containing these planes, together with the separating plane $\gamma_{i,j}$, is the plane-set of the pair (Δ_i, Δ_j) .

Let Γ be the union of all the plane-sets defined by the $m/2$ pairs of apexed triangles. Because the plane-set of each pair (Δ_i, Δ_j) consists of at most five planes and at least one unique to this pair, namely $\gamma_{i,j}$, we infer that $m/2 \leq |\Gamma| \leq 5m/2$.

Compute a $1/t$ -cutting of Γ in $O(m)$ time for some constant t to be specified later. Because t is constant, this $1/t$ -cutting splits the space into $O(1)$ convex regions, each bounded by a constant number of planes [18]. Using a side-decision algorithm (to be specified later), we can determine the region Q of the cutting that contains the solution to (P2). Because Q is the region of a $1/t$ -cutting of Γ , we know that at most $|\Gamma|/t$ planes of Γ intersect Q . In particular, at most $|\Gamma|/t$ plane-sets intersect Q and hence, at least $(t-1)|\Gamma|/t \geq (t-1)m/2t$ plane-sets do not intersect Q .

Let (Δ_i, Δ_j) be a pair whose plane-set does not intersect Q . Let Q' be the projection of Q on the plane Z_0 . Because the plane-set of this pair does not intersect Q , we know that Q' intersects neither the boundary of Δ_i nor that of Δ_j . Two cases arise:

Case 1. If either Δ_i or Δ_j does not intersect Q' , then we know that their apex function is redundant and we can drop the constraint associated with this apexed triangle.

Case 2. If $Q' \subset \Delta_i \cap \Delta_j$, then we need to decide which constraint to drop. To this end, we consider the separating plane $\gamma_{i,j}$. Notice that inside the vertical extrusion of $\Delta_i \cap \Delta_j$ (and hence in Q), the plane $\gamma_{i,j}$ has the property that if we know its side containing the solution to (P2), then one of the constraints can be dropped.

Regardless of the case, if the plane-set of a pair (Δ_i, Δ_j) does not intersect Q , then we can drop one of its constraints. Since at least $(t-1)m/2t$ plane-sets do not intersect Q , we can drop at least $(t-1)m/2t$ constraints. By choosing $t = 2$, we are able to drop at least $(t-1)m/2t = m/4$ constraints. Consequently, after $O(m)$ time, we are able to drop a constant fraction of the apexed triangles. By repeating this process recursively, we end up with a constant size problem in which we can compute the upper envelope of the functions explicitly and find the solution to (P2) using exhaustive search. Thus, the running time of this algorithm is bounded by the recurrence $T(m) = T(3m/4) + O(m)$ which solves to $O(m)$. Because $m = O(n)$, we can find the solution to (P2) in $O(n)$ time.

The last detail is the implementation of the side-decision algorithm. Given a plane γ , we need to determine in which side lies the solution to (P2). To this end, we solve (P2) restricted to γ , i.e., with the additional constraint of $(x, r) \in \gamma$. This approach was used by Megiddo [21], the idea is to recurse by reducing the dimension of the problem. Another approach is to use a slight modification of the chord-oracle described by Pollack et al. [25, Section 3].

Once the solution to (P2) restricted to γ is known, we can follow the same idea used by Megiddo [21] to find the side of γ containing the global solution to (P2). Find the apex functions that define the minimum restricted to γ . Since $\phi(x) = F_P(x)$ is locally defined by these functions, we can decide in which side the minimum lies using convexity. We obtain the following result.

Lemma 7.1. *Let R' be a convex trapezoid contained in P such that R' contains the geodesic center of P . Given the set of all apexed triangles of τ that intersect R' , we can compute the*

587 geodesic center of P in $O(n)$ time.

588 The following theorem summarizes the result presented in this paper.

589 **Theorem 7.2.** *We can compute the geodesic center of any simple polygon P of n vertices in*
590 *$O(n)$ time.*

591 8 Further work

592 Another way to compute the center of a simple polygon is to compute the farthest-point Voronoi
593 diagram of its vertices. While the best known algorithm for this task runs in $O(n \log n)$ time,
594 no lower bound is known for this instance of the problem. Therefore, it is natural to ask if the
595 farthest-point Voronoi diagram of the set of vertices of a simple polygon can be computed in
596 $O(n)$ time.

597 To generalize the result presented in this paper, we ask the following question. Given a set
598 S of m points inside of a simple polygon P with n vertices, how fast can we compute the center
599 of S inside P ? That is, the point in P that minimizes the maximum geodesic distance to a
600 point of S . While the (geodesic) farthest-point Voronoi diagram of S provides the answer in
601 $O((n + m) \log(n + m))$ time, we ask whether it is possible to compute this center in $O(n + m)$
602 time.

603 References

- 604 [1] B. Aronov. On the geodesic Voronoi diagram of point sites in a simple polygon. *Algorithmica*,
605 4(1-4):109–140, 1989.
- 606 [2] B. Aronov, S. Fortune, and G. Wilfong. The furthest-site geodesic Voronoi diagram. *Discrete*
607 *& Computational Geometry*, 9(1):217–255, 1993.
- 608 [3] T. Asano and G. Toussaint. Computing the geodesic center of a simple polygon. Technical
609 Report SOCS-85.32, McGill University, 1985.
- 610 [4] S. W. Bae, M. Korman, and Y. Okamoto. The geodesic diameter of polygonal domains.
611 *Discrete & Computational Geometry*, 50(2):306–329, 2013.
- 612 [5] S. W. Bae, M. Korman, and Y. Okamoto. Computing the geodesic centers of a polygonal
613 domain. In *Proceedings of CCCG*, 2014.
- 614 [6] S. W. Bae, M. Korman, Y. Okamoto, and H. Wang. Computing the L_1 geodesic diameter
615 and center of a simple polygon in linear time. In *Proceedings of LATIN*, pages 120–131,
616 2014.
- 617 [7] B. Chazelle. A theorem on polygon cutting with applications. In *Proceedings of FOCS*,
618 pages 339–349, 1982.
- 619 [8] B. Chazelle. Triangulating a simple polygon in linear time. *Discrete & Computational*
620 *Geometry*, 6(1):485–524, 1991.
- 621 [9] H. Djidjev, A. Lingas, and J.-R. Sack. An $O(n \log n)$ algorithm for computing the link center
622 of a simple polygon. *Discrete & Computational Geometry*, 8:131–152, 1992.

- 623 [10] H. Edelsbrunner and E. P. Mücke. Simulation of simplicity: a technique to cope with
624 degenerate cases in geometric algorithms. *ACM Transactions on Graphics*, 9(1):66–104,
625 1990.
- 626 [11] L. Guibas, J. Hershberger, D. Leven, M. Sharir, and R. E. Tarjan. Linear-time algorithms
627 for visibility and shortest path problems inside triangulated simple polygons. *Algorithmica*,
628 2(1-4):209–233, 1987.
- 629 [12] L. J. Guibas and J. Hershberger. Optimal shortest path queries in a simple polygon. *Journal*
630 *of computer and system sciences*, 39(2):126–152, 1989.
- 631 [13] H. Harborth and M. Möller. *The Esther-Klein-problem in the projective plane*. Inst. für
632 Mathematik, TU Braunschweig, 1993.
- 633 [14] D. Harel and R. E. Tarjan. Fast algorithms for finding nearest common ancestors. *SIAM*
634 *Journal on Computing*, 13(2):338–355, 1984.
- 635 [15] J. Hershberger and S. Suri. Matrix searching with the shortest-path metric. *SIAM Journal*
636 *on Computing*, 26(6):1612–1634, 1997.
- 637 [16] Y. Ke. An efficient algorithm for link-distance problems. In *Proceedings of SoCG*, pages
638 69–78, 1989.
- 639 [17] D.-T. Lee and F. P. Preparata. Euclidean shortest paths in the presence of rectilinear
640 barriers. *Networks*, 14(3):393–410, 1984.
- 641 [18] J. Matoušek. Approximations and optimal geometric divide-and-conquer. *Journal of Com-*
642 *puter and System Sciences*, 50(2):203–208, 1995.
- 643 [19] J. Matoušek. *Lectures on discrete geometry*, volume 108. Springer New York, 2002.
- 644 [20] J. Matoušek. Construction of epsilon nets. In *Proceedings of SoCG*, pages 1–10, New York,
645 1989. ACM.
- 646 [21] N. Megiddo. On the ball spanned by balls. *Discrete & Computational Geometry*, 4(1):605–
647 610, 1989.
- 648 [22] J. S. B. Mitchell. Geometric shortest paths and network optimization. In J.-R. Sack and
649 J. Urrutia, editors, *Handbook of Computational Geometry*, pages 633–701. Elsevier, 2000.
- 650 [23] B. Nilsson and S. Schuierer. Computing the rectilinear link diameter of a polygon. In
651 *Proceedings of CG*, pages 203–215, 1991.
- 652 [24] B. Nilsson and S. Schuierer. An optimal algorithm for the rectilinear link center of a recti-
653 linear polygon. *Computational Geometry: Theory and Applications*, 6:169–194, 1996.
- 654 [25] R. Pollack, M. Sharir, and G. Rote. Computing the geodesic center of a simple polygon.
655 *Discrete & Computational Geometry*, 4(1):611–626, 1989.
- 656 [26] S. Suri. *Minimum Link Paths in Polygons and Related Problems*. PhD thesis, Johns Hopkins
657 Univ., 1987.
- 658 [27] S. Suri. Computing geodesic furthest neighbors in simple polygons. *Journal of Computer*
659 *and System Sciences*, 39(2):220–235, 1989.
- 660 [28] P. Turán. On an extremal problem in graph theory. *Mat. Fiz. Lapok*, 48(436-452):137, 1941.

661 A Bounding the VC dimension

662 In this section we provide the proof of Lemma 6.3. That is, we want to prove that the set system
 663 (\mathcal{C}_R, φ) has constant VC-dimension. Recall that \mathcal{C}_R is a set of chords of P and φ is the set of all
 664 open 4-cells of P .

665 Let $A \subseteq \mathcal{C}_R$ be a subset of chords. We say that A is *shattered* by φ if each of the subsets of
 666 A can be obtained as the intersection of some $Z \in \varphi$ with A , i.e., if for each $\sigma \subseteq A$, there exists
 667 $Z \in \varphi$ such that $\sigma = Z \cap A$. The *VC-dimension* of (\mathcal{C}_R, φ) is the supremum of the sizes of all
 668 finite shattered subsets of \mathcal{C}_R .

669 Let $\mathcal{H} = \{H : H \text{ is a half-polygon of } P\}$. Because each 4-cell of P is the intersection of at
 670 most four half-polygons of P , the following result is a direct consequence of Proposition 10.3.3
 671 of [19, Chapter 10].

672 **Lemma A.1.** *If $(\mathcal{C}_R, \mathcal{H})$ has VC-dimension d , then (\mathcal{C}_R, φ) has VC-dimension $O(d)$.*

673 Let A be an arbitrary subset of \mathcal{C}_R such that $|A| = \kappa$ for some constant $\kappa > 6$ to be determined
 674 later. Recall that if no subset of \mathcal{C}_R of size κ is shattered by \mathcal{H} , then the VC-dimension of $(\mathcal{C}_R, \mathcal{H})$
 675 is at most $\kappa - 1$.

676 By Lemma A.1, it suffices to show that A is not shattered by \mathcal{H} to bound the VC-dimension
 677 of $(\mathcal{C}_R, \mathcal{H})$ and hence of (\mathcal{C}_R, φ) . We spend the remainder of this section proving that A is not
 678 shattered by \mathcal{H} .

679 A 6-cell is *strict* if it is defined as the intersection of six half-polygons, non of which are
 680 redundant, i.e., the removal of any of them modifies the 6-cell.

681 **Lemma A.2.** *If there are six chords of A supporting six half-polygons whose intersection defines
 682 a strict 6-cell σ , then A is not shattered by \mathcal{H} .*

683 *Proof.* Let C_1, \dots, C_6 be the chords supporting the six half-polygons whose intersection defines
 684 σ . Assume that C_1, \dots, C_6 appear in these order when walking clockwise along the boundary of
 685 σ . Note that any half-polygon that intersects C_1, C_3 and C_5 must intersect either C_2, C_4 or C_6 .
 686 Therefore, the set $\{C_1, C_3, C_5\} \subseteq A$ cannot be obtained as the intersection of some half-polygon
 687 $H \in \mathcal{H}$ with A . Consequently A is not shattered by \mathcal{H} . \square

688 Given two chords C_1 and C_2 of A , we say that C_1 and C_2 are *separated* if there exists a chord
 689 $S \in A$ such that C_1 and C_2 lie on different open half-polygons supported by S . In this case, we
 690 say that S *separates* C_1 from C_2 .

691 Note that if A contains two chords C_1 and C_2 that are separated by a chord S , then any half-
 692 polygon that intersects both C_1 and C_2 must also intersect S . In this case, the subset $\{C_1, C_2\}$
 693 cannot be obtained as the intersection of a half-polygon $H \in \mathcal{H}$ with A , i.e., A is not shattered
 694 by \mathcal{H} . Therefore, we assume from now on that no two chords of A are separated.

695 Let G_A be the intersection graph of A , i.e., the graph with vertex set A and an edge between
 696 two chords if they intersect. An Erdős-Szekeres type result from Harborth and Möller [13] shows
 697 that every arrangement of nine pseudo-lines determines a sub-arrangement with a hexagonal
 698 face. Thus, if G_A has a clique of size nine, then this subset of chords is a set of pseudo-lines.
 699 Therefore, it contains a subset of 6 chords that define a strict 6-cell. In this case, Lemma A.2
 700 implies that A is not shattered by \mathcal{H} . Consequently, we assume from now on that G_A has no
 701 clique of size nine.

702 Turán's Theorem [28] states that if G_A has no clique of size nine, then it has at most $(7/16)\kappa^2$
 703 edges. Let C_1 be the chord in A with the smallest degree in G_A . Notice that C_1 has degree at
 704 most $7\kappa/16$. Therefore, there are at least $9\kappa/16$ chords of A do not intersect C_1 . Consider the
 705 two half-polygons supported by C_1 and note that one of them, say P' , contains at least $9\kappa/32$
 706 chords of A that do not intersect C_1 . Let A' be the set containing these chords.

707 Let G'_A be subgraph of G_A induced by A' and let C_2 be the chord of A' with smallest degree
708 in G'_A . Because G'_A has no clique of size nine, we infer that C_2 has degree at most $7|A'|/16$.
709 Repeating the same argument, there is a set A'' of at least $9|A'|/16$ chords of A' that intersect
710 neither C_2 nor C_1 . Because we assumed that no two chords of A are separated, all chords in A''
711 must be contained in the half-polygon supported by C_2 that contains C_1 . Otherwise, C_1 and
712 some of these chords are separated by C_2 .

713 Let G''_A be the subgraph of G_A induced by A'' . Repeating the above procedure recursively
714 on G''_A and A'' four more times, we obtain a set C_1, \dots, C_6 of pairwise disjoint chords such
715 that for each $1 \leq i \leq 6$, C_1, \dots, C_{i-1} are contained in the same half-polygon supported by C_i .
716 Consequently, the set $\{C_1, \dots, C_6\}$ bounds a strict 6-cell.

The above process can be applied as long as

$$\left(\frac{9}{32}\right) \left(\frac{9}{16}\right)^4 \kappa \geq 1.$$

717 That is, as long as $|A| \geq 36$ we can always find 6 such chords defining a strict 6-cell. In this
718 case, Lemma A.2 implies that A is not shattered by \mathcal{H} . Consequently, no set of 36 chord is
719 shattered, i.e., the set system $(\mathcal{C}_R, \mathcal{H})$ has VC-dimension at most 35. By Lemma A.1, we obtain
720 the following result.

721 **► Lemma 6.3.** *The set system (\mathcal{C}_R, φ) has constant VC-dimension.*